Merging Information

Sujata Ghosh[*] and Fernando R. Velázquez-Quesada[†]

1 Introduction: the milieu

On innumerous occasions in our everyday life we are forced to make decisions on the basis of incomplete information that we acquire regarding the current state of affairs. While playing poker, we are forced to decide whether to bet without having any idea about the opponents' hands. We have *imperfect information* about the situation of the game. Scheduling cricket matches in the nor'westers season is just like a game of chance. There is no guarantee that a match would be played on the scheduled day, because of the possibility of the sudden storms. One has to depend on weather forecasts which are invariably incomplete in terms of their information content.

Thus, more often than not we are in imperfect information situations, where we do not know whether a relevant fact is the case or not. But even though we do not have the precise information, some decisions are usually needed to be taken. Then, we rely not only in what we know, but also in what we *believe*. Though we do not know if it will rain this afternoon or not, we usually find one case more plausible than the other, and we act based on our assumptions about the world.

Evidently, 'belief' is rather a dynamic notion: at this moment we may believe that it will not rain this afternoon, but our belief will change if we see the sky getting darker. Studying the dynamic nature of beliefs is the main motivation behind the emergence of *Belief Revision*, the field focussed on the process of updating beliefs, as well as revising them to consistently accept a new piece of information. It is basically the product of two converging research traditions. The first one has a computer science flavor, with its origins in the development of procedures to update databases (for example, the truth maintenance systems of [18] and, more recently, [13] and [14]); the second one is more philosophical, and has its origins in the discussion of the mechanisms by which scientific theories develop, proposing requirements for *rational* belief change (for example, the studies of [40, 41] and [32]).

But in general, when new information arrives, it does not show up just from one source: we can get information about the weather by listening to the radio, looking at some webpage, asking a friend and even by looking at how cloudy the sky is. Of course, not all the sources are equally trustworthy: we find a rain prediction of the forecast on internet more reliable than the sunny day prediction of an enthusiastic friend. Rather than revising our beliefs to accept a single incoming piece of information, we *merge* the information coming from the different sources (including our current beliefs), taking into account the different reliability of each one. This is the main motivation behind the field appropriately called *Belief Merging*.

And we can imagine an even more realistic case, where the different sources of information indicate not only beliefs about the facts in discussion but also opinions about themselves. It is possible to find a radio broadcasting stating not only that it will rain this afternoon but also that we should not trust in anyone who says the opposite.

---

[*] Department of Artificial Intelligence, University of Groningen

[†] Institute for Logic, Language and Computation, Universiteit van Amsterdam

In the present work we propose a further extension of the formal studies of belief merging, with an aim of making it closer to the reality. We consider situations where the different sources of information (considered as agents) have opinions about each other also. In section 2 we present a brief survey of the main traditions in the formal studies on merging of information. Section 3 comprises of the main contribution of this work, presenting a logic in which we can express agent's opinions about facts and other agents (LO) and then extending it to express beliefs and also preferences of agents over other agents (LOB^{-}). Some discussions are provided, pointing towards the interactive nature of these epistemic attitudes. Finally, in section 4 we discuss our general conclusions and and give pointers for further work.

2 The vast realm of approaches

In the literature one can find many proposals for postulates and procedures to revise and merge beliefs. Here we provide a brief description of the main ideas that are being nurtured in these fields. Over the years, different approaches to the problem have been proposed, some from a logical point of view but some others from a connectionist perspective. We give a small survey of the most relevant works in both the areas.

2.1 Revising vs merging

Belief Revision focusses on the process of modifying beliefs so as to accept a new piece information in a consistent manner. One of the most important traditions in the field is the AGM model introduced in [1]. Following the philosophical origins of the field, the authors discussed a set of postulates for rational revision - properties that an operator that performs revision should satisfy in order to being considered rational. The AGM approach *"analyzes belief change without committing to any particular mechanism, providing just abstract postulates on the process"* ([51]).

The AGM approach contains two asymmetries in its formulation. The first one is the precedence of incoming information over current beliefs (current beliefs should change to *accept* the new information consistently); the second one is the richer structure assigned to the belief set (a collection of formulas, sometimes with an underlying structure) compared with that of the incoming information (just a formula). While theories of non-prioritized belief change have emerged to tackle the first issue ([31, 44]), other kinds of generalizations have surfaced with works considering the aggregation of finite sets of information, all of them with similar structure and arbitrary precedence, into a collective one. Instead of *revising* beliefs, we *merge* all available information: this is the main idea behind *Belief Merging* ([34, 35, 36]).

It is also interesting to notice that the aggregation procedure in *Belief Merging* faces problems similar to those addressed in *Social Choice Theory*. Both fields consider several sources of information with different precedence (based on reliability in *Belief Merging*, priority in *Social Choice*), which provide an order over the relevant entities (beliefs about the current state of affairs in *Belief Merging*, preferences over a set of options in *Social Choice*). Links between these two disciplines have been investigated in [34, 37, 19, 21], among others.

2.2 Different approaches

As various other analogous research areas, revising and merging beliefs can be studied from two perspectives: an abstract one providing and discussing properties of ideal solutions, or a more practical one providing solutions and verifying which properties they satisfy. While providing a specific procedure to get solutions, several approaches may be considered – in the area of revising and merging of beliefs, two have been used most extensively: logical approaches, providing models and formal languages to describe the relation between the input and output of the revising and merging mechanisms, and connectionist approaches, considering the phenomena as the emergent processes of interconnected networks of simple units (usually, neural network models).

2.2.1 Logical approaches

There are several approaches that use modal logic tools. Authors like André Fuhrmann [20], Johan van Benthem [50] and Maarten de Rijke [15, 16] showed how theories of change can be analyzed with the help of *"dynamic modal logic"*. By using a multi-modal language, dynamic logic allows to express the effect of actions: formulas of the form $[a]\varphi$ indicating that φ is the case after every execution of the action a. Moreover, it allows us to build more complex actions from the basic ones – we can have formulas expressing the result of sequential or parallel composition of actions or even, non-deterministic choice between them. The following works incorporate revising (merging) operations as actions within a modal language.

Dynamic Doxastic Logic (DDL) [47, 48] was introduced *"with the aim of representing the meta-linguistically expressed belief revision operator as an object-linguistic sentence operator in the style of dynamic modal logic"* [39]. The main operations in Belief Revision are, after all, actions and hence, we can use the dynamic logic framework to describe belief change. In DDL, we have (doxastic) actions of the form $+\varphi$ for expansion by φ, $-\varphi$ for contraction and $*\varphi$ for revision. Semantically, the idea is that a belief state should not only represent the agent's beliefs but also how she would respond to new information. Based on the work of Lewis [42] and Grove [29], Segerberg proposed that a belief state can be represented with a non-empty set of theories and a doxastic action can be represented with a binary relation between belief states. Then, the effect of a doxastic action in a belief state is described as a change to another belief state following the corresponding relation.

There is another major manifestation of the "dynamic turn" in logic: *Dynamic Epistemic Logic* (DEL) is the combination of two traditions in formal logic: *Epistemic Logic* (EL) and *Dynamic semantics*. While *Epistemic Logic* is concerned with reasoning about knowledge, the main idea in *Dynamic Semantics* is that *"the meaning of a syntactic unit is best described as the change it brings about in the state of a human being or a computer"* [25]. In DEL languages, an EL language is extended with operators that describe information-changing actions. On the semantics side, such operators differ from the usual modal ones in that they are not interpreted as relations between worlds, but as operations that *modify* the whole model. In Public Announcement Logic (PAL, [46, 24, 25]), for example, the public announcement operation removes those worlds of the model where the announced formula is false.

