Physically Feasible Semantic Segmentation

Shamik Basu
University Of Bologna,Italy
shamik.basu@studio.unibo.it

Luc Van Gool
ETH Zurich
vangool@vision.ee.ethz.ch

Christos Sakaridis
ETH Zurich
csakarid@vision.ee.ethz.ch

Abstract

State-of-the-art semantic segmentation models are typically optimized in a data-driven fashion, minimizing solely per-pixel or per-segment classification objectives on their training data. This purely data-driven paradigm often leads to absurd segmentations, especially when the domain of input images is shifted from the one encountered during training. For instance, state-of-the-art models may assign the label "road" to a segment that is located above a segment that is respectively labeled as "sky", although our knowledge of the physical world dictates that such a configuration is not feasible for images captured by forward-facing upright cameras. Our method, Physically Feasible Semantic Segmentation (PhyFea), first extracts explicit constraints that govern spatial class relations from the semantic segmentation training set at hand in an offline, data-driven fashion, and then enforces a morphological yet differentiable loss that penalizes violations of these constraints during training to promote prediction feasibility. PhyFea is a plug-and-play method and yields consistent and significant performance improvements over diverse state-of-the-art networks on which we implement it across the ADE20K, Cityscapes, and ACDC datasets. Code and models will be made publicly available.

1. Introduction

Semantic segmentation is a fundamental task in computer vision and enables many downstream applications. The last decade has witnessed dramatic advances in this dense prediction task, driven by the introduction of end-to-end learned network architectures for solving it [2–4, 6, 14, 24, 27, 36, 44, 45, 47] and the development of ever stronger backbones coming from image classification [10, 15, 23, 32, 37, 46].

Indeed, designing a suitable backbone for dense semantic prediction that balances global context aggregation with preservation of fine local details has been a primary research direction in semantic segmentation. However, merely optimizing the backbone architecture neglects the *structured prediction* character of the problem. In particular, the simple optimization of cross-entropy loss terms even by state-of-the-art, foundational network architectures [6, 36] essentially poses semantic segmentation as a set of *independent* pixel-level or—at best—segment-level prediction problems and only loosely promotes *regularity* of the complete image-level outputs via the shared computation in the network across neighboring pixels. Very few deep-learning-based works [17] have focused on designing objectives that originate from the interaction between different pixels in the output predictions, but even these objectives are formulated at a rather local level of pixel neighborhoods and not across large regions of the image. To the best of our knowledge, no previous method in semantic segmentation with deep learning has considered high-level pairwise objectives at the coarse level of semantic segments.

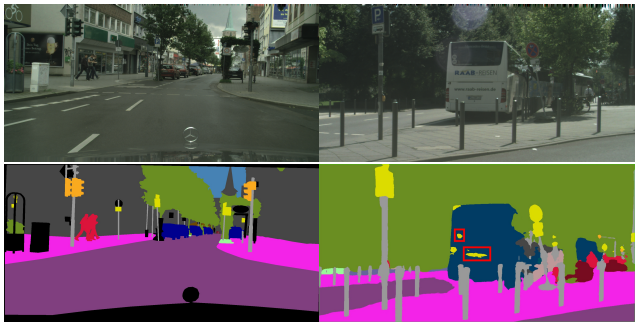

Figure 1. Left: A *traffic light* segment includes a *building* segment in a semantic annotation from the ground truth of the dataset, which is a feasible inclusion. Right: A *bus* segment includes a *traffic sign* segment (marked by red bounding boxes) in a semantic prediction, which is an **infeasible inclusion**. Our method addresses such physical infeasibilities in semantic segmentation. Best viewed on a screen and zoomed in.

Aiming at promoting conformity of semantic segmentation models to high-level constraints, we present a new method named Physically Feasible Semantic Segmentation, or PhyFea. Our method has been empirically motivated by the observation that various state-of-the-art semantic segmentation networks [6, 41, 45] make errors in their pre-

dictions which are not simple misclassifications but fundamentally violate *physical constraints* which are in force in images of real-world 3D scenes, as illustrated in Fig. 1, and which a human would easily avoid. In particular, based on their prior knowledge about the world, humans know that a *traffic sign* segment cannot be entirely included in/surrounded by a *bus* segment. We term such segment inclusions as *infeasible inclusions* and consider them as an exemplary type of physical constraint that can be exploited to provide a segmentation model with higher-level objectives based on the feasible spatial configurations of semantic segments in the output.

While it could be possible to exhaustively consider all pairs of semantic classes in the segmentation dataset at hand and conceptually determine which of these pairs correspond to such infeasible inclusions, we argue that this knowledge can be directly extracted from the data. More specifically, our method first processes the training set offline before training in a single pass and automatically identifies all class pairs for which a segment of the one class includes a segment of the other at least once in the data. These are *valid inclusion* pairs (cf. Fig. 1-left) and PhyFea should and does *not penalize* similar inclusions in the outputs of the segmentation network. By contrast, pairs of co-occurring classes for which a complete inclusion of one class by the other is never found in the data are deemed as infeasible inclusion pairs (given the large size of the training set) and they contribute to a feasibility objective that we define for optimization during training.

More specifically, the major contribution of PhyFea is a novel inclusion loss which is used as an additional objective besides standard cross-entropy for optimizing a generic segmentation network. Our inclusion loss is computed based on the softmax outputs of the employed network, similarly to cross-entropy, which allows its general application to the vast majority of deep-learning-based semantic segmentation models. The key idea of computing this loss is to identify "soft" infeasible inclusions, which occur as regions in the softmax score maps of infeasible inclusion class pairs in which the ranking of the two softmax values is inverted compared to their surroundings. We recognize that this simply amounts to a grayscale morphological area opening [13] and derive a differentiable, parameter-free and thus efficient implementation of it, which is normally included in the forward and backward pass of standard backpropagation.

We have achieved a consistent performance improvement with PhyFea across three central semantic segmentation benchmarks, i.e. ADE20K [49], Cityscapes [8], and ACDC [31], and across a variety of state-of-the-art semantic segmentation networks on which we have implemented our method, i.e. SegFormer [41], OCRNet [45], and Mask2Former [6]. Importantly, we have also observed a reduction of infeasible inclusions in our outputs.

2. Related Work

Physical priors using local correlations. Physical domain knowledge in the medical domain is implemented especially in MRA, PET, and MRI scan images concerning optimizing energy function or implementation in Markov Random Field [22] (a graphical model approach). Other approaches to implementing physical priors in the semantic segmentation model are AAF [17] (Adaptive Affinity Field) and a few similar approaches [1, 9, 12, 14, 40, 43, 50] which are very closely related to CRF (Conditional Random Field) [20]. In AAF, the predicted feature map of PSPNet [47] is compared with the input feature map using multiple-sized kernel windows using KL divergence as a loss function.

Physical priors via full supervision. Zhang et al. [39] stated that physical priors can be implemented in a neural framework in four different ways. The first strategy is to integrate constraints into data by generating synthetic training datasets. The second strategy is to design non-trainable custom layers of physical operators and pre-conditioners in the DNN architecture to modify or shape feature maps calculated within the network to make them consistent with the prior knowledge. Implementing physical priors in semantic segmentation networks is also done by modern techniques using Self-attention [34, 48]. CCNet [16] is one such example and a few other similar approaches are [11, 25, 38, 48, 50]. In CCNet, ResNet-101 [46] is used as the backbone to extract the feature map and dilated convolution is applied to increase the spatial dimension of the feature map. Then using 1-D convolution to obtain the necessary parameters to apply self-attention map the semantic structures. In this model, the key factor is applying self-attention to capture the semantic relationships between the classes.

