

From collective intentionality to intentional collectives: An ontological perspective

Action editors: Luca Tummolini and Cristiano Castelfranchi

Emanuele Bottazzi, Carola Catenacci, Aldo Gangemi *, Jos Lehmann

*Laboratory for Applied Ontology, Institute for Cognitive Sciences and Technology, National Research Council (ISTC-CNR),
Via Nomentana 56, 00161 Roma, Italy*

Received 4 April 2005; accepted 7 November 2005
Available online 7 March 2006

Abstract

This paper presents a methodology for the analysis of the entities which the discourse on collective intentionality usually refers to. We aim, in particular, at characterizing the notion of intentional collective. Based on reviews of the relevant literature, we apply three formal-ontological tools of our choice (namely, DOLCE, DnS, and DDPO) to the treatment of the notions of collection, agent, plan and collective, all underlying the concept of intentional collective. We believe that the proposed approach offers several advantages, among which its explicitness, modularity and formality. This makes it particularly suitable for a founded specification of typologies of collections and collectives, hence for contributing to both philosophic and scientific research on these topics.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Collective intentionality; Formal ontology; Social entities; Collections; Collectives; Agents; Concepts; Roles; Plans

1. Introduction

During the last decade the problems whether there exists such a thing as collective intentionality and what relationship this holds with individual intentions have been hotly debated in Philosophy of Society, Theory of Action and Social Ontology (see, for instance, Bratman, 1992; Gilbert, 1992, 1996; Searle, 1990, 1995; Tuomela, 1995, 2003a, 2003b). Despite (often deep) differences between the various existing proposals, there are a number of general assumptions that form the common ground of this debate. It is, for instance, generally accepted that the social world is intrinsically plural. It involves, trivially, many individual agents and multiple interactions among them. Moreover, and less trivially, it involves a multiplicity of non-physical entities, which are produced by the agents themselves in

order to represent and manage the complexity of their own interactions.

Reasoning along these lines, in previous work we have distinguished at least two senses in which an entity can be said to be ‘social’ (cf. Masolo et al., 2004). In the first sense, an entity is social if it is an immaterial (more precisely, non-directly extended in space) product of a community. In this sense a social entity depends on agents who constitute, make use of, communicate about, and ‘recognize’ or ‘accept’ it by means of some sort of agreement. Here the term ‘social’ is roughly synonymous of ‘conventional’ and it refers to any aspect of reality that is ‘seen’ and understood in the terms set by a historically and culturally determined conceptualization. Examples of this sense are mathematical and scientific concepts, like triangle and quark, but also common-sense concepts, like sun, inasmuch as their ‘definition’ refers to a body of knowledge shared by a community. In the second and stronger sense, an entity is social if, in addition to having a conventional nature, its very constitution involves a network of relations

* Corresponding author. Tel.: +39 0644161535; fax: +39 0644161513.
E-mail address: a.gangemi@istc.cnr.it (A. Gangemi).

and interactions among social agents. Examples of this second sense are, e.g., euro, president, and consumer, as well as International Monetary Fund, Ethical Committee, and FIAT. This second sense of ‘social’ pivots on the idea that the social world is not only plural but also organized, i.e., it involves institutions and groups characterized by internal structures and roles. Typically, each individual agent simultaneously belongs to and acts within and across a multiplicity of such groups or collectives, ranging from family to professional, cultural, economical or political groups and organizations.¹ For an ontology of social reality, the challenge consists in providing an account of at least some of the basic structures which pervade such reality.

In this paper we follow exactly this lead and try to (re)construct some of the basic notions involved in the discourse on social reality by means of formal-ontological analysis. In order to do this, we reverse the terms of the classical question – what is collective intentionality? – and target, instead, the notion of intentional collective. On the one hand, we investigate and formalize the grounds based on which we define a set of items as a collection and collected items as members of a collection. On the other hand, we propose a way to relate collections and their members to intentional notions.

The main upshots of the presented investigation are the explicitness, modularity, and formality of the notions we introduce, as well as of the methodology we follow. Explicitness, modularity, and formality are key features for any conceptually structured vocabulary that is open to testing. Such a vocabulary will only be successful if the chunks of knowledge contained in the overall structure can be easily isolated, tested on their own, and updated. This is what we do in the following sections. We do not claim, however, that our reconstructions are the only possible ones. The emphasis is on the results that can be obtained through the methodology we propose. This mainly consists in making explicit what is often taken for granted, thus providing the means for well-founded discussions and documented (dis)agreements.

The rest of the paper is structured as follows. Section 2 informally presents our main theses. Section 3 describes the methodological backbone of our investigation. As a matter of fact, our reference scientific community is that of Applied Formal Ontology – a ‘joint venture’ of Artificial Intelligence and Philosophy, which provides formal accounts of large chunks of human knowledge for use in software applications. We begin with providing indications to the (unacquainted) reader about both the sense in which the term ‘ontology’ is used in our field and the specific ontologies adopted here to conduct our investigation: the Descriptive Ontology for Linguistic and Cognitive Engineering (Masolo, Gangemi, Guarino, Oltramari, & Schnei-

der, 2003), Descriptions and Situations (Gangemi & Mika, 2003), and some extensions of these two which have been developed in the context of the DOLCE plus DnS Plan Ontology (Gangemi, Catenacci, Lehmann, & Borgo, 2004). Section 4 provides a formal-ontological account of the notion of collection in terms of what defined in Section 3, along with a typology of collections. Section 5 provides a treatment of intentional collectives in terms of the formally specified notions of collection, agent, and plan. Finally, Section 6 draws some conclusions, reports some applications of the presented work and states the direction of our current and future research.

The reported work is part of our Laboratory’s research program dedicated to social ontologies. In particular, DOLCE, DnS, and DDPO have been developed in the framework of various EU-funded projects, to which our Laboratory has participated as a partner in research on knowledge-based systems.

2. The notions at stake

As stated in Section 1, our main objective is to provide a treatment of the notion of intentional collective and use it to present a general formal framework for an ontology of social reality. Consequently, the focus of the whole paper is on collections and collectives considered as social entities. In this section, we informally present our analysis of the notions at stake. Such analysis is largely based on hypotheses, which find formal and detailed specification throughout the rest of the paper.

According to our reconstruction, collections can be seen as social objects that depend on their members at a certain time. This entails, for instance, that a collection of books in a library remains the same entity even if some books are lost and others acquired over time. Collections depend also (specifically) on the role(s) played by their members. Consider, for example, the constellation of Orion. Should the role ‘being a member of Orion’ cease to exist, the relative constellation would disappear too.

Collections must be covered by at least one role; consider, for instance, a collection of (not further specified) bones, where ‘being a bone’ is the one and only role played by the members. Collections, however, can also be (and usually are) characterized by further roles. For instance, a collection of machines in a factory – where ‘being a piece of industrial machinery’ is the covering role – can be further characterized as a collection of cutting, pasting, etc., machines.

Collections, finally, are unified by ‘theory-like’ entities that we call descriptions, which contain and specify the covering or characterizing roles of the collection.

Collectives, in our proposal, are collections of agents which are unified by the kind of descriptions that we call plans. The members of a collective are ‘held together’ by one plan which specifies a goal and (one or more) covering or characterizing role(s). Hence, in our view both a group of people running towards a common shelter because of

¹ The term collective is used here in a sense that is reminiscent of Ludwik Fleck’s epistemological observations; Fleck’s exact terms, however, were thought-collective (Denkkollektiv) and thought-style (Denkstil); cf. (Cohen & Schnelle, 1986).

a sudden storm (Searle, 1995), and a pack of hunting wolves are to be considered as examples of collectives. Another example of how a collective is unified by a plan is the staff of a publishing house (including, e.g., a project manager, some assistants, some consultants, and several editors and authors) working at the production of a book.

3. Background concepts

Before getting to the heart of the matter, we introduce in this section the formal-ontological apparatus used in our treatment of collections, collectives and related typologies. We employ three ontologies: the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE, in the following) (Masolo et al., 2003), the Descriptions and Situations ontology (DnS) (Gangemi & Mika, 2003), and the DOLCE plus DnS Plan Ontology (DDPO) (Gangemi et al., 2004). We make use of first-order logic and introduce these types of formulas: (A#) for axioms, (D#) for definitions, (T#) for theorems, and (S#) for statements.

3.1. DOLCE

DOLCE is a foundational ontology of particular (as opposed to universal) entities. It is a formalized structure of very basic categories,² conceived as conceptual containers and applied in the automatic manipulation of knowledge. This is, roughly speaking, the sense in which the term ‘ontology’ is used in Artificial Intelligence (AI, in the following). Therefore, DOLCE, as most other ontologies in AI, makes no strong claims concerning the metaphysical implications of its categories. In other words, DOLCE does not (claim to) refer to ‘true’ reality.

The domain of quantification of DOLCE comprises possibilia, i.e., possible, not only actual particulars, so that we are allowed to talk of particulars that are postulated by existentially quantified variables, even if said particulars are not explicitly introduced in the model (Masolo et al., 2003).

DOLCE top-level includes the following mutually disjoint categories (printed in bold) and relations between such categories (underlined):

Endurants: Particulars that are directly localized in space (including Objects or Substances). Objects can be either physical or non-physical. For each event in which they participate, non-physical objects depend on physical objects. Social objects are a kind of non-physical objects.

Perdurants: Particulars that are directly localized in time (including Events, States or Processes).

Endurants and Perdurants are linked by the relation of participation. Endurants get their temporal location from the perdurants they participate in. Perdurant get their spatial location from the endurants participating in hem.

Qualities: Particulars that inhere in either Endurants (as Physical or Abstract Qualities) or in Perdurants (as Temporal Qualities). Qualities correspond to ‘individualized properties’, in the sense that they inhere only in a specific particular, e.g., ‘the color of this tennis court’, ‘the velocity of this service’, etc.

Abstracts: Particulars that are neither in time nor in space. For instance, the space of values that qualities can assume (e.g., a metric space), called a Quality Space, is an abstract. Each kind of Quality is associated to a Quality Space and different quality spaces may be associated to the same kind of Quality. It should be noted that in DOLCE, Space, and Time are specific quality spaces. Different endurants or perdurants can be spatio-temporally co-localized.

