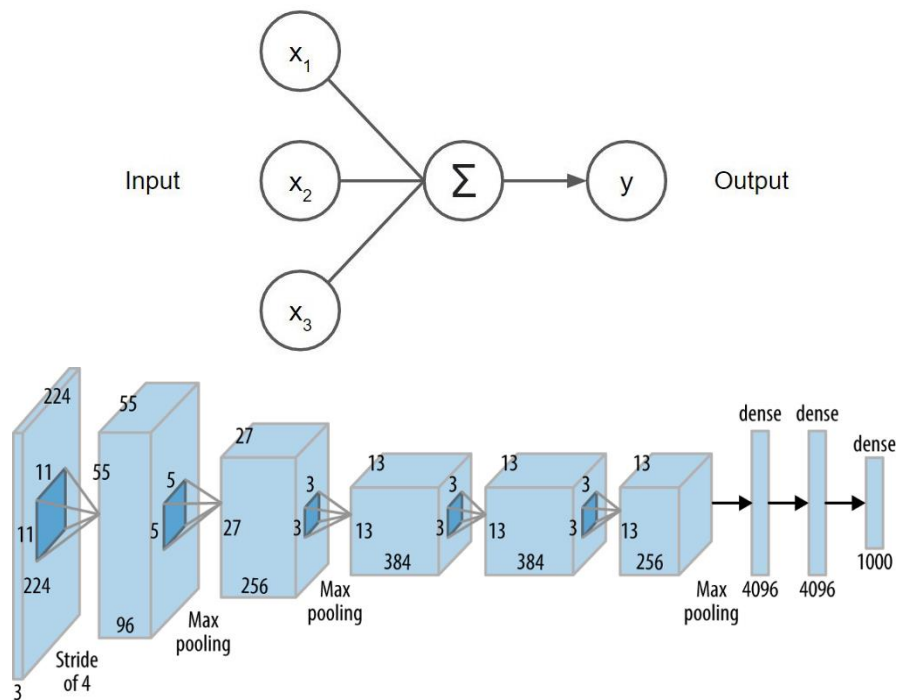
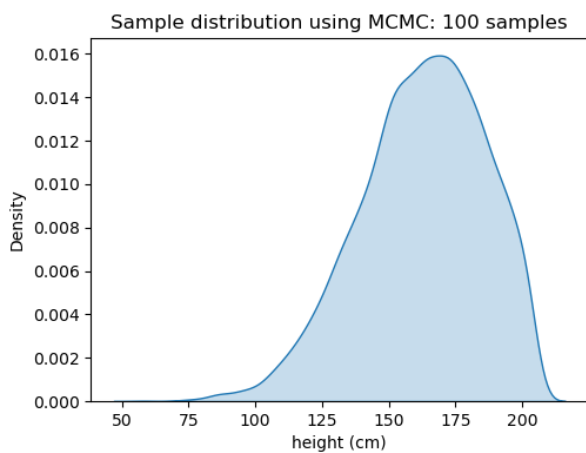