Physical priors in the loss function. Physical priors can be enforced in the loss function as regularization loss as in the Potts model [33] in a discrete domain that penalizes appearance discontinuity between neighboring pixels. The regularization term formulated in the discrete setting is biased by the discrete grid and favors curves to orient along with the grid, e.g. in horizontal and vertical or diagonal directions in a 4-connected lattice of pixels. Previous attempts have been made to implement physical priors by incorporating shape information in the architecture [18, 21, 28, 42]. Shape information is a powerful semantic descriptor for specifying targeted objects in an image. Shape priors can be modeled in three ways: geometrical [35], statistical, and physical.

There are special types of neural architectures where physical priors can be implemented in loss functions. These types of networks are also called Physics Informed Neural Networks (PINNs) [29]. PINNs are differential equation

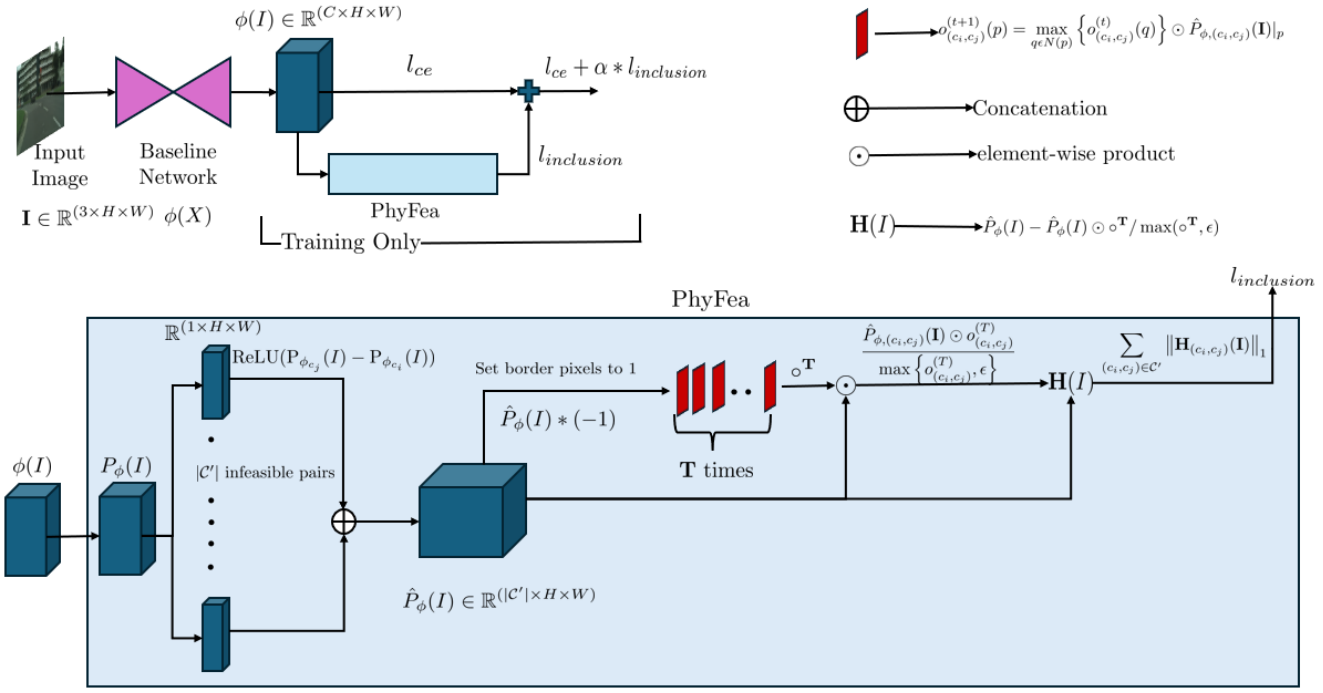

Figure 2. **Overview of PhyFea.** Top left: the complete network architecture, where the standard cross-entropy loss l_{ce} from the baseline network ϕ is added to the inclusion loss $l_{\text{inclusion}}$ computed by PhyFea. Bottom: The pipeline of PhyFea. For the softmax outputs $P_\phi(I)$ of the network and for each pair of classes $(c_i, c_j) \in \mathcal{C}'$ that signifies an infeasible inclusion where c_i cannot include c_j, we take the difference of the respective softmax scores, $\mathrm{P}_{\phi_{c_j}}(I) - \mathrm{P}_{\phi_{c_i}}(I)$, and rectify it with a ReLU. After concatenating all such $\mathcal{C}'$ rectified difference maps channel-wise into $\hat{P}_\phi(I)$, we negate the latter and set its border pixels to 1. Then an iterative operation for area opening is performed $\mathbf{T}$ times on $\hat{P}_\phi(I)$. In each iteration, we perform max-pooling with a 3x3 kernel and a stride of 1, and the result is multiplied with $\hat{P}_\phi(I)$ element-wise. The final area-opened tensor channels differ from their counterparts in $\hat{P}_\phi(I)$ only across regions of incorrect inclusions. This element-wise difference is stored in $\mathbf{H}(I)$ and the L_1 norm of the latter, capturing both the spatial extent and the intensity of infeasible inclusions, constitutes the inclusion loss $l_{\text{inclusion}}$ employed in PhyFea.

solvers where a differential equation encodes the physical constraint. The neural network approximates the solution of the equation using boundary constraints. There are studies done in the geophysics domain to implement geophysical priors and laws in deep neural networks as regularization terms in loss functions for training the DNNs.

Nosrati et al. [26] surveyed incorporating prior knowledge in medical image segmentation. This work reviews various ways to device prior knowledge information to improve image segmentation. The said physical priors are implemented with global or local optimization; as implementation of physical priors in continuous or discrete domains. Rother et al. [30] suggest ways to incorporate appearance priors in image segmentation using Gaussian mixture models [19]. In the literature, there are two ways to model the appearance: 1) learning the appearance in an adaptive manner during the segmentation procedure, and 2) knowing the appearance model before performing segmentation (e.g. by observing the appearance distribution of the training data).

Comparison of our method with the related works. All the above methods to enforce physical priors into a deep neural architecture do not take into account the semantic relation among the classes in the image. For example, in outdoor scenes, the sky should be above the road, a pole cannot be included by a bus or a sidewalk cannot be surrounded by the wall of a building. This shows that during training the model; is not being trained keeping in mind all these semantic relations among the classes. PINN although takes into account shape priors in the medical domain, it cannot extend semantic segmentation of general scenes, because we cannot enforce a fixed structure of any class in outdoor scenes unlike fixed shape anatomical structures, which are required to formulate the differential equation approximated in PINN.

