Hyperbolic Image Segmentation
Supplementary materials

Mina GhadimiAtigh[1,*], Julian Schoep[1,*], Erman Acar[2], Nanne van Noord[1], Pascal Mettes[1]
[1]University of Amsterdam, [2]Leiden University, [2]Vrije Universiteit Amsterdam

This supplementary file contains the following supporting information about Hyperbolic Image Segmentation: (i) additional quantitative evaluation on hyperbolic uncertainty, (ii) additional results for zero-label generalization, and (iii) qualitative results for zero-label generalization, (iv) the full hierarchies of all datasets.

1. Boundary information for free

The main paper provides a quantitative correlation analysis between hyperbolic uncertainty and boundary distance for 2 embedding dimensions. Figure 1 shows the histogram of correlations over all images in Pascal VOC for 256 embedding dimensions. The histogram shows that the correlation also holds for high embedding dimensions.

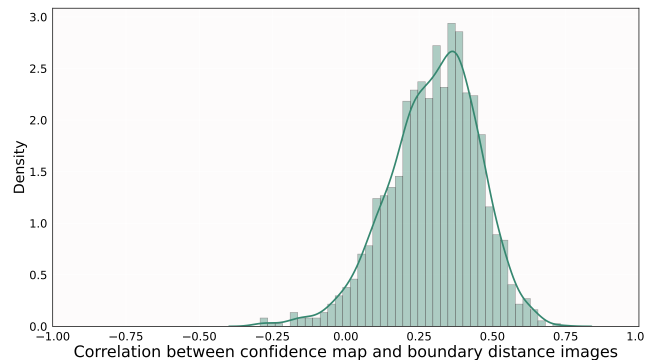

Figure 1. **Is hyperbolic uncertainty semantically meaningful?** We find that the per-pixel hyperbolic uncertainty strongly correlates with semantic boundaries for 256 embedding dimensions.

2. Further explanation on colors

While demonstrating the qualitative examples, there can be a discrepancy in segmentation colors compared to the pixel embeddings. The discrepancy is due to ambiguity in the visualization when dealing with pixels of high uncertainty. The final class logits per pixel are computed as hyperbolic distances to the gyroplanes. A pixel that looks closer to a certain gyroplane in the disk visualization can actually be closer to another gyroplane in the hyperbolic space. This holds only for gyroplanes that are close to the boundary of the Poincaré ball and for pixels with high uncertainty.

3. Zero label generalization

Table 1 provides the results for zero-label generalization on COCO-Stuff-10k for the sibling and cousin variants of the three metrics. Similar to the experiment in the main paper, the standard Euclidean architecture without hierarchical knowledge has a near-random performance, while hyperbolic embedding spaces with hierarchical knowledge performs best. The s-mIOU increases from 21.17 to 25.60 and c-mIOU increases from 21.83 to 25.72. As such, our approach performs even better for the hierarchical metrics compared to the standard metrics provided in the main paper: the standard mIOU increases 2.62, while s-mIOU and c-mIOU increase by 4.43 and 3.89.

		COCO-Stuff-10k		
Manifold	Hierarchical	S-Class Acc	S-Pixel Acc	S-mIOU
$\mathbb{R}$		2.01	0.55	0.27
$\mathbb{R}$	✓	7.89	59.96	21.17
$\mathbb{D}$	✓	**9.15**	**68.02**	**25.60**

		COCO-Stuff-10k		
Manifold	Hierarchical	C-Class Acc	C-Pixel Acc	C-mIOU
$\mathbb{R}$		3.71	0.87	0.37
$\mathbb{R}$	✓	22.91	60.89	21.83
$\mathbb{D}$	✓	**24.83**	**68.35**	**25.72**

Table 1. **Zero-label generalization** on COCO-Stuff-10k. Combining hierarchical knowledge with hyperbolic embeddings provides a more suitable foundation with respect to sibling and cousin variants of metrics.

