

MEAL: Manifold Embedding-based Active Learning

Deepthi Sreenivasaiyah, Johannes Otterbach, Thomas Wollmann
Merantix Labs GmbH, Berlin, Germany

{deepthi.sreenivasaiyah, johannes.otterbach, thomas.wollmann}@merantix.com

Abstract

Image segmentation is a common and challenging task in autonomous driving. Availability of sufficient pixel-level annotations for the training data is a hurdle. Active learning helps learning from small amounts of data by suggesting the most promising samples for labeling. In this work, we propose a new pool-based method for active learning, which proposes promising patches extracted from full image, in each acquisition step. The problem is framed in an exploration-exploitation framework by combining an embedding based on Uniform Manifold Approximation to model representativeness with entropy as uncertainty measure to model informativeness. We applied our proposed method to the autonomous driving datasets CamVid and Cityscapes and performed a quantitative comparison with state-of-the-art baselines. We find that our active learning method achieves better performance compared to previous methods.

1. Introduction

Success of computer vision methods based on machine learning depends heavily on the quality of labeled data available. Image Segmentation is one of the key sub-tasks in computer vision which has made tremendous progress over the past years [22]. Despite this progress, image segmentation methods based on Convolution Neural Networks (CNNs) [19] still pose the big hurdle of availability of sufficient pixel-level annotations for the training data [18]. In domains like autonomous driving [34] or biomedical image analysis [35, 24], large quantities of unlabeled data is available but the effort required to manually annotate the data pixel-wise is a major bottleneck [18] in terms of cost and time. In such scenarios, Active Learning (AL) is a framework where the machine learning model learns from small amounts of data and selects what data it wants to learn from [30]. In an active learning setup, a model is initially trained on a small labeled dataset. An acquisition function uses the trained model to select new samples from a pool of unlabeled data points, which are then labeled by

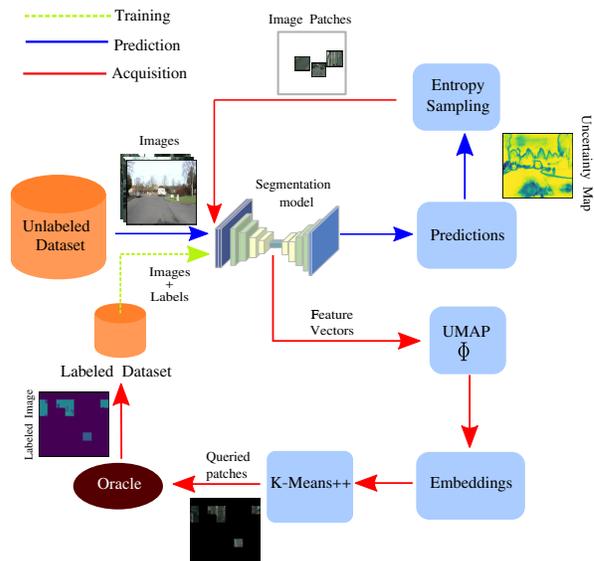


Figure 1: Overview of our proposed method for active learning: MEAL (Color best viewed online).

an expert (oracle). The newly labeled samples are added to the training dataset and a new model is trained using the updated training dataset. This process is repeated for multiple rounds with the training dataset size growing over the time to reach the required performance. Active learning has proven to drastically reduce the amount of labeled data needed to supervise a task like semantic segmentation [23, 14].

Although annotations for semantic segmentation is a costly and time consuming process compared to other visual recognition tasks, there are very few methods in active learning related to semantic segmentation especially in the field of autonomous driving [23, 14]. Currently, many AL methods operating on CNNs focus mostly on image classification tasks.

Most of the AL strategies in the literature involve evaluating either the informativeness [13, 14], which focuses on reducing the generalization error or the representative-

ness [21] that selects samples to maximize diversity of the unlabeled data. Current AL methods [2, 33, 35] that combine both of these measures do not consider the overall input dataset while designing the representativeness part. Having a trade-off between the two measures is an important consideration in active learning to select high quality and diverse samples. This trade-off could be achieved by using an exploration-exploitation framework [6]. In this framework, large regions of unlabeled data space can be explored to select diverse samples and at the same time, exploit smaller regions of unlabeled data space to select highly informative samples. In addition to combining the two measures, one can leverage the spatial coherence of an image to formulate a region-based approach to effectively deal with the large number of samples in a segmentation dataset and the costly annotation of samples for semantic segmentation.