# Chapter 1: Bayesian Inference in the Age of Deep Learning






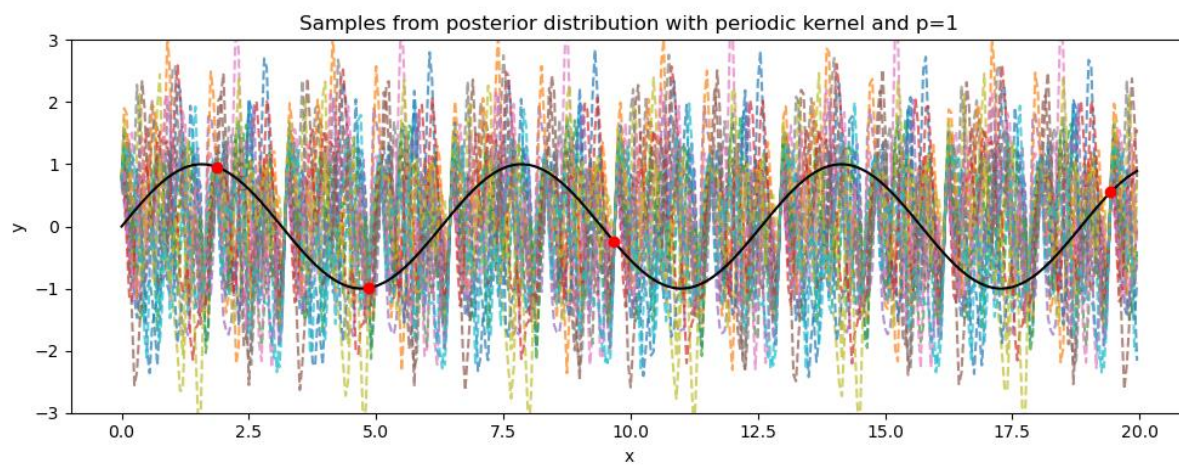


## Chapter 2: Fundamentals of Bayesian Inference

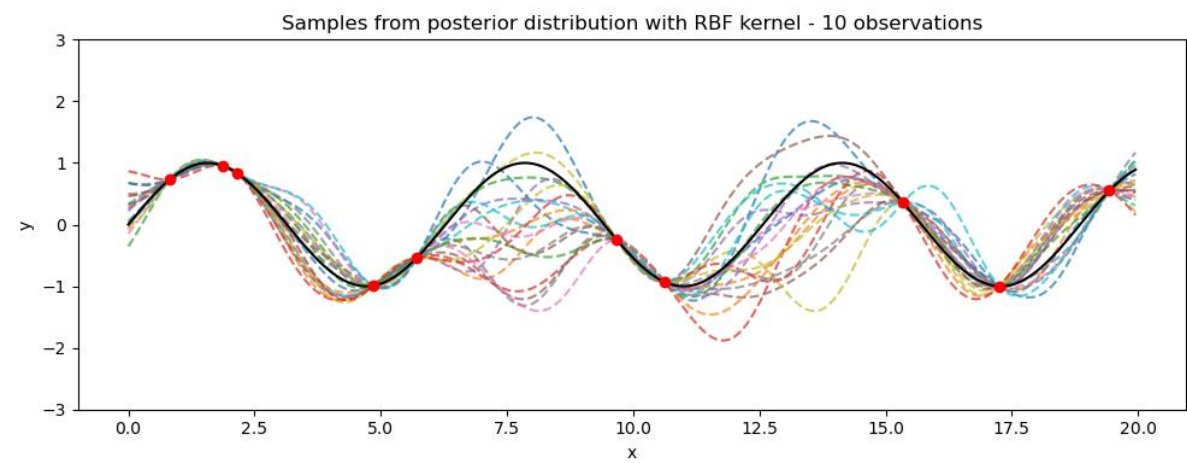