4. Qualitative zero-label results

In Figure 4, we show qualitative examples of zero-label experiments. We show outputs of DeepLabV3+ and our approach given inputs from unseen classes *cow, giraffe, river*,

[*] Equal contribution

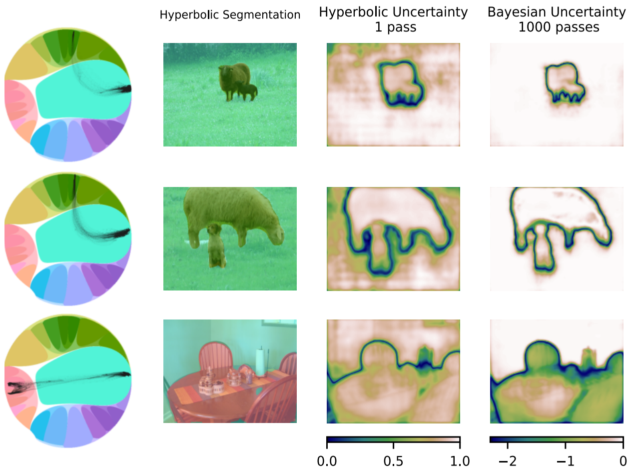

Figure 2. **Hyperbolic vs Euclidean uncertainty** for examples from Pascal VOC.

and *road* from the COCO-Stuff-10K dataset. The red points show the points classified as the target unseen class. To generate the output, we select the points for which the joint probability of animal, animal, water and ground (ancestor nodes from the dataset hierarchy) is higher than 50%.

5. Dataset hierarchies

In Figures 3, 5, and 6, we provide the full hierarchies for respectively Pascal VOC [2], COCO-Stuff-10K [1], and ADE20K [3]. We have created the hierarchies ourselves and will provide them in the final code repository.

