

Importance of Self-Consistency in Active Learning for Semantic Segmentation

S. Alireza Golestaneh
sgolesta@andrew.cmu.edu

The Robotics Institute
Carnegie Mellon University
Pittsburgh, USA

Kris M. Kitani
kkitani@cs.cmu.edu

Abstract

We address the task of active learning in the context of semantic segmentation and show that self-consistency can be a powerful source of self-supervision to greatly improve the performance of a data-driven model with access to only a small amount of labeled data. Self-consistency uses the simple observation that the results of semantic segmentation for a specific image should not change under transformations like horizontal flipping (*i.e.*, the results should only be flipped). In other words, the output of a model should be consistent under equivariant transformations. The self-supervisory signal of self-consistency is particularly helpful during active learning since the model is prone to overfitting when there is only a small amount of labeled training data. In our proposed active learning framework, we iteratively extract small image patches that need to be labeled, by selecting image patches that have high uncertainty (high entropy) under equivariant transformations. We enforce pixel-wise self-consistency between the outputs of segmentation network for each image and its transformation (horizontally flipped) to utilize the rich self-supervisory information and reduce the uncertainty of the network. In this way, we are able to find the image patches over which the current model struggles the most to classify. By iteratively training over these difficult image patches, our experiments show that our active learning approach reaches $\sim 96\%$ of the top performance of a model trained on all data, by using only 12% of the total data on benchmark semantic segmentation datasets (*e.g.*, CamVid and Cityscapes).

1 Introduction

Deep convolutional neural networks (CNNs) are hungry for large and costly labeled training samples. CNNs have been successfully applied to different areas of research in computer vision and pattern recognition, such as classification [69, 60], detection [60], and semantic segmentation [0, 0, 0, 0, 46, 51, 69]. Semantic segmentation is one of the core tasks in visual recognition and has recently witnesses a lot of progress.

Despite the remarkable performance of existing semantic segmentation models, they have a significant drawback; they need a tremendous amount of labeled data to be able to learn their large number of parameters. A big challenge in semantic segmentation is obtaining labeled data. Active learning and weakly- and semi-supervised learning [02, 61, 65, 67, 42] approaches that aim to mitigate the data demands of semantic segmentation. To circumvent the burden of labeling large amounts of data, Active Learning (AL) proposed a compelling solution by iteratively selecting the *most informative* samples to label. Active learning can significantly reduce the cost of human labeling. The goal is to focus human

annotation effort towards image regions that reduce the uncertainty of the model thus improving model performance.

In this paper, we focus on the uncertainty pool-based active learning [43, 54], which is the most practical and commonly used approach among the existing active learning methods. In this method, softmax probabilities from the model’s outputs is used to compute entropy as an uncertainty criterion for the future manual annotation [18, 24, 54]. One shortcoming of this approach is that if the model begins to overfit to the small amount of training data (*i.e.*, poor generalization) gathered through active learning, the data sample selection process is negatively impacted. An overfit model has high uncertainty for uninformative data samples and the model makes mistakes with high certainty on other data samples. Labeling such data samples during the active learning process will hinder the model from reaching best performance. A model with poor generalization also increases the sampling bias problem [1, 3, 50] in active learning, which decreases the quality of the selected data samples.

We tackle the problem of active learning for semantic segmentation, aiming to improve the robustness and generalization of the selected samples and model. Our paper makes the following contributions: i) We propose to use an equivariant transformation of the input image, in addition to use the input image, to improve the computed uncertainty. Specifically, in the active learning process, we propose to compute the uncertainty from the input image and its equivariant transformed version (in this paper we use horizontal flipping as the equivariant transformation). ii) We further propose to enforce pixel-wise self-consistency between the outputs of segmentation network for each image and its transformation to utilize the rich self-supervisory information and reduce the uncertainty of the network. In other words, given an input image and its equivariant transformation we enforce the network to output similar pixel-wise predictions. iii) We demonstrate the effectiveness of our proposed method by evaluating it on two public datasets (CamVid and Cityscapes). Our approach achieves state-of-the-art results as it is able to maintain better generalization and robustness over uncertainty estimation and model generalization compared to other methods.

The rest of the paper is organized as follows. Some of the related works are introduced in Section 2 and the proposed method is explained in Section 3. Section 4 shows experimental results and finally, Section 5 concludes the paper.

2 Related Work

Semantic segmentation assigns a class label to each image pixel. Fully Convolutional Networks (FCNs) [46] have become indispensable models for semantic segmentation. In the past few years, the accuracy of semantic segmentation has been improved rapidly [2, 2, 2, 44, 51, 59] thanks to deeply learned features [25, 39, 57] and large-scale annotations [11, 45]. Semantic segmentation is critical for a variety of applications such as autonomous driving [11, 15, 21], robot manipulation [53], visual question answering [52], and medical image analysis [40, 51, 65]. Two well-known challenges in semantic segmentation research are 1) acquiring dense pixel-wise labels, which is expensive and time-consuming, and 2) achieving good generalization and performance when a dataset has highly imbalanced classes.

