Making Complex Beliefs Tractable

Barbara Dunin-Kęplicz

Institute of Informatics
University of Warsaw

Opole, June, 2014

Based on a joint work with A. Szałas and R. Verbrugge

The Structure of Talk

- Motivations.
- Epistemic profiles and belief structures.
- Formalization.
- Implementation tools.
- Conclusions.

The merits of classical group knowledge

- *Consensus* between group participants on several levels:
  - *General knowledge*, E-KNOW_G:
    "every agent in group G knows" – consensus;
    $\text{E-KNOW}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{KNOW}(i, \varphi)$
  - *Iterations of general knowledge*, E-KNOW^k_G:
    "every agent knows that (every agent knows that (...))";
    propagation of consensus plus introspection;
  - *Common knowledge*, C-KNOW_G, informally:
    infinitely iterated stack of general knowledge operators;
    $\text{C-KNOW}_G(\varphi) \leftrightarrow \text{E-KNOW}_G(\varphi \wedge \text{C-KNOW}_G(\varphi))$.
- Important: common knowledge permits to draw common consequences from commonly known premises.

But this comes at a price.

The merits of classical group knowledge

- *Consensus* between group participants on several levels:
  - *General knowledge*, $\text{E-K}\textsc{now}_G$:
    "every agent in group G knows" – consensus;
    $\text{E-K}\textsc{now}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{K}\textsc{now}(i, \varphi)$
  - *Iterations of general knowledge*, $\text{E-K}\textsc{now}^k_G$:
    "every agent knows that (every agent knows that (...))";
    propagation of consensus plus introspection;
  - *Common knowledge*, $\text{C-K}\textsc{now}_G$, informally:
    infinitely iterated stack of general knowledge operators;
    $\text{C-K}\textsc{now}_G(\varphi) \leftrightarrow \text{E-K}\textsc{now}_G(\varphi \wedge \text{C-K}\textsc{now}_G(\varphi))$.
- Important: common knowledge permits to draw common consequences from commonly known premises.

But this comes at a price.

The merits of classical group knowledge

- *Consensus* between group participants on several levels:
  - *General knowledge*, E-KNOW_G:
    "every agent in group G knows" – consensus;
    $\text{E-KNOW}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{KNOW}(i, \varphi)$
  - *Iterations of general knowledge*, E-KNOW_G^k:
    "every agent knows that (every agent knows that (...))";
    propagation of consensus plus introspection;
  - *Common knowledge*, C-KNOW_G, informally:
    infinitely iterated stack of general knowledge operators;
    $\text{C-KNOW}_G(\varphi) \leftrightarrow \text{E-KNOW}_G(\varphi \wedge \text{C-KNOW}_G(\varphi))$.
- Important: common knowledge permits to draw common consequences from commonly known premises.

But this comes at a price.

The merits of classical group knowledge

- *Consensus* between group participants on several levels:
  - *General knowledge*, E-KNOW_G:
    "every agent in group G knows" – consensus;
    $\text{E-KNOW}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{KNOW}(i, \varphi)$
  - *Iterations of general knowledge*, E-KNOW_G^k:
    "every agent knows that (every agent knows that (...))";
    propagation of consensus plus introspection;
  - *Common knowledge*, C-KNOW_G, informally:
    infinitely iterated stack of general knowledge operators;
    $\text{C-KNOW}_G(\varphi) \leftrightarrow \text{E-KNOW}_G(\varphi \wedge \text{C-KNOW}_G(\varphi))$.
- Important: common knowledge permits to draw common consequences from commonly known premises.

But this comes at a price.

Evolution of Group Knowledge

Traditional approaches	The new approach
"What every fool knows"	Synthetic information extracted from individuals or other groups
Holistic knowledge	Selected aspects solely
Consensus	Group members not forced to adopt group conclusions: only required to obey them during the group's lifetime
Logical Omniscience	Incomplete/inconsistent beliefs allowed: information may be grabbed from others or resolved non-monotonically
Homogeneity (typically)	Heterogeneity: resoning adjusted to application domain and individualized
Reasoning intractable	Tractability of reasoning

Distributed knowledge

"What a wise person would know": pulls together the individual knowledge and draws classical conclusions from the combined information.

