
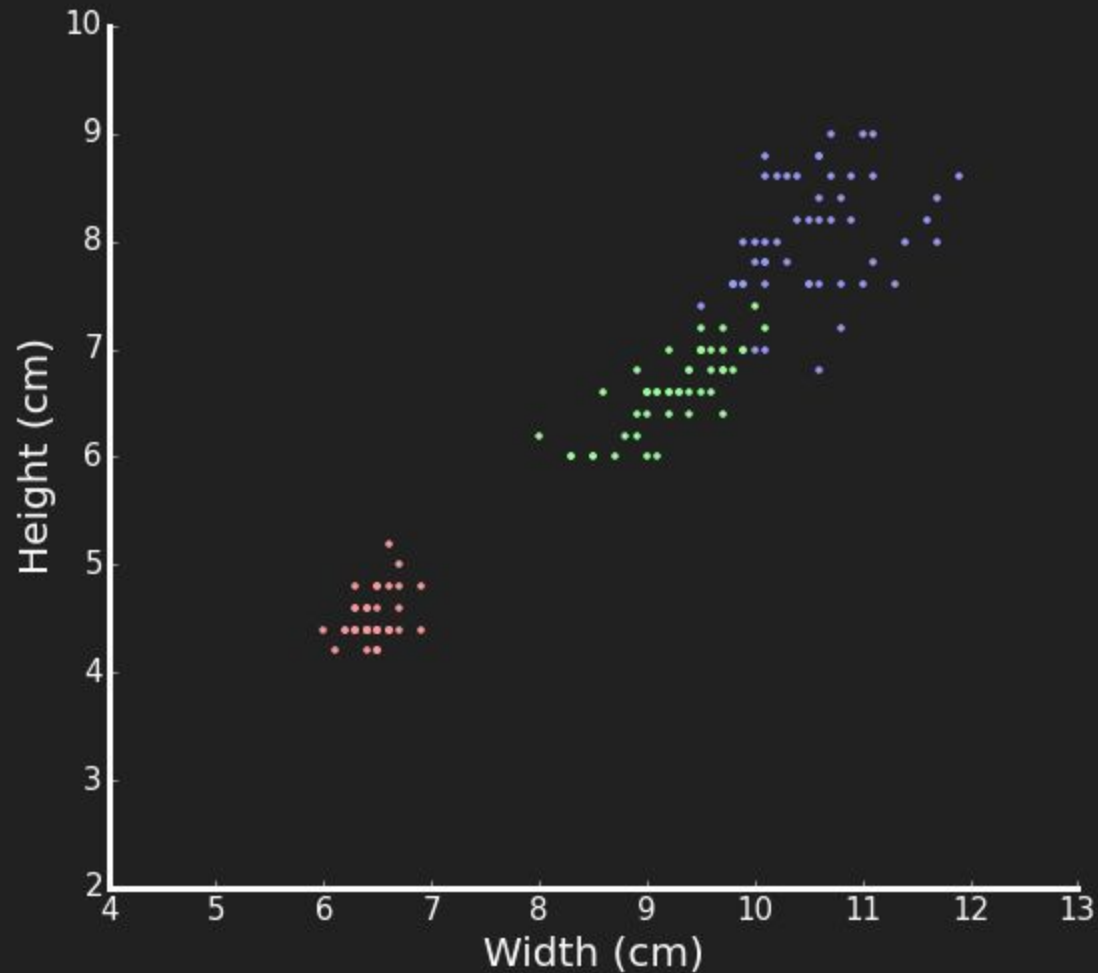


Evaluating models

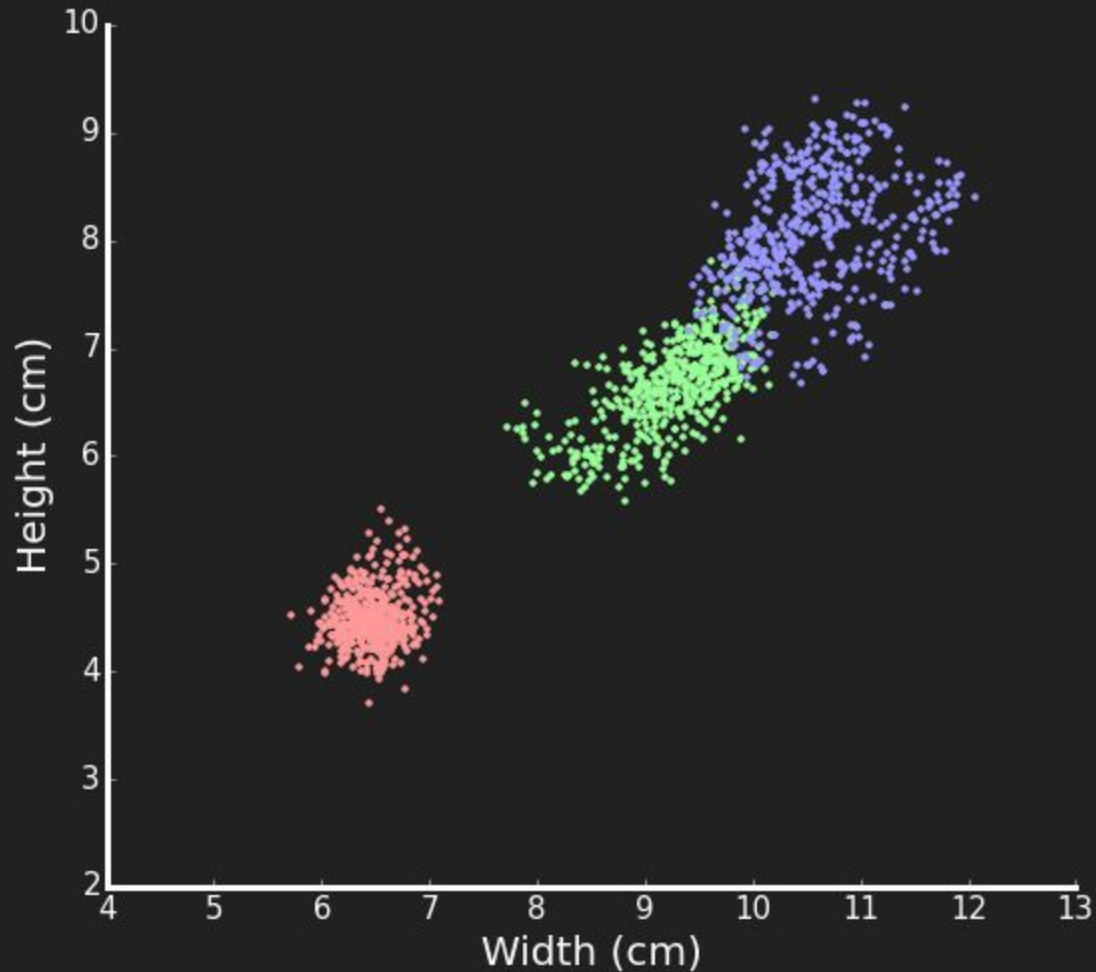
Problems with Nearest Neighbours #1: Overconfidence



