From Collective Intentionality to Intentional Collectives: An Ontological Perspective

Emanuele Bottazzi, Carola Catenacci, Aldo Gangemi, Jos Lehmann

Laboratory for Applied Ontology, ISTC-CNR, Trento & Roma, Italy

{bottazzi, catenacci}@loa-cnr.it; gangemi@istc.cnr.it; joslehmann@ip.rm.cnr.it

Abstract

This paper gives an upside-down view of the problem of collective intentionality by providing a treatment of the notion of *intentional collective*. Based on reviews of the relevant literature, we apply two formal-ontological tools of our choice (namely, DOLCE and D&S) to the definition of the notions of collection, agent, plan and collective, all underlying the concept of intentional collective. Although our results are preliminary, 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.

1. Introduction

During the last decade the problems whether there exists such a thing as collective intentionality and what relationship the latter holds with individual intentions have been hotly debated in Philosophy of Society, Theory of Action and Social Ontology (see, for instance, (Gilbert 1992; Gilbert 1996); (Searle 1990; Searle 1995); (Bratman 1992); (Tuomela 1995; Tuomela 2003a; Tuomela 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-agentive physical entities and a number 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, Vieu 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 and interactions among social agents. Examples of this second sense are e.g. *euro*, *president*, and *FIAT*. This second sense of 'social' pivots on the idea that the social world is not only plural but also organized. Typically, each individual agent simultaneously 'belongs to' and acts within and across a multiplicity of 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 put black on white some of the basic (ontological) structures of social reality. 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*, which we (re)construct by means of formal ontological analysis. 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. In addition, we sketch and discuss some preliminary typologies of both collections and collectives.

The main upshots of the presented investigation are the explicitness, modularity and formality of the notions we introduce, as well as of the very methodology we follow. Explicitness, modularity and formality are key features for any conceptually structured vocabulary that is open to testing. No matter whether such testing is of a merely conceptual nature or of an inferential (formal or computational) nature, it will only be successful if the chunks of knowledge contained in the overall structure can easily be isolated, tested on their own, and updated. This is exactly what *can* be done in the structure presented in the following sections. Therefore, we believe that our framework provides a solid ground for the treatment of many notions relating to collectives and to intentionality because, despite some disputable choices it may contain, the framework is designed for controlled expansion and incremental correction.

The rest of the paper is structured as follows. Section 2 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. First of all, then, we want to provide 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 Engeneering (Masolo, Gangemi et al. 2003), Descriptions and Situations (Gangemi and Mika 2003), and some extensions of these two. Section 3 provides a formal-ontological account of the notion of collection in terms of what defined in section 2, along with a typology of collections. Section 4 provides a treatment of intentional collectives and of their typology in terms of the formally specified notions of collection, agent and plan. Finally, section 5 draws some conclusions.

The reported work is part of our Laboratory's research program dedicated to social ontologies. The modules (DOLCE, D&S, Ontology of Plans) reused in this paper for the formalization of notions related to collectives and intentionality have been – or are being – developed within EU academic and industrial projects, in the domain of knowledge-based systems.

---

[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 and Schnelle 1986).

2. Background concepts

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

2.1 DOLCE

DOLCE is a foundational ontology of particulars (vs. universals). 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 usually used in Artificial Intelligence (AI, in the following). Therefore, DOLCE, as most other ontologies in AI, makes no strong claims concerning the ‘deep’ metaphysical implications of its categories. In other words, DOLCE does not (claim to) refer to ‘true’ reality.

The domain of quantification of DOLCE ranges over *possibilia* (possible, not only actual particulars, so that we are allowed to talk of particulars that are postulated by existentially quantified variables, despite the latters are not explicitly introduced in a model (Masolo, Gangemi et al. 2003)).

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

- **Endurants** are Particulars *directly localized in* space (including Objects or Substances). Objects can be either *physical* or *non-physical*. Non-physical objects generically depend on (agentive) physical objects[3]. *Social* objects are a kind of non-physical objects.
- **Perdurants** are Particulars *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 them.
- **Qualities** are Particulars that *inhere in* either Endurants (as Physical or Abstract Qualities) or in Perdurants (as Temporal Qualities), and they 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** are 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. Furthermore, different kinds of space and time are admitted (e.g. Galilean vs. Newtonian).

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

[3] For a discussion of the way in which ‘agentivity’ is characterized in DOLCE and D&S, and for further refinements of this notion, see paragraph 4.2 below.

Different endurants or perdurants can be spatio-temporally co-localized. Relations between instances of the same category are admitted, such as *part*, *constitution*, *connectedness*, etc.

In Fig. 1 we show DOLCE by means of a UML class diagram, assuming a description logic-like semantics (Baader and Nutt 2003) where *classes* are interpreted as *concepts* (but dashed class boxes are interpreted as *individuals*), *generalization* is interpreted as *formal subsumption*, *association* is interpreted as a binary relation with cardinality encoded for it and its inverse (where no cardinality is indicated, the default is 0..*).

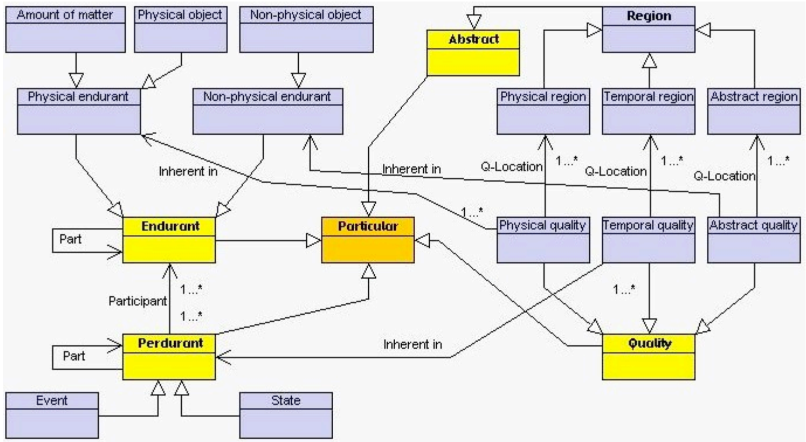

Fig.1. A UML class diagram showing the basic classes and relations of DOLCE.

2.2 D&S

The second formal ontological tool that we employ in this paper is D&S (Gangemi and Mika 2003; Gangemi, Catenacci et al. 2004), an apparatus conceived with the purpose of extending other (possibly, but not exclusively, foundational) ontologies. For instance, when using D&S 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. For example, 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 D&S *description* of social reality under a legal perspective. Such D&S 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 D&S 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 D&S, individual *constraints* and systems of constraints (theories) are *reified*, thereby becoming entities in the same domain of quantification of the entities from the ground

[4] An extension of DOLCE which does not use D&S, but introduces *descriptions* in order to treat social entities is presented in (Masolo, Vieu et al. 2004).

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

In more detail, D&S 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). This distinction may be used somewhat recursively (in the example of the legal domain, we may use the distinction to represent meta-juridical conceptualizations, i.e. *meta-norms*, or norms about norms).

