CNeuro2024

HANG ZHANG

Problem Set: *Bayesian Observers for "Suit Pursuit"*

Alice and Bob are playing a simplified card game called "Suit Pursuit." In each round, Alice secretly selects one of four suits: club, diamond, heart, or spade. Bob's goal is to guess the suit that Alice has chosen. Unknown to Bob, Alice follows a specific probabilistic strategy for selecting the suits.

Game Rules:

- Alice chooses a suit in each round according to a probabilistic rule.
- Bob tries to guess the suit based on his observations and inferential strategies.
- The suits are mapped to numbers for convenience: club (0), diamond (1), heart (2), spade (3).

Alice's Probabilistic Strategy:

Alice's suit selection depends on her previous choice. If her previous choice was suit x, she selects the suit (x+3)%4 with a probability of 0.55 and the other suits with a probability of 0.15 each. The notation (x+3)%4 represents the modulo operation, which returns the remainder when x+3 is divided by 4.

For example, if her previous choice was heart (2), she will choose diamond (1) with a probability of 0.55 because (2+3)%4 = 1. She will choose club (0), heart (2), or spade (3) with a probability of 0.15 each.

Problem 1: Visualizing Alice's Strategy

Create a transition matrix that represents Alice's probabilistic strategy. The matrix should show the probabilities of moving from one suit to another in consecutive rounds. Visualize the matrix using a heatmap or a suitable graph.

Problem 2: Developing Bayesian Observer Models for Bob

Develop several different Bayesian observer models for Bob’s potential cognitive process in guessing Alice's chosen suit. Below are three possible sets of model assumptions for your reference.

1. Offset Guess: Bob assumes that Alice's current choice is offset by a fixed number of positions from her previous choice. For example, if the offset is 2, and Alice's previous choice was diamond (1), Bob will guess spade (3). Bob knows Alice's probability of

choosing the offset suit (0.55) and the other suits (0.15), but he doesn't know the actual offset value. He tries to infer the offset from trial to trial following Bayes' theorem.

2. Adaptive Offset Guess: Similar to the Offset Guess strategy, but Bob doesn't know the probabilities of Alice choosing the offset suit and the other suits. He tries to infer both the offset and the probabilities from trial to trial following Bayes' theorem.
3. Pairwise Bayesian Comparison: Bob has a limited cognitive capacity and can only compare two hypotheses at a time. He maintains a belief about the most likely suit based on pairwise comparisons and updates his belief after each trial using Bayesian inference. Here's how the strategy works:

   a. Initialization: Bob starts with two hypotheses about Alice's suit selection strategy, denoted as H1 and H2. Each hypothesis is a possible offset value (0, 1, 2, or 3) that Alice might be using in her probabilistic strategy. Bob assigns equal prior probabilities to both hypotheses, P(H1) = P(H2) = 0.5.

   b. Likelihood Calculation: In each round, Bob calculates the likelihood of observing Alice's chosen suit (S) given each hypothesis. The likelihood is based on Alice's probabilistic strategy:

   - If the observed suit matches the predicted suit based on the hypothesis offset, the likelihood is 0.55.
   - If the observed suit does not match the predicted suit, the likelihood is 0.15.

   c. Posterior Probability Update: Bob updates the posterior probabilities of the hypotheses using Bayes' theorem:

   - P(H1|S) = P(S|H1) * P(H1) / [P(S|H1) * P(H1) + P(S|H2) * P(H2)]
   - P(H2|S) = P(S|H2) * P(H2) / [P(S|H1) * P(H1) + P(S|H2) * P(H2)]

   d. Guessing: To make a guess in each round, Bob selects the hypothesis with the higher posterior probability, subject to randomness from a softmax function.

   e. Updating Priors: After each round, Bob updates the prior probabilities of the hypotheses for the next round based on their posterior probabilities:

   - P(H1) = P(H1|S)
   - P(H2) = P(H2|S)

   f. Replacing Worse Performing Hypothesis: At each trial, there is some small probability (e.g., 0.05) that the worse performing hypothesis will be replaced by a new randomly generated hypothesis that are different from the current hypotheses. When such replacement occurs, the prior probability for the old better performing hypothesis is discounted (e.g., $P(H_{old}) = 0.8\ P(H_{old})$), and the prior probability for the new hypothesis is set to be $P(H_{new}) = 1 - P(H_{old})$.

   g. Iteration: Bob repeats steps b-f for each round, continuously updating his beliefs using Bayesian inference and replacing poorly performing hypotheses as needed.

To connect Bob’s beliefs with his final choice, we also need to specify a decision model. Let’s assume a softmax function with a temperature parameter.

Problem 3: Implementing Bayesian Observer Models

Implement these Baysian observer models. Simulate Alice's suit choices and Bob's guesses for multiple rounds (e.g., 300 rounds). Compare the performance of different models and models with different parameters in terms of their accuracy in guessing Alice's chosen suit.

Problem 4: Revealing Further Differences in Behavioral Patterns

There are a few strategies to reveal further differences between different Bayesian observer models in their predicted behavioral patterns. (1) Develop simpler descriptive models to capture the key characteristics of Bob's guessing behavior. For example, exponential curves can be fit to describe the increase of accuracy with time. (2) Examine the error patterns of different models.

Problem 5: Model Recovery Analysis

Use each model to generate 10 sets of data. Then fit each model to each dataset. See how well the true generative model can be identified and how well the true model parameters can be recovered.