

13 The Role of Joint Commitment in Intersubjectivity

Antonella CARASSA, Marco COLOMBETTI,
Francesca MORGANTI

Abstract. Since the beginning of the Nineteen-eighties, cognitive scientists have shown increasing interest in a range of phenomena, processes and capacities underlying human interaction, collectively referred to as *intersubjectivity*. The goal of this line of research is to give an account of the various forms of human interaction, and in particular of the affective, attentional and intentional determinants of joint activity. The main thesis we develop in the paper is that so far the authors interested in intersubjectivity have neglected, or at least undervalued, an important aspect of joint activity, that is, the essentially normative character of collective intentionality. Our approach to joint activity is mainly based on Margaret Gilbert's theory of plural subjects. Gilbert's general idea is that joint activities should be regarded as activities carried out by individuals who stand to one another in a special relation, called *joint commitment*, which has an intrinsically normative nature. As we shall try to show, the concept of a joint commitment is a powerful tool to explain certain specific features of joint activities. In the paper we first point out certain explanatory inadequacies of the current models of intersubjectivity, and contend that such inadequacies depend on failing to appreciate the fundamental role of normativity in collective intentionality. We briefly sketch Gilbert's theory of plural subjects, and introduce the concept of a joint commitment, and then discuss some lines along which a psychology of plural subjects may be developed.

Contents

13.1	Introduction.....	188
13.2	Intersubjectivity and deontic normativity	189
13.3	Joint commitment	191
13.4	Steps to a psychology of plural subjects	193
13.5	Conclusions.....	199
13.6	Acknowledgments.....	200
13.7	References.....	200

Support us – Ask your institution/library/department to order our books or to apply for series membership - ISSN 1566-7677

To order a book
contact by fax or
by e-mail the offices
below:



IOS Press, Inc.
4502 Rachael Manor
drive
Fairfax, VA 22032
U.S.A.
Tel.: +1 703 323 5600
Fax: +1 703 323 3668
sales@iospress.com

IOS Press
Nieuwe Hemweg 6B
1013 BG Amsterdam
The Netherlands
Tel: +31 20 688 3355
Fax: +31 20 620 3419
info@iospress.nl

IOS Press,
c/o Ohmsha, Ltd.
3-1 Kanda Nishiki-cho
Chiyoda-ku
Tokyo 101
Japan
Fax: +81 3 3233 2426
(Books only)

IOS Press/
Gazelle Book Services Ltd
White Cross Mills
Hightown
Lancaster LA1 4XS
United Kingdom
Tel.: +44 1524 68765
Fax: +44 1524 63232
sales@gazellebooks.co.uk

Enacting Intersubjectivity

A Cognitive and Social Perspective on the Study of Interactions

Volume 10 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: F. Morganti, A. Carassa and G. Riva

May 2008, approx. 280 pp., hardcover

ISBN: 978-1-58603-850-2 **NEW**

Price: US\$161 / €115 / £81

From Communication to Presence

Cognition, Emotions and Culture towards the Ultimate Communicative Experience

Festschrift in honor of Luigi Anolli

Volume 9 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: G. Riva, M.T. Anguera, B.K. Wiederhold, F. Mantovani

September 2006, 323 pp., hardcover

ISBN: 978-1-58603-662-1

Price: US\$161 / €115 / £81

Global Data Management

Volume 8 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: R. Baldoni, G. Cortese, F. Davide and A. Melpignano

July 2006, 376 pp., hardcover

ISBN: 1-58603-629-7

Price: US\$161 / €115 / £81

The Hidden Structure of Interaction

From Neurons to Culture Patterns

Volume 7 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: L. Anolli, G. Riva, S. Duncan Jr. and M.S. Magnusson

May 2005, 304 pp., hardcover

ISBN: 1-58603-509-6

Price: US\$161 / €115 / £81

Ambient Intelligence

The Evolution of Technology, Communication and Cognition Towards the Future of Human-Computer Interaction

Volume 6 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: G. Riva, F. Vatalaro, F. Davide and M. Alcañiz

January 2005, 316 pp., hardcover

ISBN: 1-58603-490-1

Price: US\$161 / €115 / £81

Being There

Concepts, Effects and Measurements of User Presence in Synthetic Environments

Volume 5 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: G. Riva, F. Davide and W.A. IJsselstein

2003, 344 pp., hardcover

ISBN: 1-58603-301-8

Price: US\$161 / €115 / £81

Say not to Say: New Perspectives on Miscommunication

Volume 3 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: L. Anolli, R. Ciceri and G. Riva

2001, 288 pp., hardcover

ISBN: 1-58603-215-1

Price: US\$161 / €115 / £81

Towards CyberPsychology

Mind, Cognition and Society in the Internet Age

Volume 2 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: G. Riva and C. Galimberti

2001, 326 pp., hardcover

ISBN: 1-58603-197-x

Price: US\$161 / €115 / £81

Communications Through Virtual Technologies

Identity, Community and Technology in the Communication Age

Volume 1 Emerging Communication: Studies on New Technologies and Practices in Communication

Edited by: G. Riva and F. Davide

2001, 292 pp., hardcover

ISBN: 1-58603-162-7

Price: US\$161 / €115 / £81

13.1 Introduction

Since the beginning of the Nineteen-eighties, cognitive scientists have shown increasing interest in a range of phenomena, processes and capacities underlying human interaction, collectively referred to as *intersubjectivity*. The view advocated by these scientists is remarkably different from the one developed within the more traditional Theory of Mind approaches, either in the “Theory Theory” or in the “Simulation Theory” versions. Through the contributions of several authors [1-10] a novel view of human interaction is being developed, that is compatible with state-of-the-art knowledge on the phylogenesis and ontogenesis of interaction capacities, with the analysis of human experience worked out by phenomenologists, and with recent findings in the field of the neurosciences.