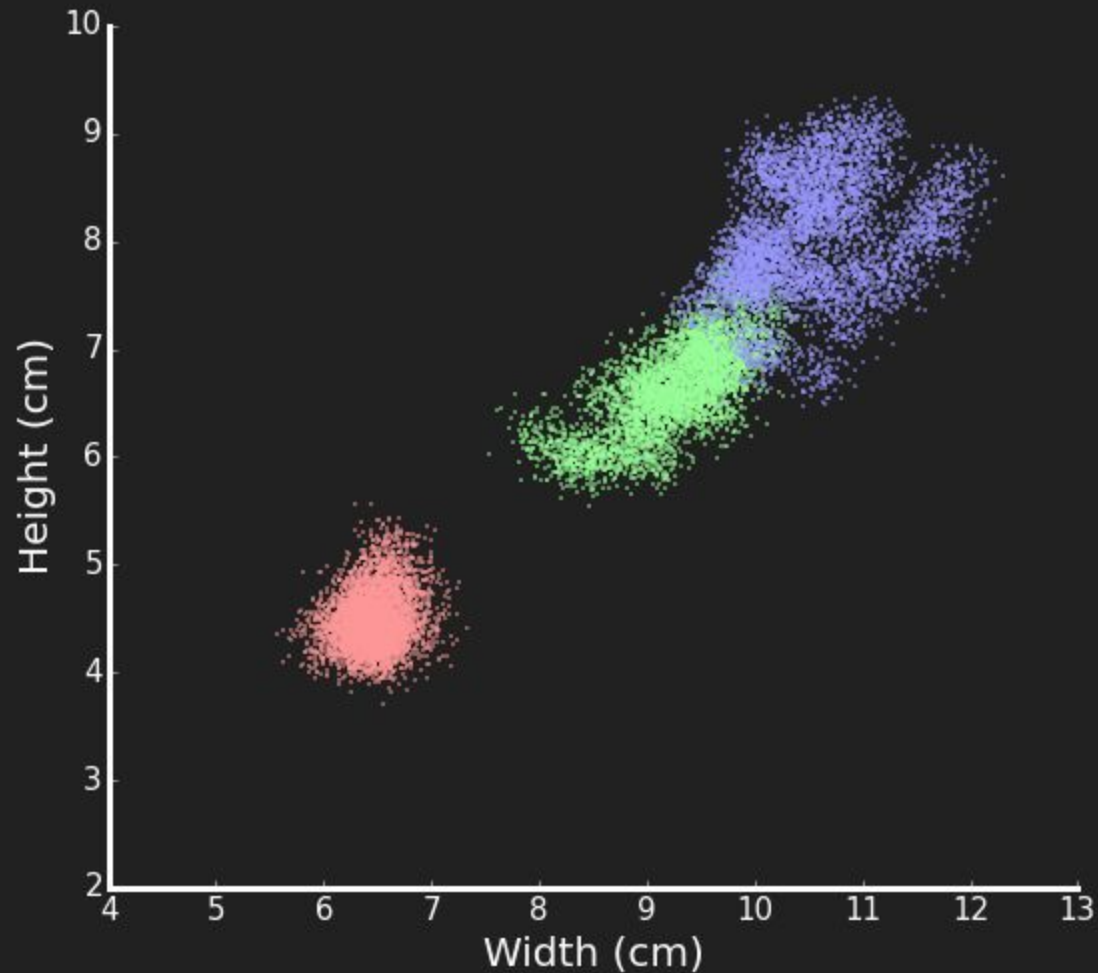
Problems with Nearest Neighbours #2: Memory & speed



