Belief Aggregation for Non-Standard Reasoners

Ondrej Majer*, Olivier Roy and Norbert Grätzl

Abstract

The goal of this paper is to study belief aggregation in non-normal modal logic. The paper has two parts. First we look abstractly at aggregated beliefs in neighborhood semantics, defined in analogy with distributed knowledge in epistemic logic. We find that for doxastic logics weaker than K several different versions of aggregated belief can be defined. We provide sound and complete axiomatizations for four of them as well as a proof system in sequent calculus. Then we turn to a probabilistic foundations for these aggregation operations. We study the logic of individual and group beliefs resulting from using the Lockean thesis on (non-standard) probabilities based on Belnap-Dunn four valued logic.

Aggregated beliefs in Neighborhood Semantics

Our starting point is a non-normal model of (categorical, i.e. non-graded) individual beliefs using neighborhood frames (see e.g. [5], [3]):

Definition 1. *A* neighborhood frame $\mathcal{F}$ *is a pair* $\langle W, n_i \rangle$, *where* W *is a set of possible worlds and* $n_i : W \to \mathcal{P}(\mathcal{P}(W))$ *is a neighborhood function, one for* i *in a given, finite set of agents. A* neighborhood model *is a triple* $\langle W, n_i, v \rangle$ *such that* $\langle W, n_i \rangle$ *is a neighborhood frame and* $v : At \to \mathcal{P}(W)$ *is a valuation function.*

As usual, $M, w \models \varphi$ means that φ (a formula of a doxastic logic with a finite set belief operators B_i) is true at the state s in a model M. Truth of atomic formulas is given by the valuation v. Truth of propositional formulas is defined in the standard way as well as the truth conditions for beliefs[1]: $M, w \models B_i\phi$ iff $\|\phi\| \in n_i(w)$, where $\|\phi\| = \{w \in W | \ M, w \models \varphi\}$

In what follows we consider neighborhood frames with the following properties:

(Monotonicity) If $X \in n(w)$ and $X \subseteq Y$ then $Y \in n(w)$.
(Unit) $W \in n(w)$.
(D) If $X \in n(w)$ then for all $W - X \notin n(w)$.

The complete logic of such frames is well-known, all three conditions above can be axiomatized modularly using an axiomatization of classical propositional logic and the following axioms:

E	If $\vdash \phi \leftrightarrow \psi$ then $\vdash B_i\phi \leftrightarrow B_i\psi$	U	If $\vdash \phi$ then $\vdash B_i\phi$
M	$\vdash B_i(\phi \wedge \psi) \to B_i\phi$	D	$\vdash B_i\phi \to \neg B_i\neg\phi$

We want to study aggregated beliefs in that setup, where aggregation is defined in analogy with distributed knowledge in epistemic logic. We say that ϕ is aggregated or distributed beliefs in group G (and write $D_G\varphi$) if it would result from putting together what the members of G believe. For distributed knowledge in normal modal logic the natural way to interpret "putting together" is the intersection of the agent's epistemic accessibility relation. In neighbourhood semantics there are more room to maneuver. Here we look at two variants:

1. Taking the intersection of *pairs of* beliefs held by some (not necessarily different) agents in the group. We denote it $\sqcap_G n_i(w)$ below. Under that reading $D_G\phi$ means that there at least two agents in the group for which ϕ would result from putting together two of their beliefs.

2. Taking the intersection of *all* beliefs of all agents in G. This boils down to close the result of the first variant under intersection. So we denote it $\sqcap^*_G n_i(w)$ below. Here $D_G\phi$ means that ϕ would result from putting together all the beliefs of the group members.

*Supported by the project SEGA: From Shared Evidence to Group Agency of Czech Science Foundation and DFG no. 16-07954J. The author would like to thank to anonymous referees for helpful comments.

[1]There are some alternative interpretations, e.g. evidence based models use a 'monotonic' semantic interpretation in which the neighbourhood requires not the truth-set of the given formula, but only a subset of it, see [2].