In this work, we present a novel region-based active learning strategy, termed Manifold Embedding-based Active Learning, or MEAL. It selects data samples using informativeness and representativeness measures combined in an exploration-exploitation framework. The overview of our active learning algorithm is represented in Figure 1. For the informativeness measure, we use the Softmax probabilities from the output layer of the model to compute the entropy of image patches. To capture the diversity, a low-level manifold mapping is learned using the feature vectors of all the input images. This manifold embedding is then used to obtain a representative sample of embedding vectors from the selected informative samples. In contrast to previous approaches, our method frames the AL problem as an exploration-exploitation task for region-based active learning in semantic segmentation and can be used in conjunction with other uncertainty methods [14, 13, 3]. The method learns a low-level manifold mapping using Uniform Manifold Approximations (UMAP) [25] to discover a non-linear latent space projection of the data in order to retrieve diverse samples using K-Means++ [1] in the low dimensional space. To obtain the low-dimensional manifold mapping, we use a MobileNetV2 [15] backbone pre-trained on ImageNet [28] that is part of our full DeepLabV3+ [9] segmentation model. We compare our approach to several uncertainty-based active learning methods on the CamVid [7] and Cityscapes [10] datasets for autonomous driving.

In summary our contributions are:

- We rephrase the informativeness-representativeness trade-off of Active Learning in terms of the exploration-exploitation framework
- We introduce UMAP-based clustering methods to improve representativeness sampling
- We combine the new clustering with existing entropy-based informativeness sampling

- We present experimental evidence for the effectiveness of the combined MEAL method based on evaluations on CamVid and Cityscapes datasets.

The remainder of this paper is structured as follows. In Section 2, we review existing AL approaches for image classification and image segmentation. Our AL strategy’s theoretical foundations are outlined and formalized in Section 3. In Section 4, we present experiments with our AL strategy followed by discussion and future work in Section 5.

2. Related Work

Active Learning is a well explored topic in machine learning problems where labeled data is scarce or difficult to obtain [30, 12]. Settles [30] provides a rich reference collection of commonly used query strategy frameworks. He outlines several different settings for active learning, among those *pool-based* approaches [30] have gained more traction in recent times which assumes that there is a small set of labeled data and a large pool of unlabeled data available.

In terms of sample selection, active learning strategies can broadly be classified into uncertainty- [20] and diversity-based methods. In the former approach, the query strategy is based on some kind of informativeness measure aiming to reduce the model’s uncertainty. The intuition is that if a model is highly uncertain about its predictions on an unlabeled sample, then accommodating such a sample into the training is beneficial for the model. Culotta & McCallum [11] propose a method which selects unlabeled samples on which the classifier is least confident. Gal et al. [13] take a different approach by adapting Bayesian deep learning into the active learning framework. Their technique, Monte Carlo Dropout, performs stochastic forward passes to approximate the posterior distribution which induces the uncertainty in weights giving rise to prediction uncertainty on the unlabeled samples. However, Beluch et al. [3] prove that their ensemble method perform better in obtaining more calibrated predictive uncertainties which forms the basis for uncertainty-based AL. Diversity-based approaches query a subset of unlabeled data that represents the whole unlabeled dataset covering the underlying data distribution. Sener et al. [29] address the problem of AL as a coreset construction by minimizing the ℓ_2 distance between the activations of fully connected layers of queried point and unlabeled dataset. Another popular approach to query representative samples is running a clustering algorithm like K-Means [16] on the unlabeled dataset and selecting the samples equally from the clusters. This concept was demonstrated in [5] using spectral clustering. In [21], the authors present an unsupervised active learning approach by selecting representative samples. The method focuses on learning a non-linear embedding to transform the input into

a latent space where samples are selected using K-Means clustering [16].

The uncertainty-based methods *exploits* the input data space to select samples close to decision boundary while completely ignoring the representativeness aspect i.e., it fails to *explore* the larger regions of input data space to select diversified samples. Hence using such a method by itself does not improve performance [6].

On the other hand, there are hybrid methods which chose to combine uncertainty and diversity-based methods [2, 33, 35]. A query strategy called Batch Active learning by Diverse Gradient Embeddings (BADGE) [2] combines both diversity and uncertainty methods for batch acquisition by measuring uncertainty through gradient embeddings and diversity through K-Means++ initialization algorithm. Further improvement was made to this method in [33], where the authors proposed to calibrate the uncertainties by incorporating a calibration metric into the training objective and add a data augmentation technique to remove confirmation bias during pseudo labelling, while retaining the core active learning querying strategy as in [2]. In semantic segmentation, this kind of hybrid method can be seen in [35] where the uncertainty is computed using bootstrapping and representative samples are selected using cosine similarity between the unlabeled data samples and informative samples.

