

Segmenting Known Objects and Unseen Unknowns without Prior Knowledge

Stefano Gasperini^{1,2,◦} Alvaro Marcos-Ramiro² Michael Schmidt²
Nassir Navab¹ Benjamin Busam¹ Federico Tombari^{1,3}

¹ Technical University of Munich ² BMW Group ³ Google

Abstract

Panoptic segmentation methods assign a known class to each pixel given in input. Even for state-of-the-art approaches, this inevitably enforces decisions that systematically lead to wrong predictions for objects outside the training categories. However, robustness against out-of-distribution samples and corner cases is crucial in safety-critical settings to avoid dangerous consequences. Since real-world datasets cannot contain enough data points to adequately sample the long tail of the underlying distribution, models must be able to deal with unseen and unknown scenarios as well. Previous methods targeted this by re-identifying already-seen unlabeled objects. In this work, we propose the necessary step to extend segmentation with a new setting which we term holistic segmentation. Holistic segmentation aims to identify and separate objects of unseen, unknown categories into instances without any prior knowledge about them while performing panoptic segmentation of known classes. We tackle this new problem with U3HS, which finds unknowns as highly uncertain regions and clusters their corresponding instance-aware embeddings into individual objects. By doing so, for the first time in panoptic segmentation with unknown objects, our U3HS is trained without unknown categories, reducing assumptions and leaving the settings as unconstrained as in real-life scenarios. Extensive experiments on public data from MS COCO, Cityscapes, and Lost&Found demonstrate the effectiveness of U3HS for this new, challenging, and assumptions-free setting called holistic segmentation. Project page: <https://holisticseg.github.io>.

1. Introduction

Since neural networks have achieved unprecedented performance in perception tasks (e.g., object detection and semantic segmentation), there has been a growing interest

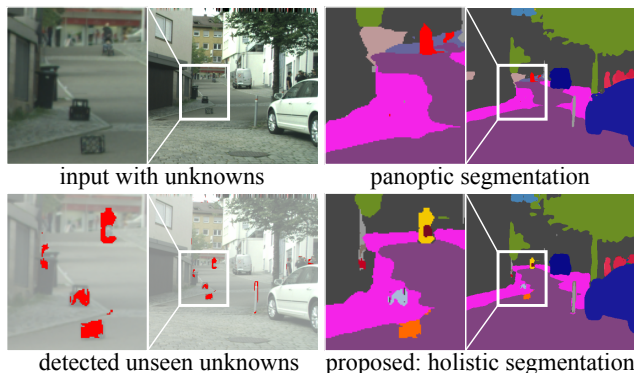


Figure 1. State-of-the-art panoptic segmentation methods [13] cannot deal with unseen classes from [54] (top right). Instead, our U3HS addresses the proposed holistic segmentation setting. U3HS finds unseen unknowns (bottom left) and separates them into instances (bottom right) without prior knowledge about unknowns.

in ensuring their safe deployment, especially important for safety-critical scenarios, such as autonomous driving and robotics [27]. Recently, several works have been proposed to improve robustness and generalization by addressing corner cases and out-of-distribution data [6, 67, 24], via domain adaptation [72], adversarial augmentations [46], simulations [1], and uncertainty estimation [56].

Due to the difficulty of collecting corner cases from the long tail of the underlying data distribution, current datasets cannot fully represent the diversity of the world, leaving its vast majority as difficult out-of-distribution (OOD) samples [34, 8]. In safety-critical applications, considering them during development and deployment is of utmost importance [3, 46], or they could cause severe damage.

Furthermore, since the powerful and popular *softmax* highly promotes the probability of the highest *logit*, state-of-the-art methods tend to be overly confident even on wrong predictions [60, 27]. In safety-critical settings, reliable confidence together with interpretability techniques [21, 80] increases trust for downstream tasks [27], e.g., trajectory prediction and path planning. Towards this end, estimating the uncertainty of a model’s output is commonly considered a key enabler for its safe applicability [41, 27].

[◦] This work was conducted while working at BMW Group.
Contact author: Stefano Gasperini (stefano.gasperini@tum.de).

While several works addressed some of these problems for image classification [60, 48, 56, 63] and object detection [50, 23], they remain primarily unexplored for dense tasks such as semantic and panoptic segmentation [5, 68]. Compared to object detection, the problem is more severe for dense tasks, where a model needs to provide a prediction for every input unit (e.g., each pixel). So unseen objects (i.e., of new, unseen categories) are systematically and wrongly assigned to one of the limited number of known classes (closed-set), as shown in Figure 1. This has led researchers towards designing new methods that work not only with the available data distribution but also with OOD samples that are not available (open-set), thereby improving robustness against unseen scenarios [40, 48, 5, 46].

Open-set panoptic segmentation [68] segments instances of unlabeled objects in addition to panoptic segmentation of known areas, i.e., the combination of semantic and instance segmentation [43]. Unlike OOD segmentation [40], segmenting unknown instances enables tracking and trajectory prediction. Prior works [68, 38, 74] tackled this problem by relying on seeing unlabeled categories during training. They learned these categories through the *void* class (i.e., unlabeled) and assumed unknowns to be within ground truth *void* regions at training time and inside *void* predictions at test time. By doing so, unknowns are transformed into learned unlabeled instances (i.e., essentially known objects) [38], constraining the open-set task. Mainly intended to segment already-seen unlabeled objects [38, 74], current works cannot deal with the wide variability of unknowns and corner cases outside the training data.

In this paper, we propose the necessary next step for panoptic segmentation to include object categories outside the training data (i.e., unseen unknowns). We term the new setting holistic segmentation. The aim is to identify and segment unseen unknowns into instances while segmenting known classes in a panoptic fashion without any external nor prior knowledge about unknowns. Unseen categories pose new challenges compared to already-seen unlabeled ones [38], requiring new solutions. Estimating the uncertainty is a key step towards finding the knowledge boundaries of a model, leaving the problem unconstrained and reducing assumptions on the training data. Therefore, we propose U3HS: Unseen Unknowns via Uncertainty estimation for Holistic Segmentation. The main contributions of this paper can be summarized as follows:

- We introduce the setting of holistic segmentation, which highlights the importance of not using prior knowledge about unknown objects (e.g., text), and leaves the setup unconstrained as in real scenarios.
- We tackle this new setting with U3HS: the first panoptic framework to deal with unseen, unknown object categories, able to segment and separate them.

- We provide uncertainty measures for the output of U3HS to further improve its safe applicability.

2. Related Work

Closed-set panoptic segmentation Combining semantic and instance segmentation, panoptic segmentation [43] distinguishes *things* (countable classes) from *stuff* (amorphous). The vast majority of methods are top-down [73, 62, 55, 42, 36, 52]: two-stage exploiting box proposals and *thing* masks from Mask R-CNN [32], and filling up *stuff* areas with a semantic branch. Bottom-up are proposal-free, e.g., Panoptic-DeepLab [13]: they segment semantically and cluster instances within *thing* regions [13, 66]. A different line of work proposed end-to-end solutions [25, 65] where instance and semantic segments are delivered directly by treating instance segmentation as a class-agnostic classification task. Others explored self-attention [36], videos [69, 11, 49], scene graphs [71, 70], multi-task learning [17, 30], neural fields [45], or text descriptions [19]. Our U3HS framework deals with unseen unknowns and extends [13] via instance-aware embeddings.

Zero-shot learning aims to predict unseen classes outside the training set [4, 81, 78] with the help of external knowledge [10, 28], e.g., a language model [79], used to build semantic spaces common between seen and unseen classes [76]. While zero-shot methods detect only unseen classes at inference time, generalized zero-shot approaches also detect seen ones [9, 58], similarly to the proposed holistic segmentation. Also **open-vocabulary** methods are zero-shot [37, 29, 75, 19]. They exploit language models such as CLIP [59] to describe unknowns. CLIP has been trained on unknown classes and is treated as an oracle, as it is assumed to be able to describe every unknown, allowing open-vocabulary and zero-shot approaches to identify them. However, this implies that unknown classes are known, e.g., to CLIP. Moreover, CLIP is not immune to corner cases and long tail samples [59]. This limits the pool of objects that these methods can recognize. As shown in Figure 2, holistic segmentation segments unseen objects too, but unlike zero-shot and open-vocabulary, it does not use any external support (e.g., text descriptions of unknowns), such that objects of unseen categories are segmented solely by learning on known ones. More recent than this work, SAM [44] is a strong foundation model. Unlike SAM, ours does not use any prompts and outputs semantic classes.

Uncertainty estimation Epistemic uncertainty is caused by the model itself, while aleatoric is due to the input [41, 27]. OOD data typically results in high epistemic due to a knowledge gap. Single deterministic approaches are sampling-free and provide predictions and uncertainty estimates with the same model [27]. Among these, DUQ [63] learns class representatives and compares them with input features. SNGP [48] improves the awareness to domain

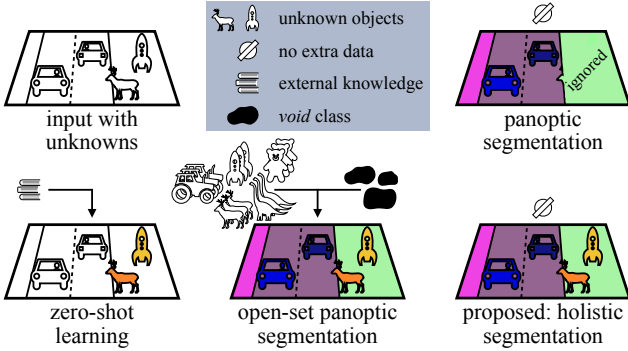


Figure 2. Comparison between closed-set (top right) [43] and open-set [38] panoptic segmentation, zero-shot learning [81], and the proposed holistic segmentation setting. While zero-shot and open-set panoptic methods commonly leverage knowledge about unknown objects, holistic segmentation does not use any priors.

shifts via weight normalization and a Gaussian process. DPN [60] predicts the parameters of a Dirichlet distribution, and uses a Dirichlet density function for each probability assignment and its uncertainty. Various works estimated uncertainty for object detection [51, 50, 23], and segmentation [57, 5, 61], improving robustness and generalization. While most compute uncertainty only to provide it as extra output [27, 61], our U3HS can be paired with any of the above techniques to find unknown objects, which it then separates into instances for holistic segmentation.

Open-set perception Open-set tasks are similar to generalized zero-shot learning [9], with the fundamental difference that here no external knowledge on the unseen classes is used [78]: inference is based only on what was learned from the training data. Uncertainty estimation helps to identify knowledge boundaries [51]. Bayesian SSD [51] uses Dropout sampling for open-set object detection. MLUC [6] tackles this for LiDAR point clouds via metric learning and unsupervised clustering. Open world recognition [2, 7] labels detected unknowns and adds them to the training set. Pham et al. [53] grouped regions perceptually to known and unknown instances, exploiting edges, boxes, and masks. Recently, several works tackled open-set semantic segmentation, telling apart unknown areas from known classes [40, 35, 31]. Among those not learning from OOD data, DML [5] uses metric learning and SML [40] acts as post-processing, standardizing the max logits, improving the class distributions. Instead, our work separates unseen, unknown objects into instances and segments known areas, addressing the proposed holistic segmentation.

Open-set panoptic segmentation The pioneering OSIS [68] was the first in this direction. Applied to LiDAR data, it exploits 3D locations to cluster unlabeled points into instances. Later, EOPSN [38] extended Panoptic FPN [42] and grouped its proposals into clusters. At training time, EOPSN clusters similar unlabeled objects across multiple

inputs. When surrounded by known segments, it labels an unlabeled object and uses it to learn to segment its instances. Instead, DDOSP [74] uses a known-unknown class discriminator and class-agnostic proposals. However, as these approaches were intended to re-identify already-seen unlabeled objects, they all rely on seeing unknown data at training time [68, 38, 74]. They all cluster into instances what falls in the predicted *void* class, learned as a fallback (i.e., must be in the training samples) and assumed to contain all unknowns. Since datasets are limited [46], by requiring to learn from unknowns and the *void* class, existing works are not designed to deal with any completely unseen object, as their pool of identifiable unknowns is also limited [38, 74]. For these reasons, they solve only part of the problem. Instead, by using no OOD data at training time and preventing learning priors for unknowns, the proposed holistic segmentation differs from the way open-set panoptic segmentation has been tackled so far [68, 38, 74]. As shown in Figure 2, our setting allows to segment without constraints and separate even unseen unknown objects. In contrast to all existing approaches, the proposed U3HS neither relies on seeing unknowns at training time nor learns the *void* class. U3HS is the first method to solve this unconstrained, assumptions-free holistic segmentation setting.

