

# Uncertainty in Real-Time Semantic Segmentation on Embedded Systems

Ethan Goan, Clinton Fookes  
Queensland University of Technology  
Queensland, Australia  
ej.goan@qut.edu.au

## Abstract

Application for semantic segmentation models in areas such as autonomous vehicles and human computer interaction require real-time predictive capabilities. The challenges of addressing real-time application is amplified by the need to operate on resource constrained hardware. Whilst development of real-time methods for these platforms has increased, these models are unable to sufficiently reason about uncertainty present when applied on embedded real-time systems. This paper addresses this by combining deep feature extraction from pre-trained models with Bayesian regression and moment propagation for uncertainty aware predictions. We demonstrate how the proposed method can yield meaningful epistemic uncertainty on embedded hardware in real-time whilst maintaining predictive performance.

## 1. Introduction

Development and capabilities of semantic segmentation models has increased dramatically, with models based on deep learning applied to domains such as autonomous vehicles [1–3], robotics [4] and human computer interaction [2, 5]. Practical application in these areas requires systems to provide real-time performance, and in safety critical domains meaningful uncertainty information is essential. Providing this uncertainty information becomes increasingly challenging for real-time operation, as obtaining this uncertainty information can considerably increase compute demands. A natural way to represent epistemic uncertainty is through a Bayesian framework, where uncertainty in the model is propagated to predictions.

Complete Bayesian inference is intractable for deep learning models, meaning expensive Monte Carlo integration is often used for approximate inference [6–8]. These sampling based techniques considerably increase the time required for prediction, making them unsuitable for real-time operation. Other research have proposed utilising uncertainty information during training [9, 10], though are

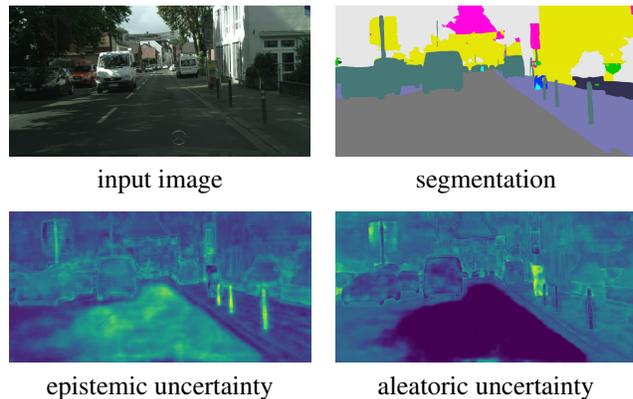


Figure 1. Example of semantic segmentation results obtainable from proposed real-time semantic segmentation method and the forms of uncertainty it permits in real-time applications.

unable to sufficiently reason about epistemic uncertainty present during predictions. Given that the primary practical application of many semantic segmentation models is performed on edge computing devices, it is crucial that we are able to deliver this uncertainty information on these platforms in real time.

This work aims to address this through the use of an analytic probabilistic module inspired by the work of the Gaussian Process treatments of deep kernel learning and Bayesian optimization [11, 12], where the task of feature extraction is computed deterministically with fixed model parameters, which are then used as inputs to a probabilistic classification module. Complexity for Gaussian processes is cubic in the number of training samples, thus making them infeasible for the intended applications. Instead we opt for parameterised probabilistic regression module combined with a moment propagating non-linearity for classification. This reduces the inference overhead for high-dimensional data whilst providing analytic results when using conjugate models, avoiding the need for expensive Monte Carlo approximations. We demonstrate how this approach can be applied to existing real-time semantic seg-

mentation models on embedded hardware whilst providing meaningful uncertainty measures with minimal compute overhead. An example of the predictions and uncertainty measures for the presented models is shown in Figure 1. The contributions of this paper are as follows,

- Propose a light-weight method for uncertainty in semantic segmentation by combining Bayesian methods within moment propagation,
- Develop meaningful real-time aleatoric and epistemic uncertainty metrics and investigation into how they can inform end users and decision protocols,
- Demonstrate how these methods can be easily adapted to pre-trained models, enabling them to provide real time uncertainty quantification on embedded devices.

## 2. Related Work

### 2.1. Deep Learning Models for Real-Time Semantic Segmentation

Whilst state-of-the-art methods in terms of predictive performance utilise modern deep-learning architectures such as transformers [13, 14] or large Fully Convolutional Networks (FCN) models [15, 16], real-time methods are restricted to a class of computationally efficient methods that allow them to be deployed in practice. Real-time semantic segmentation models utilising deep learning architectures traditionally build upon [17] by using FCN architectures consisting of both an encoder and decoder. Foundational real-time methods include ENet [18] which provides a lightweight FCN utilising skip-connections and bottleneck layers similar to ResNets [19], and max-unpooling in the decoder stage. This lightweight model has shown to provide fast predictive performance on consumer hardware at the expense of accuracy. Modern real-time models extend upon this work to design more complex encoders and decoders [20] or through the inclusion of additional auxiliary loss functions that aim to capture semantic information over a larger receptive field [21–23].

### 2.2. Uncertainty in Semantic Segmentation

Description of epistemic uncertainty in semantic segmentation methods typically rely on sampling based procedures such as Monte-Carlo Dropout [6, 8, 24, 25]. To obtain uncertainty information with these methods, a single sample must be passed through the model multiple times whilst sampling model parameters from the approximate posterior. These factors restrict their application to offline methods where considerable compute resources are available. Deterministic Uncertainty Quantification (DUQ) methods avoid the need for multiple forward passes during test time,

though requires the inclusion of an additional weight matrix for each class present in the dataset and a Gaussian Discriminant Model to be placed afterwards [26, 27]. [28] provide an uncertainty mechanism building upon Monte-Carlo Dropout that incorporates averages of an optical flow model to avoid multiple forward passes through frames, though is only applicable to video capture scenarios and relies on aggregating samples over multiple time-steps. This aggregation increases memory footprint for predictions, as prior samples from the optical flow model need to be stored, and computational load is increased due to the need to calculate uncertainty metrics over samples within the flow model. None of the methods reviewed here have been shown to offer real-time predictive performance on traditional hardware, and to the best of our knowledge no prior work has investigated real-time epistemic uncertainty estimation on resource constrained embedded devices.