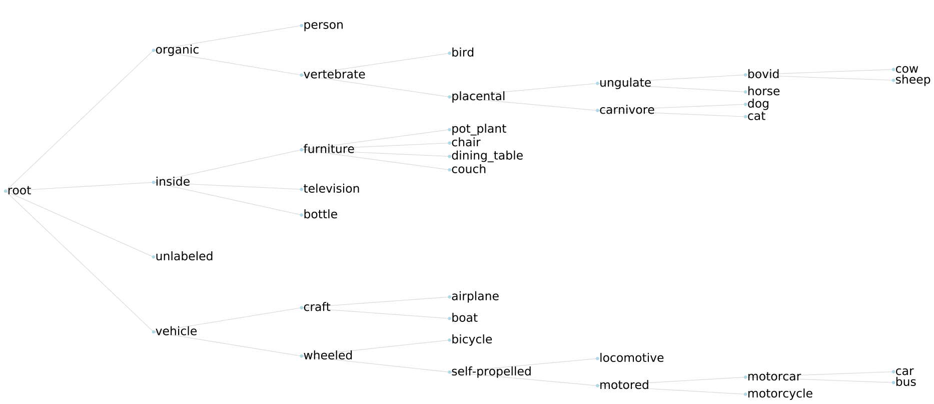

Figure 3. **Pascal VOC dataset full hierarchy.**

References

[1] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 2018. 2

[2] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. 2

[3] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. 2

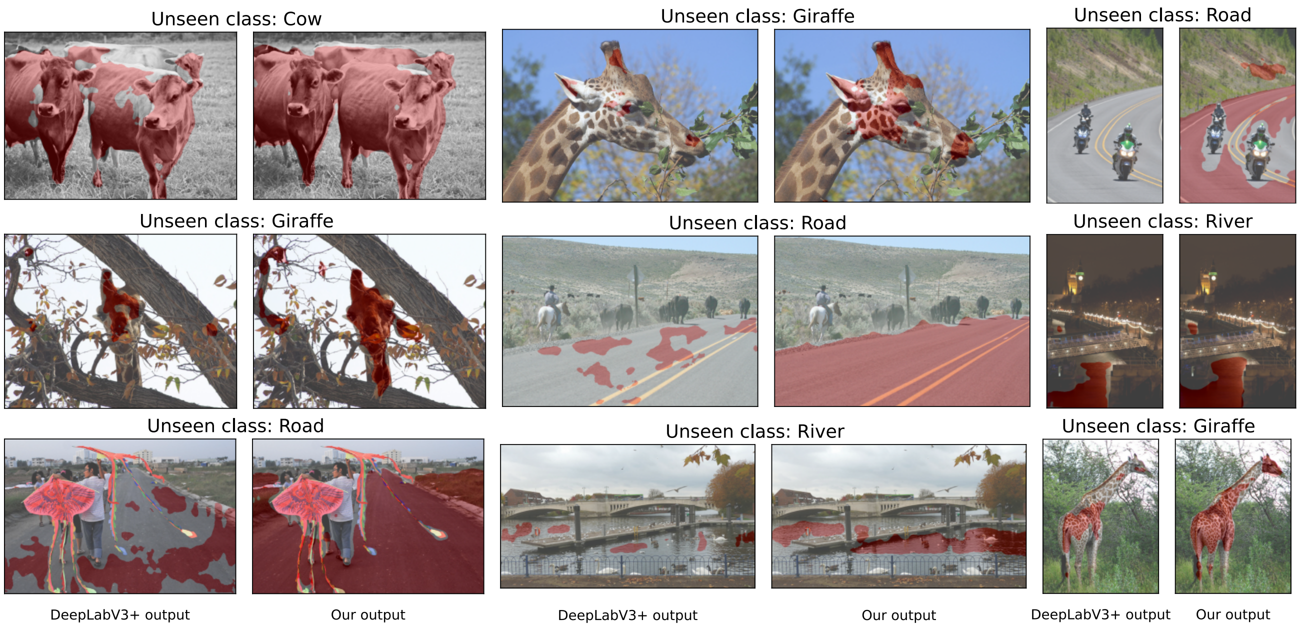

Figure 4. **Qualitative zero-label results** on COCO-Stuff-10K. The red points indicate the points classified as the target unseen class using both a standard DeepLabV3+ and the same model with hyperbolic embeddings. We find that our approach provides a more precise and complete picture of unseen classes.

