Imprecise Bipolar Belief Measures

Jonathan Lawry

Intelligent Systems Laboratory
University of Bristol

29 June, 2010

Overview

What is Bipolarity?

- Bipolarity refers to the fact that human decision making and reasoning is often based on both positive and negative information.
- This is often results in beliefs being divided into positive, negative and neutral zones on some underlying scale.
- For example, people may make choices by separately considering the advantages and disadvantages of the various alternatives.
- Choices may then be made of the bases of the relative weighting of positive and negative aspects.
- Results in cognitive psychology have pointed out the importance of bipolar reasoning in human activities.
- Indeed there is some evidence to suggest that positive and negative information are not processed in the same area of the brain.

The Bipolarity of Assertability

- As communicating agents our use of natural language requires us to make constant decisions about the assertability of propositions and sentences.
- Assertability is governed by the social practices and conventions which underpin the use of language across a population.
- Knowledge concerning the assertability of propositions is largely acquired in a distributed manner through communications and interactions with others.
- The notion of 'assertability' itself would seem to be inherently bipolar.
- This manifests itself in a distinction between those propositions which convention would deem clearly assertable, and those which convention would not classify as incorrect, or perhaps even dishonest, to assert.

A Bipolarity Example

- Consider a witness in a court of law describing a suspect as being *tall*.
- Depending on the actual height of the suspect this statement may be deemed as clearly true or clearly false, in which latter case the witness could be accused of perjury.
- However, there will also be an intermediate height range for which, while there may be doubt and differing opinions concerning the use of the description *tall*, it would not be deemed as definitely inappropriate and hence the witness would not be viewed as committing perjury.
- In other words, there may be cases where you can *get away* with using the word tall.

Exploiting Bipolarity

- Bipolarity helps us to utilize vagueness in situations where our interests and those of other agents with whom we communicate *are not the same*.
- Suppose a politician promises voters *a major reduction in the deficit*
- Most voters would then expect a reduction that could definitely be described as a major reduction.
- However, at the following election the politician can only be accused of not delivering on his promise if the reduction is definitely not a major reduction.
- This gives the politician considerable flexibility in allowing a range of borderline reductions which can acceptably be referred to as major reductions.

Preliminaries

- Let L be a language of the propositional logic with connectives $\wedge$, $\vee$ and $\neg$ and a finite set of propositional variables $P = \{p_1, \ldots, p_n\}$.
- Let SL denote the sentences of L.
- To model the bipolarity of assertability we introduce the notion of a valuation pair defined on SL.
- This consists of two binary functions $\underline{v}$ and $\overline{v}$ linked, through negation, by a duality relationship.
- The underlying idea is that $\underline{v}$ represents the strong criteria of definite assertability while $\overline{v}$ represents the weaker criteria of acceptable assertability.

Valuation Pairs

- A valuation pair for L is a pair of functions $\vec{v} = (\underline{v}, \overline{v})$ such that $\underline{v} : SL \rightarrow \{0,1\}$, $\overline{v} : SL \rightarrow \{0,1\}$ and where $\forall \theta \in SL$, $\underline{v}(\theta) = 1$ iff it is *correct* to assert θ and $\overline{v}(\theta) = 1$ iff it is *acceptable* to assert θ. Furthermore, $\underline{v}$ and $\overline{v}$ satisfy the following properties: $\forall \theta, \varphi \in SL$
- $\underline{v} \leq \overline{v}$ (*Coherence*)
- $\underline{v}(\theta \wedge \varphi) = \min(\underline{v}(\theta), \underline{v}(\varphi))$, $\overline{v}(\theta \wedge \varphi) = \min(\overline{v}(\theta), \overline{v}(\varphi))$
- $\underline{v}(\theta \vee \varphi) = \max(\underline{v}(\theta), \underline{v}(\varphi))$, $\overline{v}(\theta \vee \varphi) = \max(\overline{v}(\theta), \overline{v}(\varphi))$
- $\underline{v}(\neg\theta) = 1 - \overline{v}(\theta)$, $\overline{v}(\neg\theta) = 1 - \underline{v}(\theta)$
- The last rule is motivated by the assumption that it is *definitely correct* to assert $\neg\theta$ if and only if it is *not acceptable* to assert θ.