### 2.3. Gaussian Process Methods and Expectation Propagation

To address the limitations of existing uncertainty quantification for semantic segmentation, we will build upon the work of Deep Gaussian Process models and Expectation Propagation (EP). Deep Gaussian Process models such as [11, 12] separate the tasks of feature extraction and classification, where a Gaussian Process (GP) model is utilised for probabilistic classification. Whilst GPs reduce the need for multiple forward passes, their application requires a large number of parameters to approximate the corresponding covariance function. [29] use a similar concept for separating feature extraction and probabilistic classification for monocular depth estimation. EP methods [30, 31] similarly avoid the need for multiple forward passes, where the distribution of a random variable is approximated analytically through the different transformations. We build upon the concepts of probabilistic classification using deep features and EP to enable uncertainty quantification for semantic segmentation in real-time on embedded hardware, which we describe in the following section.

## 3. Real-time Uncertainty Quantification

We pose the quantification of epistemic uncertainty within the Bayesian framework, where we propagate uncertainty in the model parameters to predictions. For this we require a posterior of our model latent variables  $\omega$  such that,

$$p(\omega|\mathbf{X}, \mathbf{Y}) \propto p(\omega)p(\mathbf{X}, \mathbf{Y}|\omega) \quad (1)$$

where  $\mathbf{X}$  is the set of our training inputs and  $\mathbf{Y}$  is a matrix of our predictive labels. We can then form a predictive

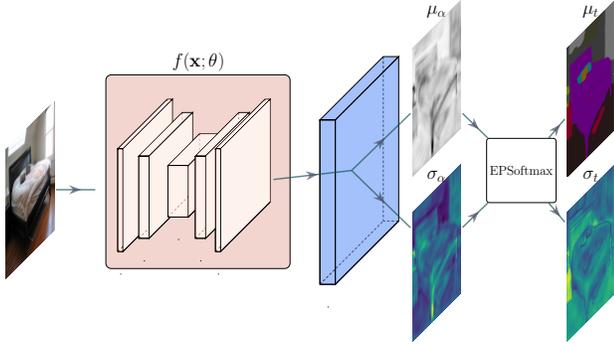


Figure 2. Summary of the proposed model. The pink box represent the deterministic feature extractor  $f(\mathbf{x}; \theta)$  that can be any pre-trained semantic segmentation network. The blue convolutional layer represents our probabilistic layer, which allows for analytic predictive inference, and is used to generate a mean and variance representation of the predictive logits  $\alpha$ . These are then passed through the EPSoftmax layer to represent final predictive categories  $\mathbf{t}$  and epistemic uncertainty information.

distribution for new test data  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$  as,

$$p(\hat{\mathbf{y}}|\hat{\mathbf{x}}, \mathbf{X}, \mathbf{Y}) = \int p(\omega|\mathbf{X}, \mathbf{Y})p(\hat{\mathbf{x}}, \hat{\mathbf{y}}|\omega)d\omega. \quad (2)$$

For the case of a neural network representing our likelihood, the integrals required for Eqns. (1) and (2) are intractable when treating all network parameters as random variables. For the case of (1), the posterior may be approximated by a tractable form through means such as variational inference, though the predictive posterior in (2) for deep neural networks remains intractable and requires expensive Monte-Carlo integration to be used.

The goal of this work is to find a simpler probabilistic representation of model predictions for semantic segmentation that provides efficient analytic computation, thus avoiding the need for expensive sampling approximations. To achieve this, we propose simplifying the probabilistic component of our model to that of linear model with respect to latent variables, where the basis function applied to the input is a neural network  $f(\cdot; \theta)$  parameterised by  $\theta$  which includes all but the last layer of the neural network. These parameters are treated as fixed and known values. We then replace the final linear layer within a given network with a probabilistic layer parameterised by our latent variables  $\omega$  for which we will perform inference. This can be seen as separating the model into a feature extraction stage with fixed parameters  $\theta$ , and the classification stage with parameters treated as random variables  $\omega$ . The latent parameters in  $\omega$  will represent the weights and bias of a final convolutional layer applied to the model. Figure 2 illustrates the

proposed model, and can be summarised as,

$$\Phi(\mathbf{x}) = f(\mathbf{x}; \theta) \quad (3)$$

$$\alpha = \Phi(\mathbf{x})\omega \quad (4)$$

$$\mathbf{t} = \text{Softmax}(\alpha), \quad (5)$$

where  $\Phi(\mathbf{x})$  represents the design matrix for input  $\mathbf{x}$  generated from a neural network with fixed parameters  $\theta$ ,  $\alpha$  and  $\mathbf{t}$  represents the corresponding logits and final predictive categorical probability respectively, and  $\omega \sim p(\omega|\mathbf{X}, \mathbf{Y}, \theta)$  is the latent variables we wish to perform inference for. We represent the multiplication between the design matrix and the probabilistic parameters as an inner product, but note that this inner product can be computed with an equivalent convolution operation. Since inference is only required for the final latent variables, we are able to leverage pretrained networks  $f(\mathbf{x}; \theta)$  for the generation of design matrices.