The goal of this line of research is to give an account of the various forms of human interaction, and in particular of the affective, attentional and intentional determinants of joint activity. Indeed, joint activity has long been a major issue for the social sciences and for analytical philosophy. Broadly speaking, the relevant theories can be classified in two groups: in the first group we have theories that attempt to give a *summative* account of joint activity, reducing it to the same building blocks underlying individual activity; the second group includes *non-summative* theories, which claim that joint activity requires certain special types of mental representations, often referred to as “collective intentionality” [11].

Most authors currently interested in intersubjectivity support some form of non-summative account. Observational and experimental results on non-human primates, human adults, and human children suggest that humans possess specific mental capacities, which enable forms of joint activity that are precluded to other primate species. A complete and coherent view of such capacities, however, is still beyond the state of the art. In this paper we aim to give a contribution to the construction of such a view. Our main thesis is that so far the authors interested in intersubjectivity have neglected, or at least undervalued, an important aspect of joint activity, that is, the essentially normative character of collective intentionality.

Our approach to joint activity is mainly based on Margaret Gilbert’s theory of plural subjects [12-15]. Gilbert’s general idea is that joint activities should be regarded as activities carried out by *plural subjects*, which can be viewed as sets of individual subjects who stand to one another in a special relation, named *joint commitment*, that has an intrinsically normative nature. As we shall try to show, the concept of a joint commitment is a powerful tool to explain certain specific features of human joint activities.

This article is structured as follows. In Section 2 we point out certain explanatory inadequacies of the current models of intersubjectivity, and contend that such inadequacies depend on failing to appreciate the fundamental role of normativity in collective intentionality. In Section 3 we briefly sketch Gilbert’s theory of plural subjects, and introduce the concept of a joint commitment. In Section 4 we discuss some lines along which a psychology of plural subjects may be developed. Finally, in Section 5 we draw some conclusions and delineate some directions for future research.

13.2 Intersubjectivity and deontic normativity

Since Trevarthen's distinction between primary and secondary intersubjectivity [1], it has become customary to differentiate among different types of intersubjectivity. For example, Stern [16] distinguishes between interaffective, interattentional, and interintentional sharing of experiences, and his distinction is taken up by other authors, like for example Ingar Brink [17]. Gärdenfors [5] advocates a similar position, but adds a fourth component, that is, representing the beliefs and knowledge of others. In general, the different components of intersubjectivity are taken to be stratified in levels, both from an evolutionary and a developmental point of view.

One of the leading themes of this area of research is to characterise human intersubjectivity with respect to the intersubjectivity of non-human primates, singling out the developmental phases at which specifically human structures and processes appear. Here we shall comment on a few works that we find representative of this approach.

In a paper on "What makes human cognition unique," Tomasello and Rakoczy [18] compare the impact on human social cognition of two key developmental moments, the first at about one year of age and the second at about four years. In the authors' terminology, the first ontogenetic step brings in *shared intentionality*, that is, the childrens' "ability to establish self-other equivalence, to take different perspectives on things, and to reflect on and provide normative judgement on their own cognitive activities" (p. 123). The second ontogenetic step, which comes after several years of "continuous interaction, especially linguistic interaction, with other persons" brings in *collective intentionality*, which ends up in the "comprehension of cultural institutions based on collective beliefs and practices such as money and marriage and government." While it is obvious that the second ontogenetic step is uniquely human, Tomasello and Rakoczy contend that a fundamental qualitative difference between human and non-human primates is already brought in by the first step, which sets the bases that make the second step possible.

One important aspect whose emergence brings from the first to the second developmental moment is *normativity*. Here we need to comment on this term, because it is used with different meanings, one of which is essential to our proposal. In the paper we are considering, the authors distinguish between *original* and *derived normativity* (p. 127). Original normativity is in fact coextensive with intentionality: every intentional state, as such, has conditions of satisfaction, and can therefore succeed or fail [19]. An intentional action, for example, may achieve or fail to achieve its purpose, and a belief may be true or false. Given that intentional states are the same thing as (mental) representations, we call this kind of normativity *representational*. Derived normativity has to do with the collectively accepted functions of artefacts. A fork is for bringing solid food to one's mouth, a switch is to turn the light on and off, and so on: functions are normative in the sense that they tell us how an artefact ought to be used. We call this kind of normativity *functional*. Besides representational and functional normativity, however, there is a third important kind of normativity, that we call *deontic*. Deontic normativity has to do with obligations and rights, in particular with *directed* obligations and rights, that is, the obligations and rights that a subject has relative to other subjects. Deontic normativity is often believed to come about

only with complex cultural products like legal systems, regulations, contracts and the like. On the contrary, we shall defend the idea that a form of deontic normativity is already there in every kind of joint activity, being a constitutive component of collective intentionality. If this is the case, representational and functional normativity, although essential for human cognition, are not sufficient to account for the normativity of collective intentionality.

A second paper we want to discuss here is Brink and Gärdenfors's work on cooperation and communication in apes and humans [6]. The authors argue that non-human primates are incapable of future-directed cooperation, which "concerns new goals that lack fixed value" and "requires symbolic communication and context-independent representations of means and goals" (p. 484). In this paper, Brink and Gärdenfors remark that one of the key aspects of cooperation, that is, the guarantee of proper compensation for one's efforts, becomes hazardous with future-directed cooperation. As the authors put it, "in the case of as yet imaginary goals, compensation becomes much more of a venture than a safe strategy" (pp. 488-489).

