

Cooperative Human–Robot Planning with Team Reasoning

Raul Hakli^{1,2}

Accepted: 30 August 2016 / Published online: 17 October 2016
© Springer Science+Business Media Dordrecht 2016

Abstract The paper studies the connections between philosophical action theory and planning methods in artificial intelligence. It proposes a method of cooperative planning in which agents select actions using a combination of planning and team reasoning. Whereas several standard approaches to multi-agent planning start from individuals' plans and try to combine them into a group plan, the proposed method starts from constructing a group plan from which the individuals derive their sub-plans. The former method is labelled I-mode planning and the latter we-mode planning. The basic idea of the proposed we-mode planning method is that the agents conceive themselves as a group agent that selects the best plan specifying the actions of all individual agents who then carry out their parts in the group plan. We-mode planning is suited for human–robot cooperation in situations in which the participants have a shared goal and a shared plan evaluation function, and can observe each others' actions. The method is expected to lead in some cases to more human-like robot behaviour and more efficient execution of joint actions in human–robot teams.

Keywords Cooperative problem solving · Multi-agent planning · Human–robot interaction · Group agency · Team reasoning · We-mode

1 Introduction

There are several connections between theories of action in philosophy and planning methods in artificial intelligence (AI) and robotics. These fields of research have influenced each other, often by philosophers presenting general ideas and making conceptual distinctions, which are then employed by computer scientists and engineers in their detailed formalizations and implemented agent systems and architectures, which are again studied and commented on by philosophers. An important development in philosophical theories of action is marked by the transition from so called belief-desire theories (or BD-theories for short) to so called belief-desire-intention theories (or BDI-theories). These theories have inspired, especially through Michael Bratman's work ([13, 15]), researchers in artificial intelligence in which several BDI-logics and BDI-agent architectures have been developed, e.g. [20, 43, 52, 61, 75, 85].

However, some recent developments in philosophical theories of action have not yet been sufficiently utilised in the AI and robotics communities, and my aim here is to draw attention to some of them. In particular, I will look at attempts to conceptualise joint action in terms of *group agents*, which may have attitudes of their own. I will try to see what consequences such a conceptualisation might have for multi-agent planning and human–robot interaction (HRI). More specifically I will focus on Raimo Tuomela's distinction between acting in the I-mode and acting in the we-mode, and argue that the distinction can be applied to multi-agent planning. Moreover, I will argue that a specific reasoning method, called team reasoning, which is currently under intensive study in philosophy and economics (see, e.g. [3, 5, 6, 48, 58, 72, 81]), can be applied in the case of artificial agents like software agents and robots. Furthermore, it can be extended from selection of actions to selection of plans.

✉ Raul Hakli
raul.hakli@helsinki.fi

¹ Department of Culture and Society, Aarhus University,
Jens Chr. Skous Vej 7, 8000 Aarhus C, Denmark

² Present Address: Department of Political and Economic
Studies, University of Helsinki, Unioninkatu 40A,
P. O. Box 24, 00014 Helsinki, Finland

Team reasoning has been proposed as an alternative to game-theoretic methods of action selection, and it has been argued to solve certain multi-agent decision problems better than standard game theory [5,6,48,72]. Moreover, team reasoning can also explain certain types of observed human behaviour that competing theories have had difficulties with, and there is some empirical evidence that people might use team reasoning when they participate in shared activities [22]. If people indeed do so, social robots exhibiting similar reasoning would be likely to appear more predictable and human-like in their behaviour, thereby enhancing interaction, especially collaboration, between humans and robots. Conceptually, team reasoning is based on the idea of group agency. In contrast to game theory, in which agents are seen as maximising their expected utilities, in team reasoning the agents conceive themselves as parts of a group agent that maximises its expected utility. They identify the profile of actions or strategies that is best for the group and then perform their part actions in that profile.

The idea of group agency can be extended to multi-agent planning as follows. The agents conceive themselves as constituting a group agent, which has all the combined capacities of the group members and which constructs a plan to reach its goal. Each individual agent in a way simulates the reasoning of such a group agent by using a method similar to team reasoning: They compare alternative multi-agent plans in relation to the group's goals using a shared evaluation function, select the best one, and derive and execute their part in the selected plan.

In ideal circumstances, this procedure will result in an optimal plan for the group of agents. In realistic situations, it needs to be augmented with monitoring, communication, and helping behaviour that ensure coordination of actions, successful execution of the plan, and replanning in cases in which carrying out the original plan becomes impossible. The method is applicable in cooperative situations in which the agents have a shared goal and certain other conditions hold. In particular, it can be applied in HRI in cases in which humans and robots cooperatively act together. This paper presents the method in broad outline and further work is required to test its usefulness empirically. Some potential benefits can be expected. For instance, using a team reasoning approach is expected to reduce need of communication and effort to resolve potential conflicts as compared to approaches in which the agents only plan for themselves and then try to combine their plans. Also it is conjectured to lead to more predictable and more human-like robot behaviour.

The paper is organized as follows: I start, in Sect. 2, by providing the background to the research and discussing connections between philosophical action theory and planning in artificial intelligence and robotics. In Sect. 3, I will describe the idea of group agents and we-mode action that will be

used to develop an account of we-mode planning in Sect. 4. I will discuss related work in Sect. 5 and possible further extensions in Sect. 6. Section 7 will conclude.

2 Background

2.1 Planning in the Single-Agent Case

Within philosophical action theory, there is an established separation between two kinds of attitudes: attitudes that represent the current state of the world (beliefs) and attitudes that motivate the agent to change the state of the world (desires). In BD-theories decision-making was conceived as practical reasoning starting from premisses concerning beliefs and desires, and resulting in conclusions concerning intentions or actions. For instance, an agent whose strongest desire is to build a tower of bricks and who believes that she can build a tower of bricks only if she takes brick A and puts it on top of brick B, would be rationally required to reason from these premisses to the practical conclusion of forming an intention to take brick A and put it on top of brick B.

This division between two kinds of attitudes seemed essential for understanding rational agency. A rational agent would be one whose beliefs are coherent and who can reason how to satisfy her strongest desires based on her beliefs about the environment. The role of intentions in this theory was then to function as a link between attitudes and actions: Having formed an intention to take a particular action, an agent will start performing it (or at least try to perform it).

Several philosophers had studied the roles of beliefs, desires, and intentions in practical reasoning and noted the importance of intentions in understanding agency (see e.g. [4,28,87]). It had been recognised that the output of practical reasoning can be seen as an intention to act and sometimes intentions were required as premisses as well, but Michael Bratman's work [13] was especially influential in establishing the crucial roles of future-directed intentions in practical reasoning: In addition to *intentions-in-action* that initiate the next action to be taken, there are also *future-directed intentions* that constrain other future actions. If I intend to do something, then I am committed to doing it and this commitment constrains my future actions: I cannot select actions that would make it impossible to satisfy the intention. Therefore, future-directed intentions must be considered as premisses for practical reasoning because they constrain future choices and pose problems for further deliberation. Bratman established a close connection between intentions and plans: plans can be viewed as certain kinds of future-directed intentions. This conceptual move paved the way for the widespread acceptance of BDI-accounts in the AI community in which planning had already been an important area of research.

Note that decision theory and game theory, which are sometimes taken to serve as foundation for multi-agent systems (MAS) [67], are based on the belief-desire framework. According to decision theory the next action to be selected by an agent is the one that gives her the highest expected utility. Game theory extends this basic idea to the case of multiple agents. Both sometimes consider possible sequences of actions in the form of decision trees or extensive games, but they lack a notion of an intention-type commitment to a sequence of actions: At each choice point, the next action is to be selected based on expected utility calculations instead of prior commitments. Some theorists criticise decision theory for lacking a notion of commitment that is independent of utility values [66].

In contrast, some others think that if commitments are to play a role in agent's decision making they should be taken into account in the utilities [84, pp. 133–136]. However, it is not clear how that should be done in the standard approach based on maximization of expected utility (see, e.g., [60]). One could try modifying the utilities by increasing the payoff of action alternatives that lead to satisfaction of the intention or by penalising those alternatives that make it impossible to satisfy the intention. The problem is that this would violate the idea that intentions *constrain* decisions and not merely serve as considerations that are weighed together with ordinary desires [13, p. 24]. Bratman et al. [15] note that such constraints are useful for resource-bounded agents because they limit the amount of practical reasoning that the agent has to do.

Since games are defined in terms of players, choices, and payoffs, it seems that if intentions cannot be modelled by modifying payoffs, the only option left to incorporate intentions into games is to introduce constraints to the available choices: The consequence of an agent's adopting an intention could be modelled by removing from the agent's set of action alternatives those that are inconsistent with the adopted intention. However, this proposal goes against the idea that intentions are revocable: In some situations, it may be reasonable to reconsider adopted intentions and this may lead to revoking them. Currently, it is not clear how to incorporate intentions into game theory.