In [6], the authors extended the PAL framework by using Kripke structures to represent not only the epistemic state but also *epistemic actions*: actions about which the agents may be incompletely informed. The epistemic state after the execution of an action is obtained by what is called *product update*, reflecting the idea that the uncertainty of an agent about a situation after an action takes place is the result of her uncertainty about the situation *before* the action and her uncertainty *about the action* ([56] provides a nice presentation of product update and other dynamic epistemic logic topics). Further extensions can be found in [8, 7], where the notion of belief has been incorporated, allowing us to describe agents with knowledge and beliefs about the situations and also about the executed actions. With this notion of *doxastic action*, it is possible to deal with both static and dynamic belief revision and also *"implement various belief-revision policies in a unified framework"* [9].

One of the main conceptual contributions of [6] was to put the description of the static situation and that of the actions at the same level. In [52], van Benthem further extends this symmetry by noting that product update (updating of the epistemic state through an epistemic action) is actually an aggregation of the (epistemic) relations of the static model and those of the action model. Aggregation is usually conceived as merging of different relations over *the same* domain; product update generalizes it by merging different relations over *different* domains. When we perform product update, two new items are built: the new domain (as the cartesian product of the previous ones) and a relation over it (based on the previous relations). In the paper, the author introduces a static modal language with modalities for the weak and strict versions of the ordering relation, whose logic is similar to the one presented in [53]; on top of which he adds the dynamic operators working as product updates following a *priority update rule*, reflecting the general idea of the Andreka et. al. approach, described below.

In [2], Andreka et. al. present an algebraic treatment for combining relations (which, in particular, can represent a plausibility order over possible situations). They define the concept of a *priority operator*: an operator that, given a family of relations with priority among them, returns a single relation representing their lexicographic combination. It is shown that priority operators are the only way of combining relations with different priorities to get a result that satisfy certain natural conditions, similar to those proposed by Arrow [4] in the context of *Social Aggregation*. Moreover, it is shown that any finitary priority operator can be expressed by the binary operators "||" (*"on the other hand"* operator, indicating the aggregation of relations with the same priority) and "/" (*"but"* operator, indicating the aggregation of relations with different priority). It should be noted how the construction of the aggregation relation is then given by a sequence of operations that defines the priority among the aggregated individual relations.

In Chapter 6 of [28], Girard presents a modal logic for order aggregation based on these priority operators. Following approaches for preference logic like [55], he presents a language that allows us to express not only individual preferences but also their aggregation as the result of operations between the corresponding relations.

2.2.2 Connectionist approaches

As we mentioned earlier, connectionist approaches consider revising and merging of beliefs as a result of some dynamics in interconnected networks (i.e., graphs) made up of simple units. At any point of time, each unit in the network is assigned a value intended to represent some aspect of that unit; this value is given by an external input depending on some combination of values of the other units. At each step, values are re-calculated, and hence the effect of the external inputs spreads to all other units in the network over time. Some of the units are considered output units and, whenever their values become stable, they are considered as the outcome of the process. There are many forms of connectionism, but the most common forms use neural network models.

An *artificial neural network* (NN) is a mathematical model based on a group of simple processing elements (neurons) whose behaviour is determined by their interconnections and their individual parameters. The NN model has its origin in biological neural networks and is typically an input-output model. They provide an interesting framework to model *reasoning*: if we consider the inputs as incoming information, the outputs can be considered as the result of a reasoning process. Then, an NN can not only represent situations with multiple sources of information (in general any finite number of inputs), but also represent their non-uniformity (by specifying different ways in which the different inputs will affect the output). This makes neural networks an attractive tool for modelling merging processes.

Analyzing logical concepts from a connectionist point of view is another stimulating area of study. An interesting analogy can be found in Gaifmann Pointer semantics ([22, 23]) and the revision theory of truth developed by Herzeberger, Gupta and Belnap ([33, 30]) which provide a semantics for sets of sentences with self-reference, by specifying stable patterns among the truth values (under subsequent revisions of truth values) of the sentences comprising the set. One of the main features of this revision theory is the *backward propagation* of truth values along the edges of some graph representing the set of sentences. A more belief-related approach is that of *Bayesian Belief Nets* (see [45, 57]), based on the idea of assigning probabilities to the units of the network and using the Bayesian rule for the propagation of values.

A related work is the *Assertion Networks* of [26]. Sources of information (considered as agents) and facts are uniformly represented as units of the network, and their interconnection represents opinions of the agents about relevant facts and also about themselves. The approach uses the graphical framework of *Bayesian Belief Nets* (though the numerals used are interpreted in a more deterministic way), and perhaps more crucially, considers the notion of *stability of revision sequences* to decide the outcome of the merging process.

The novelty of this setting is that we can now talk about opinions that agents have about other agents. It is not the case that there is a precedence among the multiple information sources; the formalism allows each one of them to have similar significance in the situation involving them. This also involves common concepts in security, like *Trust* and *Obligation* ([17, 43]).

Going back to our discussion on neural networks, one of their main drawbacks is the incapacity to provide an explanation for the underlying reasoning mechanism, that is, the incapacity to provide a *logical* description of the merging process. Several attempts have been made to provide some relation between inputs and outputs in connectionist models.

In [12], the authors discussed some of the main problems in the knowledge extraction methods of neural network models and proposed a way for these drawbacks to be amended. The basic ideas are to incorporate a partial order on the input vectors of the net and to use a number of pruning and simplification rules. For a class of networks, the *regular* ones, the algorithm is sound and complete. For the *non-regular* networks, it is shown that they contain regularities in their subnetworks, and the previous method can be applied in a decompositional fashion.

In [38], the author described *interpreted dynamical systems.* A dynamical system is in essence a neural network: we have states representing the neurons and a next-state function representing the connections between them. With an interpretation mapping, we identify formulas of the propositional language with states.

The internal dynamics of the system is given by stepwise iteration of the next-state function. External information modifies the internal flow, affecting the next state of the one to which it is plugged. The resultant dynamics of the system are described in terms of qualitative laws for which a satisfaction clause is defined. Moreover, it is shown that the resulting descriptions are related to systems of non-monotonic logic. The relation between non-monotonic logic systems and the symbolic representation of the *reasoning process* performed by neural networks is of particular relevance, as shown not only in [38], but also by [5] and [11], among others.

3 Merging opinions, preferences and beliefs

Real life communication situations are possibly the best exemplification of this merging of information. Various approaches to model this blending of information, both from the formal as well as the connectionist point of view, have been discussed extensively in the previous section. As evident from the existing literature, the formal framework provides a much more global approach regarding ‘what can be achieved’ or ‘what is the final outcome’, but does not give an account of the mutual influences and nuances of the agents involved, which has been brought to the fore by the connectionist approaches. The primary goal of this work is to provide a formal framework that describes the ‘micro-structure’ of the communication networks and capture a reasonable way of amalgamating all the information coming from different sources so as to bring out the global behavior of the system.

