
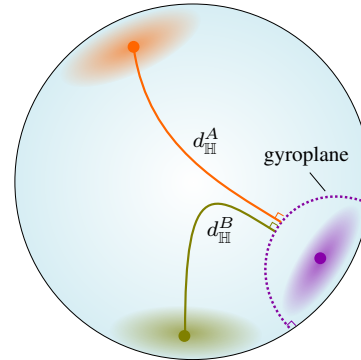


Flattening the Parent Bias: Hierarchical Semantic Segmentation in the Poincaré Ball

Simon Weber^{1,2}Barış Zöngür¹Nikita Araslanov^{1,2}Daniel Cremers^{1,2}¹Technical University of Munich²Munich Center for Machine Learning

(a) Embeddings in Euclidean space



(b) Embeddings in Poincaré ball

Figure 1. **Core idea.** Class embeddings in the Euclidean space (a) exhibit non-uniform properties of the separation margin: the average distance of a pixel embedding of one class to the decision boundaries of the other classes varies substantially (e.g. $d_E^A > d_E^B$). This creates an implicit *parent bias* in hierarchical segmentation, which prefers grouping one set of classes over the other, in terms of the parent-level segmentation accuracy. In contrast, in hyperbolic space characterized by the Poincaré ball (b), the separation margins between the class embeddings are more uniform, e.g. the hyperbolic distance of embeddings *A* and *B* of two different classes to the decision boundary (a gyroplane) of another class is approximately equal, $d_H^A \approx d_H^B$. This may explain the strong generalization of the parent-level predictions, observed in practice, in terms of the segmentation accuracy and calibration quality.

Abstract

Hierarchy is a natural representation of semantic taxonomies, including the ones routinely used in image segmentation. Indeed, recent work on semantic segmentation reports improved accuracy from supervised training leveraging hierarchical label structures. Encouraged by these results, we revisit the fundamental assumptions behind that work. We postulate and then empirically verify that the reasons for the observed improvement in segmentation accuracy may be entirely unrelated to the use of the semantic hierarchy. To demonstrate this, we design a range of cross-domain experiments with a representative hierarchical approach. We find that on the new testing domains, a flat (non-hierarchical) segmentation network, in which the parents are inferred from the children, has superior segmentation accuracy to the hierarchical approach across the board. Complementing these findings and inspired by the intrinsic properties of hyperbolic spaces, we study a more principled approach to hierarchical segmentation using the Poincaré

ball model. The hyperbolic representation largely outperforms the previous (Euclidean) hierarchical approach as well and is on par with our flat Euclidean baseline in terms of segmentation accuracy. However, it additionally exhibits surprisingly strong calibration quality of the parent nodes in the semantic hierarchy, especially on the more challenging domains. Our combined analysis suggests that the established practice of hierarchical segmentation may be limited to in-domain settings, whereas flat classifiers generalize substantially better, especially if they are modeled in the hyperbolic space.

1. Introduction

Semantic knowledge is inherently structured, and organizing it in a hierarchy is both natural and expressive. Unsurprisingly, hierarchical representations play an important role in computer vision [36, 37, 39, 59, 60, 62]. For in-

Project code: <https://github.com/tum-vision/hierahyp>

stance, we may want to assign multiple labels to each pixel in the image, rather than a single one, to encode ancestral relations (e.g. a “car” is also a “vehicle” and a “means of transport”). Adhering to a tree-based label hierarchy, this pixelwise classification task defines the so-called *hierarchical* semantic segmentation and is the subject of this work.

In the literature, recent work addresses this problem as a supervised multi-label classification task [34, 35]. In this formulation, the terminal leaf nodes and the internal nodes are modeled with individual one-*vs*-all classifiers. Remarkably, the empirical results of this approach appear to even exceed the standard supervised formulation (which only considers the leaf categories) in the evaluation of segmentation accuracy *over the leaf categories themselves*. Such an effect cannot be explained from the perspective of a learning algorithm, for which the semantic grouping of leaf nodes into parent classes is meaningless.

As our first step, we examine this phenomenon and reveal limited generalization of a state-of-the-art method for hierarchical semantic segmentation [34]. Rather surprisingly, we find that a *flat* classifier, which only learns to classify the child (leaf) categories, largely outperforms the more sophisticated prior art on *both the child and parent classes*. We formulate the sufficiency of flat classifiers under the existing formulation of the hierarchical semantic segmentation, which establishes the link between model calibration and accuracy on the hierarchical prediction task.

Moving forward, we identify an inherent bias of flat classifiers toward particular groupings of child categories into parent meta-classes, as illustrated in Fig. 1. Specifically, the Euclidean distance between a decision boundary of one class and the class embeddings of the other categories is non-uniform. However, classification errors of deep classifiers tend to occur near decision boundaries.¹ This implies that defining a parent class comprising the two classes with the lowest separation margin will tend to produce a lower error rate in the parent-level classification, if we were to combine two classes with the largest margin of separation. To mitigate this *parent bias*, we would like the decision boundary of any class to be equidistant to the embeddings of other classes. While additional regularization may be necessary to achieve this in the Euclidean space, we find that hyperbolic spaces provide such capacity naturally. We embed pixel features in the Poincaré ball instead of the Euclidean space, which allows us to alleviate the parent bias and achieve a notable improvement in segmentation accuracy and calibration on the parent-level classification task.

We summarize our contributions as follows. (i) We reveal limited generalization of prior work on the hierarchical semantic segmentation task. (ii) Through a systematic analysis, we establish the sufficiency of flat classifiers for this

¹This is evidenced by the imperfect, yet fairly high calibration quality of segmentation networks, as we will also show in the experiments.

task, which in Euclidean embedding space, however, may suffer from suboptimal accuracy on the parent classes. (iii) We show that the intrinsic properties of hyperbolic space, the Poincaré ball model, allow for mitigating this bias. (iv) We experimentally confirm the strong generalization of the Poincaré ball model, in terms of segmentation accuracy and calibration, especially on parent categories.

2. Related Work

Research on semantic segmentation spans numerous problem domains, including deep network architectures [10, 13, 46, 51], training objectives [5, 12, 52] and strategies [38, 41, 44]. Here, we are specifically interested in hierarchical semantic segmentation and, more generally, the hierarchical classification problem. Therefore, our review of related literature below will focus only on these aspects, and we refer interested readers to surveys for a comprehensive overview of semantic segmentation research [40].

