

Multi-domain semantic segmentation with overlapping labels*

Petra Bevandić[†], Marin Oršić[†], Ivan Grubišić, Josip Šarić, Siniša Šegvić
University of Zagreb, Faculty of Electrical Engineering and Computing

name.surname@fer.hr, marin.orsic@gmail.com

Abstract

Deep supervised models have an unprecedented capacity to absorb large quantities of training data. Hence, training on many datasets becomes a method of choice towards graceful degradation in unusual scenes. Unfortunately, different datasets often use incompatible labels. For instance, the Cityscapes road class subsumes all driving surfaces, while Vistas defines separate classes for road markings, manholes etc. We address this challenge by proposing a principled method for seamless learning on datasets with overlapping classes based on partial labels and probabilistic loss. Our method achieves competitive within-dataset and cross-dataset generalization, as well as ability to learn visual concepts which are not separately labeled in any of the training datasets. Experiments reveal competitive or state-of-the-art performance on two multi-domain dataset collections and on the WildDash 2 benchmark.

1. Introduction

Large realistic datasets [17, 5, 21, 40] have immensely contributed to development of dense prediction models. The reported accuracy grew rapidly over the last five years so that some of these datasets appear solved today. Unfortunately, the learned models often perform poorly in the wild [35]. Consequently, a desire emerged to measure the model performance by evaluating on multiple datasets from several domains [16, 39, 15], akin to combined events in athletics. The most straightforward approach towards that goal is simultaneous training on many datasets. Such training is likely to yield outstanding resilience due to extremely large capacity of convolutional models [36, 30].

Training on multiple datasets is especially interesting for dense prediction models due to very expensive labels [42]. However, this is not easily carried out in prac-



Figure 1. We address dense-prediction models which can learn visual concepts from multiple datasets with overlapping classes. Example: pickups are labeled as class truck in VIPER [25] (left), class car in Vistas [21] (middle) and class van in Ade20k [40] (right). We resolve the class overlap (Viper-truck vs Vistas-car vs Ade20k-van) by learning on partial labels [6].

tice since existing datasets use incompatible taxonomies [16, 15]. These incompatibilities may be caused by discrepant granularity and overlapping classes. Discrepant granularity [16] arises where some class from dataset A corresponds to several classes from dataset B. For instance the class road in Cityscapes is further divided into 8 classes in Vistas: road, bike_lane, crosswalk_plain, marking_zebra, marking_other, manhole, pothole, and service_lane. Overlapping classes occur when visual concepts get inconsistently grouped across datasets. For instance class truck in VIPER [25] includes trucks and pickups. Class car in Vistas [21] includes cars, pickups and vans. Class van in Ade20k [40] includes vans and pickups. We say that VIPER truck overlaps Vistas car (and Ade20k van) since the corresponding pixels have non-empty intersection and non-empty set difference in both directions. Fig. 1 shows that the three datasets assign pickups into three different classes.

This paper addresses two important challenges in multi-domain semantic segmentation: i) training on datasets with inconsistent labels, and ii) designing an experimental setup capable of learning hundreds of dense logits on megapixel resolution. We contribute a novel method for training semantic segmentation on datasets with discrepant granularity and overlapping classes by leveraging partial labels [6] and class-wise log-sum-prob loss. Our method builds a uni-

*This work has been supported by Croatian Science Foundation - grant IP-2020-02-5851 ADEPT, by European Regional Development Fund - grant KK.01.2.1.02.0119 DATACROSS and by VSITE College for Information Technologies who provided access to 6 GPU Tesla-V100 32GB.

[†]Equal contribution.

versal taxonomy such that each dataset-specific class can be expressed as a union of one or more universal classes. Our universal models can be seamlessly trained and evaluated on each individual dataset since probabilities of dataset classes correspond to sums of probabilities of universal classes.

Our method outperforms two baselines which ignore overlapping classes, as well as a recent approach [15] based on partial relabeling towards a closed unified taxonomy. Computational advantages of our method become decisive in large-scale multi-domain experiments where we achieve state-of-the-art performance on the RVC 2020 benchmark collection and the WildDash 2 benchmark.

2. Related Work

We consider efficient training on a collection of semantic segmentation datasets with discrepant granularities and overlapping classes. This requires a capability to learn fine-grained classes on coarse-grained labels. Thus, we review semantic segmentation, efficient architectures, learning with partial labels, and multi-domain dense prediction.

2.1. Semantic Segmentation

Deep convolutional models have spurred substantial progress in the field of semantic segmentation [19, 3, 38, 4]. Recent work improves recovery of the spatial details lost due to downsampling [3] and increases the receptive field [38]. U-Net [26] recovers the fine details by blending deep semantics with spatial information from the early layers. Further work notices that the upsampling path requires much less capacity than the recognition backbone [18, 13]. Increased subsampling enlarges the receptive field, improves efficiency and reduces the training footprint [2, 13]. Further shrinking of the training footprint can be achieved through gradient checkpointing [27, 14]. Multi-scale processing relaxes the requirements on recognition capacity and leads to competitive results with efficient backbones [24, 37]. HRNet [32] preserves fine details by sustaining high-resolution representations along the entire convolutional pipeline. Its improved variant uses attention pooling to promote pixels of the majority class [34].

2.2. Efficient inference and training

Several recent semantic segmentation architectures aim at high accuracy while supporting real-time inference [2, 22, 24, 10, 37]. Most of these approaches can be trained at a fraction of the computational budget required by widely used methods. Hence, efficient architectures may create new opportunities while making our research more inclusive and environmentally acceptable [28].