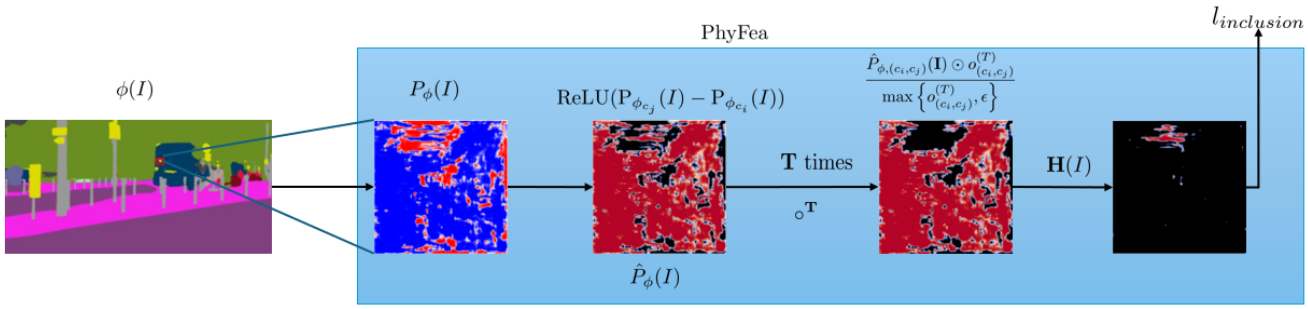

Figure 3. **Visualization of feature maps within PhyFea for a class pair corresponding to an infeasible inclusion.** In this example semantic prediction of SegFormer [41] on Cityscapes [8], *bus* infeasibly includes *traffic sign*. In the blue frame, we first show the difference map between the softmax scores of the two classes using the *coolwarm* colormap, where blue tones indicate larger scores for *traffic sign* and red for *bus*, respectively. We rectify the difference between the two softmax maps and show only the respective channel of the 3D tensor $\hat{P}_\phi(I)$, where black indicates zeros and red tones indicate positive values. The area opening operation is performed on $\hat{P}_\phi(I)$ and the regions in which $\hat{P}_\phi(I)$ is positive (red) but which are not connected to the border, i.e. *traffic sign* segments infeasibly included in *bus* segments, are opened. Exactly these segments are isolated in the final tensor $\mathbf{H}(I)$ and are used to compute our inclusion loss $l_{\text{inclusion}}$.

3. PhyFea

3.1. Method Overview

Fig. 2 shows the architectural overview of PhyFea. A 2D semantic segmentation network $\phi(X)$ is fed with an image $\mathbf{I} \in \mathbb{R}^{(3\times H\times W)}$ and produces the initial output $\phi(\mathbf{I}) \in \mathbb{R}^{(C\times H\times W)}$, where C is the number of classes. PhyFea takes $\phi(\mathbf{I})$ as input and computes a novel inclusion loss, $l_{\text{inclusion}}$. We identify infeasible inclusions for class pairs of interest via a morphological area opening operation, which is fully differentiable and is normally included in the forward and backward pass during training. Our loss $l_{\text{inclusion}}$ is then added to the standard cross-entropy loss l_{ce} for $\phi(X)$ to obtain the total loss as

$$l_{\text{total}} = l_{\text{ce}} + \alpha l_{\text{inclusion}}, 0 < \alpha < 1. \quad (1)$$

l_{total} is normally backpropagated to optimize the weights of ϕ. Thus, PhyFea is a differentiable morphological module that incorporates data-driven physical priors in the form of infeasible inclusions to improve the optimization and conformity of any segmentation network. Notably, PhyFea is fully parameter-free. In (1), α is a hyperparameter that is used to balance the inclusion loss and the standard cross-entropy loss.

3.2. Computing Infeasible Inclusion Pairs

We apply Algorithm 1 to the training set of the dataset at hand for optimizing and evaluating the segmentation model, to compute the set $\mathcal{U}$ of class pairs corresponding to feasible, or valid, inclusions. Intuitively, this set contains class pairs for which a segment of the first class includes a segment of the second one at least in one case in the training set. The set of class pairs corresponding to *infeasible inclusions*, $\mathcal{C}'$, is then simply the set-theoretic difference $\mathcal{C}' = \mathcal{P} \setminus \mathcal{U}$, where $\mathcal{P}$ is the set of all possible non-identical class pairs.

In Algorithm 1, we first compute for each annotation I in the training set two binary masks of the same dimension as I, in which pixels labeled as class p and class q are set as foreground, respectively. We then compute the connected components of class q, i.e. the class whose possibility to be included is checked, in `mask_2`. For each such connected component, we fill any holes inside it to remove any interior contour in the next step as it is irrelevant. After that, we compute the outer contour of the connected component at hand based on an 8-neighborhood and check whether it fully falls into pixels belonging to class p. If the latter is true and the length of the contour is non-negligible, we declare the case at hand as a valid inclusion and add the pair (p, q) to the set $\mathcal{U}$.

3.3. Inclusion Loss

To compute our novel inclusion loss, we operate on the softmax maps $P_\phi(\mathbf{I})$ output by the segmentation network and index them channel-wise for each of the $|\mathcal{C}'|$ infeasible inclusion pairs. Our core intuition is to identify in the pairs of softmax score maps that correspond to each infeasible inclusion type in $\mathcal{C}'$ those regions across which class c_j—the includee—has a larger score than class c_i—the includer—and around which the opposite is true. These regions indicate "soft" infeasible inclusions in the softmax predictions, which should be penalized proportionally to the area they occupy and to the variable "intensity" of the inclusion, i.e. the magnitude of the difference between the softmax scores of classes c_i and c_j at each pixel of these regions. The computation of such a loss, which we term the inclusion loss, can be directly achieved in a differentiable formula-

Algorithm 1 Computing the set $\mathcal{U}$ of feasible inclusion class pairs from the training set ground-truth annotations $\mathcal{S}$. The set $\mathcal{P}$ contains all pairs of non-identical class labels for the dataset taxonomy. We aim to check whether the first element of the pair, p, can include the second one, q. We are taking 5 pixels as the minimum threshold for a `contour` of q completely overlapping with any segment of p. Function **C** returns (i) the number of connected segments of class q in the input annotation, and (ii) a discrete image, with each pixel labeled with its connected component ID or 0 for the background. Function **F** fills the holes of the mask of the connected component at hand. Then, function **Cont** returns a binary image in which only pixels belonging to the outer contour of the respective connected component are set. All `variables` except `component` are discrete images $\in \mathbb{N}^{(1\times H\times W)}$. Function **NNZ** returns the number of non-zero pixels of its input.

Require: $\mathcal{P} = \{(p,q) : p \in \{1,\dots,C\}, q \in \{1,\dots,C\} \setminus \{p\}\}$
Require: $\mathcal{S} = \{I : I \in \mathbb{N}^{1,H,W}\}$
Ensure: Set $\mathcal{U}$ of feasible inclusion class pairs

$\mathcal{U} \leftarrow \emptyset$
for $(p,q) \leftarrow \mathcal{P}$ **do**
  for $I \leftarrow \mathcal{S}$ **do**
    `mask_1` $\leftarrow \mathbb{1}[I = p] \in \mathbb{R}^{1\times H\times W}$
    `mask_2` $\leftarrow \mathbb{1}[I = q] \in \mathbb{R}^{1\times H\times W}$
    `component` , `noOFLabels` $\leftarrow$ **C**(`mask_2`)
    for $l \leftarrow 1$ to `noOFLabels` **do**
      `comp_mask` $\leftarrow (\mathbb{1}[$`component` $= l])$
      `comp_filled` $\leftarrow$ **F**(`comp_mask`)
      `contour` $\leftarrow$ **Cont**(`comp_filled`)
      `result` $\leftarrow$ `contour` $\wedge$ `mask_1`
      if **NNZ**(`result`) $=$ **NNZ**(`contour`) $\wedge$ **NNZ**(`contour`) ≥ 5 **then**
        $\mathcal{U} \leftarrow \mathcal{U} \cup \{(p,q)\}$
      end if
    end for
  end for
end for

tion, amenable to standard backpropagation, via a grayscale morphological area opening of the aforementioned softmax difference maps. We implement this opening through an iterative algorithm, detailed in the following.