Active learning tackles the problem of finding the most crucial data in a set of unlabeled samples such that the model gains the most if it is given the label of that sample. It has been successfully used as a method for reducing labeling costs. A large number of active learning techniques have been developed in the literature [54]. There are mainly three query scenarios for active learning: (i) membership query synthesis, (ii) stream-based selective sampling, and (iii) pool-based active learning. In this work, we focus on pool-based active

learning, which is the most practical and common query scenario among different computer vision problems [6, 20, 26, 29, 33, 48, 49, 57, 58, 56]. In pool-based active learning, there exists a large unlabeled pool of data, and usually an initial small labeled set of data. At each iteration of the process, the model is trained on the labeled set, then an acquisition function chooses one or several samples (regions in this work) from the unlabeled set to be labeled by an external oracle and added to the labeled set. This process is repeated until a specific budget of labeled data is exhausted or until a certain model performance is reached.

Sampling strategies in pool-based approaches have been built upon different methods, which are surveyed in [54]. Informativeness and representativeness are two important criteria when judging the importance of each sample in the unlabeled dataset. Informativeness measures the ability of a sample to reduce the generalization error of the adopted model, and ensure less uncertainty of the model in the next iteration, while representativeness refers to whether a sample can represent the underlying structure of the unlabeled dataset [40].

Informative samples are usually identified by using the uncertainty of the model for each sample. Selecting the *true* informative and representative sample highly correlated with the generalizability of the model. If the model suffers from high variance and poor generalization, it will not identify the *most informative* sample for labeling. This causes requesting more data to be labeled to reach the desired performance. Different methods such as importance sampling [49], ensembling [46, 41], Monte Carlo (MC) dropout [47] proposed to address the aforementioned challenges.

Here we approach this problem from a different angle, our approach is not against any of the proposed methods and can be used along with them. During the active selection, we propose to compute the uncertainty from the input image and its equivariant transformed version (in this paper we use horizontal flipping as the equivariant transformation). In this way, the computed uncertainty takes account of the model generalizability for both input and its transformation. Therefore, the most uncertain sample will identify with more confidence. Furthermore, during the model’s updating stage, we propose to use pixel-wise self-consistency, in addition to cross-entropy (CE) loss, to reduce the uncertainty of the outputs of the segmentation network between each sample and its transformation.

Active learning for semantic segmentation While a variety of active learning approaches proposed for image classification, the active learning for semantic segmentation has been less explored [8, 14, 23, 34, 38, 47, 48, 56, 62, 63, 65], perhaps due to its intrinsic computational difficulties as we have to deal with much larger output spaces. In contrast to active learning for classification, in active learning for semantic segmentation, each image has pixel-wise predictions. Also, the labeling cost for each image or region could be different.

Active learning for semantic segmentation with deep neural networks has been specifically investigated in [8, 23, 34, 47, 56, 65]. In [23, 65], the authors utilized MC dropout [47] to model the pixel-wise uncertainty and focused on foreground-background segmentation of medical images. Siddiqui *et al.* [56] used multiple views for each image and proposed a method to estimate model’s uncertainty based on the inconsistency of predictions across different views. In [34, 47], the authors designed hand-crafted methods and focused on cost-effective approaches, where the cost of labeling for regions of an image is not considered equal. However, labeling cost information is not always available, which restricts their applicability. Casanova *et al.* [8] proposed an active learning strategy for semantic segmentation based on deep reinforcement learning, where an agent learns a policy to select a subset of image regions to label. Existing approaches depend entirely on the labeled regions, which had the most uncertainty during the active selection process, to update and retrain the network.

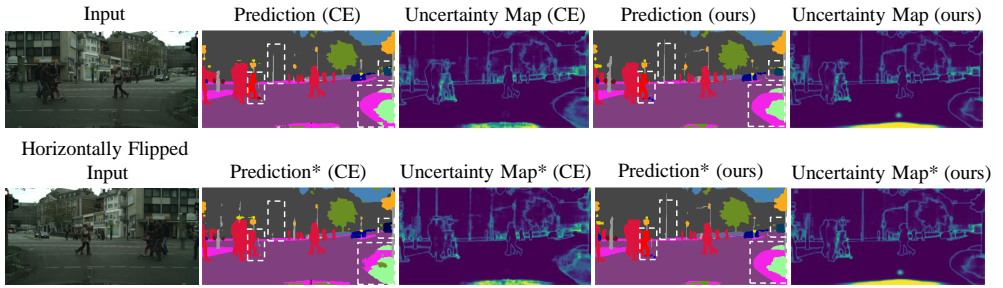


Figure 1: Inconsistencies between segmentations. The input image in the second row is horizontally flipped version of the image from the first row. While horizontal flipping is used as an augmentation method during the training, the second column shows that the model is still suffering from poor generalization. CE denotes that model only used cross-entropy loss during training. The fourth column shows the results while the model used our proposed self-consistency loss in addition to CE during the training. * means that the image is flipped horizontally for demonstration purposes. For all uncertainty maps, blue indicates low values, and yellow indicates high values.