Group knowledge

May go further: from the same knowledge, a variety of reasoning methods may lead to more far-reaching conclusions.

Important: this takes place at our approach.

Distributed knowledge

"What a wise person would know": pulls together the individual knowledge and draws classical conclusions from the combined information.

Group knowledge

May go further: from the same knowledge, a variety of reasoning methods may lead to more far-reaching conclusions.

Important: this takes place at our approach.

Some Issues

- *Reducing complexity* of communication and reasoning is essential (especially in time-critical situations).
- During cooperation consensus is superfluous.
- Information derives from different sources
- Information is *imprecise* due to:
  - limited accuracy of sensors and other devices,
  - restrictions on time and other resources,
  - unfortunate combinations of environmental conditions,
  - limited reliability/failure of physical devices, etc.

Inconsistencies

Appear on many different levels: individual, between agents, between agents and groups, between groups and groups.

Some Issues

- *Reducing complexity* of communication and reasoning is essential (especially in time-critical situations).
- During cooperation consensus is superfluous.
- Information derives from different sources
- Information is *imprecise* due to:
  - limited accuracy of sensors and other devices,
  - restrictions on time and other resources,
  - unfortunate combinations of environmental conditions,
  - limited reliability/failure of physical devices, etc.

Inconsistencies

Appear on many different levels: individual, between agents, between agents and groups, between groups and groups.

How to formally model such complicated situations?

- A shift in perspective (Dunin-Kęplicz and Szałas): from reasoning in multi-modal systems of high complexity to querying knowledge bases.
- Incosistencies as first-class citizens!
- Paraconsistent knowledge-bases.
- A variety of reasoning methods supporting the reduction of ignorance and inconsistencies.

Semantical Structures

Reflect agent's/group's belief acquisition and formation:

- an agent/group starts with *constituents*:
  sets of beliefs acquired by perception, expert supplied knowledge, communication, etc.
- constituents are transformed into *consequents* according to agent's/group's *individual epistemic profile*.
- epistemic profile characterizes agent's/group's reasoning capabilities and the manner of dealing with inconsistency or lacking information

Definition

Let $\mathbb{C} \stackrel{\text{def}}{=} \text{Fin}(\mathcal{G}(Const))$ be the set of all finite sets of ground literals over the set of constants *Const*. Then:

- by a *constituent* we understand any set $C \in \mathbb{C}$;
- by an *epistemic profile* we understand any function $\mathcal{E} : \text{Fin}(\mathbb{C}) \longrightarrow \mathbb{C}$;
- by a *belief structure over an epistemic profile* $\mathcal{E}$ we mean $\mathcal{B}^{\mathcal{E}} = \langle \mathcal{C}, F \rangle$, where: $\mathcal{C} \subseteq \mathbb{C}$ is a nonempty set of constituents; $F \stackrel{\text{def}}{=} \mathcal{E}(\mathcal{C})$ is the *consequent* of $\mathcal{B}^{\mathcal{E}}$.

Individual and Group Belief Structures [DKS]

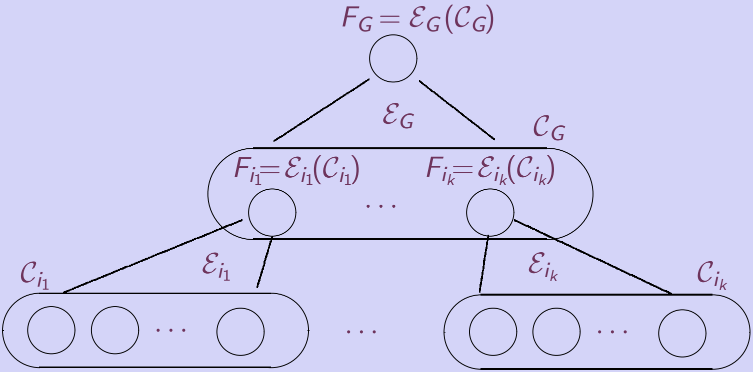

Uniformity

The same uniform approach applies to groups of groups of agents or to mixed groups of individuals and other complex topologies.