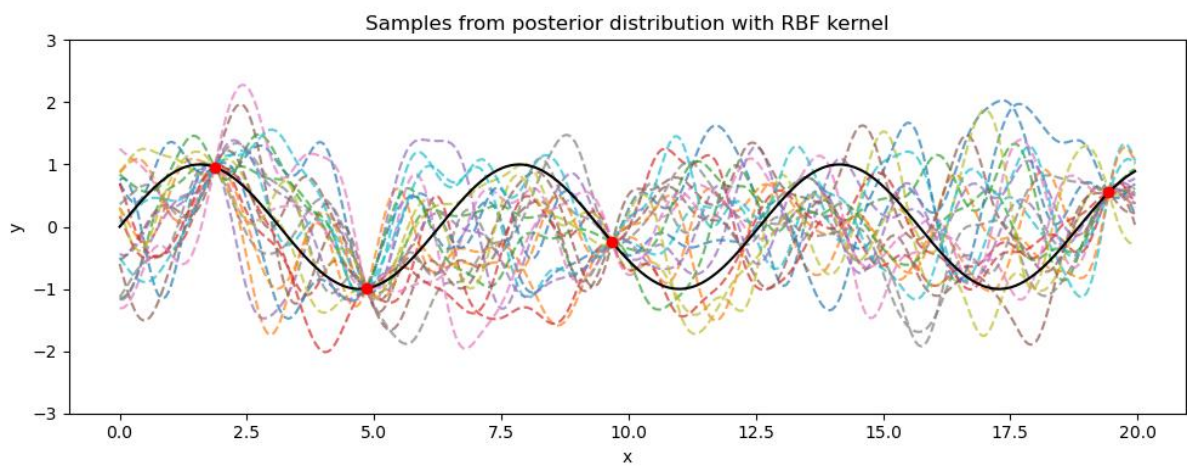

	1	2	3	4	5	6
1	2	3	4	5	6	7
2	3	4	5	6	7	8
3	4	5	6	7	8	9
4	5	6	7	8	9	10
5	6	7	8	9	10	11
6	7	8	9	10	11	12




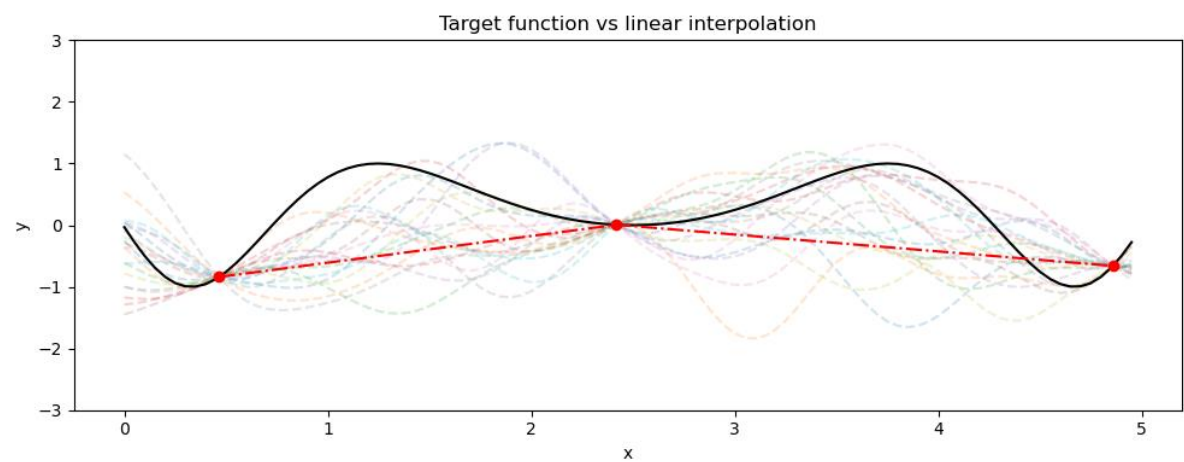
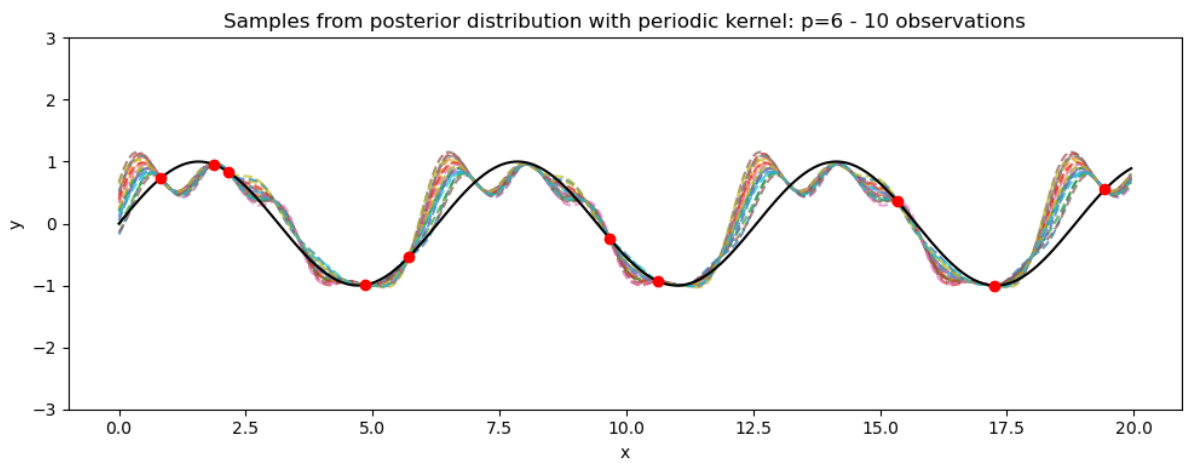