For the work presented within, we frame inference in the logit space as opposed to the final categorical distribution. This allows us to frame the inference problem as a simple linear regression model, such that our posterior of latent variables  $\omega$  with  $N$  data points is,

$$p(\omega|\mathbf{X}, \mathbf{Y}, \theta) \propto \left[ \prod_{i=0}^{N-1} \mathcal{N}(\omega|f(\mathbf{x}_i; \theta)\omega, \sigma^2) \right] p(\omega). \quad (6)$$

With a conjugate Gaussian prior placed on latent variables  $\omega$ , our posterior will also be Gaussian [32],

$$p(\omega|\mathbf{X}, \mathbf{Y}, \theta) = \mathcal{N}(\mu_\pi, \Sigma_\pi). \quad (7)$$

The advantage of representing our model in this way is that it allows us to represent predictive probabilities over the logits for a single test image  $\hat{\mathbf{X}}$  analytically as,

$$p(\hat{\alpha}|\hat{\mathbf{x}}, \mathbf{X}, \mathbf{Y}, \theta) = \int \mathcal{N}(\omega|\Phi(\hat{\mathbf{x}})\omega, \sigma) p(\omega|\mathbf{X}, \mathbf{Y}, \theta) d\omega \quad (8)$$

$$= \mathcal{N}(\hat{\alpha}|\Phi(\hat{\mathbf{x}})\mu_\pi, \sigma_N(\hat{\mathbf{x}})) \quad (9)$$

$$\sigma_N^2(\hat{\mathbf{x}}) = \sigma^2 + \hat{\mathbf{x}}^T \Sigma_\pi \hat{\mathbf{x}} \quad (10)$$

### 3.1. Inference of Latent Variables

A challenge with this modelling scheme is that since we are performing inference in the logit space, we cannot directly define a likelihood based on the data, as the labels for our data is represented as final categorical probabilities and the softmax function applied to logits to obtain this predictive probability is not bijective. To address this, we instead propose to approximate samples from the conditional posterior of our latent variables  $\omega$  using the diagonal SWAG method [33].

Given a conjugate prior, we know that the posterior of our latent variables  $\omega$  will be a Gaussian, meaning

it can be fully described using only the first two moments. To summarise these moments, we apply the diagonal SWAG method to approximately sample from our model parameters using SGD iterates to compute  $\omega_{\text{SWA}} = \sum_{i=1}^T \omega_i$ . From this, we can compute an online empirical variance estimate for our probabilistic parameters as  $\Sigma_{\text{diag}} = \text{diag}(\bar{\omega}^2 - \omega_{\text{SWG}}^2)$ . In this work, we differ from the original SWAG method, where instead of computing the empirical mean from the SWAG iterates, we use the pre-trained weights for the last convolutional layer in our network  $\bar{\omega} = \omega_{\text{pretrained}}$ . This reduces the need to record many parameter values during training, and encourages mean predictions from the probabilistic model to be similar to the point estimate network. With this empirical variance parameters, we can represent our approximate posterior over our  $N$  probabilistic parameters as a factorised Gaussian distribution,

$$p(\omega|\mathbf{X}, \mathbf{Y}, \theta) \approx \prod_{i=1}^N \mathcal{N}(\omega_i|\bar{\omega}, \Sigma_{\text{diag}}). \quad (11)$$

This allows us to perform approximate inference in the logit space by utilising the categorical labels available for our data sets. Furthermore, the use of a factorised approximation accelerates computation of predictive posterior, as the covariance matrix  $\Sigma_{\pi}$  in (10) becomes a diagonal matrix, such that,

$$\sigma_N^2(\hat{\mathbf{x}}) \approx \sigma^2 + \hat{\mathbf{x}}^T \Sigma_{\text{diag}} \hat{\mathbf{x}}. \quad (12)$$

Whilst eliminating covariance within our posterior approximation, the use of a diagonal covariance matrix considerably reduces the number of model parameters, allowing for easier adaption to low memory applications and accelerate prediction on resource constrained devices.

### 3.2. Predictive Distributions and Measures of Uncertainty

Whilst providing an analytic solution to the predictive probabilities of our logits, we ultimately wish to represent a predictive distribution for final categorical probabilities. An exact analytical solution when using the softmax activation is not known. We address this by building on the methods of moment propagation used in the works of EP, where instead of computing exact distributions, we approximate the moments of random variables after applying certain functions, and use these moments to approximate the transformed random variables as Gaussians. Similar to [31], we compute intermediate moments of the softmax using properties of the Log-Normal distribution.

For a Gaussian random variable  $\gamma \sim \mathcal{N}(\mu_{\gamma}, \sigma_{\gamma}^2)$ ,  $\beta = \exp(\gamma) \sim \text{Lognormal}(\mu_{\beta}, \sigma_{\beta}^2)$ , where the mean and variance of  $\beta$  is  $\mu_{\beta} = \exp(\mu_{\gamma} + \sigma_{\gamma}^2/2)$  and  $\sigma_{\beta}^2 = (\exp \sigma_{\gamma}^2 -$

$1) \exp(2\mu_{\gamma} + \sigma_{\gamma}^2)$  respectively. We can use these two moments to approximate the transformed random variables as a Gaussian  $\beta \sim \mathcal{N}(\mu_{\beta}, \sigma_{\beta}^2)$ . From this we can approximate the distribution of the output for each class as,

$$t_j = \frac{\exp(\alpha_j)}{\sum_{i=0}^{N-1} \exp(\alpha_i)} = \frac{y_j}{\sum_{i=0}^{N-1} y_i} \quad (13)$$

where  $\alpha_j$  is a single logit from our probabilistic layer and  $y_i \sim \mathcal{N}(\mu_{y,i}, \sigma_{y,i}^2)$  using the properties of the Log-Normal distribution. To find the mean and variance for our output  $t_i$  using moment propagation, we first find the mean and variance of the denominator. Assuming independence in our outputs, the denominator distribution in Equation (13) can be represented as,

$$\sum_{i=0}^{N-1} y_i \sim \mathcal{N}(\sum \mu_{y,i}, \sum \sigma_{y,i}^2) = \mathcal{N}(\mu_d, \sigma_d^2). \quad (14)$$

The application of moment propagation used in [31] does not aim to approximate the moments from the ratio distribution encountered within a probabilistic treatment of the softmax, and instead reduce the Gaussian random variable to a Dirichlet random variable, meaning additional computation is required to obtain class-conditional uncertainty measures. We rectify this by following ratio distribution approximation of [34], where a Gaussian approximation is derived from a Taylor expansion. With this we can approximate the mean and variance of each softmax output indexed by  $j$  as,

$$t_j \sim \mathcal{N}\left(\frac{\mu_j}{\mu_d}, \frac{\mu_j}{\mu_d} \sqrt{\sigma_j^2 + \sigma_d^2}\right), \quad (15)$$

where  $\mu_j$  and  $\sigma_j$  are found from the Log-Normal properties. We label the use of this moment propagation as the Expectation Propagation Softmax (EPSoftmax).

