

MBZUAI

Digital.Commons@MBZUAI

Computer Vision Faculty Publications

Scholarly Works

3-30-2021

Deep Gaussian Processes for Few-shot Segmentation

Joakim Johnander

Linköping University & Zenseact Ab

Johan Edstedt

Linköping University

Martin Danelljan

ETH Zurich

Michael Felsberg

Linköping University

Fahad Shahbaz Khan

Linköping University & Mohamed bin Zayed University of Artificial Intelligence

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/cvfp>



Part of the [Artificial Intelligence and Robotics Commons](#)

"Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC BY-NC-SA 4.0

Uploaded 24 March 2022"

Recommended Citation

J. Johnander, J. Edstedt, M. Danelljan, M. Felsberg, and F.S. Khan, "Deep Gaussian processes for few-shot segmentation", 2021, arXiv:2103.16549

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Computer Vision Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Deep Gaussian Processes for Few-Shot Segmentation

Joakim Johnander^{1,2} Johan Edstedt¹ Martin Danelljan³ Michael Felsberg¹ Fahad Shahbaz Khan^{1,4}
¹Computer Vision Laboratory ²Zenseact AB ³Computer Vision Lab ⁴Mohamed bin Zayed University of AI
Dept. of Electrical Engineering Sweden ETH Zurich, Switzerland UAE
Linköping University, Sweden

{joakim.johnander, johan.edstedt, michael.felsberg}@liu.se
martin.danelljan@vision.ee.ethz.ch fahad.khan@mbzuai.ac.ae

Abstract

Few-shot segmentation is a challenging task, requiring the extraction of a generalizable representation from only a few annotated samples, in order to segment novel query images. A common approach is to model each class with a single prototype. While conceptually simple, these methods suffer when the target appearance distribution is multi-modal or not linearly separable in feature space.

To tackle this issue, we propose a few-shot learner formulation based on Gaussian process (GP) regression. Through the expressivity of the GP, our approach is capable of modeling complex appearance distributions in the deep feature space. The GP provides a principled way of capturing uncertainty, which serves as another powerful cue for the final segmentation, obtained by a CNN decoder. We further exploit the end-to-end learning capabilities of our approach to learn the output space of the GP learner, ensuring a richer encoding of the segmentation mask. We perform comprehensive experimental analysis of our few-shot learner formulation. Our approach sets a new state-of-the-art for 5-shot segmentation, with mIoU scores of 68.1 and 49.8 on PASCAL-5ⁱ and COCO-20ⁱ, respectively.

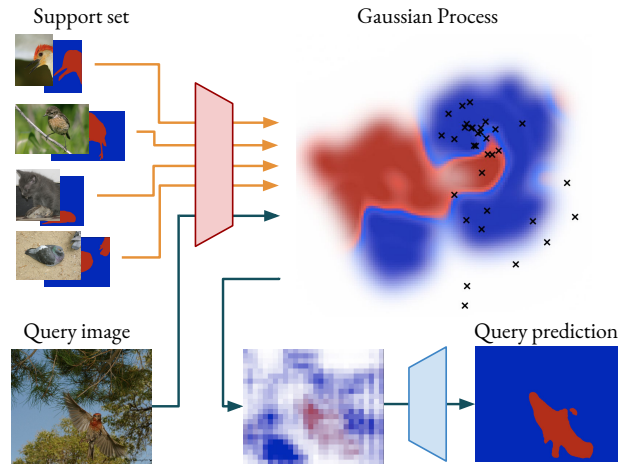


Figure 1. We use Gaussian processes to model the support set. The learnt GP is visualized, with blue areas corresponding to background and red areas to foreground. White areas and areas with low contrast correspond to highly uncertain regions, given by the high variance predicted by the GP. Our approach leverages the flexible classification boundaries and uncertainties predicted by our GP few-shot learner.

1. Introduction

Few-shot segmentation (FSS) [24] has received increased attention in recent years. The aim is to segment a query image given the support set, containing only a few annotated training samples of the class. Methods for FSS therefore need to extract information from the support set in order to segment the query image. A popular paradigm adopted by several recent works [20, 38, 39, 17] is *meta-learning*. These works train a neural network directly for the FSS-task, where the network contains inductive priors in the form of purpose-specific few-shot learner modules.

Many recently proposed few-shot learner modules [13,

27, 32, 39] extract a single prototype vector from the support set and compare it to feature vectors extracted from the query image via a dot-product. In essence, this leads to a linear classifier and thus relies on the feature space being linearly separable over the classes. However, this is not necessarily the case. In fact, recent work [1] show that even when trained for such linear separability, the resulting feature space for *unseen* classes is in many cases not linearly separable. A few works [15, 34] argue that even a single object may require different prototypes for different parts. In light of this, an attractive alternative would be to create a few-shot learning module that is more flexible, and able to model complex decision boundaries. However care must be taken to ensure that such a module does not overfit to

the support set, and additionally it must be computationally tractable and fully differentiable.

We explore using Gaussian process regression to model the support set and make predictions on the query set. In each episode, the support set is fed through an encoder to provide a set of features and corresponding values from the mask. Based on these, a GP is constructed. The GP probabilistically models the mapping between the features and the mask values and is flexible enough to act as an interpolator under certain assumptions on the noise [16]. During inference, the GP acts on query features and infers a probability distribution, the *predictive posterior*, over the mask value for each feature. All calculations made during learning and inference of the GP are differentiable and rely on operations readily integrated as a module in a neural network.

Combining GPs with deep neural networks is appealing for several reasons. The neural network provides the GP with powerful visual features. In turn, the GP predictive posterior distribution gives a principled way of combining potentially highly correlated samples. Furthermore, the covariance in the predictive posterior provides additional uncertainty information. This information may be especially important in few-shot learning. As we are given only a few training examples, we are inevitably queried with unseen appearances. Whenever the GP is queried with unseen features, it will convey this information in the form of high covariance estimates.

Since the Gaussian process is not limited to a one-dimensional output space, we go beyond predicting only a single mask value for each query feature. We encode the given support masks with a lightweight neural network. The resulting mask encoding is a multi-dimensional feature vector containing a richer representation of the mask. During inference, our Gaussian process few-shot learner makes a prediction on this output space. The predicted mask encodings, together with the uncertainty estimates, are fed to a subsequent decoder. This decoder is a neural network that combines these cues to produce a final segmentation.

Contributions: Our main contributions are as follows.

- i) We propose a few-shot segmentation learner based on Gaussian processes.
- ii) We propose to learn the output space of the GP via a neural network that encodes the mask.
- iii) We empirically show the benefits of the uncertainty information provided by the GP.
- iv) We analyze how the approach scales with larger support sets on COCO-20¹. We demonstrate its effectiveness on the PASCAL-5ⁱ and COCO-20ⁱ benchmarks, outperforming prior works in the 5-shot setting.

2. Related Work

Prototype Based Few-Shot Segmentation: Prototype-based methods aim to learn a *single* representation of a semantic class based on the support set. Shaban *et al.* [24] proposed a method of predicting the weights of a linear classifier based on the support set, which was further built upon in later works [25, 13]. Boudiaf *et al.* [4] instead perform transductive inference at test-time to predict their classifier and achieve impressive results. Instead of learning the classifier directly, Rakelly *et al.* [20] proposed to construct a global conditioning prototype from the support set and concatenate it to the query representation. This method proved successful, with a large number of subsequent works [6, 39, 32, 17, 38, 15, 2, 14]. Wang *et al.* [30] additionally introduced a probabilistic perspective and model their prototype as a single Gaussian, which they can sample from at inference-time. A fundamental limitation of the prototype based methods is the *unimodality* assumption. Some works [34, 15] try to circumvent this issue by clustering multiple prototypes to create a richer representation. However, clustering introduces extra hyperparameters as well as optimization difficulties. In contrast, our method is not restricted in any such sense, and only requires us to choose an appropriate kernel, which we will show is not a difficult task.

