

Assignment 2: Bayesian Generalization
Human and Machine Learning
Chiba Institute of Technology, School of Design & Science
Prof. Joseph Austerweil
Due Fri Jun 19, 2026 at 8:00pm

Many of the problems require making graphs. Please do not forget to label your axes and title the graph! You may plot multiple distributions on the same graph **as long as each distribution is clearly legible.**

Bayesian Generalization Framework. For this assignment, you will create your own Bayesian generalization model for six animals: Cow, Dolphin, Chicken, Seal, Penguin, and Bat. Note that there is no single correct hypothesis space! Just think of the properties shared by these animals and turn them into hypotheses (the subset of animals that have the property).

Remember that the Bayesian generalization model is trying to solve the following problem: given you observe some animal(s) that have a novel property, how likely are other animals to have that novel property? For example, if you learn that a seal has sesamoid bones, is it more likely that a cow or a dolphin also has sesamoid bones? This is the predictive probability (what you solve in problem 4). The model solves it by maintaining a set of hypothesized properties (the hypothesis space), computing how probable it is that each one of them is the novel property (the posterior probability), and then generalizing to the other animals based on the posterior probability of each hypothesis containing the other animal and the observed animals (the predictive probability). Note that “hypothesis” and “property” are used interchangeably, and that “observing an animal” is shorthand for “observing that an animal has the novel property of interest”. Also note that you are not generalizing to any animals other than the six animals named above.

1. *Defining a hypothesis space.* Write your hypotheses. For each one (or ones of the same type), have a 1–4 word description of the hypothesis. Remember that a hypothesis is a binary vector of length equal to the number of animals (so six in this case). An element of the vector is 1 if the corresponding animal has the hypothesized property and 0 if the animal does not. You can write this out in matrix form or each hypothesis individually. Make sure you include a “catch-all” hypothesis that contains every animal. You should have more than four and fewer than 63 hypotheses.
2. *Prior.* Define a prior distribution over hypotheses (a uniform one is fine). Write a 1–2 sentence justification of why you chose this prior.
3. *Posterior.* Write code that calculates the posterior probability of the hypotheses, given you know one of the animals has that property. Do this twice, once for weak sampling and once for strong sampling.

Remember that the posterior probability of a hypothesis h in hypothesis space \mathcal{H} given a set of animals \mathbf{x} that are known to be in the hypothesis, $P(h|\mathbf{x})$, is:

$$P(h|\mathbf{x}) = \frac{P(h)P(\mathbf{x}|h)}{\sum_{h' \in \mathcal{H}} P(h')P(\mathbf{x}|h')} = \frac{P(h) \prod_{n=1}^N P(x_n|h)}{\sum_{h' \in \mathcal{H}} P(h') \prod_{n=1}^N P(x_n|h')}. \quad (1)$$

Weak sampling makes minimal assumptions as to how the data are generated and merely assigns a binary label to each data point: $P(x|h) = 1$ if $x \in h$ and $P(x|h) = 0$ if $x \notin h$.

Strong sampling makes a stronger assumption — that each data point is generated uniformly at random from the set of data points consistent with the hypothesis. So, according to strong sampling, $P(x|h) = 1/|h|$ if $x \in h$ (and $P(x|h) = 0$ when $x \notin h$), where $|h|$ is the “size” of the hypothesis (the number of $x \in h$).

- (a) For a representative example, describe how the probability distribution over hypotheses changes after observing that a single animal has the property of interest (1–2 sentences). So, the animal is \mathbf{x} . Are there any differences between weak and strong sampling? If so, what are they?
 - (b) Update your code so that it computes the posterior probability of the hypotheses given you know more than one animal has a property. Add two more animals to the single animal and have the model observe that the three animals have a property. Describe how the posterior probability has changed (1–2 sentences) and how that differs between the weak and strong sampling models.
4. *Prediction.* Write code that calculates the predictive distribution: the probability that an animal has a property given a set of animals that you have observed. Recall from class that the predictive distribution is the updated probability distribution over data points after learning a subset of the data points share some property. Formally, the probability of the n -th data point x_n after observing $n - 1$ data points x_1, \dots, x_{n-1} is

$$P(x_n \text{ has property} \mid x_1, \dots, x_{n-1} \text{ have property}) = \sum_{h: x_n \in h} P(h \mid x_1, \dots, x_{n-1}).$$

¹ Note that $\sum_{h: x_n \in h}$ means the sum is only over those hypotheses that contain x_n (and not the entire hypothesis space). Plot the predictive distribution (histogram preferred) for the six animals given the one and multiple observations you did previously for both weak and strong sampling (four total histograms). Remember to label your axes and title your graph! Write a paragraph describing the results.²

5. *Break your model.* Update your hypothesis space so that it includes all possible hypotheses over six animals (all $2^6 - 1 = 63$ binary vectors of length 6 — one less because the all-zeros hypothesis does not do anything). What happens to the posterior and predictive probabilities that you calculated in problems 3 and 4? Write a few sentences describing what happened and why. Relate this to a theorem we covered in class. Note that you do not need to include every single probability calculation in this new hypothesis space, but rather only those necessary for your written explanation.

Stencils. Three starter stencils are provided in this directory; pick one and submit your completed version:

¹Note this is the probability that animal x_n has the property. There are two options — it either does or does not have the property. This is NOT the same as predicting the next animal you will observe having the property (that would normalize over the different animals rather than yes/no).

²Note that you are calculating the probability that an animal has the property given the animals you have observed to have the property, not that it will be the next animal you observe to have the property.

- `generalization.ipynb` — GenJAX (canonical, runs in Google Colab).
- `generalization_python.ipynb` — numpy + matplotlib (Python without GenJAX).
- `generalization_nosoln.Rmd` — base R + ggplot2.

A Matlab stencil is available on request — email Prof. Austerweil.

Prep reading. A textbook chapter on Bayesian generalization is in preparation and will be added to Tutorial 3 of the [Probability Tutorial](#); the lecture slides from Week 4 are the canonical reference until then.

Submission. Submit by DM or email to the instructor *one* of the following:

- your completed notebook (or knitted `.Rmd`) — it must run end-to-end and contain your figures, inline text answers, and descriptions; *or*
- a single PDF report containing your code, figures, text answers, and descriptions.