Calibrating Expressiveness of Collective Notions

Barbara Dunin-Kęplicz

Warsaw University and Polish Academy of Sciences

Based on work with Rineke Verbrugge

TARK'2011

A modern perspective of collective activity

- autonomous, intelligent cooperative systems
- teamwork (or Cooperative Distributed Problem Solving) as a paradigmatic activity
- spectacular and complex patterns of interaction

The objective

- to isolate the essential aspects of collective behavior
- to (possibly separately) characterize them
- to construct expressive enough, still possibly minimal formal model of collective behavior

A compromise between abstract model and reality is to be reached.

A modern perspective of collective activity

- autonomous, intelligent cooperative systems
- teamwork (or Cooperative Distributed Problem Solving) as a paradigmatic activity
- spectacular and complex patterns of interaction

The objective

- to isolate the essential aspects of collective behavior
- to (possibly separately) characterize them
- to construct expressive enough, still possibly minimal formal model of collective behavior

A compromise between abstract model and reality is to be reached.

Circumstances of collective behavior vary significantly w.r.t.:

- topological structure of groups (societies)
- power relations
- communication medium

Collective aspects need to be studied in detail *each and every time* when tailoring a model for a specific application.

Goal:

to diversify the expressive power of modeled notions, including formal mechanisms to *calibrate* their expressiveness.

Tuning collective notions

Circumstances of collective behavior vary significantly w.r.t.:

- topological structure of groups (societies)
- power relations
- communication medium

Collective aspects need to be studied in detail *each and every time* when tailoring a model for a specific application.

Goal:

to diversify the expressive power of modeled notions, including formal mechanisms to *calibrate* their expressiveness.

Distributed AI perspective

- agents of many different sorts (e.g. software agents, robots, UAV's: unmanned aerial vehicles)
- working together, but also with humans
- in an unstable and unpredictable environment

BGI (or BDI: beliefs, desires, intentions) systems

Our focus on mental state of cooperating participants

- beliefs (informational aspect)
- goals
- intentions (motivational aspect)

Distributed AI perspective

- agents of many different sorts (e.g. software agents, robots, UAV's: unmanned aerial vehicles)
- working together, but also with humans
- in an unstable and unpredictable environment

BGI (or BDI: beliefs, desires, intentions) systems

Our focus on mental state of cooperating participants

- beliefs (informational aspect)
- goals
- intentions (motivational aspect)

[Fagin, Halpern, Moses, Vardi] [Meyer, Van der Hoek]

Individual belief

- $\mathrm{BEL}(i, \varphi)$: agent i believes φ

General belief

- Notation: G – a group
- $\text{E-BEL}_G(\varphi)$: each agent in group G believes φ

Common belief

- $\text{C-BEL}_G(\varphi)$: everyone in G believes φ, everyone in G believes that everyone in G believes φ, etc.

[Fagin, Halpern, Moses, Vardi] [Meyer, Van der Hoek]

Individual belief

- $\mathrm{BEL}(i, \varphi)$: agent i believes φ

General belief

- Notation: G – a group
- $\mathrm{E\text{-}BEL}_G(\varphi)$: each agent in group G believes φ

Common belief

- $\mathrm{C\text{-}BEL}_G(\varphi)$: everyone in G believes φ, everyone in G believes that everyone in G believes φ, etc.

[Fagin, Halpern, Moses, Vardi] [Meyer, Van der Hoek]

Individual belief

- $\mathrm{BEL}(i, \varphi)$: agent i believes φ

General belief

- Notation: G – a group
- $\text{E-BEL}_G(\varphi)$: each agent in group G believes φ

Common belief

- $\text{C-BEL}_G(\varphi)$: everyone in G believes φ, everyone in G believes that everyone in G believes φ, etc.

Modal operators for group beliefs

Axioms & rule

- $KD45_n^C$ is the modal system $KD45_n$ plus:
- $\text{E-BEL}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{BEL}(i, \varphi)$
- $\text{C-BEL}_G(\varphi) \rightarrow \text{E-BEL}_G(\varphi \wedge \text{C-BEL}_G(\varphi))$
- From $\varphi \rightarrow \text{E-BEL}_G(\psi \wedge \varphi)$ infer $\varphi \rightarrow \text{C-BEL}_G(\psi)$ (Induction Rule).