Pointwise Few-Shot Segmentation: Closer to our method are pointwise approaches. Instead of assuming a global representation of the target class, they aim to find correspondences between the support and query set. Previous work has mostly focused on attention mechanisms [37, 35, 10]. Tian *et al.* [28] uses a mechanism similar to attention to generate a prior mask and then enrich it based on query features and a prototype. While the pointwise approach is much more flexible, we are left with the task of choosing how to optimally weigh the connections between the support and query set. Wang *et al.* [31] approached this problem by additionally re-weighting the attention weights before the softmax operation to ensure a diverse support set. However, their approach is heuristic in nature and it is unclear in their framework how to optimally re-weight the connections. In contrast to these works, our method provides a principled way of finding these correspondences, with the additional benefit of providing an uncertainty estimate which we will show greatly improves performance.

Combining GPs and Neural Networks: The combination of GPs and neural networks is not an entirely new idea. While earlier work focused on combining GPs and neural networks in the standard supervised setting [23, 33, 5], there has recently been an increased interest in applying Gaussian processes in the context of *meta-learning* [29] and *few-shot classification* [19, 26]. All these works employ the GP as a final *output layer* in the network, and optimize either the predictive or marginal log likelihood directly. We go be-

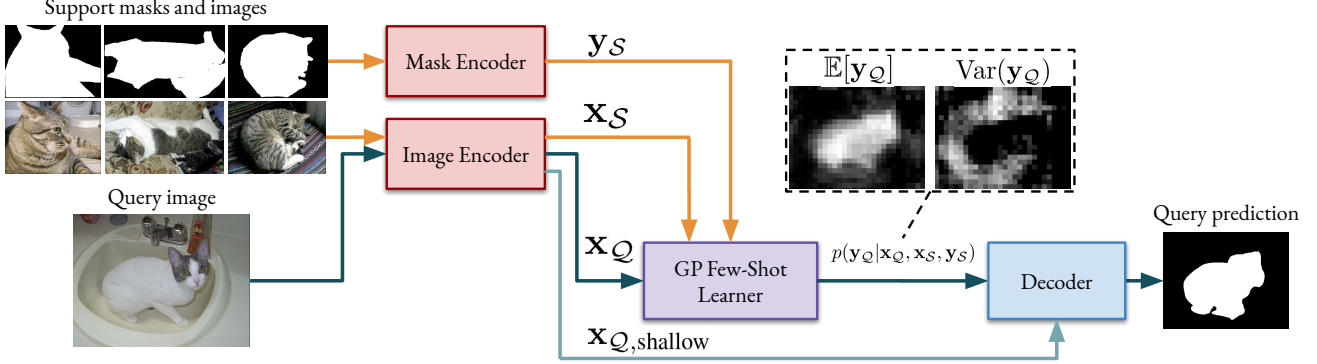


Figure 2. Overview of our approach. Support and query images are fed through an encoder to produce deep features \mathbf{x}_S and \mathbf{x}_Q respectively. The support masks are fed through another encoder to produce \mathbf{y}_S . Using Gaussian process regression, we infer the probability distribution of the query mask encodings \mathbf{y}_Q given the support set and the query features (see equations (3)-(5)). We create a representation of this distribution and feed it through a decoder. The decoder then predicts a segmentation at the original resolution.

yond this limitation and propose an internal GP model of the support features, where the posterior predictive distribution is used as input to a CNN decoder. Moreover, this allows us to learn the output space of the GP to further increase its expressive power.

3. Method

We propose an approach for few-shot segmentation where the modelling of the support set is based on Gaussian process regression. We briefly cover the few-shot segmentation problem and introduce the various components in Section 3.1. We then discuss Gaussian process regression in Section 3.2, followed by how we apply it to our task in Section 3.3. In Section 3.4 we extend the approach to encode the given support masks and in Section 3.5 we go through architectural details.

3.1. Overview

In few-shot segmentation, the aim is to learn a semantic segmentation method from a small collection of annotated images, called support set. For evaluation and training, we consider *episodes*, each comprising a labelled support set and a collection of query images on which the segmentation itself is evaluated. This poses a major challenge compared to, for instance, semantic segmentation, since the model needs to extract information from the support set and compare it to the query image. Furthermore, the training and test data contain different semantic classes. This prohibits the model from *cheating* and enforces learning from the support set.

An overview of our approach is given in Figure 2. First, the support and query images are fed through an image encoder to produce two sets of deep features, namely the support set $\{\mathbf{x}_S^i\}_i$ and the query set $\{\mathbf{x}_Q^j\}_j$. Each feature $\mathbf{x}^i \in \mathbb{R}^D$ corresponds to a spatial position in one of the im-

ages. We stack the support and query features as rows in a matrix to produce \mathbf{x}_S and \mathbf{x}_Q respectively. The support masks are fed through another encoder to produce mask encodings $\{\mathbf{y}_S^i\}_i$ — one per support feature. These are stacked into a vector \mathbf{y}_Q to match the rows of \mathbf{x}_Q . Now, the aim is to predict the mask encoding \mathbf{y}_Q^j corresponding to each query feature \mathbf{x}_Q^j . The idea in this work is to use Gaussian processes to infer the conditional probability distribution $p(\mathbf{y}_Q|\mathbf{x}_Q, \mathbf{x}_S, \mathbf{y}_S)$, *i.e.* the probability distribution of the mask encodings given the support set. A representation of this distribution is then fed to a decoder. The decoder combines the representation with shallow features extracted from the query image to predict a final segmentation.

3.2. Gaussian Processes

Gaussian process regression is a method for supervised learning that, as we shall see, has several appealing properties for few-shot segmentation. We seek a function f that relates the input to the output, *i.e.* $f(x) = y$, given the support examples $\{\mathbf{x}^i, \mathbf{y}^i\}_i$. Gaussian processes infer a probability distribution over such functions. Using this distribution, we can then infer the conditional distribution of y given previously unseen data x . In our case, the unseen data would be query features \mathbf{x}_Q^j and we would infer the distribution over the corresponding mask encoding \mathbf{y}_Q^j . The idea in Gaussian process regression is to assume that the output is jointly Gaussian with some mean μ and covariance Σ . The covariance $\text{Cov}(y, y')$ between two outputs y and y' is defined by a covariance function or *kernel* κ that is evaluated at the two corresponding inputs. That is, $\text{Cov}(y, y') = \kappa(x, x')$. The kernel is a positive definite function and selected such that if its input features are similar, the corresponding outputs are also expected to be similar.

Gaussian processes provide a principled model f with probabilistic output, containing both a point estimate as

well as uncertainty information. GPs are non-parametric in the input space and stores all given input-output pairs. As such, they have the flexibility to model multi-modal input data as well as as complex decision boundaries [16]. The wrigglyness of the decision boundaries, or inversely, their smoothness, is directly controlled with the choice of kernel function κ . The features are utilized only within κ , with linear time complexity in the feature dimensionality. This makes GPs suitable for making predictions based on deep features. Importantly, inferring the distributions over unseen data is an entirely differentiable operation. A few-shot learner based on Gaussian processes is therefore readily integrated as a module within a deep neural network.