D&S basic predicates and axioms are the following:

A **Description** is a social object which represents a conceptualization, hence it is generically dependent on some agent and communicable (Masolo, Vieu et al. 2004). Example of descriptions are regulations, plans, laws, diagnoses, projects, plots, techniques, etc.:

(A1) $\mathrm{Description}(x) \rightarrow \mathrm{NonAgentiveSocialObject}(x)$

(A2) $\mathrm{Description}(x) \rightarrow \exists y.\ \mathrm{AgentivePhysicalObject}(y) \wedge \mathrm{GenericallyDependsOn}(x,y)$

(A3) $\mathrm{Description}(x) \rightarrow \forall y.\ \mathrm{Part}(x,y) \rightarrow \mathrm{NonPhysicalObject}(y)$

Like physical objects, social ones have a lifecycle, can have parts, etc. Unlike physical objects, non-physical ones are **generically dependent on** some agentive physical object (for a discussion of the notion of 'agentivity', see paragraph 4.2 below).

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 D&S 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. Consequently, 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) $\mathrm{Situation}(x) =_{df} \mathrm{Particular}(x) \wedge (\exists y.\ \mathrm{Description}(y) \wedge \mathrm{Satisfies}(x,y)) \wedge (\exists z.\ \mathrm{Entity}(z) \wedge \mathrm{SettingFor}(z,x))$

(A4) $\mathrm{Situation}(x) \rightarrow \forall y.\ \mathrm{Part}(x,y) \rightarrow \mathrm{Situation}(y)$

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

(A5) $\mathrm{SettingFor}(x,y) \rightarrow \mathrm{Situation}(x) \wedge \mathrm{Particular}(y) \wedge \neg\mathrm{Situation}(y)$

(A6) $\mathrm{SettingFor}(x,y) \rightarrow \exists z.\ \mathrm{Perdurant}(z) \wedge \mathrm{SettingFor}(x,z)$

The **time and space of a situation** are the time and space of the particulars in the setting[5]:

[5] All 't'' variables in the formulas denote *time intervals*.

(A7) $\forall p,s,t1,t2.\ (\mathrm{Perdurant}(p) \wedge \mathrm{TimeInterval}(t1) \wedge \mathrm{TimeInterval}(t2) \wedge \mathrm{TemporalLocation}(p,t1) \wedge \mathrm{TemporalLocation}(s,t2) \wedge \mathrm{SettingFor}(s,p)) \rightarrow \mathrm{Part}(t2,t1)$

(A8) $\forall e,s,r1,r2.\ (\mathrm{Endurant}(e) \wedge \mathrm{SpaceRegion}(r1) \wedge \mathrm{SpaceRegion}(r2) \wedge \mathrm{SpatialLocation}(e,r1) \wedge \mathrm{SpatialLocation}(s,t2) \wedge \mathrm{SettingFor}(s,e)) \rightarrow \mathrm{Part}(r2,r1)$

Implicitly, (A7) and (A8) 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 Setting(x,y) can be inferred from the previous axioms.

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

(A9) $\mathrm{SAT}(x,y) \rightarrow \mathrm{Situation}(x) \wedge \mathrm{Description}(y)$

(A10) $\mathrm{SAT}(x,y) \rightarrow \exists z.\ \mathrm{Concept}(z) \wedge \mathrm{Uses}(y,z) \wedge \exists w,t.\ \mathrm{SettingFor}(x,w) \wedge \mathrm{Classifies}(z,w,t)$

This constraint is quite generic and even counterintuitive from a logical viewpoint. For specialized descriptions additional constraints should be given in order to reason over the satisfaction of candidate situations. This 'relaxed' semantics for satisfaction needs explanation.

In general, D&S does not constrain situations to include *only* particulars classified by the concepts of a description. In other words, reified satisfaction admits *redundant* models (which result to be *undecidable* from a strictly semantical point of view).

This assumption may seem logically rough, but real world uses of D&S have shown that most situations derive from legacy situations that already have an internal structure, and modifying them with the sole purpose of getting non-redundant situations seems a bad practice. For example, a *detective report* (a description) can depict a situation that may contain useless information from the point of view of a certain *legal rule* (another description) but, to a certain extent, it is important to preserve the *unity* of the reported situation, instead of 'cleaning' it up and making a new entity out of it, for the sake of merely satisfying the legal-rule description.

Under this assumption, the same situation can satisfy different descriptions that can even be unrelated. The formal consistency is given by the fact that legacy situations already satisfy other descriptions.

Moreover, D&S admits a *qualified satisfaction*: the set of concepts that 'must' classify a particular in a situation can be explicitly stated by means of a set of axioms that specialize the *satisfies* relation for a certain domain.

Summing up, reified satisfaction in D&S allows for situations that can be redundant on one hand (the respective non-reified models would be undecidable), and more restricted on the other hand (only certain non-reified models would be acceptable). Since reification allows a common domain for both ground and descriptive parts of an ontology, reified satisfaction does not lead to undecidability, and allows a custom design of the satisfiability conditions.

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 D&S, the primary ones (**role**, **course**, and **parameter**) being distinguished by the categories of particulars they classify in DOLCE:

(A11) $Defines(x,y) \rightarrow Description(x) \wedge Concept(y)$

(A12) $Classifies(x,y,t) \rightarrow Concept(x) \wedge Particular(y) \wedge TimeInterval(t)$

(A13) $Concept(x) \rightarrow NonAgentiveSocialObject(x) \wedge \exists y.\ Defines(y,x) \wedge Description(y)$

(D2) $Role(x) =_{df} Concept(x) \wedge \forall y,t.\ Classifies(x,y,t) \rightarrow Endurant(y)$

(D3) $Course(x) =_{df} Concept(x) \wedge \forall y,t.\ Classifies(x,y,t) \rightarrow Perdurant(y)$

(D4) $Parameter(x) =_{df} Concept(x) \wedge \exists y,t.\ Classifies(x,y,t) \wedge \forall y.\ Classifies(x,y,t) \rightarrow Region(y)$

Examples of roles are: *manager, student, assistant, actuator, toxic agent,* etc. Examples of courses are *routes, pathways, tasks*, etc. Examples of parameters are: *speed limits, allowed colors* (e.g. for a certain book cover), *temporal constraints*, etc.

Roles can be **specialized by** other roles, e.g. *president of the Italian republic* specializes *president of republic*:

(A14) $Specializes(x,y) \rightarrow Role(x) \wedge Role(y)$

(T1) $\forall x,y,t \exists z.\ (Classifies(x,y,t) \wedge Specializes(x,z) \wedge x \neq z) \rightarrow Classifies(z,y,t)$

Figures, or **social individuals** (either agentive or not), are other social objects defined by descriptions; differently from concepts, however, they do not classify particulars:

(A15) $Figure(x) \rightarrow SocialObject(x)$

(A16) $Figure(x) \rightarrow \exists y.\ Description(y) \wedge Defines(y,x)$

(A17) $Figure(x) \rightarrow \neg \exists y,t.\ Classifies(x,y,t)$

Examples of figures are organizations, political-geographic objects, sacred symbols, *personas*, etc.