2.3. Partial labels and log-sum-prob loss

Log-sum-prob loss allows to train unions of disjoint predictions by aggregating over classes or data. Data-wise for-

mulation has been used to decrease influence of noisy labels [41]. Their loss sums probabilities of all labels from the 3×3 neighborhood in pixels near semantic borders. On the other hand, our loss performs class-wise aggregation by summing probabilities of all universal classes corresponding to the coarse-grained label. Our motivation is to allow learning over inconsistent taxonomies, while they alleviate inaccurate annotations at semantic boundaries.

Learning from partial labels considers examples labeled with a set of classes only one of which is correct. The original work [6] assumes fairly stochastic probability $P(\mathbb{Z}|\mathbf{x}, y)$ of false labels \mathbb{Z} given the datum \mathbf{x} and its true label y . In our context, this relation is deterministic when the dataset of the datum is known. Different from [6], we formulate learning with partial labels in a principled probabilistic framework based on log-sum-prob loss.

Recent work [39] considers partial labels for multi-domain object detection. Their experiments show that log-sum-prob loss does not contribute to object detection which is likely due to the sheer asymmetry between positive and negative windows. On the other hand, our experiments indicate that log-sum-prob loss and partial labels are an effective tool for cross-dataset semantic segmentation.

2.4. Multi-domain dense prediction

Multi-domain dense prediction involves training on multiple datasets with different taxonomies. A simple baseline consists of a shared encoder and separate per-dataset decoders [12]. This is similar to training on naive concatenation of individual datasets. We find this to be suboptimal since repetition of classes (eg. Vistas-bus vs Ade20k-bus) increases the memory demands and decreases utility of predictions. Instead, we prefer to train on a shared universal taxonomy which contains much less classes and supports inference on novel datasets.

A universal taxonomy can be implemented as a hierarchy of logits for categories and classes [16, 20]. This approach can gracefully handle datasets with discrepant granularities. However, on the downside, it requires very complex training and can not train on overlapping classes.

A universal taxonomy can also be implemented as a flat vector of universal logits [15, 39]. The only such semantic segmentation approach [15] does not train on heterogeneous datasets. Instead, they propose to adapt all datasets towards a unified taxonomy of their own. This requires relabeling all datasets which do not distinguish classes from their universal taxonomy, and merging all classes which are more fine-grained than the universal taxonomy. It would be impractical for such unified taxonomy to include all dataset classes, since that would imply unacceptable relabelling effort. Consequently, such design does not allow evaluation on original datasets nor simple introduction of datasets with novel taxonomies.

As mentioned in 2.3 a recent multi-domain object-detection approach [39] considers candidate windows which may correspond either to the background or to any of the non-annotated classes. Their attempt to solve that problem with log-sum-prob loss was not successful. However, classes of coarse-grained datasets usually group visual concepts which are much more similar to each other than to a catch-all background class. Consequently, log-sum-prob training may have better a chance for success in multi-domain semantic segmentation.

Different than all previous approaches, we present a principled method for cross-dataset and multi-domain training of semantic segmentation models. We assemble a universal flat taxonomy from an almost unconstrained collection of datasets. The recovered universal taxonomy allows to seamlessly train and evaluate a universal probabilistic model on unchanged individual datasets by relying on partial labels and log-sum-prob loss. The resulting models are suitable for universal semantic segmentation in the wild.

3. Multi-domain semantic segmentation

We present a principled method for simultaneous training on semantic segmentation datasets with incompatible labeling policies. Our method builds a universal taxonomy which gathers and refines semantics of individual datasets by expressing each dataset class as a union of universal classes. Our models learn to predict universal classes while fully supporting training and evaluation according to individual datasets. We take care to contain the memory footprint of our models since effective multi-domain learning requires large crops and large batches.

3.1. Terminology and notation

We consider semantic classes c as sets of pixels in all possible images and use set notation to express relations between them. For example: $Vistas\text{-}sky = CS\text{-}sky$, $CS\text{-}road \supset Vistas\text{-}manhole$, $VIPER\text{-}truck \cap Vistas\text{-}car = WD\text{-}pickup$, $CS\text{-}road \perp CS\text{-}car \Rightarrow CS\text{-}road \cap CS\text{-}car = \emptyset$. Note that WD and CS stand for WildDash 2 and Cityscapes, respectively.

A set of mutually disjoint semantic classes defines a flat taxonomy $\mathcal{S}_d = \{c_t\}, c_i \perp c_j, \forall i, j: i \neq j$. We note that a union of flat taxonomies is not guaranteed to be a true taxonomy since members may have non-empty intersections. For example, $\mathcal{S}_{VIPER} \cup \mathcal{S}_{Vistas}$ is not a true taxonomy since $VIPER\text{-}truck \not\perp Vistas\text{-}car$ (cf. Fig. 1).

A semantic segmentation dataset consists of images and corresponding dense labels: $\mathcal{D} = \{(x_k, y_k)\}$. The labels assign semantic classes $c \in \mathcal{S}_d$.

3.2. Baseline approach: naive concatenation

We consider training a semantic segmentation model on a compound dataset $\bigcup_d \mathcal{D}_d$. This can not be accomplished

by relying on individual taxonomies \mathcal{S}_d , since that would either entail information loss (eg. remapping $Vistas\text{-}manhole$ to $CS\text{-}road$), or require expensive relabeling (eg. remapping some pixels from $CS\text{-}road$ to $Vistas\text{-}manhole$).

We propose a simple baseline approach which creates a training taxonomy as a *naive concatenation* of all individual taxonomies \mathcal{S}_d . Each training taxonomy is mapped onto the corresponding subset of universal labels. For instance, cars from Cityscapes are mapped to the class $Cityscapes\text{-}car$, while cars from Vistas are mapped to $Vistas\text{-}car$. The predictions are produced by a single softmax over $\sum_d |\mathcal{S}_d|$ logits, which allows training with the cross-entropy loss.

