Probabilistic Fitting

- Generative probabilistic model
  - Tells a story about how stochastic data comes to be
  - Darts fall around the center of the board, but where exactly?
  - Consider a model with parameters, $\boldsymbol{\theta}$
  - Consider an observation, x_i
  - We denote the probability of seeing x_i under the model by:

$$p(x_i \mid \Theta)$$

Read "given" or "conditioned on"
Restricts to the case of $\boldsymbol{\theta}$

Defined by $P(A \mid B) = \frac{P(A,B)}{P(B)}$

Probabilistic Fitting

- Multiple observations
  - Suppose we have multiple observations, in a vector $\mathbf{x}$
  - What is the probability of $\mathbf{x}$?

Probabilistic Fitting

- Multiple observations
  - Suppose we have multiple observations, in a vector $\mathbf{x}$
  - What is the probability of $\mathbf{x}$?
- If observations are independent then probability is the product of the individual observations
  - Essentially a definition, but is consistent with intuition
  - The observations are conditionally independent **given** the model
- So, the probability of x is then:

$$p(\mathbf{x} \mid \Theta) = \prod p(x_i \mid \Theta)$$

Probabilistic Fitting

- So, given the model, we have the probability of observing the data

$$p(\mathbf{x} \mid \Theta) = \prod p(x_i \mid \Theta)$$

- But what we really want is the probability of the model (parameters) given the data!
- Bayes rule comes to the rescue!