In this paper, we make use of some relations which are standard in formal ontology, e.g., part, which are characterized in DOLCE.³

3.2. DnS

The second formal-ontological tool that we employ in this paper is DnS (Gangemi et al., 2004; Gangemi & Mika, 2003; Masolo et al., 2004), an apparatus conceived with the purpose of extending other (possibly, but not exclusively, foundational) ontologies. For instance, when using DnS to extend DOLCE, what you get is DOLCE+. In DOLCE+, DOLCE plays the role of ground ontology, i.e., an ontology that is used to represent the entities in a domain, without considering their epistemological (constructive) status. Suppose that you have a ground ontology, say DOLCE itself, that contains predicates to represent entities involved in a chunk of social reality. Now, if you want to express the legal constraints imposed by norms and regulations on the domain of your ground ontology, you have to extend the latter and add to it a DnS description of social reality under a legal perspective. Such DnS description makes it possible to describe the ideal (legal) view on the behaviour of your social entities (a situation), according to a given legal system.⁴

The advantage of DnS resides here in the possibility of talking about the unity criteria of collectives, and of representing how collectives are related to other entities in complex situations.

In DnS, individual constraints and systems of constraints (theories) are reified, thereby becoming entities in the same domain of quantification of the entities from

² Most of the ontologies mentioned in the following are available in various formal languages and formats on <http://www.loa-cnr.it>.

³ See <http://dolce.semanticweb.org> for more details. The names of the DOLCE relations as introduced in the vocabulary of DnS (Gangemi et al., 2004), however, differ from those indicated in the official DOLCE document (Masolo et al., 2003), mainly because some computational logics require “unique names” for predicates of different arity, and some applications require “self-explaining” names. We maintain the DOLCE names used in DnS for this work.

⁴ An extension of DOLCE which does not use DnS, but introduces descriptions in order to treat social entities is presented in Masolo et al. (2004).

the ground ontology. Reified constraints and theories are classified as **social objects**, which hold various properties like, for instance, a(n indirect) location in space and time.

In more detail, DnS is based on a fundamental distinction between **descriptions** (for instance, in the legal domain, legal descriptions, or **conceptualizations**, which encompass laws, norms, regulations, crime types, etc.) and **situations** (again, in the legal domain, **legal facts** or **cases**, which encompass legal states of affairs, non-legal states of affairs that are relevant to the Law, and purely juridical states of affairs).

DnS basic predicates and axioms are the following:

According to the first sense of ‘social entity’ presented in Section 1, we can minimally characterize a **Social Object** as a non-physical object that requires (i) **communication events** for any event in its lifecycle; **and** (ii) **physical agents** that participate in those same events:⁵

- (A1) $\text{SocialObject}(x) \rightarrow \text{NonPhysicalObject}(x) \wedge \forall(y, t) ((\text{Perdurant}(y) \wedge \text{ParticipantIn}(x, y, t)) \rightarrow \exists(z, w) (\text{CommunicationEvent}(z) \wedge \text{PartOf}(z, y) \wedge \text{AgentivePhysicalObject}(w) \wedge \text{ParticipantIn}(x, z, t) \wedge \text{ParticipantIn}(w, z, t)))$
- (A2) $\text{CommunicationEvent}(x) \rightarrow \text{Event}(x)$

Like physical objects, social ones have a lifecycle, can have parts, etc. Unlike physical objects, non-physical ones always require some agentive physical object in the events they participate in (for a discussion of the notion of ‘agentivity’, see Section 5.2 below).

A **Description** is a social object which represents a conceptualization, hence it requires some agent in the events it participates in, and it is communicable (Masolo et al., 2004). Examples of descriptions are regulations, plans, laws, diagnoses, projects, scripts, techniques, etc.:

- (A3) $\text{Description}(x) \rightarrow \text{NonAgentiveSocialObject}(x)$

Descriptions have typical components, called *concepts* (see below). Concept types can vary according to the ground ontology that is taken into account. The version of DnS used in this paper takes DOLCE as its ground ontology.

A **Situation** is a Particular which represents a state of affairs, under the assumption that its components ‘carve up’ a view (a **setting**) on the domain of an ontology by virtue of a description. A situation aims at representing the referent of a ‘cognitive disposition’ towards a world, thus reflecting the willingness, expectation, desire, belief, etc. to carve up that world in a certain way. As a consequence, a situation has to **satisfy** a description (see below). Examples of situations, related to the examples of descriptions

above, are: facts, plan executions, legal cases, diagnostic cases, attempted projects, performances, technical actions, etc.:

- (D1) $\text{Situation}(x) =_{\text{df}} \exists y, z (\text{Satisfies}(x, y) \wedge \text{SettingFor}(x, z))$

The **satisfies** relation holds between situations and descriptions, and implies that at least some concept in a description must classify at least some particular in the situation setting:

- (A4) $\text{Satisfies}(x, y) \rightarrow \text{Situation}(x) \wedge \text{Description}(y)$
- (A5) $\text{Satisfies}(x, y) \rightarrow \exists z (\text{Concept}(z) \wedge \text{Uses}(y, z)) \wedge \exists w, t (\text{SettingFor}(x, w) \wedge \text{Classifies}(z, w, t))$

The **setting** relation holds between situations and particulars from the ground ontology. At least one perdurant must exist in the situation setting:

- (A6) $\text{SettingFor}(x, y) \rightarrow \text{Situation}(x) \wedge \text{Particular}(y) \wedge \neg \text{Situation}(y)$
- (A7) $\text{SettingFor}(x, y) \rightarrow \exists z (\text{Perdurant}(z) \wedge \text{SettingFor}(x, z))$

The **time and space of a situation** are the time and space of the particulars in the setting:⁶

- (A8) $\forall x, y, t, t' ((\text{Perdurant}(x) \wedge \text{TimeInterval}(t) \wedge \text{TimeInterval}(t') \wedge \text{TemporalLocation}(x, t) \wedge \text{TemporalLocation}(y, t') \wedge \text{SettingFor}(y, x)) \leftrightarrow \text{Part}(t', t))$
- (A9) $\forall x, y, r, r' ((\text{Endurant}(x) \wedge \text{SpaceRegion}(r) \wedge \text{SpaceRegion}(r') \wedge \text{SpatialLocation}(x, r) \wedge \text{SpatialLocation}(y, r') \wedge \text{SettingFor}(y, x)) \leftrightarrow \text{Part}(r', r))$

(A8) and (A9) state that a situation has a temporal – respectively, spatial – location that is the mereological sum of the locations of the particulars in the setting. For example, the time of *World War II* might span from the German invasion of Poland in 1939 to the Yalta conference in 1945; its space might include most of the Earth surface. Hence, the setting relation is not temporalized, because the time of $\text{Setting}(x, y)$ can be inferred from the previous axioms.

A **Concept**, like a description, is a social object, which is **defined by** a description. Once defined, a concept can be **used in** other descriptions. The **classifies** relation relates concepts to particulars (and possibly even concepts to concepts) at some time. There are several kinds of concepts reified in DnS, the primary ones (**role**,⁷ **course**, and **parameter**) being distinguished by the categories of particulars they classify in DOLCE:

⁵ An axiomatization for communication theories is under constructions. An ontology of information objects and their semiotic framework is introduced in the so-called DDIO (DOLCE + DnS + Information Ontology), cf. the DOLCE site: <http://dolce.semanticweb.org>.

⁶ All ‘t’ variables in the formulas denote time intervals.

⁷ In this paper a role is simply a concept that classifies an endurant. This definition is different from the one provided in Masolo et al. (2004), which relies crucially on the properties of anti-rigidity and foundation.

- (A10) $\text{Uses}(x, y) \rightarrow \text{Concept}(x) \wedge \text{Description}(y)$
 (A11) $\text{Defines}(x, y) \rightarrow \text{Uses}(x, y)$
 (A12) $\text{Classifies}(x, y, t) \rightarrow \text{Concept}(x) \wedge \text{Particular}(y) \wedge \text{TimeInterval}(t)$
 (A13) $\text{Concept}(x) \rightarrow \text{NonAgentiveSocialObject}(x) \wedge \exists y (\text{Defines}(y, x) \wedge \text{Description}(y))$
 (D2) $\text{Role}(x) =_{\text{df}} \text{Concept}(x) \wedge \exists y, t (\text{Classifies}(x, y, t) \wedge \forall y, t (\text{Classifies}(x, y, t) \rightarrow \text{Endurant}(y)))$
 (D3) $\text{Course}(x) =_{\text{df}} \text{Concept}(x) \wedge \exists y, t (\text{Classifies}(x, y, t) \wedge \forall y, t (\text{Classifies}(x, y, t) \rightarrow \text{Perdurant}(y)))$
 (D4) $\text{Parameter}(x) =_{\text{df}} \text{Concept}(x) \wedge \exists y, t (\text{Classifies}(x, y, t) \wedge \forall y (\text{Classifies}(x, y, t) \rightarrow \text{Region}(y)))$

Examples of roles are: manager, student, assistant, actuator, toxic agent, etc. Roles can be **specialized** by other roles, e.g., president of the Italian republic specializes president of the republic:

- (A14) $\text{Specializes}(x, y) \rightarrow \text{Role}(x) \wedge \text{Role}(y)$

From (A12) and (A14), we can derive that:

- (T1) $\forall x, y, t \exists z ((\text{Classifies}(x, y, t) \wedge \text{Specializes}(x, z) \wedge x \neq z) \rightarrow \text{Classifies}(z, y, t))$

Since descriptions and concepts are (social) objects, hence endurants, they can be classified by a role in another description. This recursivity allows to manage meta-level descriptions in DnS (e.g., a norm for enforcing norms will define a role that can classify the enforced norm), without the need for a second-order axiom.

Courses are concepts that classify perdurants; examples of courses are routes (e.g., a travel itinerary), pathways (e.g., a biochemical pathway), tasks (e.g., ‘press button 2’), etc.

Parameters classify regions within quality spaces (e.g., attributes of activities or objects); examples of parameters are: speed limits, allowed colors (e.g., for a certain book cover), temporal constraints, etc.