However, the described approach does not produce a true taxonomy, and thus leads to contention between overlapping logits (e.g. $CS\text{-}road$ vs $Vistas\text{-}road$). This wastes model capacity due to need for dataset recognition instead of promoting generalization by principled dealing with ambiguous supervision. Additionally, the class replication increases the number of logits. This may hurt recognition of rare classes and exhaust GPU memory since the predictions are produced on a large resolution. Finally, inference has to deal with starvation of overlapping logits which we alleviate with post-inference mappings (cf. 3.5).

3.3. Baseline approach: partial merge

Naive concatenation can be improved by merging classes which match exactly; for example, $Cityscapes\text{-}sky$ and $Vistas\text{-}sky$ can be merged into the universal class sky . Classes which can not be matched in a 1:1 manner remain separate (e.g. $Cityscapes\text{-}road$, $Vistas\text{-}road$ and $Vistas\text{-}marking$). This is similar to naive merge from [15], however we merge according to semantics instead of symbolic names and use post-inference mappings as described in 3.5. This may reduce the amount of capacity needed to distinguish between datasets due to eliminated contention between the merged classes. Nevertheless, partial merge also fails to produce a true taxonomy since overlapping classes remain unresolved. Hence, we encounter the same problems as presented in 3.2 although to a lesser extent.

3.4. Creating a universal taxonomy

We address disadvantages of the two baselines by building a flat universal taxonomy \mathcal{U} which will allow i) training on overlapping labels (cf. 3.6), and ii) seamless dataset-specific (cf. 3.5) or dataset-agnostic evaluation. The desired universal taxonomy \mathcal{U} should have the following properties. First, its classes should encompass the entire semantic range of individual datasets: $\bigcup_{u \in \mathcal{U}} u = \bigcup_d \mathcal{S}_d$. Second, universal classes should be disjoint: $\forall u, u' \in \mathcal{U}: u \perp u'$. Third, no universal class may have a non-empty difference towards an intersecting dataset-specific class: $\forall u \in \mathcal{U}, c \in \bigcup_d \mathcal{S}_d: (u \perp c) \vee (u \subseteq c)$. Such taxonomy will allow us to define mappings $m_{\mathcal{S}_d}: \mathcal{S}_d \rightarrow 2^{\mathcal{U}}$ from dataset-specific

classes to subsets of the universal set.

We propose to recover the universal taxonomy \mathcal{U} by iteratively transforming a multiset \mathbb{M} that contains dataset specific classes from all datasets. We apply the rules following rules until all of the classes are disjoint.

1. If two classes c_i and c_j match exactly, we merge them into a new class c' and map c_i and c_j to c' .

Example: Since WD-sky and CS-sky are equal, we merge them into a class sky, and define new mappings: WD-sky \mapsto sky, CS-sky \mapsto sky.

2. If a class c_i is a superset of a class c_j , then c_i is mapped to $\{c_j, c'_i\}$ where $c'_i = c_i \setminus c_j$.

Example: KITTI-car is a superset of car because it contains vans. We therefore add van as a new class and create a mapping KITTI-car \mapsto {car, van}.

3. If two classes overlap, $(c_i \not\subseteq c_j) \wedge (c_i \setminus c_j \neq \emptyset) \wedge (c_j \setminus c_i \neq \emptyset)$, c_i and c_j are replaced with three new disjoint classes $c'_i = c_i \setminus c_j$, $c'_j = c_j \setminus c_i$, and $c' = c_i \cap c_j$. Class c is mapped to $\{c'_i, c'\}$, while k is mapped to $\{c'_j, c'\}$.

Example: VIPER-truck contains trucks and pickups while Ade20k-truck contains trucks and trailers. We therefore replace VIPER-truck and Ade20k-truck with truck, pickup and trailer and create the following mappings: VIPER-truck \mapsto {truck, pickup} and Ade20k-truck \mapsto {truck, trailer}

3.5. Mapping predictions to target classes

In practice, universal models are evaluated on particular dataset-specific taxonomies. We therefore design evaluation mappings from dataset specific taxonomy \mathcal{S}_d to the training class set \mathcal{T} . We extend \mathcal{S}_d with a void class. In our experiments, \mathcal{T} is either one of the two baselines or the universal taxonomy \mathcal{U} , while \mathcal{S}_d corresponds to the particular dataset (e.g. $\mathcal{S}_d = \mathcal{S}_{\text{CamVid}}$). We can create evaluation mappings according to the following rules:

1. Each dataset-specific class is mapped to all training classes it overlaps with: $c^{S_d} \mapsto \{c^T \in \mathcal{T} : c^{S_d} \not\subseteq c^T\}$.
2. The void class is mapped to all training classes which do not overlap with any of the dataset-specific classes: void $\mapsto \{c^T \in \mathcal{T} : \forall c^{S_d} \in \mathcal{S}_d, c^{S_d} \perp c^T\}$.

We can use the created mappings to calculate the classification score for each dataset-specific class in each pixel:

$$S(Y^{S_d} = c|\mathbf{x}) = \sum_{t \in m_{S_d}(c)} P(Y^{\mathcal{T}} = t|\mathbf{x}). \quad (1)$$

In the case of the universal taxonomy, the above score corresponds to probability since each universal class occurs in the mapping set of exactly one dataset-specific

class. In the case of the two baselines, the score has to be renormalized because training classes may overlap with several dataset-specific classes, e.g. $S(\text{Vistas-pothole}|\mathbf{x}) = P(\text{cs-road}|\mathbf{x}) + P(\text{vistas-pothole}|\mathbf{x})$, $P(\text{Vistas-pothole}|\mathbf{x}) = S(\text{Vistas-pothole}|\mathbf{x}) / \sum S(c_j^{S_d}|\mathbf{x})$. Thus, creation of post-inference mappings for the two baselines requires almost the same effort as the design of our universal taxonomy.