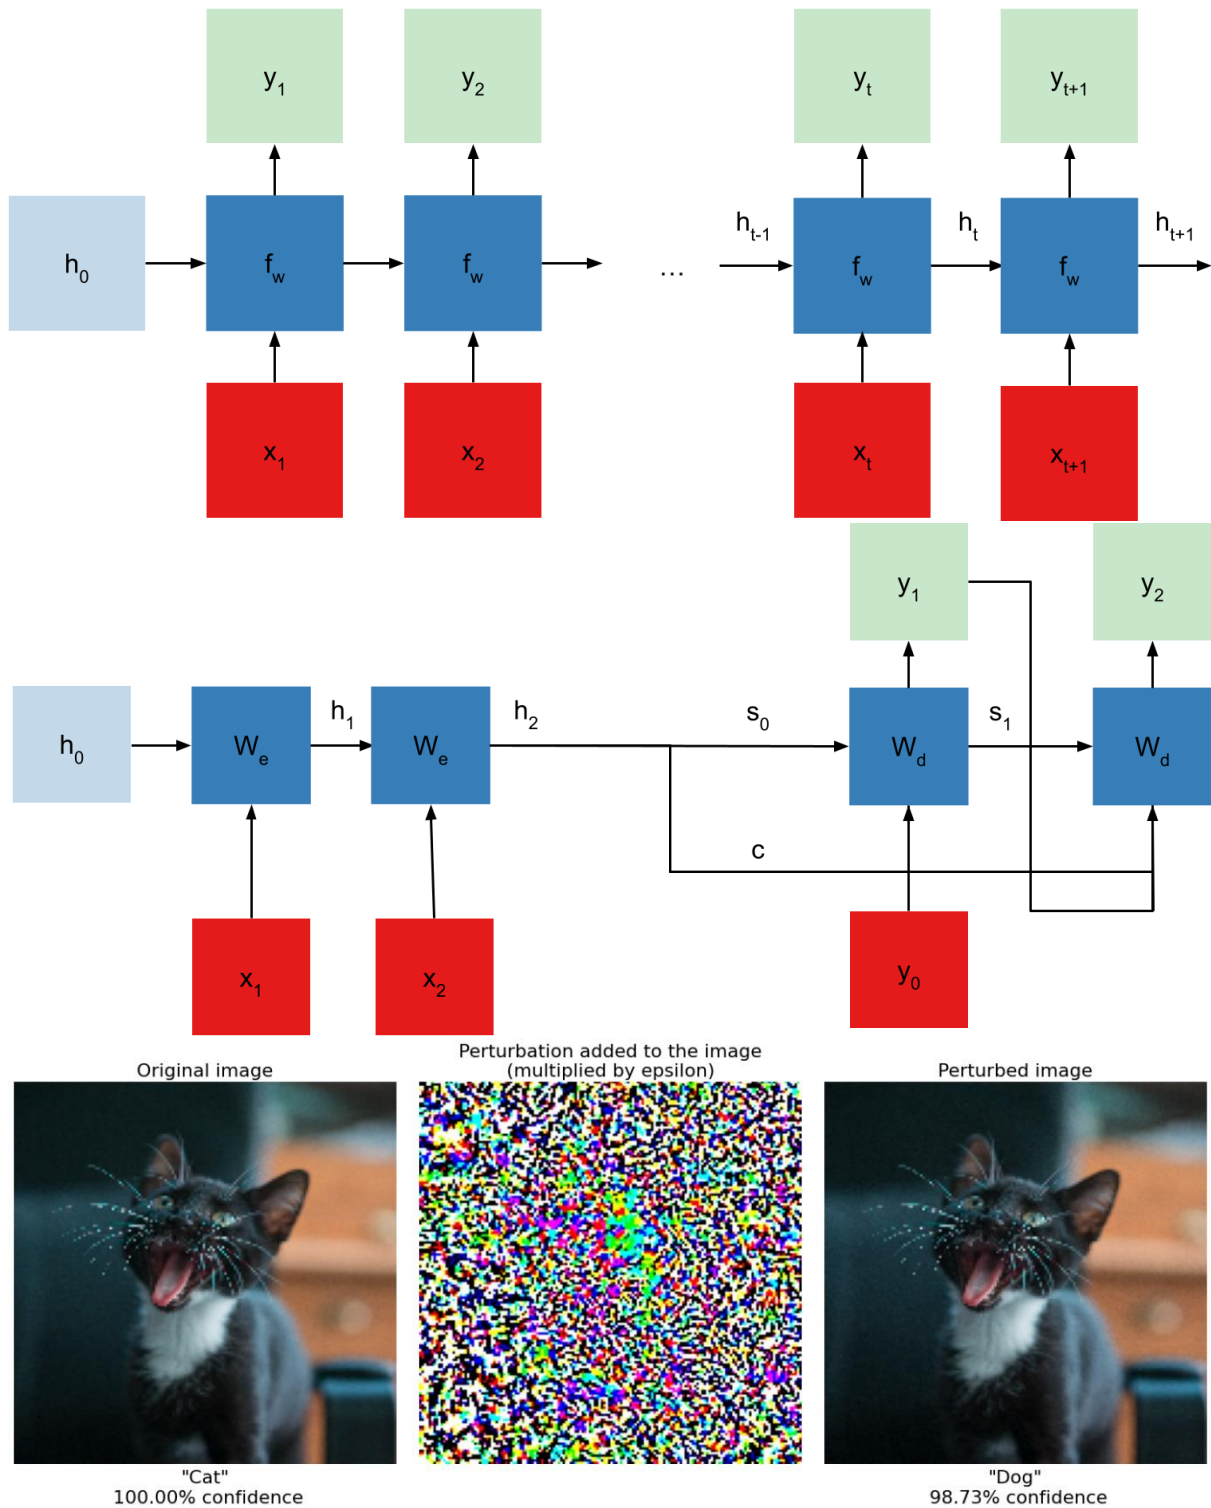






	1	2	3	4	5	6
1	2	3	4	5	6	7
2	3	4	5	6	7	8
3	4	5	6	7	8	9
4	5	6	7	8	9	10
5	6	7	8	9	10	11
6	7	8	9	10	11	12