Brink and Gärdenfors consider cooperation within a game-theoretical framework. Much of their argument is based on the difficulty of developing reliable expectations about the others' behaviour; expectations are regarded as a purely informational phenomenon, and there is little concern for the normative component of interaction. Toward the end of the paper, the authors turn their attention to aspects of cooperation that involve deontic normativity, like feelings of shame and the expectation of sanctions from the rest of the group related to defective behaviour. This line of thought, however, is not pursued to the point of considering future-directed cooperation as a form of interaction intrinsically driven by deontic normativity. As Brink and Gärdenfors remark, the core problem of future-directed cooperation is that "it will be difficult to make estimates concerning the behaviour of other agents on the basis of previous experience, since the situation is new and unknown" (p. 499). We shall argue in the rest of this paper that providing a sound basis for estimating the future behaviour of other agents is the primary function of joint commitments.

Another relevant work is Gärdenfors's article on the cognitive and communicative demands of cooperation [4], where the author presents a table of different forms of cooperation, at least three of which ("Commitment and contract", "Cooperation based on conventions", "The cooperation of *Homo oeconomicus*"; p. 20) seem to us to involve deontic normativity. Among the demands of these forms of cooperation a special place is given to symbolic communication, while the role of deontic normativity is ignored. For example, it is said (p. 14) that "to promise something only means that you intend to do it. On the other hand, when you *commit* yourself to a second person to do an action, you intend to perform the action in the future, the other person wants you to do it and intends to check that you do it, and there is joint belief concerning these intentions and desires [20]. Unlike promises, commitments can thus not arise unless the agents achieve joint beliefs and have anticipatory cognition." Two criticisms can be made to this position. The first is that promising creates obligations, and is not limited to letting someone else know what one intends to do (see for example [21]). The second is that committing to a second person to do an action cannot be analysed only in terms of epistemic and volitional states like beliefs, desires, and intentions. So, on the one hand to promising *is* committing oneself; on the other

hand, there is more to commitment than achieving joint beliefs and having anticipatory cognition.

In a series of important works, Hannes Rakoczy investigates the children's ability to construct and exploit social reality. In [22] the author interprets young children's pretend play as examples of cooperative activities involving the collective definition of fragments of social reality (understood along the lines of Searle's account [23]). Rakoczy's interpretation of pretend play comes very close to the concept of joint commitment that we shall discuss in the following sections: in Rakoczy's words, "a we-intention essentially involves some basic form of commitment to acting together, analogous to the individual commitment of actors in solitary actions, but different in that not only my own desires and intentions provide reasons for further intentions and actions, but now the collaborator's actions and intentions provide reasons for me to act accordingly in the course of the joint action" (p. 120). Still it seems to us that the deontic nature of joint commitment is not fully appreciated. As a consequence, commitments are regarded, somewhat vaguely, as "quite minimally involving an appreciation of normative inferential (reason giving) relations between collaborators' and own actions and the willingness to respect these relations in the pursuit of acting together successfully" (p. 120). We believe that the best way to characterise such "normative inferential (reason giving) relations" is to regard them as deontic relationships (i.e., directed obligations, rights, and entitlements) generated by joint commitments.

The discussion we have carried out so far suggests that deontic normativity may indeed be a fundamental component of human interaction. If this is the case, we believe, theories of intersubjectivity will have to grant deontic normativity the room it deserves. In the rest of this paper we shall try to give an initial contribution in this direction, starting from a concise introduction to Gilbert's concept of a plural subject.

13.3 Joint commitment

Gilbert's theory of joint activities is centred on the concept of a plural subject and to the strictly related normative notion of a joint commitment. The importance of normative concepts in general, and of commitment in particular, for understanding human interactions has been recognised long ago. For example, in their pioneering book Winograd and Flores [24] wanted "to counteract the forgetfulness of commitment that pervades much of the discussion (both theoretical and commonplace) about language" (p. 76). In argumentation theory, commitment-based models have been proposed and discussed since the concept of a commitment store was introduced by Hamblin [25] and later developed by Walton and Krabbe [26]. Very recently, John Searle [27] has advocated a view of human language in which deontic normativity is regarded as a basic constitutive component, side by side with representative power and syntactic compositionality.

In the current landscape, Gilbert's theory is unique in placing deontic normativity at the very heart of collective intentionality. Gilbert's idea is that all genuinely collective phenomena (like joint activities, collective beliefs, group feelings, social conventions, and so on) involve a normative component, called *joint commitment*, that turns the set of interacting subjects into a *plural subject*.

The idea of a “plural subject” may sound metaphysically suspicious, but in fact it is nothing more than a group of individuals bound by a joint commitment. In turn, for a group of individuals to be bound by a joint commitment it is necessary and sufficient for them to entertain certain mental representations.

What it means for a group of individuals to be jointly committed to doing *X* (or believing *X*, or feeling *X*, and so on) is explained by Gilbert in several books and papers (see in particular [13], Part III; [14], Chapter 4; and [15], Chapter 7). Below we briefly describe the main features of this important concept.

A subject may be individually committed to do *X*, for example as a result of a personal decision: such a decision may be rescinded, but until this does not happen the subject is committed to do *X*. Being committed to do *X* is a reason (although not a sufficient cause) for the subject to do *X*; however, in the individual case the subject is the only “owner” of the commitment, and can rescind the commitment as he or she pleases.