3.6. Learning a universal model with partial labels

We propose to learn cross-dataset models on the universal taxonomy \mathcal{U} in order to properly address overlapping labels while fully exploiting all available supervision. The resulting universal models will overcome all disadvantages of the two baselines. In particular, they will be capable to exploit partial supervision while being applicable in the wild.

Let us model the probability of universal classes as per-pixel softmax over $|\mathcal{U}|$ logits \mathbf{s} . Let the random variable Y denote a universal model prediction at that pixel, and let \mathbf{p} denote the softmax output. Then we can write: $p_u = \text{softmax}(\mathbf{s})_u = P(Y = u|\mathbf{x})$. Hence, the negative log-likelihood of the probability of a dataset-specific class c can be expressed as:

$$\mathcal{L}^{\text{NLL+}}(\mathbf{p}, c) = -\ln \sum_{u \in m_{S_d}(c)} p_u. \quad (2)$$

Note that the sum of p_u can be considered as probability since our universal classes are disjoint by design. We denote the resulting log-sum-prob form as NLL+: negative log-likelihood of aggregated probability. This loss becomes the standard negative log-likelihood when $|m_{S_d}(c)| = 1$.

In order to better understand the NLL+ loss (2) we analyze its partial derivatives with respect to logits \mathbf{s} :

$$\frac{\partial \mathcal{L}}{\partial s_u} = p_u - \mathbb{1}[u \in m_{S_d}(c)] \frac{\exp s_u}{\sum_{k \in m_{S_d}(c)} \exp s_k}. \quad (3)$$

We observe that gradients w.r.t. incorrect logits are positive and the same as in the standard case, since then the second term is zero. Furthermore, the gradients w.r.t. labeled logits are always negative since the two terms have equal numerators, while the second term has a smaller denominator. If there are two correct logits, the larger one will have an exponentially larger numerator of the second term, while the denominator will be the same. Hence the difference between the correct logits is likely to increase after the update.

This analysis indicates that NLL+ favours peaky sub-distributions over correct logits: the model is compelled to pick only one among them. In the limit, when one of the labeled logits becomes much stronger than its peers, its gradients become equal to the standard supervised case, provided that the chosen logit is correct. This suggests that partial labels may help whenever there exists some learning signal that supports recognition of individual universal classes.

Our GPU implementation represents each dataset mapping m_{s_d} as an $|S_d| \times |U|$ binary matrix with unit column sums. Probabilities of dataset-specific classes are computed by multiplying this matrix with universal probabilities.

3.7. Efficient multi-domain semantic segmentation

Multi-domain semantic segmentation is an extremely computationally intensive task since the complexity of training may be in the petaFLOP range. The training footprint of the model is an important factor since the model performance improves with large batches, large training crops and increased number of channels along the upsampling path. These considerations are a compelling reason to prefer training on universal classes instead of on less principled options such as the two baselines presented before.

This discussion suggests that efficient approaches may represent a sensible choice for such problems. We therefore base our experiments on pyramidal SwiftNet [24]. All our models use three image scales and produce dense logits at $8\times$ subsampled resolution. These logits are bilinearly upsampled to the full resolution.

We train our models with a compound loss which corresponds to the product of the boundary-aware (BA) factor and the NLL+ loss (2). The BA modulation factor prioritizes poorly classified pixels and pixels at boundaries [24]:

$$\mathcal{L}^{\text{BA}}(\mathbf{p}, c) = -\alpha \cdot e^{\gamma(1-\hat{p})} \ln \hat{p}, \text{ where } \hat{p} = \sum_{u \in m_{s_d}(c)} p_u. \quad (4)$$

We note that memory efficiency becomes extremely important when training with several hundred logits and large crops. Even a humble inference becomes a challenge in $2\times$ upsampled Vistas images (24MPx) used in multiscale inference since then the logit tensor requires almost 20GB RAM. We avoid caching multiple activations at full resolution by implementing the boundary aware NLL+ loss (4) as a layer with custom backprop. This ensures a minimal increase of the training footprint with respect to the standard NLL formulation.

4. Experiments

We present semantic segmentation experiments on models based on pyramidal SwiftNet [24]. Most experiments are based on ResNet-18 [9], except our RVC models which use DenseNet-161 [11]. All our models are pre-trained on ImageNet. We train by optimizing the BA loss (4) (either NLL or NLL+) with Adam. We attenuate the learning rate from $5 \cdot 10^{-4}$ to $6 \cdot 10^{-6}$ by cosine annealing.

We train our validation models by early stopping with respect to the average mIoU across all datasets. Most of our models may emit predictions which cannot be mapped to any of the evaluation classes. For instance, Vistas-ego-

vehicle is not labeled on Cityscapes. We map such predictions to void and count them as false negatives during calculation of the mIoU score.

We train on random square crops which we augment through resizing from $0.5\times$ to $2\times$ and horizontal flipping. The default crop size is 768. Our mini-batches favor images with multiple class instances and rare classes [27]. We also reweight the images so that the sum of image weights is equal for all of the datasets. This encourages fair representation of datasets and classes during training.

