Measuring Diversity of Preferences in a Group

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[joint work with Vahid Hashemi]

Motivation

Preferences are ubiquitous in AI: recommendation, configuration, ...
Preference handling in *MAS* is particularly interesting.

Basic intuition:

> *The less diverse the preferences in a group of agents are, the easier it should be to come to mutually acceptable decisions.*

V. Hashemi and U. Endriss. Measuring Diversity of Preferences in a Group. Proc. 21st European Conference on Artificial Intelligence (ECAI-2014).

Outline

- Formal framework, definition of the concept, examples
- Theoretical results: axiomatic method
- Experimental results: frequency of phenomena *modulo* diversity

Formal Framework

Finite set of *alternatives* $\mathcal{X}$. A *preference* is a strict *linear order* on $\mathcal{X}$. $\mathcal{L}(\mathcal{X})$ is the set of all such preference orders.

Each of a finite set of *agents* $\mathcal{N} = \{1, \ldots, n\}$ expresses a preference, giving rise to a *profile* $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(\mathcal{X})^n$.

We propose the concept of *preference diversity index* (PDI) to make judgments about which of two profiles we consider more diverse:

> A *PDI* is a function $\Delta : \mathcal{L}(\mathcal{X})^n \to \mathbb{R}^+ \cup \{0\}$, with $\Delta(\boldsymbol{R}) = 0$ for unanimous profiles $\boldsymbol{R} = (R, \ldots, R) \in \mathcal{L}(\mathcal{X})^n$.

Examples for Specific PDI's

Three options for defining the diversity of a given profile $\boldsymbol{R}$:

- *Simple support-based PDI:* number of distinct preferences in $\boldsymbol{R}$. Generalisation: count, for a given $k \leqslant m$, the number of distinct ordered k-tuples of alternatives appearing in $\boldsymbol{R}$.
- *Distance-based PDI:* measure the distance (e.g., *Kendall tau*) between any two preferences in $\boldsymbol{R}$ and then aggregate the values obtained (e.g., by computing their sum or their maximum).
- *Compromise-based PDI:* first aggregate the individual preferences (e.g., using the *Borda rule*), then compute the Kendall tau distance of each individual preference to that "compromise", and finally aggregate (e.g., add) the values obtained.

Theoretical Results

Adopt the *axiomatic method* from social choice theory to formulate and explore desirable properties of PDI's ...

- *Weak discernability:* only unanimous profiles have diversity 0
- *Anonymity* (A): order of agents does not matter
- *Neutrality* (N): order of alternatives does not matter
- *Strong discernability:* no two equally diverse profiles (unless for A/N)
- *Support invariance:* equally diverse if $\{R_1, \ldots, R_n\} = \{R'_1, \ldots, R'_n\}$
- *Independence:* $\boldsymbol{R} \succcurlyeq \boldsymbol{R'}$ implies $\boldsymbol{R} \oplus R \succcurlyeq \boldsymbol{R'} \oplus R$ for "new" R

Two results:

- **Proposition:** For $n > m!$ and $m > 2$, there can be *no* PDI that is both support-invariant and *strongly discernable*.
- **Proposition:** A PDI is weakly discernible, support-invariant, and *independent* if and only if it is the *simple support-based PDI*.

Experimental Results

Setup:

- distance-based PDI (Kendall tau, sum)
- 50 agents / 5 alternatives / 1M profiles
- x-axis: diversity (from 0 to max)

Right:

- Condorcet winners / cycles
- agreement between voting rules
- voter satisfaction

Bottom:

- y-axis: frequency of diversity x
- impartial culture *vs.* real data (AGH)