Contrary to individual commitments, a joint commitment is a commitment of *two or more subjects*, which we shall call *parties* of the joint commitment, to engage in a common enterprise as a single body. Taken together, a number of subjects jointly committed to do *X* form a plural subject of doing *X*. The main difference between individual and joint commitments is that joint commitments are not separately “owned” by their parties, but they are, so to speak, collectively owned by all parties at the same time.

Joint commitments may arise as a result of an agreement. However, explicit agreements are not necessary: according to Gilbert, what is necessary and sufficient to create a joint commitment, and thus to set up a plural subject, is that it is common knowledge of all parties that every party is ready to engage in some joint enterprise. Such common knowledge may derive from explicit agreements, but also from less structured communicative exchanges and, in many cases, from shared understanding of a culturally meaningful context.

Let us consider a few examples. Ann may say to Bob, “I’m going for a walk, would you like to come?” If Bob answers, “Yes, sure!”, then it will be common knowledge of Ann and Bob that they are both ready to engage in a walk together, and this suffices to create a joint commitment to have a walk together. In certain situations, like for example a dinner party, it will be common knowledge of all participants (without the need of specific communicative exchanges) that all parties are ready to carry out certain kinds of joint activities, like chatting or dancing, with the other participants. Indeed, joint commitments are much more common in human interaction than one may think. Even an apparently unilateral promise, like Bob saying to Ann “I promise to come visit tomorrow evening,” if accepted by the Ann creates a joint commitment, because while Bob is now obliged to do what he promised, Ann is obliged to stay at home and welcome Bob.

For our current purpose, the main feature of joint commitments is that they generate deontic relationships, like directed obligations and the correlative rights and entitlements. (A directed obligation is an obligation that a subject, the *debtor* of the obligation, owes to another subject, the *creditor* of the obligation. Every directed obligation brings about a correlative right of the creditor to the debtor.) If *n* subjects are jointly committed to do something, then every subject is obligated to all other subjects to do his or her part of the joint activity, and has the right that all other subjects do their parts. It is characteristic of joint commitments that all such obligations are *created simultaneously*, and are *interdependent* in the sense that if

one of the parties fails to fulfill one of his or her obligations, then the joint commitment is violated. What exactly this amounts to depends on a variety of circumstances, including the number of members of the plural subject. In particular, in the case of two parties the violation of an obligation by one of them rescinds the joint commitment.

According to Gilbert, every genuine case of joint activity is an activity carried out by a plural subject, and thus involves joint commitments. It is important to understand that such commitments are not imposed to the parties from the outside, but are “internal” to the joint activity. For example, when a group of people engage in a game, we do not need to assume that there is some external source of obligations that compels the participants to follow the rules of the game: rather, engaging in a game together is by itself a source of obligations.

Our brief presentation of plural subjects and joint commitments raises a number of important issues: What is the function of joint commitment? To what kind of things can people jointly commit? What kinds of joint commitments are involved in joint activities? What kinds of cognitive processes underlie joint commitment? How do people make and maintain joint commitments? Since what age are humans able to participate in joint commitments? Some of these questions are logical, in the sense that they concern the function and structure of joint commitments, and some are psychological, in the sense that they directly concern human mental capacities. In the two following sections we shall submit some initial answers to the previous questions.

13.4 Steps to a psychology of plural subjects

13.4.1 The function of joint commitments

At least since Aristotle, we understand human beings as rational animals. If we construe the concept of a reason broadly enough, humans are not the only rational species on Earth. But, based on the experimental evidence collected so far, it is generally accepted that humans are the only species that can deploy a very specific type of rationality, that is, the ability to plan their future. Given that anticipatory planning is one of the distinctive features of *Homo sapiens* [28], it is not surprising that so much attention has been devoted to it by scholars of disciplines like cognitive psychology, philosophy of mind, economy, and artificial intelligence.

The function of future-directed intentions, or *prior intentions* in Searle’s terminology [19], has been analysed, among the others, by Michael Bratman [29], who stresses their characteristic role of coordinating practical reasoning. Indeed future-directed intentions, organised into complex plans, allow human subjects to reason within stable tracks directed to specific purposes, thus avoiding the risk of being misled by fluctuating motivations.

From an analysis of current literature, it seems that most authors do not find it problematic to extend the stabilising function of intentions from individual to joint action: even the most complex forms of cooperation are assumed to require nothing more than the ability to share nested intentions and beliefs. At present, some authors are starting to see that this is not sufficient. For example, introducing contracts as a sophisticated form of human cooperation, Gärdenfors [5] states that “If we agree that I shall deliver a hen tomorrow in exchange for the axe you have

given me now, I believe that you believe that I will deliver the hen and you believe that I believe that our agreement will then be fulfilled, etc. Furthermore, a contract depends on the possibility of *future* sanctions and thus on anticipatory cognition: "If I don't deliver the hen, you or the society will punish me for breaching the agreement" (p. 20). There is an attempt, here, to reduce deontic normativity to the expectation of punishment, and thus to a purely epistemic phenomenon (a future-directed belief). However, even before we ask ourselves whether this reduction is psychologically plausible, we face a conceptual problem here, because the very concept of a punishment is deontic. Indeed a punishment is more than just a cost imposed to the subject by someone else: it is a cost *rightly* imposed to the subject by someone else.

