

# Errata and Addenda For My Articles

Ethan Goan

March 30, 2026

## Contents

<b>1</b>	<b>Uncertainty in Real-Time Semantic Segmentation on Embedded Systems</b>	<b>1</b>
<b>2</b>	<b>Piecewise Deterministic Markov Processes for Bayesian Neural Networks</b>	<b>1</b>
<b>3</b>	<b>Bayesian Neural Networks: An Introduction and Survey</b>	<b>1</b>
<b>1</b>	<b>Uncertainty in Real-Time Semantic Segmentation on Embedded Systems</b>	

Corrected in arXiv version of paper

**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** 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 single-point precision. Instead we use  $\sigma^2 \approx \mu_j/\mu_d \sqrt{\sigma_j^2 + \sigma_d^2}$ , 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, which appears to 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.

## 2 Piecewise Deterministic Markov Processes for Bayesian Neural Networks

Corrected in arXiv version of paper

**Eq. 11:**  $t_i = -b_i/a_i + \sqrt{b_i^2 + a_i^2 t_{i-1}^2 + 2a_i b_i t_{i-1} - 2a_i \log(1 - U)} / a_i$

**Bug:** When preparing my thesis, I found a bug in the code where gradients were not being appropriately scaled by the mini-batch and dataset size. Experiments rerun with results on arXiv. I also added some clarification to the event rate sampling algorithm which clarifies some of the results of how thinning work. This is captured by the  $R$  term in Algorithm 2, and we explore this more in Appendix A. Code updated on Github as well.

## 3 Bayesian Neural Networks: An Introduction and Survey

Corrected on arXiv

For the  $\alpha$  divergence on page 31, the limits to relate to forward/reverse KL divergence and Hellinger distance were not correct. It should be,  $KL(q|p) : \alpha \rightarrow 0$ ,  $KL(p|q) : \alpha \rightarrow 1$ , Hellinger:  $\alpha \rightarrow 0.5$ .