Poor generalization of CNNs Deep models with millions of parameters can suffer from poor generalization and high variance (specially in a low regime data situation), which further affects the active selection process. Overfitting and poor generalization in deep neural networks are well-known problems and an active area of research. Different approaches such as dropout [59], ensembling [70], and augmentation proposed to increase the generalizability in CNNs. In semantic segmentation, due to large models, making predictions using a whole ensemble of models is cumbersome and is too computationally expensive to allow deployment to a large number of users. Data augmentation has been successful in increasing the generalizability of CNNs. However, as shown in Fig. 1, the model can still suffer from poor generalizations for a simple transformation. It is shown in [10, 64] that although using different augmentation methods improve the generalization of CNNs, neural networks are still sensitive to equivalent perturbations in data. Semantic segmentation models are not excepted from this issue. For instance, as we show in Fig. 1, a semantic segmentation model that used image flipping as an augmentation during training still fails to have a robust prediction for an image and its flipped version. This kind of high variance can affect the active learning process directly and causes to select the wrong sample as the most uncertain/informative one.

Self-Supervised Learning is a general learning framework that relies on pretext (surrogate) tasks that can be formulated using unsupervised data. A number of self-supervised methods have been proposed, each exploring a different pretext task [13, 22, 27, 32, 58]. A pretext task is designed in a way that solving it requires learning of a useful image representation.

Inspired by the aforementioned observation and sensitivity of semantic segmentation models for equivalent transformations, we propose to utilize the model’s uncertainty for the input image and its equivariant transformation during the active sample selection process. We also propose to enforce pixel-wise self-consistency via the self-supervisory signal between each image and its equivariant transformation to increase the robustness and generalization of the model. To the best of our knowledge, all current approaches for active learning in se-

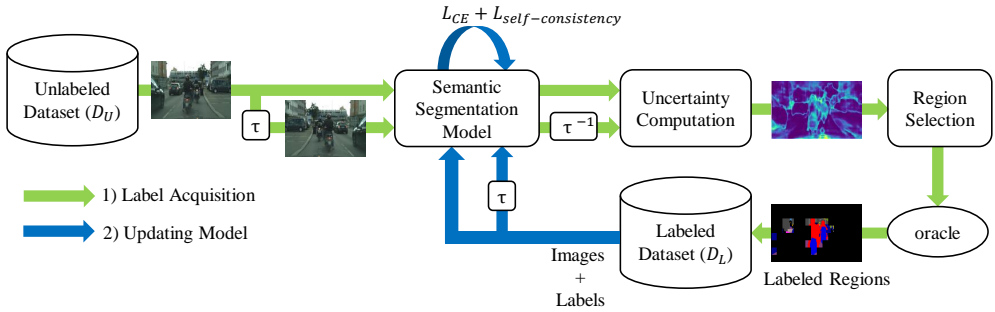


Figure 2: An overview of our proposed uncertainty pool-based active learning method. First, at each iteration of active selection, the semantic segmentation network computes the uncertainty for each image in \mathcal{D}_U and its transformation. Then, we select K regions with the most uncertainty and send them to an oracle for labeling. The labeled regions will be added to \mathcal{D}_L and removed from \mathcal{D}_U . Second, the model gets updated via the labeled regions in \mathcal{D}_L . This is repeated until the labeling budget is exhausted, or until a certain model performance is achieved.

semantic segmentation rely only on the labeled areas for updating the model. In our proposed method, we add a new regularization that penalizes the differences between the outputs of the model for the input image and its equivariant transformation in a self-supervised manner. Specifically, during the active selection process, we use both input and its equivariant transformation¹ to compute the uncertainty. During the model updating stage, in addition to the cross-entropy loss, we enforce the predictions for the input sample and its transformation to be consistent. Our proposed approach benefits from both labeled and unlabeled regions and reduces the model’s variance, which leads to a more robust sample selection. In other words, during the active learning process, not only we use the labeled regions for updating the model, but also we use the whole image (both labeled and unlabeled regions) to enforce the self-consistency between the image and its transformation. Finally, we show that our proposed pixel-wise self-consistency can also be used in the final model retraining stage, after the active learning process is finished, to improve the semantic segmentation performance. Our model is not against any of the existing active learning methods and can be used as a complementary along any of them.