Parameters are related to roles or courses by a **requisite for** relation, expressing the kind of requisites that, in given descriptions, particulars which are classified by said roles or courses should have. The relation is temporalized to suggest that a description can preserve its identity against changes of structuring among components (though there can be mandatory structures for description identity):

- (A15) $\text{RequisiteFor}(x, y, t) \rightarrow \text{Parameter}(x) \wedge \text{TimeInterval}(t) \wedge (\text{Role}(y) \vee \text{Course}(y))$

Requisites are constraints over the values of the qualities of particulars. When a situation satisfies a description that uses parameters, endurants, and perdurants that constitute the situation must have attributes that range between the boundaries stated by said parameters (in terms of DOLCE, particulars must have qualities that are mapped onto certain value ranges of regions). For example, a speed limit

of 50 kmph can be a requisite for a driving task; a satisfying situation will have any speed of driving (e.g., in an instance of driving in Rome by car) to be less or equal to 50 kmph.

The classifies relation is specialized by three subrelations: **plays**, **sequences**, and **value for**, which apply to three different categories in DOLCE (Endurant, Perdurant, and Region, from (D2-4)).⁸

- (D5) $\text{Plays}(y, x, t) =_{\text{df}} \text{Role}(x) \wedge \text{Classifies}(x, y, t)$

For instance, Ciampi is playing the role ‘president of the Italian Republic’ in 2005.

- (D6) $\text{Sequences}(x, y, t) =_{\text{df}} \text{Course}(x) \wedge \text{Classifies}(x, y, t)$

For instance, the course ‘voting in parliament’ sequences a perdurant like raising one hand (under certain circumstances).

- (D7) $\text{ValueFor}(x, y, t) =_{\text{df}} \text{Parameter}(x) \wedge \text{Classifies}(x, y, t)$

For instance, the parameter ‘speed limit of 50 kmph’ is a value for the region of speed.

As we shall see, these three relations allow us to define our notion of plan, which is essential in order to characterize the notion of collective.

3.3. DDPO: DOLCE + DnS plan ontology

DDPO specializes the concepts and relations defined in DOLCE and in DnS. Like the latter, DDPO has a very liberal domain, which includes physical and non-physical objects (social entities, mental objects and states, conceptualizations, information objects, constraints), events, states, regions, qualities, and ‘constructivist’ situations. The intended use of DDPO is to specify plans (which are a type of descriptions) at a very abstract level and independently from existing resources. DDPO allows also for the specification of tasks, namely the types of actions of which plans are made of (Gangemi et al., 2004).

The main predicates and axioms of DDPO will be given in Section 5.3, after a discussion of the notion of agentivity. We now specialize the DOLCE and DnS notions in order to introduce collections.

4. Collections

In ontology, collections are usually treated as plural entities, a controversial concept, which is subject to hot debates in the field (cf. Simons, 1987). We have chosen to treat collections in terms of the constructive boundaries of those plural entities that form themselves a whole. More

⁸ Only three categories from DOLCE have been assigned a concept type at the descriptive layer, because the resulting pattern is simpler and there is no loss of relevant knowledge, at least in applications developed until now.

intuitively, we define collections as emerging out of entities that, while retaining their identity, unity, and physical separation, are ‘kept together’ in order to form a new entity. This notion is analogous to that of discrete integral whole in Abelard (King, 2004). Our proposal, though, stresses the nature of collections as social objects and explicitly indicates what keeps a collection together: its unity criteria (at least one).

Before we introduce the formulas that define our notion of collection, we want to make a number of disclaimers. For the unacquainted reader these may provide a gentle introduction to our axiomatization. In particular, we hope that the discussion of notions that are intuitively very similar to the notion of collection – like for instance the notion of set – may provide insights in the peculiarities of each notion and in their mutual differences. On the other hand, our disclaimers may help the expert reader in positioning our proposal.

Firstly, we are clearly touching on a very difficult topic here, which is covered by a heterogeneous literature ranging from Metaphysics (Cocchiarella, 2005; King, 2004) to Logic (Russell & Whitehead, 1910; Zeman, 1982), Mathematics (Dauben, 1979; Dugac, 1976) and to Linguistics and Formal Semantics (Link, 1983; Marcus, 1993). Nevertheless, our interest is not on comparing our contribution with all existing proposals but on providing an example of how to formalize the notion of collective within a foundational ontology. Our theoretical horizon is made of DOLCE and DnS or other formal structures of a similar nature. None of these offers enough support as yet for treating collections (cf. the notion of unitary collection in Masolo et al., 2003). Therefore, our effort is mainly directed at extending DOLCE and DnS with a general characterization of collections.

Secondly, the result of this extension is not another axiomatization of sets. Collections are very different from sets in the sense of Set Theory (i.e., either Zermelo–Fraenkel plus the axiom of Choice or von Neumann–Bernays–Gödel; see Devlin, 1993; von Neumann, 1967). This is true for the following reasons: (1) a set is uniquely determined by its members, i.e., it changes when its members or its cardinality change (axiom of extension), while a collection is not, unless explicitly specified; and (2) any two sets can be summed forming a union (axiom of union), while this is not tenable for any two collections; and (3) sets do not need an identity criterion for members (axiom of specification does not apply to all sets), while collections do (there is at least one property P that is true for all members); and (4) sets can be empty or singletons, but no empty or singleton collections are allowed (at least in the axiomatization presented here); and (5) (hyper)sets can be members of themselves (anti-foundation axiom), while collections cannot; and (6) sets are abstract, having no space or time, while collections (indirectly) exist in time, and are localized.

Thirdly, our notion of collection is naturalized. By ‘naturalized’ we mean that the conceptualization of entities we commit to is embodied in cognitive agents located in

space–time, and it is due to biological, social, and cognitive evolution.⁹ An ontology of naturalized entities is quite different from one of abstract, aeternal ones. In our second disclaimer we state that collections are no sets. Could one suggest that collections, as we mean them, are nothing more than naturalized sets?

Fourthly, a collection is not a naturalized set, because (i) we assume that members have an explicit identity criterion; and (ii) we accept substitution of members of a collection while preserving its identity; and (iii) we accept changes in the cardinality of a collection while preserving its identity.

Fifthly, a collection is not a naturalized proper class, since a collection depends on its (at least two) members, while a proper class can be empty or singleton. If we assume that conceptualizations are embodied in cognitive agents, then our ‘collections’ can be seen as naturalizations of non-empty proper classes with cardinality higher than 2, and (at least one) basic properties for membership. This seems to capture the common sense intuition underlying groups, teams, collections, collectives, associations, etc.

By the above five points we hope to have sufficiently insulated our proposal from the long-standing debate (in Philosophy, Logic, and Mathematics) about the nature of sets and classes as abstract entities, their relation to so-called universals, etc.

Sixthly, overall our approach is constructivist. A collection depends on one or more social objects that provide a unity criterion for it. When a (complex of) social object(s) applies (in explicit, varied ways) to a plurality of entities, a collection appears. When such complex ceases to be conceived by any agent, or stops being applied to a plurality, then the collection dies. General criteria on the lifecycle of social objects are given in Masolo et al. (2004).

In order to provide a basis for any naturalized collection, we propose a formal version of the containment image schema informally introduced in cognitive semantics (Johnson, 1987; Lakoff & Nunez, 2000), and use it to account for the foundational intuition of a collection.

Since collections are considered here as cognitive or social objects, but they also depend on their members, their space–time behavior is peculiar. Collections can participate in actions or processes either ‘on a member basis’ or ‘on a whole basis’. For example, some cows step on a guy, and the guy recognizes a moving herd ‘stepping on him’: the herd steps on the guy ‘on a member basis’. An opposite example: in 1914, a Serbian terrorist assaulted and killed Archduke Franz Ferdinand, and Austria found Serbia (‘collectively’) guilty. In this case, Serbs were judged to have killed ‘on a whole basis’, and the collective (moral and political) responsibility was distributed over all members.¹⁰

⁹ We do not enter here the complex debate on the primacy of biology vs. society.

¹⁰ ‘Collective responsibility’ proper – i.e., the controversial issue of what e.g., Feinberg has called collective but not distributive group fault (Feinberg, 1968) – is a topic the discussion of which we postpone to future work in the field of legal ontologies.

Where, however, are the herd and the Serbs spatio-temporally located? Following what we have axiomatized for situations, we propose here that the space–time of a collection is the maximal space–time of the members when they are classified by some selected role(s). In this way, acting on a member basis holds whenever a collection participates on a whole basis. In other words, the space–time of collections is equivalent to the space–time of the members when they are classified by certain roles.

4.1. Definition of collection and membership

While we could talk in general of collections of any kind of particulars (events, objects, abstracts, etc.), here we focus on collections of endurants and, therefore, on the concepts that classify them, i.e., roles. **Collections** are defined here as social objects that have at least two endurants as their members. Said members must be classified by at least one role that gives the unity criterion to the collection. Members can change during the life of a same collection without affecting the identity criterion for the collection. We might choose a constitution relation to stress the emerging nature of collections. Intuitively, by constitution DOLCE means a relation between entities belonging to different layers of a world (typically, biological entities emerging out of chemical ones), but DOLCE limits constitution within solely the physical or the non-physical maximal layers. For this reason, we take here a more generic intuition for the emergence of collections, which is based on a very generic containee role, defined in the containment schema, which classifies each member of a collection, which on its turn plays the container role from the same schema. **CONTAINMENT** is a cognitive schema (Johnson, 1987; Lakoff & Nunez, 2000), which we represent as a description.¹¹

- (S1) Description(ContainmentSchema)
- (S2) Role(Container)
- (S3) Role(Containee)
- (S4) Defines(ContainmentSchema, Container)
- (S5) Defines(ContainmentSchema, Containee)

Based on the **CONTAINMENT** schema, we introduce a relation holding between the particulars that are classified by the **CONTAINMENT** roles:

- (D8) $\text{IsIn}(x, y, t) \rightarrow \text{Particular}(x) \wedge \text{Particular}(y) \wedge \text{Classifies}(\text{Containee}, x, t) \wedge \text{Classifies}(\text{Container}, y, t)$

IsIn is very general, and catches the intuition of whatever entity being in another, on purely schematic grounds (as a matter of fact, it is usable also for metaphorical mappings). It is antireflexive and antisymmetric:

- (A16) $\neg \text{IsIn}(x, x, t)$
- (A17) $\text{IsIn}(x, y, t) \rightarrow \neg \text{IsIn}(y, x, t)$

Notice that $\text{IsIn}(x, y, t)$ is introduced here as a primitive, because a definition would require the specification of the conditions under which a cognitive agent uses it; those conditions cannot be given unless we take into account neurobiological and cognitive developmental factors that are beyond the scope of this paper.