Agentive figures are those which can conceive descriptions, by means of some agentive physical object that *acts for* the figure (for instance, as *representative* or *delegate*). The *conceives* relation is introduced in sec. 4. 2 below:

(D5) $AgentiveFigure(x) =_{df} Figure(x) \wedge AgentiveSocialObject(x) \wedge \exists y.\ Description(y) \wedge Conceives(x,y,t)$

(A18) $(AgentiveFigure(x) \wedge Conceives(x,y,t)) \rightarrow \exists z,t.\ AgentivePhysicalObject(z) \wedge Conceives(z,y,t)$

Agentive figures are established by a society or community; hence, they can act like a physical agent, can play roles, etc.. In our ontology, this formally amounts to have at least two descriptions, one defining an agentive figure, and another defining a role played by that agentive figure:

(A19) $AgentiveFigure(x) \rightarrow Figure(x) \wedge \exists y,z,w,t.\ Description(y) \wedge Role(z) \wedge Description(w) \wedge y \neq w \wedge Defines(y,z) \wedge Defines(w,x) \wedge Classifies(z,x,t)$

Typical agentive figures are societies, organizations, and in general all socially constructed persons[6]. The notion of *agentivity* is taken here in the sense of being – directly or indirectly – able to conceive a description (for a discussion of the relation *Conceives*, see sec. 4.2 below).

Figures are not dependent on roles defined or used in the same descriptions in which the figures themselves are defined or used, but they can act because they **depute** some tasks to some of those roles, which, in turn, must classify some individual agent. In other words, when a figure is classified by some agentive role, or participates in some event, it can be classified or participate because there is someone (or something) that is classified by other roles in the descriptions that define or use the figure. The relation is temporalized in order to suggest that a figure can preserve its identity despite changes of deputed roles (even though there are cases in which the identity of a figure is inextricably bound to one - or more - of its roles):

(A20) $DeputedBy(r,f,t) \rightarrow Role(r) \wedge Figure(f) \wedge \exists c,d,t.\ Course(c) \wedge Description(d) \wedge Uses(d,r) \wedge Uses(d,f) \wedge Uses(d,c) \wedge ModalTarget(r,c,t)$

(A21) $DeputedBy(r,f,t) \rightarrow \exists r1,t1.\ Role(r1) \wedge Classifies(r1,f,t1)$

Those roles classify endurants, which result to **act for** the figure:

(A22) $ActsFor(e,f,t) \rightarrow \exists r,t1.\ Role(r) \wedge DeputedBy(r,f,t1) \wedge Classifies(r,e,t)$

(A23) $(ParticipatesIn(f,p,t) \wedge AgentiveFigure(f)) \rightarrow \exists e.\ ActsFor(e,f,t) \wedge ParticipatesIn(e,p,t)$

For example, an employee *acts for* an organization that *deputes* the role (e.g. *turner*) that *classifies* the employee. Simply put, a guy working as a turner at FIAT acts for (or *on behalf of*) FIAT, so that in actions classified by turning tasks, if FIAT participates, so necessarily does the turner[7].

In complex figures, like organizations or institutions, a *total agency* is possible (usually limited to some actions), when an endurant plays a *delegate* or a *representative* role deputed by the figure.[8] Since figures are social objects, it can happen to find figures that *act for* other figures.[9]

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 D&S (e.g. a *norm* for enforcing norms will define a role that can classify the *enforced norm*).

The *classifies* relation is specialized by three subrelations: **played by**, **sequences**, and **valued by**, which apply to three different categories in DOLCE (Endurant, Perdurant, and Region, from (D2-4))[10]:

[6] Hobbes is among the first philosophers who introduced complex relations between artificial persons and physical, individual agents (like human beings); cf. (Hobbes 1996) Many social theorists see organizations as characterized by two fundamental dimensions: roles and rules (i.e., in our terms, *descriptions*); cf. (Fales 1977), (Biddle 1979), and (Scherer 2003). Cf. also the notion of *artificial institutional agent* in (Carmo and Pacheco 2003).

[7] This treatment of the indirect agentivity of figures takes into account a fundamental peculiarity of organizations, i.e., as Ladd pointed out , their *impersonality* (Ladd 1970).

[8] Cases of full delegation or representation, however, are quite unusual, and even prohibited in some legal contexts.

[9] Indeed, this kind of situation is at work in many contemporary settings and can reach great complexity, as e.g. in financial *chinese boxes*, which can even create an *agency loop*.

10 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.

(D6) PlayedBy(x,y,t) =df Role(x) ∧ Classifies(x,y,t)
(D7) Sequences(x,y,t) =df Course(x) ∧ Classifies(x,y,t)
(D8) ValuedBy(x,y,t) =df Parameter(x) ∧ Classifies(x,y,t)

Roles or figures and courses are related by relations expressing the **modalities** that (players of) roles and figures 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):

(A24) ModalTarget(x,y,t) → (Role(x) ∨ Figure(x)) ∧ Course(y)

Modal target is the descriptive counterpart of the "participant-in" relation used in the ground ontology, i.e. modalities are *participation modes*. In other words, the ModalTarget relation can be used to reify, for instance, *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:

(A25) AttitudeTowards(x,y,t) → ModalTarget(x,y,t) ∧ Task(y) ∧ ∀e. Classifies(x,e,t) → CognitiveAgent(e)

Parameters, roles, figures or courses are related by a **requisite for** relation, expressing the kind of requisites that particulars which are classified by 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):

(A26) RequisiteFor(x,y,t) → Parameter(x) ∧ (Role(x) ∨ Figure(x) ∨ 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 50kmph* 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 *50kmph*.

A final fragment from D&S is included that states the dependence of descriptions on some information object and support. Not only does a description generically depend on the agents that conceive it, but it also generically depends on its encoding in some language and support.

Information objects (IO) are non-agentive social objects which have various semiotic properties (Eco 1997; Gangemi, Catenacci et al. 2004). Here we only axiomatize their ability to **express** a *meaning* (corresponding in D&S to a description), and their dependence on a *support* that **realizes** them:

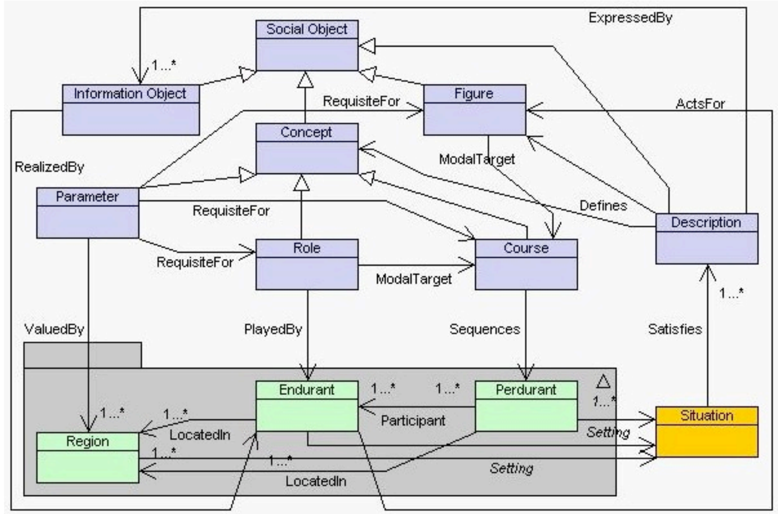

Fig. 2. A UML class diagram for D&S. The lower part of the pattern (within the grey package) is called the *ground ontology*, the higher is called the *descriptive ontology*; a situation satisfies a description if the two parts match according to the axioms specified for the concepts defined by the description.

(A27) InformationObject(x) → SocialObject(x)
(A28) InformationObject(x) → ∃y,t. Particular(y) ∧ RealizedBy(x,y,t)[11]
(A29) InformationObject(x) → ∀y,t. Expresses(x,y,t) → Description(y)
(A30) Description(x) → ∃y,t. InformationObject(y) ∧ ExpressedBy(x,y,t)

The relevance of IOs for collectives will emerge later, when considering that some collectives can exist only because their unity criterion can be communicated, shared, discussed, etc.