Bratman conceived of intentions as plans, calling his theory the planning theory of intention. The major roles of intentions are, according to Bratman [13, pp. 16–17], the following:

1. Intentions are *conduct-controlling* attitudes. This means that if a prior intention manages to survive until the time to take the action, then it will determine the course of action. Forming an intention involves a commitment to act. This is not the case for desires which are merely potential influencers of action.
2. Intentions have *stability* or *inertia*. This means that they resist reconsideration: In normal circumstances in which there is no special reason to reconsider, prior intentions persist and constrain the agent's future deliberation because the agent considers the course of action settled. (However, if need be, intentions can be revised or retracted.)
3. Intentions *pose problems* for further deliberation, and they lead to formation of new intentions. For instance, an intended end requires that the agent plans for how to achieve and adopts an intention concerning the means. Similarly, a general intention to do something requires specifying further details leading to the formation of more specific intentions.

Plans, according to Bratman [13, 29] are “intentions writ large”: They share the three properties of intentions and, additionally, have the following two characteristics:

4. Plans are typically *partial* meaning that they are not fully specified but allow for details to be filled in due course.
5. Plans typically have a *hierarchical structure* meaning that plans concerning ends embed plans concerning means, and more general plans embed more specific ones.

These features make sense for agents with bounded resources, because they enable us to coordinate our future actions both intra- and interpersonally in broad outlines without requiring all the details to be specified in advance, which could easily lead to waste of resources due to unanticipated changes in the environment. These ideas are clearly visible also in classical AI planning methods. In particular, in methods using partially ordered and hierarchical planning, steps in a plan may be refined later by replacing one high-level step by several more detailed action steps, and also the order in which they are to be performed can initially be left open and fixed later [63].

2.2 Planning in the Multi-Agent Case

Bratman has extended his theory of single-agent planning to the multi-agent case [12]. In the case of intentional group activities, Bratman talks about *shared agency* and *shared intentions*. What is important for Bratman is that shared intentions are patterns of individuals' beliefs and intentions concerning shared activities. They are thus built from ordinary beliefs and intentions attributed to individuals. Hence, they are not group intentions attributed to a group (as in Tuomela's [81] and in Gilbert's [39] theories), nor is there need to attribute specific *we-intentions* to the individuals (as in Tuomela's [81, 82] and Searle's [64, 65] theories). Bratman advocates what he calls the *continuity thesis*, according to which the move from the single-agent case to the multi-

agent case is conceptually conservative [12]: Understanding sociality and shared agency does not require radically new conceptual, metaphysical, or normative machinery beyond what is needed to account for individual planning agency. With respect to the intentional attitudes, the basic building blocks of individual beliefs, desires, and intentions, are thus sufficient to account for collective intentionality as well. Note that the strategy is clearly different from the attempt to understand collective agency by finding such analogies between individual agents and group agents that warrant the attribution of intentional attitudes to groups. This approach can be seen in Tuomela's we-mode theory [81], in Gilbert's plural subject theory [39], in List and Pettit's functionalism [55] and in Tollefsen's interpretationism [76].

Let us see how Bratman characterises shared intentions [12]. A central idea for him is that individuals can have intentions towards collective actions: I can *intend that we do* something together. For instance, I can intend that we dance tango and you can intend that we dance tango. Given that we have such intentions and certain other conditions hold under common knowledge, we have a shared intention to dance tango. And if our sub-intentions and actions are mutually responsive to these intentions, our shared intention leads to our dancing tango. This, according to Bratman, is a central case of sociality and acting together, and it is individualistic in the sense that it does not require more than attitudes attributed to individuals. However, there is a controversial bit, which comes from the idea of intending *that*. According to many philosophers, "intentions that" are derivative from "intentions to", and intentions *to* have a built in restriction that we can only intend our own doings (see, e.g., [7]). This condition is violated by the idea of intending that we do something because it is very difficult to understand what it means for me to intend to do something that involves your doings as well.

However, because plans are understood as intentions in Bratman's theory, we can describe the difference in planning terms, and this may help us understand what Bratman is after. In his theory, having a plan should not be understood merely as having a recipe for doing something but being committed to doing something, thus as having an intention [13, pp. 28–29]. According to Bratman, individuals can have plans that specify not only their own actions but also other agents' actions. When an agent considers alternative plans, these plans may be multi-agent plans in the sense that they specify what each agent is to do. It is not difficult to understand what it means for an agent to commit herself to such a plan. It entails performing the actions specified for her in the plan and possibly also keeping an eye on whether the other agents can handle their parts and maybe having a disposition to help them if needed. In terms of plans this appears unproblematic, and the idea has been endorsed also in AI approaches to multi-agent planning [41]. The problem only

appears when the situation is described in terms of intentions because according to a common understanding of intentions one can only intend one's own actions, and this is why "intentions that" have been controversial (for discussion, see, e.g., [14, 56, 83]).

Bratman is more concerned with defining the building blocks of shared intentions than presenting an algorithm for finding individual sub-plans and putting them together, but based on his discussions as well as on some of the AI literature on cooperative planning (e.g. [41]) we can try to outline how planning typically proceeds in a few main steps, assuming here that they have a shared goal:

1. Each agent finds the individual sub-plan that contributes to the shared goal and best serves their goals and intentions taking into account their previously adopted plans. The plan specifies the agent's own actions but it may involve constraints related to other agents' actions, for instance some steps may presuppose that some other agents perform complementary actions.
2. The agents negotiate and adjust their sub-plans in order to make them mesh (and to satisfy the shared goal), possibly via constructing a global multi-agent plan.
3. Each agent implements their sub-plans, and, when needed, adjusts to changes, monitors and helps others.

Such methods can be called *I-mode planning* methods: In them the individual plans (or intentions) are primary. Some methods do not even require a group-level plan but if there is one, it is produced by a method that takes the individual plans as input. *We-mode planning* would use a more top-down strategy in which the group plan (or group intention) is primary and the individual plans are produced by a method that takes the group plan as input. In the next chapter, we will motivate such an approach that is based on group agency.

3 Group Agents and We-Mode Action

Several philosophers have recently considered the possibility of group agents [55, 76, 81]. Arguments to the effect that groups satisfy conditions of agency have been presented, for example, by List and Pettit [55, p. 20]: According to them, an agent is anything that has representational states (e.g. beliefs or views), motivational states (e.g. desires or goals), and a mechanism that produces actions on the basis of these states (in essence, a method of forming intentions). Since groups, like committees, organizations, and nations, are often attributed such states and can on the basis of those states make decisions that lead their members to act, it seems to be possible to view them as agents.

In Tuomela's theory, on which I will focus here, groups act through individual group members trying to pursue the

group's goals. Group members can pursue goals in two ways: They can act in the I-mode or in the we-mode [80]. Roughly, the difference is as follows: When acting in the *I-mode*, agents select their actions on the basis of their attitudes (beliefs, desires and intentions) whereas in the *we-mode* they select their actions on the basis of what they consider as their group's attitudes: They conceive themselves as constituting a group agent that has attitudes of its own and can select between collective actions that specify each group member's actions. The individuals then do their parts as if they were the limbs of the larger agent.

This idea does not require that in addition to the individual agents there would exist another super-agent that would somehow control what the individuals do. Individuals are still autonomous, but they deliberate from the perspective of their group: They evaluate the options and deliberate, not from the I-perspective, asking "What should *I* do?", but from the we-perspective, asking "What should *we* do?". As Wooldridge and Jennings [86, p. 567] correctly note, because it is the individuals who ultimately have the ability to act, the relation between individual and collective intentional states must be made clear. However, that does not require that the relation is such that collective intentional states are defined in terms of individual intentional states which are taken to be primitive. Instead, collective intentional states can be understood as being collectively constructed by the group members who see themselves as constituting a group and who identify with it and therefore reason and act from the group's perspective. As an example, consider the following type of we-mode reasoning: I am a member of our group and we intend to build a tower together. We believe that in order to build a tower together we all have to place our bricks to the tower. Therefore, I intend to place my brick to the tower. The premisses of my practical reasoning refer to the group's irreducible attitudes but the conclusion is an individual intention. This kind of we-mode we-reasoning is what gives the connection between the group agent's attitudes and the individual intentions that are necessary for the group agent to function. The group agent is not an agent with a mind of its own, rather it is the result of the individuals' changing their point of view from first-person singular to first-person plural, from "I" to "we".

Cooperation to a shared goal is also possible in the I-mode in Tuomela's theory. If the attitudes of an individual are positively affected by the group, for instance, if they adopt the group's goal and pursue it in the way they would pursue a private goal, we can say that the agents act in *pro-group I-mode*. For example, I may reason from my intention to contribute to the shared goal to build a tower and my belief that the goal cannot be satisfied unless I place my brick to the tower, to forming an intention to place my brick to the tower. In a simple case like this there is no difference in the resulting action. However, as we will see later, in more complicated

cases, pro-group I-mode reasoning is more prone to coordination failures.