Our proposal is based on the approach of [26], where an *Assertion Network* takes the perspective of an external observer that collects opinions of relevant agents about some particular facts. Both agents and facts are represented as nodes of a graph, and opinions of the agents about themselves and the facts are represented by labelled edges. We emphasize the fact that agents are not represented by their preferences, beliefs or knowledge, but as tangible objects in the model.

The starting situation is given by some initial assignment, indicating the degree of belief (numerical values) the observer has about each agent and fact (each node). Then, just as connectionist networks, the process through which the observer *merges* the different agents’ opinions and her own initial beliefs is represented by the iterative updating of such values until a stable value is reached (which in general is not the case, e.g. *Liar* sentence).

We consider a semantic model close to the described one. It consists of a graph with two kinds of nodes, one representing agents and the other representing possible situations. We give a logical perspective here, and as such, develop a language for describing such a model. Since agents are represented by nodes, it is natural to consider a language that allows us to *name* such nodes. Our work is based on the hybrid logic approach ([3]) which provides a way to talk about individual states. In this section, we formally define the languages as well as the semantic models, and provide sound and complete axiom systems for them. The logics that we discuss here are devoid of any dynamic component, which we leave for future investigations. A short discussion of the interesting issues that arise from the notion of dynamics is in section 4.

Before proceeding further, it should be noted that the distinction made between *"belief"* and *"opinions"* is that, we understand *"opinion"* as a *first-order* concept while *"belief"* is a concept that can be extended to higher orders. In the logic of opinions, LO, that is proposed below, we can express agent's opinions about facts and about each other; opinions about opinions are meaningless for the framework. On the other hand, the logic of opinions and beliefs, LOB^-, an extension of LO, allows us to express agent's beliefs about facts and about each other, and also belief about beliefs, and so on. We also introduce formulas to express agents' preferences over other agents, which is an useful tool in representing communication situations, but aggregation of preferences so as to model the effect of group preferences is not discussed here. We leave this issue for the future.

3.1 Logic of opinions

The logic of opinions basically represents situations comprising of agents and events, together with opinions of agents about these events and also about other agents. Events are expressed by propositional variables, nominals and their boolean combinations, whereas agents are expressed by agent-names. The world-names (nominals), besides giving us uniformity by allowing us to name *all* nodes in our graph, allows to talk about agent's opinions not only about propositions that may be true in certain situations (several worlds) but also about any such complete world's descriptions. The language of this logic (LO) is given as follows:

Definition 3.1 *Let $PROP$ be a set of atomic propositions, AG be a set of agent-names and NOM be a set of world-names (nominals). Formulas φ of LO are inductively defined by:*

$$\varphi := p \mid \bot \mid i \mid \neg\varphi \mid \varphi \wedge \psi \mid [+]a : \varphi \mid [-]a : \varphi \mid (+)a : b \mid (-)a : b \mid @_i\varphi$$

where $p \in PROP$, $i \in NOM$ and $a, b \in AG$. We assume a restricted language in the sense that nesting of opinion modalities are not allowed. In formulas of the form $[+]a : \varphi$ and $[-]a : \varphi$, φ's are restricted to atomic propositions, nominals and their boolean and $@_i$ combinations.

Other connectives ($\vee$, $\rightarrow$ and $\leftrightarrow$) are defined as usual. The *diamond* version of opinion formulas $\langle + \rangle a : \varphi$ and $\langle - \rangle a : \varphi$ are defined as the duals over their *box* counterparts: $\langle + \rangle a : \varphi \leftrightarrow \neg[+]a : \neg\varphi$ and $\langle - \rangle a : \varphi \leftrightarrow \neg[-]a : \neg\varphi$, respectively.

The intended meaning of formulas of the form $[+]a : \varphi$ ($[-]a : \varphi$) is *"agent a has positive (negative) opinion about φ"*. Similarly, formulas of the form $(+)a : b$ ($(-)a : b$) are read as *"agent a has positive (negative) opinion about agent b"*. It should be noted that, in the language of LO, agent-names just appear in opinion formulas.

For the semantic model, we consider graphs with two kinds of nodes: agent-nodes representing agents and world-nodes representing possible situations. Relations between such nodes indicate the agents' opinions about possible situations and other agents. The basic link between the semantic model and the language is given by a couple of functions, a standard hybrid valuation indicating the world-nodes where elements of $PROP$ and NOM are true (with each $i \in NOM$ being true at one and only one world-node, as usual) and a naming function assigning a different agent-node to each element of AG. Formally, we have the following.

Definition 3.2 *An opinion model is a structure* $\mathcal{M} = \langle W, A, R^+, R^-, O^+, O^-, V, N\rangle$ *where*

- W *is the set of world-nodes,*
- A *is the set of agent-nodes (with* A *disjoint from* W*),*
- $R^+ \subseteq A \times W$ *is a serial binary relation from agent-nodes to world-nodes,*
- $R^- \subseteq W \times W$ *is a serial binary relation from agent-nodes to world-nodes,*
- $O^+ \subseteq A \times A$ *is a binary relation from agent-nodes to agent-nodes,*
- $O^- \subseteq A \times A$ *is a binary relation from agent-nodes to agent-nodes,*
- $V : \mathit{PROP} \cup \mathit{NOM} \to 2^W$ *is a standard hybrid valuation (that is, for each* $i \in \mathit{NOM}$, $V(i)$ *is a singleton), indicating the world-nodes where atomic propositions and nominals are true, and,*
- $N : \mathit{AG} \to A$ *is an injection assigning a different agent-node to each agent-name.*

The opinion model $\mathcal{M}$ *is* named *if every world-node in the model is the denotation of some nominal, that is, for each* $w \in W$*, there is a nominal* $i \in \mathit{NOM}$*, such that* $V(i) = \{w\}$*.*

We emphasize the fact that the set of world-nodes and that of agent-nodes are disjoint. Although $\mathcal{M}$ is a graph consisting of both these kind of nodes, atomic propositions and nominals get truth values only in world-nodes, and therefore formulas of the language are evaluated just in them. While the valuation V allows us to give truth value to atomic propositions and nominals in the standard way, the naming function N and the opinion relations (R^+ and R^- for opinions about facts, O^+ and O^- for opinions about agents) allow us to take care of opinion formulas by just looking at the outgoing arrows from the agent-nodes named after the agent-names. Negations, conjunctions and the $@_i$ operator are defined in the usual way.

Definition 3.3 *Let* $\mathcal{M} = \langle W, A, R^+, R^-, O^+, O^-, V, N\rangle$ *be an named opinion model. The truth definition for formulas* φ *in* $\mathcal{M}$ *at a world* w *is given below.*

$$\begin{array}{lll}
\mathcal{M}, w \models p & \textit{iff} & w \in V(p) \\
\mathcal{M}, w \models i & \textit{iff} & \{w\} = V(i) \\
\mathcal{M}, w \models \neg\varphi & \textit{iff} & \mathcal{M}, w \not\models \varphi \\
\mathcal{M}, w \models \varphi \land \psi & \textit{iff} & \mathcal{M}, w \models \varphi \textit{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models [+]a : \varphi & \textit{iff} & \textit{for all } u \in W \textit{ such that } N(a)R^+u,\ \mathcal{M}, u \models \varphi \\
\mathcal{M}, w \models [-]a : \varphi & \textit{iff} & \textit{for all } u \in W \textit{ such that } N(a)R^-u,\ \mathcal{M}, u \models \neg\varphi \\
\mathcal{M}, w \models (+)a : b & \textit{iff} & N(a)O^+N(b) \\
\mathcal{M}, w \models (-)a : b & \textit{iff} & N(a)O^-N(b) \\
\mathcal{M}, w \models @_i\varphi & \textit{iff} & \mathcal{M}, u \models \varphi, \textit{ where } V(i) = \{u\}
\end{array}$$

where $p \in \mathit{PROP}$, $i \in \mathit{NOM}$ *and* $a, b \in \mathit{AG}$.