Hierarchical classification with tree-like taxonomies. A multi-label classification problem is considered hierarchical if the label assignment must respect a pre-defined hierarchy [25, 29]. Hierarchical classifiers may be categorized into flat, local and global approaches. *Flat* classifiers only model the leaf nodes, thus completely ignoring the class hierarchy. Following the tree structure in the bottom-up fashion, one can infer the parent label from the predictions of its children. By contrast, *global* (or “big-bang”) methods explicitly represent each node in the tree, for example with a one-*vs*-all classifier per node [30]. Local approaches solve a number of smaller classification problems using only the local information available at each node or tree level [19, 32]. The success of these strategies appears to be domain-specific. However, it is notable that flat classifiers are generally seen as competitive baselines [4, 55, 61] – the conclusion reached in our study too. Learning individual classifiers for internal (non-leaf) nodes in the hierarchy may reduce semantically critical prediction errors [6, 22].

As a remark, hierarchical prediction has been the subject of research on problems in natural language processing and bioinformatics [29], where it is not uncommon to have large taxonomies – in the order of tens or hundreds of thousands of labels [45, 58]. By contrast, a typical size of the label space in computer vision is substantially smaller [18], especially for dense tasks, such as semantic segmentation considered here (e.g. up to 30 in Cityscapes [14]).

Hierarchical semantic segmentation. Considering a hierarchy of image segments is a classic concept in computer vision. Hierarchical image parsing helps to improve robustness to (self-)occlusions and to variation in object scale of early object recognition systems [2, 49, 54, 67]. Similar to conditional random fields (CRFs) [16], deep networks can also benefit from hierarchical representations for advanced contextual reasoning [24, 50, 63].



In contrast to the earlier work, where the hierarchy plays a facilitating role, training deep semantic segmentation networks producing a hierarchical label structure is relatively recent [34, 35]. HSSN [34], which we extensively use in our analysis, formalizes the hierarchical prediction task with auxiliary “parent” logits. Note that this implies a training objective with more decision boundaries to learn than in the standard (child-only) case, since each parent logit requires a one-*vs*-all classifier. In a follow-up work, LogicSeg [35] formulates Boolean rules describing the hierarchical constraints and maps them to a differentiable loss using fuzzy logic. While somewhat elegant, this approach does not improve over HSSN empirically in a significant way.

Hyperbolic computer vision. Deep learning on hyperbolic manifolds is still in its nascent stage [40]. In contrast to the Euclidean space, hyperbolic spaces possess properties lending themselves well to embedding hierarchical representations with minimal distortion [8, 43, 48]. Previous research concentrated on generalizing the Euclidean models operating on the hyperbolic manifold, in terms of network models [20, 23, 56] and training specifics [27]. Exploiting the properties of the hyperbolic embedding space has been of primary interest in (self-supervised) metric learning [28, 53]. Against the backdrop of this work, semantic segmentation has been studied rather marginally. In a seminal work in this domain, Atigh et al. [3] learn pixel embeddings on the Poincaré ball and reach competitive segmentation accuracy w.r.t. the more established Euclidean formulation. Concurrently, Franco et al. [21] report the correlation of the embedding norm with uncertainty in the context of active segmentation learning.

Overall, despite the recent progress, the benefits of the hyperbolic representation for semantic segmentation remain unclear. Our study of hierarchical semantic segmentation exemplifies some compelling advantages of the Poincaré ball model, both theoretically and experimentally.

3. The problem and motivation

Let us revisit the formulation of the hierarchical semantic segmentation problem from previous work [34, 35], which we follow in our study. Our label space is defined as the union of semantic categories at multiple levels of the semantic hierarchy, $\mathcal{S} := \cup_n \mathcal{S}_n$, where \mathcal{S}_0 defines the leaf classes. Learning a semantic segmentation model with the finest label space, \mathcal{S}_0 , reduces the problem to the conventional supervised scenario [10, 51], since it defines the granularity limit set by the available annotation in a given benchmark. In addition to \mathcal{S}_0 , we construct \mathcal{S}_1 by defining “meta-classes”, which *semantically* agglomerate one or more categories from \mathcal{S}_0 into a common parent class. In Cityscapes [14], for example, one defines a parent class “Human” comprising child classes “Person” and “Rider”. While one could create deep hierarchical structures, the limited annotation



(b) Segmentation accuracy does not change between semantic (Tree-A) and non-semantic hierarchies (Tree-B, Tree-C).

Figure 2. Training with non-semantic hierarchies. We train semantic segmentation models with non-semantic trees. (a) For example, we define class “Person” as a parent of a “vehicle”, which is clearly semantically meaningless. (b) Non-semantic hierarchies (“Tree-B” and “Tree-C”) do not affect the segmentation accuracy of semantic hierarchy (“Tree-A”) in a significant way.

in semantic segmentation only allows for hierarchies up to $n = 2$ – more rarely $n = 3$, in practice [34]. With the label hierarchy thus defined, our goal now is to maximize the segmentation accuracy (e.g. mIoU or mean pixel accuracy), evaluated separately for each level of the tree.

Empirical observations from previous work [34, 35] suggest that defining the meta-classes, e.g. \mathcal{S}_1 , improve the accuracy of the children categories \mathcal{S}_0 . *What explains this phenomenon?* After all, the only additional supervision signal in our new hierarchical formulation is the semantic proximity of some classes in \mathcal{S}_0 . However, this does not immediately render the learning problem easier. In fact, since we add an additional classification problem over categories in \mathcal{S}_1 , optimization may become even more difficult.

The hypothesis that the semantic proximity between child categories provides a complementary supervision signal asks for empirical validation. We design meta-classes in a semantically meaningless fashion and train a DeepLabv3+ with ResNet-101 backbone [10] following HSSN training objective [34] on Cityscapes *train* [14]. For example, we define “Car” as a child of “Human”, and “Terrain” as a sibling of “Sky”. Observing the results in Fig. 2b on Cityscapes *val*, we conclude that such semantically incoherent definitions of new meta-classes do not have much effect on the segmentation accuracy of the leaf categories (i.e. the standard 19 semantic classes in Cityscapes [14], which remained unchanged). This experiment suggests that the empirical benefits reported by Li et al. [34, 35] may not be related to the *semantic* definition of the label hierarchy.