Uniform Manifold Approximation and Projection (UMAP): UMAP [25] is a graph-based non-linear dimensionality reduction method which builds a k-neighbor graph in the original data space and finds a similar graph in the lower dimension. The underlying assumption is the existence of a manifold on which the data is uniformly distributed and preserving the structure of this manifold is a primary goal. We use UMAP to learn the lower dimensional representation of the feature vectors. The main reason is that the graph constructed under this assumption preserves both the local and global structures.

Active Learning in Semantic Segmentation: Compared to image classification, the problem of active learning is less explored in the field of semantic segmentation. Mackowiak et al. [23] make use of region-based query strategy where information map is computed as the model’s prediction, to extract high informative regions from the image. They train a cost model which predicts the number of clicks needed to annotate each image. Another recent work in semantic segmentation is [14] which selects image regions that are highly uncertain under equivariant transformation. Siddiqui et al. [32] used viewpoint entropy of multiple views of each image as an uncertainty measure based on inconsistencies in the prediction across different views. Casanova et al. [8] proposed a query strategy based on deep reinforcement learning, where an agent learns a policy to select a subset of image regions to label. These existing meth-

ods that are related to semantic segmentation make use of only uncertainty based process to query the data and retrain the model.

3. Manifold Embedding-based Active Learning

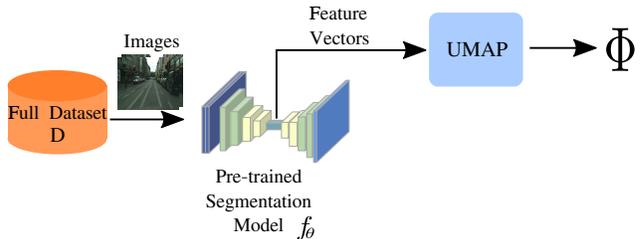


Figure 2: A low dimensional manifold representation Φ is learned with an UMAP trained on the overall unlabeled dataset D using a segmentation model f_θ loaded with pre-trained ImageNet weights (Color best viewed online).

Considering the success of querying the regions of image rather than the entire image in the active learning methods that are tailored for semantic segmentation [14, 23, 8], our method is also formulated to select image regions, based on a pre-determined tiling grid, during each acquisition step.

In our work, we present a novel active learning method that utilizes the exploration-exploitation framework in our query strategy, to achieve high performance on a segmentation model with a low number of labeled samples. The core idea of our approach *Manifold Embedding-based Active Learning (MEAL)* is to utilize the informativeness and representativeness in an exploration-exploitation framework so that the selected samples not only help in reducing the model’s uncertainty but are also highly representative for the underlying data distribution. MEAL starts by training an UMAP [25] across all input data to learn the manifold representation. This manifold representation is then used to construct a low-dimensional embedding of unlabeled samples to enable easier clustering. Converting the entire data space into this low dimensional embedding is computationally intense. To reduce the computational demands we split the sampling process into two steps. In the first step (i) we select a larger subset of highly informative samples using entropy-based uncertainty sampling and obtain the low-dimensional embeddings of these samples using the trained UMAP. In the second step (ii) we use the K-Means cluster centroid initialization algorithm [1] to compute the cluster centroids which we then add as new samples to the training dataset after labeling. In conjunction these two steps fulfill both the informativeness and the representativeness objective. A challenge with this approach is to select the right number of cluster centroids.

In the following, we describe the main components of our algorithm, uncertainty sampling and representative sampling, along with the notations in more detail.

Notations. Given the input data $x \subset X \in \mathbb{R}^{h \times w}$ where h, w are the spatial dimensions of the image, labels $c \in \{1, 2, \dots, C\}$ with C classes and label map $y \subset Y \in \mathbb{R}^{h \times w}$, let $D_L = \{x, y\}_{i=1}^{N_L}$, $D_U = \{x\}_{i=1}^{N_U}$ be the small labeled and a large pool of unlabeled data, respectively. $D \in \{D_U \cup D_L\}$ together forms the entire dataset. Let $f_\theta : X \rightarrow Y$ be the segmentation model with learnable parameters θ , and the manifold representation learned using UMAP be denoted by Φ where $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$, $d > d'$. The task of our active learning query strategy is to select N image patches at each acquisition step to be labeled by an oracle. This process is iterated for K steps and the selected samples are added to D_L dataset along with their corresponding labels.

3.1. Exploration-Exploitation Framework

When querying samples for labeling, there is often a trade-off to select samples which reduce model uncertainty or a subset of samples which represents the overall dataset. This is analogous to exploration and exploitation of the input data space. Exploitation, in the context of active learning corresponds to selecting uncertain samples from an already sampled data space while exploration refers to selecting samples in non-sampled areas of input space. Using UMAP [25] in our active learning query procedure, we learn a low dimensional manifold of the image patches from the entire input data space which preserves the local neighbourhood structure of the image patches, thus enabling to perform exploration in the lower dimensional space.