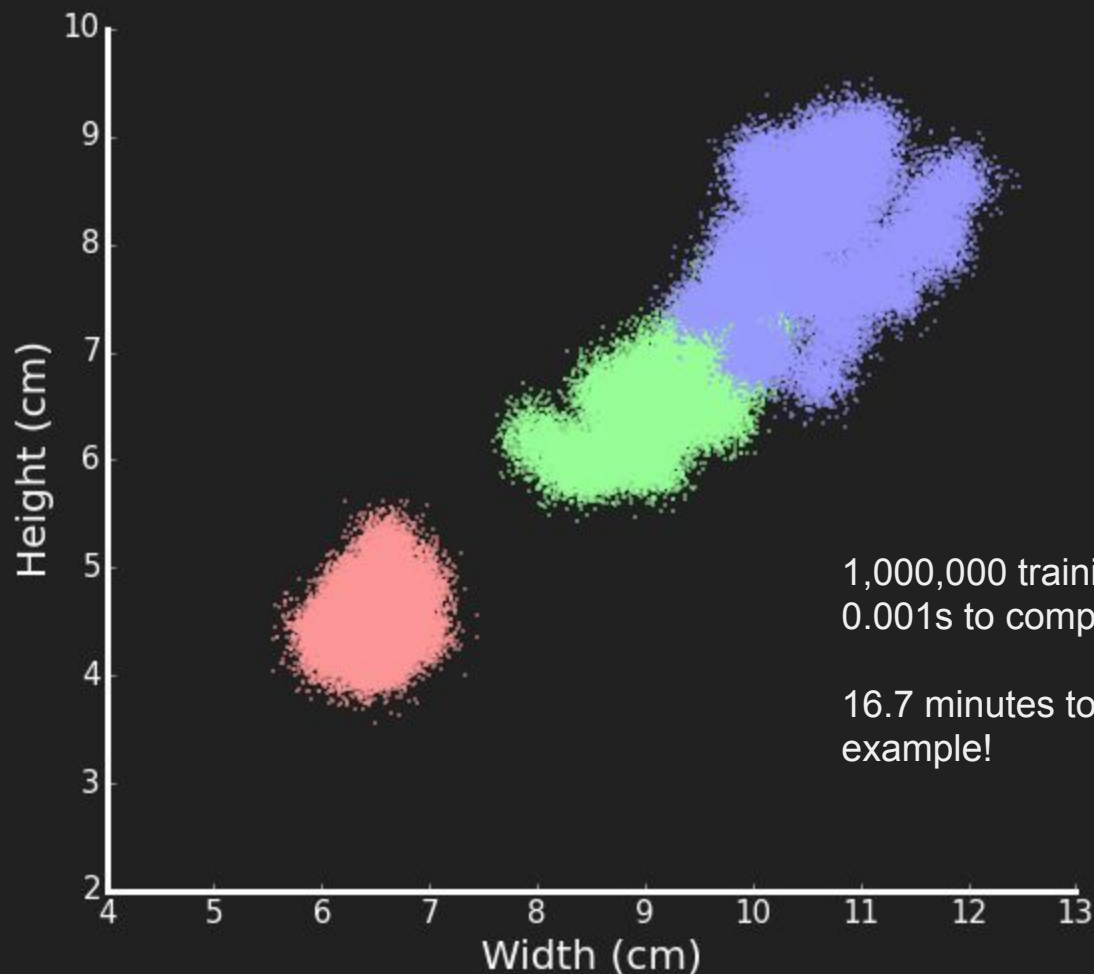
Problems with Nearest Neighbours #2: Memory & speed



Problems with Nearest Neighbours #2: Memory & speed




Problems with Nearest Neighbours #2: Memory & speed




1,000,000 training data