$\text{C-BEL}_G(\psi)$ is easy to understand, hard to a achieve.

Established by *communication + reasoning*

Modal operators for group beliefs

Axioms & rule

- $KD45_n^C$ is the modal system $KD45_n$ plus:
- $\text{E-BEL}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{BEL}(i, \varphi)$
- $\text{C-BEL}_G(\varphi) \rightarrow \text{E-BEL}_G(\varphi \wedge \text{C-BEL}_G(\varphi))$
- From $\varphi \rightarrow \text{E-BEL}_G(\psi \wedge \varphi)$ infer $\varphi \rightarrow \text{C-BEL}_G(\psi)$ (Induction Rule).

$\text{C-BEL}_G(\psi)$ is easy to understand, hard to a achieve.

Established by *communication + reasoning*

Different degrees of belief in group

[Parikh, Krasucki]

$\mathrm{BEL}(\psi)$, $\text{E-BEL}_G(\psi)$, $\text{E-BEL}_G^2(\psi)$, $\ldots$, $\text{C-BEL}_G(\psi)$
For teamwork, $\text{E-BEL}_G(\psi)$ is not sufficient: $\text{C-BEL}_G(\psi)$ needed.

Important to realize

- How wide is the spectrum of possibilities to express knowledge/beliefs of agents and teams?
- How frequently we use models of others in our everyday commonsense reasoning?
- How important in modern intelligent systems is to create and/or revise models of others' minds and reason about them?

Different degrees of belief in group

[Parikh, Krasucki]

$\mathrm{BEL}(\psi)$, $\text{E-BEL}_G(\psi)$, $\text{E-BEL}^2_G(\psi)$, $\ldots$, $\text{C-BEL}_G(\psi)$
For teamwork, $\text{E-BEL}_G(\psi)$ is not sufficient: $\text{C-BEL}_G(\psi)$ needed.

Important to realize

- How wide is the spectrum of possibilities to express knowledge/beliefs of agents and teams?
- How frequently we use models of others in our everyday commonsense reasoning?
- How important in modern intelligent systems is to create and/or revise models of others' minds and reason about them?

From our experience in modeling group behavior

- Crucial: to differentiate the scope and strength of group attitudes.
- The resulting characteristics may differ significantly, and even become logically incomparable.

Agents' *awareness* forms the main difference over various contexts of common activity

Awareness about the situation: the state of agent's:

- beliefs about itself
- beliefs about other agents
- beliefs about the environment

Various epistemic notions: from distributed beliefs to common knowledge are adequate to formalize agents' awareness

From our experience in modeling group behavior

- Crucial: to differentiate the scope and strength of group attitudes.
- The resulting characteristics may differ significantly, and even become logically incomparable.

Agents' *awareness* forms the main difference over various contexts of common activity

Awareness about the situation: the state of agent's:

- beliefs about itself
- beliefs about other agents
- beliefs about the environment

Various epistemic notions: from distributed beliefs to common knowledge are adequate to formalize agents' awareness

Tuning and awareness

From our experience in modeling group behavior

- Crucial: to differentiate the scope and strength of group attitudes.
- The resulting characteristics may differ significantly, and even become logically incomparable.

Agents' *awareness* forms the main difference over various contexts of common activity

Awareness about the situation: the state of agent's:

- beliefs about itself
- beliefs about other agents
- beliefs about the environment

Various epistemic notions: from distributed beliefs to common knowledge are adequate to formalize agents' awareness

The question

- *who* needs to know *what* in order to cooperate effectively?

Tuning awareness

- possibly minimal solution per context is searched (communication and reasoning necessary for higher levels of awareness are costly and complex)
- awareness of different aspects should be tuned separately

Outcome

A sort of logical *tuning mechanism*.

The question

- *who* needs to know *what* in order to cooperate effectively?

Tuning awareness

- possibly minimal solution per context is searched (communication and reasoning necessary for higher levels of awareness are costly and complex)
- awareness of different aspects should be tuned separately

Outcome

A sort of logical *tuning mechanism*.

The question

- *who* needs to know *what* in order to cooperate effectively?