4.1. Comparison with two baselines

We compare our universal taxonomy with two baselines: naive concatenation (cf. 3.2) and partial merge (cf. 3.3). Both baselines require complex post-inference mapping before evaluation on particular datasets. For example joint training on Cityscapes and Vistas requires: $\{\text{cs-sign, vistas-sign-front, vistas-sign-back, vistas-sign-frame}\} \mapsto \text{CS-sign}$. Some universal logits will be mapped to more than one evaluation class, eg: $\{\text{cs-road, vistas-zebra}\} \mapsto \text{Vistas-zebra}$ and $\{\text{cs-road, vistas-pothole}\} \mapsto \text{Vistas-pothole}$. The merged training classes are 1:1 mapped to evaluation classes, eg: $\text{sky} \mapsto \text{CS-sky}$ and $\text{sky} \mapsto \text{Vistas-sky}$.

We consider joint training on two pairs of road-driving datasets with incompatible taxonomies. We train for 100 epochs with batch size 18 which is the maximum that can fit into one Tesla V100 for the naive concatenation models. Our first experiment considers Cityscapes and Vistas. In this case, our universal taxonomy (cf. 3.4) corresponds to the Vistas taxonomy since it is a strict refinement of Cityscapes. Our universal taxonomy has 65 classes, which is less than the naive union (93) and the partial merge (72).

Our second experiment considers joint training on Vistas and WildDash 2 (WD2) [35]. We form WD2 minival by collecting the first 572 images and use all remaining images as WD2 minitrain. This setup is different from the first one since WD2 is much more diverse than Cityscapes. Hence, the baselines are going to have a harder time to recognize particular datasets according to the camera type or location. Additionally, WD2 has more classes than Cityscapes and a finer granularity of car types than Vistas. Consequently, our universal taxonomy has 67 classes which is less than the naive concatenation (98), and the partial merge (76).

Table 1 shows that the proposed universal taxonomy succeeds to outperform both baselines in spite of reduced capacity due to less logits. The advantage is especially evident in the case of joint training on WD2 and Vistas. We hypothesize that this occurs due to reduction of contention between overlapping classes, which allows the model capacity to be more efficiently exploited. Our universal taxonomy also outperforms the baselines on City-Vistas although the advantage is smaller than on WD2-Vistas. We explain this as follows. First, all Cityscapes images have been acquired

across a contained geographical region with the same camera model. This makes dataset detection an easy task and hence alleviates class contention within the baselines. Second, this setup entails a smaller difference between our taxonomy and the partial merge (City-Vistas: 7 logits vs WD2-Vistas: 9 logits). We notice that our universal taxonomy outperforms individual training (Cityscapes: 0.5pp, Vistas: 0.1pp [24]), as well as that training on WD2-Vistas slightly reduces Vistas performance. This indicates that WildDash 2 requires more capacity than Cityscapes.

Taxonomy	(City-Vistas)		(WD2-Vistas)	
	City	Vistas	WD2	Vistas
naive concat	76.8	44.6	55.3	43.1
partial merge	77.1	44.5	54.7	44.1
universal (ours)	77.0	44.9	56.2	44.4

Table 1. Evaluation of joint training on Cityscapes and Vistas (City-Vistas), as well as on WildDash 2 and Vistas (WD2-Vistas). Our universal taxonomy outperforms both baselines.

4.2. Cross-dataset evaluation

We explore models from Table 1 on novel road-driving datasets: CamVid test [1], KITTI train [8], BDD val [33] and IDD val [31]. We also evaluate the City-Vistas model on WildDash V2 minival and the Vistas-WD2 model on Cityscapes val. As before, we design mappings from each of the six training taxonomies to each evaluation taxonomy, and consider all unmapped logits as class void.

Tables 2 and 3 show that our taxonomy outperforms naive concatenation and partial merge on most foreign datasets, while in the remaining few cases the difference is within variance. Note that KITTI has only 200 images and is very similar to Cityscapes. As before, our contribution is greater on WD2-Vistas than City-Vistas. We connect that with Cityscapes peculiarities which discourage generalization, as discussed in 4.1.

Model	WD2	CV	KITTI	BDD	IDD
naive concat	43.3	74.1	58.9	56.7	42.4
partial merge	43.8	73.9	59.4	57.0	43.0
universal (ours)	43.9	75.3	60.5	58.0	42.8

Table 2. Cross-dataset evaluation of joint training on Cityscapes and Vistas. We evaluate the three models from Table 1 on WildDash 2 mini val, CamVid test, KITTI, BDD val, and IDD val.

4.3. Comparison with partial manual relabeling

We consider 7 datasets from different domains: Ade20k [40], BDD [33], Cityscapes [5], Coco [17], IDD [31], SUN RGBD [29] and Vistas [21]. Previous work has partially relabeled these datasets in order to make them fully compatible with a custom unified taxonomy called MSeg [15].

Model	CS	CV	KITTI	BDD	IDD
naive concat	69.0	72.7	53.6	56.1	41.6
partial merge	69.8	72.4	53.5	57.1	41.9
universal (ours)	71.4	74.9	53.0	59.0	42.6

Table 3. Cross-dataset evaluation of joint training on Vistas and WildDash 2. We evaluate the three models from Table 1 on Cityscapes val, CamVid test, KITTI, BDD val, and IDD val.

However, they had to omit 100 classes in order to contain the relabeling effort. We explore this trade-off by comparing identical models trained with: i) our universal taxonomy, original datasets and 294-way NLL+ loss, and ii) MSeg taxonomy, relabeled datasets, and 194-way NLL loss.

We resize each training image so that its smaller side is 1080 pixels, sample random square crops between 256×256 and 1024×1024 , and resize them to 512×512 . We train both models on a single Tesla V100 for 20 epochs with batch size 20. We validate the two models according to the following two protocols. The MSeg protocol only evaluates the 194 classes from the MSeg taxonomy [15], while the original protocols evaluate all classes.