Whilst having the ability to represent classification uncertainty conditional on our deep feature extraction, we also wish to form a suitable categorical distribution over outputs for classification using the conventional argmax operator. We show how we can obtain categorical probabilities for classification using the described approach.

**Proposition 1.** *With the probabilistic output in Eq. (15), we can create a valid  $k$ -dimensional categorical distribution such that,*

$$C \sim \text{Cat}(k, \mathbb{E}[\mathbf{t}]). \quad (16)$$

*Proof.* For  $\text{Cat}(k, \mathbb{E}[\mathbf{t}])$  to be a valid categorical distribution, we require that the parameters generated by  $\mathbb{E}[\mathbf{t}]$  satisfy the condition that  $\mathbb{E}[\mathbf{t}]_j \leq 1$  for  $0 \leq j < k$  and  $\sum_{i=0}^{k-1} \mathbb{E}[\mathbf{t}]_i = 1$ . We first observe that expectation for the  $j^{\text{th}}$  component is,

$$\mathbb{E}[\mathbf{t}]_j = \frac{\mu_j}{\mu_d} = \frac{\exp(\mu_j + \sigma_j^2/2)}{\sum_{i=0}^{k-1} \exp(\mu_i + \sigma_i^2/2)} = \frac{\exp(\mu_j + \sigma_j^2/2)}{\mu_d}, \quad (17)$$

where properties of the Lognormal distribution have been used and  $\mu_d$  is that of Equation (14). Similar to the traditional softmax function, the denominator is shared amongst all components, and the exponential function ensures all components are strictly positive. Following the same arguments of the softmax, this implies that  $\exp(\mu_j + \sigma_j^2/2)$  for valid index  $j$ , and that  $\sum_i^{k-1} \mathbb{E}[t]_i = 1$  as required.  $\square$

With this categorical representation, we can easily perform final classification similar to point estimate networks, whilst also delivering important uncertainty information in these predictions. In the next section, we discuss the uncertainty metrics available.

### 3.3. Computing Predictive Uncertainty

With our representation of predictive probabilities, we want to be able to reason about the different types of uncertainty present. It is common to refer to the epistemic and aleatoric uncertainty [25], where epistemic uncertainty is a representation of reducible uncertainty identified by the model, and aleatoric uncertainty is irreducible uncertainty caused by the data present. With the probabilistic model presented within, we are able to reason about both forms of uncertainty.

We propose to capture epistemic uncertainty using the variance information available in our model predictions. An advantage of our approach is that we can measure this uncertainty in both the logit and prediction space. Given we are modelling these spaces as Gaussians, we can succinctly summarise this uncertainty using the entropy of these distributions. The entropy for a D-dimensional multivariate Gaussian is,

$$\mathbb{H}[x] = \frac{1}{2} \log |\Sigma| + \frac{D}{2} (1 + \log 2\pi). \quad (18)$$

We see from Eq. (18) that the entropy depends only on the covariance information in the distribution. This is advantageous as it provides a succinct and natural form to summarise the uncertainty induced by our latent variables in our predictions. Furthermore, given that we are representing our distributions over our logits and final predictions as independent Gaussian variables, the determinant operator can be reduced to a product of the variance components for each component in the relevant distribution.

Entropy provides a succinct summary for our uncertainty over all categories of interest, though there are instances where we may wish to view uncertainty information pertaining to a single class of interest. For example, for an autonomous vehicle scenario we may be interested in observing the uncertainty for safety critical classes such as other vehicles or pedestrians. We can summarise the class conditional uncertainty utilising the corresponding variance in our marginal predictive probability in Eq. (15) for classes of interest.

It has been shown that entropy of the final categorical probability obtained from classification models can serve as an upper-bound for aleatoric uncertainty [35]. We build upon this insight in this work to allow us to approximate aleatoric uncertainty  $\mathcal{U}_A$  using the entropy  $\mathbb{H}$  of the categorical distribution represented in (16) such that,

$$\mathcal{U}_A \approx \mathbb{H}[C] \quad (19)$$

Using the method described here, we are able to reason about both aleatoric uncertainty and epistemic uncertainty present within our model. We demonstrate in the following section how we can use the described approach for real-time semantic segmentation.

## 4. Experiments

### 4.1. Experimental Setup

With our probabilistic model, inference scheme and uncertainty measures defined, we now demonstrate how they can all be combined to deliver uncertainty in real-time compute constrained devices. The NVIDIA Jetson AGX Xavier embedded device is used as the computing platform. Our aim here is to demonstrate how the proposed uncertainty mechanism can be applied to existing pretrained semantic segmentation models to endow them with the capability to compute aleatoric uncertainty measures. Code for experiments is available at <https://github.com/egstatsml/eu-seg>.

For our inference procedure, we perform training on the existing models to obtain our diagonal SWAG samples with the Ohem cross-entropy loss [36]. We perform a total of 5,000 SGD iterations, with a warmup of 1,000 iterations using a linear increase of the learning rate. After this warmup period, the learning rate remains constant. The parameters are observed every 50 SGD iterations for the remainder of training for inference utilising diagonal SWAG. These recorded parameters are then used to estimate the empirical covariance matrix for the final probabilistic convolutional layer. Our prior over our latent variables is a Gaussian distribution  $\mathcal{N}(0, 1 \times 10^{-4})$ , which is implemented using weight decay.