Going back to our few-shot learning problem, define \mathbf{y}_S and \mathbf{y}_Q as the vectors of all support mask encodings \mathbf{y}_S and query mask encodings \mathbf{y}_Q respectively. The former is observed and our aim is to predict a distribution over the latter. Here we have assumed that the mask encodings are one-dimensional for notational convenience. The GP assumes a joint distribution over the support and query samples given by,

$$\begin{pmatrix} \mathbf{y}_S \\ \mathbf{y}_Q \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \boldsymbol{\mu}_S \\ \boldsymbol{\mu}_Q \end{pmatrix}, \begin{pmatrix} \mathbf{K}_{SS} & \mathbf{K}_{SQ} \\ \mathbf{K}_{SQ}^\top & \mathbf{K}_{QQ} \end{pmatrix} \right) . \quad (1)$$

As most often the case, the prior means $\boldsymbol{\mu}_S$ and $\boldsymbol{\mu}_Q$ are assumed to be zero. Instead, the covariances $\mathbf{K}_{SS} = \kappa(\mathbf{x}_S, \mathbf{x}_S)$, $\mathbf{K}_{QQ} = \kappa(\mathbf{x}_Q, \mathbf{x}_Q)$, and $\mathbf{K}_{SQ} = \kappa(\mathbf{x}_S, \mathbf{x}_Q)$ determine the behavior of the GP. These are defined through the *kernel* $\kappa : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$. In our experiments, we adopt the commonly used *squared exponential* (SE) kernel

$$\kappa(x^m, x^n) = \sigma_f^2 \exp\left(-\frac{1}{2\ell^2} \|x^m - x^n\|_2^2\right) , \quad (2)$$

with scale parameter σ_f and length parameter ℓ .

Next, we infer the posterior distribution of the query mask encodings \mathbf{y}_Q using the support mask encodings \mathbf{y}_S . That is, we calculate the parameters of the conditional

$$\mathbf{y}_Q | \mathbf{y}_S, \mathbf{x}_S, \mathbf{x}_Q \sim \mathcal{N}(\boldsymbol{\mu}_{Q|S}, \boldsymbol{\Sigma}_{Q|S}) . \quad (3)$$

We assume that the measurements \mathbf{y}_S have been obtained with some additive i.i.d. Gaussian noise with variance σ_y^2 . This corresponds to adding a scaled identity matrix $\sigma_y^2 \mathbf{I}$ to the support covariance matrix \mathbf{K}_{SS} . Then, the parameters of the posterior are

$$\boldsymbol{\mu}_{Q|S} = \mathbf{K}_{SQ}^\top (\mathbf{K}_{SS} + \sigma_y^2 \mathbf{I})^{-1} \mathbf{y}_S , \quad (4)$$

$$\boldsymbol{\Sigma}_{Q|S} = \mathbf{K}_{QQ} - \mathbf{K}_{SQ}^\top (\mathbf{K}_{SS} + \sigma_y^2 \mathbf{I})^{-1} \mathbf{K}_{SQ} . \quad (5)$$

For a derivation of these equations, we refer the reader to [16]. With equations (3)-(5) we have calculated the distribution of the query mask encodings given our support set.

3.3. Gaussian Processes for Meta Learning

Next, we integrate a GP meta learner in a few-shot segmentation framework. First, the K support images $\{I_S^k\}_k^K$ are fed through an image encoder $F^{\text{feat}} : \mathbb{R}^{3 \times H \times W} \rightarrow \mathbb{R}^{D \times \frac{H}{16} \times \frac{W}{16}}$ to produce deep feature maps X_S^k . The feature vectors in these deep feature maps are gathered to construct our support set \mathbf{x}_S . The corresponding annotations, or support masks, are fed through the mask encoder $F^{\text{mask}} : \{0, 1\}^{H \times W} \rightarrow \mathbb{R}^{E \times \frac{H}{16} \times \frac{W}{16}}$ to produce corresponding mask encodings. These are gathered and stacked to produce the support mask encoding vector \mathbf{y}_S . Note that, in contrast to several earlier works, we model both the background and the foreground. All features are utilized.

Using the same image encoder, we calculate the query features \mathbf{x}_Q . We obtain the posterior over the query mask encodings via (3)-(5) and therefore first compute the covariance matrix blocks \mathbf{K}_{SS} , \mathbf{K}_{SQ} , and \mathbf{K}_{QQ} . We can then find the mean $\boldsymbol{\mu}_{Q|S}$ and covariance $\boldsymbol{\Sigma}_{Q|S}$ via (4) and (5). This involves the multiplication of a matrix inverse with another matrix. In practice, we use Cholesky decomposition and solve the resulting triangular systems of linear equations.

The posterior (3) yields a joint probability distribution over the mask encodings. We are mostly interested in the marginal distribution of each component y_Q^j , which contains both the expected mask encoding and its uncertainty. We obtain this by taking the mean and the diagonal of the covariance, forming a representation of the distribution as

$$z^j = [(\boldsymbol{\mu}_{Q|S})_j, (\boldsymbol{\Sigma}_{Q|S})_{j,j}] . \quad (6)$$

Here, $[\cdot]$ denotes concatenation and j is used to index the mean and covariance. Note that, within the GP, the spatial order is irrelevant. The resulting query distribution is invariant over the order of the support set elements and it is equivariant over permutations of the query features. However, after constructing z^j , we arrange them to match the original feature map extracted from the image. The result is then fed into a decoder $F^{\text{dec}} : \mathcal{X}_{Q, \text{shallow}} \times \mathbb{R}^{(E+1) \times \frac{H}{16} \times \frac{W}{16}} \rightarrow [0, 1]^{H \times W}$ that, with the help of shallow features $\mathbf{x}_{Q, \text{shallow}} \in \mathcal{X}_{Q, \text{shallow}}$, produces a final segmentation.

3.4. Mask Encoder

Thus far, we have assumed the mask encodings to be 1-dimensional values, *i.e.* $E = 1$. In the best of scenarios, the output from the GP module would be a low-resolution mask. We believe that the decoder would benefit from additional guidance. This could for instance be information about local shape or edges, anything that might help produce a fine segmentation mask. To that end, we feed the input mask through a lightweight neural network that outputs a multi-dimensional mask-encoding at the same resolution as the feature maps. This idea is not new. Recent work on

Video Object Segmentation [3] shows the benefits of learning the output feature space, albeit with a different few-shot learning module.

Gaussian processes are able to model multi-dimensional output. For simplicity and computational reasons, we will assume the covariance function to be independent and isotropic over the output dimensions. This means that the covariance becomes a block diagonal matrix. In practice, when the GP is used to infer $p(\mathbf{y}_Q | \mathbf{y}_S, \mathbf{x}_S, \mathbf{x}_Q)$ we find the mean via a series of matrix-vector multiplications (4). If the mask encodings are multi-dimensional vectors we stack them as rows in a matrix. The columns represent different dimensions, and the resulting mean is a matrix of the same size. Essentially, (4) is calculated once per dimension of the mask encoding. The resulting representation \mathbf{z}^j then contains a multi-dimensional mean and a single value for the variance. Note that while having a general covariance in the multi-dimensional case is possible, it quickly becomes computationally intractable.

3.5. Architecture Details

Following previous works we adopt a ResNet-50 [9] backbone pre-trained on ImageNet [22] as our image encoder. We dilate its last residual module by a factor of 2 and remove the terminal average pooling to provide features with stride 16. In addition, we place a single convolutional projection layer that reduces the 2048-dimensional feature map down to 512 dimensions. As mask encoder we adopt the publicly available LWL-encoder [3]. For the Gaussian process few-shot learner, we downsample the support feature maps and annotation encodings with a factor of 2, yielding a total stride of 32 compared to the original images. We use $\sigma_y^2 = 0.01$, $\sigma_f^2 = 1$, and $\ell^2 = \sqrt{512}$. As our decoder we adopt the segmentation decoder DFN proposed in [36]. We do not add the border network of DFN and we set it to upsample with the shallow feature maps output by the first and second ResNet modules. We implement the approach in PyTorch [18] and code will be made available at publication.