In our view joint commitments play, in the case of collective activities, a stabilising role analogous to that played by future-directed intentions in the case of individual actions. Joint commitments achieve this function by creating directed obligations, thus decoupling future actions from possibly fluctuating motivations. Consider the following example: by entering a suitable joint commitment, Ann and Bob may form a plural subject of mutual care. While the joint commitment is in force, Ann and Bob will be obliged to carry out appropriate actions, like proving support to each other in difficult situations, and so on. Given the joint commitment, it is not important whether Ann or Bob are continually motivated to support each other: the reason for doing so is now an obligation created by the joint commitment.

13.4.2 *The structure of joint commitments*

To what kind of things can people jointly commit? Or, in other words, what can be the *content* of a joint commitment?

The most obvious examples of joint commitments concern joint activities. For example, by jointly committing to have a walk together, Ann and Bob create obligations concerning their future behaviour. But Margaret Gilbert argues that joint commitments are more general: for example, for a group of people to entertain a collective belief means that the group constitutes a plural subject of believing something. A joint commitment to believe, say, that all men are created equal, will carry out its function in much the same way as a joint commitment to do something together: that is, by creating directed obligations to perform appropriate actions, which will be determined case by case in a context-sensitive way. In the "all men are created equal" case, for example, every party of the plural subject is obliged to act accordingly, by treating every person with equity, by reacting to blatant discriminations, and so on.

An important feature of Gilbert's non-summative treatment (see for example [13], Chapter 14) is that a plural subject may collectively believe that *p* even if not all the parties (indeed, in extreme cases, none of them) believes that *p*. The point with collective belief is not what individuals actually believe, but what are their obligations given their joint commitment to believe something.

Given the significance of affective states in intersubjectivity, it is important to understand whether people can create plural subject of *feeling* something. This seems to be the case when, for example, a team is proud of a remarkable achievement, or a group of people are sorry for a distressing event occurred to a common friend: statements like "We are proud of being the first to land on Mars"

or “*We* are so sorry your house burnt to ashes” reveal that the feeling of pride or sorrow is attributed to a plural subject. Analogously to the case of collective beliefs, a joint commitment to feeling something will carry out its function by creating obligations to perform appropriate actions, independently of the fact that the parties actually have the relevant feeling.

Recently, Margaret Gilbert suggested that also joint attention is a plural subject phenomenon [30]. The idea is that joint attention is best “understood in terms of a joint commitment to attend as a body to some particular in the environment of the parties” (p. 7). According to Gilbert, joint attention requires mutual recognition, which in turn presupposes common knowledge of co-presence. Joint commitments thus appear to be a pervasive aspect of intersubjectivity.

From the point of view of a theory of intersubjectivity, it is necessary to understand the relationships between joint commitments and psychological states. In Table 1 we propose a systematic view of all such states. We first classify psychological states into affective, attentional, and intentional states. Here the term “intentional” is to be understood as a synonym of “representational,” in line with the philosophical theory of intentionality: perceptions, beliefs, desires, and intentions are all examples of intentional states. We consider purely affective and attentional states as psychological states of a single individual: a distinction between individual and interpersonal states can be drawn only for intentional states, because interpersonality is achieved through representations.

Intentional states are classified as individual or interpersonal. Examples of individual intentional states are intending to do something (in the future), intentionally doing something (right now), perceiving something, desiring something, and so on. Interpersonal intentional states are, by definition, those intentional states of a subject whose content involves psychological states of other subjects.

There are basically three ways in which a psychological state of an individual may become interpersonal. The first way is through perception: a subject may directly perceive an affective, attentional, or intentional state of another subject. Indeed, the possibility of directly perceiving psychological states of another subject (inclusive of intentional states) is an important tenet of current theories of intersubjectivity (see for example [31]).

The second way in which a psychological state may become interpersonal is through sharing: a shared state, in our terminology, is a state that is “out in the open” (to adopt the felicitous expression used by Gilbert to describe situations of common knowledge [13]) but to which there is no joint commitment. Again, the shared state may be affective, attentional, or intentional. As an example of a shared attentional state that involves no joint commitment consider two criminals trying to kill each other and standing a few meters apart, with the only gun at their disposal lying on the ground right between the two of them. In this situation there is shared attention to the other and to the gun, but of course the two criminals do not form a plural subject of paying attention to the other and to the gun.

Finally, a psychological state may be interpersonal by being joint (or collective). By this we mean that the relevant subjects are jointly committed to entertaining such a state. As we have already remarked, the content of a joint commitment may be any affective, attentional, or intentional state.

affective		attentional		intentional		
		individual		subject A intends to do X, does X (intentionally), perceives X, believes X, desires X, etc.		
				affective	attentional	intentional
subject A has emotion X	subject A attends to object X	perceived		subject A perceives that subject B has emotion X	subject A perceives that subject B attends to object X	subject A perceives that subject B intends to do X, etc.
		interpersonal				
		shared	it is out in the open for subjects A and B that one of them (or both of them) has emotion E	it is out in the open for subjects A and B that one of them (or both of them) attend to object X	it is out in the open for subjects A and B that one of them (or both of them) intends to do X, etc.	
		joint		A and B are jointly committed to have emotion X (as a body)	A and B are jointly committed to attend to object X (as a body)	A and B are jointly committed to intend X, do X, etc. (as a body)

Table 1. A classification of psychological states.