Tuning awareness

- possibly minimal solution per context is searched (communication and reasoning necessary for higher levels of awareness are costly and complex)
- awareness of different aspects should be tuned separately

Outcome

A sort of logical *tuning mechanism*.

Group

A group is a system of agents that are somehow constrained in their mutual interactions. [Weiss]

Team

A team is a group in which the agents are restricted to having a common goal of some sort. [Weiss]

From group to team

Joint intention by a team does not consist merely of simultaneous and coordinated individual actions; to act together, a team must be aware of and care about the status of the group effort as a whole [Levesque et al.]

Group

A group is a system of agents that are somehow constrained in their mutual interactions. [Weiss]

Team

A team is a group in which the agents are restricted to having a common goal of some sort. [Weiss]

From group to team

Joint intention by a team does not consist merely of simultaneous and coordinated individual actions; to act together, a team must be aware of and care about the status of the group effort as a whole [Levesque et al.]

Group

A group is a system of agents that are somehow constrained in their mutual interactions. [Weiss]

Team

A team is a group in which the agents are restricted to having a common goal of some sort. [Weiss]

From group to team

Joint intention by a team does not consist merely of simultaneous and coordinated individual actions; to act together, a team must be aware of and care about the status of the group effort as a whole [Levesque et al.]

Future-directed intentions

- can **not** be reduced to desires (goals) and beliefs
- have to do with partial plans and enable intra- and interpersonal coordination

Intentions play an important role in practical reasoning

- drive means-end reasoning
- constrain future deliberation
- persist
- influence beliefs upon which future practical reasoning is based

Future-directed intentions

- can **not** be reduced to desires (goals) and beliefs
- have to do with partial plans and enable intra- and interpersonal coordination

Intentions play an important role in practical reasoning

- drive means-end reasoning
- constrain future deliberation
- persist
- influence beliefs upon which future practical reasoning is based

TeamLog

[Dunin-Kęplicz, Verbrugge: *Teamwork in multiagent systems. A formal approach*, Wiley 2010]:
a theory of motivational attitudes of cooperating agents.

- TEAMLOG is founded on individual and social attitudes
- TEAMLOG addresses a nontrivial problem of group attitudes: collective intention and collective commitment
- collective notions are tuned to circumstances and organizational structure of the team
- agents reason about mental attitudes of others
- in applications, assumptions regarding others are kept to a minimum (to avoid overthinking and assure flexibility)

TeamLog, a theory of teamwork

TeamLog

[Dunin-Kęplicz, Verbrugge: *Teamwork in multiagent systems. A formal approach*, Wiley 2010]:
a theory of motivational attitudes of cooperating agents.

- TeamLog is founded on individual and social attitudes
- TeamLog addresses a nontrivial problem of group attitudes: collective intention and collective commitment
- collective notions are tuned to circumstances and organizational structure of the team
- agents reason about mental attitudes of others
- in applications, assumptions regarding others are kept to a minimum (to avoid overthinking and assure flexibility)

Collective intentions in strictly cooperative groups

Postulates for a collective intention

Again, φ: the goal of the system.

- All members of the group *individually intend* φ: a general intention $\text{M-INT}_G(\varphi)$
- All members in the group *intend* other members intend φ, etc.: a mutual intention $\text{M-INT}_G(\varphi)$ (a motivational core of group intention expressing reciprocity).
- Group members are aware about this mutual intention.

By means of collective intention a loosely coupled group becomes a strictly cooperative team.

Collective intentions in strictly cooperative groups

Postulates for a collective intention

Again, φ: the goal of the system.

- All members of the group *individually intend* φ: a general intention $\text{M-INT}_G(\varphi)$
- All members in the group *intend* other members intend φ, etc.: a mutual intention $\text{M-INT}_G(\varphi)$ (a motivational core of group intention expressing reciprocity).
- Group members are aware about this mutual intention.

By means of collective intention a loosely coupled group becomes a strictly cooperative team.