The characteristic features of we-mode in Tuomela's theory are *group reasons*, *collectivity condition*, and *collective commitment*. The first feature means that the group's attitudes, especially goals, function as reasons for action for the group members. The collectivity condition holds for the group's attitudes, for instance a goal, such that it is necessarily the case that if the goal is satisfied for one member, it is satisfied for all members. Another difference between the I-mode and the we-mode is in the commitments to the action: In the I-mode the agents are only committed to themselves, but in the we-mode there is also collective commitment, meaning that each group member is committed to others to participate in the joint activity. This commitment is typically a result of a collective decision to do something together. These are just rough characterisations of the main differences between I-mode and we-mode. For precise definitions, see Tuomela's writings (e.g. [80,81]).

Some of the differences can also be characterised in game-theoretical terms [48]: I-mode action can be seen as ordinary expected-utility maximization, and in pro-group I-mode the utility function derives from the group's utility function. (There are several proposals in the literature concerning the nature of group utility [24,38,71,73,80], see also literature on social choice theory [35].) We-mode action on the other hand can be seen as the group agent's utility-maximization, that is, as the group members doing their parts in the collective action that according to the group's expectations maximizes the group's utility. This can be illustrated by viewing an agent's decision-making as employing team reasoning [5,48,72]. Team reasoning presupposes that the agents have a shared goal that they are trying to satisfy, that is, that they can be seen as maximising the same utility function. Once the agents identify the action combination (specifying the actions of all of the agents) that maximizes expected group utility, team-reasoning agents will form intentions to perform their part actions of that action combination.

As an example, consider a "blocks world" scenario in which a human being and a robot are trying to build a tower from bricks and pyramids [19]. Both agents are assumed to have available the same actions and to be able to observe each other and the state of the world. The agents have a shared goal to build a tower in which the bricks are in the correct order and one of the pyramids is on the top. Suppose that both agents have reasoned that in order to satisfy the goal of building a tower, one of them will have to start building the tower by placing the first brick in the middle of the table while the other agent waits. Let us assume that both agents can reach all the blocks but are concerned with their own resource consumption and prefer to do as little as possible. This is not essential, it is only meant as an illustration of where the individual utility functions may come from.

Table 1 Matrix representing possible actions of the agents

	Take	Wait
Take	0, 0	1, 4
Wait	4, 1	2, 2

The numbers represent the agents' utilities, that is, how much each values the outcome resulting from the combination of actions taken by the agents. The first number in each outcome is the row player's (the human being) utility

We could have as well assumed that both are very eager and want to place as many bricks as possible, in which case their preferences, and hence utilities (which are numerical representations of preferences) would be different (but would lead to similar problems). Suppose, for simplicity's sake, that the relevant actions the agents consider are "take" and "wait". We can then represent the situation as the game-theoretic matrix in Table 1 where the human is the row player (player 1) and the robot is the column player (player 2).

Because both agents are concerned with their resource consumption, each prefers the action combination in which they themselves wait and the other one takes the brick. Both think that if their preferred action combination does not happen, it is better to wait a while in the hope that the other one will reconsider in the next turn: At least it will not lead to waste of effort, but of course there is no progress on the tower building either. The next best option is to be the one who takes the brick while the other one waits, because that will contribute to the goal. The worst outcome is the one in which both simultaneously try to take the same brick because that leads to waste of resources and may not advance tower building either.

The exact numbers in the matrix do not matter. The point of the exercise is to see that agents often have mismatches in their preferences, and this may lead to inefficiency in planning and execution. Suppose, for instance, that we use standard game-theoretic ideas to build artificial agents that reason about what to do in the above situation. Then we may implement best-reply reasoning in which each agent reasons roughly as follows: On the one hand, if the other agent takes a block, it will be better for me to wait. On the other hand, if the other agent waits, it will be better for me to wait as well. Therefore, because in both cases it is better for me to wait, I should wait (independently of what the other agent does). Because the situation is symmetric, the agents will end up in the only Nash equilibrium of the game in which both agents wait. This situation is not desirable because it does not advance the goal of building a tower.

One may think that the problem arises from there being something wrong in how the interaction situation has been modelled: Maybe if we change the agents' preferences a bit, they will be able to find better ways of coordinating their

actions. After all, we have assumed that the agents are cooperative and have a shared goal, so this should be reflected in their preferences. Indeed, we will suggest that in such a case they should be seen as maximizing the same utility function. There are two problems, however. (1) Arriving at such a group utility function is non-trivial. (2) Merely changing the utility functions is not sufficient to solve the problem. In this paper, the focus is on the second problem, which will be motivated below. As a solution, team reasoning is suggested. As to the first problem, we will not try to give a general solution to it but assume that it can be solved somehow. In simple cases in which the agents only have one shared goal, the group utility function may be given in the specification of the problem to be solved by the agents. There are also methods of constructing utility functions that can take several goals into account [46]. In some special cases in which individual utilities are available and satisfy certain conditions allowing interpersonal comparison, group utility may be calculated as a function of individual utilities, e.g., by summing them or using other methods of aggregation [35]. Sometimes it may be possible to model how agents' preferences in a group influence each other [71]. In more complicated cases, the agents may need to communicate and agree on a suitable group utility function [48]. Methods like negotiation, persuasion, or argumentation may sometimes be relevant in such a process.

Suppose now that the agents agree on a group utility function and assume further that they are willing to conceive the situation in its terms. The process of changing individual utilities to group utilities has been called a *preference transformation* [6]. In the above example, we may assume that the agents ultimately agree that it is best for the group that the agent closest to the brick (here assumed to be the human being) will take it because that results in minimizing cost for the whole group. The agents then adopt these group utilities, and we have the matrix in Table 2 in which the numbers for both players in each outcome are identical.

The resulting matrix contains two Nash equilibria: one in which the robot takes the brick and the human waits and another in which the human takes the brick and the robot waits. Moreover one of the equilibria (the latter) is strictly better than the other. The resulting game resembles games like Hi-Lo which are problematic for game theory. In such games there are multiple Nash equilibria, so agents are not

Table 2 The matrix with transformed utilities

	Take	Wait
Take	0, 0	4, 4
Wait	3, 3	2, 2

The numbers still represent the agents' utilities but they are now identical with the group's valuation of each outcome

guaranteed to converge to the optimal outcome even though such games are easy for human beings to optimally solve [5, 22, 48, 73]. The problem is that the type of best-reply reasoning that tries to model how agents might end up in Nash-equilibria leads to an impasse. Think of the two agents in the above situation trying to decide what to do by considering their best replies to the other agent's possible strategies in terms of the group's utilities. Each would reason as follows: If the other one takes the brick, it is better for the group that I wait. However, if the other one waits, better for the group that I take it. The same applies to both agents: Each agent's best reply depends on the other agent's action. Because each agent's best action depends on the other agent's action, this kind of reasoning is not sufficient to decide which agent should perform which action.

Agents acting in the pro-group I-mode thus face a coordination problem, but agents acting in the we-mode may invoke a *transformation of reasoning* from that of game-theoretic best-reply reasoning to team reasoning [6]. As explained above, in team reasoning the agents consider themselves as a group agent trying to maximize its expected utility. In this particular case they would reason as follows: We intend to do what is best for the group. It is best for the group if the human takes the brick and the robot waits. The robot then infers: Hence, I will wait. And the human infers: Hence, I will take the brick. Such team reasoning leads to the human taking the brick and the robot waiting. When followed by all group members this method of reasoning leads to them reaching optimal outcomes (in terms of the group utility) when such an outcome exists. If there are several optimal outcomes, the agents still need communication or other means to solve the coordination problem.

4 We-Mode Planning

4.1 The Basic Idea

The main idea of the paper is (1) to apply the above kind of team reasoning to multi-agent planning, thereby leading to an account of we-mode planning, and (2) use that method in human–robot interaction. The general idea of we-mode planning has been described in [47], and a closely related idea is presented in [16]. An earlier attempt to apply we-mode decision-making, basically team reasoning, to the selection single actions in multi-agent systems is given in [49]. Applying the idea of team reasoning to multi-agent planning would go as follows: The agents conceive themselves as constituting a group agent, find the group plan that is evaluated as the best one for the group, and then extract, or derive, their sub-plans from the group plan and execute them. In contrast to ordinary team reasoning, the agents are not selecting merely the best action. Rather the idea is that groups are planning agents as

well: If we accept that groups can be agents, Bratman's [12] argument for the need for future-directed intentions and planning seems to apply to groups too. It seems clear that they are resource-bounded planning agents similarly to the individuals who constitute them. Groups have an even stronger need for coordination of actions than individuals because they consist of several individuals who will have to be able to act together. Planning groups select the best plan and then individual group members refine and execute their parts in it.