It follows that $\mathcal{M}, w \models \langle + \rangle a : \varphi$ iff there exists $u \in W$ such that $N(a)R^+u$ and $\mathcal{M}, u \models \varphi$, and $\mathcal{M}, w \models \langle - \rangle a : \varphi$ iff there exists $u \in W$ such that $N(a)R^-u$ and $\mathcal{M}, u \models \neg\varphi$. Note that by the above definition, the opinion formulas are either true or false in a model.

A communication situation representing agents' opinion about events and about each other can be described by finite conjunction of the modal formulas introduced above. For example, the formula $\neg[+]a : \varphi \land \neg[-]a : \varphi$ is read as *"the agent does not have any opinion about* φ*"*, whereas $([+]a : \varphi \land [-]a : \varphi)$ corresponds to a being undecided about φ. In terms of epistemic attitudes of an agent, there is a difference between *having no opinion* and *being undecided* about a certain event. One can be undecided whether to take an umbrella or not while she is going out, but she may have no opinion about who should win the next US presidential elections, as she is simply not interested in the issue. In terms of opinions concerning other agents, these attitudes are typically indistinguishable. Moreover, because of the way we have interpreted the modal formulas which is independent of the states, the intuitively inconsistent formulas like $[+]a : \varphi \wedge [+]a : \neg\varphi$ as well as $[-]a : \varphi \wedge [-]a : \neg\varphi$ are not satisfiable.

Consider the following simple example.

> Suppose Professor Calculus wants to know how good a singer Bianca Castafiore is. In come Thomson and Thompson, and convey the following.
>
> `Thomson`: "*She is a very good singer.*"
> `Thompson`: "*Aha! I do not think so. I really dislike her singing*"

This network of opinions can be represented by the following formula:

$$[+]a : \varphi \wedge [-]b : \varphi$$

with a representing Thomson, b Thompson, and φ representing the fact that "Bianca Castafiore is a good singer". The obvious question here is what would Professor Calculus infer in this situation of conflicting opinions. We will come back to this example in section 3.2.

Let us now move on to provide a sound and complete axiomatization for *LO*.

Theorem 3.4 *LO is sound and complete with respect to countable named opinion models and its validities are axiomatized by,*

a) all propositional tautologies and inference rules,

b) axioms and rules for $@_i$*:*
$\vdash @_i(p \to q) \to (@_i p \to @_i q)$
$\vdash @_i p \leftrightarrow \neg @_i \neg p$
$\vdash i \wedge p \to @_i p$
$\vdash @_i i$
$\vdash @_i j \leftrightarrow @_j i$
$\vdash @_i j \wedge @_j p \to @_i p$
$\vdash @_j @_i p \leftrightarrow @_i p$
if $\vdash \varphi$*, then* $\vdash @_i \varphi$ *for every* $i \in NOM$

c) positive opinion axioms and rules:
$\vdash [+]a : (\varphi \to \psi) \to ([+]a : \varphi \to [+]a : \psi)$ *(+normal)*
$\vdash [+]a : \varphi \to \langle + \rangle a : \varphi$ *(+ser)*
$\vdash \langle + \rangle a : i \wedge @_i \varphi \to \langle + \rangle a : \varphi$ *(+translation)*
$\vdash \langle + \rangle a : @_i \varphi \to @_i \varphi$ *(+back)*
if $\vdash \varphi$*, then* $\vdash [+]a : \varphi$*, for every* $a \in AG$ *(+gen)*

d) negative opinion axioms and rules:
$\vdash [-]a : (\neg\varphi \wedge \psi) \to ([-]a : \varphi \to [-]a : \psi)$ *(-normal)*
$\vdash [-]a : \varphi \to \langle + \rangle a : \varphi$ *(–ser)*
$\vdash \langle - \rangle a : i \wedge @_i \neg\varphi \to \langle - \rangle a : \varphi$ *(–translation)*
$\vdash \langle - \rangle a : @_i \varphi \to \neg @_i \varphi$ *(-back)*
if $\vdash \neg\varphi$*, then* $\vdash [-]a : \varphi$*, for every* $a \in AG$ *(–gen)*

e) agreement axioms:
$\vdash \langle + \rangle a : \varphi \leftrightarrow @_i \langle + \rangle a : \varphi$
$\vdash \langle - \rangle a : \varphi \leftrightarrow @_i \langle - \rangle a : \varphi$
$\vdash (+)a : b \leftrightarrow @_i (+)a : b$
$\vdash (-)a : b \leftrightarrow @_i (-)a : b$

f) name, paste and substitution rules:

if $\vdash i \to \varphi$, *then* $\vdash \varphi$ — *for* i *not occurring in* φ.

if $\vdash (@_i\langle + \rangle a : j \wedge @_j\varphi) \to \psi$, *then* $\vdash @_i\langle + \rangle a : \varphi \to \psi$ — *for* $i \neq j$ *and* j *not occurring in* φ *or* ψ

if $\vdash (@_i\langle - \rangle a : j \wedge @_j\neg\varphi) \to \psi$, *then* $\vdash @_i\langle - \rangle a : \varphi \to \psi$ — *for* $i \neq j$ *and* j *not occurring in* φ *or* ψ

if $\vdash \varphi$, *then* $\vdash \varphi\sigma$ — *where* σ *is a substitution that uniformly replaces atomic propositions by formulas, agent-names by agent-names and nominals by nominals.*

Proof. *Soundness of the axioms and rules is straightforward. The completeness proof is based on that of hybrid logic as presented in ([10]); see appendix A for details.* ■

3.2 Logic of opinions and beliefs

Let us now go back to the example we mentioned earlier, and analyze it further. Evidently, if Calculus believes Thomson more than Thompson, then he would believe that "Bianca is a good singer", otherwise he would believe that "she is not". In the retrospect, to express such decision making in presence of conflict in opinions we want to have a language where we can say that if our belief in a is more than our belief in b then we may believe in φ, and if our belief in b is more than our belief in a, we may believe in $\neg\varphi$. To incorporate the belief strength part, we add operators B_a (belief of a) and $\succeq_a$ (a's preferences among agents) for each agent a to the language of LO, to form the language of the logic of opinions and beliefs. We first present a very simple logic with a straightforward axiomatization which does not take into account the interaction between belief, opinions and preferences. We name it LOB^-, whose language is given in the following,

$$\varphi := p \mid \bot \mid i \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle + \rangle a : \varphi \mid \langle - \rangle a : \varphi \mid (+)a : b \mid (-)a : b \mid @_i\varphi \mid B_a\varphi \mid a \succeq_a a,$$

with the following restriction that φ occuring in $B_a\varphi$ should only be propositions, nominals and their boolean combinations. We use $\langle B_a \rangle\varphi$ as an abbreviation for $\neg B_a\neg\varphi$, and $b \succ_a c$ for $(b \succeq_a c) \wedge \neg(c \succeq_a b)$.

