Notes on the Lattice of Information

Matvey Soloviev

February 29, 2024

1 Relating the knowledge of different agents

So far, in a system with multiple agents, we have mostly considered their knowledge in isolation: we have not really investigated statements such as "K_1 knows more than K_2", or "K_3 is as knowledgeable as K_1 and K_2 together". You might recall that we *did* mention a notion of "what K_1 and K_2 know between each other", namely the distributed knowledge between them, $D_{\{1,2\}}$. However, we did not define this operator as the knowledge of some agent, but instead just said that $w \vDash D_{\{1,2\}}\varphi$ iff $w' \vDash \varphi$ for all w' such that (w, w') is both in $\mathcal{K}_1$ and in $\mathcal{K}_2$. Either way, it seems useful to be able to characterise when statements like that are true, particularly if we seek to use Kripke structures for security, where typical questions we want to answer may look like "if these two agents collude, could they learn more than the Average User?".

To start with, we can try to make sense of the first statement. Suppose we have fixed a model $M \in \mathcal{M}_n^{rst}$. It seems quite reasonable to say that if "K_1 knows at least as much as K_2", then if K_2 knows something, then K_1 must know it as well:

$$K_2\varphi \Rightarrow K_1\varphi. \tag{KImpl}$$

In fact, it's reasonable to take this axiom as characterising "knowing at least as much", as it is equivalent to saying that "what K_1 knows at w", $I_{1,w} \triangleq \{\varphi \in \mathcal{L}_n \mid w \vDash K_1\varphi\}$, is a superset of what K_2 knows at w, $I_{2,w}$, for any world w.

What can we say about K_1 and K_2 in a model of this axiom (a structure in which it is valid)? Not much, it turns out. For suppose that the structure is such that all the same primitive propositions are true at every world. Then every agent knows exactly the same formulas, namely those that are true everywhere, irrespective of their knowledge relation $\mathcal{K}_i$! (This is a simple consequence of knowledge generalization.)

To have any hope of saying something interesting, we need to add an extra assumption. The assumption that we make is that the language of primitive propositions is expressive enough to distinguish all the worlds. Specifically, we assume that for every world w, there exists a *characteristic formula* χ_w, which does not include any modalities, which is only true at that world:

- $w \vDash \chi_w$,
- $w' \vDash \neg\chi_w$, for all $w' \neq w$.

Given this assumption, it turns out that (KImpl) corresponds to an intuitive relation between the equivalence relations.

Theorem 1.1. *Suppose $M = (S, \pi, \mathcal{K}_1, \mathcal{K}_2)$ is a Kripke structure in $\mathcal{M}_2^{rst}$ and has characteristic formulae. Then (KImpl) is valid in M if and only if $\mathcal{K}_1 \subseteq \mathcal{K}_2$.*

Proof. **"if"**: Follows by expanding definitions. For arbitrary w and φ, we have $w \vDash K_i\varphi$ iff $w' \vDash \varphi$ for all $w' \in S$ such that $(w, w') \in \mathcal{K}_i$. So $w \vDash K_2\varphi$ implies $w' \vDash \varphi$ for all w' such that $(w, w') \in \mathcal{K}_2$. Hence $w' \vDash \varphi$ for all w' such that $(w, w') \in \mathcal{K}_1 \subseteq \mathcal{K}_2$, and hence $w \vDash K_1\varphi$.

"only if": Suppose $(w, w') \in \mathcal{K}_1$, but not $\in \mathcal{K}_2$. Then it is easily verified that $w \vDash K_2\neg\chi_{w'}$, but $w \vDash \neg K_1\neg\chi_{w'}$. So (KImpl) is not valid in M. □

So it turns out that the set inclusion order between equivalence relations characterises the logical notion of "knowing at least as much", as captured by (KImpl). What does this order look like? For four worlds:

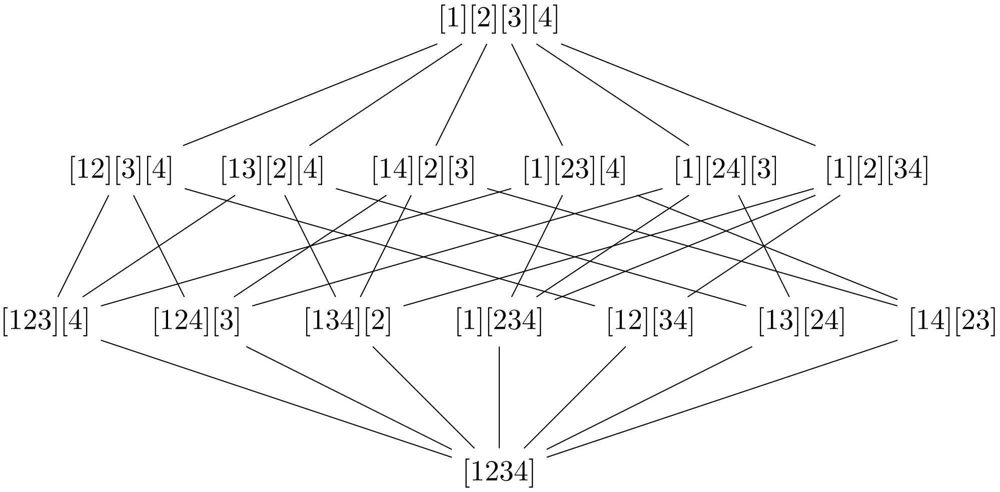

In general, the number of possible partitions is described by the *Bell number* B_n.