Connections to Kleene's Three-Valued Logic

- For (non-classical) valuation pairs satisfying coherence there are three possible values of $(\underline{v}(\theta), \overline{v}(\theta))$ for any $\theta \in SL$.
- These are $\mathbf{1} = (1, 1)$, $\mathbf{b} = (0, 1)$ and $\mathbf{0} = (0, 0)$ where $\mathbf{1}$ denotes definitely assertable, $\mathbf{b}$ denotes borderline assertable, and $\mathbf{0}$ denotes definitely not assertable.

θ	φ	$\neg\theta$	$\theta \wedge \varphi$	$\theta \vee \varphi$
1	**1**	**0**	**1**	**1**
b	**1**	**b**	**b**	**1**
0	**1**	**1**	**0**	**1**
1	**b**	**0**	**b**	**1**
b	**b**	**b**	**b**	**b**
0	**b**	**1**	**0**	**b**
1	**0**	**0**	**0**	**1**
b	**0**	**b**	**0**	**b**
0	**0**	**1**	**0**	**0**

Introducing Uncertainty

- Within the proposed bipolar framework, uncertainty concerning the sentences in L effectively corresponds to uncertainty as to which is the correct valuation pair for L.
- In practice, there are likely to be many different sources of this uncertainty, however one natural division of uncertainty types is as follows:
- Uncertainty about the linguistic conventions governing assertability for sentences in L.
- *E.g., an agent may be uncertain as to whether or not a proposition such as 'Bill is tall' is definitely or acceptably assertable even if Bill's height is known precisely.*
- Uncertainty arising from a lack of knowledge concerning the referents of sentences.
- *E.g., being uncertain about Bill's height in the proposition 'Bill is tall' or the velocity of the car in the proposition 'the car was fast'*

Bipolar Belief Measures

- Let V be the set of all valuation pairs on L.
- Let w be a probability distribution defined on V so that $w(\underline{v}, \overline{v})$ is the agent's subjective belief that $(\underline{v}, \overline{v})$ is the true valuation pair for L.
- Let $\underline{\mu} : SL \to [0,1]$ such that $\forall\theta \in SL$ $\underline{\mu}(\theta) = w(\{\vec{v} \in V : \underline{v}(\theta) = 1\}) = w(\{\vec{v} \in V : \vec{v}(\theta) = \mathbf{1}\})$
- Let $\overline{\mu} : SL \to [0,1]$ such that $\forall\theta \in SL$ $\overline{\mu}(\theta) = w(\{\vec{v} \in V : \overline{v}(\theta) = 1\}) = w(\{\vec{v} \in V : \vec{v}(\theta) = \{\mathbf{b}, \mathbf{1}\}\})$
- For $\theta \in SL$, $\underline{\mu}(\theta)$ and $\overline{\mu}(\theta)$ correspond to the agent's subjective belief that θ is definitely assertable and that θ is acceptable to assert respectively.
- $\forall\theta \in SL$ $\underline{\mu}(\theta) \leq \overline{\mu}(\theta)$
- $\forall\theta \in SL$ $\underline{\mu}(\theta \wedge \neg\theta) = 0$ and $\overline{\mu}(\theta \wedge \neg\theta) = 1$