We assume that the **collection** intuition is a special case of **CONTAINMENT**. In particular, we hold here that a collection has only endurants as members, and that it has always at least two members. The identity of a collection derives from the role that classifies its members. If this role ceases to exist, also the collection ceases to exist:

- (D9) $\text{Collection}(x) =_{\text{df}} \text{SocialObject}(x) \wedge \text{Endurant}(y) \wedge \forall t (\text{PresentAt}(x, t) \rightarrow \exists y, z (\text{IsIn}(y, x, t) \wedge \text{IsIn}(z, x, t) \wedge y \neq z))$

From (D8) and (D9), it holds that each collection has at least one role classifying its members (the containee role):

- (T2) $\text{Collection}(x) \rightarrow \exists y (\text{Role}(y) \wedge \forall z, t (\text{IsIn}(z, x, t) \rightarrow \text{Classifies}(y, z, t)))$

From (D8) and (D9), it also holds that each collection is classified by the container role, and each of its members is classified by the containee role:

- (T3) $\text{Collection}(x) \rightarrow \forall t (\text{Classifies}(\text{Container}, x, t))$
- (T4) $\text{Collection}(x) \rightarrow \forall y, t (\text{IsIn}(y, x, t) \rightarrow \text{Classifies}(\text{Containee}, y, t))$

A **membership** relation can now be defined on collections:

- (D10) $\text{Membership}(x, y, t) =_{\text{df}} \text{IsIn}(x, y, t) \wedge \text{Endurant}(x) \wedge \text{Collection}(y) \wedge \exists z (\text{Role}(z) \wedge \text{Classifies}(z, x, t))$

The membership relation is antireflexive and antisymmetric like IsIn .

Two or more collections can be extensionally equivalent and still not be the same collection, since each collection

¹¹ Currently, cognitive schemata are only informally represented in the literature. A **COLLECTION** schema, for instance, is actually included among the most important schemata listed in Johnson (1987); since, however, Johnson does not provide any analysis of this notion, we have chosen to reconstruct it on the basis of the more general and better understood **CONTAINMENT** schema. By ‘**CONTAINMENT**’ we mean here a formal schema. Therefore, our notion of containment may be applied to collections of physical as well as non-physical objects. Our Laboratory has an ongoing research on the formal representation of cognitive schemata. The choice of representing schemata as descriptions lies in the similarity between the schema metaphor in cognition (simultaneous activation of multiple intentional systems, through their physiological counterparts), and in the reification semantics of DnS (relational patterns within a reified theory). However, besides the technically seamless formalization of cognitive schemata as very general ontological descriptions, experimental validation is still unavailable.

needs a unifying description (cf. A19–21 in the following) which provides its intensional identity criterion:

$$(D11) \text{ ExtensionallyEquivalent}(x, y, t) =_{df} \text{Collection}(x) \wedge \text{Collection}(y) \wedge \forall z, t (\text{Membership}(z, x, t) \leftrightarrow \text{Membership}(z, y, t))$$

The role shared by members has a **covering** relation towards the collection. The definiens of such relation is a specialization of the theorem (T2):

$$(D12) \text{Covers}(x, y) =_{df} \text{Role}(x) \wedge \text{Collection}(y) \wedge \forall z, t (\text{Membership}(z, y, t) \rightarrow \text{Classifies}(x, z, t))$$

From (D8) and (D12), it holds that collections are covered by the container role:

$$(T5) \text{Collection}(x) \rightarrow \text{Covers}(\text{Containeer}, x)$$

The containment schema is used to define the membership relation, therefore it makes (D10) true for any arbitrary collection. On the other hand, (D9) allows for the existence of multiple roles covering a collection, as we exemplify in the following section.

4.2. Typology of collections

Several typologies of collections can be built, based for instance on member types, covering role types, etc. Since our main scope is on collectives (which, as shown later, we reconstruct as a special kind of collections), we limit collection types to one basic typology.

A **simple collection** – i.e., for instance, a collection of saxophones, or a mass of lymphocytes¹² – is a collection having only covering roles (for the definition of Characterizes, see (D18) below):

$$(D13) \text{SimpleCollection}(x) =_{df} \text{Collection}(x) \wedge \exists y (\text{Covers}(y, x) \wedge \neg \exists z (\text{Characterizes}(z, x)))$$

A **maximally generic collection** (for instance, a collection of objects selected at random) is a (simple) collection of some conceivable objects, without any further covering:

$$(D14) \text{MaximallyGenericCollection}(x) =_{df} \text{SimpleCollection}(x) \wedge \neg \exists y (\text{Covers}(y, x) \wedge y \neq \text{Containeer})$$

¹² Including masses among collections is not straightforward, since most traditional views take masses as continuous. However, the traditional criterion for membership identification, i.e., discreteness, is (implicitly) based on the perceivability of members. Since members can be either directly or instrumentally identifiable, or even inferrable, we need to enlarge the range of possible collections. On the other hand, a mass cannot be considered a collection when no member is identifiable or epistemically relevant. For instance, if we take a piece of gold without considering its constituent atoms, it cannot be a collection, it is just ‘stuff’.

Any simple collection that is not a maximally generic one needs either to specialize the role of containeer by means of further axioms, or to be covered by additional roles. For example, collections of dinosaur bones can be defined as follows:

$$(S6) \text{Specializes}(\text{Containeer\#DB}, \text{Containeer})$$

$$(A18) \forall x, t (\text{Classifies}(\text{Containeer\#DB}, x, t) \rightarrow \text{Dinosaur Bone}(x))$$

$$(D15) \text{DinosaurBonesCollection}(x) =_{df} \text{Collection}(x) \wedge \text{Covers}(\text{Containeer\#DB}, x)$$

The containee role can be specialized to limit its classification to a specific taxonomic level of classified endurants. Other collections need a covering by additional roles, for example, a collection of drugs is also covered by the medicament role:

$$(D16) \text{DrugsCollection}(x) =_{df} \text{Collection}(x) \wedge \text{Covers}(\text{Medicament}, x)$$

Another type of simple collections are **parametrized collections**, whose members must have a quality constrained by some parameter that is a requisite of their covering role(s):

$$(D17) \text{ParametrizedCollection}(x) =_{df} \text{SimpleCollection}(x) \wedge \exists y, z, v, t (\text{Covers}(y, x) \wedge \text{RequisiteFor}(z, y, t) \wedge \text{ValueFor}(v, z) \wedge \forall w (\text{Membership}(w, x, t) \rightarrow \exists q (\text{InheresIn}(q, w, t) \wedge \text{Q-Location}(q, v))))^{13}$$

For example, a crowd of people has members that have spatial positions in a range that makes them proximal (a condition traditionally used to distinguish so-called aggregates (King, 2004)).¹⁴

Organized collections introduce a different unity criterion for collections. They can be conceived as characterized by further roles played by some (or all) members of the collection, and related among them through the social objects (descriptions, collections) that either use or are covered by them:

$$(D18) \text{Characterizes}(x, y) =_{df} \text{Role}(x) \wedge \text{Collection}(y) \wedge \exists z, w, y', x', t ((\text{Description}(y') \vee \text{Collection}(y')) \wedge \text{Role}(x') \wedge z \neq w \wedge x \neq x' \wedge \text{Membership}(z, y) \wedge \text{Membership}(w, y) \wedge (\text{Uses}(y', x) \vee \text{Covers}(x, y')) \wedge (\text{Uses}(y', x') \vee \text{Covers}(x', y')) \wedge \text{Classifies}(x, z, t) \wedge \text{Classifies}(x, w, t))$$

$$(T6) \text{Characterizes}(x, y) \rightarrow \exists z (\text{Role}(z) \wedge x \neq z \wedge \text{Characterizes}(z, y))$$

$$(D19) \text{OrganizedCollection}(x) =_{df} \text{Collection}(x) \wedge \exists y, z (\text{Characterizes}(y, x) \wedge \text{Characterizes}(z, x) \wedge y \neq z)$$

¹³ Cf. Section 3.1.

¹⁴ On the other hand, if positions are reciprocally relevant (as, for instance, in a living chess setting) according to multiple roles defined by some plan or design, the collection becomes organized.

We claim that collections specifically depend on at least one description:

$$(A19) \text{ Collection}(x) \rightarrow \exists y \\ (\text{Description}(y) \wedge \text{SpecificallyDependsOn}(x, y))$$

We can therefore build a new relation of **unification** between collections and the descriptions on which they depend. Unification is axiomatized by means of sufficient conditions (A21–22), and is not temporalized, since changing the description (differently from changing some members) creates a new collection:

$$(A20) \text{ Unifies}(x, y) \rightarrow \text{Description}(x) \wedge \text{Collection}(y) \\ (A21) \text{ Covers}(x, y) \rightarrow \exists z \\ (\text{Description}(z) \wedge \text{Defines}(z, x) \wedge \text{Unifies}(z, y)) \\ (A22) \text{ Characterizes}(x, y) \rightarrow \exists z \\ (\text{Description}(z) \wedge \text{Defines}(z, x) \wedge \text{Unifies}(z, y))$$

From (A13), (D2), (D9), (D12), and (A20) we can derive that a collection must be unified by at least one description, which provides to said collection its unity criterion:

$$(T7) \text{ Collection}(x) \rightarrow \exists y (\text{Description}(y) \wedge \text{Unifies}(y, x))$$

We can imagine roles that are used by, or that cover, more than one description or collection.¹⁵ In other words, characterizing roles can be related among them through some composition (or bundle) of descriptions or collections.

With respect to dynamics, the identity of collections can survive change of some members. This behaviour ‘corresponds to’ the extensional/intensional nature of classes. For instance, if my collection of saxophones loses one member, it is still my collection of saxophones, because it respects the intensional criterion of being saxophones (possibly of a certain type).