0.001s to compute each distance

16.7 minutes to classify each test
example!


Problems with Nearest Neighbours #3: Noise



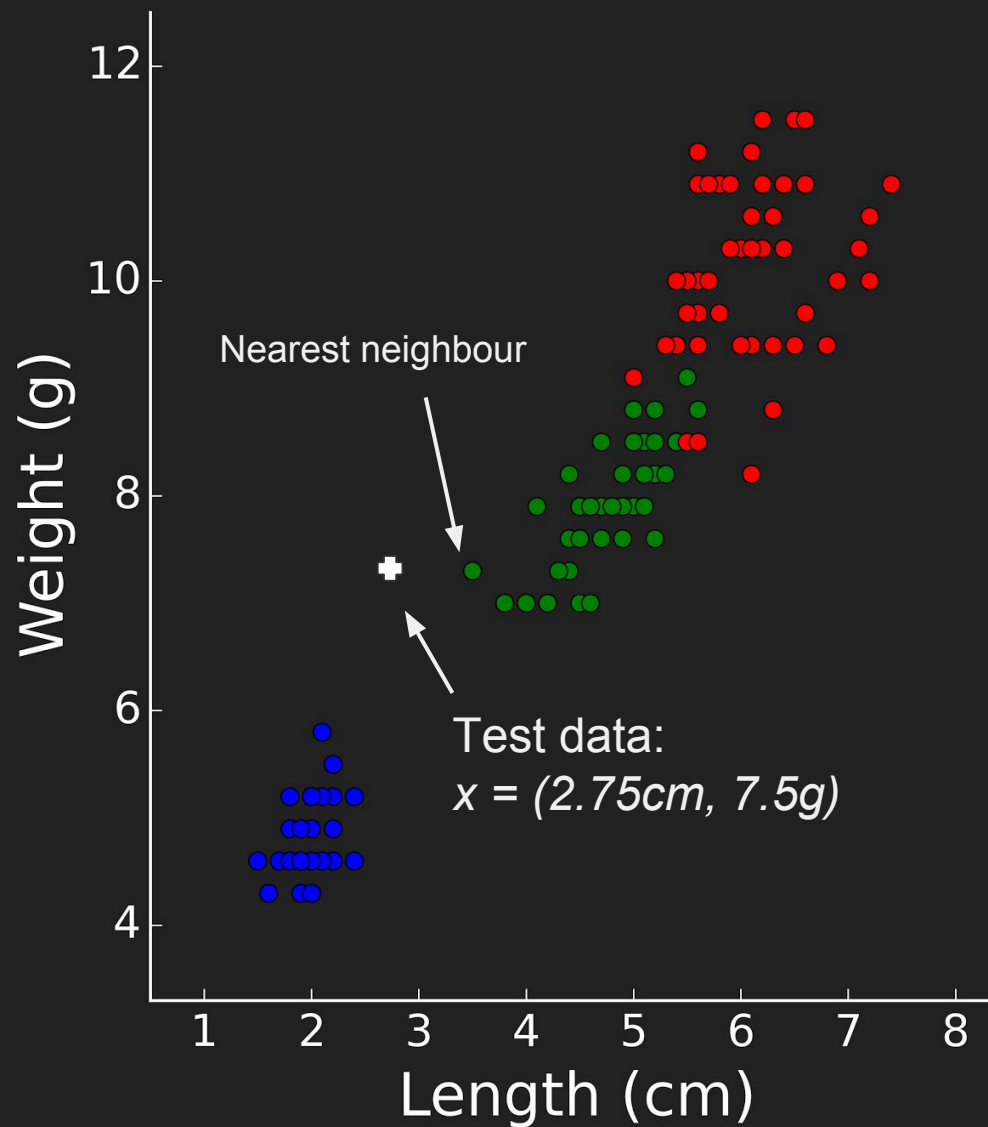
Problems with Nearest Neighbours #3: Noise