Note the difference between ordinary decision theory and decision-making based on planning. In the former, the agent selects between single actions: The best action is the one that is expected to lead to the best outcome. In the latter, the agent selects between plans consisting of multiple actions. The best plan is the one that is expected to lead to the best outcome where the expected costs are taken into account. The action to be performed is the first action of the best plan found and it may lead to a worse outcome than the best action, but its rationality is taken to derive from the rationality of the best plan. There is a sense in which planning and team reasoning are closely related decision-making methods: They both involve selecting acts on the basis of evaluation of larger units than the acts themselves. Team reasoning selects acts based on the evaluation of the joint action of which the act is a part. Planning selects acts based on the evaluation of the long-term plan of which the act is a part. In this respect they both differ from standard decision and game theories which evaluate individual acts. Because of this they go well together and the possibility of planning group agents should be considered (for a more detailed argument and discussion, see [47]).

4.2 Elements of We-Mode Planning

In order to give an account of planning for group agents, we can apply the main elements of Bratman's planning account, but now there are two levels to consider: the level of the group agent (meaning the agents' joint deliberation together) and the level of individual group members. Some of the roles Bratman identifies function mainly at the level of the group.

At the level of the group agent, the reasoning-centred roles are important: Once the group makes a decision concerning a course of action, it is committed to it, and rationality demands it to exclude from consideration other plans that are inconsistent with it. This applies also to the individual agents: The group members are collectively committed to the collectively formed intention, and ought to abstain from making plans that are inconsistent with or make it difficult for them to contribute to the execution of the adopted plan. If circumstances change during the execution of the plan, adjustments must be made. Small changes can usually be made at the individual level, but in some cases a need to reconsider the adopted plan

may arise and, due to the collective commitment, such reconsideration must take place at the group level. For example, Pollock [60, pp. 201–203] notes that reconsideration may be rational if the agent discovers an alternative, possibly preferable, plan that has not previously been given consideration. In such cases reconsideration will take place when the situation allows it. If the discovery has been made by an individual group member but the discovered plan affects the parts of the others, she will have to inform them and the reconsideration will be performed only when the group has had a chance to evaluate the alternative plan. Similarly, other agents will have to be informed if one agent finds out that a goal has been reached or it has become impossible to reach (see [53] for discussion on such social conventions in multi-agent systems and [10,20,61] for discussion on various commitment strategies).

The case of the other reasoning-centred role is similar: Once the group makes a decision concerning an end, it has to consider means to the end. The adopted intention poses problems for further deliberation. To an extent this can be done at the level of the group: The preliminary steps and sub-tasks have to be identified, their dependencies and resource demands have to be analysed in order to come up with a partial order of tasks, and some kind of agreement on how they will be allocated to group members must be reached. Detailed planning can then be carried out at the level of individuals. Of course, the individuals need to take care that the mesh between the sub-plans is maintained in the process.

Here the evaluation of plans plays a crucial role. Assuming that the group has agreed on a method of evaluating plans, the individual group members need not negotiate about which plan to adopt unless there are ties between alternative plans. If there is just one group plan that is optimal and the group members are able to find it, they can apply team reasoning to identify it and to derive their own partial plans out of it. What we get is a we-mode theory of planning which is top-down in contrast to bottom-up theories: In the we-mode theory, the planning starts from a group intention which is first specified to the level of detail that provides the roles for the members who are then able to infer their sub-plans from the general group plan. In I-mode or bottom-up theories in contrast the agents start from their individual intentions and then try to combine their sub-plans in a meshing way.

The feasible plans are evaluated using a group utility function. Generally, utility functions measure the goodness of outcomes, that is, possible states of the world. The utility of plans derives from the utility of the expected outcomes that follow their execution, taking into account the cost and uncertainty involved in the execution [60, Ch. 5–6]. The factors that affect the evaluation of plans may vary. The satisfaction of goals is a major consideration. (See [46] for the relation between utilities and goals in planning.) Other typical things to consider are the number of steps in a plan or the time taken

to implement it, use of resources, individual abilities, quality of result, probability of success, possible effects on other goals, conformity to social norms, deliberation costs, and so on. The form of the utility function is highly dependent on the application. For instance, in cases of robotic assistants for elderly people, human comfort and “legibility” of robot behaviour (meaning that the behaviour is intuitively understood as part of joint action between human and robot) are the primary evaluation criteria instead of efficiency considerations [54].

In the case of resource-bounded agents like humans, robots and groups, the deliberation costs play a role. Typically in planning literature evaluation of alternative plans is not even considered because of the high computational complexity of finding just one solution to a planning problem even in the single-agent case [63, p. 372]. And, of course, the space of all possible plans is infinite, so some limits must be imposed to the exploration. For instance, the agents may search for a fixed time limit in a distributed fashion and then communicate the best plan found or set a new limit in case there is no solution. If there is no opportunity for communication during the search they may agree beforehand on the search method or the part of the search space to be explored in the hope that all the agents will be able to find the same solution. For instance, they may use iterative deepening on the number of actions, and once the first solution is found they can continue searching the remaining plans with the same number of actions and select the best solution of those that have been found.

Assume for the moment that there is only one plan with the highest utility value and it is fully specified, including the allocation of the sub-plans to the agents. Applying the idea of team reasoning [5,72], the agents can then derive their sub-plans from the group plan. In principle, there is no need to negotiate or bargain about their individual parts, nor is there need to worry about meshing or potential conflicts because the plan is already fully specified at the group level. The individuals can then simply start executing their parts.

Of course, this is a highly idealised situation. In reality: (i) the agents’ sub-plans may be only partially specified, (ii) only agent roles may be specified but not which agent takes which role, (iii) there may be contingencies requiring re-planning during the execution, (iv) there may be ties and no unique best plan, and (v) the agents may employ different search algorithms and fail to converge to the unique best plan. Resolving such situations generally requires communication.

Case i is often a desideratum because it is more efficient to coordinate actions on a higher level and leave the refinements to the individuals. Also it allows for a higher level of privacy because the agents need not reveal detailed information about their procedures of implementation. However, it may happen

that an agent cannot find a refined plan that is independent of other agents and then communication is needed (e.g. asking for help or suggesting coordinated action). Case ii requires a method for assigning tasks to agents which may have different capacities and resources. Because there may be several ways to allocate the tasks to the agents, the team reasoning idea requires that the allocation is done in a way that is best from the group's perspective. As to iii, dealing with contingencies requires monitoring of others' action, communication and re-planning, which should lead to helping behaviour in a way that is optimal to the group.

If there are ties (case iv), they need to be broken, and direct communication is often the best way to do it: Each agent may either suggest one plan or ask the other agents which plan to adopt. If there is disagreement, negotiation and persuasion may take place. Sometimes communication can be in terms of non-linguistic actions that are perceived by others: If plans can be distinguished by the first action, one of the agents may simply communicate the selection by initiating the first action in the plan. The others will then realise that the other plans are no longer compatible with the actions taken so far and can discard them. If the plans cannot be distinguished by the first action, the selection may be postponed to a later moment in which either the draw may have been resolved or one of the agents may then take initiative and communicate the selection of the plan. Similar communication is also needed in case the agents fail to converge to the same plan (case v).

In summary, we-mode planning proceeds typically as follows:

1. The group of agents finds the best plan for the group to satisfy its goal (taking into account previously adopted plans).
2. The agents derive their sub-plans from the group plan.
3. Each agent implements their plan, monitors others, adjusts to changes, and helps if needed.

4.3 Human–Robot Teams

The case of human–robot teams differs from multi-agent planning with artificial agents in that the human behaviour cannot be programmed in advance. That creates challenges for both robots and humans who have to adapt to each other's behaviours. There are two basic ways to apply the we-mode planning method in the case of human–robot teams: One way is that the agents form real groups in which they make decisions concerning principles of plan selection together. This will require ability to communicate either in restricted natural language or via a suitable communication interface, and specific procedures for decision-making. The other way is that the robot (or robots) simulate a group situation and try to predict which plan would be selected. Here the robot may use a pre-programmed plan evaluation function and observe

whether the human's (or humans') actions are consistent with the plan ranked highest according to the function. Whenever a human's observed behaviour deviates from the predicted plan, the robot should find a new plan consistent with the actions taken and adjust the plan evaluation function so as to lead to better predictions in the future. With respect to the human side, the suggestion is not to educate people to behave in the way specified by the theory in order to be able to interact with the robot. Rather the assumption is that the we-mode theory and team reasoning tell something about how people actually make decisions in social situations, and therefore robots simulating similar decision-making mechanisms will be acting in ways that people find familiar and predictable.

For instance, in the tower building case, there may be a unique optimal plan that minimizes time by favouring taking closest objects and employing interleaving of actions. This would be a natural way for humans to go about building such a tower. Given the shared evaluation function the robot and the human can figure out the best plan independently, derive their parts, refine them, and start execution. In such a case there is no need for negotiation as long as the refinements do not introduce conflicts (like potential collisions when moving their arms in the shared space).