3. Proposed Setting: Holistic Segmentation

As shown in Figure 2, the proposed setting of holistic segmentation is a logical extension of open-set panoptic segmentation [68, 38]. To make the setting unconstrained as real-life scenarios, we aim to identify and separate **any unseen, unknown object** into instances while segmenting known classes. Other settings allow to include unknowns in the training data [68, 38, 74] (e.g., within the *void* class) or use information about them [81], and only re-identify already-seen unlabeled objects [38]. Instead, we focus on the case where **no information is available about unknowns**. Therefore, holistic segmentation is more challenging, makes no assumptions about the training data (e.g., the presence of *void*, with unknowns in it), leaves the problem unconstrained to any object, and simplifies data collection as no unknowns need to be in the training data. Formal definition and metrics follow open-set panoptic segmentation [68, 38]. As shown in Figure 2, the outputs are comparable, but the definition of unknowns differs (here unseen), as well as the inability to learn from unknowns. Unknown instances can then be used for downstream tracking, trajectory prediction, path planning, or active learning.

4. Proposed Framework: U3HS

In Figure 3, we show a representation of U3HS, targeting holistic segmentation. U3HS outputs instances of unseen unknowns by clustering instance-aware embeddings corre-

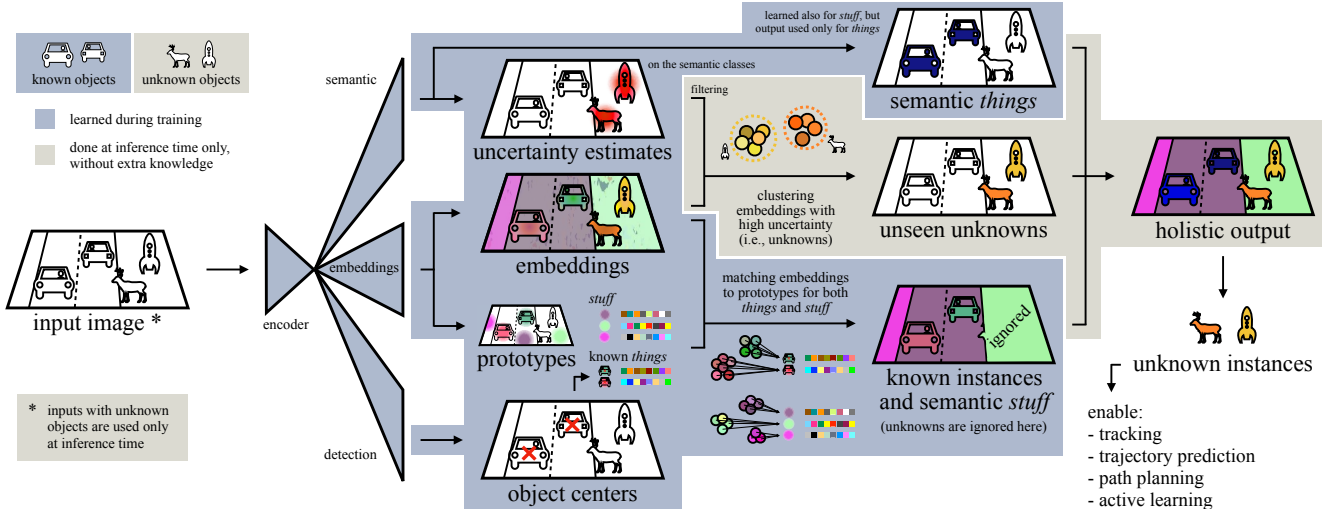


Figure 3. The proposed U3HS framework. Uncertainty is estimated in the semantic branch, and with the instance-aware embeddings, it determines unknown instances. Known instances are found via center regression and formed by grouping embeddings with their prototypes.

sponding to highly uncertain regions (Section 4.2). We use uncertainty estimation to distinguish known classes from unknowns, while embeddings are learned solely on known objects with panoptic segmentation (Section 4.1).

4.1. Panoptic Segmentation for Known Classes

Our approach for closed-set panoptic segmentation builds upon learning instance-aware embeddings. As shown in Figure 3, an encoder extracts features from an input image and propagates them to different decoders: 1) a semantic branch performing semantic segmentation and uncertainty estimation to identify unknown regions (Section 4.2); 2) a detection branch identifying object centers similarly to Panoptic-DeepLab [13]; and 3) an embeddings branch, with two separate heads, for prototypes and embeddings.

We make the embeddings instance-aware via discriminative loss functions (Section 4.3) and by concatenating the detection branch features to prototype and embeddings heads. Embeddings and detections are made also semantic-aware by concatenating the semantic logits to the last layers of the heads. Prototypes Ω are feature vectors the prototype head predicts to represent objects and *stuff* classes. While computed at each pixel, only the features at object centers C are considered *thing* prototypes, plus one for each *stuff* class. This is inspired by the size heatmap in [82]. *Thing* prototypes are $\Omega_{th} = \{(\mu_o, \sigma_o^2) \in \mathbb{R}^F \times \mathbb{R}^+ : o \in I\}$, one for each object o of all detected instances I . F is the embedding size, μ and σ^2 are mean and variance. *Stuff* prototypes are $\Omega_{st} = \{(\mu_k, \sigma_k^2) \in \mathbb{R}^F \times \mathbb{R}^+ : k \in K_{st}\}$, one for each *stuff* class $k \in K_{st}$ independently.

Similarly to [68], the embeddings head predicts embeddings $\phi_{(i,j)} \in \mathbb{R}^F$ for each pixel (i, j) , then matches them with their prototype $\omega \in \Omega$ as follows. We compute association scores $\hat{y}_{(i,j),\omega}$ for each pixel (i, j) and prototype as:

$$\hat{y}_{(i,j),\omega} = -\|\phi_{(i,j)} - \mu_\omega\|^2 / 2\sigma_\omega^2 \quad (1)$$

Compared to [68], we relax the problem by not including the term $-\frac{F}{2} \log \sigma_\omega^2$, and let the embedding variance be indirectly controlled by the final task, which naturally bounds it (shown empirically in Section 5.2). Then, we keep the prototype variance σ_ω^2 strictly positive by using *softplus*.

At inference time, for *things*, the semantic class of each instance is determined by majority voting of its semantic branch predictions, ensuring output consistency. Instead, the ID is computed from the highest score in Eq. 1. For *stuff* regions, we follow [68], determining the semantic classes by associating the pixel embeddings to the prototypes Ω_{st} via the highest scoring class from Eq. 1. This decoupling allows semantic awareness throughout the model.

4.2. Dealing with Unseen Unknown Objects

We find unknown segments by relying on uncertainty estimates, which can help identify the knowledge boundaries of a model [60, 48]. Specifically, instead of predicting the *void* class and searching in it for unknowns as in [68, 38], we estimate the uncertainty related to the semantic segmentation predictions and consider as unknown the areas with a high associated uncertainty. Although our framework can flexibly work with various uncertainty estimators (Section 5.2), here we exemplify it with DPN [60, 39], which we extended from image classification to semantic segmentation, and also improved its convergence in this context. We chose DPNs as they allow for minimal modifications at training time, i.e., replacing the *softmax* with a strictly positive activation function while providing good uncertainty estimates on OOD data without training on such data [60].

Following [60], we consider the evidence $e_k = \alpha_k - 1$ as a measure of the number of hints given by data for a pixel

to be assigned to a class $k \in K$ known classes, with α_k being the parameters of the Dirichlet distribution $Dir(\alpha)$. We compute the uncertainty as $u = K / \sum_{k=1}^K \alpha_k$. Given that the class probabilities $\mathbf{p} = \{p_k : k = [1, \dots, K]\}$ follow a simplex (i.e., are positive and sum to 1), the class assignment corresponds to a Dirichlet distribution parametrized over the evidence, as the probability density function:

$$D(\mathbf{p}|\alpha) = B(\alpha)^{-1} \prod_{k=1}^K p_k^{\alpha_k - 1} \quad (2)$$

with: $B(\alpha) = \prod_{k=1}^K \Gamma(\alpha_k) / \Gamma(\sum_{k=1}^K \alpha_k)$

where Γ is the gamma function and $B(\alpha)$ is the K -dimensional multinomial beta function [60].

We apply this to semantic segmentation by predicting a concentration parameter $\alpha^{(i,j)}$ for each pixel (i, j) , replacing the last layer with the smooth *softplus* activation function, thus converting the *logits* to a strictly positive vector, which we use as evidence $e^{(i,j)}$ in the Dirichlet distribution. We learn this distribution with the semantic loss \mathcal{L}_s minimizing the negative expected log likelihood of the correct class $Y^{(i,j)}$, for the random variable $\mathcal{X}^{(i,j)} \sim Dir(\alpha^{(i,j)})$:

$$\begin{aligned} \mathcal{L}_s^{(i,j)} &= -E[\ln \mathcal{X}_{Y^{(i,j)}}^{(i,j)}] \\ &= \psi \left(\sum_{k=1}^K \alpha_{(i,j),k} \right) - \psi(\alpha_{Y^{(i,j)}}) \end{aligned} \quad (3)$$

where ψ is the digamma function (i.e., Γ 's logarithmic derivative) and $\alpha_{(i,j),k}$ is the output of the semantic branch. Due to the difficulty of modeling the target distribution in our holistic setting, we omit the KL term used in [60], simplifying the loss design (Section 5.2). After training on the closed-set data, we consider all pixels (i, j) with an estimated uncertainty $u_{(i,j)} \geq \mu + t \cdot \sigma$ as unknown regions with μ and σ^2 being mean and variance of the uncertainties of all training pixels, and t being a hyperparameter.

Separating unknowns After finding the unknown segments, we cluster their instance-aware embeddings trained only on known objects into individual unknowns using DBSCAN [20]. We find the DBSCAN hyperparameters on the training closed-set data (Appendix). Finally, we re-assign the few DBSCAN's outliers to their originally predicted semantic class, thus ignoring their uncertainty estimates.

4.3. Learning to Find Knowns and Unknowns

We train our models with a combination of four losses. The semantic branch is optimized with $\mathcal{L}_s^{(i,j)}$ (Eq. 3) over the whole image sized $W \times H$ as:

$$\begin{aligned} \mathcal{L}_s &= \frac{1}{WH} \sum_{i,j} -E[\ln \mathcal{X}_{Y^{(i,j)}}^{(i,j)}] \\ &= \frac{1}{WH} \sum_{i,j} \psi \left(\sum_{k=1}^K \alpha_{(i,j),k} \right) - \psi(\alpha_{Y^{(i,j)}}) \end{aligned} \quad (4)$$

As in [13], the detection branch is trained with an L2 loss between predicted \hat{C} and ground truth C center heatmaps:

$$\mathcal{L}_o = \frac{1}{WH} \sum_{i,j} \left(\hat{C}^{(i,j)} - C^{(i,j)} \right)^2 \quad (5)$$

For *stuff*, we use the predicted Ω_{st} as a pseudo label to learn the prototypes Ω . For *things*, the same is done with Ω_{th} at the true instance centers. The prototype loss \mathcal{L}_p is the cross-entropy on the *softmax* of the association scores $\hat{Y}^{(i,j),\omega}$, as $\hat{Z}^{(i,j),\omega} = \exp(\hat{Y}^{(i,j),\omega}) / \sum_{\omega' \in \Omega} \exp(\hat{Y}^{(i,j),\omega'})$, with $\omega_{(i,j)}$ being the pseudo label prototype:

$$\mathcal{L}_p = \frac{1}{WH} \sum_{i,j} -\log(\hat{Z}^{(i,j),\omega_{(i,j)}}) \quad (6)$$

We learn embeddings $\phi_{(x,y)}$ with a discriminative loss [16] \mathcal{L}_d (Appendix). The overall training objective is:

$$\mathcal{L} = \lambda_1 \mathcal{L}_s + \lambda_2 \mathcal{L}_o + \lambda_3 \mathcal{L}_p + \lambda_4 \mathcal{L}_d \quad (7)$$

5. Experiments and Results

5.1. Experimental Setup

Datasets We conducted our experiments on three public datasets, namely Cityscapes [15], Lost&Found [54], and MS COCO [47]. **Cityscapes** is a popular outdoor benchmark. Recorded around 50 different cities, mainly in Germany, it contains 19 classes: 8 *things* and 11 *stuff*. We followed the standard split, with 2975 images for training and 500 as validation set, reporting all metrics on the latter. Also recorded in Germany, the **Lost&Found** dataset contains a variety of unusual OOD objects placed in the middle of the road. We selected it because: 1) it was recorded with the same sensor setup as Cityscapes, allowing seamless transfers and removing the need for fine-tuning; 2) it contains only real images; and 3) unlike similar datasets [3, 8], it provides instance annotations for unknowns. Therefore, it is a challenging complement to Cityscapes for holistic segmentation. We did not train on Lost&Found, but used it only to evaluate models trained on Cityscapes. We report all metrics on the *unknown* class of its 1202 test samples. **MS COCO** is a challenging large-scale benchmark for general image understanding, as it includes a variety of scenarios from indoor to outdoor. The 2017 panoptic split contains 80 *thing* categories, and 53 *stuff* classes. We followed EOPSN [38] by treating as unknown the least frequent 20% *thing* classes (e.g., *bear*, *frisbee*). However, instead of turning their segments into *void* and keeping their images in the training set as in [38], we removed their samples completely and regarded them as unseen unknowns. This reduced the training samples to 98112, with 117 classes to learn. We report on the 827 validation samples with unseen classes.