3.2. Learning the Dataset Manifold Approximation

The curse of dimensionality is an inherent problem when clustering high-dimensional data [4]. To address this issue and at the same time preserve the global context of the dataset, we employ UMAP [25] as a dimensionality reduction technique before clustering the data. Intuitively, this has two effects: (i) we increase the speed of the clustering algorithm by a factor proportional to the reduction in dimensionality and (ii) we increase the robustness of the clusters, since we are eliminating spurious correlations through projecting them into lower dimensionality. It should be pointed out that UMAP preserves local and global structure of the dataset topology and the resulting clusters in low dimensions are hence expected to represent relevant high dimensional clusters. For sake of simplicity, we limit the dimension of the embeddings to two in all our experiments. To train the UMAP embedding map, we make use of the entire dataset using the features obtained from feeding the pre-computed image patches through the segmentation model backbone. In a pool-based setting, we have access to the

entire unlabeled dataset at the beginning of active learning process and hence the patch-dataset includes labeled and unlabeled samples and thus information about the whole dataset is accounted for inside UMAP.

3.3. Uncertainty Sampling

For uncertainty sampling, we compute the entropy as an uncertainty measure [20]. Once the segmentation model is trained using the initial D_L labeled set, MEAL computes the prediction of all images from the unlabeled data pool D_U using the trained model f_θ . $P_c^{(u,v)}$ is the Softmax probability of the pixel (u, v) belonging to class c for an image I_i . The uncertainty map $H_i^{(u,v)}$ at each pixel position (u, v) of i^{th} image I_i is computed by aggregating probabilities of all the classes as follows.

$$H_i^{(u,v)} = - \sum_{c=1}^C P_c^{(u,v)} \log P_c^{(u,v)} \quad (1)$$

The uncertainty map will have high entropies (uncertainty) in those regions where the current model is not certain on the predictions and this sampling step selects such image regions.

Uncertainty maps have a high degree of spatial coherence, i.e. neighboring pixels are likely to have similar (small) entropy as they likely are in the same predicted class where the model is certain. To reduce labeling costs (as well as model compute intensity), we hence select image regions or patches with high entropy to be labelled. For this we divide the uncertainty map H_i into rectangular patches of fixed dimension and aggregate the entropy of all pixels in a patch to obtain one uncertainty score H_i^m for each patch m . We then select the region among all the unlabeled dataset that has the highest uncertainty score. This step is repeated for N_I times to obtain N_I informative samples. The selected regions are removed from the unlabeled dataset and processed further in the second step of the query strategy.

3.4. Representative Sampling

In the second step of our active learning query strategy, we select a representative subset from the pre-selected informative samples. To do so we obtain feature vectors from the pre-computed image patches. More formally, let $I_{i,m}$ be the m^{th} region of i^{th} image, then the corresponding feature vector

$$F_{i,m} = f_\theta(I_{i,m}) \quad (2)$$

is obtained using the trained backbone of the segmentation model f_θ . We then feed these features through the pre-computed UMAP Φ to transform them into the low-dimensional embedding space. We denote these low-dimensional embeddings by

$$G_{i,m} = \Phi(F_{i,m}). \quad (3)$$

Finally, we use the K-Means cluster centroid initialization algorithm (K-Means++ [1]) to obtain cluster centroids that are used to augment the training dataset. The detailed pseudocode of MEAL is outlined in Algorithm 1.

Algorithm 1: MEAL: Manifold Embedding-based Active Learning pseudocode

Initialize: Unlabeled set D_U , Acquisition steps K , Number of query patches N , Number of Informative patches N_I , Initial train set D_L , Unlabeled set D_U , Query set $D_Q = \emptyset$

Step 0: Train a UMAP Φ using feature vectors F computed from a segmentation model f_θ with pre-trained ImageNet weights on (D_L, D_U)

for $k \leftarrow 1$ **to** K **do**

Train a model f_θ on D_L

Query Procedure:

Step 1 : Compute pixel-wise Entropy $H_i^{(u,v)}$ of image I_i in D_U using trained f_θ

$$H_i^{(u,v)} = - \sum_{c=1}^C P_c^{(u,v)} \log P_c^{(u,v)}$$

Select N_I image patches with highest entropy

H_i

$I = \emptyset; m = 1$

repeat

$I_{i,m} \leftarrow \max(H)$

$I \leftarrow I \cup I_{i,m}$

$H \leftarrow H \setminus I_{i,m}$

until $m = N_I; m \leftarrow m + 1$

Step 2 : Compute image patch feature vectors

$F_{i,m}$ of all $I_{i,m}$ in I starting with $G = \emptyset$ as

$$F_{i,m} = f_\theta(I_{i,m})$$

Compute embeddings of the feature vector as

$$G_{i,m} \leftarrow \Phi(F_{i,m})$$

$$G \leftarrow G \cup G_{i,m}$$

Run K-Means++ on all embeddings G and select

N cluster centroids to form the query set D_Q

Update: $D_U \leftarrow D_U - D_Q, D_L \leftarrow D_L \cup D_Q$

return Final training set D_L

3.5. MEAL with Fit Transform

To interpolate between the full MEAL method and the Entropy baseline, we introduce MEAL with Fit Transform (MEAL-FT). In this variant we do not train a UMAP [25] using the entire dataset to learn a transformation Φ . Instead the transformation is learned on the informative samples only during each acquisition step and then it is used to transform those samples into low-dimensional embeddings. This modification intends to evaluate if the additional step of incorporating the global context is required. While the original method MEAL carries information of entire dataset in the form of Φ , MEAL-FT does not have this additional

information.

4. Experiments

We conducted several experiments to evaluate the performance of our method against some active learning strategies from the literature. For comparison, we use *mean Intersection over Union (mIoU)* to measure the performance of the segmentation model, trained on dataset assembled through the active learning strategies. IoU is the ratio of area of overlap between predicted segmentation mask and ground truth label and the area of union between predicted segmentation mask and ground truth label. mIoU is computed by averaging the IoUs of all the classes. We compared our method to following methods from the literature:

1. **Random Sampling (Random):** Naïve method of randomly sampling N samples in each acquisition step.
2. **Entropy Sampling (Entropy):** An uncertainty sampling strategy capturing the informativeness of samples by computing the entropy [31] as an uncertainty measure.
3. **Fully Supervised [9, 15]:** Traditional supervised method to train the segmentation model with the full labeled dataset.

We also tried comparing to EquAL [14], but found the method to perform poorly in our setting and we did not invest time in its performance tuning. To have a fair comparison with our method, we modify the above methods (except fully supervised) to query and acquire image patches rather than the complete image itself.

4.1. Datasets

We use the publicly available semantic segmentation datasets CamVid [7] and Cityscapes [10] for our experiments. During the acquisition process, instead of obtaining real annotations from an oracle, we mask the already available labels of the query dataset and reveal them once the AL algorithm selects those patches.

CamVid [7] is a dataset which is composed of street scene view images of size 360×480 with 11 semantic classes. It contains 367 (training), 101 (validation), 234 (test) images with labels. We divide each image into 16 patches of dimension 90×120 .

Cityscapes [10], a much larger dataset is also composed of real-world street view images of dimension 2048×1024 with 19 semantic classes. The training set consists of 2975 images with semantic labels and the validation set contains 500 images. Also here, we divide each image into 16 smaller patches. Hence the dimension of each patch is 128×256 since the width and height of images are bigger than CamVid dataset.

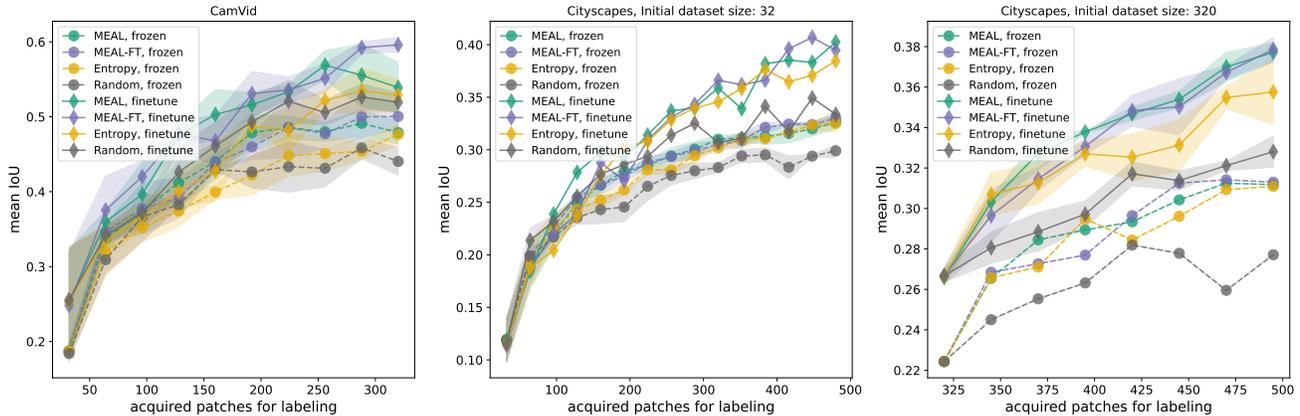


Figure 3: Performance of MEAL in comparison to other baselines for CamVid (left) and Cityscapes with small initial dataset size (center: 32) and bigger initial dataset size (right: 320). We plot the mean IoU against the number of image patches acquired in each acquisition step. For visualization reasons we do not show the fully supervised baselines, but report them as 0.73 (0.635) for Camvid with (without) backbone finetuning and 0.618 (0.497) for Cityscapes with (without) backbone finetuning. (Color best viewed online)