For each "infeasible pair" $(c_i, c_j) \in \mathcal{C}'$, we denote with $\mathrm{P}_{\phi_{c_i,c_j}}(\mathbf{I}) \in \mathbb{R}^{(2\times H\times W)}$ the concatenation of $\mathrm{P}_{\phi_{c_i}}(\mathbf{I})$ and $\mathrm{P}_{\phi_{c_j}}(\mathbf{I})$. The difference between the two channels of the latter tensor is rectified with a ReLU as $\mathrm{ReLU}(\mathrm{P}_{\phi_{c_j}}(\mathbf{I}) - \mathrm{P}_{\phi_{c_i}}(\mathbf{I})) \in \mathbb{R}^{(1\times H\times W)}$ (cf. Fig. 2). After that, all $|\mathcal{C}'|$ pairwise differences are concatenated to obtain the tensor $\hat{P}_\phi(\mathbf{I}) \in \mathbb{R}^{(|\mathcal{C}'|\times H\times W)}$.

For each $(c_i, c_j) \in \mathcal{C}'$, we initialize a running map $o^{(0)}_{(c_i,c_j)} \in \mathbb{R}^{(1\times H\times W)}$ so that its border pixels are set to 1 and the rest are equal to 0. The intuition here is that infeasibly included regions are always disconnected from the border of the image. Thus, positive values of the map $o^{(t)}_{(c_i,c_j)}$ that are gradually propagated from the border to the interior of the image via iterative 8-neighborhood-based dilation (a.k.a. max-pooling) will *never* reach the aforementioned regions if at each iteration they are multiplied with $\mathrm{ReLU}(\mathrm{P}_{\phi_{c_j}}(\mathbf{I}) - \mathrm{P}_{\phi_{c_i}}(\mathbf{I}))$, as the latter is 0 everywhere around infeasibly included regions. The above, intuitively described iterative operation can be formally written as

$$o^{(t+1)}_{(c_i,c_j)}(p) = \max_{q\in N(p)} \left\{ o^{(t)}_{(c_i,c_j)}(q) \right\} \odot \hat{P}_{\phi,(c_i,c_j)}(\mathbf{I})|_p, \quad (2)$$

where $N(p)$ denotes the structuring element of the dilation operator or equivalently the max-pooling kernel, which we set to an 8-connected neighborhood, i.e. a 3x3 kernel. We iterate (2) T times, where $T = \max\{2, \min\{H, W\}/2\}$ and H and W are the height and width of the input image. Namely, we iterate for the minimal number of times that is required for initial positive values at image borders in $o^{(0)}_{(c_i,c_j)}$ to be able to reach any pixel in the interior. The final map $o^{(T)}_{(c_i,c_j)}$ is by construction 0 at all pixels that belong to infeasibly included regions. This property implies that

$$\frac{\hat{P}_{\phi,(c_i,c_j)}(\mathbf{I}) \odot o^{(T)}_{(c_i,c_j)}}{\max\left\{ o^{(T)}_{(c_i,c_j)}, \epsilon \right\}} \quad (3)$$

constitutes the grayscale area opening we are after for the infeasible inclusion class pair (c_i, c_j) (cf. Fig. 3).

The relevant quantity for computing our inclusion loss is the difference between the original rectification output and its area opening in (3). This difference is given by

$$\mathbf{H}_{(c_i,c_j)}(\mathbf{I}) = \hat{P}_{\phi,(c_i,c_j)}(\mathbf{I}) - \frac{\hat{P}_{\phi,(c_i,c_j)}(\mathbf{I}) \odot o^{(T)}_{(c_i,c_j)}}{\max\left\{ o^{(T)}_{(c_i,c_j)}, \epsilon \right\}}. \quad (4)$$

We perform the above computations for all $(c_i, c_j) \in \mathcal{C}'$. Finally, our inclusion loss is computed as the sum of the L_1 norms of tensors $\mathbf{H}$ across all $(c_i, c_j) \in \mathcal{C}'$ via

$$l_{\text{inclusion}} = \sum_{(c_i,c_j)\in\mathcal{C}'} \left\| \mathbf{H}_{(c_i,c_j)}(\mathbf{I}) \right\|_1. \quad (5)$$

4. Experiments

4.1. Datasets

Cityscapes. Cityscapes [8] is a driving dataset, containing various road images of European cities. The dataset has 19 classes and includes 2975 training, 500 validation, and

1525 test images with fine annotations. Moreover, 20000 additional training images are coarsely annotated.

ADE20K. We have used the dataset ADE20K [49]. It was earlier used in the ImageNet scene parsing challenge in 2016. There are 150 classes and diverse scenes with 1038 image-level labels. The dataset is divided into 20K/2K/3K images for training, validation, and test set.

ACDC. The ACDC dataset [31] consists of driving scenes under adverse conditions such as fog, night, rain, and snow. It has 19 classes similar to Cityscapes. The dataset is divided into 1600 training, 406 validation, and 2000 test images.

4.2. Implementation Details

Input to PhyFea $\phi(I)$ is first interpolated to shape $\mathbb{R}^{(C\times 256\times 256)}$. The kernel size of the max-pooling used is 3×3 with a stride of 1 and a padding of 1 to keep the dimensions the same and to check for any connected components not sharing the border pixels. For ADE20K, the number of infeasible inclusion pairs $|\mathcal{C}'|$ is very large, so random sampling is applied offline in each iteration to reduce the memory footprint. PhyFea computes $l_{inclusion}$ on average for each iteration in 0.14 seconds. We report semantic segmentation performance using mean Intersection over Union (mIoU). The value of α (ref Eq. 1) is decided using a random search from the set [1e-5, 1e-10]. For all reported datasets, the value of α is taken as 1e-9. We have used 4 A100 GPUs each with memory 80GB.

Training setup. For all the baselines, we have used a batch size of 4 for the Cityscapes [8] dataset and ACDC [31] datasets and 8 for the ADE20K [49] dataset. Optimizers, number of iterations, and learning rate parameters are kept similar to what is mentioned in the respective papers of the baselines for the respective datasets. For the ACDC [31] dataset, we applied the same setup as the corresponding baselines applied on the Cityscapes [8] dataset except for the crop size of the input.

We used mmsegmentation [7] to implement PhyFea on the baseline network SegFormer-B4 [41]. Random cropping of 1080×1080 for the ACDC [31] dataset is used.

For OCRNet without SegFix [45], their official implementation is taken. For ACDC 769×769 cropping is used.

For Mask2former [6], their official implementation is taken. For the ACDC dataset, we have taken a crop size of 540×960.