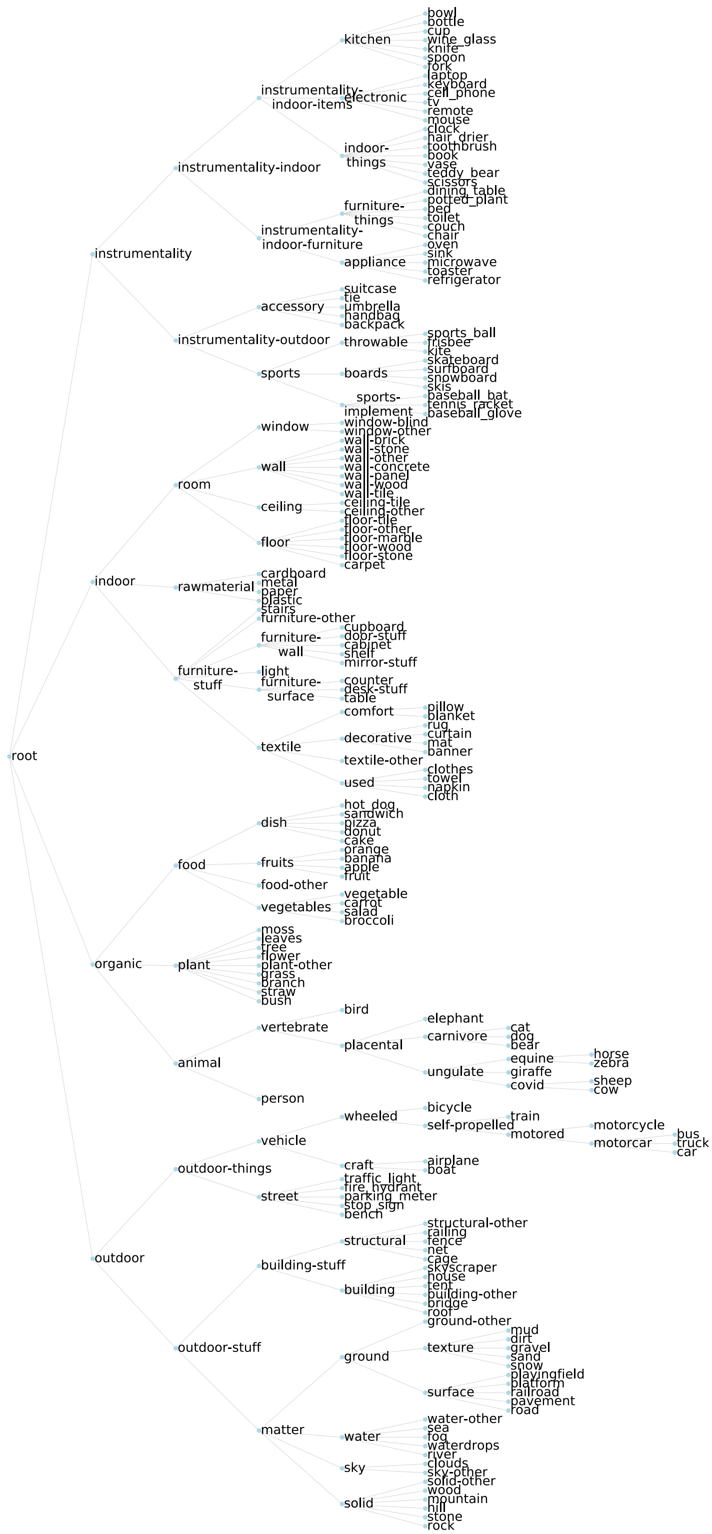

Figure 5. **COCO-Stuff-10K dataset full hierarchy.**