Evaluation metrics We evaluated the panoptic quality (PQ) metric [43] separately for known classes and unknowns, including recognition (RQ) and segmentation (SQ) qualities. We report PQ on the held-out classes of COCO [47], the unknown class of Lost&Found [54], as well as on the 19 known classes of Cityscapes [15] for both **open and closed** settings. Specifically, in open cases, models detect both knowns and unknowns, while in closed set-

Method	Assumptions	Lost&Found (<i>unseen</i>)		
		PQ	RQ	SQ
EOPSN [38]	data, <i>void</i>	0*	0*	0*
OSIS [68]	data, <i>void</i>	1.45	2.23	65.11
U3HS [ours]	none	7.94	12.37	64.24

Table 1. Segmentation of unseen unknown objects (*unknown* class) of **Lost&Found** [54] test set after training on Cityscapes [15] and transferring with no fine-tuning. *: EOPSN diverged (null TP).

tings, the same models predict only knowns, which in practice means ignoring the uncertainty estimates. By analyzing both, we explore the trade-off between detecting unknowns (open) and the in-domain performance (closed).

Network architecture All our models share the structure with Panoptic-DeepLab [13], using a ResNet50 [33] backbone and decoders following Deep-LabV3+ [12]. ResNet50 was chosen to increase reproducibility with limited resources. As described in Section 4.2, the only modification to the semantic decoder is applying the *soft-plus* activation to quantify the uncertainty. The other branches follow Panoptic-DeepLab for detecting centers and DeepLabV3+ for the embeddings, with two heads.

Implementation details For [15, 54], we used input images sized 1024×512 and batch size 16. For [47], we fed 8 images sized 640×480 . We used the Adam optimizer until convergence, with an initial learning rate of 0.001, which was reduced by 2% at each epoch. We set $t = 3$ for the uncertainty threshold (i.e., 3 times the standard deviation) and $F = 8$ for the embedding size to keep the memory low. We adjusted to the different data distribution of COCO with $t = 1$. The backbone was pre-trained on ImageNet [18]. The losses were weighted $\lambda_1 = \lambda_3 = \lambda_4 = 1$ and $\lambda_2 = 200$ [13].

Prior works We compared our U3HS with open-set panoptic works: OSIS [68], which we adapted from LiDAR point clouds to images, EOPSN and its baselines [38], and DDOSP [74]. Instead of training them directly on the unknown categories being evaluated (as in [38]), we followed their setup [68, 38, 74] by training them with the *void* class as fallback, and applied them to unseen unknowns. All methods followed this setup, except that ours ignored *void*. On COCO, we facilitated other works following the K=5% setting of EOPSN [38], thus turning 4 classes into *void*, so they learned 4 classes less than ours. We repurposed and extended a variety of uncertainty estimators [60, 63, 48] from image classification to semantic segmentation (Appendix). We then extended them to holistic segmentation by incorporating them in our U3HS framework.

5.2. Quantitative Results

Unseen unknowns, L&F Table 1 compares our U3HS with prior approaches when segmenting instances of unseen

Method	Assumptions	COCO (<i>unseen</i>)		
		PQ	RQ	SQ
EOPSN [38]	data, <i>void</i>	0.40	0.50	80.30
<i>void</i> -train [38]	data, <i>void</i>	4.40	5.90	74.80
<i>void</i> -supp. [38]	data, <i>void</i>	4.50	6.00	75.90
DDOSP [74]	data, <i>void</i>	<u>9.30</u>	<u>11.20</u>	82.50
U3HS [ours]	none	9.62	13.20	72.84

Table 2. Segmentation of unseen, unknown objects (20% least frequent, held-out classes) on the validation set of **MS COCO** [47]. All others learn *void* and are taken from [74], added trailing zero.

unknowns from Lost&Found [54]. OSIS [68] was the first to address the more limited open-set panoptic segmentation setting, followed by EOPSN [38]. However, OSIS performance fell short on PQ for unseen unknowns, proving the severe limitation of relying on unknowns at training time. By learning *void*, OSIS achieved the highest SQ, which ignores wrong predictions [43]. Instead, despite numerous attempts, EOPSN [38] did not work: it diverged as soon as the exemplars were mined, obtaining 0 true positives (TP). We attribute this to the inconsistent similarities within the *void* class of Cityscapes, compared to those across existing major classes treated as *void* (e.g., *car* in their setup). This prevented EOPSN from forming meaningful clusters from the proposal features during training [38]. Despite the similar setup to EOPSN [38], OSIS [68] could converge since it does not rely on associating unknowns across images. Our U3HS outperformed OSIS by 5.5 times on PQ.

Unseen unknowns, COCO In Table 2, we show results on the 16 held-out categories of MS COCO [47]. In this case, other works were trained following the K=5% setup of EOPSN [38], where 4 classes were learned as *void* (*car*, *cow*, *pizza*, and *toilet*). While this allows them to learn more meaningful representations of unknowns (as unlabeled), it limits the number of classes they can distinguish semantically, e.g., they cannot identify cars. For the other works, the benefit of expanding the *void* distribution by enforcing the inclusion of a variety of recurring objects (e.g., *pizza* and *cars*) is evident as it allowed EOPSN to converge, although to a low PQ on the unseen objects. DDOSP [74] delivered a PQ similar to ours, albeit requiring to turn some knowns into *void*, learning unknowns via *void* and only 113 classes. Without altering the data nor making any assumption on the training samples (e.g., the presence of unknowns within *void* as in [38, 74, 68]), our U3HS performed the best on the unseen categories, especially on RQ (i.e., ability to form instances of unknowns), while learning the whole set of 117 classes, thereby distinguishing even more classes than the other works. Avoiding data assumptions with respect to unknowns made our U3HS effective at segmenting unseen unknowns across both datasets.

Known-unknown Table 3 reports the performances in-

ID	Method	Lost&Found (<i>unseen</i>)			open Cityscapes			closed Cityscapes		
		PQ	RQ	SQ	PQ	RQ	SQ	PQ	RQ	SQ
-	Panoptic-DeepLab [13]	-	-	-	-	-	-	45.82	57.66	79.46
-	OSIS [68]	1.45	2.23	65.11	39.42	50.20	78.53	39.42	50.20	78.53
A1	[ours] baseline: semantic uncertainty	0.49	0.82	60.16	35.02	44.83	78.10	35.97	46.12	78.00
A2	A1 + relaxed embedding association	3.64	5.27	69.09	42.14	53.46	78.83	43.99	55.98	78.59
A3	A2 + prototype head = U3HS	7.94	12.37	64.24	<u>41.21</u>	<u>51.67</u>	<u>79.77</u>	46.53	58.99	78.87
A4	A3 – reassigning outliers	<u>7.85</u>	<u>12.25</u>	64.11	39.84	49.97	79.75	46.53	58.99	78.87
A5	A4 – majority voting	<u>7.85</u>	<u>12.25</u>	64.11	23.94	30.15	79.41	26.77	33.86	79.06
A6	A4 – semantic embeddings	2.33	3.48	<u>67.01</u>	35.16	43.34	81.13	35.92	44.30	81.07
U1	U3HS [ours] + <i>softmax</i> uncertainty	0.10	0.20	51.45	39.70	50.77	78.20	45.12	56.83	79.40
U2	U3HS [ours] + DUQ [63]	0.56	0.89	62.56	41.68	53.14	78.42	45.90	58.17	78.90
U3	U3HS [ours] + DPN [60]	2.09	3.30	63.43	38.90	49.56	78.49	44.91	56.95	78.85
U4	U3HS [ours] + SNGP [48]	<u>4.65</u>	<u>7.57</u>	61.49	41.02	<u>51.98</u>	<u>78.91</u>	<u>46.23</u>	<u>58.56</u>	<u>78.95</u>
A3	U3HS [ours] + improved DPN [ours]	7.94	12.37	<u>64.24</u>	<u>41.21</u>	<u>51.67</u>	79.77	46.53	58.99	78.87

Table 3. Segmentation comparison of models trained on Cityscapes [15] and transferred to the test set of Lost&Found [54] without fine-tuning. All were trained with the same constraints (e.g., ResNet50 [33], small batch, and image sizes). An ablation study (A1-A6) shows the impact of the main components of U3HS, with A3 being our full approach. A3 is paired with various uncertainty estimators (U1-U4).

domain, under open and closed settings (Section 5.1). Ideally, a method would suffer from no decrease in PQ between the two settings, meaning that its estimates are aligned with the distribution shift between knowns and unknowns. OSIS [68] does not use uncertainty estimation, so it does not have these two operating modes, resulting in identical open and closed-set outputs, as if it had only the open setting (via the prediction of *void*). Conversely, all others suffered from a reasonable decrease when extended to open-set. DUQ [63] had the smallest gap, which could be attributed to its underestimation of the uncertainty, as supported by its low scores on Lost&Found.

Closed-set In Table 3, we also compare our U3HS with Panoptic-DeepLab [13]. For a fair comparison, both approaches and all others were trained with the same backbone, image, and batch sizes (Section 5.1). As these were all smaller than those used in [13] due to the limited resources used, they resulted in a lower PQ than that reported in [13]. Nevertheless, our full approach (A3) achieved a slightly higher PQ on Cityscapes under the same setting. We attribute this to the effectiveness of the instance-aware discriminative embeddings learned by our approach, compared to the offset vectors and grouping used by Panoptic-DeepLab. As the focus is unknowns, experiments with improved training resources are out of the scope of this work.

Uncertainty In Table 3, we compare various uncertainty estimations paired to our U3HS framework (U1-U4, A3). While DUQ [63] and *softmax* underperformed compared to OSIS [68], DPN [60] and SNGP [48] achieved a higher PQ. Nevertheless, our improved DPN paired with our framework outperformed prior methods by a substantial margin (A3). For DPN and SNGP, this can be attributed to the su-

periority of our uncertainty estimates. Compared to OSIS and EOPSN, U3HS’s combination of uncertainty estimation with instance-aware embeddings was more effective than learning *void* when encountering wholly new and unseen objects, such as those found in unconstrained settings (e.g., this transfer to Lost&Found).

Ablation study Table 3 reports an ablation of the main components of our U3HS, showing their benefits for holistic segmentation. Compared to the open-set panoptic OSIS [68], with A1, we reduced assumptions not learning the *void* class, and we added a semantic branch with uncertainty for unknowns, which by itself worsened the performance. However, combining this with a relaxed embedding association (Section 4.1) for *things* and *stuff* improved all metrics (A2). A dedicated prototype head (A3, i.e., full approach) increased them even further, more than doubling the PQ on unknowns (i.e., Lost&Found). Specifically, dedicated heads allow both prototypes and embeddings to be more meaningful and expressive without sacrificing the other. A4 shows the impact of reassigning outliers (Section 4.2). While its effect was limited on unknowns, it was more significant on Cityscapes [15]. Transforming unknown predictions in standard in-domain outputs is relevant only in open settings. A5 shows the effect of majority voting to enforce consistency between the outputs (Section 4.1). This did not affect unknowns since classes are not distinguished among them, but it significantly impacted RQ and PQ on Cityscapes. Finally, A6 shows the importance of learning the embeddings according to their semantic classes. In A6, predictions are made by the dedicated semantic branch without the model learning to distinguish the embeddings semantically. Although this increased SQ, it caused a discrepancy within




Figure 4. Example predictions of U3HS on OOD data from the Lost&Found [54] test set. The model was trained on Cityscapes [15] and transferred to Lost&Found without fine-tuning. Embeddings are projected to RGB via t-SNE [64]. White arrows mark labeled unknowns.

the model outputs, decreasing RQ and PQ.