4.3. Comparison to The State of The Art

Tab. 1 compares PhyFea with a wide range of state-of-the-art methods for semantic segmentation on Cityscapes [8], ADE20K [49], and ACDC [31]. Overall, PhyFea achieves significant improvements in mIoU across all three network architectures on which we have implemented it and across all three datasets. On Cityscapes, where all examined state-of-the-art architectures already achieve very high performance, we observe a consistent 0.5–0.6% improvement in mIoU with PhyFea on all architectures, which indicates that our method can directly serve as a plug-in module for generic segmentation networks. PhyFea achieves 1.2% improvement on the baseline OCRNet [45] for the ADE20K dataset and 0.7% on the baseline SegFormer-B4. PhyFea achieves 1.1% improvement on the baseline OCRNet for the ACDC dataset and 2.17% on the baseline SegFormer-B4. We can observe that Mask2Former has the highest mIoU score across all the 3 datasets. PhyFea can gain improvement on the respective baselines on all the 3 datasets by enforcing the Infeasible Inclusion prior but PhyFea is baseline-independent. It means it can improve the image segmentation of a given baseline network on which it is ap-

Table 1. **Comparison with state-of-the-art semantic segmentation methods on ADE20K, Cityscapes and ACDC.** '*': multi-scale testing. Results are reported in mean IoU (%) on the Cityscapes test set, the ADE20K validation set, and the ACDC test set.

Method \ Dataset	Backbone	ADE20K	Cityscapes	ACDC
PSPNet [47]		44.4	78.5	-
DeepLabv2 [3]		-	71.4	55.3
DeepLabv3+ [5]		44.1	80.9	70.0
HRNet [36]		43.1	80.9	75.0
OCRNet [45]	HRNetV2-W48	45.6*	82.4*	66.5
PhyFea w/ OCRNet (ours)	HRNetV2-W48	**46.8***	**82.9***	**67.6**
SegFormer [41]	MiT-B4	50.3	82.2	67.1
PhyFea w/ SegFormer (ours)	MiT-B4	**51.0**	**82.7**	**69.2**
Mask2Former [6]	Swin-L	56.1	83.5*	77.2
PhyFea w/ Mask2Former (ours)	Swin-L	**56.7**	**84.0***	**77.8**

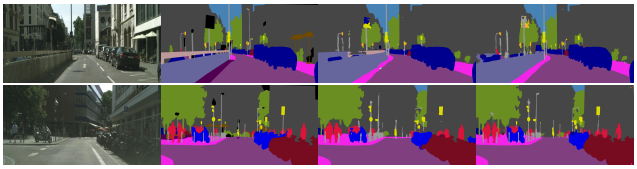

Figure 4. **Qualitative comparison on Cityscapes.** From left to right: input image, ground-truth semantic labels, and predictions of Mask2Former [6] and PhyFea.

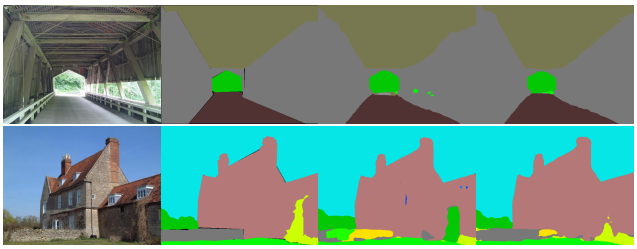

Figure 5. **Qualitative comparison on ADE20K.** From left to right: input image, ground-truth semantic labels, and predictions of OCRNet [45] and PhyFea.

Table 2. **Comparison of class-level IoU of OCRNet, SegFormer, and Mask2Former to their PhyFea-upgraded versions on the Cityscapes validation set.** Training and evaluation are performed using the complete training and validation sets, respectively.

Method	road	sidew.	build.	wall	fence	pole	light	sign	veget.	terrain	sky	person	rider	car	truck	bus	train	motorc.	bicycle	mIoU
OCRNet [45]	98.2	86.8	93.1	**65.1**	63.6	68.4	74.3	**81.2**	92.8	63.9	**95.5**	84.3	66.4	95.2	77.6	88.7	83.1	69.2	79.0	80.3
Ours	**98.3**	**86.9**	**93.4**	64.4	**64.2**	**68.5**	**74.8**	81.0	**92.9**	**64.9**	95.4	**84.6**	**67.9**	**95.4**	**86.1**	**90.8**	**83.2**	**72.3**	**80.0**	**81.3**
SegFormer [41]	98.4	87.8	93.7	68.4	65.4	**69.6**	**75.6**	**81.6**	93.1	70.6	**95.4**	84.8	68.3	95.6	81.8	90.7	83.9	73.6	**80.0**	82.0
Ours	**99.2**	**91.3**	**93.8**	**72.1**	**72.3**	56.2	73.5	79.3	**93.5**	**74.0**	93.8	**85.0**	**74.8**	**95.7**	**90.8**	**95.2**	**84.7**	**83.3**	79.1	**83.6**
Mask2Former [6]	98.6	88.7	94.0	66.7	69.6	71.6	76.2	84.3	93.4	68.4	**95.7**	86.1	70.7	96.2	89.7	92.7	**84.5**	74.6	81.5	83.3
Ours	**99.0**	**91.0**	**94.8**	**81.4**	**78.5**	**73.7**	**77.3**	**85.9**	**93.8**	**74.4**	95.6	**87.7**	**77.6**	**96.5**	**92.4**	**94.8**	63.8	**78.6**	**83.3**	**85.3**

Table 3. **Comparison of class-level IoU of OCRNet, SegFormer, and Mask2Former to their PhyFea-upgraded versions on the ACDC test set.** Training and evaluation are performed using the complete training and test sets, respectively.

Method	road	sidew.	build.	wall	fence	pole	light	sign	veget.	terrain	sky	person	rider	car	truck	bus	train	motorc.	bicycle	mIoU
OCRNet [45]	**93.4**	**77.2**	87.1	**50.0**	**45.2**	48.6	60.4	61.3	85.4	**59.4**	**95.5**	58.5	40.5	**87.4**	**63.6**	**76.3**	80.6	41.6	53.2	66.5
Ours	92.5	71.6	**87.2**	45.3	39.9	**54.2**	**70.2**	**68.2**	**86.4**	50.8	94.8	**65.2**	**45.5**	87.2	60.9	69.9	**84.5**	**50.6**	**59.9**	**67.6**
SegFormer [41]	94.6	78.3	88.4	51.6	47.4	48.0	62.1	61.2	87.0	**66.8**	95.6	59.8	39.0	87.0	63.1	73.8	79.1	38.4	52.7	67.1
Ours	**94.8**	**79.5**	**88.5**	**54.2**	**50.4**	**48.7**	**63.7**	**64.2**	**87.9**	66.2	**95.9**	**61.2**	**42.0**	**88.2**	**70.8**	**77.9**	**84.1**	**42.3**	**54.6**	**69.2**
Mask2Former [6]	**96.5**	**84.7**	**93.2**	64.7	59.5	**72.0**	**80.6**	78.9	**91.1**	**72.1**	96.7	**77.7**	44.4	**91.6**	75.9	71.0	**92.4**	58.5	65.9	77.2
Ours	96.4	84.3	92.9	**65.7**	**61.2**	66.7	78.7	**79.5**	91.0	71.9	**97.0**	77.5	**56.6**	91.3	**76.9**	**73.3**	91.6	**58.9**	**67.0**	**77.8**

plied by enforcing the physical prior.

Table 2 is the class-wise comparison of IoU on Cityscapes val set for PhyFea and the corresponding baselines. We can observe that PhyFea achieved an overall improvement in IoU scores across the classes. PhyFea with OCRNet [45] shows improvement in 16 classes from the baseline, especially in the small classes like fence, rider, motorcycle, and bicycle which are very difficult concerning the semantic segmentation task. SegFormer [41] with PhyFea shows improvement in 14 classes from the baseline and despite performing poorly in the pole class, PhyFea has managed to acquire good improvement in other small and big classes. Mask2Former [6] shows improvement in 17 classes from the baseline. Mask2Former with PhyFea has the best improvement out of the other two baselines.