3.1. Flat classifiers are strong baselines

As discussed in Sec. 2, *flat* classifiers offer a reasonable sanity check regardless of the problem domain. Indeed, we found that the HSSN model exhibits a strong bias towards the specific traffic scenes of the training domain (Cityscapes), while performing poorly on the novel domains. By contrast, a flat classifier consistently outperforms HSSN on the test datasets (*cf.* Tab. 2.). These results may not strike us as surprising and are in line with our intuition developed in the previous section. A less expected result, however, is that we also observed superior accuracy of flat classifiers on the parent categories. Recall that we represent the semantic hierarchical relations as a fixed *stationary* tree, defined by three properties: (i) Every level of the label hierarchy forms a categorical distribution. (ii) A prediction at a terminal node determines the complete hierarchical label (due to the uniqueness of the path to the root). (iii) The label hierarchy remains unchanged in training and testing. As we are dealing with a closed-world taxonomy, it is straightforward that the conditional probability of a parent class can be expressed with only the conditional probabilities of its children, which we formalize as follows:

Proposition 1. *Let y be the parent class and m_y denote the children of that class in image \mathcal{I} . Given a stationary tree defined above, the optimal class posterior $p(\hat{y} = y \mid \mathcal{I})$ for a parent node y with child class posterior $p(\hat{m}_y = m_y \mid \mathcal{I})$ is given by*

$$\begin{aligned} p(\hat{y} = y \mid \mathcal{I}) &= \sum_m p(\hat{y} = y \mid m) p(\hat{m} = m \mid \mathcal{I}) \\ &= \sum_{m_y} p(\hat{m}_y = m_y \mid \mathcal{I}). \end{aligned} \quad (1)$$

Note that prior work on hierarchical semantic segmentation [34, 35] independently models the parent posterior $p(\hat{y} = y \mid \mathcal{I})$. The above proposition states the sufficiency of predicting the child nodes independently from the parent nodes, and then inferring the parent posterior with Eq. (1).

3.2. The parent bias in Euclidean space

Class embeddings produced in the Euclidean space tend to produce a multi-modal distribution, where each class forms an independent mode (a cluster). The modes exhibit a particular rank-based arrangement and we can loosely establish that, for example, class “A” is closer to class “B”, in terms of the Euclidean distance (see *supp.* material for empirical support). Since typical classification errors occur at the decision boundaries, the spatial proximity of two class embeddings presents an inherent bias of that embedding space. For instance, if classes “A” and “B” are close in the embedding space and end up in different parent categories, this will lead to suboptimal accuracy on the parent level. Without any prior on the parent taxonomy, which is

task-specific and would lead to the parent bias, we can encourage the modes of our class embeddings to be approximately equidistant. In the Euclidean space, this would require additional regularization. However, it can emerge naturally in hyperbolic space – in the Poincaré ball model.

4. From Euclidean to hyperbolic geometry

Segmentation in Euclidean space. Let us formalize the training process of a deep network for semantic segmentation in Euclidean space. Given an image $\mathcal{I} \in \mathbb{R}^{H \times W \times 3}$, we would like to predict label $l \in \mathcal{S}$ for each pixel $i \in \{1, \dots, HW\}$, where \mathcal{S} is the set of $|\mathcal{S}|$ class labels. An encoder f_θ with parameters θ maps \mathcal{I} to a set of pixel feature embeddings $\mathbf{X} = (\mathbf{x}_i)_{i=1}^{HW} = f(\mathcal{I}) \in \mathbb{R}^{HW \times d}$. In the last layer, for each pixel i , we obtain a segmentation map, modeled as the class posterior,

$$p(\hat{l} = l \mid \mathbf{x}_i) \propto \exp(a_l^\top \mathbf{x}_i + b_l). \quad (2)$$

Here, $\{(a_l, b_l)\}_{l=1}^{|\mathcal{S}|}$ are classifier parameters and define hyperplanes for each class $l \in \mathcal{S}$ in the Euclidean space. During training, we jointly optimize for $\Theta := \{\theta, \{(a_l, b_l)\}_{l=1}^{|\mathcal{S}|}\}$ with backpropagation, minimizing the (expected) cross-entropy loss for each pixel i ,

$$\min_{\Theta} \mathbb{E}_{\mathcal{I}} [-\log p(\hat{l} = l_i^* \mid \mathbf{x}_i)], \quad (3)$$

where l_i^* is the ground-truth label for pixel i .

By analogy with the Euclidean setup, the per-pixel classification in the hyperbolic space involves gyroplanes (*cf.* Fig. 1), defined by offsets and normals [23]. Let us revisit some basic notions of hyperbolic geometry.

4.1. The Poincaré ball model

Poincaré Ball and Exponential Map. Hyperbolic geometry can be expressed with different conformal models [9]. We operate on the Poincaré ball $(\mathbb{D}_c^n, g^{\mathbb{D}_c^n})$ with $-c$ denoting the negative curvature, and $g^{\mathbb{D}_c^n}$ being the Riemannian metric associated with the manifold $\mathbb{D}_c^n = \{x \in \mathbb{R}^n \mid c\|x\| < 1\}$. $g^{\mathbb{D}_c^n}$ can be linked to the Euclidean metric tensor $g^{\mathbb{E}} = \mathbf{I}_n$ via:

$$g^{\mathbb{D}_c^n} = (\lambda_x^c)^2 g^{\mathbb{E}} = \left(\frac{2}{1 - c\|x\|^2} \right)^2 g^{\mathbb{E}}. \quad (4)$$

The exponential map between the Euclidean space \mathbb{R}^n and the Poincaré ball \mathbb{D}_c^n with anchor v is defined as:

$$\text{Exp}_v^c(x) = v \oplus_c \left(\tanh\left(\sqrt{c} \frac{\lambda_v^c \|x\|}{2}\right) \frac{x}{\sqrt{c\|x\|}} \right), \quad (5)$$

where \oplus_c is the Möbius hyperbolic addition defined as:

$$v \oplus_c w = \frac{(1 + 2c\langle v, w \rangle + c\|w\|^2)v + (1 - c\|v\|^2)w}{1 + 2c\langle v, w \rangle + c^2\|v\|^2\|w\|^2}, \quad (6)$$

for all $v, w \in \mathbb{D}_c^n$. For simplicity, v is set to the origin 0 and then the considered exponential map is:

$$\text{Exp}_0^c(x) = \tanh(\sqrt{c}\|x\|) \frac{x}{\sqrt{c}\|x\|}. \quad (7)$$

Hyperbolic Distance. The hyperbolic distance between x and z on the Poincaré ball is given by:

$$d_{\mathbb{H}}(x, z) = \text{arcosh} \left(1 + 2 \frac{\|x - z\|^2}{(1 - \|x\|^2)(1 - \|z\|^2)} \right). \quad (8)$$