Individual and Group Belief Structures [DKS]

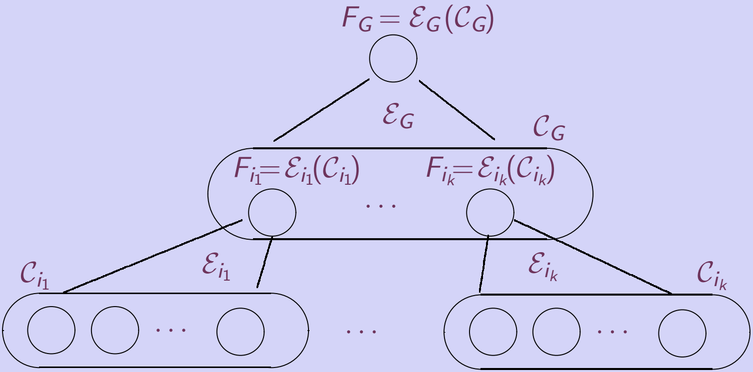

Uniformity

The same uniform approach applies to groups of groups of agents or to mixed groups of individuals and other complex topologies.

Syntax

- Classical first-order language over a given vocabulary without function symbols.
- $\mathrm{Bel}_i()$, $\mathrm{Bel}_G()$ standing for individual and group beliefs.

Semantics

Substantially differs from the classical one:

- truth values $\mathrm{T}, \mathrm{I}, \mathrm{U}, \mathrm{F}$ (true, inconsistent, unknown, false);
- is based on sets of ground literals not on relational structures.

Syntax

- Classical first-order language over a given vocabulary without function symbols.
- $\mathrm{Bel}_i()$, $\mathrm{Bel}_G()$ standing for individual and group beliefs.

Semantics

Substantially differs from the classical one:

- truth values $\mathrm{T}, \mathrm{I}, \mathrm{U}, \mathrm{F}$ (true, inconsistent, unknown, false);
- is based on sets of ground literals not on relational structures.

Definition

The *truth value* of a literal ℓ w.r.t. a set of ground literals L and valuation v, denoted by $\ell(L, v)$, is defined as follows:

$$\ell(L, v) \stackrel{\text{def}}{=} \begin{cases} \mathrm{T} & \text{if } \ell(v) \in L \text{ and } (\neg\ell(v)) \notin L; \\ \mathrm{I} & \text{if } \ell(v) \in L \text{ and } (\neg\ell(v)) \in L; \\ \mathrm{U} & \text{if } \ell(v) \notin L \text{ and } (\neg\ell(v)) \notin L; \\ \mathrm{F} & \text{if } \ell(v) \notin L \text{ and } (\neg\ell(v)) \in L. \end{cases}$$

Connectives and Quantifiers

- $\neg\mathrm{F} \stackrel{\text{def}}{=} \mathrm{T}; \neg\mathrm{U} \stackrel{\text{def}}{=} \mathrm{U}; \neg\mathrm{I} \stackrel{\text{def}}{=} \mathrm{I}; \neg\mathrm{T} \stackrel{\text{def}}{=} \mathrm{F};$
- *Truth ordering*: $\mathrm{F} < \mathrm{U} < \mathrm{I} < \mathrm{T};$
- disjunction (as well as $\exists$) is max and conjunction (as well as $\forall$) – min w.r.t. truth ordering.