4. Experiments

We conduct experiments on the FSS benchmarks PASCAL-5ⁱ [24] and COCO-20ⁱ [17]. First, we describe our experimental setup in detail. Second, we perform a thorough state-of-the-art comparison, including an experiment showing how well the GP learner scales with additional shots. This is followed by an ablation study containing both a quantitative and a qualitative comparison. Last, we analyze the choice of kernel function κ .

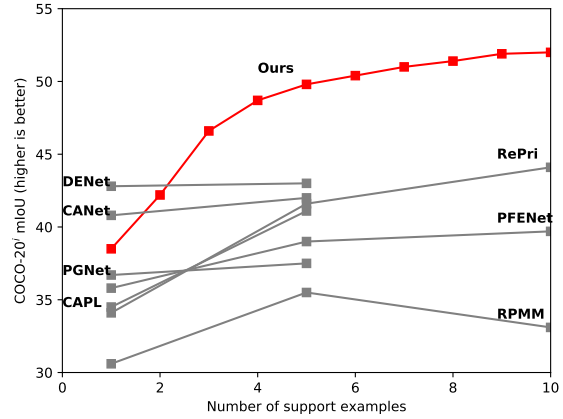


Figure 3. Performance of the proposed approach on COCO-20ⁱ for different number of shots, compared to the state-of-the-art (mIoU, higher is better). 10-shot numbers obtained from Boudiaf *et al.* [4]. As the number of shots increase from 1 to 10, the performance of the proposed approach monotonically increases. Already at 3 shots, the proposed approach outperforms the state-of-the-art in the 5-shot setting.

4.1. Experimental Setup

Datasets: We use the PASCAL-5ⁱ [24] and COCO-20ⁱ [17] benchmarks for our experiments. PASCAL-5ⁱ is composed of PASCAL VOC 2012 [7] with additional SBD [8] annotations. The dataset is split into 4 folds à 5 classes. COCO-20ⁱ [17] is built from MS-COCO [12] and is in a similar fashion split into 4 folds à 20 classes. Both benchmarks measure performance via 4-fold cross-validation. Three folds are used for training and the last fold held out for testing. Training is done on the respective training sets and testing on the respective validation sets. Thus, during the cross-validation, both the images *and* classes differ between the training and test sets.

Training Details: We train our models for 20000 and 40000 iterations for PASCAL-5ⁱ and COCO-20ⁱ respectively. We sample 8 episodes per batch and use the Adam [11] optimizer. We use a learning rate of 10^{-5} that is decayed with a factor of 0.3 halfway through training. During training, we freeze the batch normalization layers in our image encoder. We adopt the episode sampling and data augmentation used in [28]. First, all images are resized to fit in a predefined window. We use a window of size 448 for PASCAL and 512 for COCO. Then, the sample is zero-padded and augmented. The augmentation comprises scale jitter in the range $\{0.9, 1.1\}$, rotations in the range $\{-10, 10\}$, Gaussian blur, and random horizontal flips.

Evaluation: We evaluate our approach on each fold using 5000 and 20000 episodes respectively, for PASCAL-5ⁱ and COCO-20ⁱ. This follows the work of Tian *et al.* [28] in

Method	Backbone	Size	1-Shot					5-Shot				
			F-0	F-1	F-2	F-3	Mean	F-0	F-1	F-2	F-3	Mean
OSLSM [24]	VGG16	Orig	33.6	55.3	40.9	33.5	40.8	35.9	58.1	42.7	39.1	43.9
PANet[32]	VGG16	417	42.3	58.0	51.1	41.2	48.1	51.8	64.6	59.8	46.5	55.7
CANet [38]	RN50	Orig	52.5	65.9	51.3	51.9	55.4	55.5	67.8	51.9	53.2	57.1
RPMM [34]	RN50	?	55.2	66.9	52.6	50.7	56.3	56.3	67.3	54.5	51.0	57.3
PGNet ^{MS} [37]	RN50	Orig	56.0	66.9	50.6	50.4	56.0	57.7	68.7	52.9	54.6	58.5
CRNet[14]	RN50	Orig	-	-	-	-	55.7	-	-	-	-	58.8
FWB [17]	RN101	Orig	51.3	64.5	56.7	52.2	56.2	54.8	67.4	62.2	55.3	60.0
VPI [30]	RN101	?	53.4	65.6	57.3	52.9	57.3	55.8	67.5	62.6	55.7	60.4
DENet [13]	RN50	321	55.7	69.7	63.6	51.3	60.1	54.7	71.0	64.5	51.6	60.5
DAN ^{MS} [31]	RN101	?	54.7	68.6	57.8	51.6	58.2	57.9	69.0	60.1	54.9	60.5
LTM [35]	RN50	?	52.8	69.6	53.2	52.3	57.0	57.9	69.9	56.9	57.5	60.6
PFENet [28]	RN50	Orig	61.7	69.5	55.4	56.3	60.8	63.1	70.7	55.8	57.9	61.9
PPNet [†] [15]	RN50	417	47.8	58.8	53.8	45.6	51.5	58.4	67.8	64.9	56.7	62.0
CAPL [27]	RN101	Orig	-	-	-	-	57.4	-	-	-	-	63.4
RePri [‡] [4]	RN50	417	59.8	68.3	62.1	48.5	59.7	64.6	71.4	71.1	59.3	66.6
Ours	RN50	Orig	50.5	64.9	54.6	52.0	55.5	66.8	70.7	71.6	63.2	68.1

Table 1. Performance on PASCAL-5ⁱ compared to the state-of-the-art (mIoU, higher is better). *Size* is the resolution used to calculate mIoU at test-time. F-X is performance on fold X of the benchmark. ^{MS}Multi-scale inference [†]We report results without additional unlabeled data during test-time [‡]Transductive inference instead of meta-learning

Method	Backbone	Size	1-Shot					5-Shot				
			F-0	F-1	F-2	F-3	Mean	F-0	F-1	F-2	F-3	Mean
FWB [17]	RN101	Orig	17.0	18.0	21.0	28.9	21.2	19.1	21.5	23.9	30.1	23.7
VPI [30]	RN101	?	-	-	-	-	23.4	-	-	-	-	27.8
DAN ^{MS} [31]	RN101	?	-	-	-	-	24.4	-	-	-	-	29.6
RPMM [34]	RN50	?	29.5	36.8	28.9	27.0	30.6	33.8	42.0	33.0	33.3	35.5
PPNet [†] [15]	RN50	417	34.5	25.4	24.3	18.6	25.7	48.3	30.9	35.7	30.2	36.2
PFENet [28]	RN101	Orig	34.3	33.0	32.3	30.1	32.4	38.5	38.6	38.2	34.3	37.4
PGNet* [37]	RN50	321	39.5	39.7	33.9	33.5	36.7	42.4	38.9	32.4	36.5	37.5
CAPL [27]	RN101	Orig	-	-	-	-	34.5	-	-	-	-	41.1
RePri [‡] [4]	RN50	417	32.0	38.7	32.7	33.1	34.1	39.3	45.4	39.7	41.8	41.6
CANet* [38]	RN50	321	42.2	42.7	37.6	40.9	40.8	44.7	43.0	37.5	42.7	42.0
PFENet [28]	RN101	473	36.8	41.8	38.7	36.7	38.5	40.4	46.8	43.2	40.5	42.7
DENet [13]	RN50	321	42.9	45.8	42.2	40.2	42.8	45.4	44.9	41.6	40.3	43.0
Ours	RN50	Orig	37.1	40.6	38.6	37.7	38.5	48.4	53.4	49.5	48.0	49.8

Table 2. Performance on COCO-20ⁱ compared to the state-of-the-art (mIoU, higher is better). *Size* is the resolution used to calculate mIoU at test-time. F-X is performance on fold X of the benchmark. *Re-implementation by [13] ^{MS}Multi-scale inference [†]We report results without additional unlabeled data during test-time [‡]Transductive inference instead of meta-learning

which it is observed that the procedure employed by many prior works, using only 1000 episodes, yields fairly high variance. Performance is measured in terms of mean Intersection over Union (mIoU). First, the Intersection over Union (IoU) is calculated per class over all episodes in a fold. The mIoU is then found by averaging the IoU over the classes. Like the original work on FSS by Shaban *et al.* [24], we calculate the IoU on the original resolution.