After training is completed, the models are then compiled to the TensorRT format to be used on the Jetson board for calculation of inference speed metrics. During compilation, the models are converted to half-precision floating point values to accelerate computation. To measure predictive speed, we time the computation required for 1,000 forward passes. Our primary measures for evaluating model predictive performance is the mean intersection over union (mIOU) and the Macro and Micro F1 scores. Predictive speed is measured in frames-per-second (fps) on the Jetson device, and compared against traditional point estimate networks.

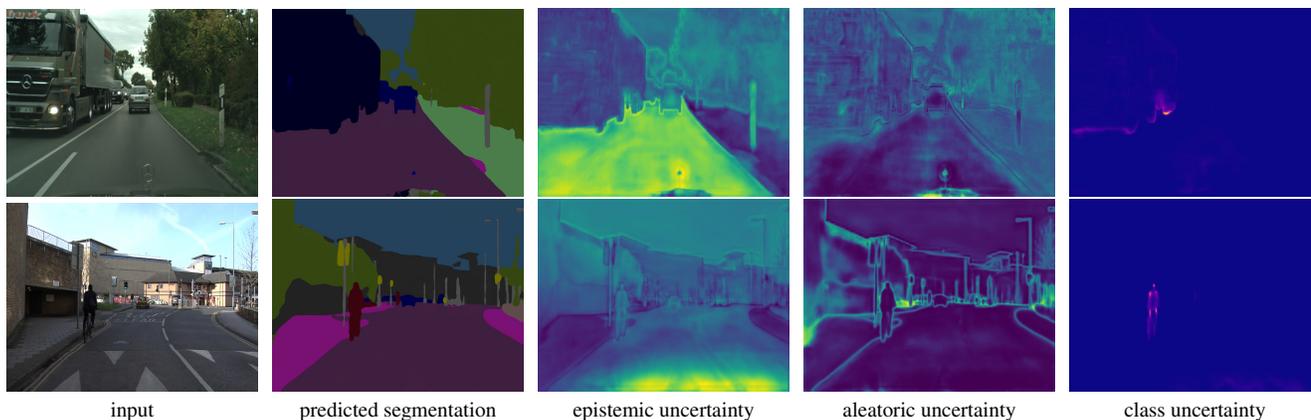


Figure 3. Examples of predictions from the proposed model for the CityScapes Dataset utilising the BiSeNetV1 backbone (top row) and the CamVid datasets using the PIDNet backbone (bottom row). From left to right, the columns represent the input, classification, epistemic uncertainty, aleatoric uncertainty, and class conditional epistemic uncertainty. For epistemic and aleatoric uncertainty, green regions indicate larger uncertainty, whilst red indicates greater uncertainty for class conditional uncertainty. For class conditional uncertainty, we show the uncertainty for the classes “truck” and “person”. Best viewed on computer screen.

## 4.2. Urban-Scene Datasets for Autonomous Driving

One area where real-time predictions is required is autonomous vehicles. Given the rapid development and need for uncertainty measures in this safety critical scenario, we perform experimentation on the CamVid [1] and CityScapes [3] datasets where pre-trained base networks are available to form our feature extractor  $f(\cdot, \theta)$ . For the CamVid datasets, we perform experimentation adding the proposed EP uncertainty mechanism to ENet [18] and PIDNet [23]. For ENet models, UnMaxpooling operations in the decoder were replaced with interpolation stages followed by a small finetuning stage to accommodate the change. This change is required as TensorRT does not support the unpooling operations. For CityScapes, a larger number of pre-trained models are available, allowing us to perform experimentation using the previous models and BiSeNetV1 [21], BiSeNetV2 [22] and PPLiteSegT and PPLiteSegB [37]. We summarise the predictive performance of these models for the CityScapes and CamVid datasets in Tables 1 and 2 respectively, where they are also compared against equivalent point estimate networks.

From the results in Table 1, we see comparable performance amongst the probabilistic models and their point estimate equivalents. Whilst the predictive speed does decrease across the Bayesian models, we see that predictive performance remains suitable for real time application. Interestingly, we note that the models that have shown faster predictions on traditional hardware such as ENet and BiSeNetV2 suffer from slower than expected predictive times. We attribute this decrease in predictive performance on the Jetson hardware to the depth-wise convolutional operations used in the BiSeNetV2 model and the excessive interpolation stages

in ENet not being as thoroughly optimised within the TensorRT framework. This highlights the importance of designing bespoke architectures for edge devices.

Table 1. Summary of predictive performance of real-time semantic segmentation methods on Jetson Xavier embedded hardware for CityScapes dataset. Uncertainty enabled models are given the “Bayes-” prefix.

Model	Macro-F1	Micro-F1	mIOU	fps
BiSeNetV1	0.8423	0.9556	0.7426	68.81
Bayes-BiSeNetV1	0.8415	0.9550	0.7408	57.359
BiSeNetV2	0.8473	0.9565	0.7476	49.502
Bayes-BiSeNetV2	0.8456	0.9558	0.7451	57.359
ENet	0.7125	0.8976	0.5795	45.1638
Bayes-ENet	0.7125	0.8974	0.5795	30.709
PIDNet	0.9150	0.9705	0.8486	84.298
Bayes-PIDNet	0.9129	0.9695	0.8454	69.367
PPLiteSegT	0.8534	0.9558	0.7578	72.5234
Bayes-PPLiteSegT	0.8481	0.9536	0.7501	65.931
PPLiteSegB	0.8666	0.9597	0.7750	56.9437
Bayes-PPLiteSegB	0.8603	0.9571	0.7656	52.022

Table 2. Summary of predictive performance of real-time semantic segmentation methods on Jetson Xavier embedded hardware for CamVid dataset. Uncertainty enabled models are given the “Bayes-” prefix.