Hyperbolic Multinomial Logistic Regression (MLR). Given a hyperbolic vector h and K classes, Ganea et al. [23] provide a geometric interpretation of the hyperbolic MLR by defining gyroplanes as:

$$H_k^c = \{h \in \mathbb{D}_c^n \mid \langle -r_k \oplus_c h, a_k \rangle = 0\}, \quad (9)$$

where $k \in \{1, \dots, K\}$ and r_k and a_k are respectively the gyroplane offset and the normal associated with class k . The hyperbolic distance between h and the gyroplane of class k is given by:

$$d_{\mathbb{H}}(h, H_y^c) = \frac{1}{\sqrt{c}} \text{arcsinh} \left(\frac{2\sqrt{c}|\langle -r_y \oplus_c h, w_y \rangle|}{(1 - c\| -r_y \oplus_c h\|^2)\|w_y\|} \right). \quad (10)$$

Based on this distance, we define the hyperbolic logit:

$$\zeta_y(h) := \frac{\lambda_{r_y}^c \|w_y\|}{\sqrt{c}} \text{arcsinh} \left(\frac{2\sqrt{c}\langle -r_y \oplus_c h, w_y \rangle}{(1 - c\| -r_y \oplus_c h\|^2)\|w_y\|} \right), \quad (11)$$

and the likelihood,

$$p(\hat{y} := y \mid h) \propto \exp(\zeta_y(h)). \quad (12)$$

4.2. Calibrated segmentation in the Poincaré ball

Following Atigh et al. [3], we perform the per-pixel classification in the hyperbolic space. We project \mathbf{X} onto the Poincaré ball with the mapping $\text{Exp}_0^c(\cdot)$ to get the hyperbolic embedding $\mathbf{H} \in \mathbb{D}^{H \times W \times n}$. Then, we optimize the obtained likelihood (cf. Eq. (12)) with the standard cross-entropy loss. Once the classification is performed in the Poincaré ball, we link the hyperbolic logits (Eq. (11)) to the confidence of the prediction, by analogy with Euclidean networks [26]. To our knowledge, while this extension of Euclidean calibration to hyperbolic networks is straightforward, we are the first to present its experimental analysis. Nevertheless, the analogy with Euclidean space has its limitations. Focusing on the hyperbolic distance, we demonstrate its concave property w.r.t. Euclidean distance. This property allows us to establish a distinguishing feature of the hyperbolic space in modeling inter-class relationships.

4.3. From concavity to flattening bias

On the hyperbolic distance. Hyperbolic geometry naturally embeds hierarchical structure [48]. However, we argue




Figure 3. **Hyperbolic distance exhibits strict concavity w.r.t. Euclidean distance.** Observe (assuming $\alpha = 1$ for simplicity) that the hyperbolic distance has sublinear (logarithmic) growth. In practical terms, this implies that the difference in the distance between two hyperbolic representations as x increases will diminish, while it remains constant in Euclidean space.

here that the hyperbolic space lends itself well also for *flat* classification. We observe that during training, the hyperbolic embeddings of the same class tend to be pushed to the periphery of the Poincaré ball and onto the same side of the associated gyroplane [3]. Studying inter-class hyperbolic distance (cf. Eq. (10)), we further find that embeddings of one class are approximately equidistant to the embeddings of any other class. This implies that, in contrast to the Euclidean space, there is no parent bias in the Poincaré ball, *i.e.* there is no preference for a specific grouping of child categories into parents. The following proposition provides the formal argument explaining our observation:

Proposition 2. *The hyperbolic distance between two embeddings h_1, h_2 is strictly concave in the Euclidean distance between h_1 and h_2 .*

Proof. We can write the hyperbolic distance $d_{\mathbb{H}}$ between two embeddings in the Poincaré ball as a function of the Euclidean distance $d_{\mathbb{E}}(h_1, h_2)$. The derivative of $d_{\mathbb{H}}$ with respect to $d_{\mathbb{E}}(h_1, h_2)$ is then

$$\frac{2}{\sqrt{(1 - \|h_1\|^2)(1 - \|h_2\|^2)} \sqrt{1 + \frac{d_{\mathbb{E}}^2}{(1 - \|h_1\|^2)(1 - \|h_2\|^2)}}}, \quad (13)$$

which is strictly decreasing in $d_{\mathbb{E}}$. \square

Fig. 3 illustrates this proposition. Given the concave property of the hyperbolic distance, we postulate that the Poincaré ball formulation facilitates distance uniformity between classes. Since the hyperbolic embeddings are pushed to the border of the Poincaré ball during training, the norms $\|h_1\|$ and $\|h_2\|$ in Eq. (8) are close to 1. Therefore, we operate at the high-end spectrum of the domain in Fig. 3 (see supp. material for further details and empirical support).

| Dataset | Method | cwECE | | | |
|------------|-----------------|-------------|--------------|--------------|--------------|
| | | Level 1 | | Level 0 | |
| Cityscapes | HSSN* [34] | 0.90 | – | 5.97 | – |
| | HSSN [34] | 0.79 | 0.79 | 6.85 | 5.09 |
| | Flat-Euc (ours) | 0.52 | 0.60 | 4.40 | 3.97 |
| | Flat-Hyp (ours) | 0.66 | 0.73 | 4.35 | 4.09 |
| Mapillary | HSSN* [34] | 4.26 | – | 17.39 | – |
| | HSSN [34] | 4.47 | 4.35 | 18.44 | 15.59 |
| | Flat-Euc (ours) | 2.84 | 3.62 | 11.03 | 11.26 |
| | Flat-Hyp (ours) | 2.88 | 3.38 | 11.44 | 11.42 |
| IDD | HSSN* [34] | 7.85 | – | 17.74 | – |
| | HSSN [34] | 7.86 | 7.56 | 19.39 | 17.24 |
| | Flat-Euc (ours) | 5.94 | 7.24 | 12.77 | 13.43 |
| | Flat-Hyp (ours) | 6.05 | 6.27 | 13.37 | 10.87 |
| ACDC | HSSN* [34] | 8.12 | – | 23.54 | – |
| | HSSN [34] | 6.78 | 8.40 | 21.18 | 20.92 |
| | Flat-Euc (ours) | 6.55 | 5.96 | 16.63 | 16.49 |
| | Flat-Hyp (ours) | 6.34 | 6.03 | 17.58 | 12.64 |
| BDD | HSSN* [34] | 8.68 | – | 26.37 | – |
| | HSSN [34] | 8.82 | 6.77 | 29.13 | 23.71 |
| | Flat-Euc (ours) | 7.40 | 7.05 | 21.44 | 20.95 |
| | Flat-Hyp (ours) | 6.55 | 6.09 | 20.71 | 18.58 |
| Wilddash | HSSN* [34] | 14.28 | – | 32.32 | – |
| | HSSN [34] | 14.95 | 11.48 | 33.21 | 28.23 |
| | Flat-Euc (ours) | 13.53 | 13.16 | 23.25 | 23.93 |
| | Flat-Hyp (ours) | 9.52 | 10.04 | 20.17 | 18.45 |