4.2. Implementation Details

For the segmentation model, we used DeepLabV3+ [9] with a MobileNetV2 [15] backbone pre-trained on ImageNet dataset [28] and implemented using PyTorch [27]. The UMAP implementation is based on [26] and we use the features extracted from the MobileNetV2 [15] backbone as training inputs. We train the segmentation model using cross-entropy loss and the Adam optimizer [17] with the PyTorch default values. After each acquisition step we train the model for 25 epochs and for 32 epochs after the final acquisition step, and chose the best model among them. Note that we retrain the model from scratch after each acquisition step. In all the experiments, the batch size is set to 8, the learning rate is set to 0.001 and no learning rate schedule is used. Moreover we do not use any weight decay and no data augmentation techniques were employed. We repeated each experiment three times with different random choices for the initial labeled dataset D_L . At the end, we reported the average value of mean IoU score on the validation dataset along with the standard deviation.

The fully supervised baselines were trained with the full training data of CamVid and Cityscapes, respectively. We trained the model for 60 epochs with a batchsize of two; note that we are here using the full images and not just the patches. Moreover we used a learning rate of 0.0008 and used a learning rate decay of 0.1 every 20 epochs. We report the mIoU score on the validation dataset for both setting freezing and finetuning the MobileNetV2 backbone.

4.3. Experimental Results

Based on some initial hyperparameter searches, we settle on using two initial randomly sampled images and split each into 16 patches. The total amount of 32 labelled patches form our initial dataset. At each acquisition step, we (i) split the input images into 16 patches and (ii) query 32 patches from the unlabeled set based on the active learning strategy at hand. In MEAL (MEAL-FT) this acquisition is based on first extracting the 200 most informative patches, before extracting the final 32 patches using the clustering of UMAP embeddings. We do this in order to guarantee diversity in the acquired samples at the end of each acquisition step. These final queried samples are added to the training set along with the corresponding labels and a new model is trained on this updated training set. We also experiment with a larger initial sample of 20 images and 16 patches each in the case of Cityscapes. In this case, we modify the query sample size to 25 patches. Moreover we run ablations with finetuning and freezing the MobileNetV2 backbone. All results are summarized in Table 1.

In Figure 3 (left), we plot the average mean IoU on the CamVid validation set with respect to the effective training dataset size after new patches have been acquired after each acquisition step. At the end of all acquisition steps, the total number of patches acquired correspond to approximately 5% of the total training dataset size. We see that finetuning the backbone gives a marginal lift given the increased numbers of parameters that need to be trained. However in both cases, MEAL and MEAL-FT outperform the baselines given by Entropy Sampling and Random Sampling. We refer to Table 1 for quantitative results.

	finetune	# params	Init. bs	qs	perc.	Random	Entropy	MEAL	MEAL-FT	Supervised
CamVid	no	4.0e6	32	32	5%	0.509	0.516	0.558	0.573	0.635
CamVid	yes	5.8e6	32	32	5%	0.52	0.528	0.54	0.596	0.73
Cityscapes	no	4.0e6	32	32	1%	0.298	0.325	0.326	0.330	0.497
Cityscapes	yes	5.8e6	320	25	1%	0.328	0.357	0.377	0.378	0.618
Cityscapes	no	4.0e6	320	25	1%	0.277	0.310	0.312	0.313	0.497
Cityscapes	yes	5.8e6	32	32	1%	0.333	0.384	0.402	0.395	0.618

Table 1: Summary of mean IoU on the validation sets of CamVid and Cityscapes after the final acquisition step. *Perc.* shows the fraction of the acquired dataset compared to the full dataset size, *Init. bs* is the initial dataset size for the seed labeled dataset, *qs* is the query size and *# params* gives the number of trainable parameters. Best method is highlighted in bold.

