Friend or Foe - Likelihood Ratio Tests and CIs

Set Up/Reminder of Example

Recall the friend or foe example: 16 ten month old babies were shown a sequence of videos where a helpful toy assisted a character in going up a hill, or an unhelpful toy pushed the character down the hill. In the experiment, $x = 14$ out of 16 babies chose the helpful toy to play with when given a choice.

$X \sim \text{Binomial}(16, \theta)$

1. Likelihood Ratio Test

Let's conduct a test of the claim that ten month old babies are not capable of processing the contents of the video; they choose their toy at random.

$H_0 : \theta = 0.5$

$H_A : \theta \neq 0.5$

(a) Write down the form of the likelihood ratio statistic for this test. Your answer should only involve the random variable X and numbers, but you do not need to simplify your expression other than any clear cancellations.

(b) Find the observed value of the likelihood ratio statistic. (You may want to use R to do your calculations.)

(c) Using a large-sample approximation to the sampling distribution of the likelihood ratio statistic, find the p-value for this test. (You will need to use R to do this, using one of the `pchisq` or `qchisq` functions.)

2. Approximate Confidence Interval by Inverting the Likelihood Ratio Test

(a) Suppose you were conducting a likelihood ratio test of the null hypothesis $H_0 : \theta = \theta_0$ with a significance level of $\alpha = 0.05$. For what values of the likelihood ratio statistic would you reject the null hypothesis? This will be in terms of a critical value for the test, w^*. What is w^*? (You will need to use R to find this, using one of the `pchisq` or `qchisq` functions.)

(b) Illustrate how a confidence interval could be obtained by inverting the likelihood ratio test.

Here is a plot of the likelihood function for θ:

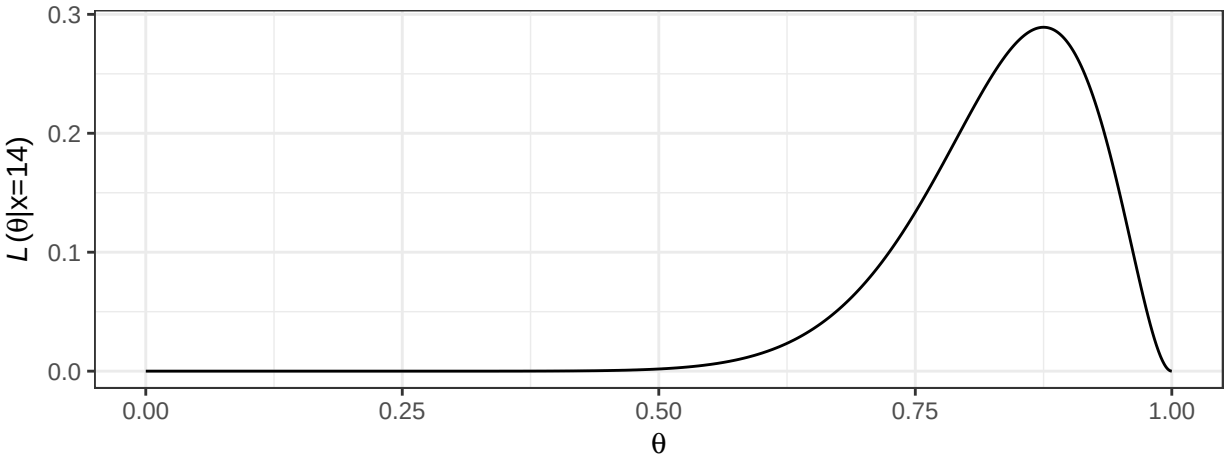

Here is a plot of the likelihood function for θ, but scaled by dividing by the value of the likelihood function at its maximum:

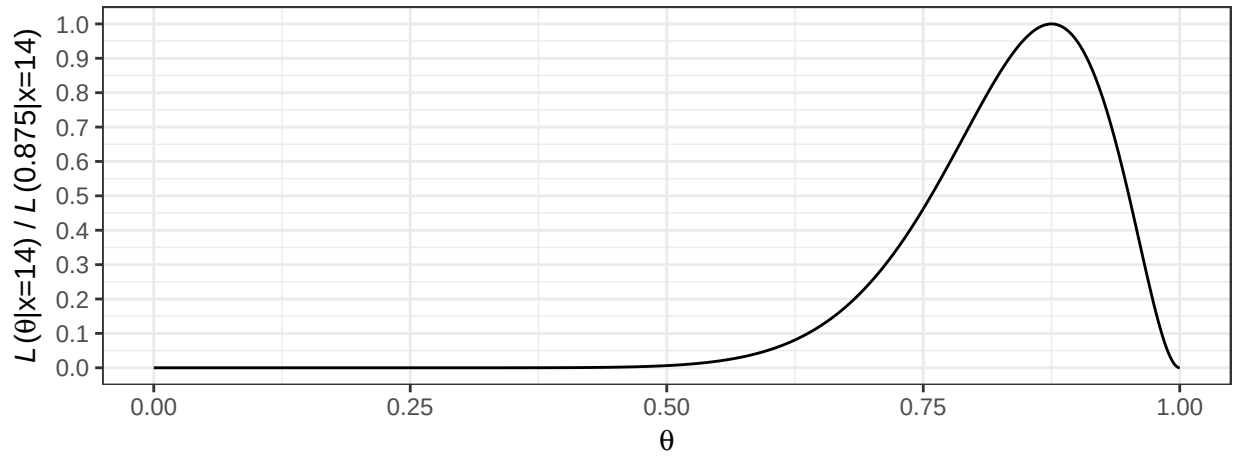

- On one of the two figures above (which one?), draw a line at the critical value w^* from part (a) (approximately - this does not need to be exact).
- Using that line as a reference, indicate the values θ_0 for which you would reject the null hypothesis that $\theta = \theta_0$.
- Now indicate the values of θ that would fall in an approximate 95% confidence interval for θ obtained by inverting the likelihood ratio test.

(c) Find the endpoints of your confidence interval using R.

Doing this exactly would require finding the roots of a degree 16 polynomial in θ. You can do that if you want, but an approximate approach that's good enough would be to calculate the likelihood ratio statistics at a grid of possible values for θ between 0 and 1, and then find out which values of θ are in your interval by comparing to the critical value w^*.