The model becomes $\mathcal{M} = \langle W, A, R^+, R^-, O^+, O^-, \{R_a : a \in A\}, \{\geq_a : a \in A\}, V, N\rangle$, where for each a, R_a is a serial, transitive and Euclidean relation on W and $\geq_a$ is a reflexive, transitive and connected relation over A. We name the model as *bop model.* The truth definition is given by,

$\mathcal{M}, w \models b \succeq_a c$ iff $N(b) \geq_{N(a)} N(c)$.

$\mathcal{M}, w \models B_a\varphi$ iff for all $w' \in W$ such that $wR_{N(a)}w'$, $\mathcal{M}, w' \models \varphi$.

If Professor Calculus is expressed as c, then the earlier situations can now be expressed as:

$$((a \succ_c b) \wedge [+]a : \varphi \wedge [-]b : \varphi) \to \langle B_c \rangle\varphi$$
$$((b \succ_c a) \wedge [+]a : \varphi \wedge [-]b : \varphi) \to \langle B_c \rangle\neg\varphi$$

Let us now provide a complete axiomatization for LOB^-.

Theorem 3.5 *LOB is sound and complete with respect to countable named bop models and its validities are axiomatized by,*

a) LO axioms and rules

b) if $\vdash \varphi$ *then* $\vdash B_a\varphi$

c) belief axioms:

$$B_a(p \to q) \to (B_a p \to B_a q)$$
$$\langle B_a \rangle @_i p \to @_i p$$
$$B_a i \to \langle B_a \rangle i$$
$$B_a i \to B_a B_a i$$
$$\neg B_a i \to B_a \neg B_a i$$

d) paste rule:

if $\vdash @_i \langle B_a \rangle j \wedge @_j \psi \to \varphi$, *then* $\vdash @_i \langle B_a \rangle \psi \to \varphi$, *where,* $i \neq j$ *and* j *not occurring in* φ *or* ψ

d) belief order axioms:

$$b \succeq_a b$$
$$(b \succeq_a c) \wedge (c \succeq_a d) \to (b \succeq_a d)$$
$$(b \succeq_a c) \vee (c \succeq_a b)$$
$$(b \succeq_a c) \leftrightarrow @_i(b \succeq_a c).$$

Proof. *Soundness of the axioms and rules is once again straightforward. For the completeness proof, which is an extension of that of LO, see section B.* ■

It is evident that the logic LOB^- is suitable for expressing communication situations, but to reason in these situations we definitely need some axioms or rules that brings out the correspondence between opinions, preferences and beliefs, which this logic does not provide for.

In what follows we make some preliminary ventures into providing bop models with different properties depicting some such intuitive correspondences and finding out the logical validities that follow from such properties of the models.

Consider $\mathcal{M} = \langle W, A, R^+, R^-, O^+, O^-, \{R_a : a \in A\}, \{\geq_a : a \in A\}, V, N \rangle$ to be a *bop model*. Let us now list some intuitive interactive properties in this model, with the validities they ensure.

- Positive and negative opinions may lead to preferences: $b \in Ran_{O^+}(a)$ and $c \in Ran_{O^-}(a)$ imply $b \geq_a c$.

  validity: $(+)a : b \wedge (-)a : c \to b \succeq_a c$.

- Preferences help in decision making under conflicting opinions: $a \geq_c b$ implies for each $w \in W$, $\exists w' \in (Ran_{R^+}(a) \backslash Ran_{R^-}(b))$, such that $wR_c w'$, and $\exists w'' \in (Ran_{R^-}(a) \backslash Ran_{R^+}(b))$, such that $wR_c w''$.

  validities: $(a \succ_c b \wedge [+]a : \varphi \wedge [-]b : \varphi) \to \langle B_c \rangle \varphi$,
  $(b \succ_c a \wedge [+]a : \varphi \wedge [-]b : \varphi) \to \langle B_c \rangle \neg \varphi$.

It can be easily seen that we can model the communication situation expressed in the example in section 3.1, and reason about it in *bop models* satisfying one of the conditions given earlier. But this is a very simple case of interaction between opinions, belief and preferences. As is evident, a careful and systematic study remains to be done to bring out all such plausible interactive properties. This also gives rise to the following interesting issue:

Question What would be a complete axiomatization of the full interactive logic of beliefs, opinions and preferences?

To give the readers an idea about what we intend to model, let us refer back to the example in the previous subsection, and make it a little more complicated:

> While Professor Calculus was mulling over whether to believe Thomson more than Thompson, Captain Haddock enters the scene and observes:
>
> `Haddock`: "*Thomson and Thompson are a pair of jokers. Do not rely on any of them*"

Obviously, if Calculus had a high opinion of Haddock, then he will not be able to come to any conclusion regarding the singing quality of Bianca, whatever Thomson and Thompson say. The model should be able to reason about these detailed intricacies also.

4 Conclusion

The paper brings out the 'micro-structure' of the communication networks describing the mutual influences of the agents over themselves as well as the events involved, together with their effect on the different epistemic attitudes of the agents, viz. beliefs, opinions, preferences, in a logical framework. Thus, a first attempt is made towards capturing the hitherto unheralded territory of the mutual enhancing and dampening effects of communication between agents from the logical point of view, which have been aptly described in the connectionist approaches.

It is natural, for example, to ask for opinions to lead to preferences, that is, if an agent a has a positive opinion about another, say b but a negative opinion about some other c, then she should prefer b over c. Such properties are reasonable when considering the outcome of the merging, and they can occur anywhere in the merging process, be as an initial information, or an intermediate step or as a final outcome. As it stands, LOB^- describes a very general (logical) framework with which we can express *static* opinions, preferences and beliefs of agents, without particular restrictions on the relations between them. What we need to do now is to find a *complete* list of desired interactive properties between beliefs, opinions and preferences (belief strengths) that the bop model should satisfy, and then add a *dynamic* component in the existing framework to take care of the merging process. This dynamic component should modify the model (i.e., should modify opinions, preferences and beliefs) with reasonable operations in order to get all the desired properties mentioned above. As in the works of [26, 27] (in connectionist systems, as well), the dynamicity does not need to be just one action, but a collective iteration of operations; that will end whenever we have reached a model with the desired properties, that is, an iteration of some model-changing operations that end in a *stable* situation.

Clearly enough, this is not the only kind of dynamics that can be studied. Consistent opinions, beliefs and preferences may also change due to new incoming information. One of the many possibilities is to introduce *suggestions*; announcements about beliefs made by some agent: if Professor Calculus has a good opinion about Captain Haddock, then any suggestion of the latter will modify the beliefs of the former. Some further avenues of investigation are mentioned below.

Aggregation of preferences Together with individual preferences, the notion of group preferences as introduced in [28], which is an emerging area of study can well be incorporated into the framework provided and will add in the expressiveness of the logic. A detailed study of the various notions of preferences as well as their properties and their effects in communication with a focus on the interplay of preferences and beliefs in such situations is also a project worth pursuing.