Table 3 is the class-wise comparison of IoU on the ACDC test set for PhyFea and the corresponding baselines. Here also we can observe that PhyFea achieved an overall improvement in IoU scores across the classes, especially over the SegFormer baseline. PhyFea achieved significant gains across all the baselines for the small classes like motorcycle, bicycle, and rider.

In Fig. 6 we can observe the training losses of baseline Mask2Former and when re-trained with PhyFea on the Cityscapes and ACDC datasets. At first, in both of the graphs we can see the loss with PhyFea is slightly higher than the baseline loss because, at first the $l_{inclusion}$ loss is very high up to 20k iterations roughly. The training loss tends to decrease after 20k iterations with PhyFea implementation slightly. Mask2Former implementation with PhyFea does not lose the ability to generalization on the test set. So we can confirm that Mask2Former with PhyFea is not over-fitting on the train set.

Cityscapes results. In Fig. 4, we note in the prediction of the basic SegFormer a segment of truck included by building, which is an infeasible inclusion. PhyFea, on the contrary, successfully opens (in morphological terminology) this truck segment. Moreover, in the SegFormer prediction, a segment of a motorcycle is completely included by a bicycle on the right of the image. PhyFea managed to completely open the motorcycle segment.

ADE20K results. In Fig. 5, in the baseline prediction in the top row, a segment of vegetation is incorrectly included by another class, while PhyFea successfully opens the former segment. We can also observe that the overall condition of the prediction image after PhyFea re-training is not degraded from the baseline prediction. In the bottom row, a segment of the sea class is included by building class which is an infeasible inclusion. PhyFea managed to successfully open the included segment.

ACDC results. In Fig. 7 we can see in baseline prediction that some segments of the sidewalk class are included by the wall class. And PhyFea has successfully opened those

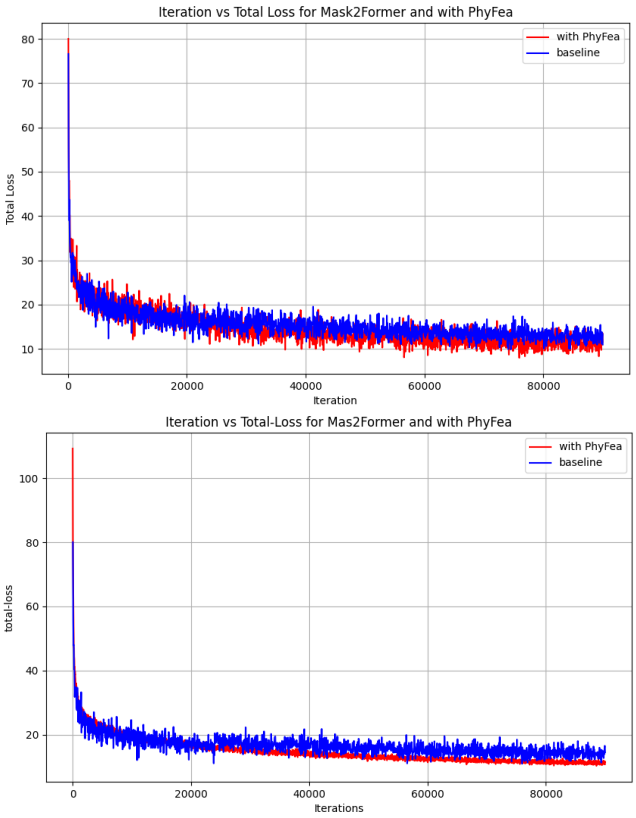

Figure 6. Training loss graph for Mask2Former and when re-trained with PhyFea on Cityscapes(**top**) and ACDC(**bottom**) datasets.

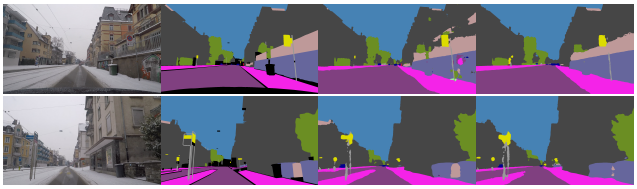

Figure 7. **Qualitative comparison on ACDC.** From left to right: input image, ground-truth semantic labels, and predictions of baseline SegFormer network and PhyFea (ours).

Table 4. **Ablation Study performed on PhyFea with Mask2Former-SwinB on Cityscapes val set.** Number of Iterations used for max-pooling Kernels 5×5 and 7×7 is 60k.

Kernel-size	Stride	Padding	T	mIoU
3×3	1	1	128	85.3
5×5	1	2	60	81.8
7×7	1	3	3	79.8

segments. In the bottom row, we can observe that a wall segment is included by the building class which falls under the infeasible inclusion. PhyFea successfully opened the included segment in the final re-trained image.

Table 5. **Feasible and infeasible inclusion pair statistics for different datasets.** The statistics are computed on the training set of each dataset. Infeasible pair counts exclude the non-cooccurring pairs.

Dataset \ Type of class pair	Feasible	Infeasible
Cityscapes (19 classes)	162	180
ADE20K (150 classes)	11070	11280
ACDC (19 classes)	118	224

4.4. Ablation Study

Tab 4 is the ablation study of PhyFea on the Cityscapes dataset using the Mask2Former as the baseline. We have shown the mIOU score on the val set. We can observe that by increasing the kernel size with a decrease in the number of iterations for max-pooling the mIOU score tends to decrease. We have kept the stride as 1 to compute for every possible position. We adjusted the padding appropriately to keep the dimensions the same for each iteration. The motivation behind the decrease in the mIOU score is that a bigger structuring element i.e. the max-pooling kernel, enables the propagation of the border pixels to the interior portion of the feature map which should otherwise be disconnected and as a result, the infeasible included segment will not be opened correctly. Thus, inclusion loss will be computed with a very high value penalizing non-infeasible included segments.

4.5. Empirical Analysis

The feasible inclusion pairs are computed offline individually from the training set of each dataset using Algorithm 1. The complete count of the feasible and infeasible pairs of inclusion are given in Table 5. Every dataset is divided into train/val/test set. Algorithm 1 is applied to the train set of every dataset even when we jointly re-trained the models on $train + val$ sets.

5. Conclusion and Future Work

In this paper, rather than building a new semantic segmentation architecture, we have attempted to enhance any existing state-of-the-art network by enforcing a physical, data-driven prior into the network. Our method, PhyFea, constitutes a light-weight, plug-and-play module for semantic segmentation, which can be coupled with any network architecture to improve accuracy and conformity to our exemplary inclusion prior. We thus believe it is a first yet firm step towards physically-informed semantic perception. Overall, we are successful in attaining our goal of physical feasibility in segmentation both in the qualitative and quantitative sense.