Definition

The *truth value* of a literal ℓ w.r.t. a set of ground literals L and valuation v, denoted by $\ell(L, v)$, is defined as follows:

$$\ell(L,v) \stackrel{\text{def}}{=} \begin{cases} \mathrm{T} & \text{if } \ell(v) \in L \text{ and } (\neg\ell(v)) \notin L; \\ \mathrm{I} & \text{if } \ell(v) \in L \text{ and } (\neg\ell(v)) \in L; \\ \mathrm{U} & \text{if } \ell(v) \notin L \text{ and } (\neg\ell(v)) \notin L; \\ \mathrm{F} & \text{if } \ell(v) \notin L \text{ and } (\neg\ell(v)) \in L. \end{cases}$$

Connectives and Quantifiers

- $\neg\mathrm{F} \stackrel{\text{def}}{=} \mathrm{T}; \neg\mathrm{U} \stackrel{\text{def}}{=} \mathrm{U}; \neg\mathrm{I} \stackrel{\text{def}}{=} \mathrm{I}; \neg\mathrm{T} \stackrel{\text{def}}{=} \mathrm{F};$
- *Truth ordering*: $\mathrm{F} < \mathrm{U} < \mathrm{I} < \mathrm{T}$;
- disjunction (as well as $\exists$) is max and conjunction (as well as $\forall$) – min w.r.t. truth ordering.

Remark

Constituents represent the world as perceived by group members. If $\mathcal{B}^{\mathcal{E}} = \langle \mathcal{C}, F \rangle$ is a belief structure then $\mathrm{Bel}()$-free formulas are evaluated in the union of constituents $\bigcup_{C \in \mathcal{C}} C$.

Definition

- The *truth value of a* $\mathrm{Bel}()$*-free formula* α w.r.t. belief structure $\mathcal{B}^{\mathcal{E}} = \langle \mathcal{C}, F \rangle$ and valuation v, is defined by:

$$\alpha(\mathcal{B}^{\mathcal{E}}, v) \stackrel{\text{def}}{=} \alpha(\bigcup_{C \in \mathcal{C}} C, v).$$

Remark

Constituents represent the world as perceived by group members. If $\mathcal{B}^{\mathcal{E}} = \langle \mathcal{C}, F \rangle$ is a belief structure then $\mathrm{Bel}()$-free formulas are evaluated in the union of constituents $\bigcup_{C \in \mathcal{C}} C$.

Definition

- The *truth value of a* $\mathrm{Bel}()$*-free formula* α w.r.t. belief structure $\mathcal{B}^{\mathcal{E}} = \langle \mathcal{C}, F \rangle$ and valuation v, is defined by:

$$\alpha(\mathcal{B}^{\mathcal{E}}, v) \stackrel{\text{def}}{=} \alpha\Big(\bigcup_{C \in \mathcal{C}} C, v\Big).$$

Definition

- $\mathrm{Bel}(\mathrm{T}) \stackrel{\mathrm{def}}{=} \mathrm{T} \quad \mathrm{Bel}(\mathrm{I}) \stackrel{\mathrm{def}}{=} \mathrm{I} \quad \mathrm{Bel}(\mathrm{U}) \stackrel{\mathrm{def}}{=} \mathrm{U} \quad \mathrm{Bel}(\mathrm{F}) \stackrel{\mathrm{def}}{=} \mathrm{F}.$
- $\mathrm{Bel}(\alpha)(\mathcal{B}^{\mathcal{E}}, v) \stackrel{\mathrm{def}}{=} \mathrm{Bel}(\alpha(F, v))$ when α is $\mathrm{Bel}()$-free (note that $\alpha(F, v)$ is a truth value).
- When $\mathrm{Bel}()$ operators are nested in α then $\alpha(\mathcal{B}^{\mathcal{E}}, v)$ is evaluated starting from the innermost occurrence of $\mathrm{Bel}()$, which is then replaced by the obtained truth value, etc.

Epistemic profile

Defines a schema of reasoning and dealing with conflicting and missing information.

Completing imperfect knowledge of agents and groups

- non-monotonic reasoning;
- defeasible reasoning;
- heuristic methods;
- methods inspired by argumentation theory.

Derivatives

Used to define epistemic profiles in a better structured and more readable manner.

Epistemic profile

Defines a schema of reasoning and dealing with conflicting and missing information.

Completing imperfect knowledge of agents and groups

- non-monotonic reasoning;
- defeasible reasoning;
- heuristic methods;
- methods inspired by argumentation theory.

Derivatives

Used to define epistemic profiles in a better structured and more readable manner.

Building Epistemic Profiles

Epistemic profile

Defines a schema of reasoning and dealing with conflicting and missing information.