In case there are more than one optimal plan, one can be selected by communication or by initiating action: In the tower building case, for instance, if both are ready to take a pyramid and finish the tower, either agent may just start doing it enabling the other to recognize the intention, drop the inconsistent alternative and continue.

As another example scenario, consider collaborative human–robot exploration for search and rescue operations (see, e.g. [40]). There the objective for the team is to explore an area and, for instance, locate a missing person. The team members should choose their paths in a way that maximises the area to be thoroughly explored. For example, they should choose different routes around an obstacle. In order to do that, the area can be divided into topological classes and the robots can wait for the humans to have made their first moves and then select their trajectories in complementary classes [40]. This is consistent with the idea of selecting the best group plan, but it is not necessary to wait for the humans; instead robots may use initiation of action to signal that they have identified an optimal plan. They may then take initiative and select their homotopy classes as specified in the best group plan and start moving accordingly. The assumption behind this is that the human team members will be able to recognise the robots' intentions and choose their trajectories in complementary classes.

Monitoring of the other's actions is usually needed to ensure that the chosen plan is properly implemented. Conflicts have to be avoided and, in case one agent deviates from the best plan, either communication or replanning must take place and a new best plan consistent with the actions taken so

far will be adopted. Applying team reasoning in replanning can automatically initiate helping behaviour. For instance, a robot may realise that the human has problems in doing its part in placing an object to a tower because the action takes longer than estimated. This may trigger replanning and lead to finding out that an alternative plan in which the robot supports the tower while the human places the object is better than waiting as in the original plan. In the exploration case, for example, if the human has not appeared from behind an obstacle as expected, the robot can reason that not everything has gone according to the plan. The robot can then change its plan in order to find out what has happened and possibly offer help. Of course, in the case of multiple robots, they should team reason to find out which one of them is in the best position to do that in order to minimise the costs for the group.

4.4 Conditions of We-Mode Planning

There are several conditions that are assumed to hold in we-mode planning. The scenario is assumed to be cooperative: The agents must have shared goals instead of private ones. In AI terms, it means that we are focussing on cooperative problem solving (CPS) rather than proper multi-agent systems in which the agents have their private goals or utility functions. Of course, the existence of private goals or utility functions does not make we-mode planning impossible, but the agents need to be able to put them aside and cooperate for the shared goal. One potential problem for the we-mode theory is that it may be difficult for human beings to work for shared goals without any concern for private goals. In the case of robots, however, the situation may be easier because we can design and implement whichever kinds of robots we want, for example, robots that act in the we-mode.

When viewed in terms of intentional states presupposed by the account, we-mode planning does not require “intentions that” that Bratman’s theory employs. Instead it relies on the notions of group intention and we-intention [80,81]. These notions are controversial as well but, again, they are quite easy to understand in terms of plans: Group intention means a group agent’s plan that specifies what each group member will do, and a we-intention means an individual’s sub-plan that specifies what that agent will do as her part in the group plan. If we-mode planning is implemented in robots with an ability to reason about mental states (see, e.g., [19]), they will need to represent not only their own intentions and the other agents’ intentions, but also intentions in the first person plural, that is, “our intentions”.

If the aim is to use we-mode planning in a way that enables the agents to make decisions independently then it is necessary to assume that there exists a shared plan evaluation function that is known by all the individual agents together with a shared (and restricted) space of possible plans. There

are many possible ways to arrive at such a function. For instance, the agents may agree on a function in a preliminary phase by using argumentation or voting or some other group decision-making methods (e.g. those discussed in [42]), or it may be given to them from an external source.

We-mode planning can also be used without a pre-existing evaluation function, but then the agents will need to communicate and agree on the plan (and the refinements later on) to be adopted. Again, there are many ways to do that. One way is to employ something like Tuomela’s bulletin board method [80, pp. 85–92] in which the agents propose plans and others accept or reject these proposals. This is still different from I-mode planning: The agents suggest candidate plans to be considered by the group instead of offering plans to which they are already committed to.

In a bit more detail, we-mode planning can be presented schematically as follows:

1. Agreement on the plan evaluation criteria, or plan evaluation function, and the constraints on acceptable plans (this happens at the group level and it may also be left implicit and combined with step 2).
2. Finding the best plan and committing to it (this may happen jointly or sometimes, given that step 1 has been carried out in sufficient detail, also individually).
 - 2.1 Finding a partial action plan (consistent with previous plans).
 - 2.2 Refining the plan (applying team reasoning when applicable).
 - 2.3 Scheduling and allocating actions to agents.
3. Derivation of individual sub-plans (either at the level of group or individuals).
4. Refinement of individual sub-plans (remembering that mesh must be maintained).
5. Execution of individual sub-plans, monitoring, adjusting, helping others, replanning when needed.

4.5 Expected Benefits

There are some anticipated benefits of we-mode planning: In principle, via the use of team reasoning, we-mode planning guarantees optimality of plans because the best plan is always selected. In practice, however, time constraints and possible changes during execution complicate matters. We-mode planning is expected to lead to reduced need for communication, bargaining and argumentation, because it relies on shared goals and because team reasoning resolves coordination problems in cases in which there is only one optimal equilibrium. The assumption of shared evaluation function also guides coordination in case of emerging situations that require new decisions to be made.

One important benefit of we-mode planning is that implementing team reasoning in robots may lead to more human-like behaviour and therefore make it easier for humans to interact with robots. There are some social psychological studies that provide some evidence that human beings in fact sometimes use team reasoning in decision-making situations involving multiple agents [22, 23]. If team reasoning is in fact used by humans, then arguably robots that use similar decision-making procedures in situations of human–robot interaction have an improved ability to coordinate their actions with human beings leading to more human-like behaviour that is predictable and appears familiar to human partners. In many simple everyday situations, it is obvious to people what each team member should do, and team reasoning is an attempt to operationalise such obviousness.

It has been noted that teams composed of human beings are able to cooperate and perform complicated joint actions in pursuit of shared goals with relative ease as compared to human–robot teams. One line of research suggests that robots should be better equipped with so called mind-reading capacities that enable them to recognise human intentions, beliefs, and other attitudes, and thereby adjust their actions in response to the anticipated actions of humans. However, such approaches, which build on an idea that human beings have a “theory of mind” that allows them to reason about each others’ mental states and eventually activities, have been criticised as well. Arguments against theories employing “folk-psychological” concepts like beliefs, desires and intentions have been presented [17, 62], and one alternative way of conceptualising interaction situations claims that the contextual features of the social situation play a much larger role than attribution of mental states: We predict other peoples’ actions based on their social roles: In a supermarket I do not need to construct complicated practical syllogisms to be able to predict that some people will collect stuff and then go to the cashier and that one person will collect their money, because based on my experiences I already know that some people are customers and some people are working at the supermarket and they will do what people in these roles as customers and employees generally do in supermarkets.

The question whether or to what extent people rely on reasoning about mental states or social roles is empirical, and this paper does not take a stand on that issue. However, the proposal given in this paper can be seen as making relatively small demands in terms of mind-reading capacities (even though it was presented in terms of relatively complicated collectively intentional concepts like group intentions and we-intentions). This is because the actual reasoning that it postulates, team reasoning, does not require people to make complicated predictions about other peoples’ doings, as argued, e.g., by Elisabeth Pacherie [58, 59]. The selection of actions can often be done without any mind-reading, on the basis of the social situation, in particular, the features of

the shared goal and the various means of attaining that goal. It suffices to understand that we are doing something and then to figure out what that requires of me. I need not reason about all the possible action alternatives of other agents; I can just presuppose that they will do their parts in the plan that best serves our shared goal. In the supermarket case, the shared goal is that I manage to get the things I need and the seller gets the adequate payment and all this happens in the most convenient way for all of us. Of course, I may have to change my actions in case other people turn out not to be doing theirs but such situations will typically arise when I observe the behaviour of others, not as a result of reasoning about their mental states.

The method may also give efficiency gains in complicated human–robot situations in which it is not obvious which multi-agent action combination leads to best results. This is because the proposed method not only allows the robot to select appropriate actions in response to human beings’ anticipated actions which is the common intention-recognition approach often adopted in HRI research. The we-mode method also allows the robot to take initiative towards the shared goal. This may improve the efficiency of human–robot joint action because robots may be faster in calculating the optimal multi-agent action whereas the human beings may be more skilled in recognizing the intentions of the robot: Instead of the robots trying to figure out what the humans are up to, in some cases it might be better to reverse their roles and let the robots take the lead and let the humans follow.