Table 1. **Calibration quality (cwECE).** We train DeepLabV3+/ResNet-101 and OCRNet/HRNet-W48 on Cityscapes (train) and report the results for six datasets. HSSN* corresponds to the pretrained model provided by the authors, only available for DeepLabV3+.

Consequently, changes in the class embeddings have a diminished effect on the distance in the Poincaré ball, whereas in the Euclidean space, this relationship is linear. Demonstrating a practical consequence, the next section experimentally confirms that the hyperbolic space leads to strong parent-level generalization in semantic segmentation.

5. Experiments

Datasets. Different from prior work, we perform our analysis by testing the models on 6 datasets: Cityscapes [15], Mapillary [42], IDD [57], ACDC [47], BDD [64] and Wilddash [66]. Note that since we train our models on Cityscapes, the datasets with a larger visual domain shift (ACDC, BDD and Wilddash) are more challenging.

Metrics. We compare the models in terms of segmentation accuracy and calibration quality, on both the child and parent nodes (19 and 7 classes, respectively). To evaluate the accuracy, we use the mean Intersection-over-Union (mIoU), the mean accuracy over classes (mAcc) and the average pixel accuracy (aAcc). We follow Kull et al. [33] to derive a calibration metric by comparing the difference between the confidence and accuracy of class predictions. We

report the class-wise expected calibration error (cwECE).

Models. We use DeepLabV3+ with ResNet-101 backbone [11] and OCRNet with HRNet-W48 backbone [65], with the backbones pre-trained on ImageNet [17]. We train each model on Cityscapes *train* with fine annotations [15] for 80K iterations. Leveraging the child class posterior probabilities (Level 0), flat models compute the parent posterior (Level 1) using Eq. (1).

Implementation details. For a fair comparison, we follow HSSN [34] to set the training hyper-parameters. For the hyperbolic networks, we use the optimization method in [23]. We optimize the offsets with Riemannian SGD [7], and set the learning rate to 0.0001. The normals are optimized with SGD in the Euclidean space, with a learning rate 0.001. The projection onto the Poincaré ball uses the Geopt library [31], setting curvature to $c = 1$.

5.1. Quantitative results

Tab. 1 and Tab. 2 report the calibration quality and the segmentation accuracy, respectively, for the hierarchical training (HSSN [34]), and the flat Euclidean (Flat-Euc) and hyperbolic models (Flat-Hyp). To ensure a fair comparison, we train all models with a consistent codebase and training schedule. For reference, we also report the results for HSSN* with DeepLabV3+ using the weights provided by the authors [1]. A pre-trained model for OCRNet is not available. Level 1 and Level 0 refer to the parent (classes S_1) and child (classes S_0) nodes in the hierarchical tree, respectively. In the context of our study, we are particularly interested in Level 1.

Calibration quality. Referencing Tab. 1, we inspect the calibration quality of DeepLabV3+/ResNet-101 (shaded in red). For child nodes (Level 0), the flat classifiers exhibit calibration quality on par with or better than HSSN. For parent nodes (Level 1), Flat-Hyp is better calibrated than Flat-Euc on challenging datasets ACDC/BDD/Wilddash (-0.21/-0.85/-4.01). Notably, the gap grows toward the most challenging testbeds, BDD and Wilddash. The Poincaré ball model outperforms HSSN on five datasets, and, notably, for the most challenging datasets – BDD and Wilddash (-2.27/-5.43 w.r.t. HSSN).

Similarly for OCRNet/HRNet-W48 (shaded in blue, cf. Tab. 1), for child nodes (Level 0), the Poincaré ball is better calibrated by a large margin for the datasets with a large domain shift, IDD/ACDC/BDD/Wilddash (-2.56/-3.85/-2.37/-5.48 w.r.t. Flat-Euc; -6.37/-8.28/-5.13/-9.78 w.r.t. HSSN). For parent nodes (Level 1), Flat-Hyp is better calibrated than its competitors for datasets Mapillary/IDD/BDD/Wilddash (-0.24/-0.97/-1.04/-3.12 w.r.t. Flat-Euc; -0.97/-1.29/-0.68/-1.44 w.r.t. HSSN).

Segmentation accuracy. Let us examine the segmentation accuracy of DeepLabV3+/ResNet-101 in Tab. 2. For child nodes (Level 0), flat classifiers Flat-Euc