3 Method

We describe our active learning method for incorporating self-consistency over equivariant transformations of the input, to achieve high performance using only a small amount of labeled data. Formally, let \mathcal{D}_U and \mathcal{D}_L denote the unlabeled and labeled datasets, respectively. Our goal is to achieve the highest performance possible through iteratively adding new labeled data until we reach a certain budget (*i.e.*, maximum number of labeled data) or until a desired level of performance is achieved. Let f_ϕ represent the segmentation network with learnable parameters ϕ . We select small image patches to label from unlabeled image regions such that they will maximize the performance of the segmentation model. The initial \mathcal{D}_L can contain a small amount of labeled data or it can be empty.

Fig. 2 shows an overview of our proposed method. At each iteration t , the goal of

¹In this paper we use horizontal flipping as the equivariant transformation.

active learning is to select an unlabeled image patch to be labeled and add it to the training dataset \mathcal{D}_L . To identify which unlabeled image patch to label, we use uncertainty sampling. Uncertainty sampling selects image regions with high uncertainty according to our current model. Regions with high uncertainty correspond to image regions that are difficult for the semantic segmentation network to classify. Let \mathbf{x}_i denote the i^{th} image in \mathcal{D}_U . At the i^{th} iteration we have two outputs corresponding to the original image and a transformed version of the image (*i.e.*, flipped image):

$$I_{i,0}^{(u,v)}(c) = f_i(\mathbf{x}_i; \phi) \text{ and } I_{i,1}^{(u,v)}(c) = \tau^{-1}(f_i(\tau(\mathbf{x}_i); \phi)), \quad (1)$$

$$P_{i,j}^{(u,v)}(c) = \text{softmax}(I_{i,j}^{(u,v)}(c)), \quad (2)$$

$$H_i^{(u,v)} = -\sum_j \sum_c P_{i,j}^{(u,v)}(c) \cdot \log(P_{i,j}^{(u,v)}(c)), \quad (3)$$

where $I_{i,0}^{(u,v)}(c)$ denotes the output logits of pixel (u, v) belonging to class c obtained from the network f_ϕ for an input image \mathbf{x}_i . τ and τ^{-1} denote an equivariant transformation and its inverse, respectively, where $\tau^{-1}(\tau(\mathbf{x}_i)) = \mathbf{x}_i$. In our experiments, we use horizontal flipping as the equivariant transformation. $P_{i,j}^{(u,v)}(c)$ is the softmax probability of pixel (u, v) belonging to class c obtained from $I_{i,j}^{(u,v)}(c)$, where $j \in \{0, 1\}$. Interestingly, for many semantic segmentation networks, flipping the image will result in slightly different results such that $I_{i,0}^{(u,v)}(c) \neq I_{i,1}^{(u,v)}(c)$. In Eq. 3, $H_i^{(u,v)}$ denotes the pixel-wise entropy (estimated uncertainty) for i^{th} image, where it is computed by aggregating the posterior uncertainties over network predictions for image \mathbf{x}_i and its transformation. $H_i^{(u,v)}$ captures the uncertainty of the network not only for input \mathbf{x}_i , but also for its transformation.

To exploit the structure and locality in semantic segmentation, we argue for selecting regions to be labeled instead of whole images. In particular, we opt for using fixed sized rectangular regions. Given H_i as the pixel-wise entropy for image i , we compute the next region to label by dividing H_i into M fixed-size rectangular regions denoted by H_i^m where $m \in M$. We then look for the region among all the unlabeled data that has the highest entropy:

$$(s, n) = \underset{(i,m)}{\text{arg max}} H_i^m, \quad (4)$$

where H_s^n has the maximum uncertainty among the unlabeled data. The selected region will not be considered for further selection. Eq. 4 will be repeated K times, where K is the number of samples that are set to be labeled at each iteration by an oracle.

In our experiments, during the label acquisition process, instead of using a real annotator, we simulate annotation by using the ground truth annotation of the region as the annotation from the oracle. Images which contained labeled regions are then added to the labeled dataset. If an image is fully labeled, it will be removed from the unlabeled dataset. The labeled dataset, therefore, is comprised of a lot of images, each with a subset of their regions labeled. The unlabeled regions of these images are marked with the ignore label. For each image x_l where $l \in |\mathcal{D}_L|$ we have the ground truth y_l . After each label acquisition iteration, the network will get updated via the labeled data. In all of the existing active learning works for semantic segmentation, the network gets updated via the labeled data to minimize the cross-entropy loss. However, for each image, there are unlabeled regions with rich information that can be exploited to reduce the uncertainty of the network. In this work, we propose to use both labeled and unlabeled regions to increase the self-consistency of the network. At