The results are slightly different on Cityscapes as shown in Figure 3 (center, right). Due to time constraints we did not repeat the experiments over several random seeds, but limited ourselves to one seed to show trendlines. Here the lift obtained by additionally finetuning the MobileNetV2 backbone is significant. This is true for small and large initial batch sizes. In the case of large initial dataset size, we can also select fewer samples per acquisition step and quickly outperform the baseline strategies with fewer initial samples. This indicates that the initial gradients of the active learner are important and strongly influence later acquisition steps. Moreover, we find that MEAL and MEAL-FT do not outperform the Entropy Sampling strategy, except in the case of backbone finetuning and large initial dataset size. We hypothesize that this could be due to the initial hyperparameters, as we transferred the CamVid settings to Cityscapes. This indicates the importance of well chosen hyperparameters in the case of a frozen backbone, but also shows that finetuning the backbone renders the method less sensitive to the choice of hyperparameters. For quantitative results we refer to Table 1.

5. Discussion & Future Work

In our work, we conceptualized a new active learning query strategy *MEAL*, within an exploration-exploitation framework for image segmentation. The method combines informativeness measures, based on entropy sampling, with representative measures, based on clustering methods. Our method extracts low-dimensional manifold embeddings from feature vectors of unlabeled data points obtained using UMAP which maintains the structural similarity between the data points in lower dimensions. This allows us to efficiently cluster the datapoints using a K-Means++ approach. We conducted experiments on semantic segmentation datasets of different sizes and analyzed the performance of our algorithm. The experiments on CamVid and Cityscapes dataset outperformed existing methods that were adopted to fit to our image segmentation task and validate the correctness of this approach. Moreover, we showed

that finetuning the backbone of the segmentation network improved the method but is not uniform over datasets, which is likely due to the usefulness of the representation for the dataset at hand. We note that while MEAL-FT was designed as an intermediate ablation step between MEAL and the Entropy Sampling method we do not see the expected interpolating behavior. This could be due to a too large sample size for selecting the informative subsample of patches. In this case, the large subsample already contains enough representative power of the full dataset and hence we would not see a loss in global context.

There are various directions that can be investigated for further improvement of the approach which we leave for future work. As the current experiments did not perform exhaustive hyperparameter tuning it would be worthwhile to understand the dependence of the algorithm on factors such as the initial batch size, the query size as well as the UMAP embedding dimension. Moreover, the use of UMAP as a dimensionality reduction technique is a sound choice, but certainly not the only one and a more exhaustive comparison to existing active learning approaches for image segmentation should be done. Finally, as seen in the work by [33], prediction calibration that provides reliable gradient embeddings needed for querying, could be incorporated in the uncertainty sampling step of our method to obtain calibrated Softmax values.

6. Author contributions

D. Sreenivasaiah and T. Wollmann contributed to the design of the method and research. D. Sreenivasaiah implemented the approach. D. Sreenivasaiah and J. Otterbach performed the analysis. T. Wollmann supervised the work. All authors discussed the results and wrote the manuscript.

Acknowledgements

The research leading to these results is funded by the Federal Ministry for Economic Affairs and Energy (BMWi) within the project "AI Safeguarding – Methods and measures for safeguarding AI-based perception functions for

automated driving” (#19A19005Y). The authors would like to thank the consortium for the successful cooperation.