It is customary to define notation for an ordering operator so that "less" looks like "less knowledge", which requires flipping the direction of set inclusion: we write $K_1 \sqsubseteq K_2$ iff $\mathcal{K}_1 \supseteq \mathcal{K}_2$. The top element of the above Hasse diagram is then "greatest" under $\sqsubseteq$, while being "least" under $\subseteq$.

2 The Lattice of Information

Now that we have defined this partial order, we would like to tease out some more ways of relating the knowledge of agents. One natural thing to do with partial orders is to ask about *least upper bounds* for some collection of points (a term that is often used for the elements on which an order is defined; here, our points are partitions of the set of worlds), which are the least points that are greater than every point in that collection; and, dually, *greatest lower bounds*. Not all partial orders come with either, or both; but when they do, they are called *lattices*.

Definition 2.1. A *lattice* $(S, \leq, \wedge, \vee)$ is a partially ordered set $(S, \leq)$ together with two binary operators $\wedge$ (*meet*) and $\vee$ (*join*), which give the greatest lower bound and the least upper bound respectively.

(Note relation between operator visuals and what they stand for!) Lubs and Glbs turn out to exist for our information ordering $\sqsubseteq$. The resulting lattice is called the *Lattice of Information* (LoI) on a set S, $(\mathsf{Part}(S), \sqsubseteq, \sqcup, \sqcap)$.

Under our interpretation of the ordering $\sqsubseteq$ as "more"="more knowledge", a least upper bound $K_1 \sqcup K_2$ of K_1 and K_2 would correspond to a relation that encodes the *least* knowledge that one can have while having *at least as much* knowledge as either K_1 or K_2; and the greatest lower bound $K_1 \sqcap K_2$ would correspond to the *most* knowledge that one can have while having *no more* than either.

It may not be unexpected that the join of K_1 and K_2 in fact represents the result of *joining* the knowledge of K_1 and K_2, that is, knowledge that includes that in K_1 (as it is an upper bound on K_1) and that in K_2 (as it is an upper bound on K_2) but no more (as it is a *least* upper bound). This is what we have previously encountered as *distributed knowledge.*

Theorem 2.2. *Suppose $M = (S, \pi, \mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3)$ is a Kripke structure in $\mathcal{M}_3^{rst}$ and has characteristic formulae. Then $D_{1,2}\varphi \Leftrightarrow K_3\varphi$ is valid in M for all φ if and only if $K_3 = K_2 \sqcup K_1$.*

Proof. We first establish that $K_1 \sqcup K_2 = \mathcal{K}_1 \cap \mathcal{K}_2$. Statement follows. □

What about their meet? Unlike in the case of the join, the definition of the meet is not actually so straightforward, because the union of two equivalence relations is not in general an equivalence relation. However, it turns out that the minimum amount of work we can do to "fix up" the union – which is to form the transitive closure – does in fact give us the greatest lower bound, and this turns out to correspond to the notion of *common knowledge* that we have seen before.

Theorem 2.3. *Suppose $M = (S, \pi, \mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3)$ is a Kripke structure in $\mathcal{M}_3^{rst}$ and has characteristic formulae. Then $C_{1,2}\varphi \Leftrightarrow K_3\varphi$ is valid in M for all φ if and only if $K_3 = K_2 \sqcap K_1$.*

Homework problems

1. A lattice can also be defined algebraically in terms of the behaviour of $\wedge$ and $\vee$, by asserting that $a \vee (a \wedge b) = a \wedge (a \vee b) = a$ for all elements a, b. Show that this definition is equivalent to the glb/lub definition. Explain intuitively what this axiom means in the Lattice of Information.

2. A *complete lattice* is one in which joins and meets $\bigwedge S$ and $\bigvee S$ can be formed for infinite sets S as well. Show that, assuming there exists a maximum and minimum element (a $\top$ and $\bot$ such that $\top \geq a \geq \bot$ for all a), existence of infinite-set joins in fact implies existence of infinite-set meets, and vice versa.

3. We have established that distributed and common knowledge can be modelled as the knowledge of a particular special agent. There are some modalities for which this is not possible. In the following questions, you should prove or disprove.

   (a) Recall the "everybody knows" modality E_G, which we used in defining common knowledge. For general K_1 and K_2, does there exist an equivalence relation K_3 such that $E_{\{1,2\}}\varphi \Leftrightarrow K_3\varphi$? If not, is this possible with a general (non-equivalence) relation K_3? What if K_1 and K_2 are themselves not equivalence relations?

   (b) The "somebody knows" modality S_G is defined by $S_G\varphi \Leftrightarrow \exists i \in G.\, K_i\varphi$. For general K_1 and K_2, does there exist an equivalence relation K_3 such that $S_{\{1,2\}}\varphi \Leftrightarrow K_3\varphi$? If not, is this possible with a general (non-equivalence) relation K_3? What if K_1 and K_2 are themselves not equivalence relations?