each iteration, after each acquisition step, the semantic segmentation network will retrain as following:

$$\min_{\phi} L_{\text{total}} \quad \text{where } L_{\text{total}} = L_{\text{CE}} + L_{\text{self-consistency}}, \quad (5a)$$

$$L_{\text{CE}} = - \sum_j \sum_{h,w,c} y_l \cdot \log(P_{i,j}), \quad (5b)$$

$$L_{\text{self-consistency}} = \frac{2}{W \times H \times C} \| I_{i,0}^{(u,v)}(c) - I_{i,1}^{(u,v)}(c) \|_2^2, \quad (5c)$$

where L_{CE} is the standard supervised pixel-wise cross-entropy loss, where it is evaluated only for the regions that are labeled. Eq. (5c) is the self-consistency loss, which is used to utilize the information between the outputs of the model for input image and its corresponding transformation. In Eq. (5c), W , H , C denote width, height, and depth of the logits, respectively. Enforcing the self-consistency between the outputs of model for the input and its transformation helps the network generalize better and have less variance.

After meeting the budget amount, the active selection process concludes with the retraining of the network with the final dataset \mathcal{D}_L , we refer that as the final retraining stage.

4 Results

4.1 Datasets

We evaluate our method on two semantic segmentation datasets: CamVid and Cityscapes.

CamVid dataset [1] composed of street scene view images, where each images has 360×480 resolution and contains 11 semantic classes. CamVid has 367, 101 and 233 images for train, validation, and test set, respectively. **Cityscapes** dataset [10] is a large dataset that focuses on semantic understanding of urban street scenes. The dataset contains 2048×1024 images with fine annotations across different cities, different seasons, and varying scene layout and background. The dataset is annotated with 30 categories, of which 19 categories are included for training and evaluation (others are ignored). The training, validation, and test set contains 2975, 500, and 1525 fine images, respectively.

4.2 Implementation Details

In our experiments, for all the methods, we use the same segmentation network (ResNet50 [16] + FCN [23]).² To analyze the true potential of each active learning method, for all the experiments, for the active sample selection process, the segmentation network is initialized with random weights. After the active learning process is finished, we retrain the networks (i.e. the final retraining stage) with the final dataset \mathcal{D}_L while the ResNet50 backbone is initialized with weights from a model that was pre-trained on the ILSVRC 1000-class classification task [17]. For all the experiments, we use random horizontal flipping for the data augmentation. For the final retraining stage, we use 60 epochs. We set the batch size to five and four for CamVid and Cityscapes, respectively. For optimization, we use the adaptive moment estimation optimizer (ADAM) [18] with $\beta_1 = 0.9$, $\beta_2 = 0.999$, we set the initial learning rate and weight decay to $5e-4$ and $2e-4$, respectively. Learning rate decays is set to 0.1, and it applied after every 20 epochs. For all the experiments, we report the average of the five different runs.

²Since we do not change the network architecture, any other segmentation network can be used.




Figure 3: Active learning performance for our proposed method and other baselines. (a) and (b) demonstrate the mIoU results for CamVid and Cityscapes datasets, respectively. The horizontal solid lines at the top represent the performance with the entire training set labeled. The dashed line represents 96% of the maximum model performance (Fully-Supervised). Budget denotes the percentage of unlabeled data labeled during the active learning process.

For both datasets, we use their training sets as \mathcal{D}_U . For CamVid we evaluate the performance on the test set and report the mIoU. We use 360×480 images as an input of the network; each image is divided into 25 non-overlapping regions of dimension 72×96 . We chose $K = 1$ region per step. For Cityscapes we report the mIoU on the validation set since the test set is not available publicly. We resize the images to 512×1024 because of the computational complexity. Each image is divided into 32 non-overlapping regions of dimension 128×128 . We chose $K = 5$ regions per step.

4.3 Experiments

Fig. 3 demonstrates the mIoU results achieved by the segmentation model against the percentage of datasets labeled used to train the model. We evaluate our proposed method for different budget sizes against different baselines on two datasets. During the active selection, we use the same settings for all the methods. After collecting the labeled data, we retrain in the same environment and hyperparameters.

Fig. 3 summarizes the results of our experiments. In our experiments, at the final re-training stage, Fully-Supervised, Random, Entropy, and EquAL baselines are retrained via CE loss. Fully-Supervised+ and EquAL+ baselines are retrained via CE loss as well as the proposed pixel-wise self-consistency loss (Eq. 5).