Note that neither of these guarantees consistent aggregation. In case some of the agent's beliefs are mutually inconsistent, then the group will believe in contradiction: $\emptyset$ will be in the aggregated neighbourhood. If the agent's individual beliefs are otherwise closed under logical consequences (i.e. the neighbourhood are monotonic), then inconsistent aggregation results in explosion. This motivates looking at the result of applying the two operations above in neighbourhoods that are not necessarily upward closed. So taken together we consider four possible scenarios, which all result in different, non-equivalent logics:

	$\sqcap_G n_i(w)$	$\sqcap^*_G n_i(w)$
Monotonic n_i	Variation 1	Variation 2
Arbitrary n_i	Variation 3	Variation 4

In all these variations we use a standard truth condition for aggregated beliefs:

$$M, w \models D_G\phi \text{ iff } \|\phi\| \in n_D(w)$$

The two different understanding of aggregation that we just presented will correspond to two different definitions of n_D. Let A be the set of agents and G one of its non-empty subsets. Let:

$$\sqcap_G n_i(w) = \{X \cap Y : \text{ for } X \in n_i(w), Y \in n_j(w) \text{ for some agents } i, j \in G\}$$

We then define $\sqcap^*_G n_i(w)$ as the smallest set that contains $\sqcap_G n_i(w)$ but is such that if X and Y are in $\sqcap^*_G n_i(w)$ then their intersection is there as well.

Variation 1 This is the version that Eric Pacuit proposes in [5]. The resulting neighbourhood will be monotonic, and contain the unit. It might not, however, satisfy D, nor will it be closed under conjunction. Observe that necessitation (U) for D_G follows from $\wedge'$ and U for the individual beliefs. The usual inclusion axiom that is used in the axiomatization of distributed knowledge for normal beliefs ($B_i\phi \to D_G\phi$) is an instance of $\wedge'$.

Variation 2 This makes the resulting aggregated beliefs a normal modality. Like before, necessitation for D_G follows from necessitation for the individual B_i's. Observe, however, that despite being normal this modality still doesn't validate the D axiom, even if all the individual beliefs do.

Variation 3 We start by considering binary aggregation, i.e. the first variant in the introduction, of individual beliefs that only satisfy the E rule. We are looking at arbitrary neighbourhood functions for the individual beliefs, which might not contain the unit, might not be upward closed or might even be completely empty. Leaving out monotonicity has the direct consequence that D_G will not be closed under logical consequence either. Otherwise the logic of D_G stays the same.

Variation 4 We finally consider the case were we close the aggregated beliefs under conjunctions, but where the individual beliefs are represented by arbitrary neighbourhood. The resulting logic is what one would expect. The loss of monotonicity for the individual beliefs already makes D_G non-normal, but one direction of $\wedge$ remains valid. With this direction in hand we can use the weaker $B_i\phi \to D_{\{i\}}\phi$ in order to capture the relation between individual and aggregated beliefs.

We show that the following axioms for Variation 1 and 3 respectively plus the axioms for belief are together sounds and complete w.r.t. neighbourhood frames with D_G such that $n_D(w)$ is defined as $\sqcap_G n_i(w)$. The same statement holds for Variations 2 and 4 where $n_D(w)$ is defined as $\sqcap^*_G n_i(w)$.

	Variation 1		Variation 2
REG	If $\vdash \phi \to \psi$ then $\vdash D_G\phi \to D_G\psi$	$\wedge$	$\vdash D_G(\phi \wedge \psi) \leftrightarrow D_G\phi \wedge D_G\psi$
$\wedge'$	$\vdash B_i\phi \wedge B_j\psi \to D_{\{i,j\}}(\phi \wedge \psi)$	$Incl'_1$	$\vdash B_i\phi \to D_{\{i\}}\phi$
Incl	$\vdash D_G\phi \to D_{G'}\phi$ whenever $G \subseteq G'$	*Incl*	$\vdash D_G\phi \to D_{G'}\phi$ whenever $G \subseteq G'$