3. Collections

In order to formalize collectives with reference to a foundational ontology, and since in DOLCE and D&S there is not much support for them (cf. the notion of *unitary collection* in (Masolo, Gangemi et al. 2003)), we must extend the reused ontologies with a general characterization of ***collections**, as instances of some specialization of an existing DOLCE category*. We are touching here a difficult topic, with a heterogeneous literature ranging from metaphysics (Cocchiarella 2004; King 2004) to logic (Russell and Whitehead 1910; Zeman 1982), mathematics (Dugac 1976; Dauben 1979), and - more recently - linguistics and formal semantics (Link 1983; Marcus 1993). Since our main focus is not to significantly contribute to such formidable and interdisciplinary issue, we have tried to insulate ourselves from crossfire by excluding at least some of the entities we do not aim at describing.

[11] Information objects are reifications of pure information as social objects, hence they are assumed to be in space-time, and realized by some entity.

Collections are considered as *plural entities*, an open and debated topic in ontology (cf. (Simons 1987)). However, we do not address this wider dimension of the problem here. On the contrary, we face it from the viewpoint of the *constructive boundaries of those plural entities that form themselves a whole*. In simpler terms, we talk 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), but we will treat it from a different perspective, including social objects and talking explicitly of their unity criteria. Firstly, we distinguish collections from sets (the mentioned axioms are from axiomatizations of set-theory (ZFC, NBG) (von Neumann 1967; Devlin 1993):

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;
2. any two sets can be summed forming a union (axiom of union), while this is not tenable for any two collections;
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);
4. sets can be empty or singletons, but no empty or singleton collections are allowed;
5. (hyper)sets can be members of themselves (anti-foundation axiom), while collections do not;
6. the elements of a set have no part or constitution relation to the set, while collection elements have;
7. it is possible to conceive of sets the elements of which are parts of a same something, while this does not hold for collections (although spare parts can form a collection);
8. sets are abstract, having no space or time, while collections exist in time, and are localized.

In order to further distinguish our notion of collection from that of other authors, we declare our bias towards a *naturalized ontology*. 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.[12] It results that an ontology of naturalized entities is quite different from one of abstract, aeternal ones.

According to what stated above, a collection is neither a set nor a 'set naturalized in space-time', since i) we exclude collections of *physically connected parts*, ii) we assume that members have an explicit identity criterion, iii) we accept substitution of members of a collection while preserving its identity, and iv) we accept changes in the cardinality of a collection while preserving its identity.

Moreover, a collection cannot be a 'proper class naturalization' either, since a collection depends on its (at least two) members, while a proper class can be empty or a singleton.

If we assume that conceptualizations are embodied in cognitive agents, our 'collections' can be seen as naturalizations in space-time of non-empty proper classes with a ≥ 2 cardinality, and (at least one) basic properties for membership. This seems to capture the common sense intuition underlying groups, teams, collections, collectives, associations, etc.

Our ontology is therefore (provisionally) insulated from the issues arising in the long-standing debate (in philosophy, logic, and mathematics) on the nature of sets and classes as *abstract* entities, their relation to so-called *universals*, etc.

[12] We do not enter here the complex debate on the primacy of biology vs. society.

We say that a collection is *constituted by* its members, i.e. the *membership* relation defined on collections is a constitution relation. Endurants constituting a collection are either mereotopologically *unconnected* (e.g. statues in a statuary) or *weakly connected* (e.g. a pile of plates).

We also defend a *constructivist* position: a collection depends on one or more social objects that provide a *unity criterion* for it. When a (complex of) social objects 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, Vieu et al. 2004).

In order to provide a strong basis for any *naturalized* collection, we propose a formal version of the *containment image schema* informally introduced in cognitive semantics (Johnson 1987; Lakoff and 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. Participation of collections in actions or processes can be done '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, some Serbian terrorists 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 across all members[13].

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.

3.1 Definition of membership and collection

Collections are defined here as social objects (either agentive or not) which, although not defined by a description, (generically, one-sidedly, and temporarily) depend on member entities and (specifically, one-sidedly and constantly) depend on *concepts*, hence indirectly on descriptions; in some cases, collections can depend also on *figures*. 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*.

(D9) $\text{Collection}(x) =_{df} \text{SocialObject}(x) \wedge \exists r.\ \text{Role}(r) \wedge \forall w,t.\ \text{ConstituentOf}(w,x,t) \rightarrow \text{Classifies}(r,w,t) \wedge \exists y,z,t1.\ \text{Endurant}(y) \wedge \text{Endurant}(z) \wedge y \neq z \wedge \text{ConstituentOf}(y,x,t1) \wedge \text{ConstituentOf}(z,x,t1) \wedge \text{Classifies}(r,y,t1) \wedge \text{Classifies}(r,z,t1)$

A **membership** relation is defined on collections:

(D10) $\text{Membership}(e,c,t) =_{df} \text{ConstituentOf}(e,c,t) \wedge \text{Endurant}(e) \wedge \text{Collection}(c) \wedge \exists r.\ \text{Role}(r) \wedge \text{Classifies}(r,e,t)$

[13] '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.

In other words, a collection is a social object, the members of which are all classified by the same role, and which has at least two endurants as actual members.

Two or more collections can be *extensionally* equivalent and still not be the same collection. Each collection needs a unifying description which provides its *intensional* identity criterion:

(D11) ExtensionallyEquivalent(x,y) =df Collection(x) ∧ Collection(y) ∧ ∀z,t. Membership(z,x,t) ↔ Membership(z,y,t)

Members of collections cannot be parts of the same endurant:

(A31) Collection(x) → ¬∃z. Endurant(z) ∧ ¬ArbitrarySum(z) ∧∀y,t. Membership(y,x,t) → ProperPartOf(y,z,t)

The role shared by members has a **covering** relation towards the collection. The *definiens* of such relation is a theorem that follows from (D9-10):

(D12) Covers(r,c) $=_{df}$ Role(r) ∧ Collection(c) ∧ ∀w,t. Membership(w,c,t) → Classifies(r,w,t)

A purely logical property of collections is their **homogeneity**. We can say that a collection is *homogeneous* (at a time t) if all members have the same *leaf type* from the current ground ontology:

(A32) (Homogeneous(x,t) ∧ Collection (x)) ↔ ∀y,t. Membership(y,x,t) → (φ(y) ∧ ¬∃ψ. ψ(y) → φ(y))

Consider, for instance, a collection of musical instruments. Obviously, such a collection admits different kinds of musical instruments: let us say that our collection is made up of saxophones, drums and guitars. Now, take DOLCE as ground ontology. In DOLCE there are no predicates like Saxophone(x), Drum(x), or Guitar(x); all the endurants that constitute our collection are simply NonAgentivePhysicalObjects (NAPO in the following). NAPO, in DOLCE, is a leaf type, so the collection of musical instruments shows the property of being homogeneous with respect to DOLCE. If, on the other hand, we take the very same collection and a more specialized ground ontology, like, suppose, an ontology of musical instruments, which includes the types 'saxophone', 'drum', and 'guitar', the collection loses the property of being homogeneous.

As stated above, a *concept* is defined by a description and can classify some particular (a role being a concept that classifies only endurants), while a *figure* is defined by a description, but cannot classify any entity, and must act by means of something else. A *collection*, on the other hand, is not defined by a description, and cannot classify any particular, but has members that are classified by at least one and the same *role*.