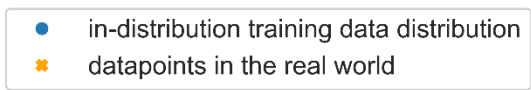



## Chapter 3: Fundamentals of Deep Learning















mean-pooling




max-pooling





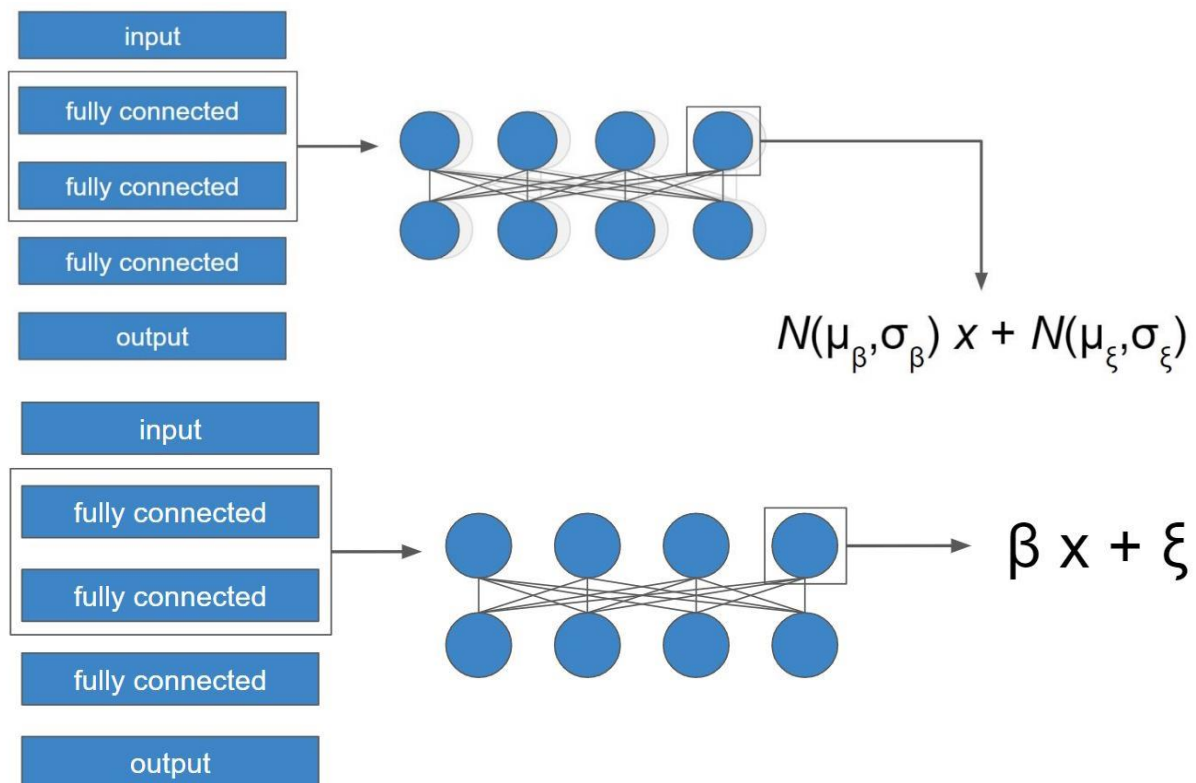
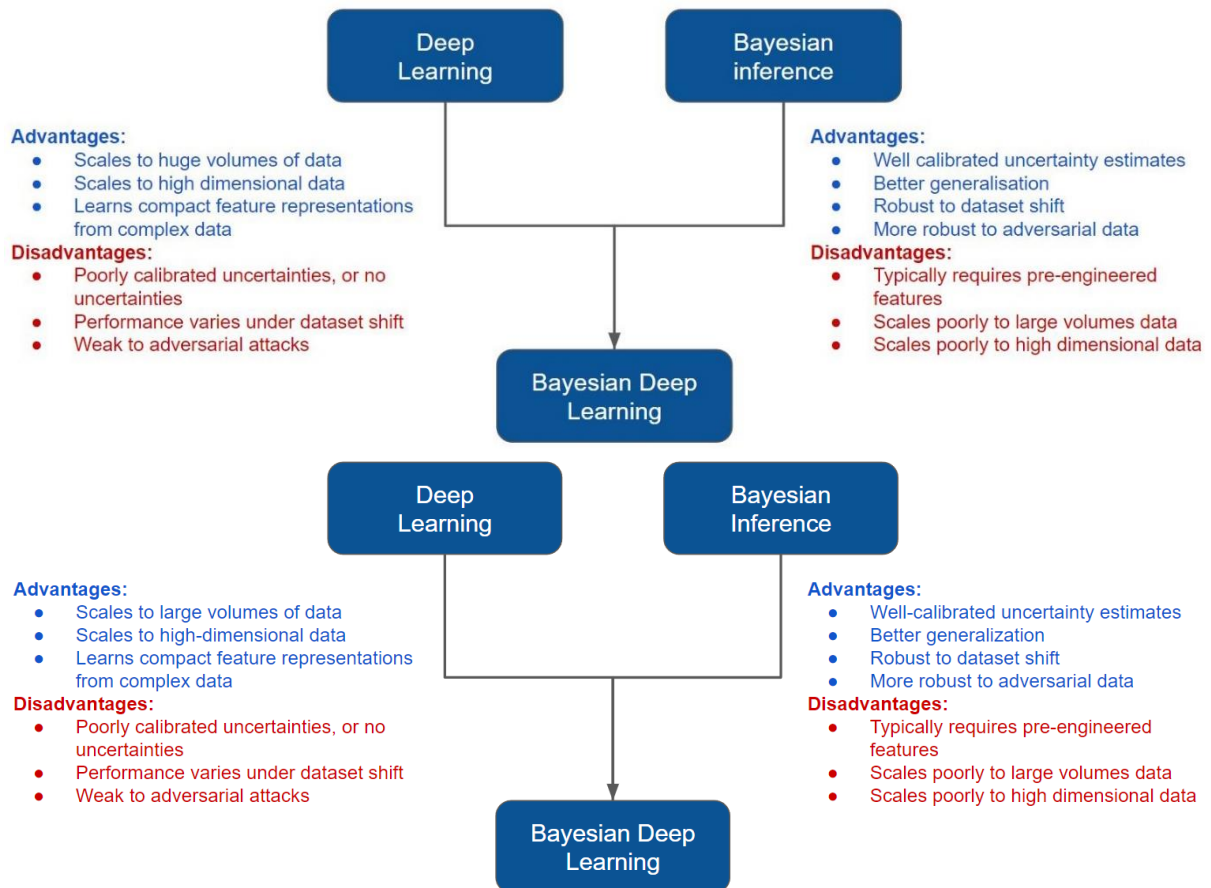












## Chapter 4: Introducing Bayesian Deep Learning














Input value	Ideal parameter	Predicted output value	Target output value	Loss
4.8	5.00	$(4.8 * 4.92) = 23.616$	24.00	0.384
5.1	4.90	$(5.1 * 4.92) = 25.092$	24.99	0.102
4.9	4.80	$(4.9 * 4.92) = 24.108$	23.52	0.588
4.4	4.8	$(4.4 * 4.92) = 21.648$	21.12	0.528
5.2	5.1	$(5.2 * 4.92) = 25.584$	26.52	0.936
Average parameter	4.92		Average loss	0.5076



Input value	Ideal parameter	Predicted output value	Target output value	Error
4.80	5.00	$(4.80 * 4.92) = 23.62$	24.00	0.38
5.10	4.90	$(5.10 * 4.92) = 25.09$	24.99	0.10
4.90	4.80	$(4.90 * 4.92) = 24.11$	23.52	0.59
4.40	4.80	$(4.40 * 4.92) = 21.65$	21.12	0.51
5.20	5.10	$(5.20 * 4.92) = 25.58$	26.52	0.94
Mean	4.92	24.01	24.03	0.51
Standard deviation	0.12	1.37	1.78	0.27

Input value	Ideal parameter	Predicted output value	Target output value	Error
4.80	2.88	$(4.80 * 6.38) = 30.60$	24.00	6.60
5.10	6.40	$(5.10 * 6.38) = 32.52$	24.99	7.53
4.90	3.50	$(4.90 * 6.38) = 31.24$	23.52	7.72
4.40	3.20	$(4.40 * 6.38) = 28.05$	21.12	6.93
5.20	15.90	$(5.20 * 6.38) = 33.16$	26.52	6.64
Mean	6.38	31.11	24.03	7.08
Standard deviation	4.93	1.77	1.78	0.46