Problems with Nearest Neighbours #4: Scaling



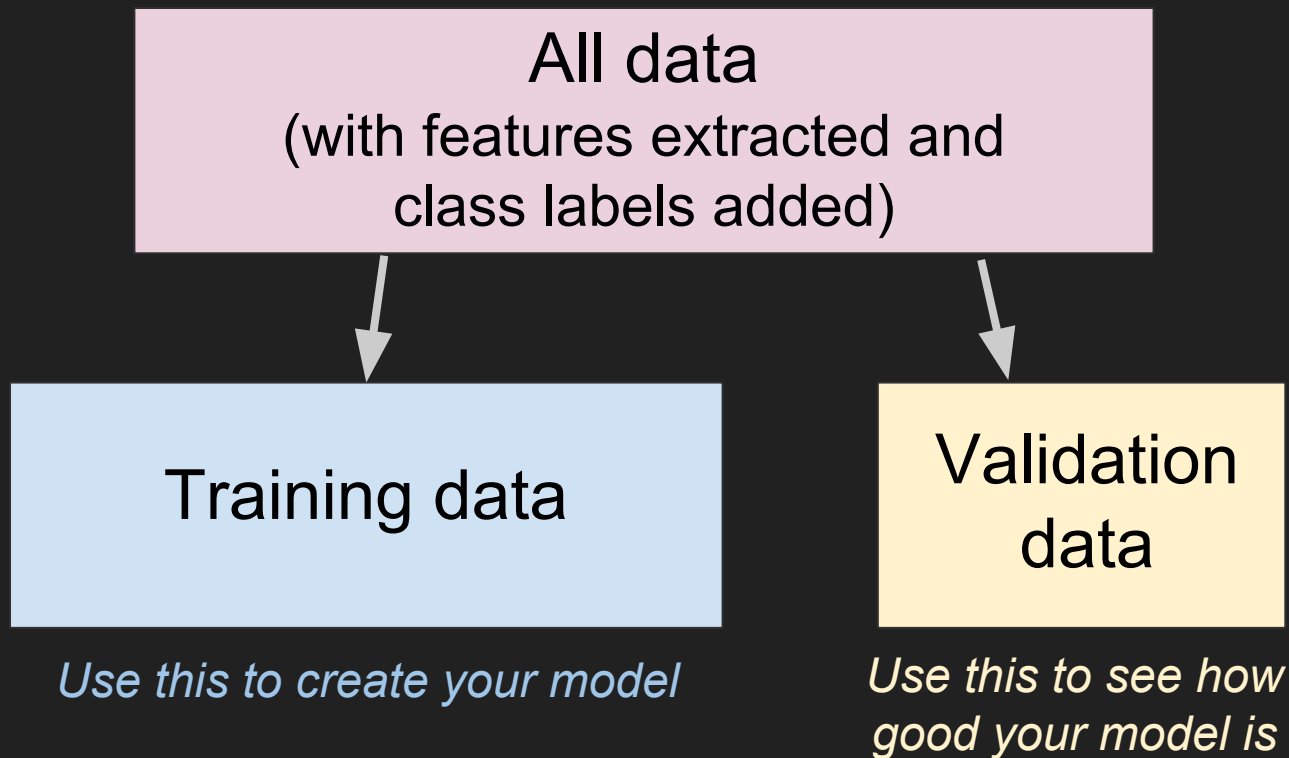
Problems with Nearest Neighbours #4: Scaling