Figures and collections are social *individuals*, while concepts are not. We may say that collections are *emergent* social individuals because, unlike figures and concepts, they do not need to be explicitly *defined by* a description.

Collections can be covered by roles of any generality. In the maximal case, a collection of *some conceivable objects*, a very generic *containee role*, defined in the *containment schema*, covers the collection, which on its turn plays the *container* role from the same

schema. **Containment schema** is a *cognitive schema* (Lakoff and Nunez 2000), which we represent as a *description*:[14]

(S1) Description(ContainmentSchema)
(S2) Role(Container)
(S3) Role(Containee)
(S4) Defines(ContainmentSchema, Container)
(S5) Defines(ContainmentSchema, Containee)
(A33) ∀c,t. Classifies(Container,c) → Collection(c)
(A34) ∀e,t. Classifies(Containee,e) → Endurant(e)
(A35) Collection(c) → Covers(Containee,c)

The containment schema could be also used to define the membership relation, since it makes (D10) true for any arbitrary collection. On the other hand, (D9) allows for the existence of multiple roles covering a collection.

3.2 Typology of collections

Several typologies of collections can be built, based on e.g. member types, covering role types, etc. Since our main scope is on *collectives* (collection with only agents as members), we limit collection types to one basic typology.

A **simple collection** (for instance, *a collection of saxophones,* or *a mass of lymphocytes*[15]) is a collection having only covering roles:

(D13) SimpleCollection(c) $=_{df}$ Collection(c) ∧ ∃r. Covers(r,c) ∧ ¬∃s. Characterizes(s,c)

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) MaximallyGenericCollection(c) $=_{df}$ SimpleCollection(c) ∧ ¬∃r. Covers(r,c) ∧ r≠Containee

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

(S6) Specializes(Containee#DB, Containee)
(A36) ∀e,t. Classifies(Containee#DB,e) → DinosaurBone(e)
(D15) DinosaurBonesCollection(c) $=_{df}$ Collection(c) ∧ Covers(Containee#DB,c)

The containee role can be specialized to any 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:

[14] 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.

[15] Including masses among collections is not straightforward, since most traditional views take masses as *continuous*. However, the traditional criterion for membership identification (*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".

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

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) $\text{ParametrizedCollection}(c) =_{df} \text{SimpleCollection}(c) \wedge \exists r,p,v,t.\ \text{Covers}(r,c) \wedge \text{RequisiteFor}(p,r,t) \wedge \text{ValueFor}(v,p) \wedge \forall e.\ \text{Membership}(e,c,t) \rightarrow \exists q.\ \text{InheresIn}(q,e,t) \wedge \text{Q-Location}(q,v)$

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)[16].

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 (figures, descriptions, collections) that either use or depute or are covered by them:

(D18) $\text{Characterizes}(r,c) =_{df} \text{Role}(r) \wedge \text{Collection}(c) \wedge \exists e,f,o,s,t.\ (\text{Figure}(o) \vee \text{Description}(o) \vee \text{Collection}(o)) \wedge \text{Role}(s) \wedge e \neq f \wedge r \neq s \wedge \text{Membership}(e,c) \wedge \text{Membership}(f,c) \wedge (\text{Uses}(o,r) \vee \text{Deputes}(o,r) \vee \text{CoveredBy}(o,r)) \wedge (\text{Uses}(o,s) \vee \text{Deputes}(o,s) \vee \text{CoveredBy}(o,s)) \wedge \text{Classifies}(r,e,t) \wedge \text{Classifies}(s,f,t)$

(T2) $\text{Characterizes}(r,c) \rightarrow \exists s.\ \text{Role}(s) \wedge r \neq s \wedge \text{Characterizes}(s,c)$

(D19) $\text{OrganizedCollection}(c) =_{df} \text{Collection}(c) \wedge \exists r,s.\ \text{Characterizes}(r,c) \wedge \text{Characterizes}(s,c) \wedge r \neq s$

From previous definitions and theorems, we can claim that collections specifically depend on some description:

(A37) $\text{Collection}(c) \rightarrow \exists d.\ \text{Description}(d) \wedge \text{SpecificallyDependsOn}(c,d)$

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 (A38-40), and is not temporalized, since changing the description (differently from changing some members) creates a new collection:

(A38) $\text{Unifies}(x,y) \rightarrow \text{Description}(x) \wedge \text{Collection}(y)$

(A39) $\text{Covers}(x,y) \rightarrow \exists d.\ \text{Description}(x) \wedge \text{Defines}(d,x) \wedge \text{Unifies}(d,y)$

(A40) $\text{Characterizes}(x,y) \rightarrow \exists d.\ \text{Description}(x) \wedge \text{Defines}(d,x) \wedge \text{Unifies}(d,y)$

(A41) $(\text{Characterizes}(x,y) \wedge \exists f.\ \text{Deputes}(f,x)) \rightarrow \exists d.\ \text{Description}(x) \wedge \text{Uses}(d,f) \wedge \text{Unifies}(d,y)$

From (A13), (D2), (D9), (D12) and (A39) we can derive that a collection must be unified by at least one description, which provides to said collection its unity criterion:

(T3) $\text{Collection}(c) \rightarrow \exists d.\ \text{Description}(d) \wedge \text{Unifies}(d,c)$

[16] 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 can imagine roles that are used by, deputed by, or that cover more than one description, figure, or collection[17]. In other words, characterizing roles can be related among them through some *composition* (or *bundle*) of descriptions, figures, or collections. We expect to extend our axiomatization to compositions and bundles in the near future.

With respect to dynamics, the identity of collections can resist change of some members. This behaviour 'corresponds to' the extensional/intensional nature of classes. For instance, if my collection of saxophones looses 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 exactly by me/you. Consider that 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'.

4. Towards collectives

Now that we have an explicit notion of collection, we want to use it for making a step towards the definition of social entities, starting with collectives. 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 (Sartre 1982) considers collectives as intermediate entities between collections and social groups. A similar view has been entertained, too, both in sociology and linguistics ((French 1984) and (Borschev and Partee 2001)).

We consider collectives to be something more than collections, since they are composed by agents, 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 (Searle 1990; Bratman 1992; Gilbert 1992; Searle 1995; Tuomela 2003b) .

Moreover, the classical literature is also strongly committed to the notion of 'we-mode': in a social group, agents think or speak about themselves in terms of 'we'. Some of the authors (Searle 1990; Gilbert 1992; Searle 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 analysis is not strictly necessary in order to define collectives.

[17] 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.

Our own definition of collectives is built around a 'descriptive' interpretation of the notions of intentionality (sec. 4.1), agent (sec. 4.2), and – especially – plan (sec. 4.3). This committment 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 (Cohen and Levesque 1990; Bratman 1992), 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[18]. 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 initial phase of our investigation we concentrate on the means-end type of rationality, but in Sect. 4.4 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.

4.1 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 (Brentano 1924), who gave new life and meaning to the medieval notion of *intentio*[19] 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 (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" [20]. This ordinary use of 'intention', on the other hand, seems closer to that typically made in Belief-Desire-Intention (BDI in the following) approaches, 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 the behavior of agents (Ferrario and 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 and 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.