| Dataset | Method | Level 1 | | | | | | Level 0 | | | | | |
|------------|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | | mIoU | | mAcc | | aAcc | | mIoU | | mAcc | | aAcc | |
| | | | | | | | | | | | | | |
| Cityscapes | HSSN* [34] | 90.82 | – | 94.92 | – | 97.35 | – | 81.62 | – | 87.90 | – | 96.16 | – |
| | HSSN [34] | 90.68 | 90.27 | 94.47 | 94.25 | 97.29 | 97.20 | 80.30 | 80.21 | 86.11 | 86.62 | 95.93 | 95.98 |
| | Flat-Euc (ours) | 90.96 | 90.57 | 95.24 | 94.80 | 97.64 | 97.52 | 80.89 | 79.82 | 87.53 | 87.34 | 96.56 | 96.27 |
| | Flat-Hyp (ours) | 90.91 | 90.47 | 95.08 | 94.67 | 97.60 | 97.48 | 80.36 | 79.28 | 86.83 | 86.80 | 96.41 | 96.26 |
| Mapillary | HSSN* [34] | 83.26 | – | 89.37 | – | 94.70 | – | 62.32 | – | 72.53 | – | 90.37 | – |
| | HSSN [34] | 81.89 | 79.09 | 88.35 | 86.19 | 93.77 | 91.91 | 59.34 | 59.09 | 70.37 | 70.62 | 89.53 | 87.54 |
| | Flat-Euc (ours) | 83.11 | 80.27 | 90.73 | 89.14 | 93.98 | 93.7 | 63.94 | 60.47 | 76.44 | 74.62 | 90.47 | 90.12 |
| | Flat-Hyp (ours) | 81.87 | 81.89 | 90.15 | 89.23 | 93.41 | 94.9 | 60.34 | 58.66 | 74.88 | 73.55 | 89.60 | 90.96 |
| IDD | HSSN* [34] | 79.03 | – | 84.37 | – | 94.46 | – | 58.33 | – | 70.05 | – | 91.57 | – |
| | HSSN [34] | 78.52 | 77.69 | 84.00 | 83.01 | 94.45 | 94.29 | 55.21 | 58.14 | 67.37 | 67.21 | 90.57 | 91.79 |
| | Flat-Euc (ours) | 81.27 | 79.30 | 87.08 | 85.03 | 95.29 | 94.90 | 61.64 | 58.50 | 73.70 | 71.65 | 92.22 | 91.89 |
| | Flat-Hyp (ours) | 80.98 | 79.83 | 86.70 | 85.25 | 95.54 | 95.27 | 58.76 | 59.01 | 71.55 | 72.88 | 91.97 | 92.14 |
| ACDC | HSSN* [34] | 65.56 | – | 78.60 | – | 86.62 | – | 42.97 | – | 57.33 | – | 81.62 | – |
| | HSSN [34] | 73.45 | 69.43 | 82.00 | 78.78 | 90.64 | 87.41 | 52.71 | 49.20 | 62.62 | 61.40 | 84.95 | 82.10 |
| | Flat-Euc (ours) | 75.18 | 71.29 | 83.86 | 83.47 | 91.80 | 90.38 | 54.34 | 51.90 | 66.75 | 64.56 | 86.79 | 85.99 |
| | Flat-Hyp (ours) | 72.90 | 70.88 | 83.66 | 82.58 | 91.78 | 90.20 | 47.72 | 49.31 | 62.46 | 67.19 | 85.37 | 85.85 |
| BDD | HSSN* [34] | 73.53 | – | 82.25 | – | 89.60 | – | 48.32 | – | 60.54 | – | 86.76 | – |
| | HSSN [34] | 74.28 | 73.66 | 81.88 | 81.10 | 90.11 | 90.08 | 47.75 | 48.22 | 57.96 | 60.18 | 86.85 | 87.26 |
| | Flat-Euc (ours) | 76.27 | 74.47 | 84.64 | 83.52 | 91.22 | 90.62 | 51.54 | 50.04 | 64.58 | 63.12 | 88.52 | 88.05 |
| | Flat-Hyp (ours) | 76.49 | 76.62 | 85.12 | 84.34 | 91.91 | 92.12 | 49.64 | 49.60 | 63.21 | 64.37 | 89.00 | 89.31 |
| Wilddash | HSSN* [34] | 57.20 | – | 71.42 | – | 76.85 | – | 36.55 | – | 50.61 | – | 71.62 | – |
| | HSSN [34] | 58.60 | 59.80 | 71.65 | 71.22 | 76.49 | 79.74 | 37.07 | 39.01 | 50.03 | 50.86 | 71.12 | 74.42 |
| | Flat-Euc (ours) | 58.98 | 57.11 | 74.15 | 72.64 | 76.75 | 75.74 | 39.38 | 39.97 | 55.97 | 54.31 | 71.84 | 70.84 |
| | Flat-Hyp (ours) | 62.53 | 62.52 | 77.02 | 75.77 | 81.91 | 81.88 | 40.57 | 39.48 | 57.17 | 57.57 | 76.29 | 76.35 |

Table 2. **Segmentation accuracy (mIoU, mAcc).** We train DeepLabv3+/ResNet-101 and OCRNet/HRNet-W48 on Cityscapes *train* and test them on six datasets. HSSN* corresponds to the pretrained model given by the authors, only available for DeepLabV3+.

and Flat-Hyp substantially outperform HSSN on Mapillary/IDD/ACDC/BDD/Wilddash, in terms of mAcc, aAcc and mIoU. For parent nodes (Level 1), the Poincaré ball model outperforms the Euclidean model and HSSN, in terms of aAcc for Mapillary/IDD/BDD/Wilddash. For the most challenging datasets, BDD/Wilddash, Flat-Hyp clearly reaches the best results in terms of mIoU (+0.22/+3.55 w.r.t. Flat-Euc; +2.21/+3.93 w.r.t. HSSN) and mAcc (+0.48/+2.87 w.r.t. Flat-Euc, +3.24/+5.37 w.r.t. HSSN). Even on the less challenging datasets (Cityscapes/Mapillary/IDD/ACDC), where Flat-Euc outperforms the hyperbolic model on the parent-level predictions, the difference in mIoU between the child and parent nodes in the Euclidean and the hyperbolic model reduces (from -0.53/-3.60/-2.88/-6.62 to -0.07/-1.24/-0.29/-2.28) in favor of the hyperbolic model. This observation supports our analytical analysis of the parent bias in Sec. 4.3.

Similarly for OCRNet/HRNet-W48, we observe that Flat-Hyp shows the best mAcc for IDD/ACDC/BDD/Wilddash on Level 0, by a significant margin. For parent nodes (Level 1), Flat-Hyp outperforms hierarchical training and Flat-Euc in terms of aAcc on Mapillary/IDD/BDD/Wilddash. For the most challenging datasets, Flat-Hyp also clearly outperforms its competitors in terms of mIoU (+2.15/+5.41 w.r.t. the Flat-Euc;

+2.96/+2.72 w.r.t. HSSN, on BDD/Wilddash) and mAcc (+0.82/+3.13 w.r.t. Flat-Euc; +3.24/+4.55 w.r.t. HSSN, on BDD/Wilddash). Similarly to DeepLabV3+, on less challenging datasets, Cityscapes/Mapillary/IDD/ACDC, the difference in mIoU scores from child to parent levels is larger for Flat-Hyp. Unlike DeepLabV3+, the hyperbolic model outperforms the Euclidean model (+1.62) on the Mapillary dataset.