On the other hand, collections can change identity while preserving the same members. What about, for instance, your collection of saxophones, which is covered or characterized by the same roles as mine? They result to be identical under the sole intensional criterion suggested by the solution above. In this case, we must postulate additional constraints, like exact restrictions on role playing. For instance, my/your collection will require that a role owner be played only by me/you. Consider that these additional constraints do not have an impact on the organization of a collection: for instance, owner is not played by any member of the collection. Such solution is similar to that applied to industrial products: serial number, quality check date, registration date, etc., which are actually regions (values) of the overall collection (not of members).

Notice, however, that if I state that all and only the saxophones of a certain brand in a certain serial-number range can be members of my collection, and if one of the saxophones gets lost, then my collection ceases to exist. This happens because I am using the identity criterion of the members as unity criterion for the collection: the collection is ‘maximally specified’.

Let us recapitulate here the network of relations among the main notions introduced thus far. A concept is defined by a description and can classify some particular (a role being a concept that classifies only endurants). A collection, on the other hand, is not directly defined by a description, and cannot classify any particular, but has members that are classified by at least one and the same role. Collections are social individuals, while concepts are not. We may say that collections are emergent social individuals because, unlike concepts, they do not need to be directly defined by a description.

The detailed distinctions on collections exemplified above are not only for the sake of formal-ontological rigor. They also drop a hint on the method by which one can discover why different, apparently similar theories about collectives and collective intentionality can actually result incompatible.

5. Towards collectives

Now that we have an explicit notion of collection, we want to use it for making a step towards the definition of collectives as social entities. Despite the fact that traditional literature on collective intentionality is usually not committed to this kind of entities – preferring, instead, the notion of ‘social group’ – our move finds some theoretical support in the literature from various other fields. For instance, Sartre considers collectives as intermediate entities between collections and social groups (Sartre, 1982). A similar view has been entertained, too, both in sociology (French, 1984) and linguistics (Borshev & Partee, 2001).

We consider collectives to be something more than collections, since they are composed by agents who participate in the activities specified by a plan, but something less than social groups, because, for example, they can exist even in absence of mutual beliefs or joint intentions among agents, which are requirements for the entities treated by the classical literature on collective intentionality (Bratman, 1992; Gilbert, 1992; Searle, 1990, 1995; Tuomela, 2003b). Moreover, the classical literature is also strongly committed to the notion of ‘we-mode’: in a social group, individual agents think or speak about themselves in terms of ‘we’. Some of the authors (Gilbert, 1992; Searle, 1990, 1995) consider this ‘we’ as a primitive; others (Bratman, 1992; Castelfranchi, 2003; Tuomela, 2003b), on the contrary, believe in the reducibility of this we-modality to a complex composition of I-modalities. Our position, however, is that this sort of psychological notions are not strictly necessary

¹⁵ Unifying descriptions of a collection can be: (a) those which define covering or characterizing roles; and (b) those which use said roles (defined elsewhere), but whose unifying function is explicitly stated.

in order to define collectives in general, although they can be introduced to propose a particular interpretation of collectives for some purpose.

Our definition of collectives is built around a ‘descriptive’ interpretation of the notions of intentionality (Section 5.1), agent (Section 5.2), and – especially – plan (Section 5.3). This commitment to the notion of plan gives our approach a teleological flavour. In this sense, we comply with some of the classical work in Philosophy and AI (Bratman, 1992; Cohen & Levesque, 1990), which assumes that the kind of rationality usually guiding actors in a society is a means–end rationality, and that the latter plays a crucial role in many contexts of (contemporary) productive societies.¹⁶ There are, of course, other models of rational interactions in a society, as pointed out in Weber (1968), and recognized at least since Aristotle’s description of ‘akrasia’ (Rorty, 1986). In this phase of our investigation we concentrate on the means–end type of rationality, but in Section 6 we suggest how plans in social interaction are usually intertwined with other, possibly conflicting, plans, as well as with other descriptions. Next phases of our research will concentrate on such ‘bundles’ of descriptions. Notice that if one wants to axiomatize a notion of collective based on the primitives introduced so far, but without making it plan-based, this is possible by specializing a different notion of collective out of that of collection.

5.1. Our approach to intentionality

Intentionality is still a debated notion in philosophy. From a historical point of view, the first modern account of intentionality is due to Brentano (1924), who gave new life and meaning to the medieval notion of intentio¹⁷ and used it to distinguish between physical and psychical phenomena. Following Searle (as representative, on this topic, of the received view in Philosophy of Mind), we take intentionality to be “that feature of [mental] representations by which they are about something or directed at something” (Searle, 1995). Intentionality is thus the requisite for entertaining intentional mental states: beliefs, desires, fears, or making hypotheses are different types of intentional states, but they all share the feature of being about something. As noted by Searle (1983), ‘intentionality’ in this wider, philosophical sense is not to be confused with what is ordinarily called an ‘intention’. The German language is less ambiguous in this respect, since it distinguishes between Intentionalität and Absicht, the latter corresponding, for instance, to what expressed by a sentence like “I intend to go to the movies tonight”.¹⁸ This

ordinary use of ‘intention’, on the other hand, seems closer to that typically made in Belief-Desire-Intention approaches (BDI in the following), where ‘intentions’ (as representations of the goals an agent is committed to achieve) are considered to be the third type of mental states which, together with beliefs and desires, plays a crucial role in the modeling of agent behavior (Ferrario & Oltramari, 2004).

Providing an ontology of mental states, however, is definitely beyond the aim of this paper. Although a first move towards such objective has actually been done in our laboratory (Ferrario & Oltramari, 2004), and although our ontology of descriptions should ideally ‘correspond’ to an ontology of mind, there is still not enough agreement either in Cognitive Sciences or Philosophy of Mind on the nature of mental entities, and the currently available primitives are not sufficient for developing typologies of mental states and their mutual relations – hence, to handle intentionality and intentions in formal ontological terms. In the DOLCE+ framework, we will consider the descriptive equivalent of a type of mental states which undoubtedly include intentions, namely plans. Due to the above-mentioned lack of a sufficiently developed ontology of mind, however, we cannot provide a one-to-one correspondence between the two ontologies as yet. Therefore, we will characterize intentional agents and collectives within our ontology of descriptions alone.

5.2. Agents

Our characterization of agentivity, however, takes into account some of the main features attributed to (intentional) agents in the philosophical, AI and (mostly) BDI literature. As reported in Ferrario and Oltramari (2004), agents are generally characterized by their being oriented at producing some results; they perceive their environment and act on it in order to achieve their goals. In particular, goal-directed agents are “endowed with [...] internal anticipatory and regulatory representations of action results” (Castelfranchi, 1998). From the perspective of DOLCE+, this supports, and is consistent with, the assumption that an agentive physical object is able to conceive descriptions.¹⁹ On the other hand, an agent is considered to be intentional (or rational (Wooldridge, 2000)) when not only it builds a (mental) representation of the goal, but also a representation of the action necessary to its achievement, and of the resulting consequences. Finally, another central distinctive features of intentional agents is considered to be their ability for social interaction, i.e., the fact that they act in and on an environment where external stimuli are originated also (and mostly) by other agents. This picture (and, in general, the close link it establishes between intentionality,

¹⁶ For an account of rationality and social reality, see Searle (2001), and also Pettit (2003).

¹⁷ For a treatment of intentionality in medieval philosophy cf., for instance, Perler (2003).

¹⁸ For an introduction to the different senses of ‘intentionality’, cf. also Jacob (2003).

¹⁹ We assume that descriptions can be expressed either by natural-language propositions or by other (not necessarily language-like) encodings.

social dimension, and planning activities) seems to leave room for a distinction between, and characterization of, two levels of agentivity.

As stated in Section 3.2, in DOLCE and DnS descriptions (like all non-physical objects) require some agentive physical object in the events they participate in. We have further characterized the relationship between a description and an agent (see below) in the following axiom:

$$(A23) \text{ Conceives}(x, y, t) \rightarrow ((\text{AgentivePhysicalObject}(x) \vee \text{AgentiveSocialObject}(x)) \wedge \text{Description}(y) \wedge \forall z ((\text{Perdurant}(z) \wedge \text{ParticipantIn}(y, z, t)) \rightarrow \exists w (\text{CommunicationEvent}(w) \wedge \text{PartOf}(w, y) \wedge \text{ParticipantIn}(y, w, t) \wedge \text{ParticipantIn}(x, w, t))))))$$

Besides physical agents, this axiom refers also to agentive social objects, a DOLCE category meant to include agentive ‘figures’ which are established by a society or community and can act like physical agents, can play roles, and so on. Examples of agentive social objects are, for instance, organizations like the FIAT Company or the International Monetary Fund. We assume that conceptions can be held by agentive social objects as well, through the agentive physical objects they depend on. Organizations are a topic which is currently under investigation in our Laboratory (Bottazzi & Ferrario, 2005; Gangemi et al., 2004). In order not to burden our proposal excessively, however, we do not pursue this notion further in this paper. All we say here is that a description requires some agent in the events it participate in, which is (at some time) able to **conceive** it.

Agentivity in DOLCE is not (explicitly) defined, but by means of DnS we can now define it as follows:

$$(D20) \text{ AgentivePhysicalObject}(x) =_{\text{df}} \text{PhysicalObject}(x) \wedge \exists y, t (\text{Description}(x) \wedge \text{Conceives}(x, y, t))$$

In simple words, this first level of agentivity is defined in (D20) in a wide sense as implying conception. A conception only requires intentionality in Brentano’s terms (i.e., the ability to represent something to oneself).