5.3. Qualitative Results

Figure 4 shows example predictions of the proposed U3HS. The images illustrate the setting difficulty and provide examples of the OOD objects of **Lost&Found** [54]. These are often small and hard to see, hidden in the shade or far away from the camera. As seen in the quantitative results (Section 5.2), our U3HS could distinguish instances of unknowns (e.g., stroller in Figure 1), albeit leaving room for improvement. While unknowns correctly triggered high uncertainty estimates, their necessary filtering (third col.) sometimes left too few pixels, if any, on unknowns, leading to missed predictions. However, this is to be expected without any access to OOD data. Furthermore, without distinguishing between unknown *things* and unknown *stuff*, also structures (e.g., fence in the lower image) were given an ID. Nevertheless, thanks to our learned instance-aware embeddings, these were not further subdivided but formed a single large instance (e.g., blue in the lower output). Separate unusual *stuff* regions had the same effect, e.g., the structures around the trees in the upper image. This proves that instances are not simply created by separating disjoint OOD segments but are formed using the learned embeddings. As shown in Figure 4, the embeddings are closely coupled with

the uncertainty estimates and the outputs.

Figure 5 reports predictions of U3HS on samples of **MS COCO** [47] containing two held-out classes (i.e., *bear* and *frisbee*). Remarkably, U3HS was able to separate both bears and frisbees into individual instances despite their high inter-class similarity and not having accessed any information about them. This is thanks to the uncertainty estimation and instance-aware embeddings of our U3HS.

Data considerations and limitations Lost&Found [54] introduces a significant domain shift from Cityscapes. By placing real OOD objects on the road, the authors had to choose unusual scenarios (Figure 4), causing the whole scenes to be OOD. This leads to high uncertainty estimates also on a few known areas. As we do not use any OOD data, nothing constrains high uncertainty to unknown segments, decreasing PQ. A similar issue occurs in COCO, albeit less severely, thanks to more training data. However, COCO has no dedicated unknown class, so it had to be extracted from the set of known ones. Nevertheless, results show that uncertainty is highly valuable, allowing to leave the settings unconstrained. U3HS would mainly benefit from improvements in uncertainty estimates, embeddings descriptiveness, and their clustering. So, learning-based clustering [25, 26] could be advantageous.

The **Supplementary Material** includes more details on the proposed holistic segmentation setting, U3HS and the baselines, as well as additional results, including the trade-off between in-domain and OOD performances, failure cases and qualitative comparisons.

6. Conclusion

In this paper we introduced holistic segmentation: a new setting addressing completely unseen unknown objects in unconstrained scenarios. Additionally, we presented U3HS: the first solution for this new problem. Thanks to its uncertainty estimation and instance-aware learned embeddings, U3HS identifies and separates instances of completely unseen unknowns without any information about them, while segmenting known regions. Extensive experiments on multiple datasets showed the effectiveness of U3HS.

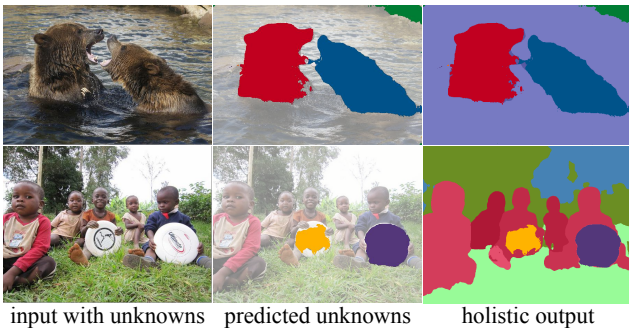


Figure 5. Example predictions of U3HS on OOD data from the COCO [47] validation set. The model had never seen images containing *bear* or *frisbee* (part of the held-out classes), nor had any information about them. Colors represent the predicted instances.

A. Supplementary Material

In this appendix, we include further details and results. Specifically, Sections A.1, A.2 and A.3 provide deeper insights on the proposed setting, the method, and the experimental setup, respectively, while Sections A.4 and A.5 contain more results, both quantitative and qualitative.

A.1. Additional Details on the Setting

In this section, we describe the benefits of the proposed holistic segmentation setting in greater detail, considering both the impact on downstream tasks and the differences with other perception tasks addressing unknown objects.

The proposed holistic segmentation setting aims to segment any unseen, unknown objects without prior knowledge about the unknowns while segmenting known areas. In this context, "unseen unknowns" means any object of any category outside the known classes learned during training, such as the sheep in Figure 6 for a method trained on, e.g., Cityscapes [15], as well as unidentified and distorted parts following a car accident.

A.1.1 Motivation

The importance of identifying unseen unknowns arises from safety-critical scenarios, such as autonomous driving, where ignoring them can lead to dangerous consequences when simply using the predicted segments for downstream tasks, e.g., path planning. This is shown in the top right of Figure 6. Since even large-scale datasets are limited representations of the real world, there will always be corner cases and long tail samples which are problematic for standard models [46]. Therefore, it is crucial to identify these cases and then deal with them safely via downstream tasks.

OOD segmentation [40], i.e., segmenting unknown areas as a whole and not identifying individual instances of unknowns, flags the presence of something unknown in the input. In the context of downstream tasks, such as path planning, a single OOD segment (bottom left in the figure) could trigger an alert state, leading to a potential stop, which is a safe state. However, given that OOD segmentation does not separate unknown objects into instances, once found, it is unclear whether they are moving or static, which means that it would be difficult for path planning. Instead, by segmenting instances of unseen unknowns (bottom right in the figure), holistic segmentation allows tracking unseen objects and estimating their trajectory, leading to a safe path. This motivates the instance segmentation of unknowns, which brings benefits similar to those of instance segmentation compared to semantic segmentation for known objects.

Also critical is the ability to deal with any unseen, unknown object category and not be restricted to a limited subset of them. This is of utmost importance to address the wide variability of objects and scenarios encountered

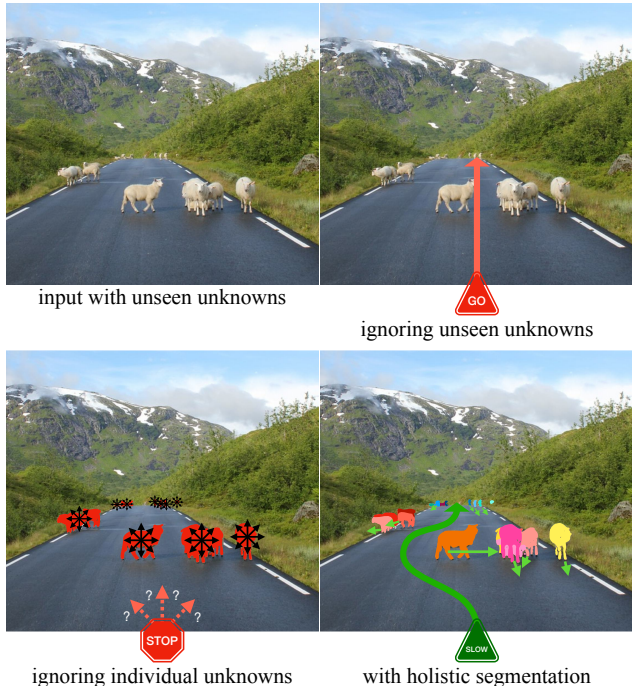


Figure 6. Motivation diagram for identifying unseen, unknown objects on a sample of [8] considering path planning as a downstream task, and hypothesizing that *sheep* is not part of the training data (i.e., unseen unknown). Shown segments are not predictions. State-of-the-art approaches dangerously ignore unknowns (top right) [13]. OOD segmentation does not identify instances of unknowns (bottom left) [40], making it difficult for downstream tasks as the unknowns cannot be tracked, and their trajectory cannot be predicted. The proposed setting (bottom right) identifies individual unseen unknowns, leading to a safe path.

in the real world. While previous settings focused on re-identifying already-seen objects [38, 75], we design holistic segmentation specifically to address any unseen category.

A.1.2 Comparison with Other Settings

As shown in Table 4, compared with other tasks and settings also dealing with unknown objects, the proposed holistic segmentation makes no assumptions about the unknown objects, allowing one to segment any objects. Instead, zero-shot and open-vocabulary approaches assume that text descriptions of unknown objects are available [37, 75]. Open-set panoptic segmentation methods assume unknowns are confined within *void* regions at training and test time [68, 38]. In the latter case, *void* may not be available or not sufficiently large and diverse (as in Cityscapes [15]), depending on the training data. Due to their construction, both of these setups inherently restrict the pool of recognizable objects to those for which text descriptions are available through a vision-language model (open-vocabulary) or to those present within their own training set (open-set panop-

Setting	Data Assumptions	Identifiable Objects	Not Identifiable Objects
open-set panoptic segm. [38, 68]	unknowns are already in the training data, within <i>void</i> areas	known and unlabeled objects present in the training data as <i>void</i>	categories outside of the training data [38]
open-vocabulary, zero-shot [75, 37]	the underlying language model knows about every unknown	objects known by the language model [59]	categories outside of the training data of the language model [59]
holistic segmentation [ours]	none	any known and unknown (i.e., unseen) object	none

Table 4. Comparison of tasks and settings dealing with instances of unknown objects. The second column (Data Assumptions) is related to unknowns. The 2 rightmost columns represent the objects that are theoretically identifiable or not, given the setting.

tic).

For example, for the scene in Figure 6, because of its setup, EOPSN [38] cannot identify any *sheep* unless a vast amount of images containing *sheep* is part of its training data (with *sheep* being labeled as *void*, or directly as a dedicated class *sheep*). Open-vocabulary methods would rely on the fact that a language model [59] already knows about *sheep* to be able to identify them in the image. While the concept of *sheep* is relatively simple and could be assumed to be known by a large language model, there is no guarantee that such a model would know about every possible object and scene that can be encountered in real life (e.g., unidentified pieces on the road following a car accident), meaning that open-vocabulary approaches cannot deal with long tail samples from the distribution of the natural world, simply because their language model cannot process them.

Again, given that datasets include by definition only a fraction of the diversity of the world [46], also datasets to test the ability of a model to identify unknowns are limited [54, 8, 3], containing only a small amount of the possible objects and situations that can be encountered in real life. Therefore, to operate reliably in real unconstrained scenarios, it is of utmost importance not to have limitations on the types of recognizable objects, which should go beyond those found in existing datasets. Instead, relying on a language model to identify unknowns is equivalent to shifting the unknown problem to a different model. As shown with CLIP by Radford et al. [59], large language models also have issues with OOD samples. For example, unidentified broken car parts lying on the ground after an accident would be difficult to describe, so it would be problematic for language models. Thus, when given inputs that are unseen and unknown to the underlying language model, open-vocabulary and zero-shot methods would fail to identify the unknown objects. Furthermore, existing open-set panoptic works rely on the presence of unknowns (intended as unlabeled) directly in the training data through the *void* class. This highlights the need for a new and unconstrained solu-

tion.

For these reasons, the critical differences between the proposed holistic segmentation setting and previous tasks are that holistic segmentation is not constrained in terms of the types of unknown objects that are identifiable and that holistic segmentation does not assume the presence of unknowns in the training data, thereby segmenting **any unseen, unknown object without any prior knowledge about unknowns**. Limited by design by either the unknowns that are known to the underlying language model (e.g., open-vocabulary) or the unknowns that are directly present in the training data (e.g., open-set panoptic segmentation), previous tasks do not enable the identification of any instance of unknowns and rely on prior knowledge about unknowns and their data distribution (e.g., through CLIP [59] or by learning *void*).

A.2. Additional Details on the Method

Loss functions As described in Section 4.3, the proposed method is trained with a combination of losses: a semantic loss \mathcal{L}_s , an object detection loss \mathcal{L}_o , a prototype loss \mathcal{L}_p , and a discriminative loss \mathcal{L}_d . The discriminative loss is aimed at learning meaningful embeddings. It is composed of three different terms [16], namely variance \mathcal{L}_{va} to attract elements towards the mean, distance \mathcal{L}_{di} to push away different groups, and regularization \mathcal{L}_{re} to prevent the divergence of clusters from the origin:

$$\begin{aligned}
 \mathcal{L}_d &= \lambda_{41}\mathcal{L}_{va} + \lambda_{42}\mathcal{L}_{di} + \lambda_{43}\mathcal{L}_{re} \\
 \mathcal{L}_{va} &= \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \frac{1}{N_\omega} \sum_{a=1}^{N_\omega} [\|\mu_\omega - \phi_a\| - \delta_v]_+^2 \\
 \mathcal{L}_{di} &= \frac{1}{|\Omega|(|\Omega|-1)} \sum_{\omega_A \in \Omega} \sum_{\omega_B \in \Omega} [2\delta_d - \|\mu_{\omega_A} - \mu_{\omega_B}\|]_+^2 \\
 \mathcal{L}_{re} &= \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \|\mu_\omega\|
 \end{aligned} \tag{8}$$

where: $|\Omega|$ is the number of prototypes, N_ω is the number of embeddings associated to the prototype ω , μ_ω is the mean embedding of the cluster related to ω , $\|\cdot\|$ is the L2 distance,

$[x]_+ = \max(0, x)$ is the hinge (i.e., until which threshold the terms are active [16]), $\omega_A \neq \omega_B$, and we follow [16] for the hyperparameters, e.g., $\lambda_{41} = \lambda_{42} = 1$ and $\lambda_{43} = 0.001$.