The top section of Table 4 presents evaluation according to the original protocol. Our approach prevails on most datasets due to capability to predict all 294 classes. The bottom section of the table shows that training with NLL+ remains competitive even when we evaluate only on the 194 classes covered by the MSeg taxonomy. This suggests that our universal taxonomy represents a flexible alternative to custom taxonomies, especially in view of adding new datasets to the training collection (cf. 4.5).

4.4. Recognition of unlabeled concepts

This experiment explores whether our method is able to recognize unlabeled concepts. We partition Cityscapes train into two splits with approximately equal number of images and similar overall distribution of classes. The first split gathers images from Aachen to Hanover. It relabels all trucks, buses and cars into the class called 'four-wheels-vehicle'. The second split gathers images from Jena to Zurich and relabels bicycles, motorcycles and cars into the class 'personal-vehicle'. Both splits have 17 classes. Buses, trucks, motorcycles and bicycles are labeled as standalone classes in only one of the two splits (roughly, half of the images), while cars are never labeled as a standalone class. Our universal taxonomy corresponds to 19 Cityscapes classes. The dataset class four-wheel-vehicle maps to {car, bus, truck}. The dataset class personal-vehicle maps to {car, bicycle, motorcycle}. All other dataset classes map to themselves.

We validate four models for 19-way dense prediction. The baseline model ignores all pixels of aggregate classes and uses the standard NLL loss. Two models train on aggre-

Evaluation protocol	Taxono	Ade20k	BDD	Cityscapes	COCO	IDD	SUN RGBD	Vistas
Original	Universal (ours)	31.0	58.5	72.6	35.4	54.4	41.7	39.1
	MSeg	23.3	59.4	72.6	30.3	42.6	40.2	26.1
MSeg	Universal (ours)	34.5	58.5	72.6	36.3	53.0	41.7	45.4
	MSeg	34.3	59.4	72.6	34.9	55.6	40.2	43.6

Table 4. Multi-domain experiments with SNp-RN18 on the seven MSeg datasets [15]. We train a NLL+ model on original labels, and compare it with a NLL model which trains on manually relabeled images according to the MSeg taxonomy [15]. Both models are evaluated on validation subsets of Ade20k, BDD, Cityscapes, Coco, IDD, SUN RGB-D and Vistas. We consider all unmapped logits as class void.

gate classes through our NLL+ loss (4) and its variant which replaces the sum with max. The oracle uses the original Cityscapes train. We train all four models for 250 epochs with batch size 14 on GTX1080 while oversampling trains.

Table 5 reveals that NLL+ performs much closer to the oracle than to the NLL baseline. There is a substantial advantage for classes with only half non-ambiguous labels (bus, truck, motorcycle, bicycle), and especially on cars which are recognized as a novel concept at the intersection of two training labels. NLL-max [39] performs only slightly better than the NLL baseline, while NLL+ appears as a method of choice in presence of class overlap. The two baselines succeed to detect cars due to post-inference mapping, however NLL+ outperforms them substantially.

Model						mIoU
NLL baseline	0	54.2	43.9	32.5	60.3	58.7
NLL-max	0	9.6	40.8	1.6	75.7	61.8
naive concat	91.1	61.4	42.2	39.0	72.3	67.6
partial merge	92.4	54.3	55.1	39.0	74.3	71.4
NLL+	93.6	73.3	66.6	46.4	75.4	74.3
oracle	94.4	82.9	72.9	62.2	76.5	76.2

Table 5. Training on relabeled Cityscapes: one split aggregates cars, buses and trucks, while the other aggregates cars, bicycles and motorcycles. We validate NLL+ training, its NLL-max variant, and the NLL baseline which ignores all aggregate labels.

4.5. Large-scale multi-domain challenge

We present a large-scale multi-domain experiment on a benchmark collection proposed at the Robust Vision Challenge RVC 2020. The challenge requires submitting a single semantic segmentation model with less than 300 logits to seven benchmarks: Ade20k, Cityscapes, KITTI, Vistas, Scannet [7], Viper, and WildDash 2.

We recover a taxonomy of 192 universal classes according to the procedure from 3.4. It turns out that Vistas labels vehicle windows as vehicles, while VIPER labels these pixels with what is seen behind. A consistent resolution of this issue would effectively double the number of logits since most universal classes would require addition of their twin

class (eg. person vs person_through_glass). Consequently, we introduce simplifying assumptions such as VIPER-car = Vistas-car and Vistas-car \perp VIPER-person in order to reduce the training footprint of the universal model.

We increase the model capacity by using the DenseNet-161 [11] backbone and setting the upsampling width to 384 channels. We decrease the memory footprint with custom backprop and gradient checkpointing [27] in order to enable training on batches of $8 \times 768 \times 768$ crops per GPU.

We train on seven datasets of the RVC 2020 collection on 6 Tesla V100 32GB for 100 epochs. We set the boundary modulation to 1 (minimum) for all ScanNet crops in order to alleviate noisy labels. We evaluate on six scales and two horizontal flips. Table 6 compares our model SNp_DN161 with the two valid submissions to RVC 2020 [23]. Fig. 2 presents qualitative performance of our model.

Model	ADE	CS	KIT	MV	SN	VIP	WD
MSeg1080	33.2	80.7	62.6	34.2	48.5	40.7	35.2
SN_rm152pyrx8	31.1	74.7	63.9	40.4	54.6	62.5	45.4
SNp-DN161	30.8	77.9	68.9	44.6	53.9	64.6	46.8

Table 6. Performance evaluation on RVC 2020 benchmarks Ade20k (ADE), Cityscapes (CS), KITTI (KIT), Vistas (MV), ScanNet (SN), VIPER (VIP) and WildDash 2 (WD). We compare the two valid RVC submissions (top) with our model (bottom).