A second, stronger sense of agentivity involves the explicit conceiving of plans (see below). As stated in the previous section, this complies with the BDI paradigm, when it attributes to **cognitive agents** the ability of self-representing beliefs, desires, and intentions:

$$(D21) \text{ CognitiveAgentivePhysicalObject}(x) =_{\text{df}} \text{AgentivePhysicalObject}(x) \wedge \exists y, t (\text{Plan}(y) \wedge \text{Conceives}(x, y, t))$$

The way cognitive agents create, choose, or transform their conceptualizations (the nature of intentionality) is extremely diversified. We do not enter here this difficult area, leaving it to future investigation. We need, however, some preliminary distinction in order to relate **agents** and descriptions that represent those conceptualizations. In

order to simplify our formulas and try to comply with the common-sense polysemy of ‘agent’, we define it here as a catch-all class, encompassing either agentive physical objects or agentive social objects:

$$(D22) \text{ Agent}(x) =_{\text{df}} \text{AgentivePhysicalObject}(x) \vee \text{AgentiveSocialObject}(x)$$

We also introduce a restricted class for cognitive agents:

$$(D23) \text{ CognitiveAgent}(x) =_{\text{df}} \text{CognitiveAgentivePhysicalObject}(x) \vee (\text{AgentiveSocialObject}(x) \wedge \exists y (\text{CognitiveAgentivePhysicalObject}(y) \wedge \forall z ((\text{Perdurant}(z) \wedge \text{ParticipantIn}(x, z, t)) \rightarrow \exists w (\text{CommunicationEvent}(w) \wedge \text{PartOf}(w, z) \wedge \text{ParticipantIn}(x, w, t) \wedge \text{ParticipantIn}(y, w, t))))))$$

An important relation between agents and descriptions is **creation**, implying that a description is specifically dependent on a given (cognitive) agent:

$$(A24) \text{ Creates}(x, y) \rightarrow \text{CognitiveAgent}(x) \wedge \text{Description}(y) \wedge \text{SpecificallyDependsOn}(y, x) \wedge \exists t (\text{Conceives}(x, y, t) \wedge \neg \exists x', t' (t' < t \wedge \text{Conceives}(x', y, t')))$$

Another important relation between agents and descriptions is **adoption**, requiring (at least) creation and previous conceiving:

$$(A25) \text{ Adopts}(x, y, t) \rightarrow \text{Conceives}(x, y, t) \wedge \text{CognitiveAgent}(x) \wedge \text{Description}(y) \wedge \exists z (\text{CognitiveAgent}(z) \wedge \text{Creates}(z, y))$$

$$(A26) \text{ Adopts}(x, y, t) \rightarrow \exists t' (t' > t \wedge \text{Conceives}(x, y, t'))$$

5.3. Plans

Before introducing our notion of collective, we present here some axioms for plans.²⁰

We assume a **plan** to be a description that represents an action schema.²¹ A plan is conceived by a cognitive agent, defines or uses at least one task (a kind of course of actions) and one role (played by agents), and has at least one goal as a proper part:

$$(A27) \text{ Plan}(x) \rightarrow \text{Description}(x)$$

$$(A28) \text{ Plan}(x) \rightarrow \exists y, t (\text{Conceives}(y, x, t) \wedge \text{CognitiveAgent}(y))$$

$$(A29) \text{ Plan}(x) \rightarrow \exists y (\text{Task}(y) \wedge \text{Uses}(x, y))$$

$$(A30) \text{ Plan}(x) \rightarrow \exists y ((\text{Role}(y) \wedge \forall z, t (\text{Classifies}(y, z, t) \rightarrow \text{Agent}(z)) \wedge \text{Uses}(x, y)))$$

$$(A31) \text{ Plan}(x) \rightarrow \exists y (\text{Goal}(y) \wedge \text{ProperPart}(x, y))$$

²⁰ Gangemi et al. (2004). See Section 3.3 above.

²¹ A good candidate for formalizing plans is the SOURCE-PATH-GOAL schema (Johnson, 1987).

Examples of plans include: a way to prepare an espresso in the next 5 min, a company's business plan, a military air campaign, a car maintenance routine, a project to start a relationship, etc.

Tasks are courses that are (mostly) used to sequence activities, or other perdurants that can be under the control of a planner. They are defined by a plan, but can be used by other kinds of descriptions.

According to DnS, tasks, as all courses, are connected to roles by a relation that we call 'modal target', expressing the **modalities** that, in given descriptions, roles can have towards a course. The relation is temporalized to suggest that a description can preserve its identity against changes of structuring among components (though there can be mandatory structures for description identity):

$$(A32) \text{ ModalTarget}(x, y, t) \rightarrow (\text{Role}(x) \wedge \text{Course}(y) \wedge \text{TimeInterval}(t))$$

ModalTarget is the descriptive counterpart of the 'participant-in' relation used in DOLCE, i.e., modalities are participation modes. In other words, the 'modal target' relation can be used to reify, among others, alethic, epistemic, or deontic operators. For example, a person is usually obliged to drive in a way that prevents her from hurting other people; or a person can have the right to express her ideas. A subclass of modal-target relations representing dispositional attitudes towards courses is called AttitudeTowards, and it holds only when roles are played by cognitive agents:²²

$$(A33) \text{ AttitudeTowards}(x, y, t) \rightarrow (\text{ModalTarget}(x, y, t) \wedge \text{Course}(y) \wedge \forall z (\text{Classifies}(x, z, t) \rightarrow \text{CognitiveAgent}(z)))$$

In the context of DDPO, tasks can be considered as shortcuts for plans, since at least one role played by agents (possibly a different agent from the one initiating the task) implies a 'desire attitude' towards them:

$$(D24) \text{ DesireTowards}(x, y, t) =_{df} \text{AttitudeTowards}(x, y, t) \wedge \exists z, w, t (\text{Agent}(z) \wedge \text{Classifies}(x, z, t) \wedge \text{Uses}(w, x) \wedge \text{Uses}(w, y) \wedge \text{Conceives}(z, w, t))$$

$$(D25) \text{ Task}(x) =_{df} \text{Course}(x) \wedge \exists y, z (\text{Plan}(y) \wedge \text{Defines}(y, x) \wedge (\text{Role}(z) \wedge \forall w, t ((\text{Classifies}(z, w, t) \rightarrow \text{Agent}(w)) \wedge \text{Uses}(y, z) \wedge \text{DesireTowards}(z, x, t))))$$

The notion of **Goal** is more complicated, due to the widespread polysemy it suffers from. Here a goal is considered to be a desire (another kind of description) that is a part of a plan.

Desires in general are characterized as defining or using at least one role classifying an agent, and at least one

course. The role is played by the agent in a desire mode towards the course:

$$(A34) \text{ Desire}(x) \rightarrow \text{Description}(x)$$

$$(A35) \text{ Desire}(x) \rightarrow \exists y, t. \text{Conceives}(y, x, t) \wedge \text{CognitiveAgent}(y)$$

$$(A36) \text{ Desire}(x) \rightarrow \exists y, z ((\text{Role}(y) \wedge \forall w, t (\text{Classifies}(y, w, t) \rightarrow \text{Agent}(w)) \wedge \text{Course}(z) \wedge \text{Uses}(x, y) \wedge \text{Uses}(x, z) \wedge \text{DesireTowards}(y, z, t)))$$

For example, a desire to start a relationship can become a goal to start a relationship if someone decides to takes action, or actually takes action, or let someone else take action on her behalf with the purpose of starting the relationship.

We are proposing here a restrictive notion of **goal** that relies upon its desirability by some agent, which does not necessarily play a role in the execution of the plan the goal is part of. For example, an agent can have an attitude towards some task defined in a plan, e.g., duty towards, which is different from desiring it (desire towards). We might assume that a goal is usually desired by the creator or beneficiary of a plan. The minimal constraint for a goal is to be a proper part of a plan:

$$(D6) \text{ Goal}(x) =_{df} \text{Desire}(x) \wedge \exists y (\text{Plan}(y) \wedge \text{ProperPart}(y, x))$$

A **main goal** can be defined as a goal that is part of a plan but not of one of its subplans (i.e., it is a goal, but not a subgoal in that plan):

$$(D27) \text{ MainGoal}(x, y) =_{df} \text{ProperPart}(x, y) \wedge \text{Plan}(x) \wedge \text{Goal}(y) \wedge \neg \exists x' (\text{Plan}(x') \wedge \text{ProperPart}(x, x') \wedge \text{ProperPart}(x', y))$$

A **subgoal** (relative to a plan) is a goal that is a part of a subplan:

$$(D28) \text{ Subgoal}(x, y) =_{df} \text{Part}(x, y) \wedge \text{Goal}(y) \wedge \text{Plan}(x) \wedge \exists z (\text{Plan}(z) \wedge \text{ProperPart}(z, x))$$

A goal is not necessarily a part of the main goal of the plan it is a subgoal of. E.g., consider the goal: being satiated; eating food can be a subgoal of the plan that has being satiated as its main goal, but it is not a part of being satiated.

In interesting cases, supergoals can be created in order to support the adoption of a subgoal (see Section 5.2 above).

In order to describe these cases, we need to specialise the adoption relation. Goals and plans can in fact be adopted with different constraints:

$$(D29) \text{ AdoptsGoal}(x, y, t) =_{df} \text{Adopts}(x, y, t) \wedge \text{CognitiveAgent}(x) \wedge \text{Goal}(y) \wedge \forall z ((\text{Task}(z) \wedge \text{Uses}(y, z)) \rightarrow \text{DesireTowards}(x, z, t))$$

²² See Section 5.2 for more on agents, and Section 5.3 for more on tasks.

$$(D30) \text{AdoptsPlan}(x, y, t) =_{df} \text{Adopts}(x, y, t) \wedge \text{CognitiveAgent}(x) \wedge \text{Plan}(y)$$

In those interesting cases, given a plan and its main goal, e.g., some service to be delivered, it is a common practice to envisage the supergoals of the main goal that can be more clearly desirable from e.g., prospective users of a service (for example, a claim like the following generates a supergoal for the service's goal: our service will improve your life). In these cases, goal adoption and plan adoption are taken as if the following theorem would be undebatably sustainable, i.e., that goal adoption implies adopting all its subgoals:

$$(T8) ? (\text{AdoptsGoal}(x, y, t) \wedge \text{Subgoal}(y, z)) \rightarrow \text{AdoptsGoal}(x, z, t)$$

5.4. Collectives

As anticipated in Section 2, on the basis of the notions introduced so far we can now define a very general notion of collective as follows:

$$(D31) \text{CollectiveBS}(x) =_{df} \text{Collection}(x) \wedge \forall y, t ((\text{Membership}(y, x, t) \rightarrow \text{Agent}(y)) \wedge \exists z (\text{Plan}(z) \wedge \text{Unifies}(z, x)))$$

The intuition we want to model in D31 is that a collective in the broadest sense (CollectiveBS) – i.e., capable of being specialized according to different theories – is a collection of agents that is unified by a plan (see Fig. 1). So – following an example given by Sartre (1982) – people waiting at a bus stop are a collective: we can describe their behavior in terms of a plan that specifies the role of passenger, the task of hanging-around-the-bus-stop, and the goal of taking-a-bus.