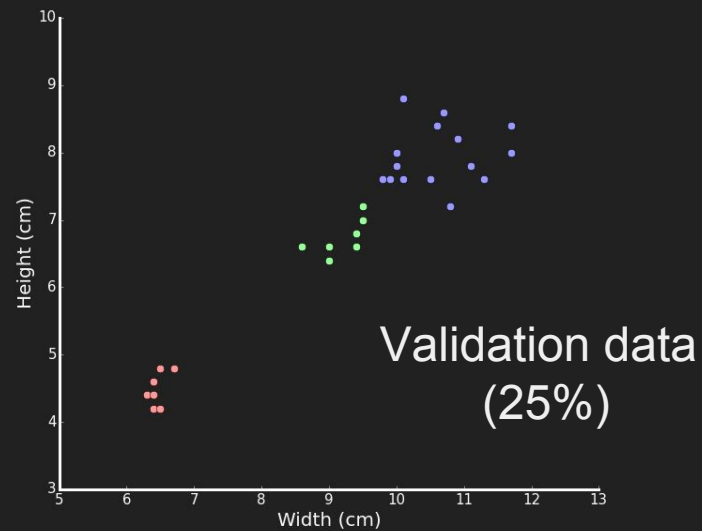
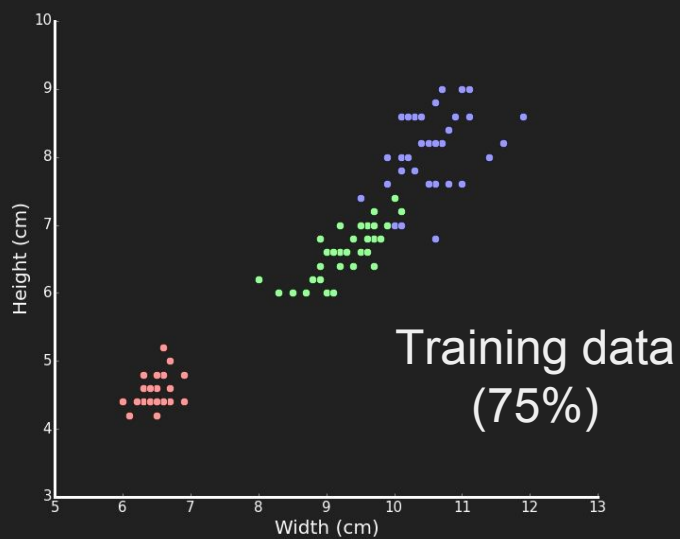

Problems with Nearest Neighbours #4: Scaling



Training / validation split



Training / validation split



How good is your model?

The 'correct' answer

Ground truth
1
1
2
2
1
3
3
1
3
3

1 = polygonia c-album, 2 = maniola jurtina, 3 = pyronia tithonus

How good is your model?

- 10 butterflies in validation set, 6 were classified correctly
- Accuracy = 0.6 or 60%
- Error = 0.4 or 40%