References

[1] Liang-Chieh Chen, Jonathan T Barron, George Papandreou, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In *CVPR*, 2016. 2

[2] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L Yuille. Attention to scale: Scale-aware semantic image segmentation. In *CVPR*, 2016. 1

[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, 2018. 6

[4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 1

[5] 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

[6] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1280–1289, 2022. 1, 2, 6, 7

[7] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. https://github.com/open-mmlab/mmsegmentation, 2020. 6

[8] 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 (CVPR)*, 2016. 2, 4, 5, 6

[9] Henghui Ding, Xudong Jiang, Ai Qun Liu, Nadia Magnenat Thalmann, and Gang Wang. Boundary-aware feature propagation for scene segmentation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6818–6828, 2019. 2

[10] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1

[11] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation, 2019. 2

[12] Lixue Gong, Yiqun Zhang, Yunke Zhang, Yin Yang, and Weiwei Xu. Erroneous pixel prediction for semantic image segmentation. *Computational Visual Media*, 8:165–175, 2022. 2

[13] Rafael Corsino González, Richard E. Woods, and Barry R. Masters. Digital image processing, third edition. *Journal of Biomedical Optics*, 14:029901, 2009. 2

[14] Meng-Hao Guo, Cheng-Ze Lu, Qibin Hou, Zhengning Liu, Ming-Ming Cheng, and Shi-Min Hu. Segnext: Rethinking convolutional attention design for semantic segmentation. *Advances in Neural Information Processing Systems*, 35:1140–1156, 2022. 1, 2

[15] 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. 1

[16] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. CCNet: Criss-Cross Attention for Semantic Segmentation . In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 603–612, Los Alamitos, CA, USA, 2019. IEEE Computer Society. 2

[17] Tsung-Wei Ke, Jyh-Jing Hwang, Ziwei Liu, and Stella X. Yu. Adaptive affinity fields for semantic segmentation. In *Computer Vision – ECCV 2018: 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part I*, page 605–621, Berlin, Heidelberg, 2018. Springer-Verlag. 1, 2

[18] Alexander Kirillov, Evgeny Levinkov, Bjoern Andres, Bogdan Savchynskyy, and Carsten Rother. InstanceCut: From Edges to Instances with MultiCut . In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7322–7331, Los Alamitos, CA, USA, 2017. IEEE Computer Society. 2

[19] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, page 109–117, Red Hook, NY, USA, 2011. Curran Associates Inc. 3

[20] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, page 282–289, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. 2

[21] Justin Liang, Namdar Homayounfar, Wei-Chiu Ma, Yuwen Xiong, Rui Hu, and Raquel Urtasun. PolyTransform: Deep Polygon Transformer for Instance Segmentation . In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9128–9137, Los Alamitos, CA, USA, 2020. IEEE Computer Society. 2

[22] Ziwei Liu, Xiaoxiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Deep learning markov random field for semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 40(8):1814–1828, 2017. 2

[23] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows . In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, Los Alamitos, CA, USA, 2021. IEEE Computer Society. 1

[24] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation . In *2015 IEEE Conference on Computer Vision and Pattern Recog-*

nition (CVPR), pages 3431–3440, Los Alamitos, CA, USA, 2015. IEEE Computer Society. 1

[25] Jiajia Ni, Jianhuang Wu, Jing Tong, Zhengming Chen, and Junping Zhao. Gc-net: Global context network for medical image segmentation. *Computer Methods and Programs in Biomedicine*, 190:105121, 2020. 2

[26] Masoud S Nosrati and Ghassan Hamarneh. Incorporating prior knowledge in medical image segmentation: a survey. *arXiv preprint arXiv:1607.01092*, 2016. 3

[27] George Papandreou, Liang-Chieh Chen, Kevin Murphy, and Alan L Yuille. Weakly- and semi-supervised learning of a dcnn for semantic image segmentation. In *ICCV*, 2015. 1

[28] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2

[29] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. 2

[30] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "grabcut": Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, 2004. 3

[31] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The Adverse Conditions Dataset with Correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 6

[32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 1

[33] Andrea Tirelli, Danyella O. Carvalho, Lucas A. Oliveira, José P. de Lima, Natanael C. Costa, and Raimundo R. dos Santos. Unsupervised machine learning approaches to the q-state potts model. *The European Physical Journal B*, 95(11), 2022. 2

[34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. 2

[35] Luminita Vese and Tony Chan. A multiphase level set framework for image segmentation using the mumford and shah model. *International Journal of Computer Vision*, 50, 2004. 2

[36] 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. *TPAMI*, 2019. 1, 6

[37] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 548–558, 2021. 1

[38] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. 2

[39] Xinming Wu, Jianwei Ma, Xu Si, Zhengfa Bi, Jiarun Yang, Hui Gao, Dongzi Xie, Zhixiang Guo, and Jie Zhang. Sensing prior constraints in deep neural networks for solving exploration geophysical problems. *Proceedings of the National Academy of Sciences*, 120(23):e2219573120, 2023. 2

[40] Xiaoyang Xiao, Yuqian Zhao, Fan Zhang, Biao Luo, Lingli Yu, Baifan Chen, and Chunhua Yang. Baseg: Boundary aware semantic segmentation for autonomous driving. *Neural Networks*, 157:460–470, 2023. 2

[41] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Neural Information Processing Systems (NeurIPS)*, 2021. 1, 2, 4, 6, 7

[42] Jimei Yang, Brian Price, Scott Cohen, Honglak Lee, and Ming-Hsuan Yang. Object contour detection with a fully convolutional encoder-decoder network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 193–202, 2016. 2

[43] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Springer-Verlag*, page 334–349, Berlin, Heidelberg, 2018. Springer-Verlag. 2

[44] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations*, 2016. 1

[45] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *ECCV*, 2020. 1, 2, 6, 7

[46] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Haibin Lin, Zhi Zhang, Yue Sun, Tong He, Jonas Mueller, R Manmatha, et al. Resnest: Split-attention networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2736–2746, 2022. 1, 2

[47] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid Scene Parsing Network . In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 6

[48] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Psanet: Point-wise spatial attention network for scene parsing. In *Computer Vision – ECCV 2018*, pages 270–286, Cham, 2018. Springer International Publishing. 2

[49] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ADE20K dataset. *International Journal of Computer Vision*, 127:302–321, 2019. 2, 6

[50] Zhen Zhu, Mengdu Xu, Song Bai, Tengteng Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic

segmentation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 593–602, 2019. 2

Physically Feasible Semantic Segmentation

Supplementary Material

6. Additional Qualitative Comparisons

ADE20K. In Fig. 8 showing additional qualitative results on the ADE20K val set, we observe that in the prediction of the basic OCRNet network on the first example building includes sky, which is rectified in the PhyFea-upgraded version of the network. In the second example, we observe that a segment of bed is infeasibly included in a cupboard segment, which is corrected by our PhyFea model.

ACDC. In the first example of Fig. 9, we observe that a segment of sky is completely included in a building, which is a physical anomaly and is solved by PhyFea.

Cityscapes. In Fig. 10, we note that a rider segment corresponding to the rider's arm is infeasibly included in a bus segment in the prediction of Mask2Former [6], whereas PhyFea manages to correctly connect this part of the arm with the rest of the rider's body. In the second example, a sky segment is infeasibly included in a building segment in the Mask2Former prediction, while PhyFea avoids this error and better segments distant people which appear very small in the image.

7. Analysis of Prediction Feasibility

We define a metric named mean infeasibility normalized frequency (mINF) on an evaluation set using:

$$\text{mINF} = \frac{1}{|\mathcal{C}'|} \sum_{(c_i,c_j)\in\mathcal{C}'} \frac{f_{\text{inclusion}}(c_i,c_j)}{f_{\text{co-occur}}(c_i,c_j)} \tag{6}$$

We compute $f_{\text{co-occur}}(c_i, c_j)$, i.e., the frequency of co-occurrence of the classes in the infeasible class pair (c_i, c_j) in the set of predicted labelings on the evaluation set. $f_{\text{inclusion}}(c_i, c_j)$ is the frequency of the occurrence of infeasible inclusions for class pair (c_i, c_j) in the set of predicted labelings on the evaluation set. An mINF of 0 signifies the complete absence of infeasible inclusions in the predicted labelings. We report mINF scores in Tab. 6 for SegFormer and Mask2Former models on Cityscapes and ACDC. These results show that PhyFea can almost entirely solve the anomaly of infeasible Inclusions. It achieved an mINF of almost 0

Table 6. **Comparison of prediction feasibility with state-of-the-art models on Cityscapes and ACDC.** Results are reported with our mINF metric (%, lower is better) on the Cityscapes test set and the ACDC test set.

