



What you will learn: chance & distributions

Probability for ML • Learning map 1 of 2


1 Outcomes, events & rules




2 Random variables & spread




3 Discrete models




4 Continuous chance



5 Normal distribution



6 Events & waiting




What you will learn: evidence & decisions

Probability for ML • Learning map 2 of 2

1 Joint, conditional & independent


We can look at events together, focus on a subset, and see when events are independent.



Joint • Marginal • Conditional • Independence

2 Bayes & base rates


We update our belief when we see evidence, and the starting rate (base rate) matters.



The starting rate matters.

3 Naive Bayes & smoothing

Use simple features and assume they are independent within each class.




Zero counts can be a problem

Zero counts		Smoothing	
Feature	Count	Feature	Count
"Offer" in not spam	0	"Offer" in not spam	1

Add 1 to BOTH present and absent counts.

5 Decisions & expected cost

Use probabilities and costs to decide what to do.



6 Check probability quality

We can assess if predicted probabilities are sensible and well calibrated.