Clustering unseen unknowns As described in Section 4.2, we use DBSCAN [20] to cluster the embeddings of unknown regions into individual unknown objects. Specifically, DBSCAN has multiple advantages: it does not need the number of clusters as input (which is unknown in our case), it is effective and very fast, has a low memory footprint, and distinguishes outliers (Table 3 shows the impact of this feature with A3-A4). Although other traditional clustering methods (e.g., Mean Shift, Affinity Propagation, Birch) are theoretically applicable in our setting, they come with drawbacks (e.g., have high memory requirements, do not output outliers, are significantly slower, or tend to deliver sub-optimal results). On the other hand, popular approaches that require the number of clusters as input cannot be applied in our settings (e.g., K-Means). Hence, DBSCAN was selected.

A.3. Additional Details on the Experimental Setup

Clustering parameters DBSCAN requires two parameters: $minPts$ (number of points in a neighborhood to count as a core point) and ϵ (maximum neighborhood size). To find such parameters, we trained a model (i.e., on Cityscapes [15] or MS COCO [47]), then selected $(minPts, \epsilon)$ with a simple grid search maximizing PQ on a random subset of the known dataset (i.e., without unknowns). Towards this end, we formed instances as follows: ignoring the detection output (i.e., using only the embeddings) and determining their class via majority voting from the semantic output. In particular, when finding $(minPts, \epsilon)$, this means treating the embeddings of knowns as if they were unknowns (apart from their semantic class), assuming that the model treats them similarly. It is essential to consider that the parameters were selected on the known objects (i.e., from Cityscapes or COCO), despite DBSCAN being used only for separating unknowns (i.e., in the Lost&Found [54] dataset or the held out classes of COCO). We did this to maintain the unknowns completely unseen (i.e., only as test set), as in real scenarios.

Comparisons with previous works As described in Section 5.1, we compared our uncertainty-based solution with prior works learning the *void* class and tested various uncertainty estimators within our proposed framework. The following paragraphs further detail how these other methods were trained.

Following our setup of training on Cityscapes [15] and transferring to Lost& Found [54] without any fine-tuning, prior works addressing open-set panoptic segmentation (i.e., OSIS [68] and EOPSN [38]) were trained by learning the *void* class of Cityscapes, unlike our U3HS. This unlabeled class comprises all pixels that do not fulfill the




Figure 7. Trade-off between known (i.e., Cityscapes [15] validation set) and unknown (i.e., Lost&Found [54] test set) performance introduced by OSIS [68], compared to our approach, both using SNGP [48] and our improved DPN (i.e., [ours] full, in blue). The different data points are obtained by varying the parameter t .

requirements to be part of one of the standard 19 annotated classes. Some of these *void* pixels are systematic, e.g., the back side of traffic signs and street lights (excluding poles). By exploiting the variability within *void*, the models learn the extra class decision boundary as a fallback covering anything far from the other classes. To do so, OSIS learns a constant U representing such boundary. In particular, we adapted OSIS from LiDAR point clouds to RGB images, applying it to each pixel instead of point and changing architecture accordingly. As for ours and all other models in this work, we used a ResNet50 [33] as backbone and decoders following the structure of DeepLabV3+ [12]. Moreover, to keep the GPU memory low, we used the same $F = 8$ for the embedding size as in our U3HS. For the experiments on MS COCO [47], we followed the K=5% setup of EOPSN [38], turning 4 classes into *void* to facilitate prior works learning on *void*, by ensuring a diverse distribution of its pixels, as they now cover a diverse set of classes (e.g., *pizza* and *car*).

For the other uncertainty estimation approaches evaluated in this work, we used the authors descriptions and implementations, adapting [63, 48] from image classification to semantic and panoptic segmentation. For DML [5], we used the authors best hyperparameters, therefore a variance loss weight $\gamma_{VL} = 0.01$, and weights $\beta = 20$ and $\gamma = 0.6$. For SML [40], we did not employ the boundary suppression, as it did not improve the results. This might be due to Lost&Found [54] being annotated only for the OOD objects and a coarse road segment. For DUQ [63], we used an embedding dimension of $m = 8$, due to constrained training resources, same as our $F = 8$. Then, we used length scale $\sigma^2 = 0.3$ and exponential smoothing factor $\gamma = 0.999$. For SNGP [48], we again used an embedding dimension $D = 8$

Uncertainty method	L&F (<i>unseen</i>)		open CS mIoU
	AP	FPR ₉₅ ↓	
<i>softmax</i>	16.72	22.88	71.77
MC Dropout [22]	11.22	13.94	68.31
DML [5]	3.14	83.04	69.86
DUQ [63]	5.43	26.64	68.78
DPN [60]	5.43	19.79	66.99
SML [40]	16.91	51.67	<u>70.69</u>
SNGP [48]	<u>22.70</u>	12.02	70.68
improved DPN [ours]	25.44	19.10	70.10

Table 5. Comparison of open-set semantic segmentation on Lost&Found [54] test set of uncertainty estimators based on DeepLabV3+ [12] and trained only on Cityscapes (CS) [15].

(due to the limited training resources), no layer norm for the embeddings, an exponential smoothing factor $\gamma = 0.99$ for updating Σ and 50 samples for Monte Carlo averaging to estimate the uncertainty.

MS COCO As described in the main paper, given that there is no official set of unknown classes for MS COCO [47], we treat as unknown the least frequent 20% known classes. These classes are: *baseball bat, bear, fire hydrant, frisbee, hair drier, hot dog, keyboard, microwave, mouse, parking meter, refrigerator, scissors, snowboard, stop sign, toaster, and toothbrush*. We held out all training samples where any of these 16 classes appeared, such that they were completely unseen to the models.

Ablation study for holistic segmentation With reference to Table 3, A1 is our baseline, which was built upon OSIS [68]. As OSIS, A1 included learned instance-aware embeddings, but unlike OSIS, it featured a semantic decoder delivering semantic segmentation and uncertainty estimates based on the semantic output via our improved DPN. Moreover, as for all our models, A1 did not learn the *void* class (unlike OSIS). A2 featured the relaxed score for the embedding association (described in Section 4.1), which lets the variance be indirectly controlled by the final task (i.e., the loss \mathcal{L}_p , Section 4.3). Unlike A1 and A2, which had a shared head between embeddings and prototypes (i.e., as in OSIS), A3 introduced a dedicated prototype head. In practice, this meant having more layers fully dedicated to the embeddings and the prototypes separately instead of sharing the computation until a later stage. Therefore, this allowed for more expressive and purposed features. A4 did not reassign to the known classes the outliers obtained from clustering unknowns via DBSCAN. Therefore, these pixels were kept unknown and shared the same instance ID. A5 did not perform majority voting (Section 4.1). This meant directly assigning the semantic classes predicted by the semantic branch to all known instance pixels instead of enforcing coherence within an instance. This caused the instances to be fragmented according to how many seman-

Configuration			L&F (<i>unseen</i>)		open CS mIoU
Ref.	Activ.F.	KL	AP	FPR ₉₅ ↓	
[60]	<i>exp</i>	yes	5.43	19.89	66.99
[ours]	<i>softplus</i>	yes	3.43	25.97	64.36
[ours]	<i>softplus</i>	no	25.44	19.10	70.10

Table 6. Ablation study on uncertainty estimates for open-set semantic segmentation. Models trained only on Cityscapes [15].

tic classes they contained, decreasing RQ. Finally, A6 predicted the semantic classes for *stuff* areas directly from the semantic prediction branch instead of matching the embeddings with *stuff* prototypes as in A1-A5 (Section 4.1).

A.4. Additional Quantitative Results

Trade-off between known and unknown Figure 7 shows the trade-off between the performance on known and unknown for our framework, both with SNGP [48] and our improved DPN, compared to that of OSIS [68]. The different data points were extracted by evaluating the outputs at different thresholds t , namely [2, ..., 5], and ignoring the uncertainty estimates entirely (i.e., closed-set, reported where PQ Lost&Found is 0). The hyperparameter t directly affects how high the uncertainty estimates must be for their associated pixels to be considered unknown. This has an impact on the performance on open-set Cityscapes [15] and Lost&Found [54], since changing in output what is considered unknown alters what is regarded as in-domain (i.e., known) as well. OSIS [68] does not have such a hyperparameter as it considers unknown everything predicted as *void*. Overall, it can be seen that our proposed framework offers a better trade-off in both configurations (red and blue) than that of OSIS [68] (yellow). Furthermore, using our full approach (i.e., our framework with our improved DPN) typically gave the best trade-off between known and unknown without compromising the metrics too much (blue).

Unknowns in semantic segmentation In Table 5, we compare the ability of a wide variety of uncertainty estimators (i.e., [60, 63, 48, 40, 5], and MC Dropout with 25 runs [22]) to find unknowns in a semantic setting on Lost&Found [54], after training on Cityscapes [15]. This meant retraining all methods under the same conditions while also extending DPN [60], DUQ [63], and SNGP [48] to semantic segmentation. Semantic models (Tables 5 and 6) used smaller crops sized 512×256 compared to the other experiments. For uncertainty estimation, we evaluated the ability to identify unknowns reporting the AP on the unknown class [54], as well as the false positive rate at the recall 95 (FPR₉₅). For semantic segmentation on Cityscapes, we computed the mIoU. As seen in Table 3, DUQ [63] and DPN [60] performed worse than SNGP [48]. MC Dropout [22] underperformed *softmax*, probably due to the contrasting opinions from 25 forward passes. Our

ResNet depth	F	PQ	RQ	SQ
18	2	33.0	42.3	77.9
18	4	38.9	49.8	78.0
18	8	41.3	52.7	78.3
18	16	42.3	53.8	78.7
18	32	<u>42.1</u>	<u>53.6</u>	<u>78.5</u>
50	8	47.7	60.4	79.0

Table 7. Different embedding dimensions F on closed-set panoptic segmentation on the validation set of Cityscapes [15]. The first column indicates the depth of the ResNet [33] backbone used (i.e., 18 for ResNet18).

method was the best at finding unknowns (AP) with high-quality uncertainty estimates (FPR_{95}). Table 5 also reports the mIoU on Cityscapes (CS), showing that all methods introduce a trade-off between OOD and in-domain outputs, as overestimating the uncertainty decreases the in-domain mIoU. Balancing these two complementary aspects is not trivial, with our approach and SNGP managing it best.

Ablation on uncertainty estimation Table 6 compares the DPN [60] we adapted from image classification to semantic segmentation with our extension. Our improvements were oriented to simplify the training process and help convergence. First we applied the *softplus* activation function to the last semantic layer, instead of *exp* as in DPN [60]. We chose *softplus* because it grows slower than *exp* and it is smooth, differentiable everywhere, and monotonic. This significantly improved the training stability at the cost of a reduced quality of the uncertainty estimates. Finally, due to the complexity of modeling the target distribution in our setting, omitting the KL term used by DPN [60] further stabilized training and boosted the performance on all metrics.

Impact of embedding size and architecture Table 7 shows the effect of different embeddings dimensions F on a smaller ResNet18 [33]. In the rest of this work, all experiments used $F = 8$ and ResNet50, as in the last line of the table, due to constrained training resources. The embedding dimension directly affects the learning capability of the model. Since the instance-aware embeddings are a critical part of the output, a smaller F is linked to inexpressive embeddings that cannot be as discriminative as those from a larger F . Therefore, increasing F improved all metrics except for the larger $F = 32$. This can be attributed to the small ResNet18 backbone being already saturated at $F = 16$, unable to extract rich and detailed features for the larger embeddings to exploit. With a larger model (e.g., ResNet101), even higher embedding dimensions F might be beneficial. Table 7 shows that our proposed approach, given less constrained resources, could deliver better results when using an embedding dimension higher than the $F = 8$ employed across this work. The table also shows the comparison between ResNet18 and ResNet50, with the latter

Method	Clustering	PQ	RQ	SQ
U3HS [ours]	Mean Shift	2.71	3.91	69.22
U3HS [ours]	DBSCAN	9.36	14.83	63.14

Table 8. Transfer from Cityscapes [15] to Lost&Found-300 [54] test set (i.e., on the first 300 samples, see Section A.4). DBSCAN (our choice) is compared to Mean Shift to cluster the embeddings of unknown areas.

delivering over 15% higher PQ at the same $F = 8$. This shows how our proposed approach would perform with a larger backbone.