Pair of Sets Representation of Valuation Pairs

- We can represent a valuation pair by a pair of disjoint subsets of propositions.
- Let $\mathcal{D}^{\vec{v}} = \{p_i \in P : \underline{v}(p_i) = 1\}$
- And let $\mathcal{C}^{\vec{v}} = \{p_i \in P : \underline{v}(\neg p_i) = 1\}$
- For any pair of sets (F, G) where $G \subseteq F^c$ there is a unique valuation pair $\vec{v} \in V$, for which $\mathcal{D}^{\vec{v}} = F$ and $\mathcal{C}^{\vec{v}} = G$. Let this be denoted $\vec{v}_{(F,G)} = (\underline{v}_{(F,G)}, \overline{v}_{(F,G)})$ as follows:
- $\underline{v}_{(F,G)}(p_i) = 1$ iff $p_i \in F$, $\underline{v}_{(F,G)}(\neg p_i) = 1$ iff $p_i \in G$
- Also, by duality $\overline{v}_{(F,G)}(p_i) = 1$ iff $p_i \notin G$ and $\overline{v}_{(F,G)}(\neg p_i) = 1$ iff $p_i \notin F$.
- Let $P = \{p_1, p_2, p_3, p_4\}$ and $(\mathcal{D}^{\vec{v}}, \mathcal{C}^{\vec{v}}) = (\{p_1, p_2\}, \{p_4\})$
- Then $\vec{v}$ is such that $\vec{v}(p_1) = \vec{v}(p_2) = (1, 1)$, $\vec{v}(p_3) = (0, 1)$ and $\vec{v}(p_4) = (0, 0)$.

Introducing Partial Knowledge

- In many scenarios knowledge regarding the assertability conventions of a population will be partial or imprecise.
- For example, monitoring an individual agent involved in a dialogue is only likely to provide partial information about his/her use of assertability conventions.
- From the perspective of valuation pairs this can mean that for certain propositional variables p_j, the values of $\underline{v}(p_j)$ and $\overline{v}(\neg p_j)$ are *unknown* or the values of $\underline{v}(\neg p_j)$ and $\overline{v}(p_j)$ are *unknown*.
- Consequently only lower and upper bounds (relative to $\subseteq$) can be determined for $\mathcal{D}^{\vec{v}}$ and $\mathcal{C}^{\vec{v}}$ which motivates the following definition of *imprecise valuation pairs*.

Imprecise Valuation Pairs

- An imprecise valuation is a pair of lower and upper sets $\Theta = (\langle \underline{F}, \overline{F} \rangle, \langle \underline{G}, \overline{G} \rangle)$ where $\underline{F} \subseteq \overline{F} \subseteq P$ and $\underline{G} \subseteq \overline{G} \subseteq P$, representing the following constraints on a coherent valuation pair $\vec{v}$:

$$\underline{F} \subseteq \mathcal{D}^{\vec{v}} \subseteq \overline{F} \text{ and } \underline{G} \subseteq \mathcal{C}^{\vec{v}} \subseteq \overline{G}$$

- Let $K^{\Theta} = \{(F, G) : \underline{F} \subseteq F \subseteq \overline{F},\ \underline{G} \subseteq G \subseteq \overline{G},\ \text{and } G \subseteq F^c\}$
- For example: Let $P = \{p_1, p_2, p_3\}$ and assume partial knowledge $\underline{v}(p_1) = 1$ and $\overline{v}(p_2) = 0$.
- From this we can infer $\{p_1\} \subseteq \mathcal{D}^{\vec{v}} \subseteq \{p_1, p_3\}$ and $\{p_2\} \subseteq \mathcal{C}^{\vec{v}} \subseteq \{p_2, p_3\}$.
- This is represented by the imprecise valuation pair $\Theta = (\langle \{p_1\}, \{p_1, p_3\} \rangle, \langle \{p_2\}, \{p_2, p_3\} \rangle)$.
- $K^{\Theta} = \{(\{p_1\}, \{p_2\}), (\{p_1, p_3\}, \{p_2\}), (\{p_1\}, \{p_2, p_3\})\}$