5.2. Qualitative results

In Fig. 4, we visualize an example of semantic segmentation on Level 1, comparing HSSN to our Euclidean and hyperbolic networks. We observe that HSSN mislabels parts of the “Building” as a “Vehicle”. Notably, the confidence of this incorrect prediction is high. By contrast, both Euclidean and the hyperbolic networks largely predict the “Building” correctly, although with higher uncertainty than HSSN. Rather curiously, the hyperbolic network exhibits a spatially smoother confidence map, which suggests a higher leverage of spatial correlations in the Poincaré ball model.

5.3. Discussion

Overall, the empirical results of segmentation accuracy and calibration strongly support our analysis in Sec. 3 and Sec. 4.1. On the one hand, the hierarchical training in HSSN




Figure 4. **Qualitative results on Willdash.** We show parent-level segmentation results for HSSN (left), Flat-Euc (middle), and Flat-Hyp (right) models by using label predictions (top), pixel-level accuracy maps (middle row), and confidence maps (bottom).

leads to poor performance on novel domains, presumably due to the large bias from the training on Cityscapes. By contrast, our flat classifiers, be they Euclidean or hyperbolic, are much more resilient to the domain shift, both in terms of calibration quality and segmentation accuracy, which is in line with our intuitive analysis in Sec. 3. The hyperbolic model exhibits a larger gap between the accuracy of the child and the parent nodes, compared to the Euclidean model. In cases, where both have comparable scores for child nodes, the hyperbolic model significantly outperforms the Euclidean model for parent nodes, and similarly for calibration. This suggests the advantage of the Poincaré ball in flattening the Euclidean parent bias, as we conjectured in Sec. 4.3. Importantly, these conclusions are consistent for both DeepLabV3+ and OCRNet.

Limitations. We have only evaluated on datasets with a strong focus on traffic scenes. This limitation comes from the lack of datasets with other taxonomic structures and of comparable visual complexity. Our analysis also focuses on the Poincaré ball, although this is not the only possible realization of the hyperbolic space. An extension of our study to other conformal models, such as the Lorentz model, offers exciting avenues for future research.

6. Conclusion

Our empirical investigations show that semantically meaningful hierarchical relations might not be the primary driver of the reported improvements in segmentation accuracy. *Contrario*, our cross-domain experiments reveal that flat segmentation networks, in which the parent categories are inferred from children, outperform hierarchical approaches consistently, for parent nodes most notably. However, we show that flat classifiers may suffer from the parent bias in the Euclidean space. By contrast, the Poincaré ball model exhibits more uniform properties between class representations. We demonstrate that a flat hyperbolic classifier coupled with a straightforward bottom-up inference generalizes surprisingly well across unseen test domains. It tends to outperform, for parent categories, the Euclidean representation in terms of the segmentation accuracy and calibration on the more challenging datasets. To our knowledge, our work is also the first empirical analysis of dense calibration with hyperbolic networks. We hope that our study will encourage future efforts toward more accurate and calibrated semantic segmentation models, extending beyond the currently mainstream Euclidean representation.

Acknowledgement. This work was supported by the ERC Advanced Grant SIMULACRON.

References

- [1] <https://github.com/lingorX/HieraSeg/tree/main/Pytorch>. [Online; accessed 28-March-2024]. 6
- [2] Pablo Arbelaez, Michael Maire, Charless C. Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE TPAMI*, 33(5):898–916, 2011. 2
- [3] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Noord, and Pascal Mettes. Hyperbolic image segmentation. In *CVPR*, 2022. 3, 5
- [4] Rohit Babbar, Ioannis Partalas, Éric Gaussier, and Massih-Reza Amini. On flat versus hierarchical classification in large-scale taxonomies. In *NIPS*, 2013. 2
- [5] Maxim Berman, Amal Rannen Triki, and Matthew B. Blaschko. The Lovász-Softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *CVPR*, 2018. 2
- [6] Luca Bertinetto, Romain Müller, Konstantinos Tertikas, Sina Samangooei, and Nicholas A. Lord. Making better mistakes: Leveraging class hierarchies with deep networks. In *CVPR*, 2020. 2
- [7] Silvere Bonnabel. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013. 6
- [8] Martin R Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer Science & Business Media, 2013. 3
- [9] James W. Cannon, William J. Floyd, Walter R. Parry, and et al. Hyperbolic geometry. In *Flavors of Geometry*, 1997. 4
- [10] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 2, 3
- [11] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6
- [12] Bowen Cheng, Ross B. Girshick, Piotr Dollár, Alexander C. Berg, and Alexander Kirillov. Boundary iou: Improving object-centric image segmentation evaluation. In *CVPR*, 2021. 2
- [13] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021. 2
- [14] 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 *CVPR*, 2016. 2, 3
- [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 *CVPR*, 2016. 6
- [16] Marius Cordts, Timo Rehfeld, Markus Enzweiler, Uwe Franke, and Stefan Roth. Tree-structured models for efficient multi-cue scene labeling. *IEEE TPAMI*, 39(7):1444–1454, 2017. 2
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [18] Ivica Dimitrovski, Dragi Kocev, Suzana Loskovska, and Saso Dzeroski. Hierarchical annotation of medical images. *Pattern Recognit.*, 44(10-11):2436–2449, 2011. 2
- [19] Roman Eisner, Brett Poulin, Duane Szafron, Paul Lu, and Russell Greiner. Improving protein function prediction using the hierarchical structure of the gene ontology. In *IEEE CIBCB*, 2005. 2
- [20] Aleksandr Ermolov, Leyla Mirvakhabova, Valentin Khruikov, Nicu Sebe, and Ivan V. Oseledets. Hyperbolic vision transformers: Combining improvements in metric learning. In *CVPR*, 2022. 3
- [21] Luca Franco, Paolo Mandica, Konstantinos Kallidromitis, Devin Guillory, Yu-Teng Li, and Fabio Galasso. Hyperbolic active learning for semantic segmentation under domain shift. *arXiv:2306.11180 [cs.CV]*, 2023. 3
- [22] Andrea Frome, Gregory S. Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomáš Mikolov. Devise: A deep visual-semantic embedding model. In *NIPS*, 2013. 2
- [23] Octavian-Eugen Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. In *NIPS*, 2018. 3, 4, 5, 6
- [24] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 2
- [25] Allan D Gordon. A review of hierarchical classification. *Journal of the Royal Statistical Society: Series A (General)*, 150(2):119–137, 1987. 2
- [26] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *CVPR*, 2016. 5
- [27] Yunhui Guo, Xudong Wang, Yubei Chen, and Stella X. Yu. Clipped hyperbolic classifiers are super-hyperbolic classifiers. In *CVPR*, 2022. 3
- [28] Joy Hsu, Jeffrey Gu, Gong Her Wu, Wah Chiu, and Serena Yeung. Capturing implicit hierarchical structure in 3D biomedical images with self-supervised hyperbolic representations. In *NeurIPS*, 2021. 3
- [29] Carlos Nascimento Silla Jr. and Alex Alves Freitas. A survey of hierarchical classification across different application domains. *Data Min. Knowl. Discov.*, 22(1-2):31–72, 2011. 2
- [30] Svetlana Kiritchenko, Stan Matwin, Richard Nock, and A. Fazel Famili. Learning and evaluation in the presence of class hierarchies: Application to text categorization. In *Adv. in Art. Intell.*, 2006. 2
- [31] Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian optimization in pytorch, 2020. 6
- [32] Daphne Koller and Mehran Sahami. Hierarchically classifying documents using very few words. In *ICML*, 1997. 2
- [33] Meelis Kull, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. *NIPS*, 32, 2019. 6