	Variation 3		Variation 4
$\wedge'$	$\vdash B_i\phi \wedge B_j\psi \to D_{\{i,j\}}(\phi \wedge \psi)$	$\wedge - I$	$\vdash D_G\phi \wedge D_G\psi \to D_G(\phi \wedge \psi)$
		$Incl'_1$	$\vdash B_i\phi \to D_{\{i\}}\phi$
Incl	$\vdash D_G\phi \to D_{G'}\phi$ whenever $G \subseteq G'$	*Incl*	$\vdash D_G\phi \to D_{G'}\phi$ whenever $G \subseteq G'$

Foundation in Non-Standard Probabilities

This part is aimed at looking for a foundation of the different notions of aggregated categorical beliefs above in terms of (non-standard) credence. The connection between partial and full belief is given by the Lockean thesis: a proposition φ is fully believed if the degree of belief in φ is sufficiently high i.e. above a given threshold $r > 1/2$ (see [4] for a recent discussion).

Probability functions are usually defined for a σ-algebra of subsets of a given set Ω. In logical contexts, however, it is often more natural to define probability functions directly for a propositional language. A probability function (for L - the language of classical propositional logic) is a function $p : L \to R$ satisfying the following constraints:

1. Non-negativity. $p(\varphi) \geq 0$ for all $\varphi \in L$.
2. Tautologies. If $\models \varphi$ then $p(\varphi) = 1$.
3. Finite additivity. If $\models \neg\varphi \wedge \psi$, then $p(\varphi \vee \psi) = p(\varphi) + p(\psi)$.

The starting point for individual beliefs is to use a simple form of the Lockean thesis[2]:

$$B\varphi \text{ if and only if } p(\varphi) \geq r > 1/2$$

It is easy to see that the Lockean belief operator B based on classical probabilities satisfies the conditions $\mathsf{E}, \mathsf{M}, \mathsf{D}, \mathsf{U}$ from the previous section. It is not normal however, because $p(\varphi) \geq r$ and $p(\psi) \geq r$ does not imply $p(\varphi \wedge \psi) \geq r$.

Our main motivation for introducing a non-standard framework is to represent agents with possibly inconsistent partial beliefs. Non-standard probabilities (see [6]) are supposed to be a generalization of the classical ones in the same spirit as the Belnap-Dunn four valued logic is a generalization of the classical logic. A formula in Belnap-Dunn logic might be neither true nor false (a truth value gap) or both true and false (a truth value glut). The probabilistic counterpart allows for gaps and gluts having non-zero probabilities.

Definition 2 (Non-standard probabilities)**.** *A (non-standard) probability space is a pair $\langle \mathcal{L}, p \rangle$, where $\mathcal{L}$ is the set of formulas of a propositional logic L and p is a (non-standard) probability measure, that is, a function from $\mathcal{L}$ into the real numbers satisfying:*

1. *for all $\varphi \in \mathcal{L}$, $0 \leq p(\varphi) \leq 1$,*
2. *for all $\varphi, \psi \in \mathcal{L}$, if $\varphi \models_L \psi$ then $p(\varphi) \leq p(\psi)$,*
3. *for all $\varphi, \psi \in \mathcal{L}$, $p(\varphi \wedge \psi) + p(\varphi \vee \psi) = p(\varphi) + p(\psi)$.*

The condition 3. makes the framework non-standard as the probability of a formula and its negation are not complementary in the usual sense, but are connected by a much weaker condition: $p(\varphi \wedge \neg\varphi) + p(\varphi \vee \neg\varphi) = p(\varphi) + p(\neg\varphi)$. This allows for $p(\varphi \wedge \neg\varphi) > 0$ (positive probability of gluts) and $p(\varphi \vee \neg\varphi) < 1$, i.e. $1 - p(\varphi \vee \neg\varphi) > 0$ (positive probability of gaps).