Impact of the clustering method In Table 8 we compare two popular clustering methods within our U3HS framework, namely DBSCAN and Mean Shift [14]. Due to the very high computation effort and memory required by Mean Shift, we opted for the following setup for this experiment. First, instead of a standard CPU implementation, we used a parallelized CUDA version of the algorithm [77]. Then, due to the still very high memory requirements, specific samples of Lost&Found caused memory issues. Therefore, we reduced the size of the test set of Lost&Found [54] to its first 300 samples (Lost&Found-300), which were not problematic. These 300 samples are sufficient to indicate the effect of using Mean Shift instead of DBSCAN. Table 8 shows the superiority of DBSCAN for this setting, with a 3.5x higher PQ and 3.8x better RQ. In particular, RQ should be the focus as we compare instance segmentation of unknowns.

Additional details on EOPSN vs. OSIS As described in Section 5.2, EOPSN always diverged on Cityscapes despite numerous attempts, leading to null true positives, as shown in Table 1. OSIS did not suffer from this issue: EOPSN’s mining strategy requires associating similar “unknowns” across different inputs, but OSIS operates frame-by-frame. Since *void* pixels are unstructured and undefined in Cityscapes, EOPSN’s association fails. As this process is a fundamental step of its training procedure, it makes EOPSN diverge and leads to null scores due to the absence of true positives. Instead, in EOPSN’s setup (i.e., re-identifying unlabeled objects seen during training), such associations can be made across the instances of the classes EOPSN’s authors treat as *void* or “unknown” (e.g., all cars in their setup, as in Table 2). Therefore, on MS COCO (Table 2), EOPSN managed to identify a few true positives, thereby scoring more than 0 for PQ and RQ and significantly more for SQ as it considers only the IoU of matched segments (TP) and not the wrong predictions (FP and FN).

Varying number of unknowns Both EOPSN and DDOSP showed that their performance drops across the board by increasing the amount of *void* classes to detect (i.e., K , pseudo-unknowns), especially for unknowns. As shown in Table 1, they perform poorly also with $K = 0$.

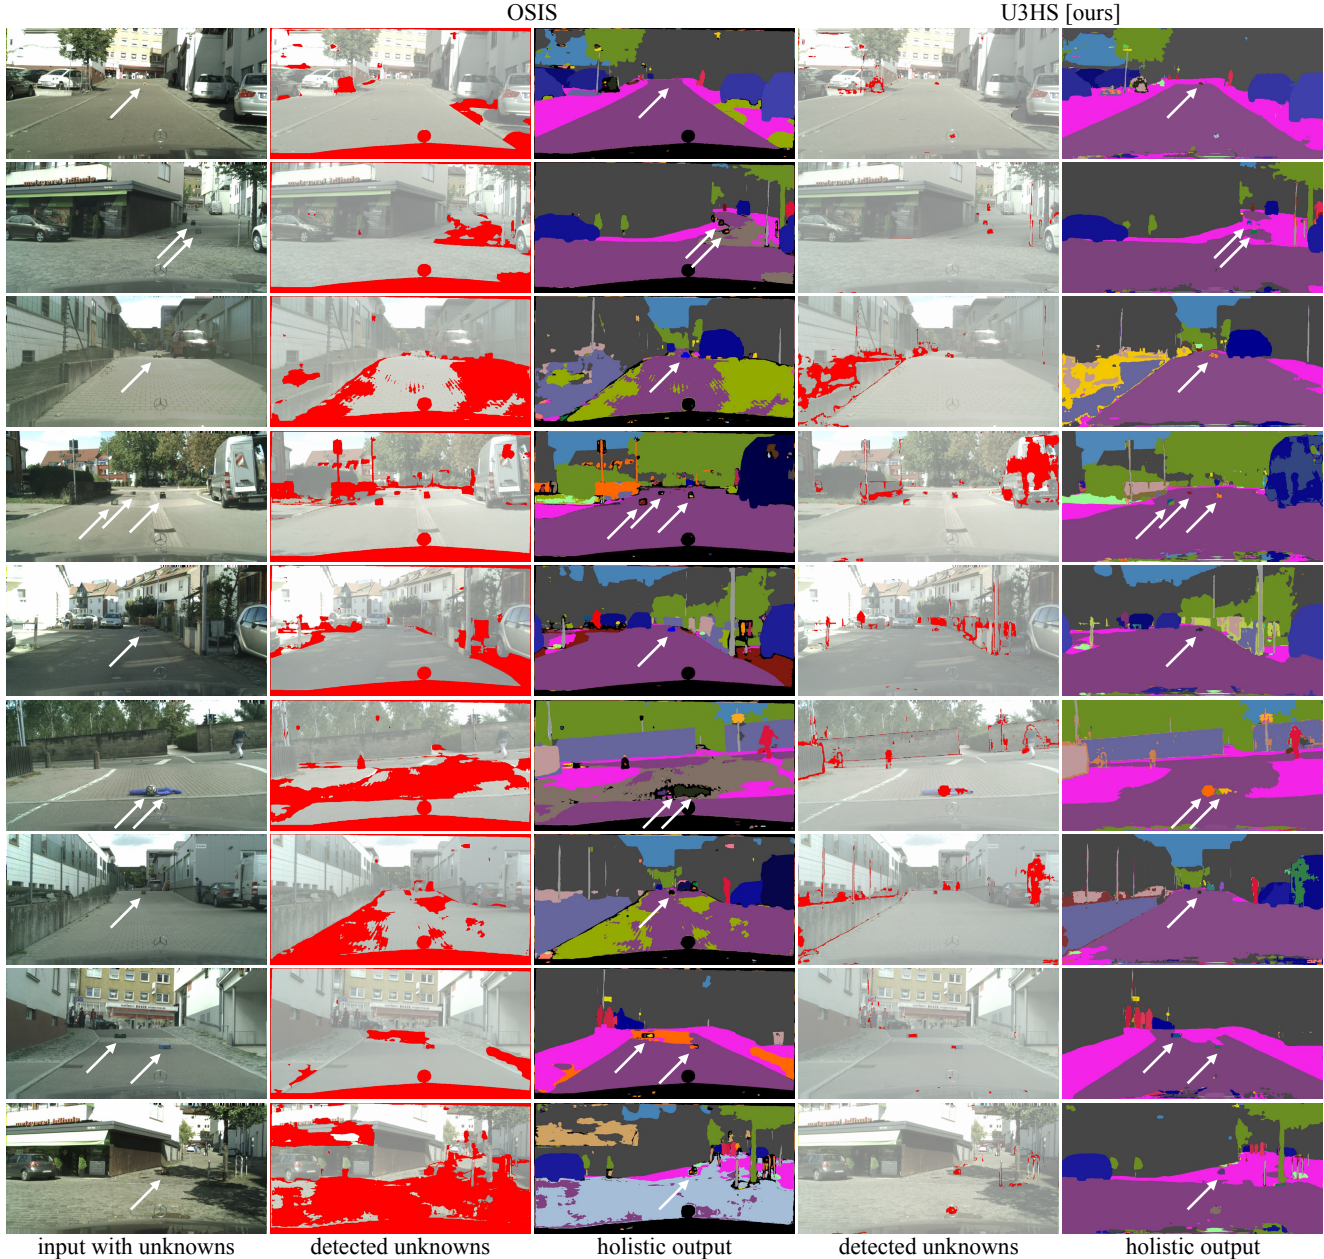


Figure 8. Example predictions of OSIS [68] and the proposed U3HS on unknown categories from the test set of Lost&Found [54]. The models were trained on Cityscapes [15] and transferred to Lost&Found without any fine-tuning. OSIS found unknowns as the *void* class (learned during training), while our U3HS discovered them via uncertainty estimation. Black regions in OSIS’s outputs, including around unknowns, represent pixels predicted as part of the unknown instance of the ego vehicle bonnet: since the bonnet is labeled as *void* in the training set, OSIS learned it as such and it turned it into an unknown instance at inference time. White arrows mark labeled OOD objects.

Instead, our U3HS does not rely on *void*, so it is unaffected by K or what is assigned to *void*. This means that U3HS’s performance varies only slightly with different K s, as for random initialization. Moreover, whether turning known classes into *void* (MS COCO, Table 2, with $K = 5\%$) or not (Lost&Found, Table 1, $K = 0$), our method outperforms prior works despite letting the others learn from unknowns via *void*. Figure 8 shows this qualitatively. Thanks

to uncertainty estimation, our setup has an edge with unseen categories (e.g., Table 1). Increasing K for open-set works (i.e., treating more classes as *void*) means reducing the number of classes that can be segmented semantically. Therefore, the proposed setting is more practical than open-set panoptic and open-vocabulary because, for ours, no unknowns need to be part of the training of any model (simpler data collection), and ours detects any unseen categories.

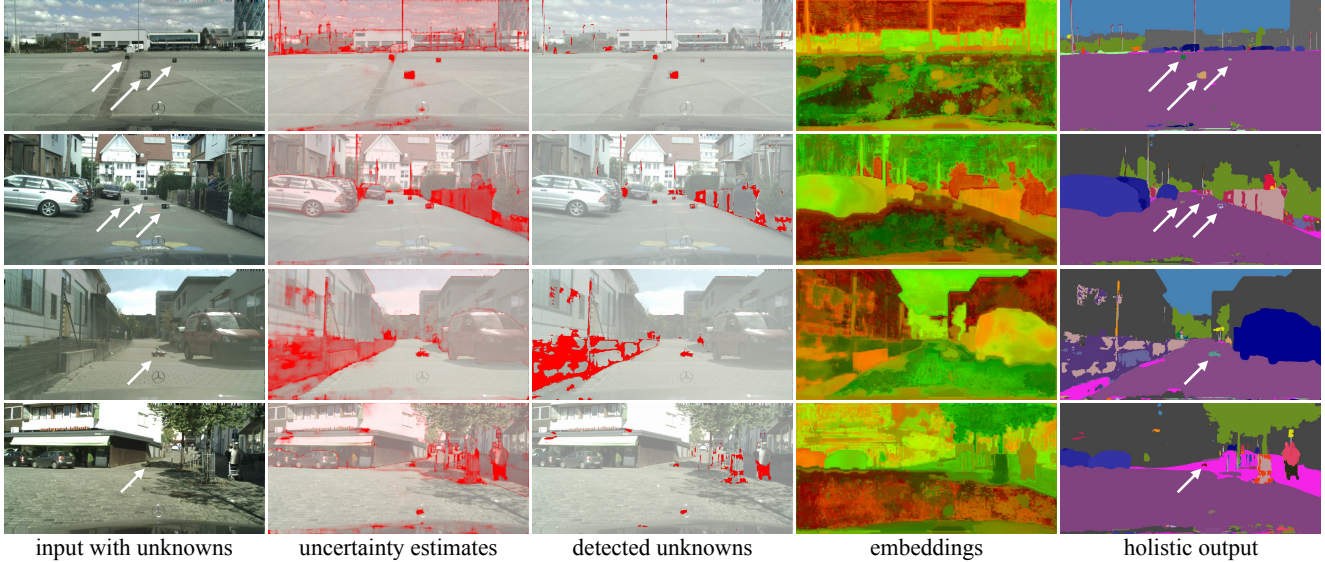


Figure 9. Additional example predictions of the proposed U3HS on unknown categories from the test set of Lost&Found [54]. The model was trained on Cityscapes [15] and transferred to Lost&Found without any fine-tuning. White arrows mark labeled OOD objects.

A.5. Additional Qualitative Results

A.5.1 Qualitative Comparison

Figure 8 shows a comparison of the predictions of the proposed U3HS with the prior work OSIS [68], as well as the regions each predicted as unknown, on a set of samples from Lost&Found [54]. In particular, ours found unknowns as segments estimated as highly uncertain, and OSIS found them as the pixels predicted to be part of the learned *void* class.