Model	Macro-F1	Micro-F1	mIOU	fps
ENet	0.6932	0.9116	0.5950	42.2916
Bayes-ENet	0.6702	0.8990	0.5698	30.825
PIDNet	0.8683	0.9485	78.4699	73.7831
Bayes-PIDNet	0.8661	0.9471	0.7816	59.978

With our predictive performance measured quantita-

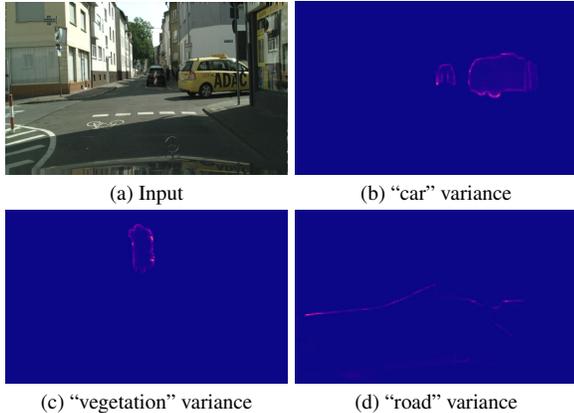


Figure 4. Visualisation of class-conditional uncertainty for a sample from CityScapes dataset. Red regions indicate an increase in uncertainty. Predictions performed with PIDNet. Best viewed on a computer screen.

tively, we now investigate the uncertainty measures qualitatively. In Section 3.3, we state how the proposed modelling method is capable of producing measures for aleatoric, epistemic and class conditional epistemic uncertainty. Figure 3 illustrates the predictive performance over CityScapes and CamVid datasets and the relevant uncertainty measures. We can see from these figures that the measured epistemic uncertainty from entropy in Gaussian representation of our output (Eqn 18) is largest within the objects of interest, whilst the aleatoric uncertainty measured from the entropy of the final categorical distribution (Eqn 19) is concentrated around the edges of these objects. We further see from the class conditional uncertainty measures that when targeting individual classes, the uncertainty is targeted towards the edge of the objects for the class of interest. We explore additional examples of the class-conditional uncertainty measures in Figure 4 to further verify this result.

### 4.3. Semantic Segmentation of General Objects

Robotics and human computer interaction systems are another class of practical application where semantic segmentation methods are required to operate in real-time on edge devices. Given the complexity of operating with humans or the dynamic environments robotics systems encounter, it is important that segmentation methods used within these scenarios are able to reason about uncertainty present. We perform additional experimentation using the challenging ADE20k [5] and CoCoStuff [2] datasets to assert the suitability of the proposed probabilistic module for data seen within these scenarios. These experiments build upon the available pre-trained BiSeNetV1 and BiSeNetV2 models to form  $f(\cdot, \theta)$ , with results shown in Tables 3 and 4. Visualisations of uncertainty measures is shown in Figure 5, and examples of class conditional uncertainty in Figure 6.

Table 3. Summary of predictive performance of real-time semantic segmentation methods on Jetson Xavier embedded hardware for CoCoStuff dataset.

Model	Macro-F1	Micro-F1	mIOU	fps
BiSeNetV1	0.4254	0.6295	0.3077	72.4150
Bayes-BiSeNetV1	0.4256	0.6300	0.3079	43.3340
BiSeNetV2	0.4042	0.6025	0.2851	58.6345
Bayes-BiSeNetV2	0.4047	0.6027	0.2855	41.2032

Table 4. Summary of predictive performance of real-time semantic segmentation methods on Jetson Xavier embedded hardware for ADE20k dataset.

Model	Macro-F1	Micro-F1	mIOU	fps
BiSeNetV1	0.4859	0.7643	0.3513	62.4800
Bayes-BiSeNetV1	0.4861	0.7642	0.3514	54.2450
BiSeNetV2	0.4490	0.7593	0.3211	59.5024
Bayes-BiSeNetV2	0.4494	0.7591	0.3215	41.8055

These results further demonstrate the ability of the proposed uncertainty mechanism to provide comparable predictive accuracy to the corresponding point-estimate networks on embedded hardware, whilst also providing a range of meaningful uncertainty estimates. We see in Figure 6 that class conditional accuracy has started to become much more diffuse, though is still concentrated around the edges of the objects of interest. This can be attributed to the greater complexity and diversity offered within the ADE20k and CoCoStuff datasets, and highlights the importance of these uncertainty measures for these challenging scenarios.

## 5. Discussion and Future Research

Experimental results in Section 4 have demonstrated how the proposed uncertainty propagation method can provide meaningful uncertainty estimates when included with a variety of existing real-time segmentation networks. From these results we highlight areas for future research and how to further enhance the ability to capture uncertainty within real-time semantic segmentation models.

The proposed approach builds upon existing segmentation models to enable them to capture uncertainty present within parameters in the classification layer and propagate these uncertainties to meaningful predictive measurements. Since the incorporation of the proposed uncertainty mechanism is included post-training of the feature extractor  $f(\cdot, \theta)$ , any uncertainty present during training of the base network is not explicitly captured. Recent methods have proposed training mechanisms to fit semantic-segmentation models that incorporate aleatoric uncertainty into the training procedure [38–40]. Incorporation of these training schemes with the proposed uncertainty propagation module could allow meaningful epistemic uncertainty esti-