Axioms for mutual and collective intention

- $\text{E-INT}_G(\varphi)$: "every agent in G intends φ".
- $KD45_n^{\text{M-INT}_G}$ is the modal system $KD45_n$ plus:
- $\text{E-INT}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{INT}(i, \varphi)$
- $\text{M-INT}_G(\varphi) \leftrightarrow \text{E-INT}_G(\varphi \wedge \text{M-INT}_G(\varphi))$
- From $\varphi \rightarrow \text{E-INT}_G(\psi \wedge \varphi)$ infer $\varphi \rightarrow \text{M-INT}_G(\psi)$ (Induction Rule).

Definition of collective intention

$$\text{C-INT}_G(\varphi) \leftrightarrow \text{M-INT}_G(\varphi) \wedge awareness_G(\text{M-INT}_G(\varphi))$$

Axioms for mutual and collective intention

- $\text{E-INT}_G(\varphi)$: "every agent in G intends φ".
- $KD45_n^{\text{M-INT}_G}$ is the modal system $KD45_n$ plus:
- $\text{E-INT}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{INT}(i, \varphi)$
- $\text{M-INT}_G(\varphi) \leftrightarrow \text{E-INT}_G(\varphi \wedge \text{M-INT}_G(\varphi))$
- From $\varphi \rightarrow \text{E-INT}_G(\psi \wedge \varphi)$ infer $\varphi \rightarrow \text{M-INT}_G(\psi)$ (Induction Rule).

Definition of collective intention

$$\text{C-INT}_G(\varphi) \leftrightarrow \text{M-INT}_G(\varphi) \wedge \mathit{awareness}_G(\text{M-INT}_G(\varphi))$$

Example of a collective intention

- Two violinists, a and b, have studied together and toyed with the idea of giving a concert together someday.
- Later this becomes more concrete: $\mathrm{INT}(a, \varphi)$ and $\mathrm{INT}(b, \varphi)$, where $\varphi =$ "a and b perform the solo parts of the Bach Double Concerto".
- After communicating with each other about this, they start practising together.
  A mutual intention $\mathrm{M\text{-}INT}_G(\varphi)$ is now in place for $G = \{a, b\}$, plus a collective belief about this, so $\mathrm{C\text{-}INT}_G(\varphi)$.
- An opportunity appears: Carnegie Hall plans a concert for Christmas Eve, including the Bach Double Concerto.

Example of a collective intention cntd.

- Now a, b refine their collective intention to $\text{C-INT}_G(\psi)$, where $\psi =$ "a and b perform the solo parts of the Bach Double Concerto at the Christmas Eve concert in Carnegie Hall".
- a, b are chosen to be the soloists, and both sign the appropriate contract.
  Because they do this together, they have common knowledge, not merely collective belief, of their mutual intention:
  $\text{M-INT}_G(\psi) \wedge \text{C-KNOW}_G(\text{M-INT}_G(\psi))$.
- Common knowledge can be justified if needed, and a commonly signed contract provides a perfect basis for this.

a, b have developed a very strong variant of collective intention

Example of a collective intention cntd.

- Now a, b refine their collective intention to $\text{C-INT}_G(\psi)$, where $\psi =$ "a and b perform the solo parts of the Bach Double Concerto at the Christmas Eve concert in Carnegie Hall".
- a, b are chosen to be the soloists, and both sign the appropriate contract.
  Because they do this together, they have common knowledge, not merely collective belief, of their mutual intention: $\text{M-INT}_G(\psi) \wedge \text{C-KNOW}_G(\text{M-INT}_G(\psi))$.
- Common knowledge can be justified if needed, and a commonly signed contract provides a perfect basis for this.

a, b have developed a very strong variant of collective intention

Collective commitment

- the *key* concept in TeamLog
- subject to calibration: various building blocks tuned separately

Subjects related to agents' autonomy

- collective responsibility
- collective decision making
- collective planing
- collective revision making
- hiding classified issues

Collective commitment

- the *key* concept in TEAMLOG
- subject to calibration: various building blocks tuned separately

Subjects related to agents' autonomy

- collective responsibility
- collective decision making
- collective planing
- collective revision making
- hiding classified issues

The distinction de dicto vs. de re

- $\bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{C-BEL}_G(\text{COMM}(i,j,\alpha))$ — detailed awareness
- $\text{C-BEL}_G(\bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i,j,\alpha))$ — global awareness

TeamLog's complexity

Complexity

- The key notions are highly complex infinite concepts: its satisfiability problem is ExpTime-complete.
- Domain-specific knowledge helps to tailor TeamLog to the circumstances, reducing the complexity by applying weaker forms of awareness.