Dynamic notions We have already talked about introducing the notion of dynamicity to the framework in the form of 'suggestions' which affect the opinions of agents regarding certain events. A host of existing notions of dynamicity can be added to the language of LOB^- so as to study their effects. To name a few, the notions of soft beliefs [51], preference upgrade [54], lexicographic upgrade [51], elite change [51], agent promotion [28] and others. One can also study the effects of the change in opinions about agents, change in opinion about events on that of agents and vice versa which will bring this study closer to that in [27].

A Completeness for *LO*

The completeness proof follows the idea of that appearing in chapter 7.3 [10]. Let Σ be a consistent set of formulas of LO, with $PROP$, NOM and AG the sets of atomic propositions, nominals and agent-names, respectively. We extend Σ to a named and pasted maximal consistent set Σ^+. We define the sets of formulas Δ_i (for $i \in NOM$) and Δ_a (for $a \in AG$) as follows:

$$\begin{array}{lcl} \Delta_i & := & \{\varphi \mid @_i\varphi \in \Sigma^+\} \\ \Delta_a & := & \{\langle + \rangle a : \varphi \mid \langle + \rangle a : \varphi \in \Sigma^+\} \cup \{\langle - \rangle a : \varphi \mid \langle + \rangle a : \varphi \in \Sigma^+\} \cup \\ & & \{(+)a : b \mid (+)a : b \in \Sigma^+\} \cup \{(-)a : b \mid (-)a : b \in \Sigma^+\} \end{array}$$

We call Δ_i a *named world yielded by* Σ^+, and Δ_a a *named agent yielded by* Σ^+. Note that for each $a \in AG$, Δ_a is non-empty, because of the *gen* and *ser* axioms. So for different agent-names a and b, Δ_a and Δ_b are different sets."

From Σ^+, we build the model $\mathcal{M} = \langle W, A, R^+, R^-, O^+, O^-, V, N \rangle$ as follows:

- $W := \{\Delta_i \mid \Delta_i$ is a named world yielded by $\Sigma^+\}$,
- $A := \{\Delta_a \mid \Delta_a$ is a named agent yielded by $\Sigma^+\}$,
- $\Delta_a R^+ \Delta_i$ iff for all non-opinion formulas φ, $\varphi \in \Delta_i$ implies $\langle + \rangle a : \varphi \in \Delta_a$,
- $\Delta_a R^- \Delta_i$ iff for all non-opinion formulas φ, $\neg\varphi \in \Delta_i$ implies $\langle - \rangle a : \varphi \in \Delta_a$,
- $\Delta_a O^+ \Delta_b$ iff $(+)a : b \in \Delta_a$,
- $\Delta_a O^- \Delta_b$ iff $(-)a : b \in \Delta_a$,
- $V(x) := \{\Delta_i \in W \mid x \in \Delta_i\}$ for $x \in (PROP \cup NOM)$,
- $N(a) := \Delta_a$ for $a \in AG$.

Note that W and A are disjoint. Also, because of the axioms, V is an hybrid valuation assigning a singleton to every nominal i (see definition 7.26 of [10]); moreover, N is an injection (in fact, a bijection in this case). We first show that R^+ and R^- are serial. Then we move onto proving existence lemmas for these relations and finally the truth lemma.

R^+ is serial: Consider any Δ_a. Since it is non-empty, we have that, for some formula φ, $@_i\langle + \rangle a : \varphi \in \Sigma^+$ (by *agreement* axiom and definition of Δ_a). So, as Σ^+ is *pasted*, $@_i\langle + \rangle a : j \wedge @_j\varphi \in \Sigma^+$ for some j. Then, $\langle + \rangle a : j \in \Sigma^+$. Now consider any propositional variable, nominal or their boolean combination, ψ in Δ_j. Then $@_j\psi \in \Sigma^+$. This implies that $\langle + \rangle a : \psi \in \Sigma^+$, (by *+translation* axiom) and so, $\langle + \rangle a : \psi \in \Delta_a$. So, $\Delta_a R^+ \Delta_j$.

R^- is serial: Consider any Δ_a. Since it is non-empty, we have that, for some formula φ, $@_i\langle - \rangle a : \varphi \in \Sigma^+$ (by *agreement* axiom and definition of Δ_a). So, as Σ^+ is *pasted*, $@_i\langle - \rangle a : j \wedge @_j\varphi \in \Sigma^+$ for some j. Then, $\langle - \rangle a : j \in \Sigma^+$. Now consider any propositional variable, nominal or their boolean combination, ψ such that $\neg\psi \in \Delta_j$. Then $@_j\neg\psi \in \Sigma^+$. This implies that $\langle - \rangle a : \psi \in \Sigma^+$, (by *–translation* axiom) and so, $\langle - \rangle a : \psi \in \Delta_a$. So, $\Delta_a R^- \Delta_j$.

The proofs of the existence lemmas for both R^+ and R^- follow similarly with the additional fact to note that, if $@_j\varphi \in \Sigma^+$, then $\varphi \in \Delta_j$. We now prove the truth lemma.

The clauses for atomic propositions, nominals, negation, conjunctions and the satisfaction operator @ are standard. For opinion formulas, we have

$$\begin{array}{llll} \mathcal{M}, \Delta_i \models \langle + \rangle a : \varphi & \text{iff} & \exists\, \Delta_j \in W \text{ s.t. } N(a) R^+ \Delta_j \text{ and } \mathcal{M}, \Delta_j \models \varphi & \text{def. of} \models \\ & \text{iff} & \exists\, \Delta_j \in W \text{ s.t. } \Delta_a R^+ \Delta_j \text{ and } \varphi \in \Delta_j & \text{def. of } N \text{ and I.H.} \\ & \text{iff} & \langle + \rangle a : \varphi \in \Delta_a & \text{def. of } R^+ \text{ and } \textit{existence lemma} \\ & \text{iff} & \langle + \rangle a : \varphi \in \Sigma^+ & \text{def. of } \Delta_a \\ & \text{iff} & @_i\langle + \rangle a : \varphi \in \Sigma^+ & \text{agree axiom} \\ & \text{iff} & \langle + \rangle a : \varphi \in \Delta_i & \text{def. of } \Delta_i \end{array}$$

$$\begin{array}{llll}
\mathcal{M}, \Delta_i \models \langle - \rangle a : \varphi & \text{iff} & \exists\, \Delta_j \in W \text{ s.t. } N(a) R^- \Delta_j \text{ and } \mathcal{M}, \Delta_j \models \neg\varphi & \text{def. of } \models \\
& \text{iff} & \exists\, \Delta_j \in W \text{ s.t. } \Delta_a R^- \Delta_j \text{ and } \neg\varphi \in \Delta_j & \text{def. of } N \text{ and I.H.} \\
& \text{iff} & \langle - \rangle a : \varphi \in \Delta_a & \text{def. of } R^- \text{ and } \textit{existence lemma} \\
& \text{iff} & \langle - \rangle a : \varphi \in \Sigma^+ & \text{def. of } \Delta_a \\
& \text{iff} & @_i \langle - \rangle a : \varphi \in \Sigma^+ & \text{agree axiom} \\
& \text{iff} & \langle - \rangle a : \varphi \in \Delta_i & \text{def. of } \Delta_i
\end{array}$$