4.2. State-of-the-Art Comparison

We compare the proposed approach to the state-of-the-art on PASCAL-5ⁱ and COCO-20ⁱ and show the results in Table 1 and Table 2 respectively. Following prior works, we report results given a single support example, *1-shot*, and given five support examples, *5-shot*. On PASCAL-5ⁱ, our approach obtains a respectable 1-shot performance of 55.5 mIoU. In the 5-shot setting, our approach outperforms the state-of-the-art with an mIoU of 68.1. This corresponds to an absolute mIoU gain of 1.5 compared to previous best method, RePri [4], and a 4.7 gain compared to the previous

best meta-learning method, CAPL [27].

On COCO-20ⁱ, our approach obtains a competitive 1-shot performance of 38.5 mIoU. In the 5-shot setting, the top performing approaches [27, 4, 38, 28, 13] obtain 41-43 mIoU. These include prototype based approaches [27, 38, 13], a point-based approach [28], and an approach not based on meta-learning [4]. Our approach obtains 49.8 mIoU, an absolute increase of 6.8 mIoU. This clearly demonstrates the effectiveness of the GP-based few-shot learner.

An important property for few-shot segmentation is to scale with more data. We therefore analyze the effectiveness of the proposed approach as the number of shots increases. The results are shown in Figure 3. Our approach is trained for 5 shots, except for the 1-shot result in which case we train for 1 shot. The performance of the proposed approach monotonically increases as the number of shots increases. At three shots, the GP-based few-shot learner significantly outperforms the state-of-the-art, and the performance keeps increasing all the way to 10 shots. While DENet [13] and CANet [38] provides high 1-shot performance, the performance increases only slightly when instead given 5 support examples. Both these methods rely on single prototypes. We believe that this representation struggles to model multiple support images. The GP, in contrast, is expressive enough to accurately model the given support set.

4.3. Ablation Study

Here, we analyze the impact of the key components in our proposed architecture. Quantitative results are shown in Table 3, while qualitative results are shown in Figure 4.

Mean: First, we alter equation (6) to utilize only the mean estimate of the GP few-shot learner. While this utilizes the predictive power of the GP, we do not gain the full benefit of the additional uncertainty information. This already performs competitively with an mIoU of 63.1. We refer to this version as our baseline.

Variance: Then, we add the variance as described by equation (6), leading to an impressive absolute gain of 3.3 mIoU. This shows the benefit of incorporating uncertainty information from our GP.

Covariance: We further experiment with incorporating more advanced representations of the predictive covariance. Here, for each z^j we collect the covariance of its spatial neighbours in a 5×5 window. This leads to a minor performance increase of 0.44.

Mask Encoder: Last, we add the mask encoder proposed in [3]. This enables our GP to make predictions on a richer output space. This addition gives us an absolute mIoU gain of 1.3. We therefore adopt this configuration as our final model.

Configuration	mIoU
Mean (Baseline)	63.1
Mean + Variance	66.4
Mean + Covariance	66.8
Mean + Covariance + Mask Encoder (Ours)	68.1

Table 3. Performance for different versions of our approach on the PASCAL-5ⁱ benchmark. Most notably, adding uncertainty information, in the form of variance or covariance, provides a significant gain in performance.

Method	mIoU
Linear	65.1
RQ	68.1
SE (Ours)	68.1

Table 4. Comparison between different covariance functions κ on the PASCAL-5ⁱ benchmark. The performance is stable between the squared exponential and rational quadratic kernels, while it significantly decreases for the linear kernel.

4.4. Choice of Kernel

We experiment with different covariance functions κ . In particular we compare the performance of the squared exponential (SE), rational quadratic (RQ), and linear (Lin) kernels. As discussed previously we use $\ell^2 = \sqrt{D}$ for all our experiments where applicable. We present our results in Table 4. Overall, the Gaussian process few-shot learner generalizes well to different kernel functions κ . This is perhaps not very surprising as the feature space on which the kernel act is learnt during training.

Squared Exponential: The squared exponential kernel is defined as

$$\kappa_{SE}(x, y) = \exp\left(-\frac{\|x - y\|_2^2}{2\ell^2}\right). \quad (7)$$

This is equivalent to the full method discussed in our ablation analysis and achieves 68.1 mIoU.

Rational Quadratic: The rational quadratic kernel is defined as

$$\kappa_{RQ}(x, y) = \left(1 + \frac{\|x - y\|_2^2}{2\alpha\ell^2}\right)^{-\alpha}. \quad (8)$$

In our experiments we use $\alpha = 1$. We found that it achieves similar performance to the squared exponential with an mIoU of 68.1.

Linear: The linear kernel is defined as

$$\kappa_{Lin}(x, y) = x^T y. \quad (9)$$

Using a linear kernel leads to a significant performance decrease of 3.0. The decrease in performance is in line with our previous reasoning since a Gaussian process with a linear kernel is equivalent to Bayesian linear regression [21].



5. Conclusion

We have proposed a few-shot learner based on Gaussian process regression for the few-shot segmentation task. The GP models the support set in deep feature space and its flexibility permits it to capture complex feature distributions. It makes probabilistic predictions on the query image, providing both a point estimate and additional uncertainty information. These predictions are fed into a CNN decoder that predicts the final segmentation. The resulting approach obtains state-of-the-art performance on PASCAL-5ⁱ and COCO-20ⁱ. It scales well with larger support sets during inference, even when trained for a fixed number of shots. With three or more shots, the approach outperforms the state-of-the-art on the COCO-20ⁱ benchmark.

Acknowledgement: This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation; ELLIIT; and the ETH Zürich Fund (OK).