Completing imperfect knowledge of agents and groups

- non-monotonic reasoning;
- defeasible reasoning;
- heuristic methods;
- methods inspired by argumentation theory.

Derivatives

Used to define epistemic profiles in a better structured and more readable manner.

(Distributed) Derivatives

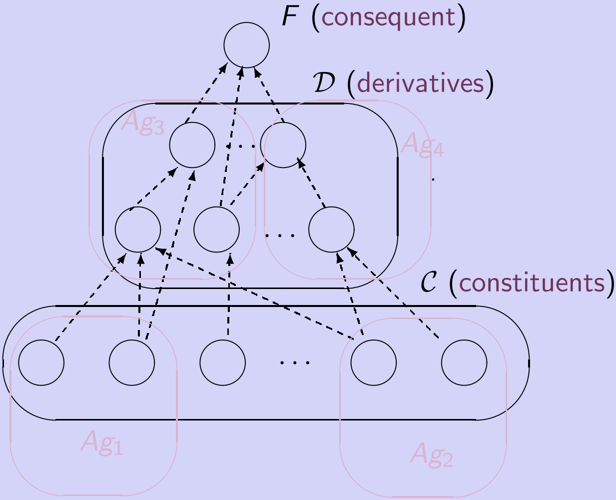

(Distributed) Derivatives

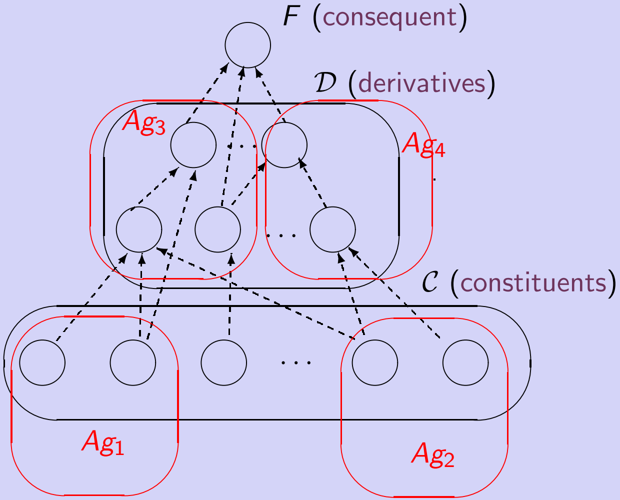

Group specific completing imperfect knowledge

Social procedures help to establish different types of group knowledge or belief:

- public announcements;
- different voting methods;
- methods involving power relations.

Applied methods need to be reflected in individual or group epistemic profile.

Group specific completing imperfect knowledge

Social procedures help to establish different types of group knowledge or belief:

- public announcements;
- different voting methods;
- methods involving power relations.

Applied methods need to be reflected in individual or group epistemic profile.

Conflicting Information [BDKSV]

Resolving Conflicts

- Conflicting information may be resolved on the individual or group level in a similar way.
- Disambiguation methods are highly context, application-dependent and individualized.

Some strategies as to timing

- "Killing inconsistency at the root": to solve them asap;
- "Living with inconsistency": postpone disambiguation to the last possible moment (or even forever);
- Solving inconsistency whenever relevant information appears.

Conflicting Information [BDKSV]

Resolving Conflicts

- Conflicting information may be resolved on the individual or group level in a similar way.
- Disambiguation methods are highly context, application-dependent and individualized.

Some strategies as to timing

- "Killing inconsistency at the root": to solve them asap;
- "Living with inconsistency": postpone disambiguation to the last possible moment (or even forever);
- Solving inconsistency whenever relevant information appears.

4QL [MS], see 4ql.org)

- Possibly many, perhaps distributed information sources.
- Four logical values (T, F, I, U).
- Unrestricted negation (in premisses and conclusions of rules).
- Introduces simple tools: rules, modules and external literals to formulate and enrich (lightweight versions of) (Local) Closed World Assumption, autoepistemic reasoning, default reasoning, defeasible reasoning, etc.
- PTime complexity of computing queries.
- Captures all tractable queries.

