CS530 HOMEWORK 4, DUE 11/17

Justify your answers: if you do the problems by hand or calculator, show your intermediate formulas; if you do the problems by computer, show your code. Each problem in this homework is worth the same for grading.

Every day Waldo bikes to work through the same 4-way intersection. He heads north and sees the traffic light in that direction: red, yellow, or green. He has also been stopping at the intersection to check carefully whether there is traffic in the east-west direction. The following table lists his observations on 10 past days. Each column represents a day. The random variable L represents the northbound traffic light; its possible values are r, y, and g. The random variable T represents east-west traffic; its possible values are t and f.

L	r	g	r	r	r	g	g	y	g	r
T	f	f	f	f	t	f	f	f	f	t

In problems 1 through 3, assume that Waldo's mental model of the intersection is the following Bayes net.

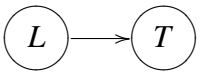

Watch Waldo learn!

(1) What is Waldo's maximum-likelihood estimate of the (conditional) probability distributions P(L) and P($T \mid L$)? What if Waldo applies add-1 smoothing?

(2) Suppose that Waldo has a uniform prior over these distributions (not an unreasonable prior for someone who has never seen a traffic light before the observations above took place). What is Waldo's Bayesian estimate of the distributions? What is Waldo's maximum-a-posteriori estimate of them?

Problems 3 through 5 are about the 11th day. On the 11th day, Waldo arrives at the same intersection as usual. He sees that the light is yellow, but he cannot check for east-west traffic unless he stops at the intersection. Stopping takes time (utility -3). If he bikes into the intersection when there is east-west traffic, he will be injured (utility -10).

(3) Should Waldo stop, according to the maximum-likelihood estimate from problem 1? Should Waldo stop, according to the Bayesian estimate from problem 2? (Note that Waldo's "output" here (whether to stop) is not the same variable as the Bayes net's "output" (whether there is east-west traffic).)

(4) Suppose that Waldo's mental model of the intersection is the following Bayes net instead.

That is, suppose that Waldo didn't realize that L affects T. Should Waldo stop, according to the maximum-likelihood estimate of the distributions? Should Waldo stop, according to the Bayesian estimate of them? Can this Bayes-net structure express the same set of joint probability distributions P(L, T) as the previous structure? Is this a difference in restriction bias or preference bias?

(5) Suppose that Waldo's mental model of the intersection is the following Bayes net instead.

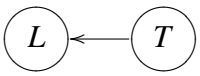

That is, suppose that Waldo thought that T affects L rather than the other way around. Should Waldo stop, according to the maximum-likelihood estimate of the distributions? Should Waldo stop, according to the Bayesian estimate of them? Given that this Bayes-net structure can express the same set of joint probability distributions $P(L, T)$ as the first structure, why are these recommended actions different from those in problem 3? Is this a difference in restriction bias or preference bias?