[18] For an account of rationality and social reality, see (Searle 2001); and (Pettit 2003).
[19] For a treatment of intentionality in medieval philosophy cf., for instance, (Perler 2003)
[20] For an introduction to the different senses of 'intentionality', cf. also (Jacob 2003)

4.2 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 related 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[21]. 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, social dimension, and *planning activities*) seems to leave room for a distinction between, and characterization of, two levels of agentivity.

As stated in section 2.2, in DOLCE and DOLCE+ descriptions (like all non-physical objects) are **generically dependent on** some agentive physical object. We have further characterized the relationship between a description and an agent (see below) in the following axiom:

(A42) $\mathrm{Conceives}(x,y,t) \rightarrow \mathrm{GenericallyDependsOn}(y,x) \wedge \mathrm{Agent}(x) \wedge \mathrm{Description}(y)$

Hence, a description generically depends on some agent, which is (at some time) able to **conceive** it.

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

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

In simple words, this first level of *agentivity* is defined in (D20) in a wide sense as implying *conception* (to be characterized in a dedicated – but not developed as yet – ontology of mind). 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 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) $\mathrm{CognitiveAgentivePhysicalObject}(x) =_{df} \mathrm{AgentivePhysicalObject}(x) \wedge \exists y,t.\ \mathrm{Plan}(y) \wedge \mathrm{Conceives}(x,y,t)$

Conceptions can be held by *agentive social objects* as well, through the cognitive agentive physical objects they depend on:

[21] Notice that descriptions are expressed by some information object (see (A29)), which can be a natural-language proposition as well as some other (not necessarily language-like) encoding.

(A43) $(\text{Conceives}(x,y,t) \wedge \text{AgentiveSocialObject}(x)) \rightarrow \exists z,t.$ $\text{CognitiveAgentivePhysicalObject}(z) \wedge \text{GenericallyDependsOn}(x,z) \wedge$ $\text{Conceives}(z,y,t)$

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

We also introduce a restricted class for *cognitive agents:*

(D23) $\text{CognitiveAgent}(x) =_{df} \text{CognitiveAgentivePhysicalObject}(x) \vee$ $(\text{AgentiveSocialObject}(x) \wedge \exists y. \text{CognitiveAgentivePhysicalObject}(y) \wedge$ $\text{GenericallyDependsOn}(x,y))$

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

(A44) $\text{Creates}(x,y) \rightarrow \forall t. \text{Conceives}(x,y,t) \wedge \text{SpecificallyDependsOn}(y,x) \wedge$ $\text{CognitiveAgent}(x) \wedge \text{Description}(y)$

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

(A45) $\text{Adopts}(x,y,t) \rightarrow \text{Conceives}(x,y,t) \wedge \text{CognitiveAgent}(x) \wedge \text{Description}(y) \wedge$ $\exists z. \text{CognitiveAgentivePhysicalObject}(z) \wedge \text{Creates}(z,y)$

(A46) $\text{Adopts}(x,y,t) \rightarrow \exists t1. >(t1,t) \wedge \text{Conceives}(x,y,t1)$

4.3 Plans

Before discussing our typology of collectives, we introduce here some axioms for *plans*[22].

A **plan** is a description that represents an *action schema*. 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:

(A47) $\text{Plan}(x) \rightarrow \text{Description}(x)$

(A48) $\text{Plan}(x) \rightarrow \exists y,t. \text{Conceives}(y,x,t) \wedge \text{CognitiveAgent}(y)$

(A49) $\text{Plan}(x) \rightarrow \exists y. \text{Task}(y) \wedge \text{Uses}(x,y)$

(A50) $\text{Plan}(x) \rightarrow \exists c. (\text{Role}(c) \wedge \forall a,t. \text{Classifies}(c,a,t) \rightarrow \text{Agent}(a)) \wedge \text{Uses}(x,c)$

(A51) $\text{Plan}(x) \rightarrow \exists g. \text{Goal}(g) \wedge \text{ProperPart}(x,g)$

[22] A *plan ontology* is currently under development in our laboratory (see (Gangemi, Catenacci et al. 2004)).

Examples of plans include: a *way to prepare an espresso in the next five minutes*, a *company's business plan*, a *military air campaign*, a *car maintenance routine*, a *plan 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.

Tasks can be considered as *shortcuts* for plans, since at least one role played by agents has a 'desire attitude' towards them (possibly different from the one that puts the task into action):

(D24) $\text{DesireTowards}(x,y,t) \rightarrow \text{AttitudeTowards}(x,y,t) \wedge \exists e,d,t. \text{Agent}(e) \wedge$ $\text{Classifies}(x,e,t) \wedge \text{Uses}(d,x) \wedge \text{Uses}(d,y) \wedge \text{Conceives}(e,d,t)$

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

Tasks can be complex, and ordered according to an abstract succession relation. Tasks can relate to concrete actions or decision making; the latter deals with typical flowchart content. A task is different both from a flowchart node, and from an action or a class of actions.

A **complex task** is a task that has at least two other tasks as components:

(D26) $\text{ComplexTask}(x) =_{df} \text{Task}(x) \wedge \exists y,z. \text{Task}(y) \wedge \text{Task}(z) \wedge y \neq z \wedge$ $\text{Component}(x,y) \wedge \text{Component}(x,z)$

The primary ordering relation for tasks is **direct successor**; its transitive version is called **successor**. Notice that *successor* relations are abstract, and do not include a temporal ordering, although the usual correspondence within sequenced perdurants is a *temporal* relation (*precedes* or *overlaps*), and sometimes a *causal* relation. The distinction is clear when we consider two tasks having a direct successor relation holding for them, while the actions sequenced by them could temporally overlap:

(A52) $\text{DirectSuccessor}(x,y) \rightarrow \text{Particular}(x) \wedge \text{Particular}(y)$

(A53) $\text{Successor}(x,y) \rightarrow \text{Particular}(x) \wedge \text{Particular}(y) \wedge \forall z,w,k.$ $(\text{DirectSuccessor}(z,w) \wedge \text{DirectSuccessor}(w,k)) \rightarrow \text{Successor}(z,k)$

DirectSuccessor is irreflexive, antisymmetric, and intransitive. Successor is irreflexive, antisymmetric, and transitive.

An **elementary task** is a an atomic task:

(D27) $\text{ElementaryTask}(x) =_{df} \neg\exists y. \text{Component}(x,y) \wedge \text{Task}(y)$

An **action task** is an elementary task that sequences non-planning activities, like: moving, exercising forces, gathering information, etc. Planning activites are mental events involving some *rational* event:

(D28) $\text{ActionTask}(x) =_{df} \neg\exists y. \text{Sequences}(x,y) \wedge \text{PlanningActivity}(y)$

For example, "*eat your watermelon slice*" is an action task.

A **control task** is an elementary task that sequences a planning activity, e.g. an activity aimed at (cognitively or via simulation) anticipating other activities. Therefore, control tasks have usually at least one direct successor task (the *controlled* one), with the exception of *ending tasks* (see below):

(D29) $ControlTask(x) =_{df} Task(x) \wedge (\forall y.\ Sequences(x,y) \rightarrow PlanningActivity(y)) \wedge \exists z.\ Task(z) \wedge DirectSuccessor(x,z)$

The reification of control constructs allows to represent procedural knowledge into the same ontology including controlled action. Besides cognitive transparency and independency from a particular grounding system, a further advantage offered by reification is to enable the representation of *coordination* tasks and their relation to roles defined in the same plan.