- [34] Liulei Li, Tianfei Zhou, Wenguan Wang, Jianwu Li, and Yi Yang. Deep hierarchical semantic segmentation. In *CVPR*, 2022. 2, 3, 4, 6, 7
- [35] Liulei Li, Wenguan Wang, and Yi Yang. LogicSeg: Parsing visual semantics with neural logic learning and reasoning. In *ICCV*, 2023. 2, 3, 4
- [36] Zhiheng Li, Wenxuan Bao, Jiayang Zheng, and Chenliang Xu. Deep grouping model for unified perceptual parsing. In *CVPR*, 2020. 1
- [37] Xiaodan Liang, Hongfei Zhou, and Eric Xing. Dynamic-structured semantic propagation network. In *CVPR*, 2018. 1
- [38] Pauline Luc, Camille Couprie, Soumith Chintala, and Jakob Verbeek. Semantic segmentation using adversarial networks. *arXiv:1611.08408 [cs.CV]*, 2016. 2
- [39] Panagiotis Meletis and Gijs Dubbelman. Training of convolutional networks on multiple heterogeneous datasets for street scene semantic segmentation. In *IEEE IV*, 2018. 1
- [40] Pascal Mettes, Mina Ghadimi Atigh, Martin Keller-Ressel, Jeffrey Gu, and Serena Yeung. Hyperbolic deep learning in computer vision: A survey. *arXiv:2305.06611 [cs.CV]*, 2023. 2, 3
- [41] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high- and low-level consistency. *IEEE TPAMI*, 43(4):1369–1379, 2021. 2
- [42] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, pages 4990–4999, 2017. 6
- [43] Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *NIPS*, 2017. 3
- [44] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *CVPR*, 2020. 2
- [45] Ioannis Partalas, Aris Kosmopoulos, Nicolas Baskiotis, Thierry Artières, George Paliouras, Éric Gaussier, Ion Androutsopoulos, Massih-Reza Amini, and Patrick Gallinari. LSHTC: A benchmark for large-scale text classification. *arXiv:1503.08581 [cs.IR]*, 2015. 2
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 2
- [47] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *ICCV*, 2021. 6
- [48] Frederic Sala, Christopher De Sa, Albert Gu, and Christopher Ré. Representation tradeoffs for hyperbolic embeddings. In *ICML*, 2018. 3, 5
- [49] Paul Schnitzspan, Mario Fritz, Stefan Roth, and Bernt Schiele. Discriminative structure learning of hierarchical representations for object detection. In *CVPR*, 2009. 2
- [50] Abhishek Sharma, Oncel Tuzel, and David W. Jacobs. Deep hierarchical parsing for semantic segmentation. In *CVPR*, 2015. 2
- [51] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE TPAMI*, 39(4):640–651, 2017. 2, 3
- [52] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sébastien Ourselin, and M. Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *MICCAI Workshops*, 2017. 2
- [53] Didac Suris, Ruoshi Liu, and Carl Vondrick. Learning the predictability of the future. In *CVPR*, 2021. 3
- [54] Jasper R. R. Uijlings, Koen E. A. van de Sande, Theo Gevers, and Arnold W. M. Smeulders. Selective search for object recognition. *IJCV*, 104(2):154–171, 2013. 2
- [55] Jack Valmadre. Hierarchical classification at multiple operating points. In *NeurIPS*, 2022. 2
- [56] Max van Spengler, Erwin Berkhout, and Pascal Mettes. Poincare resnet. In *ICCV*, 2023. 3
- [57] Girish Varma, Anbumani Subramanian, Anoop Namboodiri, Manmohan Chandraker, and CV Jawahar. Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *WACV*, 2019. 6
- [58] Celine Vens, Jan Struyf, Leander Schietgat, Saso Dzeroski, and Hendrik Blockeel. Decision trees for hierarchical multi-label classification. *Mach. Learn.*, 73(2):185–214, 2008. 2
- [59] Wenguan Wang, Zhijie Zhang, Siyuan Qi, Jianbing Shen, Yanwei Pang, and Ling Shao. Learning compositional neural information fusion for human parsing. In *ICCV*, 2019. 1
- [60] Wenguan Wang, Hailong Zhu, Jifeng Dai, Yanwei Pang, Jianbing Shen, and Ling Shao. Hierarchical human parsing with typed part-relation reasoning. In *CVPR*, 2020. 1
- [61] Xiaolin Wang and Bao-Liang Lu. Flatten hierarchies for large-scale hierarchical text categorization. In *ICDIM*, 2010. 2
- [62] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. 1
- [63] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas M. Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *CVPR*, 2022. 2
- [64] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2020. 6
- [65] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *ECCV*, 2020. 6
- [66] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steininger, and Gustavo Fernandez Dominguez. Wilddash-creating hazard-aware benchmarks. In *ECCV*, 2018. 6
- [67] Long Zhu, Yuanhao Chen, Alan L. Yuille, and William T. Freeman. Latent hierarchical structural learning for object detection. In *CVPR*, 2010. 2