5 Related Work

The distinction between I-mode and we-mode planning is related to but distinct from the more familiar distinction between distributed and centralised methods of multi-agent planning (see, e.g. [53]). In I-mode planning the individual plans are constructed by the individual agents, but combining and revising them can be done either in a distributed or centralised fashion. In we-mode planning the group plans can, but need not, be constructed by a central controller. They can be constructed in a distributed fashion, for instance, by agents’ proposing partial solutions to a global planning problem without necessarily committing to these partial solutions as their plans (in a related fashion as in blackboard systems [25]), or they may search for complete solutions in parallel and then find out which one is the best. In cases in which the group cannot plan together it is possible that each agent plans in isolation but for the whole group. Assuming that they employ the same planning method (same search algorithm in the same restricted search space using the same plan evaluation function), they may expect to converge to the same plan and start executing their subplans in it. In any case,

in we-mode planning the individuals' plans concerning the satisfaction of the shared goals are always derived from the group plan.

The distinction is more akin to the division of multi-agent planning methods into ones in which coordination is prior to local planning and ones in which local planning is prior to coordination [32]: Representatives of the former are methods based on social laws and conventions [53, 68], methods based on organizational structuring [29, 69], and methods that use predefined protocols, such as the contract net [70]. These methods are similar to the we-mode method in that they have been inspired by the study of coordination mechanisms, like norms and practices, that are at work in actual human societies [79], and they are compatible with, and sometimes even required by, the we-mode theory. However, as in actual human societies, social structures like laws, norms, and conventions, guide, regulate, and enable various human actions, but they still leave open the possibility that people, when they act in concordance with these social structures, may think and act in the I-mode or in the we-mode. For instance, an individual may follow a social norm for prudential and self-interested reasons, for instance to avoid a sanction, not because she takes it to be best for her group. Hence, we can say that we-mode planning entails priority of coordination to local planning but not vice versa.

Many of the proposed approaches to multi-agent planning within distributed artificial intelligence (DAI) and multi-agent systems fall in the category of I-mode planning in this classification because they employ a bottom-up strategy of starting from the plans of individual agents and then use some sort of plan merging method to coordinate them into a global plan [34] or into a set of coordination plans that functions as if there were a global plan [2]. Multi-agent plan coordination is a closely related I-mode approach: A multi-agent plan coordination problem, as defined by Cox and Durfee [26], is a problem of transforming a set of individual agent plans into a consistent multi-agent plan by identifying and resolving interactions between the plans. They present a search algorithm that searches through the space of possible resolutions of inconsistencies to produce an optimal solution that minimises the total number of action steps in the multi-agent plan. Similarly, [77] defines a plan coordination method that takes the agents' separate plans and revises them in a way that improves either the agents' joint profit or their individual profits. Then these revised plans are used to construct a new distributed plan that replaces the agents' old plans.

Also partial global planning (PGP) methods [32, 33] start from the agents' individual plans: The agents exchange information about their local plans on the level of abstraction required for coordination. They then combine the information received from others' plans with their own to obtain a representation of all agents' activities in the form of partial global plans.

Nissim et al. [57] present a method that combines local planning of actions with constraint satisfaction where the constraints specify consistency requirements between the agents' public actions. The agents first find relaxed local plans (with a limit to the number of public actions) that satisfy their sub-goals but ignore dependencies to other agents' plans (in particular, preconditions). The agents then use these plans to define the constraints (preconditions and temporal dependencies), and a distributed constraint satisfaction algorithm is used to find a consistent assignment for the variables, employing backtracking and iterative deepening on the allowed number of public actions in the local plans.

Dimopoulos et al. [30] present μ -SATPLAN, which extends the classical propositional satisfiability planner SATPLAN. The agents aim at finding a joint plan by first building a plan that satisfies their individual goals (which are taken to be necessary for satisfying the global goal) and then sending that to other agents, who will then build plans that satisfy their goals and are consistent with the plan of the other agent. Once a solution has been found, the agents still try to improve it by taking turns in suggesting the initial plan and the other agents filling in their parts.

In the FMAP method [78], the agents jointly construct a partially-ordered multi-agent plan by exploring a multi-agent search tree in which the partial plans of the nodes are iteratively built by agents proposing refinements over unexplored nodes. This cooperative method of plan construction seems closely related to the we-mode approach even though it is not conceptualised in terms of group agents. Also it does not employ team reasoning. Instead, the agents use a heuristic evaluation function to select which of their alternative refinements to propose. This may effect the quality of the resulting plan since a refinement is evaluated only by the proposing agent who may not have a complete view of the plan. For reasons of privacy this may be necessary, and the situation may be similar also in we-mode planning when the agents' sub-plans are refined locally.

Other approaches closer to the spirit of we-mode planning are those in which the agents jointly construct and commit themselves to a multi-agent plan and then go about implementing their parts in it. A well-known example is the SharedPlans approach developed by Barbara Grosz and her collaborators [41–44]. Also methods employing the idea of joint intentions are similar in this respect [20, 21, 75, 86]. However, these approaches do not conceptualise planning in terms of a group agent which has beliefs, goals, and intentions of its own. In the SharedPlans, the group's decision-making process starts from the individuals' intentions concerning shared activities (Bratmanian "intentions that"), and the individuals are committed to reaching an agreement about a shared plan and updating their individual plans accordingly [42]. Also many teamwork accounts based on joint intentions ultimately define joint intentions in terms of individuals' atti-

tudes [21, 86]. As a result these approaches do not apply team reasoning to the selection of actions or plans as is suggested in this paper.

Boella et al. [9] consider cooperation to the group's utility extending the decision-theoretic planning methods of [46] with shared plans and a group utility function. The agents intend to do their parts in the shared plan and consider the shared group utility in selecting their actions and may also engage in helping behaviour. All this is similar to we-mode planning, but their idea is less idealised, because the agents may have mixed motives: They are not considering only the group utility but rather a combination of group utility and private utilities. The agents can give up the cooperation in cases it becomes too costly in terms of personal goal satisfaction. This also entails that the actual planning process cannot be done fully from the group's point of view. Instead, the agents add actions to a plan by predicting how the other group members would continue from the results of alternative actions and then by evaluating possible results in terms of how they contribute to the group's goal.

Dunin-Kępicz and Verbrugge present a very sophisticated formal framework for teamwork [31, pp. 37–39]. Their definitions of collective intentions differ from the definitions used in the methods mentioned above in that they involve infinitely nested individual intentions and beliefs. Hence collective intention for them is an infinitary concept, but still reducible and thus conceptually different from the concept of group intention adopted here. Nevertheless, their approach to planning seems quite similar to the we-mode approach at least in the broad outline: The group starts from a collective intention to reach a goal. The goal is divided into a sequence of subgoals. For each subgoal, a means-end analysis provides actions needed to realize them. Finally, the actions are allocated to the group members and a temporal structure for the actions is devised [31, pp. 85–86]. They do not discuss team reasoning but they do analyze potential failures and present a method for re-planning and adjusting collective commitments as a response to changes in a dynamic environment.

Finally, there are some approaches that share with the we-mode approach the idea that a group of agents with a shared goal can be seen as a single agent which is trying to find a plan that satisfies its goal [27] or maximises its utility [11]. Under such an understanding, the multi-agent planning problem can be mapped into a single-agent planning problem, after which more efficient standard planning algorithms can be applied. Boutilier [11] presents the idea in the framework of decision-theoretic planning where uncertainty is present, whereas Crosby et al. [27] consider a more traditional setting with deterministic actions. In the latter, the multi-agent planning problem is first transformed, taking into account concurrency constraints arising from several agents manipulating the same object, into a single-agent planning problem in which agents

perform actions individually, one at a time. Once a solution to that problem is found, concurrency is re-introduced to the plan by combining consecutive actions into parallel or joint actions in accordance with the constraints. In the decision-theoretic setting, the problem is modelled as a multi-agent Markov decision process, and methods are presented for constructing an optimal policy for the group given the joint utility function. Boutilier [11] discusses coordination problems but only in the case of pure of pure coordination problems (in which there are several optimal outcomes) suggesting that cases like Hi-Lo are not problematic, similarly to the we-mode approach. As potential solutions to pure coordination cases, predefined social conventions or learned coordination are suggested. Pure coordination cases are similarly problematic for we-mode planning and the same methods can be applied.

In addition to work on teamwork and planning in DAI and MAS, there is another field that is closely related to the present work, namely human-aware task planning (HATP) [1, 18, 45]. Whereas multi-agent planning typically assumes that the actions of all agents can be planned, the idea of HATP is that only the actions of artificial agents can be planned but the possible actions of human beings must be taken into account already at the planning stage. Such approaches emphasise methods of action prediction and intention recognition. Current HATP methods fall in the I-mode category, because the plans for the robot and the human beings are strictly separated and there are no group plans. As the term implies, the planning is made only for the robot but possible plans of human(s) are considered and used as input for the planner. A we-mode approach to HATP would construct a plan for the human–robot team (even if the plan would only be followed by the robots).