References

- [1] Kelsey Allen, Evan Shelhamer, Hanul Shin, and Joshua Tenenbaum. Infinite mixture prototypes for few-shot learning. In *International Conference on Machine Learning*, pages 232–241. PMLR, 2019. 1
- [2] Reza Azad, Abdur R Fayjie, Claude Kauffmann, Ismail Ben Ayed, Marco Pedersoli, and Jose Dolz. On the texture bias for few-shot CNN segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2674–2683, 2021. 2
- [3] Goutam Bhat, Felix Järemo Lawin, Martin Danelljan, Andreas Robinson, Michael Felsberg, Luc Van Gool, and Radu Timofte. Learning what to learn for video object segmentation. In *European Conference on Computer Vision*, pages 777–794. Springer, 2020. 5, 7, 11
- [4] Malik Boudiaf, Hoel Kervadec, Ziko Imtiaz Masud, Ismail Ben Ayed, Jose Dolz, and Pablo Piantanida. Few-shot segmentation without meta-learning: A good transductive inference is all you need? *arXiv preprint arXiv:2012.06166*, 2020. 2, 5, 6, 7, 12
- [5] Roberto Calandra, Jan Peters, Carl Edward Rasmussen, and Marc Peter Deisenroth. Manifold Gaussian Processes for regression. In *Proceedings of the International Joint Conference on Neural Networks*, 2016. 2
- [6] Nanqing Dong and Eric P. Xing. Few-shot semantic segmentation with prototype learning. In *British Machine Vision Conference 2018, BMVC 2018*, 2019. 2
- [7] Mark Everingham, Luc Van Gool, Christopher K.I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 2010. 5
- [8] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE International Conference on Computer Vision*, 2011. 5
- [9] 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. 5, 11
- [10] Tao Hu, Pengwan Yang, Chiliang Zhang, Gang Yu, Yadong Mu, and Cees GM Snoek. Attention-based multi-context guiding for few-shot semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8441–8448, 2019. 2
- [11] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015. 5
- [12] 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. 5
- [13] Lihao Liu, Junyi Cao, Minqian Liu, Yong Guo, Qi Chen, and Minghui Tan. Dynamic Extension Nets for Few-shot Semantic Segmentation. pages 1441–1449. Association for Computing Machinery (ACM), 2020. 1, 2, 6, 7
- [14] Weide Liu, Chi Zhang, Guosheng Lin, and Fayao Liu. CR-Net: Cross-Reference Networks for Few-Shot Segmentation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 4164–4172, 2020. 2, 6
- [15] Yongfei Liu, Xiangyi Zhang, Songyang Zhang, and Xuming He. Part-Aware Prototype Network for Few-Shot Semantic Segmentation. In *European Conference on Computer Vision*, pages 142–158, 2020. 1, 2, 6
- [16] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012. 2, 4
- [17] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2019-Octob, pages 622–631, 2019. 1, 2, 5, 6
- [18] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 5
- [19] Massimiliano Patacchiola, Jack Turner, Elliot J. Crowley, and Amos Storkey. Bayesian meta-learning for the few-shot setting via deep kernels. In *Advances in Neural Information Processing Systems*, 2020. 2
- [20] Kate Rakelly, Evan Shelhamer, Trevor Darrell, Alexei Efros, and Sergey Levine. Conditional networks for few-shot semantic segmentation. In *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*, 2018. 1, 2
- [21] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. 2006. 7
- [22] 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. *IJCV*, pages 1–42, April 2015. 5

- [23] Ruslan Salakhutdinov and Geoffrey Hinton. Using deep belief nets to learn covariance kernels for Gaussian processes. In *Advances in Neural Information Processing Systems 20 - Proceedings of the 2007 Conference*, 2009. 2
- [24] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *British Machine Vision Conference 2017, BMVC 2017*, 2017. 1, 2, 5, 6
- [25] Mennatullah Siam, Boris Oreshkin, and Martin Jagersand. AMP: Adaptive masked proxies for few-shot segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2019-Octob, pages 5248–5257, 2019. 2
- [26] Jake Snell and Richard Zemel. Bayesian few-shot classification with one-vs-each pólya-gamma augmented gaussian processes. In *International Conference on Learning Representations*, 2021. 2
- [27] Zhuotao Tian, Xin Lai, L Jiang, Michelle Shu, Hengshuang Zhao, and Jiaya Jia. Generalized Few-Shot Semantic Segmentation. *arXiv preprint arXiv:2010.05210*, 2020. 1, 6, 7
- [28] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior Guided Feature Enrichment Network for few-shot segmentation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 2020. 2, 5, 6, 7, 11, 12
- [29] Prudencio Tossou, Basile Dura, Mario Marchand, François Laviolette, and Alexandre Lacoste. Adaptive deep kernel learning. *arXiv preprint arXiv:1905.12131*, 2019. 2
- [30] Haochen Wang, Yandan Yang, Xianbin Cao, Xiantong Zhen, Cees Snoek, and Ling Shao. Variational prototype inference for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 525–534, 2021. 2, 6
- [31] Haochen Wang, Xudong Zhang, Yutao Hu, Yandan Yang, Xianbin Cao, and Xiantong Zhen. Few-Shot Semantic Segmentation with Democratic Attention Networks. In *European Conference on Computer Vision*, pages 730–746, 2020. 2, 6
- [32] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. PANet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2019-Octob, pages 9196–9205, 2019. 1, 2, 6
- [33] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016*, 2016. 2
- [34] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype Mixture Models for Few-Shot Semantic Segmentation. In *European Conference on Computer Vision*, pages 763–778. Springer, 2020. 1, 2, 6, 12
- [35] Yuwei Yang, Fanman Meng, Hongliang Li, Qingbo Wu, Xiaolong Xu, and Shuai Chen. A new local transformation module for few-shot segmentation. In *International Conference on Multimedia Modeling*, pages 76–87. Springer, 2020. 2, 6
- [36] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Learning a Discriminative Feature Network for Semantic Segmentation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2018. 5, 11
- [37] Chi Zhang, Guosheng Lin, Fayao Liu, Jiushuang Guo, Qingyao Wu, and Rui Yao. Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2019-Octob, pages 9586–9594, 2019. 2, 6
- [38] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. CANET: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2, 6, 7
- [39] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas S Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE Transactions on Cybernetics*, 50(9):3855–3865, 2020. 1, 2

Appendix

In this appendix, we first provide additional details in Section A. Then, we supply additional results in Section B. Last, we show additional qualitative examples and visualizations of the Gaussian process in Section C.

A. Additional Implementation Details

We provide code for the Gaussian Process inference, the neural network layers that make up the modules used in our approach, and details on how we sample our episodes.

A.1. Code for Gaussian Process

Pseudo-code for the GP kernel is shown in Listing 1. For brevity and clarity, we omit device casting and simplify the solve implementation. In practice, we use the standard triangular solver in PyTorch, `torch.triangular_solve`.

A.2. Neural Network Architecture

We supplement the neural network layers utilized in our approach in table 5.

```
def GP(x_q, y_s, x_s, sigma_y, kernel):
    """ Produces the predictive
    posterior distribution of the GP

    Args:
    x_q: deep query features (B,Q,D)
    y_s: support mask features (B,S,M)
    x_s: deep support features (B,S,D)
    sigma_y: mask standard deviation
    kernel: the kernel function,
    """
    B, S, D = x_s.shape
    I = torch.eye(S)
    K_ss = kernel(x_s, x_s) # (B, S, S)
    K_qq = kernel(x_q, x_q) # (B, Q, Q)
    K_sq = kernel(x_s, x_q) # (B, Q, S)
    L_ss = torch.cholesky(K_ss +
        sigma_y**2*I) # (B, S, S)
    mu_q = K_qs @ solve(L_ss.T,
        solve(L_ss, y_s)) # (B, Q, M)
    v = solve(L_ss, K_sq) # (B, S, Q)
    cov_q = K_qq - v.T @ v # (B, Q, Q)
    return mu_q, cov_q
```

Listing 1: PyTorch implementation of the Gaussian Process utilized in the proposed approach. Here, the learning and inference is combined in a single step. The `@` operator denotes matrix multiplication and `solve` the solving of a linear system of equations. The `.T` is the batched matrix transpose.