References

- [1] David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA21)*, 2007. 2, 3, 5
- [2] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *Proc. International Conference on Learning Representations (ICLR)*, 2020. 2, 3
- [3] W. H. Beluch, T. Genewein, A. Nurnberger, and J. M. Kohler. The power of ensembles for active learning in image classification. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9368–9377, 2018. 2
- [4] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, Berlin, Heidelberg, 2006. 4
- [5] Zalán Bodó, Zsolt Minier, and Lehel Csató. Active learning with clustering. In Isabelle Guyon, Gavin Cawley, Gideon Dror, Vincent Lemaire, and Alexander Statnikov, editors, *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS) - Active Learning and Experimental Design workshop*, pages 127–139, 2011. 2
- [6] Alexis Bondu, Vincent Lemaire, and M. Boullé. Exploration vs. exploitation in active learning : A bayesian approach. In *Proc. International Joint Conference on Neural Networks (IJCNN)*, pages 1 – 7, 2010. 2, 3
- [7] Gabriel J. Brostow, Jamie Shotton, Julien Fauqueur, and Roberto Cipolla. Segmentation and recognition using structure from motion point clouds. In *Proc. European Conference on Computer Vision (ECCV)*, pages 44–57, 2008. 2, 5
- [8] Arantxa Casanova, Pedro O. Pinheiro, Negar Rostamzadeh, and Christopher J. Pal. Reinforced active learning for image segmentation. *CoRR*, abs/2002.06583, 2020. 3
- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. *CoRR*, abs/1802.02611, 2018. 2, 5, 6
- [10] 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 *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5
- [11] Aron Culotta and Andrew McCallum. Reducing labeling effort for structured prediction tasks. In *Proc. Conference on Artificial Intelligence (AAAI)*, page 746–751. AAAI Press, 2005. 2
- [12] Sanjoy Dasgupta. Two faces of active learning. *Theor. Comput. Sci.*, 412:1767–1781, 2011. 2
- [13] Y Gal, R Islam, and Z Ghahramani. Deep bayesian active learning with image data. In *Proc. International Conference on Machine Learning (ICML)*, volume 3, pages 1923–1932, 2017. 1, 2
- [14] S. Alireza Golestaneh and Kris M. Kitani. Importance of self-consistency in active learning for semantic segmentation. *CoRR*, abs/2008.01860, 2020. 1, 2, 3, 5
- [15] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. 2, 5, 6
- [16] Xin Jin and Jiawei Han. K-means clustering. In Claude Sammut and Geoffrey I. Webb, editors, *Encyclopedia of Machine Learning*, pages 563–564. Springer US, 2010. 2, 3
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *Proc. International Conference on Learning Representations (ICLR)*, 2015. 6
- [18] Adriana Kovashka, Olga Russakovsky, Li Fei-Fei, and Kristen Grauman. Crowdsourcing in computer vision. *CoRR*, abs/1611.02145, 2016. 1
- [19] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1
- [20] David D. Lewis and William A. Gale. A sequential algorithm for training text classifiers. *CoRR*, abs/cmp-lg/9407020, 1994. 2, 4
- [21] Changsheng Li, Handong Ma, Zhao Kang, Ye Yuan, Xiaoyu Zhang, and Guoren Wang. On deep unsupervised active learning. In Christian Bessiere, editor, *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2626–2632. International Joint Conferences on Artificial Intelligence Organization, 2020. 2
- [22] Xiaolong Liu, Zhidong Deng, and Yuhuan Yang. Recent progress in semantic image segmentation. *Artificial Intelligence Review*, 52:1089–1106, 2019. 1
- [23] Radek Mackowiak, Philip Lenz, Omair Ghori, Ferran Diego, Oliver Lange, and Carsten Rother. CEREALS - cost-effective region-based active learning for semantic segmentation. *CoRR*, abs/1810.09726, 2018. 1, 3
- [24] Martin Maška, Vladimír Ulman, David Svoboda, Pavel Matula, Cristina Ederra, Ainhoa Urbiola, T. Espaa, Subramanian Venkatesan, Deepak Balak, Pavel Karas, T. Bolckov, S. Étreitov, Craig Carthel, Stefano Coraluppi, Nathalie Harder, Karl Rohr, Klas Magnusson, J. Jaldn, Helen Blau, and Carlos Ortiz-de Solorzano. A benchmark for comparison of cell tracking algorithms. *Bioinformatics*, 30:1609–1617, 2014. 1
- [25] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *CoRR*, abs/1810.03426, 2020. 2, 3, 4, 5
- [26] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. Umap: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3:861, 2018. 6
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H.

- Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 8024–8035. Curran Associates, Inc., 2019. [6](#)
- [28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *Proc. International Journal of Computer Vision (IJCV)*, 115:211–252, 2015. [2](#), [6](#)
- [29] O. Sener and S. Savarese. Active learning for convolutional neural networks: A core-set approach. *CoRR*, abs/1708.00489, 2018. [2](#)
- [30] B. Settles. *Active Learning*. Morgan & Claypool Publishers, 2012. [1](#), [2](#)
- [31] C. E. Shannon. A mathematical theory of communication. *Mobile Computing and Communications Review*, page 3–55, 2001. [5](#)
- [32] Yawar Siddiqui, Julien Valentin, and Matthias Nießner. Viewal: Active learning with viewpoint entropy for semantic segmentation. *CoRR*, abs/1911.11789, 2019. [3](#)
- [33] Bindya Venkatesh and Jayaraman J. Thiagarajan. Ask-n-learn: Active learning via reliable gradient representations for image classification. *CoRR*, abs/2009.14448, 2020. [2](#), [3](#), [7](#)
- [34] W. Wang, Y. Fu, Z. Pan, X. Li, and Y. Zhuang. Real-time driving scene semantic segmentation. *IEEE Access*, 8:36776–36788, 2020. [1](#)
- [35] Lin Yang, Yizhe Zhang, Jianxu Chen, Siyuan Zhang, and Danny Z. Chen. Suggestive annotation: A deep active learning framework for biomedical image segmentation. *CoRR*, abs/1706.04737, 2017. [1](#), [2](#), [3](#)