From the images, it can be seen how for the most part, OSIS managed to learn a relatively good class boundary around the *void* class, as it was typically able to predict the OOD objects as unknown via *void*. This is interesting as it shows how OSIS can potentially work with challenging unseen unknowns. However, the same figure also shows the substantial limitations of learning and predicting *void* due to the assumptions about the data distributions this entails. In the first image, OSIS completely ignored the unknown object, assigning it to the *road* class, while in the fifth image, it detected the toy as *car*. In contrast, in the last picture, OSIS predicted almost everything as unknown. This proves how the binary aspect introduced by predicting the *void* class (a pixel is either unknown, by being *void*, or known, if another class) does not cope well with the diversity and unpredictability of the scenes in unconstrained real-world settings. Specifically, predicting the *void* class severely relies on the closed-set training data, as the success of such a method is directly related to the diversity of the *void* class seen during training, which is limited as it cannot correctly sample the long tail of the data distribution [46].

Nevertheless, as shown already in Section 5, estimating the uncertainty allows to properly cope with unknown objects by adding an extra layer of prediction. Contrary to the idea of prior works (Section 2) of predicting unknowns via the *void* class, which directly competes with the other semantic classes for being part of the output, uncertainty estimates go on top of the standard semantic predictions. Although this complicates dealing with multiple network outputs, it offers a wider spectrum and deeper insights since the uncertainty could be ignored or considered with various thresholds depending on the situation (Figure 7), for the same trained model and output. Since estimating the uncertainty aims at smoothly quantifying the domain gap from the training data, we believe it is better suited to highly unpredictable unseen real-world scenarios as in holistic segmentation settings.

Furthermore, Figure 8 shows the capability of each method to identify instances of unknown OOD objects. For both approaches, this is related to the clustering of embeddings corresponding to those pixels predicted as unknown, via *void* (OSIS), or as highly uncertain (ours). In particular, OSIS tended to over-fragment unknown objects into several small instances, as seen in the fourth, sixth, eighth, and last images. This again proves our modifications' effectiveness when dealing with the embeddings, as described in Section 4 and evaluated in Table 3. Additionally, OSIS could not distinguish the two neighboring OOD objects in the sixth image. Moreover, OSIS often improperly assigned large regions to the same unknown instance. Similarly to ours, OSIS considers every unknown segment as part of an instance. By learning and predicting the *void* class, during training, OSIS learned to precisely segment the bonnet of





Figure 10. Example predictions of the proposed U3HS on OOD data containing the held out classes of MS COCO [47]. Held out classes (unseen unknowns) in the samples: *frisbee*, *microwave*, *toothbrush*, *baseball bat*, and *bear*. All samples are equally resized. Input, embeddings, detected unknowns, and holistic output are shown.

the ego car (labeled as *void* in Cityscapes [15]). However, at test time on Lost&Found it could not tell the ego vehicle bonnet apart from a wide variety of pixels. This was the case for the unknown object in the seventh image, which was entirely assigned to the same instance as the bonnet or many other segments around knowns and unknowns (colored in black). The ego vehicle bonnet unknown instance (black) often surrounded other predicted unknown instances (e.g., in the second, fourth, sixth, eighth, and last images).

A benefit of estimating the uncertainty is the ability to account for a wide array of unusual regions. This is valuable for downstream tasks, e.g., trajectory prediction and path planning. Specifically, uncertainty estimates by the proposed U3HS were high on the stroller in Figure 1, as well as in Figure 4 on the walking assistance device on the left of the upper image and the cart pushed by the man waving on the right of the bottom image in Figure 9, none of which were labeled as unknown in the dataset [54], as they



input with unknowns

detected unknowns

input with unknowns


detected unknowns

Figure 11. Example predictions of the proposed U3HS on OOD data containing the held out classes of MS COCO [47]. Held out classes (unseen unknowns) in the samples: *microwave*, *parking meter*, *stop sign*, *baseball bat*, *toothbrush*, *toaster*, *snowboard*, *fire hydrant*, and *frisbee*. Other unknown objects are included in the samples, such as the umbrella and the rice cooker (i.e., not part of the known classes). All samples are equally resized. Input and detected unknowns are shown.

were not part of the objects manually placed by the authors. In Figure 8, this is repeated from a different perspective on the stroller in the background of the second image, the unusual van with the open doors in the fourth image, and the duffel bag in the sixth. By learning and predicting *void*, OSIS ignored these unusual regions as it lacks the flexibility and granularity that our U3HS offers by estimating the uncertainty.

A.5.2 Additional Results on Unknowns

Lost&Found Figure 9 shows additional qualitative outputs. Once again, it can be seen how challenging the proposed holistic segmentation setting is. As in the predictions of Figure 4, the model can distinguish most unknown objects. It can be seen how specific areas of the images trigger higher uncertainty estimates. This is the case of the fences in the



input with unknowns

detected unknowns

input with unknowns

detected unknowns

Figure 12. Failure predictions of the proposed U3HS on OOD data containing the held out classes of MS COCO [47]. Held out classes (unseen unknowns) in the samples: *frisbee*, *keyboard*, *mouse*, *refrigerator*, *hot dog*, *toothbrush*, and *scissors*. Other unknown objects are included in the samples, such as bag and comb (i.e., not part of the known classes). All samples are equally resized. Input and detected unknowns are shown.




Figure 13. Failure predictions of the proposed U3HS on OOD data from the test set of Lost&Found [54], with a transfer from Cityscapes [15]. White arrows mark missed OOD objects as the estimated uncertainty was relatively low and filtered out.

second and third images of Figure 9, as well as unknown objects not part of the OOD labels of Lost&Found, such as the cart on the right of the bottom image, as previously mentioned. As previously seen, *stuff* structures (e.g., fences) are assigned to a single coherent instance ID throughout the whole image. At the same time, unusual objects (e.g., the cart in the last picture) have their dedicated ID. Figure 9 also provides some examples of unusual scenes present in the Lost&Found [54] dataset, posing significant challenges compared to Cityscapes [15].

MS COCO Figures 10 and 11 show qualitative outputs of the proposed U3HS on the held out classes of MS COCO [47]. The images report the vast diversity of the dataset, ranging from outdoor scenes to indoor close-ups. Remarkably, U3HS delivered precise segments for unknown objects, correctly segmenting instances of individual unknowns despite their similarity with other objects of the same type in the same input, with reasonable estimations of known classes too (Figure 10). Due to the difficulty of this problem, only a handful of segments are perfect, leaving room for improvements for known and unknown objects. Thanks to its strong uncertainty estimation capabilities, the proposed U3HS not only identified the held-out classes but also other unknown objects which are not part of the set of known classes, such as the rice cooker and the umbrella on the right of Figure 11 in the third and bottom rows respectively. This shows the efficacy of our method on a wide variety of scenarios typical of the real world.

A.5.3 Failure Cases

MS COCO While the proposed U3HS delivers reasonable estimates in various settings, from indoor to outdoor, the problem at hand is highly challenging, and its predictions are not perfect, as confirmed by the quantitative results. Figure 12 reports failure cases on a set of challenging samples of MS COCO. It can be seen that while U3HS identified unknowns reasonably, it often missed those objects that are part of the evaluated held-out categories. A series of issues cause this. In some instances, the integrated uncertainty estimates could not fully discover the unseen unknowns, e.g.,

only partially detected for both *refrigerators* at the left of the third and fourth rows. We also noticed systematic issues with certain classes, such as *keyboard*, *mouse* (both in the right of the fourth row), *hot dog* (in the fifth row), and *scissors* (right of the last row). This could be attributed to these objects being semantically relatively close to known classes, such as *sandwich* for *hot dog*. Small objects were hard to see and therefore ignored by our method. This is the case of the *toothbrush* behind the *cat* in the last row of the figure. U3HS also had difficulties telling apart from one another very close and similar unknowns, e.g., the central *frisbees* in the right of the second row differing for the logo’s color. However, it could successfully separate neighboring objects on multiple occasions, such as the containers on the ground of the left image in the fourth row (not evaluated in this experiment). Moreover, we noticed difficulties with particularly unusual inputs and cluttered environments, e.g., on the left of the second row. With other inputs, the uncertainty of U3HS was also triggered on known objects, such as the *dogs* on the top right (albeit correctly separated into two instances), or the *cabinets* on the right of the third row. Since several held-out classes contained everyday kitchen-related items (e.g., *refrigerator*, *microwave*, and *toaster*), or typical desk objects (e.g., *keyboard* and *mouse*), the model could see only a handful of kitchens and offices (i.e., those images where none of these held out objects appeared), which severely impacted its ability to handle these situations appropriately. This is a limitation of holding out samples from MS COCO, compared to using a dedicated separate dataset, such as Lost&Found [54]. Furthermore, given that U3HS estimates the model uncertainty, more training data covering a wider variety of scenarios could be beneficial to further reduce the uncertainty on the known classes and improve the known-unknown boundary of U3HS.

Lost&Found Figure 13 shows failure cases caused by the necessary filtering of the uncertainty estimates. While the uncertainty was triggered by a variety of unusual areas, including the vast majority of unknown objects, its a priori filtering (based on closed-set training data, Section 4.2) sometimes caused the unknown object to be completely undetected. Although this filtering is aimed at removing low

uncertainty areas which are probably in-domain (e.g., the fence in the upper image), it could inadvertently remove proper OOD objects (e.g., those marked by the white arrows). This is related to the trade-off shown in Figure 7, so keeping more unknowns (i.e., lower threshold t) reduces the in-domain performance. Nevertheless, in the embeddings visualizations, the model correctly isolated the entire marked box in the lower image and precisely segmented the cardboard box in the upper one. However, the two unknown objects were not detected, due to the difficulty of merging multiple outputs and interpreting uncertainty estimates without access to OOD data. It should be considered that the proposed U3HS does not distinguish between the uncertainty for unknown objects and that of unusual known classes. The difference might lie in the amount of uncertainty corresponding to these regions, hence the filtering via the threshold t to attempt telling apart completely unknown from unusual, which remains highly challenging without using any information about unknowns at training time, as in our setup.

References

- [1] Sara Beery, Yang Liu, Dan Morris, Jim Piavis, Ashish Kapoor, Neel Joshi, Markus Meister, and Pietro Perona. Synthetic examples improve generalization for rare classes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 863–873, 2020. [1](#)
- [2] Abhijit Bendale and Terrance Boulton. Towards open world recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1893–1902, 2015. [3](#)
- [3] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. The Fishyscapes benchmark: Measuring blind spots in semantic segmentation. *Springer International Journal of Computer Vision*, 129(11):3119–3135, 2021. [1](#), [5](#), [10](#)
- [4] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [5] Jun Cen, Peng Yun, Junhao Cai, Michael Yu Wang, and Ming Liu. Deep metric learning for open world semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15333–15342, 2021. [2](#), [3](#), [11](#), [12](#)
- [6] Jun Cen, Peng Yun, Junhao Cai, Michael Yu Wang, and Ming Liu. Open-set 3D object detection. In *International Conference on 3D Vision (3DV)*, pages 869–878. IEEE, 2021. [1](#), [3](#)
- [7] Jun Cen, Peng Yun, Shiwei Zhang, Junhao Cai, Di Luan, Mingqian Tang, Ming Liu, and Michael Yu Wang. Open-world semantic segmentation for LiDAR point clouds. In *Proceedings of the European Conference on Computer Vision*, pages 318–334. Springer, 2022. [3](#)
- [8] Robin Chan, Krzysztof Lis, Svenja Uhlemeyer, Hermann Blum, Sina Honari, Roland Siegwart, Pascal Fua, Mathieu Salzmann, and Matthias Rottmann. SegmentMeIfYouCan: A benchmark for anomaly segmentation. In *Neural Information Processing Systems - Datasets and Benchmarks Track*, 2021. [1](#), [5](#), [9](#), [10](#)
- [9] Wei-Lun Chao, Soravit Changpinyo, Boqing Gong, and Fei Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *Proceedings of the European Conference on Computer Vision*, pages 52–68. Springer, 2016. [2](#), [3](#)
- [10] Jiaoyan Chen, Yuxia Geng, Zhuo Chen, Ian Horrocks, Jeff Z. Pan, and Huajun Chen. Knowledge-aware zero-shot learning: Survey and perspective. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 4366–4373, 2021. [2](#)
- [11] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D Collins, Ekin D Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 695–714. Springer, 2020. [2](#)
- [12] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 801–818. Springer, 2018. [6](#), [11](#), [12](#)
- [13] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-DeepLab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12475–12485, 2020. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [9](#)
- [14] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002. [13](#)
- [15] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016. [5](#), [6](#), [7](#), [8](#), [9](#), [11](#), [12](#), [13](#), [14](#), [15](#), [16](#), [19](#)
- [16] Bert De Brabandere, Davy Neven, and Luc Van Gool. Semantic instance segmentation with a discriminative loss function. *arXiv preprint arXiv:1708.02551*, 2017. [5](#), [10](#), [11](#)
- [17] Daan de Geus, Panagiotis Meletis, Chenyang Lu, Xiaoxiao Wen, and Gijs Dubbelman. Part-aware panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5485–5494, 2021. [2](#)
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [6](#)
- [19] Zheng Ding, Jieke Wang, and Zhuowen Tu. Open-vocabulary panoptic segmentation with MaskCLIP. *arXiv preprint arXiv:2208.08984*, 2022. [2](#)