An experimental open source interpreter Inter4QL can be downloaded from 4ql.org.

Rules

The form:

Conclusion :– Premises.

- Conclusion is a literal.
- Premises are first-order formulas extended by allowing to refer to multiple information sources (modules).
- The meaning: premises imply conclusion.

Well-supported model (intuition)

Intuitively, a *well-supported model* is a model where each literal has value T or I iff this is forced by a finite derivation starting from facts.

Theorem

Any set of rules has a unique well-supported model which can be computed in deterministic polynomial time w.r.t. the size of the database.

Well-supported model (intuition)

Intuitively, a *well-supported model* is a model where each literal has value T or I iff this is forced by a finite derivation starting from facts.

Theorem

Any set of rules has a unique well-supported model which can be computed in deterministic polynomial time w.r.t. the size of the database.

4QL: Layered and Modular Structure

The architecture

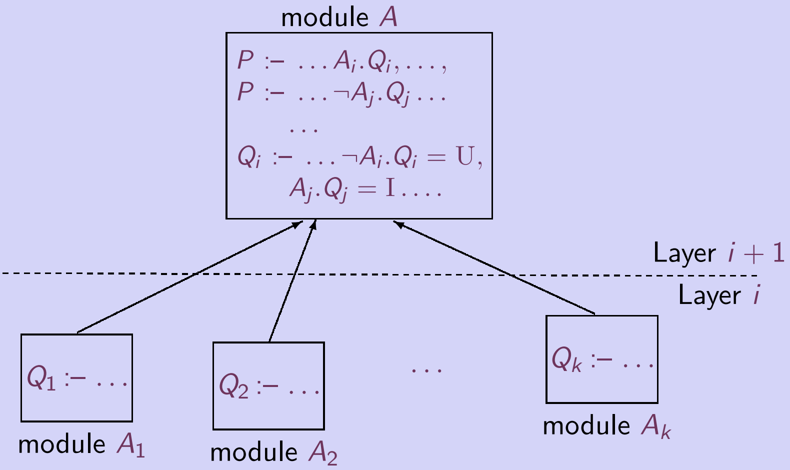

External literals: the tool for expressing nonmonotonic rules

- An *external literal* is of one of the forms:
  $$A.R,\ \neg A.R,\ A.R \ \mathrm{IN}\, T,\ \neg A.R \ \mathrm{IN}\, T,$$
  where:
  - A is a module and R is a relation in A;
  - $T \subseteq \{\mathrm{T}, \mathrm{F}, \mathrm{I}, \mathrm{U}\}$ (if $T = \emptyset$ then $\ell \ \mathrm{IN}\, T$ is F).
- An external literal $A \ldots$ may only appear in a module B, provided that module A is in a strictly lower layer than B.
- We write $\ell = v$ rather than $\ell \ \mathrm{IN}\{v\}$.

Defeasible Reasoning

Consider the following defeasible rules reflecting buyer's requirements as to apartments:

$$r1 : size(X, large) \Rightarrow acceptable(X)$$
$$r2 : \neg pets\_allowed(X) \Rightarrow \neg acceptable(X)$$

with priorities $r2 > r1$.

Example Continued: Formalization in 4QL

A module, say B, contains rules:

$$\begin{array}{lcl} acceptable(X) & :- & size(X, large). \\ \neg acceptable(X) & :- & \neg pets\_allowed(X). \end{array}$$

Rules in a module other than B resolve possible inconsistencies according to the required priority:

$$\begin{array}{lcl} acceptable(X) & :- & B.acceptable(X) = \mathrm{T}. \\ \neg acceptable(X) & :- & B.acceptable(X) = \mathrm{I}. \end{array}$$

Example Continued: Formalization in 4QL

A module, say B, contains rules:

$$\begin{array}{lcl} acceptable(X) & :- & size(X, large). \\ \neg acceptable(X) & :- & \neg pets\_allowed(X). \end{array}$$