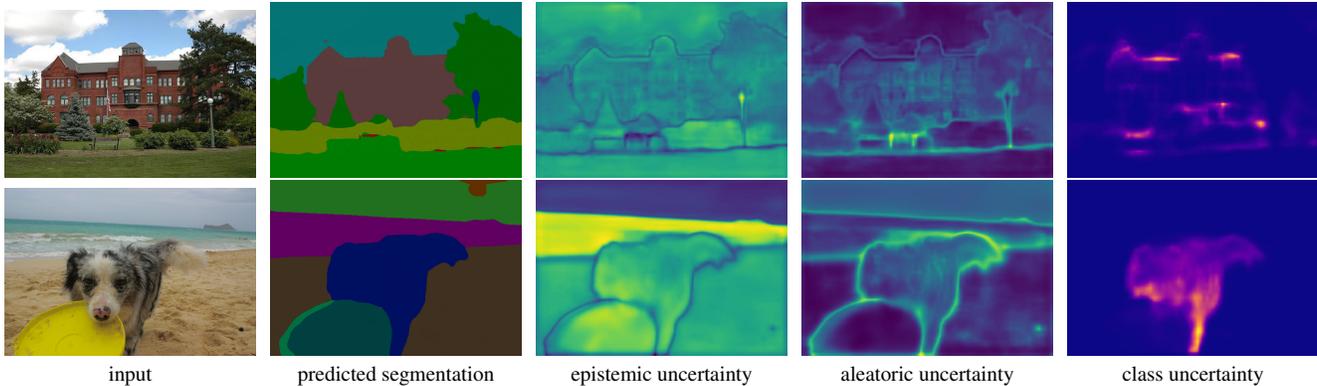


Figure 5. Examples of predictions from the proposed model for the ADE20k Dataset (top row) and the CoCoStuff dataset (bottom row) when utilising the BiSeNetv1 backbone. From left to right, the columns represent the input, classification, epistemic uncertainty, aleatoric uncertainty, and class conditional epistemic uncertainty. For epistemic and aleatoric uncertainty, green regions indicate larger uncertainty, whilst red indicates greater uncertainty for class conditional uncertainty. For class conditional uncertainty, we show the uncertainty for the classes “building” and “dog”. Best viewed on computer screen.

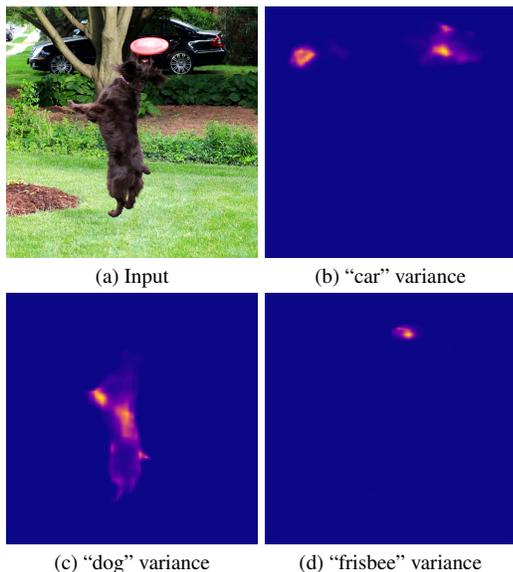


Figure 6. Visualisation of class-conditional uncertainty for a sample from CoCoStuff dataset. Red regions indicate an increase in uncertainty. Predictions performed with BiSeNetV2. Best viewed on computer screen.

mates to be found whilst encouraging the base network to have a more accurate representation of aleatoric uncertainty present for resource constrained hardware.

Whilst the vast majority of work on modelling uncertainty within deep-learning has focused on finding predictive uncertainty estimates, there has been relatively little research on how best to utilise these measures in practice despite the existence of a rich and mature theory to address these important challenges [41, 42]. Given the rising prevalence of semantic segmentation systems to safety critical en-

vironments such as autonomous vehicles, medical imaging or human-computer interaction, it is crucial that sensible and safe decisions are made by the machine learning system. Optimal and safe decisions require the incorporation of meaningful uncertainty information. The sensitivity to context of different application poses further challenges associated with optimal and safe decision making for automated systems. Given the need to process and interpret uncertainty information present, the additional processing of this information will incur a computational overhead. This computational burden is especially important for real-time embedded systems, and we highlight as an important area for future research in practical uncertainty aware semantic segmentation systems.

## 6. Conclusion

Within this research, we targeted the need for meaningful uncertainty information for semantic segmentation on resource constrained hardware. We proposed a combination of a deterministic feature extractor with a probabilistic regression module, that can be combined with an EP softmax module to generate final predictive outputs and uncertainties. We demonstrate how the proposed probabilistic module can be included with existing pre-trained networks. Evaluation of these models on embedded hardware using multiple datasets demonstrate how the proposed probabilistic scheme can mitigate the expensive compute traditionally required for a complete probabilistic treatment of deep segmentation models to provide real-time performance on resource constrained hardware. Qualitative evaluation of obtained uncertainty measures demonstrate how the different forms of uncertainty obtained can be used to identify different regions of interest.