For example, a *manager* that coordinates the execution of several related activities can be represented as a role with a *responsibility* (defined as a combination of duties and rights) towards a control task that has some complex task as a direct successor.

A plan can have several **proper parts** (regulations, goals, laws), including other plans.

If a plan uses a figure, that figure is defined by a constitutive description. If a plan defines a figure, the related constitutive description is a proper part of the plan:

(A54) $ConstitutiveDescription(x) \rightarrow Description(x)$

(T4) $\forall x,f.\ (Plan(x) \wedge Figure(f) \wedge Uses(x,f) \wedge \neg Defines(x,f)) \rightarrow \exists y.\ ConstitutiveDescription(y) \wedge Defines(y,f)$

(T5) $\forall x,f.\ (Plan(x) \wedge Figure(f) \wedge Defines(x,f)) \rightarrow \exists y.\ ConstitutiveDescription(y) \wedge ProperPart(x,y) \wedge Defines(y,f)$

For example, some plans define *temporary* figures, such as *teams* or *task forces*, whose lifecycle starts and ends within the plan lifecycle.

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:

(A55) $Desire(x) \rightarrow Description(x)$

(A56) $DesireTowards(x,y) \rightarrow AttitudeTowards(x,y)$

(A57) $Desire(x) \rightarrow \exists y,t.\ Conceives(y,x,t) \wedge CognitiveAgent(y)$

(A58) $Desire(x) \rightarrow \exists y,z.\ (Role(y) \wedge \forall a,t.\ Classifies(y,a,t) \rightarrow Agent(a)) \wedge Course(z) \wedge Uses(x,y) \wedge Uses(x,z) \wedge DesireTowards(y,z,t)$

For example, a *desire to start a relationship* can become a *goal to start a relationship* if someone *takes action* - or lets 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 not necessarily plays 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 say 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:

(D30) $Goal(x) =_{df} Desire(x) \wedge \exists p.\ Plan(p) \wedge ProperPart(p,x)$

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):

(D31) $MainGoal(p1,x) =_{df} ProperPart(p1,x) \wedge Plan(p1) \wedge Goal(x) \wedge \neg \exists p2.\ Plan(p2) \wedge ProperPart(p1,p2) \wedge ProperPart(p2,x)$

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

(D32) $Subgoal(x,y) =_{df} Part(x,y) \wedge Goal(y) \wedge Plan(x) \wedge \exists z.\ Plan(z) \wedge ProperPart(z,x)$

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*.

Nonetheless, we can also conceive of an **influence** relation between a goal and the main goal of the plan the first goal is a subgoal of:

(D33) $InfluenceOn(x,y) =_{df} Goal(x) \wedge Goal(y) \wedge \exists z.\ Plan(z) \wedge Subgoal(z,x) \wedge MainGoal(z,y)$

By using the previous definitions, we can also define a **disposition** relation between the roles used in a plan having a main goal, and the influenced goal:

(D34) $DispositionTo(x,y) =_{df} (Role(x) \wedge \forall a,t.\ Classifies(x,a,t) \rightarrow Agent(a)) \wedge Goal(y) \wedge \exists p,g.\ Plan(p) \wedge Goal(g) \wedge ProperPart(p,g) \wedge Uses(p,x) \wedge Goal(g) \wedge InfluenceOn(g,y)$

For example, the role *eater* can have a disposition to *being satiated*, meaning that a person playing the role of *eater* that adopts that plan can act in order to be satiated.

Disposition relation is useful to account for those cases in which a task addressed by a role is not *internal* to the plan, but the plan is a subplan of another one in which that task is represented as a full-fledged goal.

In interesting cases, supergoals can be created in order to support the adoption of a subgoal.

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

(D35) $AdoptsGoal(x,y,t) =_{df} Adopts(x,y,t) \wedge CognitiveAgent(x) \wedge Goal(y) \wedge \forall z.\ (Task(z) \wedge Uses(y,z)) \rightarrow DesireTowards(x,z,t)$

(D36) $AdoptsPlan(x,y,t) =_{df} Adopts(x,y,t) \wedge CognitiveAgent(x) \wedge Plan(y)$

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 *super*goals 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:

(T6) $?\ (AdoptsGoal(x,y,t) \wedge Subgoal(y,z)) \rightarrow AdoptsGoal(x,z,t)$

Plan executions are situations that proactively satisfy a plan, meaning that the plan anticipates its execution:

(D37) $PlanExecution(x) =_{df} Situation(x) \wedge \exists y. Plan(y) \wedge SAT(x,y) \wedge \exists t. PresentAt(y,t) \wedge \neg PresentAt(x,t)$

Subplan executions are parts of the whole plan execution:

(A59) $\forall p1,p2,s1,s2. (Plan(p1) \wedge Plan(p2) \wedge ProperPart(p1,p2) \wedge SAT(p1,s1) \wedge SAT(p2,s2)) \rightarrow ProperPart(s1,s2)$

A **goal situation** is a situation that satisfies a goal:

(D38) $GoalSituation(x) =_{df} Situation(x) \wedge \exists y. Goal(y) \wedge SAT(x,y)$

Contrary to the case of subplan executions, a goal situation is not part of a plan execution:

(A60) $GoalSituation(x) \rightarrow \forall y,p,s. (Goal(y) \wedge SAT(x,y) \wedge Plan(p) \wedge ProperPart(p,y) \wedge SAT(s,p)) \rightarrow \neg ProperPart(s,x)$

In other words, it is not true in general that any situation satisfying a part of a description is also part of the situation that satisfies the whole description:

(T7) $\forall p1,p2,s1 \neg \forall s2. (Plan(p1) \wedge Plan(p2) \wedge ProperPart(p1,p2) \wedge SAT(p1,s1) \wedge SAT(p2,s2)) \rightarrow ProperPart(s1,s2)$

This helps to account for the following cases:

- Execution of plans containing *abort* or *suspension* conditions (the plan would be satisfied even if the goal has not been reached, see below)
- *Incidental* satisfaction, as when a situation satisfies a goal without being intentionally planned (but anyway desired).

4.4 Collectives

Let us define a collective as a collection of *agents:*

(D39) $Collective(c) =_{df} Collection(c) \wedge \forall x,t. Membership(x,c,t) \rightarrow Agent(x)$

Similarly to all collections, collectives are covered or characterized by roles and eventually unified by some description.

Similarly to what we have done for 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*. In particular, we make use of *plans* to that purpose.

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

Plans can be framed in a wider descriptive context (e.g. regulations, local constraints, etc.), therefore collective action results to emerge from the 'bundle' of descriptions that unifies the collective.

Whereas this bundle is explicitly stated ('anticipated'), like in a closed set of tasks that describe, for instance, the possible actions for a figure, there exists a unique, communicable motivation (the plan defining the tasks) for the collective action.[23]

On the contrary, whereas the bundle of descriptions is not anticipated, the collective action is an *epiphenomenon*, or something that dynamically appears out of local conditions.

A preliminary typology of collectives is introduced (sec. 4.5) that mainly exploits the presence of a *plan* and of its inner structure (its *goal* and *tasks*) as unifying criterion for collective action. 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 build our typology.