Rules in a module other than B resolve possible inconsistencies according to the required priority:

$$\begin{array}{lcl} acceptable(X) & :- & B.acceptable(X) = \mathrm{T}. \\ \neg acceptable(X) & :- & B.acceptable(X) = \mathrm{I}. \end{array}$$

Example: Robot Rescue Scenario [DKSV]

The Scenario

- There is fire in certain regions, resulting in a high temperature in these regions and their neighborhoods.
- A surveillance team $team = \{r_1, \ldots, r_k\}$ $(k > 1)$ of robots is formed, whose group beliefs include the one that searching for victims is more important than preserving robots.
- An example of a group belief can be:

  – enter the affected region and search for victims unless it is certain that operating in the region is impossible (e.g., temperature $> 80^o C$). (1)

Example: Robot Rescue Scenario

To formalize beliefs we use the following relations, where R represents regions:

- $temp(R, T)$: temperature in R is T;
- $risk(R)$: situation in R is risky;
- $allowed(R)$: entering R is allowed (perhaps also in a risky situation);
- $search(R)$: search for victims in R.

Each agent (robot) is equipped with its individual knowledge base, so it has individual beliefs about these relations.

Example: Robot Rescue Scenario

Geographic information system (GIS)-based information about subregions and robots' locations is available via:

- $close(P, R)$: robot P is close to R;
- $subreg(S, R)$: S is a subregion of R.

Constituents of *team*

- Consequents of each robot $r_1, \ldots, r_k$ in *team*.
- The *gis* module.

Epistemic Profile of *team*

To define *team*'s epistemic profile, we use derivatives:

- *allClose*, containing the relation *risk*, calculated according to votes of all agents close to a given region;
- *safe*, containing the relation *allowed*, stating that searching a given region is allowed (no certainty of damaging robots there).

Example: Robot Rescue Scenario

Constituents of *team*

- Consequents of each robot $r_1, \ldots, r_k$ in *team*.
- The *gis* module.

Epistemic Profile of *team*

To define *team*'s epistemic profile, we use derivatives:

- *allClose*, containing the relation *risk*, calculated according to votes of all agents close to a given region;
- *safe*, containing the relation *allowed*, stating that searching a given region is allowed (no certainty of damaging robots there).

Selected Rules for Derivatives: module *allClose*

$$risk(R) :- \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) = \mathrm{T}\} > \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) \neq \mathrm{T}\}.$$

$$\neg risk(R) :- \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) = \mathrm{T}\} \leq \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) \neq \mathrm{T}\}.$$

Selected Rules for Derivatives: module *safe*

$$\neg allowed(R) :- \exists r \in team(\quad gis.close(r,R) = \mathrm{T} \wedge r.temp(R,T) = \mathrm{T} \wedge T > 80).$$

Selected Rules for Derivatives: module *allClose*

$$
\begin{aligned}
risk(R) \mathrel{:-} & \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) = \mathrm{T}\} > \\
& \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) \neq \mathrm{T}\}. \\
\neg risk(R) \mathrel{:-} & \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) = \mathrm{T}\} \leq \\
& \#\{r \in team \mid gis.close(r,R) = \mathrm{T} \wedge r.risk(R) \neq \mathrm{T}\}.
\end{aligned}
$$

Selected Rules for Derivatives: module *safe*

$$
\begin{aligned}
\neg allowed(R) \mathrel{:-} \exists r \in team(\quad & gis.close(r,R) = \mathrm{T} \wedge \\
& r.temp(R,T) = \mathrm{T} \wedge T > 80).
\end{aligned}
$$

Consequent of *team*

The *team*'s consequent can be defined by rules:

$$risk(R) \text{ :– } allClose.risk(R). \tag{2}$$

$$\neg risk(R) \text{ :– } allClose.(\neg risk(R)) \wedge safe.allowed(R) \neq \mathrm{F}. \tag{3}$$

$$search(R) \text{ :– } safe.allowed(R) \neq \mathrm{F}. \tag{4}$$

Of course, robots may have individual beliefs about *risk* and *search*(R) contradicting (2)–(4). These inconsistencies can be resolved by a rule concluding that a robot cannot search regions where it cannot operate without being damaged.