## References

- [1] Gabriel J Brostow, Julien Fauqueur, and Roberto Cipolla, “Semantic object classes in video: A high-definition ground truth database,” *Pattern Recognition Letters*, vol. 30, no. 2, pp. 88–97, 2009. [1](#), [6](#)
- [2] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari, “Coco-stuff: Thing and stuff classes in context,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [1](#), [7](#)
- [3] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223. [1](#), [6](#)
- [4] Andres Milioto, Philipp Lottes, and Cyrill Stachniss, “Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in cnns,” in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 2229–2235. [1](#)
- [5] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba, “Scene parsing through ade20k dataset,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [1](#), [7](#)
- [6] Jishnu Mukhoti and Yarin Gal, “Evaluating bayesian deep learning methods for semantic segmentation,” *arXiv preprint arXiv:1811.12709*, 2018. [1](#), [2](#)
- [7] Clément Dechesne, Pierre Lassalle, and Sébastien Lefèvre, “Bayesian u-net: Estimating uncertainty in semantic segmentation of earth observation images,” *Remote Sensing*, vol. 13, no. 19, 2021. [1](#)
- [8] Michael Kampffmeyer, Arnt-Borre Salberg, and Robert Jenssen, “Semantic segmentation of small objects and modeling of uncertainty in urban remote sensing images using deep convolutional neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2016, pp. 1–9. [1](#), [2](#)
- [9] Prabhu Teja Sivaprasad and Francois Fleuret, “Uncertainty reduction for model adaptation in semantic segmentation,” in *2021 IEEE/CVF Conference On Computer Vision And Pattern Recognition, Cvpr 2021*. IEEE, 2021, number CONF, pp. 9608–9618. [1](#)
- [10] Zhedong Zheng and Yi Yang, “Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation,” *International Journal of Computer Vision*, vol. 129, no. 4, pp. 1106–1120, 2021. [1](#)
- [11] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing, “Deep kernel learning,” in *Artificial intelligence and statistics*. PMLR, 2016, pp. 370–378. [1](#), [2](#)
- [12] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams, “Scalable bayesian optimization using deep neural networks,” in *International conference on machine learning*. PMLR, 2015, pp. 2171–2180. [1](#), [2](#)
- [13] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao, “Vision transformer adapter for dense predictions,” *arXiv preprint arXiv:2205.08534*, 2022. [2](#)
- [14] Yuhui Yuan, Xiaokang Chen, Xilin Chen, and Jingdong Wang, “Segmentation transformer: Object-contextual representations for semantic segmentation,” 2020. [2](#)
- [15] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al., “InternImage: Exploring large-scale vision foundation models with deformable convolutions,” *arXiv preprint arXiv:2211.05778*, 2022. [2](#)
- [16] Jun Fu, Jing Liu, Jie Jiang, Yong Li, Yongjun Bao, and Hanqing Lu, “Scene segmentation with dual relation-aware attention network,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 6, pp. 2547–2560, 2020. [2](#)
- [17] Jonathan Long, Evan Shelhamer, and Trevor Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440. [2](#)
- [18] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello, “Enet: A deep neural network architecture for real-time semantic segmentation,” *arXiv preprint arXiv:1606.02147*, 2016. [2](#), [6](#)
- [19] 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*, 2016, pp. 770–778. [2](#)

- [20] Haochen Wang, Xiaolong Jiang, Haibing Ren, Yao Hu, and Song Bai, “Swiftnet: Real-time video object segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1296–1305. 2
- [21] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang, “Bisenet: Bilateral segmentation network for real-time semantic segmentation,” *CoRR*, 2018. 2, 6
- [22] Changqian Yu, Changxin Gao, Jingbo Wang, Gang Yu, Chunhua Shen, and Nong Sang, “Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation,” *CoRR*, 2020. 2, 6
- [23] Jiacong Xu, Zixiang Xiong, and Shankar P Bhattacharyya, “Pidnet: A real-time semantic segmentation network inspired from pid controller,” *arXiv preprint arXiv:2206.02066*, 2022. 2, 6
- [24] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla, “Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding,” in *British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017*. 2017, BMVA Press. 2
- [25] Alex Kendall and Yarin Gal, “What uncertainties do we need in bayesian deep learning for computer vision?,” *Advances in neural information processing systems*, vol. 30, 2017. 2, 5
- [26] Joost Van Amersfoort, Lewis Smith, Yee Whye Teh, and Yarin Gal, “Uncertainty estimation using a single deep deterministic neural network,” in *International conference on machine learning*. PMLR, 2020, pp. 9690–9700. 2
- [27] Jishnu Mukhoti, Joost van Amersfoort, Philip HS Torr, and Yarin Gal, “Deep deterministic uncertainty for semantic segmentation,” *arXiv preprint arXiv:2111.00079*, 2021. 2
- [28] Po-Yu Huang, Wan-Ting Hsu, Chun-Yueh Chiu, Ting-Fan Wu, and Min Sun, “Efficient uncertainty estimation for semantic segmentation in videos,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 520–535. 2
- [29] Chao Qu, Wenxin Liu, and Camillo J Taylor, “Bayesian deep basis fitting for depth completion with uncertainty,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 16147–16157. 2
- [30] Brendan J Frey and Geoffrey E Hinton, “Variational learning in nonlinear gaussian belief networks,” *Neural Computation*, vol. 11, no. 1, pp. 193–213, 1999. 2
- [31] Jochen Gast and Stefan Roth, “Lightweight probabilistic deep networks,” *CoRR*, 2018. 2, 4
- [32] kevin Murphey, *Machine learning, a probabilistic perspective*, MIT Press, Cambridge, MA, 2012. 3
- [33] Wesley J. Maddox, Pavel Izmailov, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson, “A simple baseline for bayesian uncertainty in deep learning,” in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, Eds., 2019, pp. 13132–13143. 3
- [34] Eloísa Díaz-Francés and Francisco J Rubio, “On the existence of a normal approximation to the distribution of the ratio of two independent normal random variables,” *Statistical Papers*, vol. 54, no. 2, pp. 309–323, 2013. 4
- [35] Lewis Smith and Yarin Gal, “Understanding measures of uncertainty for adversarial example detection,” in *Uncertainty in AI*, 2018. 5
- [36] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick, “Training region-based object detectors with online hard example mining,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 761–769. 5
- [37] Juncai Peng, Yi Liu, Shiyu Tang, Yuying Hao, Lutao Chu, Guowei Chen, Zewu Wu, Zeyu Chen, Zhiliang Yu, Yuning Du, et al., “Pp-liteseg: A superior real-time semantic segmentation model,” *arXiv preprint arXiv:2204.02681*, 2022. 6
- [38] Francois Fleuret et al., “Uncertainty reduction for model adaptation in semantic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9613–9623. 7
- [39] Yuxi Wang, Junran Peng, and ZhaoXiang Zhang, “Uncertainty-aware pseudo label refinery for domain adaptive semantic segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9092–9101. 7

- [40] Qianyu Zhou, Zhengyang Feng, Qiqi Gu, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma, “Uncertainty-aware consistency regularization for cross-domain semantic segmentation,” *Computer Vision and Image Understanding*, vol. 221, pp. 103448, 2022. 7
- [41] Morris H DeGroot, *Optimal statistical decisions*, John Wiley & Sons, 2005. 8
- [42] Mykel J Kochenderfer, Tim A Wheeler, and Kyle H Wray, *Algorithms for decision making*, MIT press, 2022. 8