Fig. 3 demonstrates that our proposed methods, EquAL and EquAL+, outperforms the other baselines for the same percentage of labeled data. We observe that for the low budget size, in both of the datasets, all the baselines suffer from the cold start problem [28, 35]. In contrast, our method achieves significantly better performance in the low regime data, which we believe is due to the better generalization. Particularly, in CamVid, for 8% budget we improved the entropy baseline by 2% (EquAL) and 7% (EquAL+). In Cityscapes, for 1% budget we improved the entropy baseline by 3% (EquAL) and 9% (EquAL+). As we move to larger budgets the networks can achieve better generalization. EquAL achieves 96% of the maximum model performance with 12% labeled data in both datasets, while the runner-up method does so with 16% and 13% data on the respective datasets, which corresponds to 14 and 29 additional images, respectively. Furthermore, we show that by using the self-consistency loss in addition to the CE loss we can consistency achieves better performance.

To measure the diversity of the labeled data, for 12% budget in the Cityscapes, we further compute the entropy of the class distribution for the data in the final labeled set; our model achieves higher entropy (2.376) comparing to the entropy baseline (2.335), which means more diversity in the selected samples.

Table 1 provides per-class IoU comparison between our model and various baselines, where the budget is fixed to 12%. All the models except EquAL+ uses CE loss during the final retraining stage. Table 1 shows that our proposed method, EquAL, outperforms other methods on 5 classes as well as mIoU and achieves competitive results for the other classes. EquAL+ denotes our model when our self-consistency loss is used in addition to CE loss during the final retraining stage and leveraged both labeled and unlabeled regions. In EquAL+, we regularize the dark knowledge of the model and reduces the high variance, which leads to a significantly better performance.

Table 1: Per-class performance comparison results for Citydataset while setting the labeling budget to 12%. Bold entries are the best and second-best performers.

Methods	Road	SideWalk	Building	Wall	Fence	Pole	Traffic Light	Traffic Sign	Vegetation	Terrain
Random	96.64	75.51	88.19	32.51	34.61	36.09	38.80	58.25	89.03	53.48
BALD [14]	96.22	74.08	87.71	33.68	39.73	40.78	44.90	40.67	89.13	52.79
Vote Entropy [14]	96.22	73.50	87.08	34.97	42.32	42.76	44.91	61.32	89.24	54.18
RALIS [8]	96.19	73.24	86.25	33.56	42.28	43.51	51.12	60.31	89.05	51.72
Entropy	96.11	73.72	87.78	33.76	40.40	44.50	44.78	61.32	89.22	53.77
EquAL (Ours)	96.05	73.05	87.66	40.94	42.31	43.79	46.22	61.58	89.14	56.15
EquAL+ (Ours)	97.01	77.90	89.06	41.74	44.93	46.07	51.78	65.44	89.88	58.33
	Sky	Person	Rider	Car	Truck	Bus	Train	Motorcycle	Bicycle	mIoU
Random	90.18	52.67	32.51	88.47	24.62	41.93	31.26	29.11	60.15	55.47
BALD [14]	89.81	65.73	39.52	89.54	34.65	55.59	36.14	31.27	63.33	58.17
Vote Entropy [14]	89.94	69.31	40.59	89.14	37.09	54.92	36.57	32.16	65.73	60.01
RALIS [8]	89.32	69.23	42.14	89.51	39.50	57.22	35.51	31.42	64.12	60.27
Entropy	90.53	68.81	41.09	89.52	37.87	54.83	36.15	31.03	65.51	60.03
EquAL (Ours)	90.58	68.10	40.75	89.32	37.39	58.63	35.03	31.27	64.49	60.65
EquAL+ (Ours)	92.05	72.30	47.74	91.06	48.53	63.75	36.98	42.19	68.11	64.46

Fig. 4 provides qualitative comparison between our model and entropy baseline for the same budget (12%). We observe that the entropy baseline (third column) fails to segment sidewalk and terrain comparing to our proposed model (forth column), EquAL. Moreover, we show in the fifth column (EquAL+) that we can even improve our results further by adding the pixel-wise self-consistency loss.

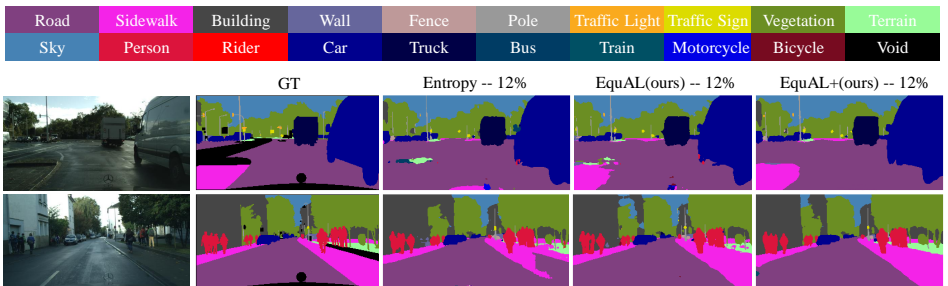


Figure 4: Qualitative results on the Cityscapes dataset. All the models are retrained after they use active learning with budget limit of 12%. The third column represent the entropy baseline, and last two columns represent our proposed models.