The axioms above clearly validate M for categorical beliefs. As in the classical case we do not have normality and given the above regarding gluts and gaps, we do lose both U (necessitation), and D (consistency). Lockean beliefs satisfy:

E	If $\vdash \phi \leftrightarrow \psi$ then $\vdash B_i\phi \leftrightarrow B_j\psi$
M	$\vdash B_i(\phi \wedge \psi) \to B_i\phi$

A standard way of getting a group attitude in this setup is to aggregate the individual probabilities and then apply the Lockean thesis again. We start with a simple linear weighted averaging: assume we have a group G of agents with individual probability functions p_i and equal weights. Then the aggregated probability $p_{G'}$ for a $G' \subseteq G$ is defined as

$$p_{G'}(\varphi) = \sum_{i \in G'} 1/|G'| \cdot p_i(\varphi), \text{ and aggregated belief as } D^{ag}_{G'}\varphi \text{ iff } p_{G'}(\varphi) \geq r$$

It is easy to see that $p_{G'}$ satisfies 1.- 3. from the previous definition. We compare the properties of the group belief operators from the previous section to those of the non-standard Lockean operator. D^{ag}_G satisfies regularity (REG), because if $\vdash \phi \to \psi$, then $\phi \models \psi$, and according to 2. in the definition of probability $p(\varphi) \leq p(\psi)$ as well as $p_{G'}(\varphi) \leq p_{G'}(\psi)$ for every G'. Distribution over conjunction ($\wedge$) holds for the same reason.

All variations from the previous part preserve an aggregated belief when the set of agents increases ($Incl$). In general this does not hold for $D^{ag}_{G'}$, the reason is that the probabilistic aggregation

[2]This form of Lockean thesis does not allow us to define higher order beliefs, we address this topic in a future research.

procedure takes into account not only belief of the agents, but in a sense also their disbelief – the agents added to the original group might push the new group average below the threshold. The weak form of individual belief inclusion ($Incl'_1$) $\vdash B_i\phi \to D^{ag}_{\{i\}}\phi$ holds, because $p_{\{i\}}(\varphi) = p_i(\varphi)$. The remaining properties ($\wedge'$) and ($\wedge - I$) do not hold in the Lockean setup. The properties of the non-standard Lockean beliefs are summarized in the following table

Lockean belief	
$Incl$	$\vdash D^{ag}_{G}\phi \to D^{ag}_{G'}\phi$ whenever $G \subseteq G'$
REG	If $\vdash \phi \to \psi$ then $\vdash D^{ag}_{G}\phi \to D^{ag}_{G}\psi$
$\wedge - E$	$\vdash D^{ag}_{G}(\phi \wedge \psi) \to D^{ag}_{G}\phi \wedge D^{ag}_{G}\psi$

Future work

This topic presents an initial stage of a work in progress and there are many ways to proceed. There are more methods of merging of both categorical and partial beliefs to be explored. We will also concentrate on a closer comparison of the categorical and Lockean part and try to establish some correspondence results. Another topic to deal with is the relation of our aggregation procedures for non-standard reasoners to judgment aggregation.

References

[1] Thomas Ågotnes, Yì N. Wáng. Resolving Distributed Knowledge. TARK 2015: 31-50, 2015.

[2] Johan van Benthem and Eric Pacuit. Dynamic logics of evidence-based beliefs. *Studia Logica* 99(1): 61-92, 2011.

[3] Helle Hvid Hansen, Clemens Kupke and Eric Pacuit. Neighbourhood Structures: Bisimilarity and Basic Model Theory. Logical Methods in Computer Science 5(2) 2009.

[4] Hannes Leitgeb. The stability theory of belief. *Philosophical Review*, 123(2): 131–171, 2014.

[5] Eric Pacuit. *Neighborhood semantics for modal logic: An introduction*, Course Notes for ESSLLI 2007, 2007.

[6] Graham Priest. *In Contradiction: A Study of the Transconsistent*, second/expanded edition, Oxford: Clarendon Press, 2006.

[7] Floris Roelofsen. Distributed knowledge. *Journal of Applied Non-Classical Logics* 17(2): 255-273, 2007.