## Chapter 5: Principled Approaches for Bayesian Deep Learning

Encoder




$x$

Encoding



$z = f_e(x)$

Decoder  
Encoder



$\hat{x} = f_d(z)$

$x$

Encoding



$z \in \mathbb{R}^n$

Decoder



$\hat{x} = f_d(z)$

$x$

Encoder

$\mu$

$\sigma$

Encoding

$$z \approx N(\mu, \sigma)$$

Decoder

$$\hat{x} = f_d(z)$$

$x$

Encoder

$\varepsilon$

$\mu$

$\sigma$

$$\varepsilon \approx N(0, 1)$$

Encoding

$$z = \mu + \sigma \odot \varepsilon$$

Decoder

$$\hat{x} = f_d(z)$$

$x$

$\mu$

$\sigma$

$\mu$

$\sigma$

$\varepsilon$

$\mu$

$\sigma$


$\mu$

$\sigma$


$$\varepsilon \approx N(0, 1)$$

$$\hat{y} = \mu + \sigma \odot \varepsilon$$

## Chapter 6: Using the Standard Toolbox for Bayesian Deep Learning









Initial network without dropout



Inference 1




Inference 2

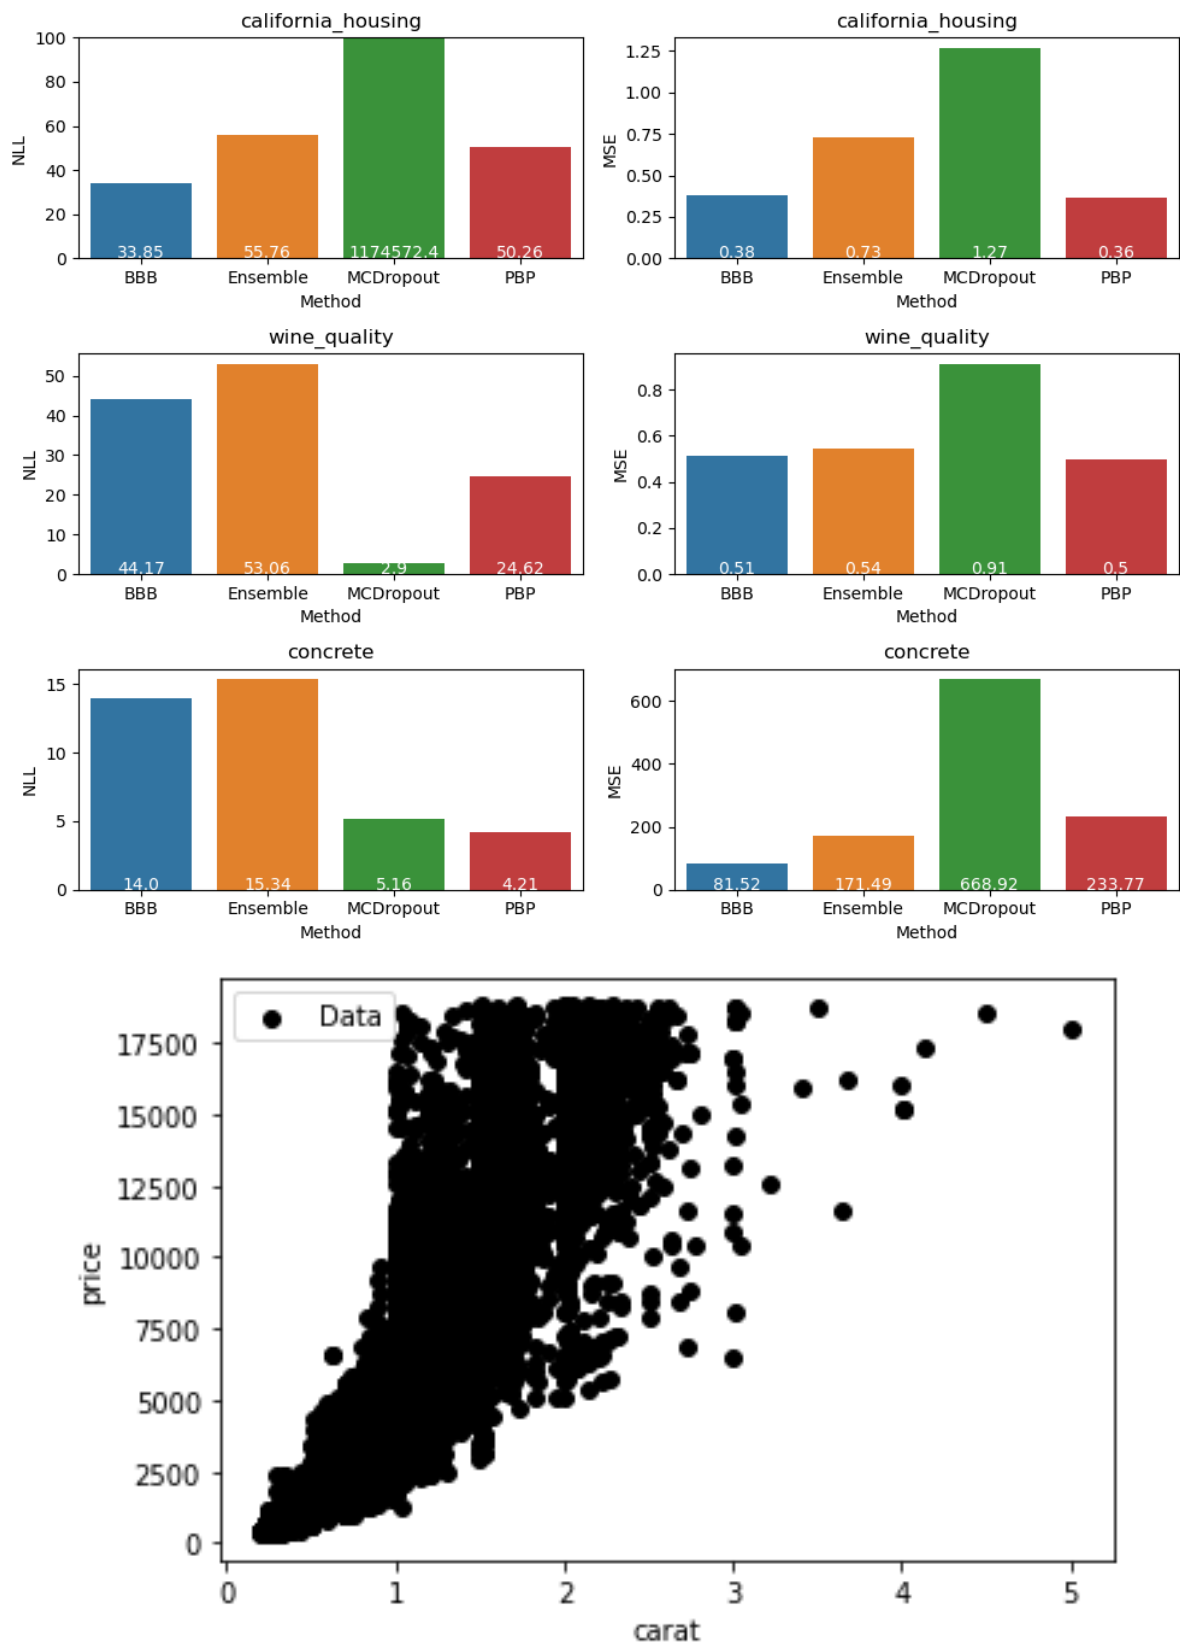





Inference 3









## Chapter 7: Practical Considerations for Bayesian Deep Learning















Selection start	246,195,100 ns
Selection extent	3,773,894,659 ns

☐ Start
 ☐ CPU Duration
 ☒ Duration
 ☐ CPU Self Time
 ☐ Self Time



## Tools (8)


memory\_profile

pod\_viewer

tensorflow\_stats

tf\_data\_bottleneck\_analysis

trace\_viewer



## Chapter 8: Applying Bayesian Deep Learning

Original image



Shot noise, level 1



Shot noise, level 5



0



1



2



3



4



5



6




7


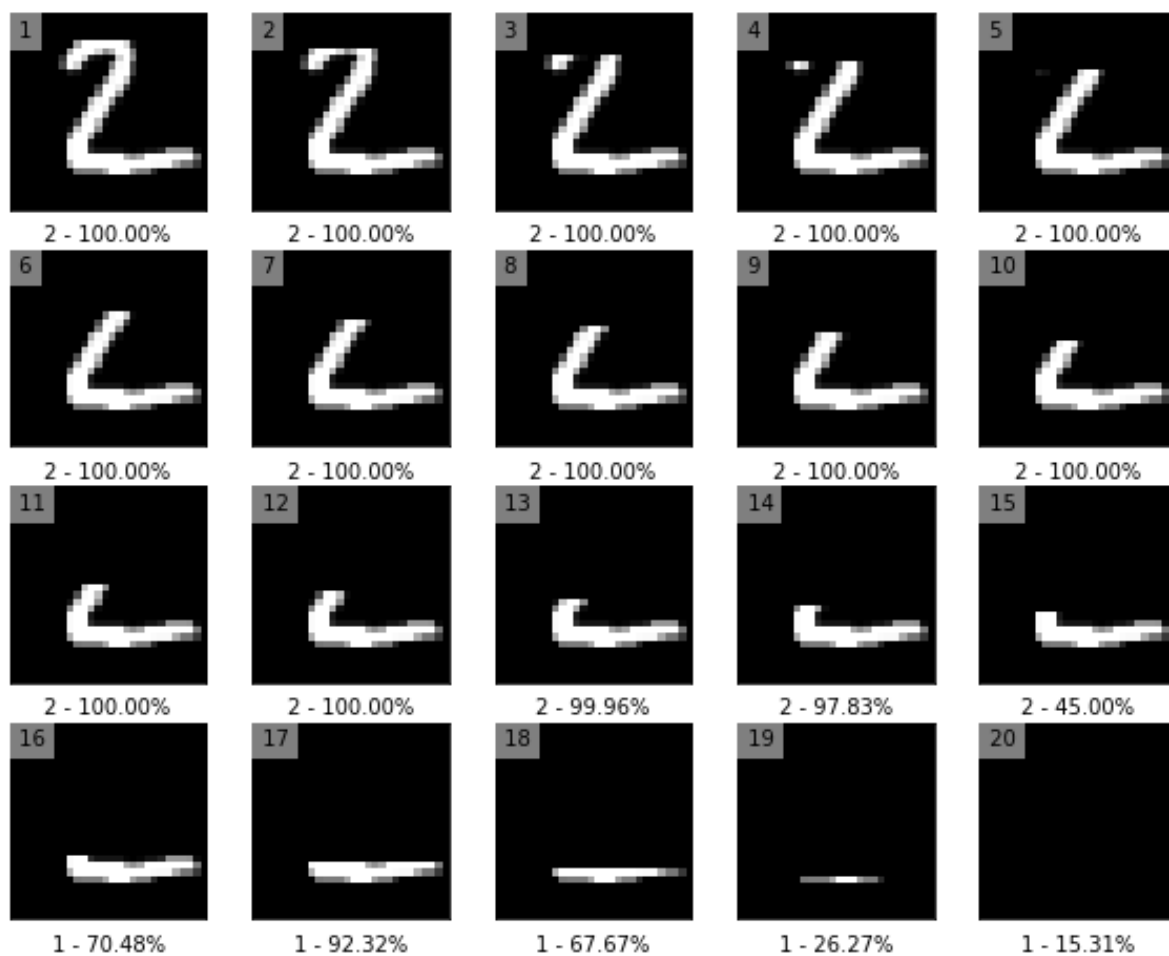



8




9








Starting position





Possible action 1

The agent gets closer to the target, so receives a positive reward.