Method \ Dataset	**Cityscapes**	**ACDC**
SegFormer [41]	4.0	8.0
PhyFea w/ SegFormer (ours)	**0.1**	**0.2**
Mask2Former [6]	3.2	7.3
PhyFea w/ Mask2Former (ours)	**0.1**	**0.1**

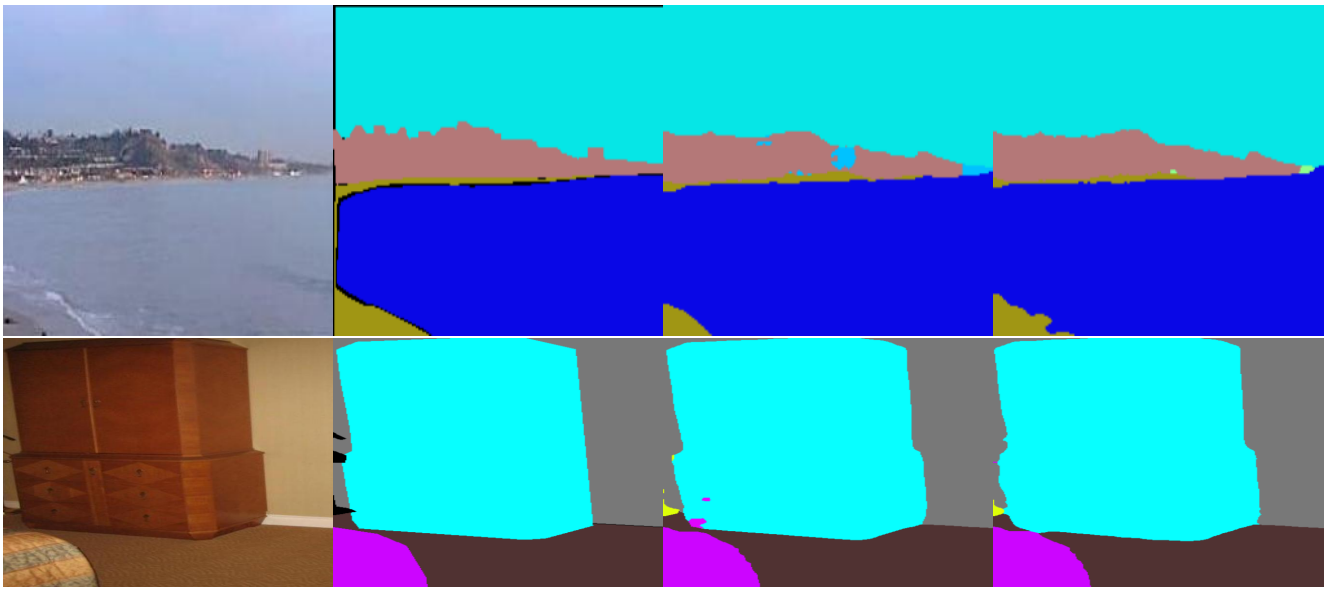

Figure 8. **Additional qualitative comparison on ADE20K.** From left to right: input image, ground-truth semantic labels, and predictions of OCRNet [45] and PhyFea. Best viewed on a screen and zoomed in.

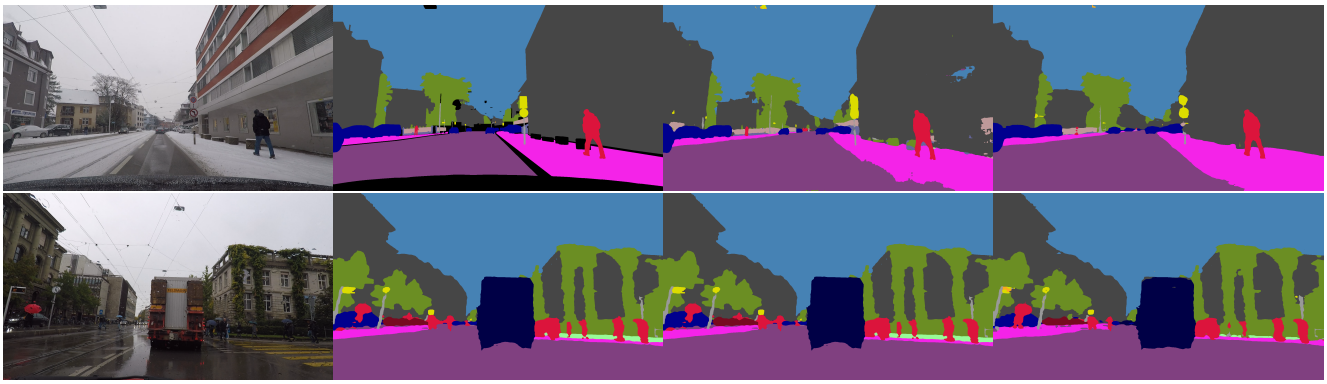

Figure 9. **Additional qualitative comparison on ACDC.** From left to right: input image, ground-truth semantic labels, and predictions of SegFormer [41] and PhyFea. Best viewed on a screen and zoomed in.

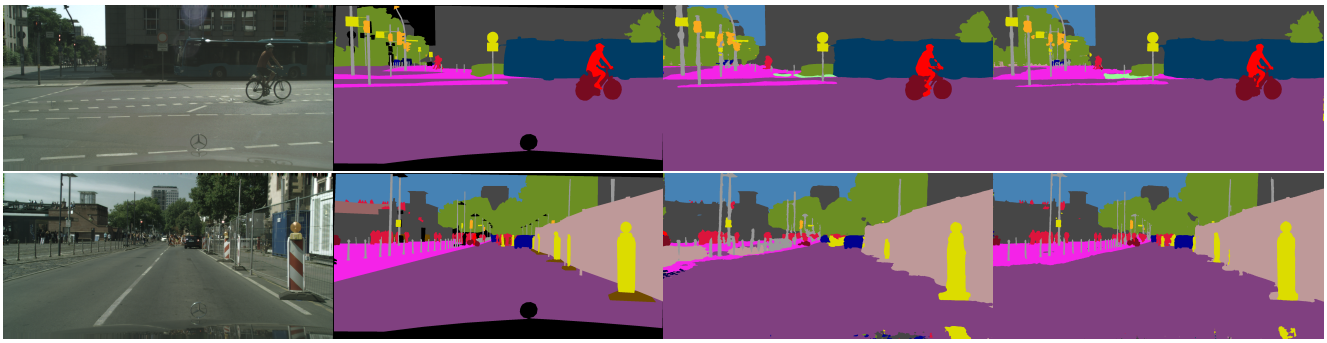

Figure 10. **Additional qualitative comparison on Cityscapes.** From left to right: input image, ground-truth semantic labels, and predictions of Mask2Former [6] and PhyFea. Best viewed on a screen and zoomed in.