MLE estimators are not always unbiased

Very simple example:

Suppose $\{x_i\}$ are N measurements of an unknown quantity μ with unknown resolution σ. What are the MLE estimates of μ and σ ?

$$\mathcal{L} = \prod p(x_i) = \prod \frac{1}{\sigma\sqrt{2\pi}} e^{-(x_i-\mu)^2/(2\sigma^2)}$$

$$\log \mathcal{L} = -\sum \frac{(x_i-\mu)^2}{2\sigma^2} - \sum \log\sigma - \sum \log\sqrt{2\pi}$$

$$\frac{\partial \log \mathcal{L}}{\partial \mu} = 0 \quad \rightarrow \quad \sum (x_i - \mu) = 0 \quad \rightarrow \quad \boxed{\hat{\mu} = \frac{\sum x_i}{N} = \overline{x_i}}$$

This is unbiased

$$\frac{\partial \log \mathcal{L}}{\partial \sigma} = 0 \quad \rightarrow \quad \sum \frac{(x_i-\mu)^2}{\sigma^3} - \sum \frac{1}{\sigma} = 0$$

$$\rightarrow \quad \boxed{\hat{\sigma}^2 = \frac{\sum (x - \overline{x_i})^2}{N}}$$

This is biased

Plot MLE form 10K toy experiments, N=10, μ=0, σ=1

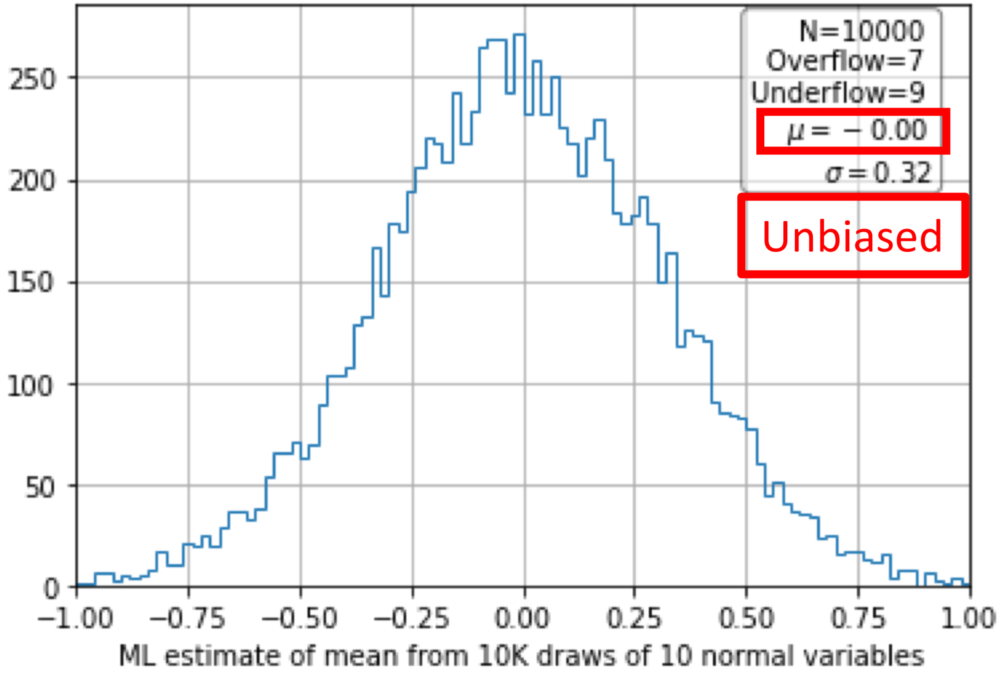

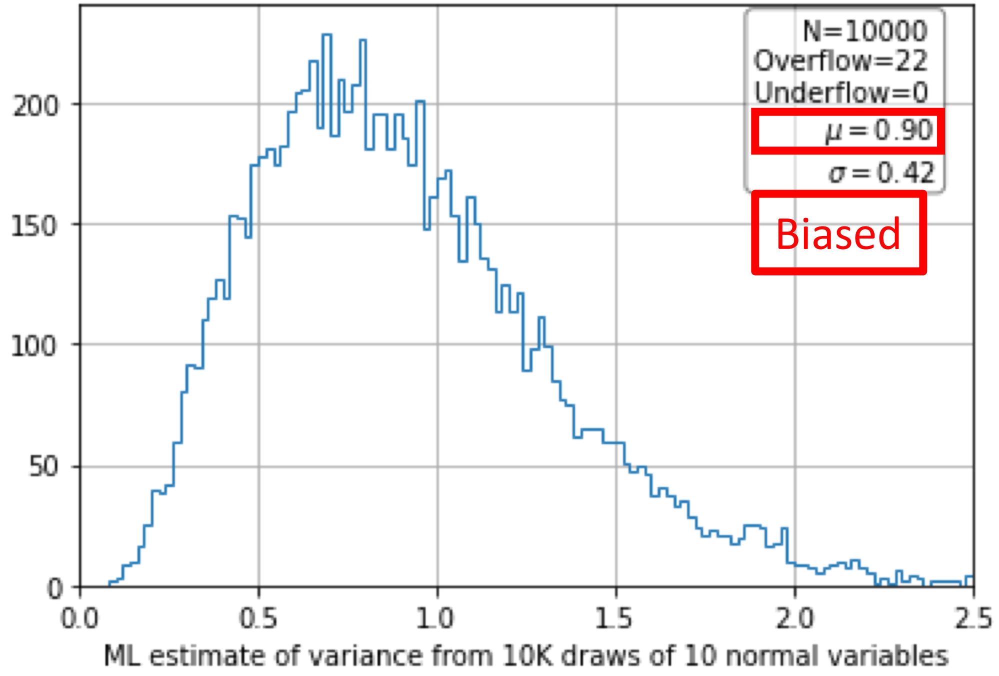

In case you are wondering , the unbiased estimator of σ is

$$\hat{\sigma}^2 = \frac{\sum(x - \overline{x_i})^2}{N-1}$$