


## Predicting house prices

To find the relation between  
 $X$  and  $Y$ , we need  
 $m$  and  $b$



Formulae:

$$b = \frac{\sum y \sum x^2 - \sum x \sum xy}{n \sum x^2 - (\sum x)^2}$$

$$m = \frac{n \sum xy - \sum x \sum y}{n \sum x^2 - (\sum x)^2}$$

"Regression"

25

"point estimates"

1.9



$x$	$y$	$xy$	$x^2$	$y^2$
...	...	...	...	...
$\Sigma =$				

We do something very similar in machine learning!

But!

$m$  and  $b$  are both "point estimates".

They say nothing about the uncertainty!



↗  
 The problem is  $\{m_1, b_1\}$  here and  $\{m_2, b_2\}$  here  
 are equally "confident".

This is a major problem in classical statistics!

Model parameters are learned through observations alone  
and they are point estimates!


(multiverse shenanigans)  
 x-square  
 t-testing

So, what's the solution?

Bayesian inference!


- Start with a prior distribution
- Update prior to posterior distribution based on evidence

Let's ignore the slope for now and focus on "b"





Our data will eventually look like this



Ground truth




$f(b)$




$f(b)$




$f(b)$





$f(b)$



Notice that :

"b" still has a distribution. We have the uncertainty quantified — at all times!