Combining Bipolarity, Partial Knowledge and Uncertainty

- Suppose we have a number of possible imprecise valuation pairs $\mathcal{IV} = \{\Theta_1, \ldots, \Theta_k\}$.
- If we then assume a weighting on the members of $\mathcal{IV}$ as represented by a probability distribution w, then this naturally generates lower and upper bipolar belief measures as follows:
- $\underline{\mu}_*(\theta) = w(\{\Theta_i : \forall (F,G) \in K^{\Theta_i} \ \underline{v}_{(F,G)}(\theta) = 1\})$
- $\underline{\mu}^*(\theta) = w(\{\Theta_i : \exists (F,G) \in K^{\Theta_i} \ \underline{v}_{(F,G)}(\theta) = 1\})$
- $\overline{\mu}_*(\theta) = w(\{\Theta_i : \forall (F,G) \in K^{\Theta_i} \ \overline{v}_{(F,G)}(\theta) = 1\})$
- $\overline{\mu}^*(\theta) = w(\{\Theta_i : \exists (F,G) \in K^{\Theta_i} \ \overline{v}_{(F,G)}(\theta) = 1\})$
- Let $\underline{\mu}$ and $\overline{\mu}$ be obtained by replacing each Θ_i with a single valuation pair in K^{Θ_i}.
- Then $\underline{\mu}_*(\theta) \leq \underline{\mu}(\theta) \leq \underline{\mu}^*(\theta)$ and $\overline{\mu}_*(\theta) \leq \overline{\mu}(\theta) \leq \overline{\mu}^*(\theta)$

Some Results

- For any $\theta \in SL$ it holds that:
- $\underline{\mu}_*(\theta) = w(\{\Theta_i : \underline{v}_{(\underline{F}_i,\underline{G}_i)}(\theta) = 1\})$
- $\overline{\mu}_*(\theta) \geq w(\{\Theta_i : \underline{v}_{(\overline{F}_i,\overline{G}_i)}(\theta) = 1\})$
- $\underline{\mu}^*(\theta) \leq w(\{\Theta_i : \underline{v}_{(\overline{F}_i,\overline{G}_i)}(\theta) = 1\})$
- $\overline{\mu}^*(\theta) = w(\{\Theta_i : \overline{v}_{(\underline{F}_i,\underline{G}_i)}(\theta) = 1\})$
- If l is a literal corresponding to p_i or $\neg p_i$ for some $p_i \in P$, then:
- $\overline{\mu}_*(l) = w(\{\Theta_i : \overline{v}_{(\overline{F}_i,\overline{G}_i)}(l) = 1\})$
- $\underline{\mu}^*(l) = w(\{\Theta_i : \underline{v}_{(\overline{F}_i,\overline{G}_i)}(l) = 1\})$

From Imprecise to Precise Valuation Pairs

- In some circumstances it may not be sufficient to determine lower and upper measures from $\mathcal{IV}$ and instead precise values of μ^+ and μ^- may be needed.
- This requires us to make additional assumptions above and beyond the information provided by each imprecise valuation pair.
- We propose three such additional assumptions in order to generate, maximally conservative, maximally positive and maximally negative measures.
- We also investigate the resulting measures obtain by replacing every imprecise valuation pair in $\mathcal{IV}$ by a precise valuation pair according to each of these assumptions.

Maximally Conservative Measures

- Underlying assumption: Adopt the *least decisive* precise valuation pair consistent with the imprecise valuation pair.
- In other words, for any literal l we should assume, where possible, that l is not definitely assertable (i.e. $\underline{v}(l) = 0$) but that it is acceptable to assert l (i.e. $\overline{v}(l) = 1$).
- $\underline{\mu}_{mc}(\theta) = w(\{\Theta_i : \underline{v}_{(\underline{F}_i, \underline{G}_i)}(\theta) = 1\})$
- $\overline{\mu}_{mc}(\theta) = w(\{\Theta_i : \overline{v}_{(\underline{F}_i, \underline{G}_i)}(\theta) = 1\})$
- $\vec{v}_{(\underline{F}_i, \underline{G}_i)}$ is the valuation pair consistent with Θ_i which *minimizes* both the number of propositional variables and negated propositional variables deemed to be definitely assertable.
- By duality this same valuation pair also *maximizes* the number of propositional and negated propositional variables which are deemed acceptable to assert.