$$\begin{array}{llll}
\mathcal{M}, \Delta_i \models (+) a : b & \text{iff} & N(a) O^+ N(b) & \text{def. of } \models \\
& \text{iff} & \Delta_a O^+ \Delta_b & \text{def. of } N \\
& \text{iff} & (+) a : b \in \Delta_a & \text{def. of } O^+ \\
& \text{iff} & (+) a : b \in \Sigma^+ & \text{def. of } \Delta_a \\
& \text{iff} & @_i (+) a : b \in \Sigma^+ & \text{agree. axiom} \\
& \text{iff} & (+) a : b \in \Delta_i & \text{def. of } \Delta_i
\end{array}$$

$$\begin{array}{llll}
\mathcal{M}, \Delta_i \models (-) a : b & \text{iff} & N(a) O^- N(b) & \text{def. of } \models \\
& \text{iff} & \Delta_a O^- \Delta_b & \text{def. of } N \\
& \text{iff} & (-) a : b \in \Delta_a & \text{def. of } O^- \\
& \text{iff} & (-) a : b \in \Sigma^+ & \text{def. of } \Delta_a \\
& \text{iff} & @_i (-) a : b \in \Sigma^+ & \text{agree. axiom} \\
& \text{iff} & (-) a : b \in \Delta_i & \text{def. of } \Delta_i
\end{array}$$

Because of the name and paste rules as well as the seriality of R^+ and R^-, the model $\mathcal{M}$ is a named opinion model. Now, we also have that $\Sigma^+ = \Delta_k$ for some nominal k. Hence, by truth lemma, $\mathcal{M}, \Delta_k \models \Sigma$, so our original consistent set of formulas is satisfiable.

B Completeness for LOB^-

As before, let Σ be a consistent set of formulas of LOB, with $PROP$, NOM and AG the sets of atomic propositions, nominals and agent-names, respectively. We extend Σ to a named and pasted maximal consistent set Σ^+. We define the sets of formulas Δ_i as before, while Δ_a (for $a \in AG$) is defined as follows:

$$\begin{array}{lll}
\Delta_a & := & \{\langle + \rangle a : \varphi \mid \langle + \rangle a : \varphi \in \Sigma^+\} \cup \{\langle - \rangle a : \varphi \mid \langle + \rangle a : \varphi \in \Sigma^+\} \cup \\
& & \{(+) a : b \mid (+) a : b \in \Sigma^+\} \cup \{(-) a : b \mid (-) a : b \in \Sigma^+\} \cup \\
& & \{b \succeq_a c \mid b \succeq_a c \in \Sigma^+\}
\end{array}$$

From Σ^+, we build the model $\mathcal{M} = \langle W, A, R^+, R^-, O^+, O^-, \{R_{\Delta_a} : \Delta_a \in A\}, \{\geq_{\Delta_a} : \Delta_a \in A\}, V, N \rangle$, as earlier with R_{Δ_a}, and $\geq_{\Delta_a}$ (for each Δ_a) defined as follows:

- $\Delta_i R_{\Delta_a} \Delta_j$ iff for all formulas φ, $\varphi \in \Delta_j$ implies $\langle B_a \rangle \varphi \in \Delta_i$,
- $\Delta_b \geq_{\Delta_a} \Delta_c$ iff $b \succeq_a c \in \Delta_a$.

The proof of existence lemma for the R_{Δ_a}'s is similar to that in Lemma 7.27 in [10]. The reflexivity, transitivity and connectedness of the relations $\geq_{\Delta_a}$ follow from the definition of the corresponding Δ_a's. Let us now focus on the remaining part of the truth lemma.

$$\begin{array}{llll}
\mathcal{M}, \Delta_i \models \langle B_a \rangle \varphi & \text{iff} & \exists\, \Delta_j \in W \text{ s.t. } \Delta_i R_{N(a)} \Delta_j \text{ and } \mathcal{M}, \Delta_j \models \varphi & \text{def. of } \models \\
& \text{iff} & \exists\, \Delta_j \in W \text{ s.t. } \Delta_i R_{\Delta_a} \Delta_j \text{ and } \varphi \in \Delta_j & \text{def. of } N \text{ and I.H.} \\
& \text{iff} & \langle B_a \rangle \varphi \in \Delta_i & \text{def. of } R_{\Delta_a} \text{ and } \textit{existence lemma}
\end{array}$$

$$\begin{array}{llll}
\mathcal{M}, \Delta_i \models b \succeq_a c & \text{iff} & N(b) \geq_{N(a)} N(c) & \text{def. of } \models \\
& \text{iff} & \Delta_b \geq_{\Delta_a} \Delta_c & \text{def. of } N \\
& \text{iff} & b \succeq_a c \in \Delta_a & \text{def. of } \geq_{\Delta_a} \\
& \text{iff} & b \succeq_a c \in \Sigma^+ & \text{def. of } \Delta_a \\
& \text{iff} & @_i (b \succeq_a c) \in \Sigma^+ & \text{axiom} \\
& \text{iff} & b \succeq_a c \in \Delta_i & \text{def. of } \Delta_i
\end{array}$$

Completeness follows as in the case LO, noting the fact that, any set of pure formulas Σ (i.e. formulas without propositional variables), when added to an extension of $\mathcal{K}^{+}_{\mathcal{H}(@)}$ [49], becomes complete with respect to the frames definable by Σ.

References

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

[2] Hajnal Andréka, Mark D. Ryan, and Pierre-Yves Schobbens. Operators and laws for combining preference relations. *Journal of Logic and Computation*, 12(1):13–53, 2002.

[3] Carlos Areces and Balder ten Cate. Hybrid logics. In Patrick Blackburn, Johan van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical reasoning*. Elsevier, 2005.

[4] Kenneth J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58:328–346, 1950.

[5] Christian Balkenius and Peter Gärdenfors. Nonmonotonic inferences in neural networks. In J. A. Allen, R. Fikes, and E. Sandewall, editors, *Principles of knowledge representation and reasoning*, pages 32–39, San mateo, CA, 1991. Morgan Kaufmann.

[6] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements, common knowledge and private suspicious. Technical Report SEN-R9922, CWI, Amsterdam, 1999.

[7] Alexandru Baltag and Sonja Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 165:5–21, 2006.

[8] Alexandru Baltag and Sonja Smets. Dynamic belief revision over multi-agent plausibility models. In Wiebe van der Hoek and Michael Wooldridge, editors, *Proceedings of LOFT 2006*, pages 11–24, Liverpool, UK, 2006. University of Liverpool. Available at http://www.vub.ac.be/CLWF/SS/loft.pdf.

[9] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, *Logic and the Foundation of Game and Decision Theory (LOFT7)*, volume 3 of *Texts in Logic and Games*, pages 13–60. Amsterdam University Press, 2008.

[10] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. CUP, 2001.

[11] Reinhard Blutner and Paul David Doherty. Nonmonotonic logic and neural networks. Available at http://citeseer.ist.psu.edu/160749.html.

[12] Artur d'Avila Garcez, Krysia Broda, and Dov Gabbay. Symbolic knowledge extraction from trained neural networks: A sound approach. *Artificial Intelligence*, 125(1–2):155–207, January 2001.

[13] Johan de Kleer. An assumption-based truth maintenance system. *Artificial Intelligence*, 28(2):127–162, 1986.

[14] Johan de Kleer. A perspective on assumption-based truth maintenance. *Artificial Intelligence*, 59(1-2):63–67, 1993.