This notion, of course, is far more general than the one given in the classical literature on collective intentionality, for instance by Gilbert (1992) & Tuomela (1995)). It does not require what is considered to be a

crucial feature of social groups: a mutual awareness of the common goal on the members' side. In our opinion, the broad category we are proposing allows for the modeling of many important sociological, socio-political, and economical notions – as, for instance, the notions of 'social class' and 'consumers' – that would not be included otherwise. Members of the upper class, for instance, are not necessarily aware of their respective goals, nor of the membership of one another. Nonetheless, from a sociologist's point of view this collective does exist. Each member pursues the same goal (suppose, for the sake of simplicity, 'to-increase-profits' or 'to-gain-political-power') and plays a specific role in this plan. In the same way, for example, different people buying a certain product or using a certain service are not necessarily aware of being members of a specific collective of 'consumers' but they are conceptualized as such in the context, for instance, of marketing strategies.

Moreover, since goals in our reconstruction are proper parts of plans, and having a plan is a sufficient condition to have a goal, it is possible to have a plan that specifies only its goal and the role(s) that the agents play, i.e., leaving empty – so to speak – the slot for tasks. This kind of plan is the most general one that can unify a collective, where the means to reach the ends are not explicitly stated.

The example of people waiting at the bus stop fits the notion of CollectiveBS: there is a plan ('waiting for the bus to come') that unifies the collective by means of a role (the passenger role). This role is played by all the members of the collective. Notice that the role can be defined by the plan, but it could also simply be used. For example, the passenger role could be defined in another description, representing a social practice related to people that use buses.

A collective in a narrow sense can now be introduced as a collective in a broad sense that is also characterized by one (or more) role(s):

$$(D32) \text{CollectiveNS}(x) =_{df} \text{CollectiveBS}(x) \wedge \exists y \text{Characterizes}(y, x)$$

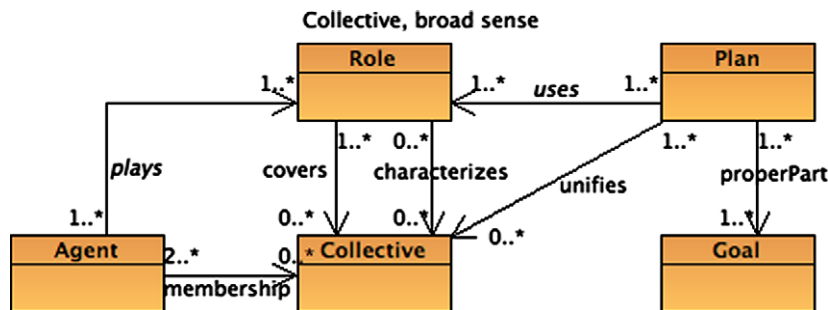


Fig. 1. A so-called “class diagram” of the definition of a broad sense of collective. Boxes represent classes (unary predicates), arrows represent associations (binary predicates). The numbers represent the cardinality of the associations when applied to said classes (e.g., “at least 1 plan unifies a collective, but at least 0 collectives are unified by a plan). The names of associations in *italics* indicate that the cardinalities shown are “local”, i.e., refer to the application of that association within the axiomatization of the class collective (in a broad sense).




Fig. 2. A so-called “class diagram” of the definition of a narrow sense of collective. Boxes represent classes (unary predicates), arrows represent associations (binary predicates). Again, names of binary predicates in *italics* indicate that the cardinalities shown are “local”, i.e., refer to the application of that predicate within the axiomatization of the class collective (in a narrow sense).

Collectives as defined in D32 are more tightly dependent on plans (see Fig. 2). Let us consider the example of the staff of a publishing house (Section 2). Besides the role played by each member of the collective – in this case employee – there must be one (or more) role(s) that characterize(s) the collective in question: project manager, editor, etc. These roles must be defined by the plan – for instance, the plan of publishing a textbook with the goal of increasing profits – and not simply used by it. In this narrow sense, plans create the roles for the collection.

According to our proposal, both broad- and narrow-sense collectives are intentional, since (i) their members are agents; and (ii) said members play at least one (covering) role in a plan; and (iii) the collectives are unified by plans that are conceived by a(t least one) cognitive agent. Put it shortly, a collective is intentional if there is a plan, and if the agentive members of a collection play the covering or characterizing roles of that plan.

6. Conclusions: applications and future work

We have presented a formal-ontological constructive account of intentional collectives based on complex axiomatizations of the following notions: collection, agent, plan, collective, description (adopted from DnS and DDPO), and various other foundational notions (adopted from DOLCE). Moreover, by applying a reification mechanism (DnS), we have made sure that all the needed notions are characterized by means of first-order axioms, which implies that we have a single domain of quantification for all entities and their relationships. This constitutes an important step towards software applicability.

For example, conceptual models of databases or other information systems aimed at supporting collaborative work (e.g., workflow and corporate knowledge management systems), community negotiation or characterization, etc., greatly benefit from a general, application-independent ontology of collectives. The vision of a Semantic Web (<http://www.w3.org/2001/sw/>), with globally annotated information resources, distributed but composable Web

Services, etc., largely depends on the ability to build precise and sensitive ontologies that reflect the lifecycle of the social reality that such information systems try to support. Within the Semantic Web programme, our ontology of collectives is being used for conceptual models of software (Oberle, 2005), content objects (Behrendt, Gangemi, Maass, & Westenthaler, 2005), electronic services (Oberle, Mika, Gangemi, & Sabou, 2004), and knowledge collectives, in the context of internationally funded projects and collaborations of our Laboratory (e.g., the EU project Metokis, <http://metokis.salzburgresearch.at>).

Our generic definition of collective opens up ways to distinguish between the diversified ways of acting collectively. In recent work, we have introduced²³ a typology of collectives that considers three possible sets of criteria: (1) the kind of criteria introduced for collections (internal organization); or (2) the amount and type of awareness shared by agents about a plan, its goal, its subplans, or its tasks; or (3) the kind of social practices related to the collective (e.g., workflows, sport, politics, religion, private relationships or interests, etc.). Notice that while the collectives identified through the first two sets of criteria, although lacking a rich lexical counterpart, can be axiomatized with a good detail using the methodology presented in this paper, the third set is much less predictable, but it does have a rich lexical counterpart. For example, in WordNet (Fellbaum, 1998), more than 500 senses refer to kinds of collectives.

Another direction we are following is based on a forthcoming typology of plans based on the formalization of the source-path-goal schema from cognitive linguistics (Johnson, 1987; Lakoff & Nunez, 2000).

Concerning the first set of criteria, similarly to what we have done for building our typology of collections, we distinguish simple and organized collectives. On the other hand, we need a finer-grained set of criteria for figuring out where collective action comes from.

²³ Cf. <http://dolce.semanticweb.org>, and look for the latest version of the DOLCE-Lite-Plus ontology library in the OWL language.

In collectives, roles are played by agents. Since agents can be involved in – and/or conceive – plans, roles can be assigned modalities or attitudes (participation modes) towards tasks that can sequence actions.

In many cases, plans are part of a wider descriptive context (including e.g., regulations, local constraints, etc.), therefore collective action emerges from the ‘bundle’ of descriptions that unifies the collective.

In other cases, whereas the bundle of descriptions is not anticipated, the collective action is an epiphenomenon, i.e., something that dynamically appears out of local conditions.

The typology we have built using the first two sets of criteria mainly exploits the inner structure (its goal and tasks) of the unifying plan. The prior existence of this plan, its conceivability by the members of the collective, and the amount, modes, and types of existence and conceivability are the criteria used.

To this purpose, other notions are introduced, in particular: **bringing about**, **control**, and **trust**.

Agents bring about²⁴ a collective when they create its unifying plan; for instance, consider a governmental agency bringing about a collective through a constitutive legal norm, and its related regulations (Gangemi, Prisco, Sagri, Steve, & Tiscornia, 2003):

(A37) $\text{BringsAbout}(x, y) \rightarrow \text{CognitiveAgent}(x) \wedge \text{Collective}(y) \wedge \exists z (\text{Plan}(z) \wedge \text{Creates}(x, z) \wedge \text{Unifies}(z, y))$

Agents control a collective when they conceive a meta-level plan involving the plan that unifies the collective:

(A38) $\text{Controls}(x, y, t) \rightarrow \text{Agent}(x) \wedge \text{Collective}(y) \wedge \exists z, w, r (\text{Plan}(z) \wedge \text{Unifies}(z, y) \wedge \text{Plan}(w) \wedge \text{Conceives}(x, w, t) \wedge \text{Uses}(w, r) \wedge \text{Plays}(z, r, t))$

As an example, consider a judge providing guidelines on how to interpret the regulations released for an institutional collective.

Trust is a more elusive notion (cf. the large literature, e.g., Castelfranchi, 2000, 2001; Castelfranchi & Falcone, 2000). In the axiomatization we are currently building, trust could be directed at members, at those who brought about the collective, at controllers, or at plans. Trust could be about truth, validity, or plausibility of a description, as well as about known reliability, disposition to follow norms and plans, etc. At a very general level, we see it as another form of conceivability over descriptions. At a different – and more specific – level, trust could be treated as a special kind of social relationship, linked to the notion of communities of trust.

In the applications mentioned above, our axiomatization of collectives serves several purposes, including: (1)

the formal description of communities and complex organizations – e.g., formal encoding of roles, plans, and the related groups and communities of agents, that makes it possible to simulate scenarios, to negotiate among peers, to assist executives in taking decisions, to study the dynamics of social groups, etc.; and (2) the formal description of plural resources for planning tasks – e.g., for building a workflow management system of a corporation that dynamically allocates work and collaboration according to profiles (either pre-defined or emerging out of work practices). Another task is the management of policies for access to software services, according to profiles (that can unify a simple collection of agents), or to communities of users; and (3) the automatic classification of plural agents – e.g., for discovering groups of persons behaving in ways that can be jointly considered for security reasons, one can create predefined (or can learn emerging) plans, and can draw inferences from the observation of sparse behaviors in a spatio-temporal context; and (4) the representation of class expressions in a first-order language, e.g., if one needs to restrict a set of guidelines to a professional figure (e.g., nurses), the intended relation is between a guideline and the (ideal or local) collective of nurses, for example according to a national registry that authorizes nurses to do their job.