We note that our model significantly outperforms the baseline submission MSeg1080.RVC [15] which is unable to predict any classes outside their closed unified taxonomy. This suggests advantage of open universal taxonomies such as ours, similarly as in 4.2. We believe that multi-domain training is instrumental for achieving robust performance in the real world. This is confirmed by the state-of-the-art performance of our RVC model on the WildDash 2 benchmark which is currently the only dataset which specifically targets expected points of failure of dense prediction models.

5. Conclusion

We have presented a principled method for cross-dataset training of semantic segmentation models. Our method allows seamless training and evaluation on multi-

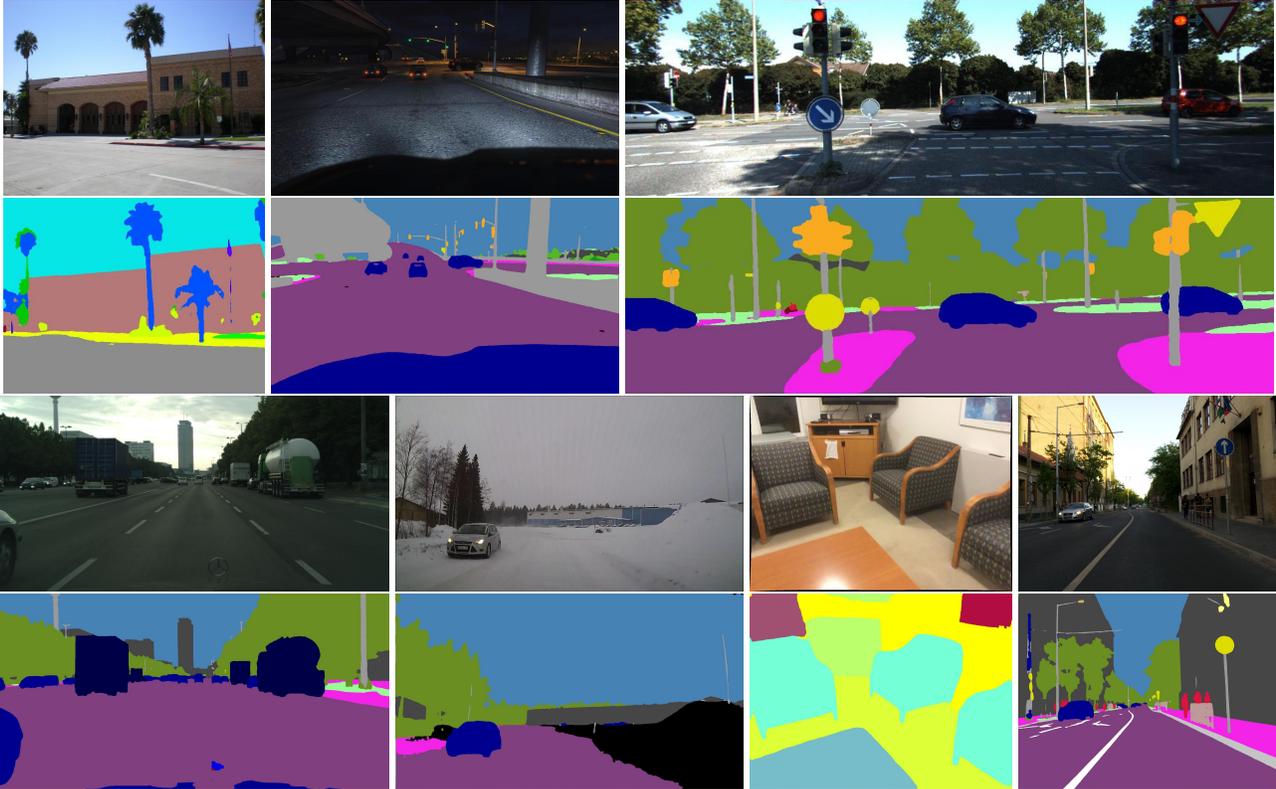


Figure 2. Qualitative performance of our SNp-DN161 submission to RVC 2020. Rows 1 and 3 show input images while rows 2 and 4 show the model predictions. Images belong to (top to bottom, left to right) ADE20k, Viper, KITTI, Cityscapes, WildDash, ScanNet and Vistas.

domain datasets with discrepant granularity and overlapping classes. Unlike simpler approaches such as multi-head prediction with shared features and dataset detection, our method recovers a universal taxonomy which transcends semantics of individual datasets. This allows transparent inference on novel datasets and promises thriving performance in the wild.

Our method expresses each dataset class as a union of disjoint universal classes. Hence, probabilities of dataset classes correspond to sums of probabilities of universal classes. This allows us to train a universal model on dataset-specific labels through negative log-likelihood of aggregated probability which we refer to as NLL+. Such training can be viewed as learning from partial labels since we train fine-grained universal classes on coarse-grained labels.

We compare our universal models to two baselines which ignore the overlap between dataset-specific logits. Our universal models outperform these two baselines in within-dataset generalization and cross-dataset validation. Experiments on relabeled Cityscapes indicate that our method is able to transcend semantics of individual datasets by learning an unlabeled universal concept at the intersection of overlapping dataset classes.

Validation on the MSeg collection shows that our method is competitive to standard NLL training on manually rela-

beled data. Our advantage becomes substantial when evaluating on original datasets where unified taxonomies have to drop classes in order to alleviate the relabeling effort.

Our method achieves the best aggregate performance on the 7 benchmarks from the RVC 2020 challenge. The corresponding universal taxonomy has only 192 classes, which is a considerable improvement over 311 classes in the naive concatenation. We see such multi-domain training as a method of choice for machine vision in uncontrolled environments. Indeed, our RVC models set a new state of the art on WildDash 2 — the toughest road-driving benchmark, while outperforming other approaches by a large margin.