Team structure

- *coordinator* – coordinates teamwork between subteams of G.
- Helicopter with a *pilot* – directly accountable to the coordinator, communicates as equal with the *UAVs*.
- Several subteams $G_1, \ldots G_k \subseteq G$ work in parallel. Each of these subteams G_i consists of:
  - UAV_i – responsible for assigned sectors
  - n_i identical robots $rob_{i_1}, \ldots, rob_{i_{n_i}}$ – responsible to their UAV_i.

Team hierarchy

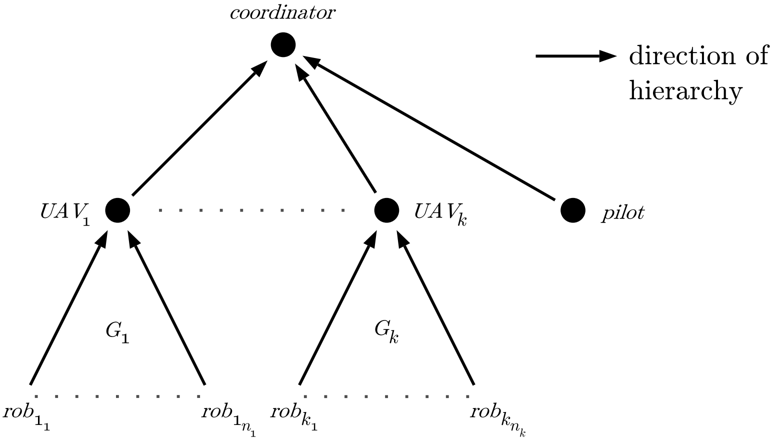

Adjusting collective intention to the case-study

Robots – two cases

1. Only individual actions are performed.
2. Limited form of cooperation: teams of two robots.

Robots – two cases for intentions

1. A general intention E-INT_G about the goals is enough.
2. E-INT^2_G is enough to form two-robot teams that are not competitive internally!

Adjusting collective intention to the case-study

Robots – two cases

1. Only individual actions are performed.
2. Limited form of cooperation: teams of two robots.

Robots – two cases for intentions

1. A general intention E-INT_G about the goals is enough.
2. E-INT^2_G is enough to form two-robot teams that are not competitive internally!

Robots: minimal levels of awareness and group intention

Robots – two cases for beliefs

1. General belief about every group intention $\text{E-BEL}_G(\text{E-INT}_G(\varphi))$.
2. E-BEL^2_G suffices to allow deliberation about other robots' intentions and beliefs, especially $\text{E-BEL}^2_G(\text{E-INT}^2_G(\varphi))$.

UAV: minimal levels of awareness and group intention

The *UAVs* – two cases for intentions

Within the team *UAV* must make sure that all agents are motivated to do their tasks.

1. $\mathrm{INT}(UAV, \mathrm{E\text{-}INT}_G(\varphi))$ is required w.r.t. the subteam group intention $\mathrm{E\text{-}INT}_G(\varphi)$,
2. $\mathrm{INT}(UAV, \mathrm{E\text{-}INT}^2_G(\varphi))$ is required w.r.t. the level of subteam group intention $\mathrm{E\text{-}INT}^2_G(\varphi)$.

The *UAVs* – beliefs

For *UAVs* to work with each other, they need at least $\mathrm{E\text{-}BEL}^2_G$ of other $UAVs'$ intentions.

UAV: minimal levels of awareness and group intention

The *UAVs* – two cases for intentions

Within the team *UAV* must make sure that all agents are motivated to do their tasks.

1. $\text{INT}(UAV, \text{E-INT}_G(\varphi))$ is required w.r.t. the subteam group intention $\text{E-INT}_G(\varphi)$,
2. $\text{INT}(UAV, \text{E-INT}^2_G(\varphi))$ is required w.r.t. the level of subteam group intention $\text{E-INT}^2_G(\varphi)$.

