

An Introduction to Multi-Agent Statistics

Mathias Winther Madsen

August 2, 2014

Abstract

Building on ideas from probabilistic dynamic epistemic logic, the aim of this workshop contribution is to present the practical details of how to build and use Bayesian models that can represent multi-agent uncertainty, including probability distributions over (other people's) probability distributions. The key insight that underlies the approach taken here is that probabilities like $P_A(\varphi)$ are random variables and can be treated as such. By keeping carefully track of what various agents know under various circumstances, this way of looking at the issue provides enough structure to support queries about higher-order probabilistic belief statements.

Recent decades have seen a surge in two kinds of computational reasoning, sophisticated dynamic logics for multi-agent and higher-order uncertainty [1, 4], and probabilistic reasoning using causal Bayesian networks, often in conjunction with approximation methods like Markov Chain Monte Carlo or Gibbs sampling [2]. The purpose of my contribution to the workshop is to explain and strengthen the connection between these two styles of computational reasoning by presenting a system that integrates the core insights from dynamic logic into an explicitly statistical framework.

In this system, a series of statements about uncertain events can be translated systematically into a statistical model that includes information about which agents possess what pieces of information. This posterior model can be queried for higher-order questions such as the probability that a specific agent assigns to the beliefs of another agent.

The principles underlying this dynamic style of model construction come almost directly out of the literature on dynamic epistemic logic, especially its probabilistic incarnations [4]. My strategy throughout has been to use the simplest possible version of these ideas but to augment them with a useful statistical interpretation.

Because of space constraints, this note is not intended to completely specify the nature of the system, but rather to provide a suggestive and representative overview of its features. This makes the current abstract less formal than some (including perhaps myself) may prefer, but I hope the loss of precision is made up by the gain in intuitive clarity.

1 Multi-Agent Probability Models

Conventional single-agent probability theory has the following characteristics:

- The agent starts with a prior over the entire sample space.
- New information comes in the form of propositions which, if they are informative, are inconsistent with certain parts of the sample space.
- This information is incorporated in the model by deleting the obsolete parts of the sample space and then renormalizing the remaining part so that it still contains a total probability of 1.
- At any particular time, we can query such a model for the posterior probability the agent assigns to some particular proposition, or about the expected value of of some random variable.

In multi-agent probability theory with visible announcements, the situation is slightly different:

- The agents all start with a prior over the same sample space.
- Information is then given in the form of “announcements,” that is, partitions of the sample space. Each equivalence class in this partition corresponds to a proposition which will be communicated to the agent when the actual world falls in that class. This allows us to let all other agents know that we relay a certain piece of information to an agent without revealing the actual content of the message.
- The information an agent receives is incorporated into his or her knowledge representation by cutting up the sample space into information cells, one for each possible message, and then renormalizing these cells so that they form local posterior probability measures.
- Because the agents may have different pieces of information, they can ask meaningful questions about the probability that other agents assign a specific probability to a certain event, or compute the expected value of such probability assignments, etc.

Using these principles, we can translate a set of prior probability distributions and uncertain announcements into a posterior model. The relevant computations that can be performed on this model are discussed briefly below.

2 Whispering and Note-Passing

If we want to model secret or uncertain announcements, this picture has to be complicated a bit more: We then have to introduce random variables whose values are announcements (i.e., partitions of the sample space). Using such

a representation, we can allow for the possibility that an agent A possesses different kinds of information in different possible worlds.

A convenient way of thinking about this situation [3] is that each possible world in the sample space contains all the information necessary to completely settle the two questions

1. What are the ontic facts about the universe, as expressed by the values of the random variables X, Y, Z, \dots ?
2. What are the epistemic facts about the agents, as expressed by the list of variables that they know the value of?

To make this a bit more concrete, suppose we have an ontic description of the world which only contains a single coin flip X . In a single-agent scenario, we could model this situation by a sample space with just two possible worlds. However, if we want to introduce an uncertain announcement into this model, we have to introduce a second variable Y which keeps track of what the receiving agent knows in any particular possible world.

We could, for instance, announce the value of X to the agent when $Y = 1$, and announce a tautology (“ $1 = 1$ ”) when $Y = 0$, as in Figure 1. In that case, our agent will be able to distinguish the worlds $(X, Y) = (0, 1)$ and $(1, 1)$, but not the worlds $(0, 0)$ and $(1, 0)$.




Figure 1: Posterior uncertainty model for an agent who is told the value of a fair coin flip X if and only if another fair coin flip Y comes up heads ($Y = 1$).

3 Types and Type Transitions

Probability theory, be it multi-agent or single-agent, contains entities of two kinds, random variables and propositions. Random variables are functions from the sample space into some domain, typically the real number line; propositions are functions from the sample space into the truth value set $\{T, F\}$. Because the sample space is assumed to be equipped with a probability measure, the random variables have expected values, and the propositions have probabilities.

Given a set of random variables, we can construct new ones using arithmetic operations like addition, multiplication, and the like. Similarly, we can construct new propositions out of old ones using the logical operations of conjunction,




Figure 2: Syntactic parse of “ $P_A(P_B(X = 0) < 1/3 | X + 1 > 3) > 1/2$.” Random variables are shown as circles, propositions as squares.

negation, etc. These two families of operations thus map objects onto other objects of the same type (ignoring some trivial arity issues).

There are, however, also operations that cross the type boundary. Inequalities and other tests can convert random variables into propositions, as in “ $X \leq Y$ ” or “ $X = 3$ ” (cf. Fig. 2). Reversely, we can convert propositions into random variables by forming indicator functions: For instance, “ $X > 2$ ” is a proposition and thus has a truth value in each possible world; but “ $I(X > 2)$ ” is (in the notation I will use) a random variable that takes the value 0 or 1.

The most conceptually important operations of this kind, however, are the probabilities modalities. When φ is a proposition, $P_A(\varphi)$ is a random variable; but unlike conventional random variables, its value in a particular possible world depends on facts outside that world. Specifically, its value at a world ω is the conditional probability A assigns to the extension of φ given the information available at world ω . Computing this number involves measuring the intersection of the information cell around ω with the extension of φ .

These probabilities will usually be different in different information cells — either because the content of an announcement can differ between cells, or because the actual announcement itself can differ. However, the value of $P_A(\varphi)$ in ω is wholly determined once we know which of A ’s information cells ω falls in, and if we are, for instance, computing its expected value, this means that we only have to perform one computation per information cell.

References

- [1] Alexandru Baltag, Lawrence S. Moss, and Sławomir Solecki. The Logic of Public Announcements, Common Knowledge, and Private Suspicions. In *Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 43–56, 1998.
- [2] Thomas L. Griffiths, Charles Kemp, and Joshua B. Tenenbaum. Bayesian models of cognition. *Cambridge Handbook of Computational Cognitive Modeling*, pages 59–100, 2008.
- [3] Andreas Herzig and Tiago De Lima. Epistemic Actions and Ontic Actions: A Unified Logical Framework. In Solange Oliveira Rezende Jaime Simão Sichman, Helder Coelho, editor, *Advances in Artificial Intelligence – IBERAMIA-SBIA 2006*, pages 409–418. 2006.
- [4] Johan van Benthem, Jelle Gerbrandy, and Barteld Kooi. Dynamic Update with Probabilities. *Studia Logica*, 93(1):67–96, 2009.