To study the effects of different equivariant transformations, in Table 2, we study effects

of employing five different equivariant transformations during the active learning process while using CamVid. It is not surprising to see that translation and horizontal flipping are the only two that improves the performance. We believe that this behavior is because transformations such as rotation or vertical flipping create images with conditions that do not exist in the dataset.

Table 2: Ablation study for different equivariant transformations during the active learning process.

CamVid				
Models	Budget			
	8%	12%	16%	20%
Entropy (Baseline)	0.5451	0.6111	0.6192	0.6311
EquAL (Rotation 90 degree)	0.5158	0.5855	0.5879	0.6036
EquAL (Rotation 180 degree)	0.5278	0.5982	0.6005	0.6072
EquAL (Random Translation 16 pixels)	0.5472	0.6121	0.6193	0.6294
EquAL (Vertical Flipping)	0.5266	0.5925	0.5987	0.6027
EquAL (Horizontal Flipping)	0.5673	0.6182	0.6274	0.6341

5 Conclusion

In this paper, we proposed an effective method for active learning for semantic segmentation. We propose to use an equivariant transformation of the input image, in addition to use the input image, to improve the computed uncertainty during the active learning process. We also proposed a pixel-wise self-consistency loss, which takes advantage of self-supervisory information and regularizes the inconsistency of the model predictions between the input image and its transformation. Our proposed method yields a more robust active sample selection as well as better segmentation performance. Specifically, our method achieves $\sim 96\%$ of maximum achievable network performance using only 12% labeled data on CamVid and Cityscapes datasets. Finally, by using our proposed pixel-wise self-consistency we improve the segmentation performance by 2-9%.

As future work, we highlight the possibility of designing an optimal stopping method that could help accelerate the active sample selection process, in this way the active learning algorithms do not require to iterate completely through the unlabeled set to select the most uncertain samples. Furthermore, if the community provides ground-truth information about human annotation costs for more datasets, it will be an important direction to develop cost-estimate and cost-effective methods to reduce the cost of labeling for different regions even further.

6 Acknowledgement

This project was sponsored by Caterpillar Inc.

References

- [1] Francis R Bach. Active learning for misspecified generalized linear models. In *Advances in neural information processing systems*, pages 65–72, 2007.
- [2] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.

- [3] Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance weighted active learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 49–56, 2009.
- [4] Gabriel J Brostow, Julien Fauqueur, and Roberto Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2): 88–97, 2009.
- [5] Arantxa Casanova, Pedro O Pinheiro, Negar Rostamzadeh, and Christopher J Pal. Reinforced active learning for image segmentation. *arXiv preprint arXiv:2002.06583*, 2020.
- [6] Devendra Singh Chaplot, Helen Jiang, Saurabh Gupta, and Abhinav Gupta. Semantic curiosity for active visual learning. *arXiv preprint arXiv:2006.09367*, 2020.
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [8] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [9] 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.
- [10] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999, 2016.
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [12] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1635–1643, 2015.
- [13] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015.
- [14] Suyog Dutt Jain and Kristen Grauman. Active image segmentation propagation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2864–2873, 2016.
- [15] Andreas Ess, Tobias Müller, Helmut Grabner, and Luc J Van Gool. Segmentation-based urban traffic scene understanding. Citeseer.

- [16] Yoav Freund, H Sebastian Seung, Eli Shamir, and Naftali Tishby. Selective sampling using the query by committee algorithm. *Machine learning*, 28(2-3):133–168, 1997.
- [17] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 740–755, 2016.
- [18] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1183–1192. JMLR. org, 2017.
- [19] Ravi Ganti and Alexander Gray. Upal: Unbiased pool based active learning. In *Artificial Intelligence and Statistics*, pages 422–431, 2012.
- [20] Yonatan Geifman and Ran El-Yaniv. Deep active learning with a neural architecture search. In *Advances in Neural Information Processing Systems*, pages 5974–5984, 2019.
- [21] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012.
- [22] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [23] Marc Gorriz, Axel Carlier, Emmanuel Faure, and Xavier Giro-i Nieto. Cost-effective active learning for melanoma segmentation. *arXiv preprint arXiv:1711.09168*, 2017.
- [24] Manuel Haußmann, Fred Hamprecht, and Melih Kandemir. Deep active learning with adaptive acquisition. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 2470–2476. AAAI Press, 2019.
- [25] 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.
- [26] Fabian Caba Heilbron, Joon-Young Lee, Hailin Jin, and Bernard Ghanem. What do i annotate next? an empirical study of active learning for action localization. In *European Conference on Computer Vision*, pages 212–229. Springer, 2018.
- [27] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in Neural Information Processing Systems*, pages 15663–15674, 2019.
- [28] Neil Houlsby, José Miguel Hernández-Lobato, and Zoubin Ghahramani. Cold-start active learning with robust ordinal matrix factorization. In *International Conference on Machine Learning*, pages 766–774, 2014.
- [29] Peiyun Hu, Zachary C Lipton, Anima Anandkumar, and Deva Ramanan. Active learning with partial feedback. *arXiv preprint arXiv:1802.07427*, 2018.