Maximally Positive (Negative) Measures

- Maximally Positive Measures:
- $\underline{\mu}_{mp}(\theta) = w(\{\Theta_i : \underline{v}_{(\overline{F}_i, \underline{G}_i)}(\theta) = 1\})$
- $\overline{\mu}_{mp}(\theta) = w(\{\Theta_i : \overline{v}_{(\overline{F}_i, \underline{G}_i)}(\theta) = 1\})$
- Maximally Negative Measures:
- $\underline{\mu}_{mn}(\theta) = w(\{\Theta_i : \underline{v}_{(\underline{F}_i, \overline{G}_i)}(\theta) = 1\})$
- $\overline{\mu}_{mn}(\theta) = w(\{\Theta_i : \overline{v}_{(\underline{F}_i, \overline{G}_i)}(\theta) = 1\})$
- $\vec{v}_{(\overline{F}_i, \underline{G}_i)}$ ($\vec{v}_{(\underline{F}_i, \overline{G}_i)}$) is the valuation pair consistent with Θ_i which simultaneously *maximizes* both the number of propositional variables (negated propositional variables) deemed to be definitely assertable and the number of propositional variables (negated propositional variables) deemed to be acceptable to assert.

Some More Results

- $\forall \theta \in SL$, $\underline{\mu}_{mc}(\theta) = \underline{\mu}_*(\theta)$ and $\overline{\mu}_{mc}(\theta) = \overline{\mu}^*(\theta)$
- $\underline{\mu}_*(p_j) = \underline{\mu}_{mc}(p_j) = \underline{\mu}_{mn}(p_j) = w(\{\Theta_i : p_j \in \underline{F}_i\})$
- $\underline{\mu}^*(p_j) = \underline{\mu}_{mp}(p_j) = w(\{\Theta_i : p_j \in \overline{F}_i\})$
- $\overline{\mu}_*(p_j) = \overline{\mu}_{mn}(p_j) = w(\{\Theta_i : p_j \notin \overline{G}_i\})$
- $\overline{\mu}^*(p_j) = \overline{\mu}_{mc}(p_j) = \overline{\mu}_{mp}(p_j) = w(\{\Theta_i : p_j \notin \underline{G}_i\})$
- $\underline{\mu}_*(\neg p_j) = \underline{\mu}_{mc}(\neg p_j) = \underline{\mu}_{mp}(\neg p_j)$
- $\underline{\mu}^*(\neg p_j) = \underline{\mu}_{mn}(\neg p_j)$
- $\overline{\mu}_*(\neg p_j) = \overline{\mu}_{mp}(\neg p_j)$
- $\overline{\mu}^*(\neg p_j) = \overline{\mu}_{mc}(\neg p_j) = \overline{\mu}_{mn}(\neg p_j)$.

Agent Dialogues

- Let $\mathcal{A} = \{a_1, \ldots, a_m\}$ be a population of communicating agents.
- In this simplified model each agent in $\mathcal{A}$ is able to *assert* propositions or negated propositions from P and to *condemn* or to *agree with* the assertions of others.
- PAS_i denotes the propositional variables asserted by a_i
- NAS_i denotes the negated propositional variables asserted by a_i
- PCD_i denotes the propositional variables condemned by a_i
- NCD_i denotes the negated propositional variables condemned by a_i
- PAG_i denotes the propositional variables agreed with by a_i
- NAG_i denotes the negated propositional variables agreed with by a_i

Learning from Agent Assertions

- It is assumed that the assertions, agreements and condemnations made by each agent a_i is consistent with a consistent valuation pair $\vec{v}_i$ in the following sense: For literal $l = \pm p_j$
- If a_i asserts l then $\overline{v}(l) = 1$
- If a_i condemns another agent for asserting l then $\underline{v}(\neg l) = 1$
- If a_i agrees with another agent for asserting l then $\underline{v}(l) = 1$
- Under these assumptions we have:
- $NCD_i \cup PAG_i \subseteq \mathcal{D}^{\vec{v}_i} \subseteq (NAS_i)^c \cap (PCD_i)^c \cap (NAG_i)^c$
- $PCD_i \cup NAG_i \subseteq \mathcal{C}^{\vec{v}_i} \subseteq (PAS_i)^c \cap (NCD_i)^c \cap (PAG_i)^c$
- The assertions, agreements and condemnations of each agent identifies an imprecise valuation pair.

