


PHASE: UNDERSTAND

Bayesian inference: the whole learning map

Use observations to update the chances assigned to possible explanations.




Why keep several possible explanations?



A few deliveries do not tell us the true on-time chance exactly.

What is probability?

A probability is a number from 0 to 1. It tells how likely something is.



1 List possible on-time chances

What on-time chance could each candidate have?

Candidate A



On-time chance
0.25
(often late)

Candidate B



On-time chance
0.75
(often on time)



These are two possible explanations for the deliveries?

2 Observe deliveries

We see two deliveries, and both arrive on time.



On time



On time



Yes. Now we use these observations to update which explanation is more likely.

3 Update which explanation is more likely

Start with equal chance for each candidate.

A: 0.5

B: 0.5



After seeing two on-time deliveries, update the chances:

A: 0.1

B: 0.9

Candidate B becomes much more likely, because on-time deliveries are expected under B.

4 Predict the next delivery

Use the updated chances to predict the next delivery.

$$P(\text{on time}) = 0.1 \times 0.25 + 0.9 \times 0.75 = 0.70$$



On time
70% chance



Late
30% chance

A simple example: how we update our guess

1 Start: equal chances

A
50%

B
50%

We start with the same chance for both candidates.

2 Observe: two on time



On time



On time

We see two deliveries, and both arrive on time.

3 Compare: more likely under B



Two on-time deliveries are much more likely if candidate B is correct.

4 Update: A 10%, B 90%

A
10%

B
90%

We update our chances based on what we saw.

5 Predict: 70% on time



On time
70% chance


Late
30% chance

A complete learning map

From a simple example to bigger ideas. Each topic answers a more complex question.


1 Bayesian updates

Use new evidence to update chances.



2 Graphical models

Draw which quantities depend on which.



3 Markov chains

The next state uses the current state.

D: Dry

W: Wet


4 Markov chain Monte Carlo (MCMC)

Sample possible model settings using a Markov chain.




5 Variational inference

Fit a simpler probability description.




6 Hidden Markov model (HMM)

Infer a hidden state from readings.



7 Gaussian process

Predict a curve with uncertainty.



8 Diffusion models

Learn to reverse added noise.



9 Decisions

Compare the average costs of possible choices.