The *UAVs* – beliefs

For *UAVs* to work with each other, they need at least E-BEL^2_G of other *UAVs'* intentions.

Coordinator: minimal levels of awareness and group intention

The coordinator – intentions

Similarly, one level of intention more than the *UAVs* suffices to ensure the proper level of motivations: $\mathrm{INT}(\mathit{coordinator}, \mathrm{INT}^2_G(\varphi)))$.

The coordinator – beliefs

The coordinator sees the team as a collection of cooperating subteams: Therefore $\mathrm{BEL}(\mathit{coordinator}, \text{E-BEL}^2_G(\text{E-INT}^2_G(\varphi)))$ w.r.t. every group intention $\text{E-INT}^2_G(\varphi)$.

Substantial question

To what extend cognitive science can help in these issues?

Coordinator: minimal levels of awareness and group intention

The coordinator – intentions

Similarly, one level of intention more than the *UAVs* suffices to ensure the proper level of motivations: $\mathrm{INT}(\mathit{coordinator}, \mathrm{INT}^2_G(\varphi)))$.

The coordinator – beliefs

The coordinator sees the team as a collection of cooperating subteams: Therefore $\mathrm{BEL}(\mathit{coordinator}, \text{E-BEL}^2_G(\text{E-INT}^2_G(\varphi)))$ w.r.t. every group intention $\text{E-INT}^2_G(\varphi)$.

Substantial question

To what extend cognitive science can help in these issues?

Coordinator: minimal levels of awareness and group intention

The coordinator – intentions

Similarly, one level of intention more than the *UAVs* suffices to ensure the proper level of motivations: $\mathrm{INT}(\mathit{coordinator}, \mathrm{INT}^2_G(\varphi)))$.

The coordinator – beliefs

The coordinator sees the team as a collection of cooperating subteams: Therefore $\mathrm{BEL}(\mathit{coordinator}, \text{E-BEL}^2_G(\text{E-INT}^2_G(\varphi)))$ w.r.t. every group intention $\text{E-INT}^2_G(\varphi)$.

Substantial question

To what extend cognitive science can help in these issues?

Another possibility to calibrate collective concepts

Graded concepts

- In BGI graded concepts express e.g.
  - the strength of beliefs
  - the importance of goals
  - the degree of commitment to intentions.
- By reinterpreting beliefs, goals and intentions and then group attitudes, a graded version of TEAMLOG has been constructed
  [Dunin-Kęplicz, Nguyen and Szałas]

Lowering the complexity of reasoning

Reasoning via querying deductive/knowledge databases

The tradeoff:

- complexity of computing queries
- expressiveness of query language

A candidate rule query language: 4QL (http:\\4ql.org)

- 4QL [Małuszyński and Szałas] is a general purpose DATALOG$^{\neg\neg}$ rule language
- 4QL addresses lack of knowledge and inconsistencies
- 4QL has PTIME data complexity and captures PTIME

A shift from the multimodal BGI model to a 4QL-based BGI model

- by reinterpreting beliefs, goals and intentions
- by reinterpreting group attitudes

Lowering the complexity of reasoning

Reasoning via querying deductive/knowledge databases

The tradeoff:

- complexity of computing queries
- expressiveness of query language

A candidate rule query language: 4QL (http:\\4ql.org)

- 4QL [Małuszyński and Szałas] is a general purpose DATALOG$^{\neg\neg}$ rule language
- 4QL addresses lack of knowledge and inconsistencies
- 4QL has PTIME data complexity and captures PTIME

A shift from the multimodal BGI model to a 4QL-based BGI model

- by reinterpreting beliefs, goals and intentions
- by reinterpreting group attitudes

Lowering the complexity of reasoning

Reasoning via querying deductive/knowledge databases

The tradeoff:

- complexity of computing queries
- expressiveness of query language

A candidate rule query language: 4QL (http:\\4ql.org)

- 4QL [Małuszyński and Szałas] is a general purpose DATALOG$^{\neg\neg}$ rule language
- 4QL addresses lack of knowledge and inconsistencies
- 4QL has PTIME data complexity and captures PTIME

A shift from the multimodal BGI model to a 4QL-based BGI model

- by reinterpreting beliefs, goals and intentions
- by reinterpreting group attitudes