Example Dialogue

- Let $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$ and let $P = \{p_1, p_2, p_3, p_4\}$. The agents' communications are summarized in the following table.

Agent	PAS_i	NAS_i	PCD_i	NCD_i	PAG_i	NAG_i
a_1	$\{p_1, p_3\}$	$\emptyset$	$\{p_4\}$	$\emptyset$	$\{p_3\}$	$\emptyset$
a_2	$\{p_2\}$	$\{p_4\}$	$\{p_4\}$	$\emptyset$	$\emptyset$	$\{p_4\}$
a_3	$\{p_1\}$	$\{p_4\}$	$\{p_3\}$	$\emptyset$	$\emptyset$	$\{p_2\}$
a_4	$\{p_2, p_4\}$	$\emptyset$	$\emptyset$	$\{p_4\}$	$\{p_4\}$	$\emptyset$

- From this data we generate the following imprecise valuation pairs:
- $\Theta_1 = (\langle\{p_3\}, \{p_1, p_2, p_3\}\rangle, \langle\{p_4\}, \{p_2, p_4\}\rangle)$
- $\Theta_2 = (\langle\emptyset, \{p_1, p_2, p_3\}\rangle, \langle\{p_4\}, \{p_1, p_3, p_4\}\rangle)$
- $\Theta_3 = (\langle\emptyset, \{p_1\}\rangle, \langle\{p_2, p_3\}, \{p_2, p_3, p_4\}\rangle)$
- $\Theta_4 = (\langle\{p_4\}, \{p_1, p_2, p_3, p_4\}\rangle, \langle\emptyset, \{p_1, p_3\}\rangle)$

Example Dialogue: 2

- Now assuming equal weighting to agents results in a uniform distribution w on $\{\Theta_1, \Theta_2, \Theta_3, \Theta_4\}$ we obtain the following lower and upper bipolar belief measures:

Proposition	$\underline{\mu}_*$	$\underline{\mu}^*$	$\overline{\mu}_*$	$\overline{\mu}^*$
p_1	0	1	$\frac{1}{2}$	1
p_2	0	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{3}{4}$
p_3	$\frac{1}{4}$	$\frac{3}{4}$	$\frac{1}{4}$	$\frac{3}{4}$
p_4	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{2}$

- Selecting precise valuation pairs gives:

Proposition	$\underline{\mu}_{mc}$	$\overline{\mu}_{mc}$	$\underline{\mu}_{mp}$	$\overline{\mu}_{mp}$	$\underline{\mu}_{mn}$	$\overline{\mu}_{mn}$
p_1	0	1	1	1	0	$\frac{1}{2}$
p_2	0	$\frac{3}{4}$	$\frac{3}{4}$	$\frac{3}{4}$	0	$\frac{1}{2}$
p_3	$\frac{1}{4}$	$\frac{3}{4}$	$\frac{3}{4}$	$\frac{3}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
p_4	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$

Conclusions

- We have proposed valuation pairs as a possible framework for representing bipolar assertability conventions.
- This approach has been extended by introducing imprecise valuation pairs to represent partially or imprecisely defined assertability conventions.
- Allowing for uncertainty then naturally results in lower and upper bounds on bipolar belief measures.
- The potential of this extended framework has then been illustrated by its application to an idealised dialogue between agents, each making assertions, agreements and condemnation based on their individual valuation pairs.
- Information inferred from this dialogue can then be represented by imprecise valuation pairs together with lower and upper bounds on bipolar belief measures.