Only some intersubjective processes involve shared intentional states, and an even smaller fraction involve joint intentional states. Among these, however, we find a very significant category of intersubjective processes, that is, joint activities which, in particular, presuppose joint intentions. It is important not to confuse our distinction between shared and joint psychological states with other kinds of distinctions, like for example the one between coordination, collaboration, and cooperation. All types of joint activities involve some kind of dependency between the actions performed by the different parties as part of the joint activity. The difference between coordination, collaboration, and cooperation concerns what we could call *degree of coupling*: while in the case of coordination, typically based on a loosely synchronised execution of individual plans, coupling is kept to a minimum, cooperation involves a very high degree of coupling, achieved through the collective execution of a common plan. However, all such types of joint activities involve joint commitments, even if their contents will be different for different kinds of joint activities. Suppose for example that Ann and Bob decide to have dinner together at Bob's apartment at 8 pm. As both of them are very busy, they will separately buy some ready-made food: Ann will get the entrées and the wine, and Bob will take care of the main course. Ann and Bob are now bound by a joint commitment that generates at least the following obligations: that Ann gets the entrées and the wine, goes to Bob's apartment around 8 pm, and then has dinner with Bob; that Bob gets the main course, will be at his apartment around 8 pm, and then has dinner with Ann. The first part of the joint activity, when Ann and Bob separately get the food, has a very low degree of coupling; in spite of this, however, there is a genuine joint commitment binding Ann and Bob to act as agreed.

13.4.3 Cognitive requirements

A plural subject is a group of people bound by a joint commitment. In turn, the members of the group are bound by a joint commitment if, and only if, they have certain psychological states. But what kind of psychological states are involved?

As we have remarked in Section 2, many authors agree that some form of commitment is essential at least for the most complex types of joint activity. There is, however, no attempt to explain what it takes for a subject to commit to a course of action.

Given that joint commitments involve deontic normativity, it is tempting to consider them as a case of moral thought. This, however, may not prove a fruitful approach, because the deontic relationships produced by a joint commitment appear to be different from moral obligations. In our opinion, a major difference between moral obligations and the obligations of joint commitments is that, contrary to the former, the latter are intentionally created by people. To clarify the difference, suppose that Bob, motivated by his moral conviction that one should care after the ill, agrees with an elderly neighbour of his that he will soon visit her at the hospital. While one may dispute whether visiting his neighbour was really a moral obligation of Bob's, there is no doubt that after promising Bob is obliged to do so. Even if Bob changes his idea about the moral obligation of caring after the ill, he will still be obliged, because he freely committed his will by making an agreement.

In any case, it is clear that the ability to enter into joint commitments presupposes the ability to understand obligations, rights, entitlements, and the like. We believe that such ideas cannot be reduced to non-deontic psychological states, like beliefs and intentions. Being obliged to do *X* is more than just expecting that if one does not do *X* something bad will happen. Suppose for example that Bob, together with a group of clients of the local branch of his bank, is caught in a robbery and is ordered by a masked guy to sit on the floor and stay still. Bob knows that something bad will happen if he tries to escape, and in some sense of the word we can actually say that he is *obliged* to sit on the floor and stay still. However, this obligation cannot be considered as a deontic relationship between Bob and the masked criminal.

The problem of finding suitable primitives to which all deontic ideas can be reduced has long been considered in such fields as the philosophy of law and deontic logic. In [23], John Searle defends the idea that all deontic relationships can be defined in terms of one primitive, like for example obligation. This means that any being capable of entertaining thoughts of the kind "I am obliged to ..." would be able to represent all deontic relationships. A different approach, developed for the first time by Anderson in the field of deontic logic [32], is to reduce deontic notions like obligation and right to a lower-level concept, like violation. To understand this idea, suppose again that Ann and Bob agreed that Bob will visit Ann at her summer cottage next Sunday. Bob, in particular, is now obliged to Ann to go to Ann's cottage next Sunday. This idea may take the following form: "If I do not go to Ann's cottage next Sunday, then I make a violation to Ann." What seems to be sufficient to have joint commitments is therefore a concept of *directed violation*, that is, of a violation relative to some individual.

A different approach is taken by Margaret Gilbert, who proposes to understand the obligations of joint commitments in terms of "owing" (see [15], Chapter 11).

The general idea is that once a joint commitment has been created, every party owes certain actions to the other parties; symmetrically, all parties “own,” even if they do not yet possess, the actions that are owed to them. It may indeed be the case that the concept of owing can be reduced to the more primitive notion of violation we have previously introduced. But it may also be the other way round: the concept of owing may be a psychological primitive, on which more complex aspects of social cognition are based. In any case, we think that only empirical research may settle this issue.

Whether joint commitments are based on a primitive notion of directed violation, or an a primitive notion of owing, it would be extremely interesting to discover at what age human beings are capable of building the relevant representations. Since the publication of Kohlberg’s pioneering paper on moral stages [33], much research has been carried out on the development of moral reasoning, but situations of joint commitment have not been a primary concern. Monika Keller and colleagues [34] reported on some experiments in which children were asked to reason on situations in which an agreement between a child and his mother was either fulfilled or violated, and found that even children of about three years of age were able to correctly detect situations of agreement violation. This kind of experiments, though, rather than testing whether children are able to engage in joint commitments in first person, test the children’s ability to reason on third-person situations of joint commitments. Moreover, due to the cognitive complexity of the experimental task, such experiments can be run only on children of at least three years of age. However, recent literature on the early development of sociality (like [18, 22, 35]) suggest that certain fundamental social abilities show up considerably earlier.

Recently Maria Gräfenhain and colleagues [36] reported on an experiment aimed to identify the presence of joint commitments in social play contexts. The preliminary results show that the deontic implications of joint commitment begin to emerge at two years, and are clearly established by three years of age. Of course, further research is needed before we have a clear picture of the ontogeny of joint commitment.

13.4.4 The life cycle of plural subjects

As everything on earth, plural subjects have a beginning, a period of life, and an end – that is, a life cycle. Describing all possible life cycles of plural subjects is beyond the scope of this article. In what follows we shall just sketch a few important points.