Image Encoder		
conv1	Conv2d	$64 \times 224 \times 224$
bn1	BatchNorm2d	$64 \times 224 \times 224$
relu1	ReLU	$64 \times 224 \times 224$
maxpool	MaxPool2d	$64 \times 112 \times 112$
layer1	3x BottleNeck	$256 \times 112 \times 112$
layer2	4x BottleNeck	$512 \times 56 \times 56$
layer3	6x BottleNeck	$1024 \times 28 \times 28$
layer4	3x BottleNeck	$2048 \times 28 \times 28$
appearance	Conv2d	$512 \times 28 \times 28$
Mask Encoder		
conv1	Conv2d	$16 \times 224 \times 224$
bn1	BatchNorm2d	$16 \times 224 \times 224$
relu1	ReLU	$16 \times 224 \times 224$
maxpool	MaxPool2d	$16 \times 112 \times 112$
layer1	BasicBlock	$32 \times 56 \times 56$
layer2	BasicBlock	$64 \times 28 \times 28$
conv2	Conv2d	$64 \times 28 \times 28$
bn2	BatchNorm2d	$64 \times 28 \times 28$
relu2	ReLU	$64 \times 28 \times 28$
Decoder		
conv1	Conv2d	$256 \times 28 \times 28$
upsample1	Upsample	$256 \times 56 \times 56$
cab1	CAB	$256 \times 56 \times 56$
rrb1	RRB	$256 \times 56 \times 56$
upsample2	Upsample	$256 \times 112 \times 112$
cab2	CAB	$256 \times 112 \times 112$
rrb2	RRB	$2 \times 112 \times 112$
upsample3	Upsample	$2 \times 448 \times 448$

Table 5. All neural network blocks used by our approach. The rightmost column shows the dimensions of the output of each block, assuming a 448×448 input resolution. The image encoder is from He *et al.* [9]; the mask encoder from Bhat *et al.* [3]; and the decoder from Yu *et al.* [36]. The BottleNeck and BasicBlock blocks are from He *et al.* [9], and the CAB and RRB blocks from Yu *et al.* [36]. See their works for additional details.

A.3. Episode Sampling

The way we sample episodes during training and evaluation follows prior work [28]. During training, we select a query image at random from the dataset, with the condition that it contains one of the classes considered. During evaluation, we go through the dataset images in sequence to select a query image. Next, the class is selected at random from one of the classes contained in the query image. Last, a support set is constructed. We make sure that (i) the support set does not contain the query images; and (ii) that the support set does not contain the same image more than once.

Method	1-Shot					5-Shot				
	F-0	F-1	F-2	F-3	Mean	F-0	F-1	F-2	F-3	Mean
RPMM [34]	36.3	55.0	52.5	54.6	49.6	40.2	58.0	55.2	61.8	53.8
PFENet [28]	43.2	65.1	66.5	69.7	61.1	45.1	66.8	68.5	73.1	63.4
RePRI [4]	52.8	64.0	64.1	71.5	63.1	57.7	66.1	67.6	73.1	66.2
Ours	48.7	65.6	61.6	71.6	61.9	65.8	71.7	71.9	80.2	72.4

Table 6. The results of our approach in a COCO-20ⁱ to PASCAL transfer experiment (mIoU, higher is better). The approach is trained on a fold of COCO-20ⁱ training set and tested on the PASCAL validation set. The testing folds are constructed to include classes not present in the training set, and thus not the same as PASCAL-5ⁱ.

Fold-0	Fold-1	Fold-2	Fold-3
Airplane, Boat, Chair, Dining Table, Dog, Person	Bicycle, Bus, Horse, Sofa	Bird, Car, Potted Plant, Sheep, Train, TV-monitor	Bottle, Cat, Cow, Motorcycle

Table 7. The classes used for testing in the COCO-20ⁱ to PASCAL transfer experiment. This split is different from that of PASCAL-5ⁱ in order to avoid overlap between the training and testing classes.

B. Additional Results

We supply three additional results. First, we show the results of a domain transfer experiment. Next, we provide the performance over different support set sizes also on PASCAL-5ⁱ. Last, we report the runtimes of our approach.

B.1. COCO-20ⁱ to PASCAL Transfer

We show the effects of a domain transfer from COCO-20ⁱ to PASCAL in table 6. This experiment follows the domain transfer experiment supplemented by Boudiaf *et al.* [4]. First, our approach is trained on each of the four folds of COCO-20ⁱ. Next, we test each of the four versions on PASCAL, using only the classes held-out during training. We list the classes of each fold in table 7. Our approach obtains a performance of 61.9 mIoU in the 1-shot setting and 72.4 mIoU in the 5-shot setting. Note that these results are not directly comparable to our result on the PASCAL-5ⁱ as the folds are not the same. However, the proposed approach obtains competitive performance in the 1-shot setting and sets a new state-of-the-art in the 5-shot setting. This is a clear example where the proposed approach manages to transfer between different datasets.

B.2. Performance over Support Set Size

We report the performance of the proposed approach on PASCAL-5ⁱ for different support set sizes. As in our original results, we use a model trained on 5 shots for all our results, except for the 1 shot case where we use a model trained on 1 shot. Additionally we compare the aggregated results to previous methods in Table 9.

B.3. Runtimes

We show the runtime of our method in Table 10. We partition the timing into different parts. The Gaussian process (GP) is split into two. One part preparing the support set matrix and computing the matrix inverse in (4) and (5) of

Num. Shots	Fold-0	Fold-1	Fold-2	Fold-3	Mean
1-S	50.5	64.9	54.6	52.0	55.5
2-S	57.5	64.8	55.5	50.4	57.1
3-S	62.1	68.7	64.4	57.2	63.1
4-S	64.7	70.2	69.4	59.8	66.0
5-S	66.8	70.7	71.6	63.2	68.1
6-S	67.2	71.5	72.3	64.1	68.8
7-S	67.6	71.8	74.6	65.1	69.8
8-S	67.9	72.9	75.0	66.2	70.5
9-S	68.2	73.0	75.9	66.3	70.8
10-S	68.4	73.5	76.5	67.0	71.4

Table 8. Performance of our model when evaluated at different numbers of shots on the PASCAL-5ⁱ benchmark (mIoU, higher is better).

Method	1-S	5-S	10-S
RPMM [34]	56.3	57.3	57.6
PFENet [28]	60.8	61.9	62.1
RePRI [4]	59.7	66.6	68.1
Ours	55.5	68.1	71.4

Table 9. Performance on PASCAL-5ⁱ for the 1-, 5- and 10-shot settings averaged over 4 folds (mIoU, higher is better).

the main paper, and another part computing the mean and covariance given the pre-computed inverse.

The timings are measured in the 1-shot and 5-shot settings on images from COCO-20ⁱ of 512×512 resolution, using a single episode at a time. We make use of the python-utility `time.perf_counter()` and GPU-synchronization `torch.cuda.synchronize()`. We run our approach on a single NVIDIA V100 for 1000 episodes and report the average timings of each part.

Function	1-shot time	5-shot time
Image encoder on support	14	34
Mask encoder on support	2	3
GP preparation on support	13	129
Image encoder on query	12	14
GP inference on query	11	12
Decoder	4	5
Total	56	197

Table 10. Runtimes of the different functions in our approach, measured in milliseconds (ms). Timings are for a single episode, averaged over 1000 episodes, on 512×512 size images from COCO-20ⁱ.

C. Qualitative Results

We provide additional qualitative results on COCO-20ⁱ and visualizations of the Gaussian process output on PASCAL-5ⁱ.

C.1. Additional Examples

In Figure 5 we show qualitative results on COCO-20ⁱ. Note that in the last row, the annotation is not complete. Only one out of three rows of books placed above some urinals are annotated. Our approach instead correctly segments out all three rows of books.