HATP methods may be used in various situations: the human and the robot may have a shared goal, the robot may assist the human to achieve her goal, or the robot may have a goal of its own that it tries to achieve in the presence of humans (without preventing them to achieve their goals). In this paper the focus is on cases in which the robots and the human beings have a shared goal that they are all aware of. The human beings are autonomous and cannot be controlled by the planning algorithm, but they are not independent of it either: The robot or robots can use the planning algorithm to find ways to achieve the shared goal and suggest them to the human beings. Hence the task for robots is not only to predict and recognise human actions but possibly to suggest alternative actions as well.

6 Future Work

The method of we-mode planning presented here is merely a sketch and needs to be further refined in order for it to be

implemented and evaluated in practical scenarios. In addition there are several ways to extend the idea further. One way would be to allow the robot to learn the cost function from its experiences with the human collaborator. This would alleviate the need to have a previously made agreement on the evaluation function and it would also be more flexible because humans are often expected to deviate from the best plan in selecting their own actions and find innovative solutions in the course of plan execution. As we have sketched the method, the robot would be able to detect these situations and adjust its behaviour accordingly, but it would be a further step to learn from systematic deviations and adapt its behaviour to the human's way of doing things. This would require a learning algorithm to adjust the weights of the parameters in the evaluation function so as to have a better match between the suggested plan and the observed human behaviour. There are several alternative learning methods that could be used for such a purpose. One interesting approach is to train the robot with repeated interactions with a person whose evaluation of the quality of the interaction is then used as a feedback for the robot's training [51]. This kind of adaptability could also be used for personalising the robot behaviour relative to the collaboration partners.

The idea of using team reasoning in plan selection can also be employed in context of intention recognition. Most work on intention recognition has focussed on intentions in action, that is, on recognizing one action only. If we can reliably recognize one action, we can use the plan evaluation function to select the most plausible plan (or a set of most plausible plans) with a matching first action. This would lead to a method of recognizing future-directed intentions, that is, a method of plan recognition.

Another interesting future prospect would be to try to find a way of combining action learning and high-level planning methods like the one proposed here. The idea would be that instead of starting by a preprogrammed set of action descriptions, the robot would learn what kinds of consequences its actions have. A possible way to do that would use some sort of action-based or process-based representations that keep track of the changes in terms of the agent's action possibilities, or affordances [37], instead of traditional representations that try to directly model changes in the environment and are thus more vulnerable to frame problems [8]. There are now promising approaches that try to model an agent's learning of the consequences of its actions, for instance, a method called developmental learning [36]. These methods try to model the agent's development and are based on the idea of internal motivation in which the agent's curiosity leads it to try various actions which then differ in the way the agent finds performing these actions, and eventually action-sequences, internally rewarding. These ideas of internal motivation and affordance-based representation distinguish the approach from more

familiar reinforcement learning methods [74]. An open question that remains is how to bridge the gap between a method of learning the consequences of actions in concrete interaction with the environment guided by internal motivation and building abstract action plans targeted to satisfy a goal which is possibly externally given or collectively accepted by several agents. One possible starting point would be to build sociality in the agents so that they would be internally motivated to do things together with others, similarly to human beings.

7 Conclusions

There are several connections between philosophical theories of action and AI accounts of planning. Traditional theories of action were belief-desire accounts but they have to a large extent been replaced by belief-desire-intention accounts that are more amenable to planning than mere selection of the next action. However, decision theory and game theory still build on traditional desire-belief accounts, and there is no straightforward way to model intentions in decision-theoretic and game-theoretic formalisms. These theories have also been criticised for their inability to account for group actions. Team reasoning is one attempt to incorporate group action to game theory (see esp. [5]). Team reasoning is designed for the selection of the next action in a group context: The group members figure out the best joint action for the group and then perform their parts in that action. In the context of planning, the idea of team reasoning can be extended to the selection of plans: The agents figure out the best group plan, derive their sub-plans from that plan, and start executing them.

This leads to a planning method that can be called we-mode planning, employing Tuomela's [80] distinction between I-mode and we-mode. In contrast, many traditional planning accounts can be called I-mode accounts as they start from the plans of individual agents who then try to combine their plans by negotiating and adjusting their individual plans in order to make them mesh. I-mode planning is necessary for competitive scenarios in which the agents may have their private goals in addition to shared goals, we-mode planning suits well for cases with shared goals with no competition.

Acknowledgements This work has been supported by the Velux Foundation and the Academy of Finland. I thank S.M. Amadae, Pekka Mäkelä and three anonymous referees whose comments and suggestions significantly improved the quality of the paper. I also thank the PEN-SOR group in Aarhus and the audiences at the 2nd Workshop Towards a Framework for Joint Action (in conjunction with HRI 2015, Portland, USA, March 2, 2015) and at the Advanced Telecommunications Research Institute (ATR, Kyoto, Japan, August 28, 2015) for useful questions and comments.

References

1. Alami R, Clodic A, Montreuil V, Sisbot EA, Chatila R (2006) Toward human-aware robot task planning. In: AAAI spring symposium: to boldly go where no human-robot team has gone before, pp 39–46
2. Alami R, Robert F, Ingrand F, Suzuki S (1994) A paradigm for plan-merging and its use for multi-robot cooperation. *IEEE Int Conf Syst Man Cybern* 1:612–612
3. Amadae SM, Lempert D (2015) The long term viability of team reasoning. *J Econ Methodol* 22:462–478
4. Anscombe E (1957) *Intention*. Basil Blackwell, Oxford
5. Bacharach M (1999) Interactive team reasoning: a contribution to the theory of co-operation. *Res Econ* 53:117–147
6. Bacharach M (2006) *Beyond individual choice: teams and frames in game theory*. Princeton University Press, Princeton
7. Baier A (1970) Act and intent. *J Philos* 67:648–658
8. Bickhard MH, Terveen L (1995) *Foundational issues in artificial intelligence and cognitive science: an impasse and solution*. North-Holland, Amsterdam, The Netherlands
9. Boella G, Damiano R, Lesmo L (2000) Cooperation and group utility. In: *Intelligent agents VI. Agent theories, architectures, and languages*. Springer, pp 319–333
10. Boella G, Pigozzi G, Slavkovik M, Torre L (2011) Group intention is social choice with commitment. In: Vos Marina, Fornara Nicoletta, Pitt Jeremy V, Vouros George (eds) *Coordination, Organizations, institutions, and norms in agent systems VI: COIN 2010 international workshops, COIN@AAMAS 2010, Toronto, May 2010, COIN@MALLOW 2010, Lyon, France, August 2010, Revised selected papers*. Springer, Berlin, pp 152–171
11. Boutilier C (1996) Planning, learning and coordination in multi-agent decision processes. In: *Proceedings of the 6th conference on theoretical aspects of rationality and knowledge*. Morgan Kaufmann Publishers Inc, San Francisco, pp. 195–210
12. Bratman M (2014) *Shared agency: a planning theory of acting together*. Oxford University Press, Oxford
13. Bratman ME (1987) *Intention, plans, and practical reason*. Harvard University Press, Harvard
14. Bratman ME (1999) *Faces of intention: selected essays on intention and agency*. Cambridge University Press, Cambridge
15. Bratman ME, Israel DJ, Pollack ME (1988) Plans and resource-bounded practical reasoning. *Comput Intell* 4(3):349–355
16. Butterfill SA (2015) Planning for collective agency. In: SA Butterfill, Misselhorn C, (eds) *Collective agency and cooperation in natural and artificial systems*. Synthese Library
17. Churchland PM (1981) Eliminative materialism and the propositional attitudes. *J Philos* 78(2):67–90
18. Cirillo M, Karlsson L, Saffiotti A (2010) Human-aware task planning: an application to mobile robots. *ACM Trans Intell Syst Technol (TIST)* 1(2):15
19. Clodic A, Alami R, Chatila R (2014) Key elements for human-robot joint action. In: Seibt J, Hakli R, Nørskov M (eds) *Sociable robots and the future of social relations: proceedings of RoboPhilosophy*. IOS Press, Amsterdam
20. Cohen PR, Levesque H (1990) Intention is choice with commitment. *Artif Intell* 42:213–261
21. Cohen PR, Levesque HJ (1991) Teamwork. *Noûs* 25:487–512
22. Colman A, Pulford B, Rose J (2008) Collective rationality in interactive decisions: evidence for team reasoning. *Acta Psychol* 128(2):387–397
23. Colman A, Pulford B, Rose J (2008) Team reasoning and collective rationality: piercing the veil of obviousness. *Acta Psychol* 128(2):409–412
24. Copp D (1995) *Morality, normativity, and society*. Oxford University Press, Oxford
25. Corkill DD (1991) Blackboard systems. *AI Expert* 6(9):40–47
26. Cox JS, Durfee EH (2005) An efficient algorithm for multiagent plan coordination. In: *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pp. 828–835. ACM
27. Crosby M, Jonsson A, Rovatsos M (2014) A single-agent approach to multiagent planning. In: Schaub T, Friedrich G, O’Sullivan B (eds) *Eur conf artif intell*. IOS Press, Amsterdam, pp 237–242
28. Davidson D (1980) *Essays on actions and events*. Clarendon Press, Oxford
29. Dignum V, Dignum F (2001) Modelling agent societies: coordination frameworks and institutions. In: *Progress in artificial intelligence*. Springer, Berlin, , pp. 191–204
30. Dimopoulos Y, Adnan Hashmi M, Moraitis P (2012) μ -SATPLAN: multi-agent planning as satisfiability. *Knowl Based Syst* 29:54–62
31. Dunin-Kępicz B, Verbrugge R (2011) *Teamwork in multi-agent systems: a formal approach*. Wiley, New York
32. Durfee Ed, Zilberstein S (2013) Multiagent planning, control, and execution. In: Weiss G (ed) *Multiagent systems*, 2nd edn. MIT Press, Cambridge, pp 485–545
33. Durfee EH, Lesser VR (1987) Using partial global plans to coordinate distributed problem solvers. In: *Proceedings of the 10th international joint conference on Artificial intelligence*, vol. 2. Morgan Kaufmann Publishers Inc, San Francisco, pp. 875–883
34. Ephrati E, Rosenschein JS (1994) Divide and conquer in multi-agent planning. In: *Proceedings of the international AAAI conference*, vol. 1, p. 80
35. Gaertner W (2009) *A primer in social choice theory*, revised edn. Oxford University Press, Oxford
36. Georgeon O, Ritter F (2012) An intrinsically-motivated schema mechanism to model and simulate emergent cognition. *Cogn Syst Res* 15–16:73–92
37. Gibson JJ (1971) *The ecological approach to visual perception*. Houghton-Mifflin, Boston
38. Gilbert M (2001) Collective preferences, obligations, and rational choice. *Econ Philos* 17:109–119
39. Gilbert M (2013) *Joint commitment: how we make the social world*. Oxford University Press, Oxford
40. Govindarajan V, Bhattacharya S, Kumar V (2016) Human-robot collaborative topological exploration for search and rescue applications. In: *Distributed autonomous robotic systems*. Springer, Berlin, pp. 17–32
41. Grosz BJ (1996) Collaborative systems (AAAI-94 Presidential Address). *AI Mag* 17:67
42. Grosz BJ, Hunsberger L (2006) The dynamics of intention in collaborative activity. *Cogn Syst Res* 7:259–272
43. Grosz BJ, Kraus S (1996) Collaborative plans for complex group action. *Artif Intell* 86:269–357
44. Grosz BJ, Sidner C (1990) Plans for discourse. In: Cohen P, Morgan JL, Pollack ME (eds) *Intentions in communication*. MIT Press, Cambridge
45. Guitton J, Warnier M, Alami R (2012) Belief management for HRI planning. *BNC@ ECAI 2012*, p. 27
46. Haddawy P, Hanks S (1998) Utility models for goal-directed, decision-theoretic planners. *Comput Intell* 14(3):392–429
47. Hakli R, Mäkelä P (2016) Planning in the we-mode. In: Preyer G, Peter G (eds) *Critical essays on the philosophy of Raimo Tuomela with his responses*, *Studies in the philosophy of sociality*. Springer, Berlin
48. Hakli R, Miller K, Tuomela R (2010) Two kinds of we-reasoning. *Econ Philos* 26:291–320
49. Hakli R, Tuomela R (2003) Cooperation and we-mode preferences in multi-agent systems. In: Dittrich P, Kim J (eds) *Proceedings of the first international workshop on social life (SOLI 2003)*. ECAL 2003, Dortmund, 14–17 Sept, University of Dortmund