As we have already remarked in Section 2, the joint commitment that constitutes a plural subject may be created through an explicit agreement or may come to exist as an implicit consequence of the parties’ interaction. For example, at a dancing party two persons may just start dancing together without prior agreement: the joint activity they engage in will imply a joint commitment to dance together at least for a while.

Margaret Gilbert suggests that the necessary and sufficient condition for a group of people to form a plural subject is that it is out in the open (i.e., common knowledge) that all members of the group are ready to engage in some common enterprise. Often, the readiness to engage in the common enterprise will mature through a more or less lengthy phase of negotiation.

A plural subject exists as long as the underlying joint commitment is in force. During this period the parties of the plural subject are bound by a network of deontic relationships, produced by the joint commitment in a context-dependent way. Such deontic relationships may be classified into two classes: basic and derivative. The basic deontic relationships are the directed obligations, rights, entitlements, and so on that are directly related to carrying out the common enterprise. For example, if Ann and Bob agreed that Bob will visit Ann at her summer cottage next Sunday, then Ann is obliged to Bob to be at her summer cottage next Sunday, Bob has the correlative right to Ann that Ann be at her summer cottage next Sunday, Bob is entitled to go to Ann's summer cottage next Sunday, and so on. The derivative deontic relationships concern the management of the joint commitment in the face of violations by the parties of the plural subject. For example, in case Ann is not at her cottage next Sunday, Bob has the derivative entitlement to rebut; or, if after their agreement Ann discovers it will be impossible for her to be at her summer cottage next Sunday, she has the derivative obligation to tell Bob and to provide a suitable justification.

A plural subject may come to an end in many different ways. In some cases, the underlying joint commitment will have a well-defined deadline: consider for example the joint commitment of moving a table together, which terminates when the action is completed. In other cases the deadline will be only vaguely defined, and consequently the termination of the joint commitment will require some form of explicit or implicit negotiation. As an example, consider the joint commitment of going for a walk together: given that "a walk" is a vague concept, sooner or later the parties will start negotiating the end of the common enterprise, for example by saying "I start feeling tired now" or "I'm afraid I have to go back now, I have to dress up for dinner." A plural subject may also come to an end due to a violation by one of the parties. In the case of two parties, a violation by one of them is sufficient to wipe out the joint commitment, thus freeing the other party of all obligations. With more than two parties the situation is more complex, and we shall not try to deal with it here.

13.5 Conclusions

In this article we have argued that joint activities involve a particular form of deontic normativity, that following Margaret Gilbert we call joint commitment. Joint commitments arise when a number of subjects make it overt that they are ready to engage in a common enterprise, and generate deontic relationships (directed obligations, rights, and entitlements) among these subjects. By creating such deontic relationships, joint commitments play an essential role in stabilising interaction, which is particularly relevant to anticipatory planning.

More work needs to be done before we can form a satisfactory picture of the deontic normativity of joint commitments as part of the general phenomenon of human intersubjectivity. Below we mention some issues that seem to us to be important.

At the theoretical level, we think that the relationship between joint commitments and moral obligations is in need of clarification. Intuitively, the deontic normativity of joint commitments appears to be distinct from moral

normativity. However, what this difference exactly amounts to, and what are the relationships between commitment and morality is still unclear.

At the empirical level, there seems to be at least four areas in which it would be interesting to carry out experimental work. First, research on the ontogenesis of joint commitment, which as we have seen has already started, may contribute to our understanding of the development of sociality; moreover, considering results in the light of the available literature on moral development may help to understand the relationships between the normativity of commitments and moral normativity. Second, the analysis of adult interactions may clarify important aspects of the life-cycle of plural subjects and the relationships between joint commitments and what we have called the degree of coupling of collective activities. Third, the analysis of narratives may shed light on the affective side and on the first-person perspectives of joint commitment. Finally, it would be interesting to find out how certain types of cognitive and/or relational disorders, due to brain injuries or neurological disorders, influence the human capacity to engage in joint commitments.

13.6 Acknowledgments

We are grateful to Margaret Gilbert for kindly engaging with us in a stimulating electronic discussion on joint commitment.

13.7 References

- [1] C. Trevarthen, Communication and cooperation in early infancy: A description of primary intersubjectivity. In M. Bullowa (Ed.), *Before speech: The beginning of interpersonal communication*, (pp. 321–347). Cambridge: Cambridge University Press, 1979.
- [2] M. Tomasello, M. Carpenter, J. Call, T. Behne & H. Moll, Understanding and sharing intentions: The origins of human cognition. *Behavioral and Brain Sciences*, 28, 675–735, 2005.
- [3] S. Gallagher, *How the body shapes the mind*. New York: Oxford University Press, 2005.
- [4] P. Gärdenfors, The cognitive and communicative demands of cooperation. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson (Eds.), *Hommage à Wlodek: Philosophical papers dedicated to Wlodek Rabinowicz*, 2007. Online: [http://www. fil.lu.se/hommageawlodek](http://www.fil.lu.se/hommageawlodek).
- [5] P. Gärdenfors, Evolutionary and developmental aspects of intersubjectivity. In H. Liljenström & P. Århem (Eds.), *Consciousness transitions: Phylogenetic, ontogenetic and physiological aspects*, Amsterdam: Elsevier, to appear. Online: <http://project.sol.lu.se/sedsu/publications/2006-Gardenfors-EvolIntersubj.pdf>.
- [6] I. Brink & P. Gärdenfors, Co-operation and communication in apes and humans. *Mind & Language*, 18, 484–501, 2003.
- [7] J. Zlatev, What's in a schema? Bodily mimesis and the grounding of language. In B. Hampe (Ed.), *From perception to meaning: Image schemas in cognitive linguistics*, (pp. 313–342). Berlin: Mouton.
- [8] S. Bråten (Ed), *On being moved: From mirror neurons to empathy*. Amsterdam/Philadelphia: John Benjamins, 2007.
- [9] S. Hurley & N. Chater (Eds), *Perspectives on imitation: From cognitive neuroscience to social science, 1*. Cambridge, MA: MIT Press, 2005.
- [10] S. Hurley & N. Chater (Eds), *Perspectives on imitation: From neuroscience to social science, 2*. Cambridge, MA: MIT Press, 2005.
- [11] D. Tollefsen, Collective intentionality, *The Internet Encyclopedia of Philosophy*, 2004. Online: <http://www.iep.utm.edu/c/coll-int.htm>.
- [12] M. Gilbert, *On social facts*. London and New York: Routledge, 1989.
- [13] M. Gilbert, *Living together: Rationality, sociality, and obligation*. Lanham: Rowman & Littlefield, 1996.