- [30] Sheng-Jun Huang, Rong Jin, and Zhi-Hua Zhou. Active learning by querying informative and representative examples. In *Advances in neural information processing systems*, pages 892–900, 2010.
- [31] Mostafa S Ibrahim, Arash Vahdat, Mani Ranjbar, and William G Macready. Weakly supervised semantic image segmentation with self-correcting networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2020.
- [32] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [33] Chieh-Chi Kao, Teng-Yok Lee, Pradeep Sen, and Ming-Yu Liu. Localization-aware active learning for object detection. In *Asian Conference on Computer Vision*, pages 506–522. Springer, 2018.
- [34] Tejaswi Kasarla, G Nagendar, Guruprasad M Hegde, V Balasubramanian, and CV Jawahar. Region-based active learning for efficient labeling in semantic segmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1109–1117. IEEE, 2019.
- [35] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017.
- [36] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [37] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European conference on computer vision*, pages 695–711. Springer, 2016.
- [38] Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua. Introducing geometry in active learning for image segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2974–2982, 2015.
- [39] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [40] Weicheng Kuo, Christian Häne, Esther Yuh, Pratik Mukherjee, and Jitendra Malik. Cost-sensitive active learning for intracranial hemorrhage detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 715–723. Springer, 2018.
- [41] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6405–6416. Curran Associates Inc., 2017.

- [42] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5267–5276, 2019.
- [43] David D Lewis and William A Gale. A sequential algorithm for training text classifiers. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 3–12, 1994.
- [44] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017.
- [45] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [46] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [47] Radek Mackowiak, Philip Lenz, Omair Ghori, Ferran Diego, Oliver Lange, and Carsten Rother. Cereals-cost-effective region-based active learning for semantic segmentation. *The British Machine Vision Association BMVC*, 2018.
- [48] Dwarikanath Mahapatra, Behzad Bozorgtabar, Jean-Philippe Thiran, and Mauricio Reyes. Efficient active learning for image classification and segmentation using a sample selection and conditional generative adversarial network. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 580–588. Springer, 2018.
- [49] Alejandro Pardo, Mengmeng Xu, Ali Thabet, Pablo Arbelaez, and Bernard Ghanem. Baod: Budget-aware object detection. *arXiv preprint arXiv:1904.05443*, 2019.
- [50] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [52] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [53] Max Schwarz, Anton Milan, Arul Selvam Periyasamy, and Sven Behnke. Rgb-d object detection and semantic segmentation for autonomous manipulation in clutter. *The International Journal of Robotics Research*, 37(4-5):437–451, 2018.

- [54] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [55] Jingyu Shao, Qing Wang, and Fangbing Liu. Learning to sample: an active learning framework. *arXiv preprint arXiv:1909.03585*, 2019.
- [56] Yawar Siddiqui, Julien Valentin, and Matthias Nießner. Viewal: Active learning with viewpoint entropy for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9433–9443, 2020.
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [58] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5972–5981, 2019.
- [59] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [60] Masashi Sugiyama. Active learning for misspecified models. In *Advances in neural information processing systems*, pages 1305–1312, 2006.
- [61] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [62] Alexander Vezhnevets, Joachim M Buhmann, and Vittorio Ferrari. Active learning for semantic segmentation with expected change. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3162–3169. IEEE, 2012.
- [63] Sudheendra Vijayanarasimhan and Kristen Grauman. What’s it going to cost you?: Predicting effort vs. informativeness for multi-label image annotations. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2262–2269. IEEE, 2009.
- [64] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017.
- [65] Lin Yang, Yizhe Zhang, Jianxu Chen, Siyuan Zhang, and Danny Z Chen. Suggestive annotation: A deep active learning framework for biomedical image segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 399–407. Springer, 2017.
- [66] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 93–102, 2019.

- [67] Licheng Yu, Xinlei Chen, Georgia Gkioxari, Mohit Bansal, Tamara L Berg, and Dhruv Batra. Multi-target embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6309–6318, 2019.
- [68] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE international conference on computer vision*, pages 1476–1485, 2019.
- [69] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [70] Zhi-Hua Zhou. *Ensemble methods: foundations and algorithms*. CRC press, 2012.