Examples of Policies to Resolve Inconsistencies [DKSV]

- If a group member evaluates $\mathit{search}(R)$ as T, then do search:

$$\mathit{search}(R) :\!- \exists r \in \mathit{team}[r.\mathit{search}(R) = \mathrm{T}]. \tag{5}$$

- Search if no group member claims that one should not:

$$\mathit{search}(R) :\!- \forall r \in \mathit{team}[r.\mathit{search}(R) \neq \mathrm{F}]. \tag{6}$$

- Search if at least one group member claims one should search and no group member claims that one should not:

$$\mathit{search}(R) :\!- \exists r \in \mathit{team}[r.\mathit{search}(R) = \mathrm{T}] \wedge (6). \tag{7}$$

Example Policies for Solving Inconsistencies

With authority or outside expert

A group belief identified with the leader's or an expert's belief:

$$risk(R) \coloneq expLead.risk(R) = \mathrm{T}.$$
$$\neg risk(R) \coloneq expLead.risk(R) \in \{\mathrm{U}, \mathrm{I}, \mathrm{F}\}.$$

- We differentiate agent's/group's characteristics via individual and group epistemic profiles.
- The framework:
  - suits real-world applications;
  - allows for natural handling of inconsistencies and gaps in beliefs by using paraconsistent and nonmonotonic reasoning;
  - ensures a uniform modeling of individual and group beliefs, where group is a generic concept consisting of individual agents, groups of agents, groups of groups of agents, etc.

Tractability!

4QL guarantees tractability and is sufficient to implement all epistemic profiles and belief structures constructible in deterministic polynomial time.

- We differentiate agent's/group's characteristics via individual and group epistemic profiles.
- The framework:
  - suits real-world applications;
  - allows for natural handling of inconsistencies and gaps in beliefs by using paraconsistent and nonmonotonic reasoning;
  - ensures a uniform modeling of individual and group beliefs, where group is a generic concept consisting of individual agents, groups of agents, groups of groups of agents, etc.

Tractability!

4QL guarantees tractability and is sufficient to implement all epistemic profiles and belief structures constructible in deterministic polynomial time.

Belief structures and epistemic profiles

- B. Dunin-Kęplicz, A. Szałas: *Epistemic Profiles and Belief Structures*. KES-AMSTA, LNCS 7327, 360-369, 2012.
- B. Dunin-Kęplicz, A. Szałas: *Taming Complex Beliefs*. LNCS Trans. on Comput. Collective Intelligence 11 (LNCS 8065), 2013.
- B. Dunin-Kęplicz, A. Szałas: *Paraconsistent Distributed Belief Fusion*. IDC, 6th Int. Symposium on Intelligent Distributed Computing, Studies in Computational Intelligence 446, 59-69, 2013.
- B. Dunin-Kęplicz, A. Strachocka, A. Szałas, R. Verbrugge: *Perceiving Speech Acts under Incomplete and Inconsistent Information*. KES AMSTA, Frontiers of Artificial Intelligence and Applications 252, 255-264, IOS Press, 2013.
- B. Dunin-Kęplicz, A. Strachocka, A. Szałas, R. Verbrugge: *A Paraconsistent Approach to Speech Acts*. ArgMAS'2012, 9th Int. Workshop on Argumentation in Multi-Agent Systems, 59-78, 2012.

4QL (see also `4ql.org`)

- J. Małuszyński, A. Szałas: *Partiality and Inconsistency in Agents' Belief Bases*. KES-AMSTA, Frontiers of Artificial Intelligence and Applications 252, 3-17, IOS Press, 2013.
- A. Szałas: *How an Agent Might Think*. Logic Journal of IGPL 21(3), 515-535, 2013.
- J. Małuszyński, A. Szałas: *Living with Inconsistency and Taming Nonmonotonicity*. Datalog Reloaded, LNCS 6702, 384-398, 2011.
- J. Małuszyński, A. Szałas: *Logical Foundations and Complexity of 4QL, a Query Language with Unrestricted Negation*. Journal of Applied Non-Classical Logics 21(2), 211-232, 2011.