50. Hintikka J (2007) Socratic epistemology: explorations of knowledge-seeking by questioning. Cambridge University Press, Cambridge
51. Ikemoto S, Amor HB, Minato T, Jung B, Ishiguro H (2012) Physical human-robot interaction: mutual learning and adaptation. *IEEE Robot Autom Mag* 19:24–35
52. Jennings NR (1995) Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artif Intell* 75:195–240
53. Jennings NR (1993) Commitments and conventions: the foundation of coordination in multi-agent systems. *Knowl Eng Rev* 8:223–250
54. Kirsch A, Kruse T, Akin Sisbot E, Alami R, Lawitzky M, Brščić D, Hirche S, Basili P, Glasauer Stefan (2010) Plan-based control of joint human-robot activities. *KI-Künstl Intell* 24(3):223–231
55. List C, Pettit P (2011) Group agency: the possibility, design, and status of corporate agents. Oxford University Press, Oxford
56. Miller S (2001) Social action: a teleological account. Cambridge University Press, Cambridge
57. Nissim R, Brafman RI, Domshlak C (2010) A general, fully distributed multi-agent planning algorithm. In: van der Hoek W, Kaminka GA, Lespérance Y, Luck M, Sen S (eds) 9th international conference on autonomous agents and multiagent systems (AAMAS 2010), Toronto, May 10–14, 2010, Vol. 1–3, pp. 1323–1330. IFAA-MAS
58. Pacherie E (2011) Framing joint action. *Rev Philos Psychol* 2:173–192
59. Pacherie E (2013) Intentional joint agency: shared intention lite. *Synthese* 190:1817–1839
60. Pollock JL (1995) Cognitive carpentry: a blueprint for how to build a person. The MIT Press, Cambridge
61. Rao AS, Georgeff MP (1991) Modeling rational agents within a BDI-architecture. In: Allen J., Fikes R., Sandewall E. (eds) Proceedings of the 2nd international conference on principles of knowledge representation and reasoning, pp. 473–484, San Mateo, CA, USA, Morgan Kaufmann
62. Ratcliffe M (2007) Rethinking commonsense psychology: a critique of folk psychology, theory of mind and simulation. Palgrave Macmillan, Houndmills
63. Russell SJ, Norvig P (2010) Artificial intelligence: a modern approach, 3rd edn. Pearson Education, Upper Saddle River
64. Searle JR (1990) Collective intentions and actions. In: Cohen PR, Morgan J, Pollack ME (eds) Intentions in communication. MIT Press, Cambridge, pp 401–415
65. Searle JR (1995) The construction of social reality. The Free Press, New York
66. Sen A (2002) Rationality and freedom. Harvard University Press, Cambridge
67. Shoham Y, Leyton-Brown K (2009) Multiagent systems: algorithmic, game-theoretic, and logical foundations. Cambridge University Press, New York
68. Shoham Y, Tennenholtz M (1995) On social laws for artificial agent societies: off-line design. *Artif Intell* 73(1):231–252
69. Sims M, Corkill D, Lesser V (2008) Automated organization design for multi-agent systems. *Auton Agents Multi-Agent Syst* 16(2):151–185
70. Smith RG, Davis R (1981) Frameworks for cooperation in distributed problem solving. *Syst Man Cybern IEEE Trans* 11(1):61–70
71. Stirling WC (2013) A game-theoretic social model for multiagent systems. In: 2013 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 2718–2723. IEEE
72. Sugden R (1993) Thinking as a team: towards an explanation of nonselfish behavior. *Soc Philos Policy* 10:69–89
73. Sugden R (2000) Team preferences. *Econ Philos* 20:175–204
74. Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. MIT Press, Cambridge
75. Tambe M (1997) Towards flexible teamwork. *J Artif Intell Res* 7:83–124
76. Tollefsen DP (2015) Groups as agents. Polity Press, Cambridge
77. Tonino H, Bos A, de Weerd M, Witteveen C (2002) Plan coordination by revision in collective agent based systems. *Artif Intell* 142:121–145
78. Torreño A, Onaindia E, Sapena Ó (2014) Fmap: distributed cooperative multi-agent planning. *Appl Intell* 41:606–626
79. Tuomela R (2002) Philosophy of social practices: the collective acceptance view. Cambridge University Press, Cambridge
80. Tuomela R (2007) The philosophy of sociality. Oxford University Press, Oxford
81. Tuomela R (2013) Social ontology: collective intentionality and group agents. oxford University Press, Oxford
82. Tuomela R, Miller K (1988) We-intentions. *Philos Stud* 53:115–137
83. Velleman JD (1997) How to share an intention. *Philos Phenomenol Res* LVII:29–50
84. Weirich P (2004) Realistic decision theory: rules for nonideal agents in nonideal circumstances. Oxford University Press, Oxford
85. Wooldridge M (2000) Reasoning about rational agents. MIT Press, Cambridge
86. Wooldridge M, Jennings NR (1999) The cooperative problem-solving process. *J Logic Comput* 9(4):563–592
87. von Wright GH (1971) Explanation and understanding. Cornell University Press, Ithaca

Raul Hakli works as a postdoctoral researcher at the University of Helsinki, Finland. He has previously worked as an associate professor at Aarhus University, Denmark. Hakli holds a Ph.D. degree in theoretical philosophy and a M.Sc. degree in computer science. His interests include philosophy of social robotics, philosophy of social science, collective epistemology, collective intentionality and social ontology.