[15] Maarten de Rijke. *Extending Modal Logic*. PhD thesis, Institute for Logic, Language and Computation (Universiteit van Amsterdam), 1993. ILLC Dissertation Series DS-1993-04.

[16] Maarten de Rijke. Meeting some neighbours. In Jan van Eijk and Albert Viser, editors, *Logic and Information Flow*, pages 170–196. The MIT Press, Cambridge, MA, 1994.

[17] Robert Demolombe and Emiliano Lorini. A logical account of trust in information sources. In R. Falcone, S. Barber, J. Sabater-Mir, and M. P. Singh, editors, *Eleventh International Workshop on Trust in Agent Societies*, LNAI, 2008. To appear.

[18] Jon Doyle. A truth maintenance system. *Artificial Intelligence*, 12(3):231–272, 1979.

[19] Daniel Eckert and Gabriella Pigozzi. Belief merging, judgment aggregation and some links with social choice theory. In James P. Delgrande, Jérôme Lang, Hans Rott, and Jean-Marc Tallon, editors, *Belief Change in Rational Agents*, volume 05321 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2005.

[20] André Fuhrmann. Reflective modalities and theory change. *Synthese*, 81(1):115–134, October 1989.

[21] Dov M. Gabbay, Gabriella Pigozzi, and Odinaldo Rodrigues. Common foundations for belief revision, belief merging and voting. In Giacomo Bonanno, James P. Delgrande, Jérôme Lang, and Hans Rott, editors, *Formal Models of Belief Change in Rational Agents*, volume 07351 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2007.

[22] Haim Gaifman. Operational pointer semantics: solution to self-referential puzzles i. In *TARK '88: Proceedings of the 2nd conference on Theoretical Aspects of Reasoning about Knowledge*, pages 43–59, San Francisco, CA, USA, 1988. Morgan Kaufmann Publishers Inc.

[23] Haim Gaifman. Pointers to truth. *Journal of Philosophy*, 5(89):223–261, 1992.

[24] Jelle Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, Institute for Logic, Language and Computation (Universiteit van Amsterdam), 1999. ILLC Dissertation Series DS-1999-01.

[25] Jelle Gerbrandy. Dynamic epistemic logic. In Lawrence S. Moss, Jonathan Ginzburg, and Maarten de Rijke, editors, *Logic, language and computation*, volume 2, pages 67–84. Center for the Study of Language and Information, Stanford, CA, USA, 1999.

[26] Sujata Ghosh, Benedikt Löwe, and Erik Scorelle. Belief flow in assertion networks. In Uta Priss, Simon Polovina, and Richard Hill, editors, *Proceedings of the 15th International Conference on Conceptual Structures (ICCS 2007), Sheffield, UK*, volume 4604 of *LNAI*, pages 401–414, Heilderberg, July 2007. Springer-Verlag.

[27] Sujata Ghosh and Fernando R. Velázquez-Quesada. Expressing belief flow in Assertion Networks.

[28] Patrick Girard. *Modal Logic for Belief and Preference Change*. PhD thesis, Department of Philosophy (Stanford University), Stanford, CA, USA, February 2008. ILLC Dissertation Series DS-2008-04.

[29] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, May 1988.

[30] A. Gupta and Nuel D. Belnap. *The Revision Theory of Truth*. The MIT Press, Cambridge, MA, 1993.

[31] Sven Ove Hansson. A survey of non-prioritized belief revision. *Erkenntnis*, 50(2-3):413–427, 1999.

[32] William L. Harper. Rational conceptual change. In *Proceedings of the Meeting of the Philosophy of Science Association (PSA 1976)*, volume 2, pages 462–494. Philosophy of Science Association, 1977.

[33] Hans G. Herzberger. Naive semantics and the liar paradox. *The Journal of Philosophy*, 79:479–497, 1982.

[34] Sébastien Konieczny. *Sur la logique du changement: révision et fusion de bases de connaissance.* PhD thesis, Laboratoire d'Informatique Fondamentale de Lille, Universit de Lille 1, November 1999.

[35] Sébastien Konieczny and Ramón Pino-Pérez. On the logic of merging. In *Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 488–498, 1998.

[36] Sébastien Konieczny and Ramón Pino-Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

[37] Sébastien Konieczny and Ramón Pino-Pérez. Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory. *European Journal of Operational Research*, 160(3):785–802, 2005.

[38] Hanner Leitgeb. Interpreted dynamical systems and qualitative laws: from neural networks to evolutionary systems. *Synthese*, 146(1-2):189–202, August 2005.

[39] Hannes Leitgeb and Krister Segerberg. Dynamic doxastic logic: why, how, and where to? *Synthese*, 155(2):167–190, 2007.

[40] Isaac Levi. Subjunctives, dispositions, and chances. *Synthese*, 34:423–455, 1977.

[41] Isaac Levi. *The Enterprise of Knowledge.* MIT Press, Cambridge, MA, 1980.

[42] David K. Lewis. *Counterfactuals.* Blackwell, Cambridge, Massachusetts, 1973.

[43] Emiliano Lorini and Robert Demolombe. Trust and norms in the context of computer security: toward a logical formalization. In R. Van der Meyden and L. Van der Torre, editors, *International Workshop on Deontic Logic in Computer Science (DEON), Luxembourg 2008.* Springer-Verlag, 2008. To appear.

[44] T. Meyer, A. K. Ghose, and S. Chopra. Non-prioritized ranked belief change. *Journal of Philosophical Logic*, 32(4):117–143, 2003.

[45] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems : Networks of Plausible Inference.* Morgan Kaufmann, Santa Mateo, CA, USA, September 1988.

[46] J. A. Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.

[47] Krister Segerberg. Belief revision from the point of view of doxastic logic. *Logic Journal of IGPL*, 3(4):535–553, 1995.

[48] Krister Segerberg. Two traditions in the logic of belief: bringing them together. In Hans Jürgen Ohlbach and Uwe Reyle, editors, *Logic, Language and Reasoning. Essays in Honour of Dov Gabbay*, pages 135–147. Kluwer Academic Publishers, 1999.

[49] Balder ten Cate. *Model theory for extended modal languages.* PhD thesis, Institute for Logic, Language and Computation (Universiteit van Amsterdam), 2005. ILLC Dissertation Series DS-2005-01.

[50] Johan van Benthem. Semantic parallels in natural language and computation. In *Logic Colloquium '87*, pages 331–375, Amsterdam, North-Holland, 1989. Elsevier Science Publishers.

[51] Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.

[52] Johan van Benthem. Priority product update as social choice. Available at http://dare.uva.nl/document/93918, 2007.

[53] Johan van Benthem, Patrick Girard, and Olivier Roy. Everything else being equal: a modal logic approach to ceteris paribus preferences. *Journal of Philosophical Logic*, 2008. To appear.

[54] Johan van Benthem and Fenrong Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 17(2):157–182, 2007.

[55] Johan van Benthem, Sieuwert van Otterloo, and Olivier Roy. Preference logic, conditionals and solution concepts in games. In *Festschrift for Krister Segerberg*. University of Uppsala, 2005.

[56] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library Series*. Springer, 2007.

[57] Jon Williamson. *Bayesian nets and causality: philosophical and computational foundations.* Oxford University Press, 2005.