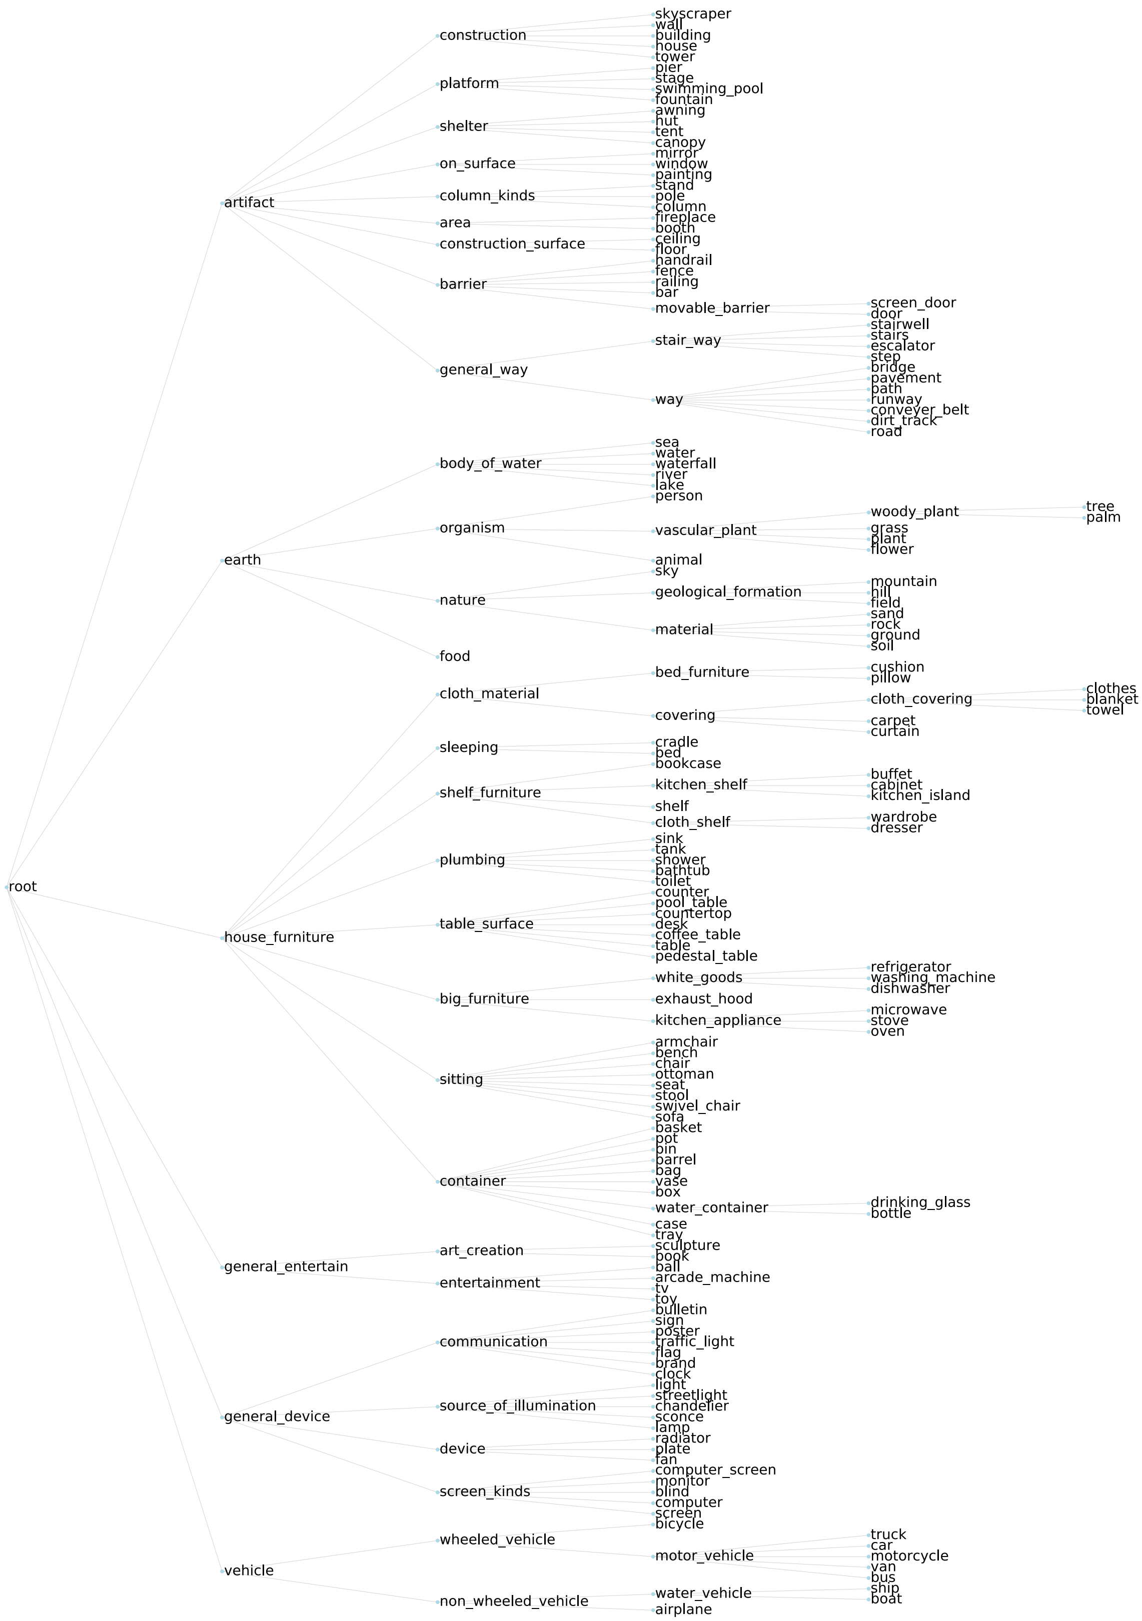

Figure 6. **ADE20K dataset full hierarchy.**