- [20] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231, 1996. [5](#), [11](#)
- [21] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2950–2958, 2019. [1](#)
- [22] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016. [12](#)
- [23] Stefano Gasperini, Jan Haug, Mohammad-Ali Nikouei Mahani, Alvaro Marcos-Ramiro, Nassir Navab, Benjamin Busam, and Federico Tombari. CertainNet: Sampling-free uncertainty estimation for object detection. *IEEE Robotics and Automation Letters*, 7(2):698–705, 2021. [2](#), [3](#)
- [24] Stefano Gasperini, Patrick Koch, Vinzenz Dallabetta, Nassir Navab, Benjamin Busam, and Federico Tombari. R4Dyn: Exploring radar for self-supervised monocular depth estimation of dynamic scenes. In *International Conference on 3D Vision (3DV)*, pages 751–760. IEEE, 2021. [1](#)
- [25] Stefano Gasperini, Mohammad-Ali Nikouei Mahani, Alvaro Marcos-Ramiro, Nassir Navab, and Federico Tombari. Panoster: End-to-end panoptic segmentation of LiDAR point clouds. *IEEE Robotics and Automation Letters*, 6(2):3216–3223, 2021. [2](#), [8](#)
- [26] Stefano Gasperini, Magdalini Paschali, Carsten Hopke, David Wittmann, and Nassir Navab. Signal clustering with class-independent segmentation. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3982–3986. IEEE, 2020. [8](#)
- [27] Jakob Gawlikowski, Cedric Rovile Njjeutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342*, 2021. [1](#), [2](#), [3](#)
- [28] Yuxia Geng, Jiaoyan Chen, Xiang Zhuang, Zhuo Chen, Jeff Z Pan, Juan Li, Zonggang Yuan, and Huajun Chen. Benchmarking knowledge-driven zero-shot learning. *Elsevier Journal of Web Semantics*, 75:100757, 2023. [2](#)
- [29] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *Proceedings of the European Conference on Computer Vision*, pages 540–557. Springer, 2022. [2](#)
- [30] Kratarth Goel, Praveen Srinivasan, Sarah Tariq, and James Philbin. QuadroNet: Multi-task learning for real-time semantic depth aware instance segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 315–324, 2021. [2](#)
- [31] Matej Grcić, Petra Bevandić, and Siniša Šegvić. Densehybrid: Hybrid anomaly detection for dense open-set recognition. In *Proceedings of the European Conference on Computer Vision*, pages 500–517. Springer, 2022. [3](#)
- [32] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2961–2969, 2017. [2](#)
- [33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. [6](#), [7](#), [11](#), [13](#)
- [34] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. [1](#)
- [35] Jie Hong, Weihao Li, Junlin Han, Jiyang Zheng, Pengfei Fang, Mehrtash Harandi, and Lars Petersson. GOSS: Towards generalized open-set semantic segmentation. *Springer The Visual Computer*, pages 1–14, 2023. [3](#)
- [36] Rui Hou, Jie Li, Arjun Bhargava, Allan Raventos, Vitor Guizilini, Chao Fang, Jerome Lynch, and Adrien Gaidon. Real-time panoptic segmentation from dense detections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8523–8532, 2020. [2](#)
- [37] Dat Huynh, Jason Kuen, Zhe Lin, Jiuxiang Gu, and Ehsan Elhamifar. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7020–7031, 2022. [2](#), [9](#), [10](#)
- [38] Jaedong Hwang, Seoung Wug Oh, Joon-Young Lee, and Bohyung Han. Exemplar-based open-set panoptic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1175–1184, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [9](#), [10](#), [11](#)
- [39] Taejong Joo, Uijung Chung, and Min-Gwan Seo. Being Bayesian about categorical probability. In *International Conference on Machine Learning*, pages 4950–4961. PMLR, 2020. [4](#)
- [40] Sanghun Jung, Jungsoo Lee, Daehoon Gwak, Sungha Choi, and Jaegul Choo. Standardized Max Logits: A simple yet effective approach for identifying unexpected road obstacles in urban-scene segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15425–15434, 2021. [2](#), [3](#), [9](#), [11](#), [12](#)
- [41] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? *Advances in Neural Information Processing Systems*, 30, 2017. [1](#), [2](#)
- [42] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019. [2](#), [3](#)
- [43] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. [2](#), [3](#), [5](#), [6](#)
- [44] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. [2](#)
- [45] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas J Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic Neural

- Fields: A semantic object-aware neural scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12871–12881, 2022. [2](#)
- [46] Alexander Lehner, Stefano Gasperini, Alvaro Marcos-Ramiro, Michael Schmidt, Mohammad-Ali Nikouei Mahani, Nassir Navab, Benjamin Busam, and Federico Tombari. 3D-VField: Adversarial augmentation of point clouds for domain generalization in 3D object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17295–17304, 2022. [1](#), [2](#), [3](#), [9](#), [10](#), [15](#)
- [47] 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 *Proceedings of the European Conference on Computer Vision*, pages 740–755. Springer, 2014. [5](#), [6](#), [8](#), [11](#), [12](#), [16](#), [17](#), [18](#), [19](#)
- [48] Jeremiah Z. Liu, Zi Lin, Shreyas Padhy, Dustin Tran, Tania Bedrax-Weiss, and Balaji Lakshminarayanan. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. In *Advances in Neural Information Processing Systems*, 2020. [2](#), [4](#), [6](#), [7](#), [11](#), [12](#)
- [49] Jiaxu Miao, Xiaohan Wang, Yu Wu, Wei Li, Xu Zhang, Yunchao Wei, and Yi Yang. Large-scale video panoptic segmentation in the wild: A benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21033–21043, 2022. [2](#)
- [50] Dimity Miller, Feras Dayoub, Michael Milford, and Niko Sünderhauf. Evaluating merging strategies for sampling-based uncertainty techniques in object detection. In *Proceedings of the International Conference on Robotics and Automation*, pages 2348–2354. IEEE, 2019. [2](#), [3](#)
- [51] Dimity Miller, Lachlan Nicholson, Feras Dayoub, and Niko Sünderhauf. Dropout sampling for robust object detection in open-set conditions. In *Proceedings of the International Conference on Robotics and Automation*, pages 3243–3249. IEEE, 2018. [3](#)
- [52] Rohit Mohan and Abhinav Valada. EfficientPS: Efficient panoptic segmentation. *Springer International Journal of Computer Vision*, 129(5):1551–1579, 2021. [2](#)
- [53] Trung Pham, Thanh-Toan Do, Gustavo Carneiro, Ian Reid, et al. Bayesian semantic instance segmentation in open set world. In *Proceedings of the European Conference on Computer Vision*, pages 3–18. Springer, 2018. [3](#)
- [54] Peter Pinggera, Sebastian Ramos, Stefan Gehrig, Uwe Franke, Carsten Rother, and Rudolf Mester. Lost and Found: Detecting small road hazards for self-driving vehicles. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1099–1106, 2016. [1](#), [5](#), [6](#), [7](#), [8](#), [10](#), [11](#), [12](#), [13](#), [14](#), [15](#), [16](#), [19](#)
- [55] Lorenzo Porzi, Samuel Rota Buló, Aleksander Colovic, and Peter Kotschieder. Seamless scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8277–8286, 2019. [2](#)
- [56] Janis Postels, Hermann Blum, Yannick Strümpfer, Cesar Cadena, Roland Siegwart, Luc Van Gool, and Federico Tombari. The hidden uncertainty in a neural networks activations. *arXiv preprint arXiv:2012.03082*, 2020. [1](#), [2](#)
- [57] Janis Postels, Francesco Ferroni, Huseyin Coskun, Nassir Navab, and Federico Tombari. Sampling-free epistemic uncertainty estimation using approximated variance propagation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2931–2940, 2019. [3](#)
- [58] Farhad Pourpanah, Moloud Abdar, Yuxuan Luo, Xinlei Zhou, Ran Wang, Chee Peng Lim, Xi-Zhao Wang, and QM Jonathan Wu. A review of generalized zero-shot learning methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [2](#)
- [59] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [2](#), [10](#)
- [60] Murat Sensoy, Lance M. Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. In *Advances in Neural Information Processing*, pages 3183–3193, 2018. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [12](#), [13](#)
- [61] Kshitij Sirohi, Sajad Marvi, Daniel Büscher, and Wolfram Burgard. Uncertainty-aware panoptic segmentation. *IEEE Robotics and Automation Letters*, 8(5):2629–2636, 2023. [3](#)
- [62] Konstantin Sofiiuk, Olga Barinova, and Anton Konushin. AdaptIS: Adaptive instance selection network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7355–7363, 2019. [2](#)
- [63] Joost Van Amersfoort, Lewis Smith, Yee Whye Teh, and Yarin Gal. Uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, pages 9690–9700. PMLR, 2020. [2](#), [6](#), [7](#), [11](#), [12](#)
- [64] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(11), 2008. [8](#)
- [65] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. MaX-DeepLab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5463–5474, 2021. [2](#)
- [66] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-DeepLab: Stand-alone axial-attention for panoptic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 108–126. Springer, 2020. [2](#)
- [67] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE Transactions on Knowledge and Data Engineering*, 2022. [1](#)
- [68] Kelvin Wong, Shenlong Wang, Mengye Ren, Ming Liang, and Raquel Urtasun. Identifying unknown instances for autonomous driving. In *Proceedings of the Conference on Robot Learning*, pages 384–393. PMLR, 2020. [2](#), [3](#), [4](#), [6](#), [7](#), [9](#), [10](#), [11](#), [12](#), [14](#), [15](#)

- [69] Sanghyun Woo, Dahun Kim, Joon-Young Lee, and In So Kweon. Learning to associate every segment for video panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2705–2714, 2021. 2
- [70] Shun-Cheng Wu, Keisuke Tateno, Nassir Navab, and Federico Tombari. Incremental 3D semantic scene graph prediction from RGB sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5064–5074, 2023. 2
- [71] Shun-Cheng Wu, Johanna Wald, Keisuke Tateno, Nassir Navab, and Federico Tombari. SceneGraphFusion: Incremental 3D scene graph prediction from RGB-D sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7515–7525, 2021. 2
- [72] Xinyi Wu, Zhenyao Wu, Hao Guo, Lili Ju, and Song Wang. DANNet: A one-stage domain adaptation network for unsupervised nighttime semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15769–15778, 2021. 1
- [73] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. UPSNet: A unified panoptic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8818–8826, 2019. 2
- [74] Hai-Ming Xu, Hao Chen, Lingqiao Liu, and Yufei Yin. Two-stage decision improves open-set panoptic segmentation. *arXiv preprint arXiv:2207.02504*, 2022. 2, 3, 6
- [75] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *Proceedings of the European Conference on Computer Vision*, pages 736–753. Springer, 2022. 2, 9, 10
- [76] Kai Yi, Xiaoqian Shen, Yunhao Gou, and Mohamed Elhoseiny. Exploring hierarchical graph representation for large-scale zero-shot image classification. In *Proceedings of the European Conference on Computer Vision*, pages 116–132. Springer, 2022. 2
- [77] Le You, Han Jiang, Jinyong Hu, C Hwa Chang, Lingxi Chen, Xintong Cui, and Mengyang Zhao. GPU-accelerated faster Mean Shift with Euclidean distance metrics. In *Proceedings of the Computers, Software, and Applications Conference*, pages 211–216. IEEE, 2022. 13
- [78] Zhongqi Yue, Tan Wang, Qianru Sun, Xian-Sheng Hua, and Hanwang Zhang. Counterfactual zero-shot and open-set visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15404–15414, 2021. 2, 3
- [79] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. LiT: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18123–18133, 2022. 2
- [80] Yang Zhang, Ashkan Khakzar, Yawei Li, Azade Farshad, Seong Tae Kim, and Nassir Navab. Fine-grained neural network explanation by identifying input features with predictive information. *Advances in Neural Information Processing Systems*, 34:20040–20051, 2021. 1
- [81] Ye Zheng, Jiahong Wu, Yongqiang Qin, Faen Zhang, and Li Cui. Zero-shot instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2593–2602, 2021. 2, 3
- [82] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 4