

# Errata: Uncertainty in Real-Time Semantic Segmentation on Embedded Systems

Ethan Goan

Errata Corrected on arXiv version of paper [1].

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

**Eq.10 and 12:** are missing  $\Phi$ , they should read

$\sigma_N^2(\hat{\mathbf{x}}) = \sigma^2 + \Phi(\hat{\mathbf{x}})^T \Sigma_\pi \Phi(\hat{\mathbf{x}})$  and  $\sigma_N^2(\hat{\mathbf{x}}) \approx \sigma^2 \Phi(\hat{\mathbf{x}})^T \Sigma_{\text{diag}} \Phi(\hat{\mathbf{x}})$

**Eq.15:** The approximation from ref [34] in the paper provides the Taylor Series approximation of  $\mathcal{N}\left(\frac{\mu_j}{\mu_d}, \frac{\mu_j^2}{\mu_d^2} (\sigma_j^2 / \mu_j^2 + \sigma_d^2 / \mu_d^2)\right)$ . The issue with this is the square of the  $\mu_d$  terms, as this is at risk of numerical overflow in half floating-point precision. Instead we use in  $\sigma^2 \approx \mu_j / \mu_d \sqrt{\sigma_j^2 + \sigma_d^2}$ , which empirically followed a similar trend and was less of a risk for numerical overflow (though not immune).

In Table 1, predictive performance metrics for PIDNet used the pre-trained weights from the official repo for the testing set of CityScapes. I initially thought this was due to me evaluating at 1024x512 instead of 2048x1024 like the original paper did, but it that those weights have included the validation data during training. The mIoU for PIDNet and Bayes-PIDNet with the weights for the validation set are 0.761 and 0.758 respectively.

## References

- [1] E. Goan and C. Fookes. Uncertainty in real-time semantic segmentation on embedded systems. *arXiv preprint arXiv:2301.01201*, 2022.