C.2. Visualization of Gaussian Process Output

We provide a visualization of the Gaussian process output in Figure 6. In the top row, the approach correctly segments the person class. The background in this example contains a large horse. There are no horses in the support set and the Gaussian process is uncertain in that region, reporting high variance. In the second row, we instead aim to segment the horse-class. The support set contains multiple horses similar to that in the query image, and the Gaussian process accurately reports the horse as foreground with low variance. In the third row, our approach misses parts of the plants. These parts are marked with high variance, despite fairly similar examples being present in the support set. In the fourth row, our approach segments a small part of the background, but otherwise performs well. In the fifth row, a carriage is deemed similar to a bicycle and segmented as such. The Gaussian process reports a high variance for the persons but a low variance for large parts of the carriage.



Figure 5. Additional qualitative results of our approach on COCO-20ⁱ. We show the five support images with annotations overlaid (left), the query image and annotation (center), and predictions made by our model (right). Note that the images are cropped for aesthetic reasons.

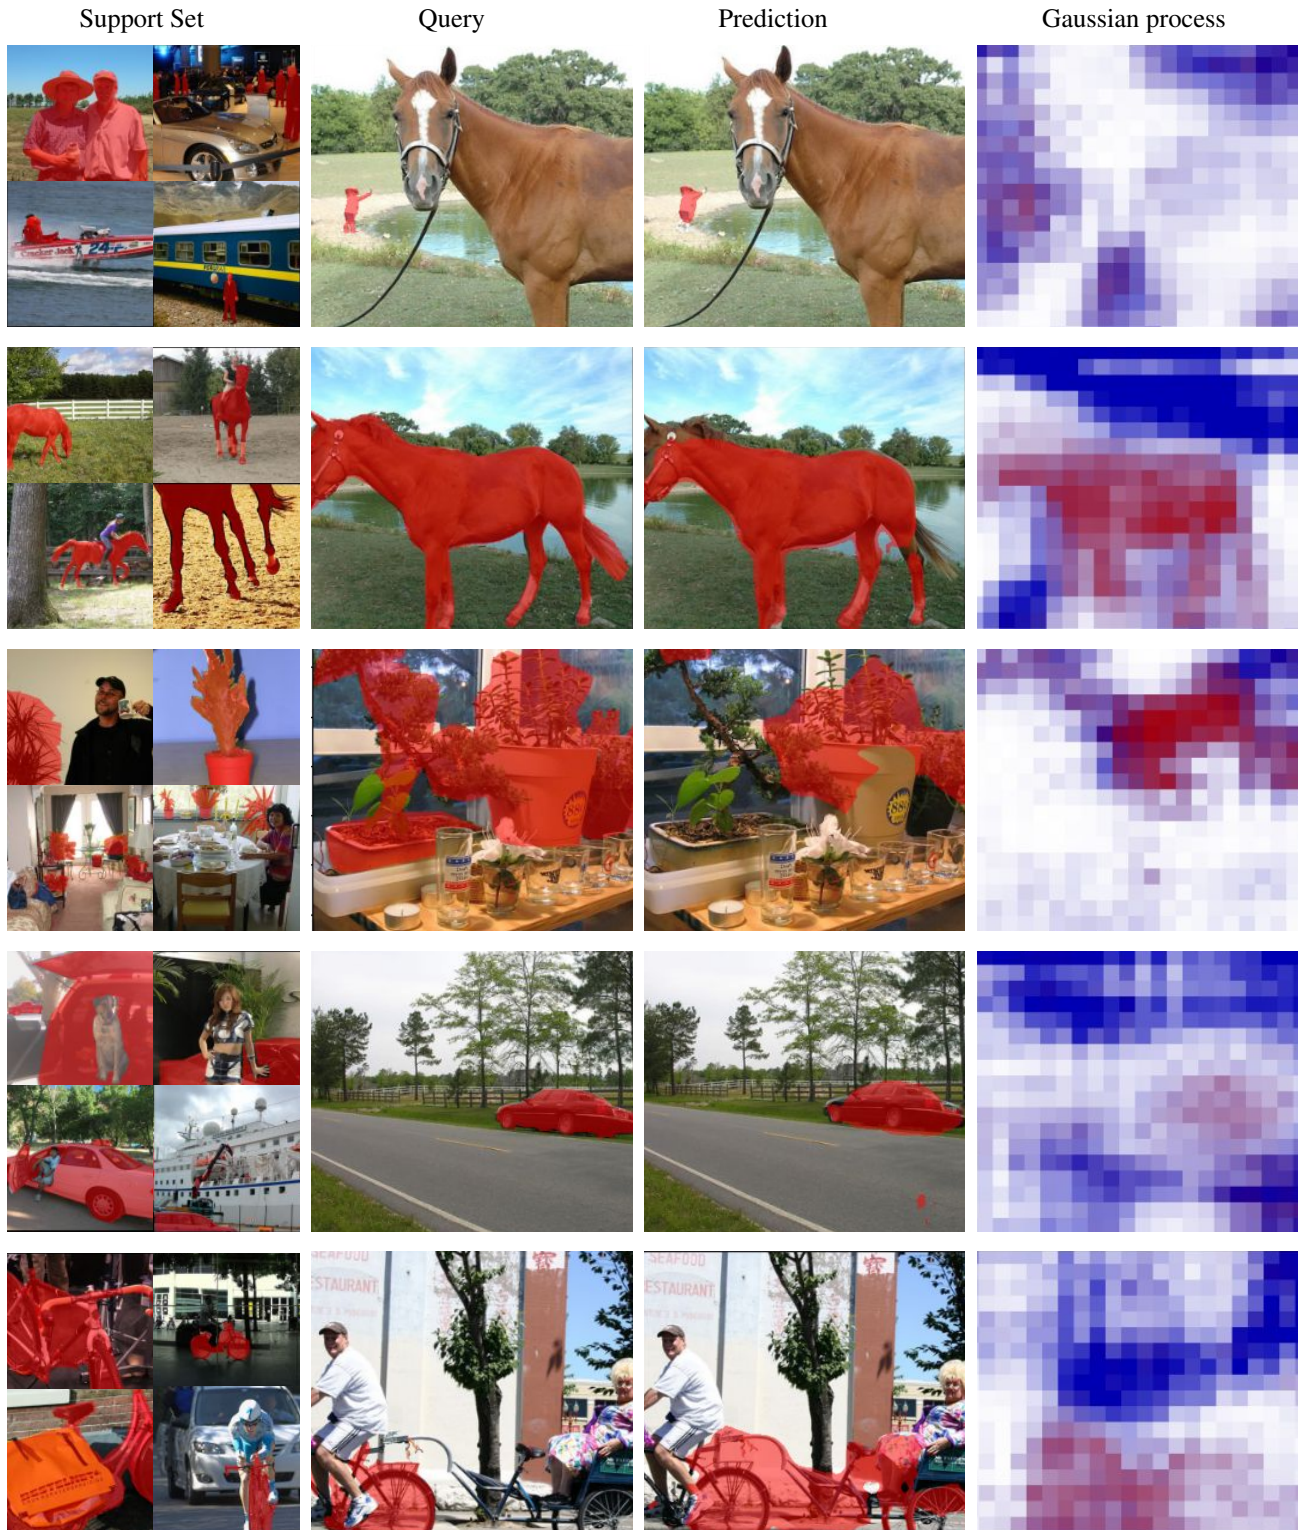


Figure 6. Visualization of the Gaussian process module output on PASCAL-5ⁱ. Here, we use a version of our approach (*Mean + Covariance* in Table 3 of the main paper) with a simple mask encoder to increase interpretability of the Gaussian process. The approach is given four support images. We show the mean $\mu_{Q|S}$ in red-blue. Areas with high variance are shown in white. Images have been cropped for aesthetic reasons.