- [14] M. Gilbert, *Sociality and responsibility: New essays in plural subject theory*. Lanham: Rowman & Littlefield, 2000.
- [15] M. Gilbert, *A theory of political obligation*. Oxford: Clarendon Press, 2006.
- [16] D. N. Stern, *The interpersonal world of the infant*. New York: Basic Books, 1985.
- [17] I. Brinck, The role of intersubjectivity for the development of intentional communication. In J. Zlatev, T. Racine, C. Sinha & E. Itkonen (Eds), *The shared mind: Perspectives on intersubjectivity*, Amsterdam: John Benjamins, to appear. Online: <http://www.fil.lu.se/publicationfiles/pp103.doc>
- [18] M. Tomasello & H. Rakoczy, What makes human cognition unique? From individual to shared to collective intentionality. *Mind & Language*, 18, 121–147, 2003.
- [19] J. R. Searle, *Intentionality: An essay in the philosophy of mind*. Cambridge: Cambridge University Press, 1983.
- [20] B. Dunin-Kepliz & R. Verbrugge, A tuning machine for cooperative problem solving. *Fundamenta Informatica*, 21, 1001–1025, 2001.
- [21] J. R. Searle, *Speech acts: An essay in the philosophy of language*. Cambridge: Cambridge University Press, 1969.
- [22] H. Rakoczy, Pretend play and the development of collective intentionality. *Cognitive Systems Research*, 7, 113–127, 2006.
- [23] J. R. Searle, *The construction of social reality*, New York: Free Press, 1995.
- [24] T. Winograd & F. Flores, *Understanding computers and cognition: A new foundation for design*. Norwood: Ablex, 1986.
- [25] C. L. Hamblin, *Fallacies*. London: Methuen, 1970.
- [26] D. N. Walton & E. C. W. Krabbe, *Commitment in dialogue: Basic concepts of interpersonal reasoning*. Albany: State University of New York Press, 1995.
- [27] J. R. Searle, What is language: Some preliminary remarks. In S. L. Tsohatzidis, *John Searle's philosophy of language: Force, meaning, and mind*, (pp. 15–48). Cambridge: Cambridge University Press, 2007.
- [28] A. Gulz, *The planning of action as a cognitive and biological phenomenon*, Lund: Lund University Cognitive Studies 2, 1991.
- [29] M. E. Bratman, *Intention, plans, and practical reason*, Cambridge, MA: Harvard University Press, 1987.
- [30] M. Gilbert, Mutual recognition, common knowledge, and joint attention. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson (Eds.), *Hommage à Wlodek: Philosophical papers dedicated to Wlodek Rabinowicz*, 2007. Online: [http://www. fil. lu.se/hommageawlodek](http://www.fil.lu.se/hommageawlodek).
- [31] S. Gallagher & D. Hutto, Understanding others through primary interaction and narrative practice. In J. Zlatev, T. Racine, C. Sinha & E. Itkonen (Eds), *The shared mind: Perspectives on intersubjectivity*, Amsterdam: John Benjamins, to appear. Online: <http://pegasus.cc.ucf.edu/~gallaghr/gall&Hutto07.pdf>
- [32] A. R. Anderson, A reduction of deontic logic to alethic modal logic. *Mind*, 67, 100–103, 1958.
- [33] L. Kohlberg, Moral stages and moralization: The cognitive developmental approach. In T. Lickona (Ed.), *Moral development and behavior: theory, research, and social issues*, (pp. 31–53). New York: Holt, Rinehart & Winston, 1976.
- [34] M. Keller, M. Gummerum, X.-T. Wang, & S. Lindsay, Understanding perspectives and emotions in contract violation: Development of deontic and moral reasoning. *Child Development*, 75, 614–635, 2004.
- [35] H. Rakoczy, Play, games, and the development of collective intentionality. In M. Sabbagh & C. Kalish (Eds.), *New directions in child and adolescent development*, Special issue on “Conventionality,” to appear. Online: http://email.eva.mpg.de/~rakoczy/texte/rakoczy_2007_pretense_conventions.pdf
- [36] M. Gräfenhain, T. Behne, M. Carpenter & M. Tomasello, Young children's understanding of joint action and joint commitment in social play contexts. Poster presented at the *2007 Biennial Meeting of the Society for Research in Child Development (SRCD)*, Boston (MA), March 26–April 1, 2007.