Our results provide us, on the one hand, with preliminary indications about the intuitiveness and/or the plausibility of our axiomatization and, on the other hand, with new research questions. In particular, our ongoing and future research addresses the following issues concerning the applicability of our proposal to areas where collective/intentional concepts play a role: (1) Supporting the understanding and/or the representation of social and institutional reality. And (2) More specifically, how much development would be needed for a comprehensive ontology of organizations? Our taxonomy of plans and collectives allows, in principle, for the characterization of different types of organizations, from very simple to very complex ones. What else is needed for making this a viable option? And (3) Handling unambiguous sharing of plans and negotiation of meaning. This is a hot topic in ICT (Information Communication Technology) applications. And (4) Our modular approach to intentional collectives addresses a problem which is often overlooked in the theoretical literature: each different exemplar from a variety of collective entities relates differently to intentions. We believe that distributed AI and sociological research can benefit from an explicit, formal representation of their working assumptions concerning collectives and intentions.

Acknowledgments

We would like to thank Nicola Guarino, Claudio Masolo, Stefano Borgo, Laure Vieu, Roberta Ferrario, Pawel Garbacz, Gilles Kassel, and Geri Steve for their

²⁴ BringAbout is used here in a different way with respect to deontic logic; see, e.g., Carmo and Pacheco (2003).

help and precious comments on previous versions of this paper.

References

- Behrendt, W., Gangemi, A., Maass, W., & Westenthaler, R. (2005). Towards an ontology-based distributed architecture for paid content. In: A. Gomez-Perez, & J. Euzenat (Eds.), *Proceedings of the second European semantic web conference*. Springer.
- Borschev, V., & Partee, B. H. (2001). Genitive modifiers, sorts, and metonymy. *Nordic Journal of Linguistics*, 24(2), 140–160.
- Bottazzi, E., & Ferrario, R. (2005). A path to an ontology of organizations. Paper presented at the *EDOC international workshop on vocabularies, ontologies and rules for the enterprise (VORTE 2005)*, Enschede, The Netherlands.
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327–341.
- Brentano, F. (1924). *Psychologie vom empirischen Standpunkt. 2. Aufgabe, hrsg. von O. Kraus. Meiner* (Eng. trans. ed. by L.L. McAlister, Psychology from an Empirical Standpoint). London: Leipzig.
- Carmo, J., & Pacheco, O. (2003). A role based model for the normative specification of organized collective agency and agents interaction. *Autonomous Agents and Multi-Agent Systems*, 6(2), 145–148.
- Castelfranchi, C. (1998). Modelling social action for AI agents. *Artificial Intelligence*, 103(1), 157–182.
- Castelfranchi, C. (2000). Founding agent's 'Autonomy' on dependence theory. Paper presented at the *ECAI-2000: 14th European conference on artificial intelligence*, Berlin, Germany.
- Castelfranchi, C. (2001). Information agents: The social nature of information and the role for trust. In F. Zambonelli (Ed.), *Cooperative Information Agents V, 5th International Workshop, CIA 2001* (Vol. 2182, pp. 208–210). Modena, Italy: Springer.
- Castelfranchi, C. (2003). Grounding we-intention in individual social attitudes: On social commitment again. In K. Miller (Ed.), *Realism in action – Essays in the philosophy of social sciences*. Dordrecht: Kluwer.
- Castelfranchi, C., & Falcone, R. (2000). Trust and control: A dialectic link. *Applied Artificial Intelligence*, 14, 799–823.
- Cocchiarella, N. (2005). Denoting concepts, reference, and the logic of names, classes as many, groups and plurals. *Linguistics and Philosophy*, 28(2), 135–179.
- Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42, 213–261.
- Cohen, R. S., & Schnelle, T. (Eds.). (1986). *Cognition and fact – materials on ludwik fleck*. Dordrecht: Reidel.
- Dauben, J. W. (1979). *Georg Cantor: His mathematics and philosophy of the infinite*. Princeton: Princeton University Press.
- Devlin, K. (1993). *The joy of sets: Fundamentals of contemporary set theory*. New York: Springer-Verlag.
- Dugac, P. (1976). *Richard Dedekind et les fondements des mathématiques*. Paris: J. Vrin.
- Feinberg, J. (1968). Collective responsibility. *Journal of Philosophy*, 21, 222–251.
- Fellbaum, C. (1998). *WordNet. An electronic lexical database*. MIT Press.
- Ferrario, R., & Oltramari, A. (2004). Towards a computational ontology of mind. Paper presented at the *Formal ontology in information systems, Proceedings of the international conference FOIS 2004*, Torino (in press).
- French, P. A. (1984). *Collective and corporate responsibility*. New York: Columbia University Press.
- Gangemi, A., Catenacci, C., Lehmann, J., & Borgo, S. (2004). Task taxonomies for knowledge content: EU 6FP METOKIS Project D07. Available from <http://metokis.salzburgresearch.at>; updated version available at <http://www.loa-cnr.it/Publications.html>.
- Gangemi, A., & Mika, P. (2003). Understanding the semantic web through descriptions and situations. Paper presented at the *International conference ODBASE03*, Catania, Sicily (Italy).
- Gangemi, A., Prisco, A., Sagri, M. T., Steve, G., & Tiscornia, D. (2003). Some ontological tools to support legal regulatory compliance, with a case study. Paper presented at the *WORM03 workshop at ODBASE03 conference*, Catania, Italy.
- Gilbert, M. (1992). *Social facts*. Princeton, New Jersey: Princeton University Press.
- Gilbert, M. (1996). *Living together: Rationality, sociality, and obligation*. Lanham, Md, London: Rowman & Littlefield Publishers.
- Jacob, P. (2003). Intentionality. In E.N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2003 Edition). Available from <http://plato.stanford.edu/archives/fall2003/entries/intentionality>.
- Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination, and reason*. University of Chicago Press: Chicago, London.
- King, P. (2004). The metaphysics of Peter Abelard. In K. Guilfooy (Ed.), *The Cambridge companion to Abelard*. Cambridge, UK, New York: Cambridge University Press.
- Lakoff, G., & Nunez, R. E. (2000). *Where mathematics comes from: How the embodied mind brings mathematics into being*. New York: Basic Books.
- Link, G. (1983). The logical analysis of plurals and mass terms: A lattice-theoretical approach. In A. von Stechow (Ed.), *Meaning, use and interpretation of language* (pp. 302–323). Berlin: Walter de Gruyter.
- Marcus, R. B. (1993). Classes, collections, assortments, and individuals. In R. B. Marcus (Ed.), *Modalities: Philosophical essays* (pp. 90–100). New York: Oxford University Press.
- Masolo, C., Gangemi, A., Guarino, N., Oltramari, A., & Schneider, L. (2003). *WonderWeb Deliverable D18: The WonderWeb Library of Foundational Ontologies*.
- Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Ferrario, R., Gangemi, A., et al. (2004). Social roles and their descriptions. Paper presented at the *Proceedings of the ninth international conference on the principles of knowledge representation and reasoning*, Whistler, Canada.
- Oberle, D., Mika, P., Gangemi, A., & Sabou, M. (2004). Foundations for service ontologies: Aligning OWL-S to DOLCE. In S. Staab, & P. Patel-Schneider (Eds.), *Proceedings of the World Wide Web conference (WWW2004)*, Semantic Web Track.
- Oberle, D. (2005). *Semantic management of middleware*. Springer.
- Perler, D. (2003). *Théories de l'intentionnalité au moyen âge*. Paris: Vrin.
- Pettit, P. (2003). Groups with minds of their own. In F. F. Schmitt (Ed.), *Socializing metaphysics: The nature of social reality*. Lanham, Md, Oxford: Rowman & Littlefield Publishers.
- Rorty, A. O. (1986). Self-deception, akrasia and irrationality. In J. Elster (Ed.), *The multiple self*. Cambridge, UK: Cambridge University Press.
- Russell, B., & Whitehead, A. N. (1910). *Principia Mathematica*. Cambridge, UK: Cambridge University Press.
- Sartre, J. P. (1982). *Critique of dialectical reason* (Verso pbk. ed.). London, New York, NY: Verso/NLB.
- Searle, J. R. (1983). *Intentionality. An essay in philosophy of mind*. Cambridge, UK: Cambridge University Press.
- Searle, J. R. (1990). Collective intentions and actions. In M. E. Pollak (Ed.), *Intentions in communication* (pp. 401–415). Cambridge, Massachusetts: MIT Press.
- Searle, J. R. (1995). *The construction of social reality*. New York: The Free Press.
- Searle, J. R. (2001). *Rationality in action*. Cambridge, MA: MIT Press.
- Simons, P. (1987). *Parts: A study in ontology*. Oxford: Clarendon Press.
- Tuomela, R. (1995). *The importance of us: A philosophical study of basic social notions*. Stanford, CA: Stanford University Press.
- Tuomela, R. (2003a). Collective acceptance, social institutions, and social reality. *The American Journal of Economics and Sociology*, 62(1), 123–165.
- Tuomela, R. (2003b). The We-Mode and the I-Mode. In F. Schmitt (Ed.), *Socializing metaphysics: The nature of social reality*. Lanham, Md: Rowman and Littlefield.
- von Neumann, J. (1967). An axiomatization of set theory. In J. Van Heijenoort (Ed.), *From Frege to Gödel; a source book in*

- mathematical logic, 1879–1931*. Cambridge: Harvard University Press.
- Weber, M. (1968). *Economy and society; an outline of interpretive sociology*. New York: Bedminster Press.
- Wooldridge, M. J. (2000). *Reasoning about rational agents*. Cambridge, Mass, London: MIT Press.
- Zeman, J. J. (1982). Peirce on abstraction. *The Monist*, 65, 211–222.