The recovered universal models can be used as a tool for evaluating automatically recovered taxonomies. The authors would like to encourage further work on advanced evaluation datasets that will reward cross-dataset training and promote development of smarter and more robust visual systems. An ideal future dataset should combine rich taxonomies with domain shift, anomalous objects, out-of-domain context, and adversarial input. Other suitable avenues for future work include automatic recovery of universal taxonomies, extending the developed framework towards outlier detection and open-set recognition, and exploring knowledge-transfer potential of multi-domain universal models.

References

- [1] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(12):2481–2495, 2017. 6
- [2] Ping Chao, Chao-Yang Kao, Yu-Shan Ruan, Chien-Hsiang Huang, and Youn-Long Lin. Hardnet: A low memory traffic network. In *ICCV*, pages 3551–3560, 2019. 2
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 2
- [4] 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. 2
- [5] 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. 1, 6
- [6] Timothee Cour, Ben Sapp, and Ben Taskar. Learning from partial labels. *The Journal of Machine Learning Research*, 12:1501–1536, 2011. 1, 2
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 7
- [8] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *Int. J. Robotics Res.*, 32(11):1231–1237, 2013. 6
- [9] 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. 5
- [10] Yuanduo Hong, Huihui Pan, Weichao Sun, and Yisong Jia. Deep dual-resolution networks for real-time and accurate semantic segmentation of road scenes. *CoRR*, abs/2101.06085, 2021. 2
- [11] Gao Huang, Zhuang Liu, Geoff Pleiss, Laurens Van Der Maaten, and Kilian Weinberger. Convolutional networks with dense connectivity. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 5, 7
- [12] Tarun Kalluri, Girish Varma, Manmohan Chandraker, and CV Jawahar. Universal semi-supervised semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5259–5270, 2019. 2
- [13] Ivan Kreso, Josip Krapac, and Sinisa Segvic. Ladder-style densenets for semantic segmentation of large natural images. In *ICCV CVRSUAD*, pages 238–245, 2017. 2
- [14] Ivan Krešo, Josip Krapac, and Siniša Šegvić. Efficient ladder-style densenets for semantic segmentation of large images. *IEEE Transactions on Intelligent Transportation Systems*, 2020. 2
- [15] John Lambert, Zhuang Liu, Ozan Sener, James Hays, and Vladlen Koltun. Mseg: A composite dataset for multi-domain semantic segmentation. In *CVPR*, 2020. 1, 2, 3, 6, 7
- [16] Xiaodan Liang, Hongfei Zhou, and Eric Xing. Dynamic-structured semantic propagation network. In *CVPR*, pages 752–761, 2018. 1, 2
- [17] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, pages 740–755, 2014. 1, 6
- [18] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 2
- [19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [20] Panagiotis Meletis and Gijs Dubbelman. Training of convolutional networks on multiple heterogeneous datasets for street scene semantic segmentation. In *Intelligent Vehicles Symposium*, pages 1045–1050, 2018. 2
- [21] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. Mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, pages 5000–5009, 2017. 1, 6
- [22] Dong Nie, Jia Xue, and Xiaofeng Ren. Bidirectional pyramid networks for semantic segmentation. In *ACCV*, pages 654–671, 2020. 2
- [23] Marin Orsic, Petra Bevandic, Ivan Grubisic, Josip Saric, and Sinisa Segvic. Multi-domain semantic segmentation with pyramidal fusion. *arXiv preprint arXiv:2009.01636*, 2020. 7
- [24] Marin Oršić and Siniša Šegvić. Efficient semantic segmentation with pyramidal fusion. *Pattern Recognition*, page 107611, 2021. 2, 5, 6
- [25] Stephan R. Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *ICCV*, pages 2232–2241, 2017. 1
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2
- [27] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5639–5647, 2018. 2, 5, 7
- [28] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. Green AI. *Commun. ACM*, 63(12):54–63, 2020. 2
- [29] Shuran Song, Samuel P. Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, pages 567–576, 2015. 6

- [30] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 843–852. IEEE Computer Society, 2017. [1](#)
- [31] Girish Varma, Anbumani Subramanian, Anoop M. Nambodiri, Manmohan Chandraker, and C. V. Jawahar. IDD: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *WACV*, pages 1743–1751, 2019. [6](#)
- [32] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 2020. [2](#)
- [33] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. BDD100K: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018. [6](#)
- [34] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *ECCV*, 2020. [2](#)
- [35] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steiner, and Gustavo Fernandez Dominguez. Wilddash - creating hazard-aware benchmarks. In *ECCV*, 2018. [1](#), [5](#)
- [36] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [1](#)
- [37] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. Icnnet for real-time semantic segmentation on high-resolution images. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part III*, volume 11207 of *Lecture Notes in Computer Science*, pages 418–434. Springer, 2018. [2](#)
- [38] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [2](#)
- [39] Xiangyun Zhao, Samuel Schuster, Gaurav Sharma, Yi-Hsuan Tsai, Manmohan Chandraker, and Ying Wu. Object detection with a unified label space from multiple datasets. In *ECCV*, pages 178–193, 2020. [1](#), [2](#), [3](#), [7](#)
- [40] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, pages 633–641, 2017. [1](#), [6](#)
- [41] Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8856–8865, 2019. [2](#)
- [42] Aleksandar Zlateski, Ronnchai Jaroensri, Prafull Sharma, and Frédo Durand. On the importance of label quality for semantic segmentation. In *CVPR*, pages 1479–1487, 2018. [1](#)