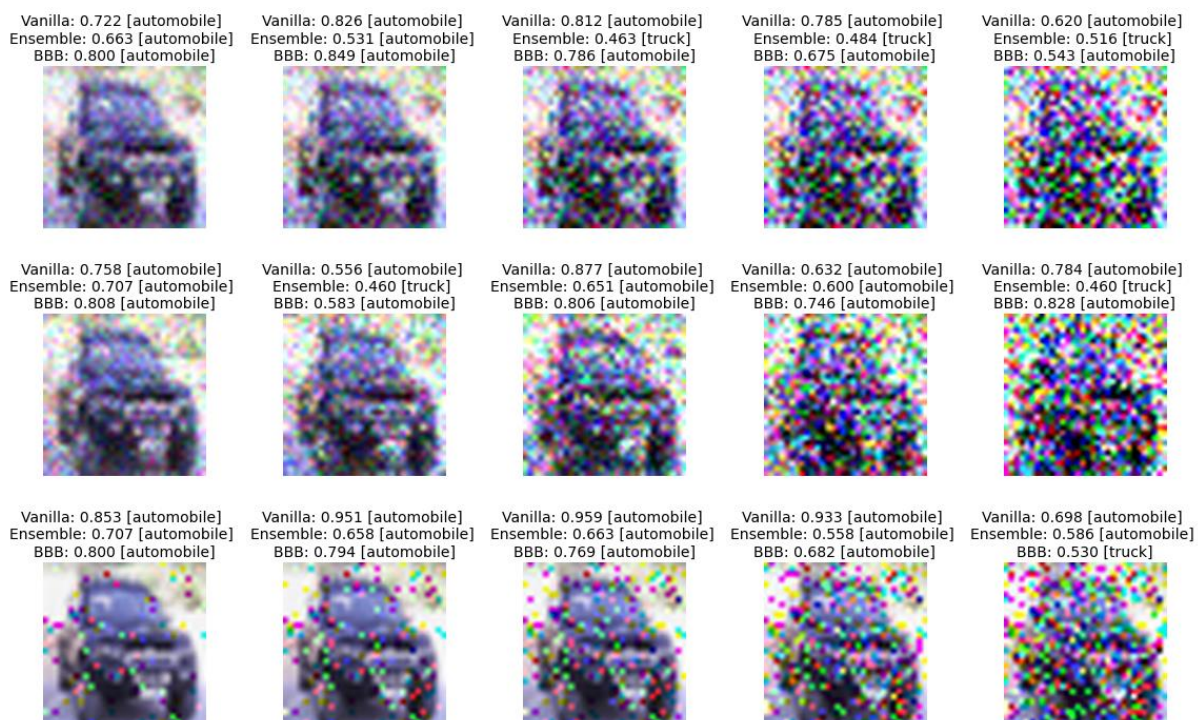


Possible action 2


The agent gets further from the target, so receives a negative reward.



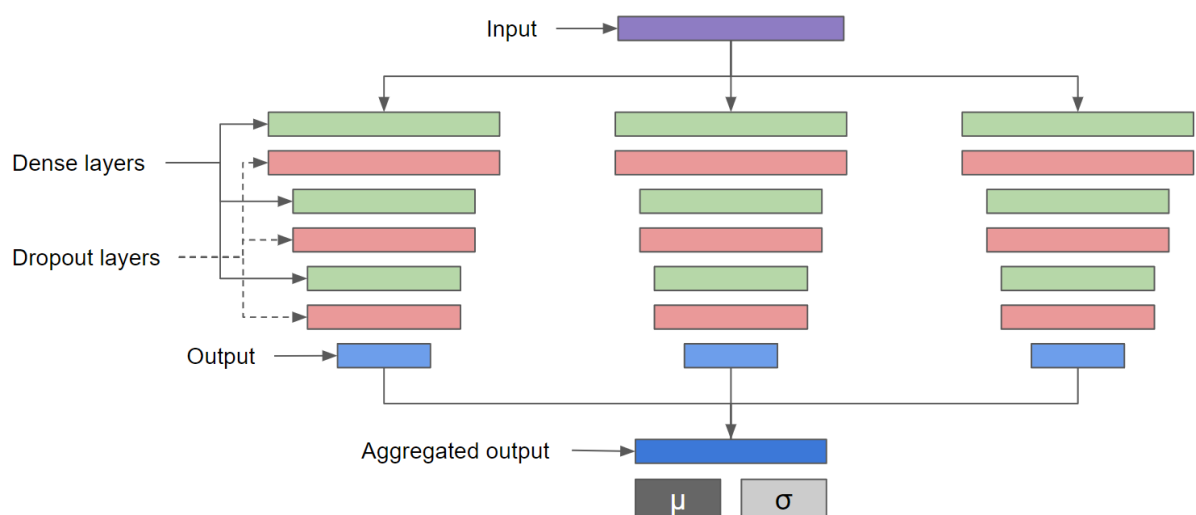













## Chapter 9: Next Steps in Bayesian Deep Learning





- Data point selected for model
- Mean predictions from model

