(1) Terms/concepts to know: mapping problem, fast mapping, cross-situational learning, Bayesian inference, posterior probability, likelihood, prior probability, sequential updating, positive examples, whole object constraint, taxonomic constraint, subordinate, basic, superordinate, suspicious coincidence, lexical contrast

(2) While fast mapping may sound like a good strategy in theory, why is it unlikely to be easy to carry out in real world situations? (Hint: Think about how many potential referents there are in a real world situation.)

(3) What evidence is there that infants can do cross-situational learning in experimental scenarios?

(4) As with fast mapping, why might cross-situational learning be more difficult in realistic scenarios, as compared with experimental scenarios where it has been shown to be present in infants? Is there any reason to believe that realistic scenarios (which have more potential referents) might be better for a cross-situational learner than scenarios where only a few word-referent pairings are presented?

(5) What are some factors that appear helpful for learners who use cross-situational learning to figure out word meaning? Why might they be helpful? (Hint: Does partial word knowledge help? Does repetition help? Does having a child's perspective of the world seem to help?)

(6) Give one example of overlapping concepts. How could this complicate a cross-situational learning strategy?

(7) What does it mean to have a graded inference about word meaning? How does this fit with the idea of cross-situational learning and Bayesian inference?

(8) How do constraints like the whole object constraint and the taxonomic constraint help a Bayesian learner? (Hint: Think about what these constraints do to the hypothesis space of possible word meanings.)

(9) Bayesian learners automatically implement a sensitivity to "suspicious coincidences", which is particularly useful when one hypothesis is a subset of another hypothesis. Give one example where a word-meaning hypothesis is a subset of another word-meaning hypothesis. Which would a Bayesian learner choose if it had encountered a number of suspicious coincidences?

(10) How do we know that children were conservative in how they made generalizations in the Xu & Tenenbaum (2007) experiment? (Hint: Think about their behavior on the one

example condition and the three subordinate example condition.) What was the difference between the one example condition and the three subordinate example condition, when we look at children's generalization behavior? Does this fit with the idea that children are sensitive to suspicious coincidences? Why or why not?

(11) Are children sensitive to how the data they learn from are selected? How do their generalizations differ when they think the data are sampled randomly vs. when they think the data are not sampled randomly? (Hint: Think about this in relation to suspicious coincidences.)

(12) Is there any evidence that the way children learn the meaning of adjectives is consistent with Bayesian inference?

(13) Can a Bayesian learner incorporate the idea of lexical contrast? How?

(14) What is one problem with the Bayesian learning account when we look at very early word learning (such as the word learning that occurs under 3 years of age)? (Hint: Think about how fast the Bayesian learner learns.) What are some ideas about how to reconcile a Bayesian learning account with very young children's word learning?

(15) Does children's sensitivity to suspicious coincidences (at least as measured by the noun generalization experiment of Xu & Tenenbaum 2007) remain constant over time? How does it change? (Hint: Think about children who know fewer category members vs. children who know more category members vs. adults.) How could language experience impact older children's generalization tendencies?