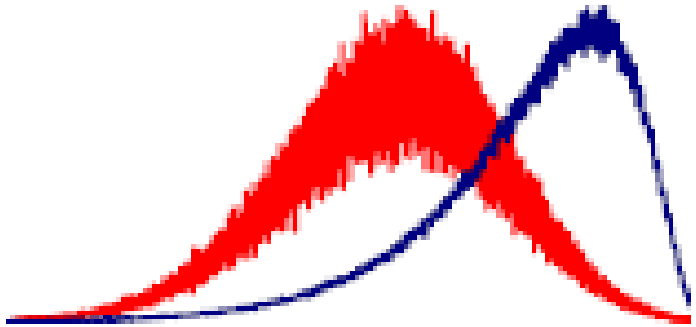

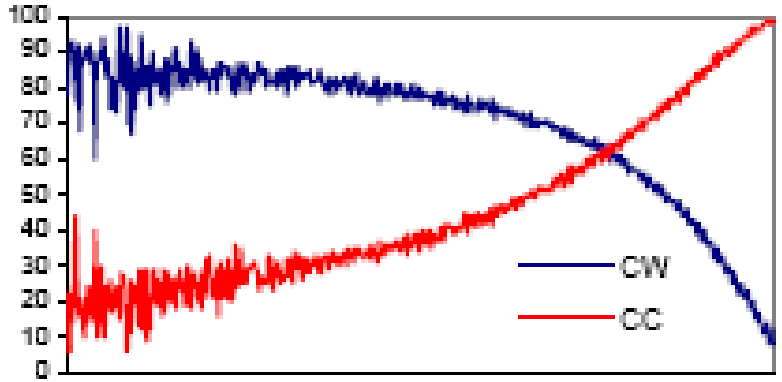

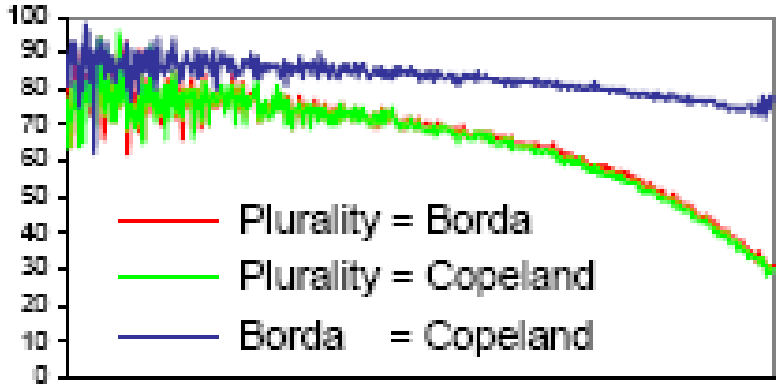

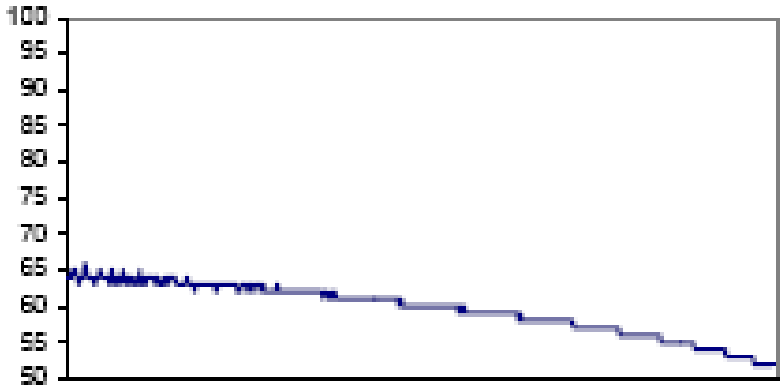

Last Slide

We have introduced a formal framework for studying the important concept of diversity of preferences in a group.

- *General* notion of PDI (preference diversity index)
- Three proposals for (families) of *specific* PDI's
- *Axioms* for "good" PDI's (*impossibility* / *characterisation* results)
- Experimental results: *discern real* from *synthetic data*
- Experimental results: *choice-theoretic effects* depend on diversity

Many opportunities for future research:

- More/better *specific PDI's* to use in practice
- More/better *axioms*, more *characterisation* results
- Analytically *prove* facts about experimentally observed *trends*
- Use PDI's to structure *experimental work* in preference handling

V. Hashemi and U. Endriss. Measuring Diversity of Preferences in a Group. Proc. 21st European Conference on Artificial Intelligence (ECAI-2014).