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

Agents *bring about* a collective when they *create* its unifying plan; for instance, consider a governmental agency bringing about a collective through a constitutive norm, and its related regulations (Gangemi, Prisco et al. 2003):

(A61) $BringsAbout(x,y) \rightarrow CognitiveAgent(x) \wedge Collective(y) \wedge \exists z. Plan(z) \wedge Creates(x,z) \wedge Unifies(z,y)$

Agents *control* a collective when they *conceive* a meta-level *plan involving the plan unifying the collective*. For instance, consider a judge providing guidelines on how to interpret the regulations released for an institutional collective:

(A62) $Controls(x,y,t) \rightarrow Agent(x) \wedge Collective(y) \wedge \exists z,w,r. Plan(z) \wedge Unifies(z,y) \wedge Plan(w) \wedge Conceives(x,w,t) \wedge Uses(w,r) \wedge Plays(z,r,t)$

Trust can be directed to members, to those who brought about the collective, to controllers, or to plans. Trust is a very difficult issue (cf. the large literature, e.g. Castelfranchi (Castelfranchi 2000; Castelfranchi and Falcone 2000; Castelfranchi 2001). Trust can be about *truth*, *validity*, or *plausibility* of a description, as well as about *known reliability*, *disposition* to follow norms and plans, etc. We preliminarily treat it as another form of *conceivability* over descriptions. A different notion of trust will be treated as a special kind of *social relationship* (see below), linked to the notion of *communities of trust*. On the other hand, since trust on agents is based on social descriptions in which those agents can be involved in, we tentatively propose to encompass all trust notions under a special conceivability relation:

(A63) $Trusts(x,y,t) \rightarrow Conceives(x,y) \wedge Agent(x) \wedge Description(y)$
(A64) $Trusts(x,y,t) \rightarrow \exists t1. >(t1,t) \wedge Conceives(x,y,t1)$

4.5 Typology of collectives

Collectives can be classified according to different property kinds. The first one is definitely the **type** of members (e.g. physical persons, boys, cows, left-handers, etc.). Types are used in traditional classifications. For example, biological collectives can be

[23] Since a bundle of descriptions is needed in order to understand the origin of collective action, postulating a figure for each occurrence of a collective is tempting, but too strong, although very useful, as in cases of 'social engineering', marketing techniques (brands, logos, testimonials, etc.).

distinguished from ecological and social collectives, based on the (biological or social) properties ascribed to members.[24]

Biological collectives can be divided into various kinds (genetic, taxonomic, epidemiological, etc.). Biological properties produce either crisp or fuzzy/probabilistic types, and so-called *simple collectives* (see below) can be defined on them.

On the other hand, ecological and social collectives seem to be more resistant to a flat description in terms of simple properties. In most cases, the competence of members to conceive plans (in the generic sense outlined here of *action schemas*) creates the possibility of *being member* of such collectives.

For example, ecological collectives are not based only on the physical properties of the organisms, but they require also that organisms *interact with* the environment in an *effective* way, sometimes without completely conceiving a plan, but making it emerge casually or spontaneously (see below).

Social collectives are more obviously based on action schemas. They can be distinguished into neighborhood, geographic (at various granularities), ethnic, linguistic, commercial, industrial, scientific, political, religious, institutional, administrative, professional, sportive, interest-based, stylistic, devotional, etc.

The typologies just listed seem to be based on the *domain* the collectives pertain to. In what sense plans unify those collectives? They are probably based on bundles of descriptions that are too complex to be handled with plans alone. Therefore, it would be a long way to identify their properties, because the related social practices would need to be singled out in advance.

We introduce here some properties of collectives that are not strictly dependent on those domains, leaving to future investigation the formal linking of collectives to social practices.

We can conceive of **organized** collectives as opposed to **simple** ones by applying the same distinction we have used for collections in general:

(D40) SimpleCollective(x) $=_{df}$ Collective(x) ∧ ¬∃y,r. Description(y) ∧ Unifies(y,x) ∧ Role(r) ∧ Uses(y,r) ∧ Characterizes(r,x)

(D41) OrganizedCollective(x) $=_{df}$ Collective(x) ∧ ∃y,r. Description(y) ∧ Unifies(y,x) ∧ Role(r) ∧ Uses(y,r) ∧ Characterizes(r,x)

However, differently from what applied to collections, we will used the presence and structure of a unifying plan in order to further characterize kinds of collectives. A preliminary consideration is that plan unification can have two senses. The first one only takes into account the action schemas executed by the members, who do not necessarily interact in a 'global' way. In other words, the roles played by members *cover* the collective, because they are (dispositionally) played by each member.

The second sense is richer, and assumes that the unifying (maximal) plan uses roles that *characterize* the collective.

The first sense of plan unification is applicable to a subclass of simple collectives:

(D42) SimplePlannedCollective(x) $=_{df}$ SimpleCollective(x) ∧ AgentiveSocialObject(x) ∧ ∃y,r. Plan(y) ∧ Unifies(y,x) ∧ Role(r) ∧ Uses(y,r) ∧ Covers(r,x)

The second sense of plan unification applies to intentional collectives proper:

(D43) IntentionalCollective(x) $=_{df}$ Collective(x) ∧ AgentiveSocialObject(x) ∧ ∃y,r. Plan(y) ∧ Unifies(y,x) ∧ Role(r) ∧ Uses(y,r) ∧ Characterizes(r,x)

[24] Notice that biological, ecological and social collectives can be extensionally equivalent.

With respect to these two senses, it is not trivial to understand how the traditional typologies could be reconsidered. For example, a *neighborhood collective* could be conceptualized as *simple planned* by city administrators, and as *intentional* by a sociologist studying social interactions in urban areas. In these cases, we introduce a relation between two *extensionally equivalent* collectives, one simple planned, and the other intentional, and assuming that the second logically depends on the first. Hence, we could say that the second one **redescribes** the first:

(D44) Redescribes(x,y) $=_{df}$ IntentionalCollective(x) ∧ SimplePlannedCollective(y) ∧ ExtensionallyEquivalent(x,y) ∧ ∃t1,t2. PresentAt(x,t1) ∧ PresentAt(y,t2) ∧ >(t1,t2)

An *intentional collective* acts intentionally because its members act, and because it is unified by a plan that is conceived by some cognitive agent. Therefore, there is nothing special in a collective being intentional: it is just a matter of having a plan and agentive members playing its characterizing roles. What is special is the distinction between the diversified ways of acting collectively.

We postulate that intentional collectives have always a **maximal unifying plan** containing all unifying plans, and at least two of them:

(A65) IntentionalCollective(x) → ∃y. Plan(y) ∧ Unifies(y,x) ∧ ∀z. (z≠y ∧ Plan(z) ∧ Unifies(z,x)) → Part(y,z)

(D45) MaximalPlan(x) $=_{df}$ Plan(x) ∧ ∃y. IntentionalCollective(y) ∧ Unifies(x,y) ∧ ∀z. (z≠x ∧ Plan(z) ∧ Unifies(z,y)) → ProperPart(x,z)

Once we have introduced intentional collectives, and maximized their unifying plans, we can start introducing new distinctions. Firstly, is the unifying maximal plan **negotiated**, or potentially **conflicting**?

(D46) NegotiatedPlan(x) $=_{df}$ MaximalPlan(x) ∧ ¬∃p,q,y. Plan(p) ∧ Plan(q) ∧ ProperPart(x,p) ∧ ProperPart(x,q) ∧ Conflicts(p,q)

(D47) ConflictingPlan(x) $=_{df}$ MaximalPlan(x) ∧ ∃p,q,y. Plan(p) ∧ Plan(q) ∧ ProperPart(x,p) ∧ ProperPart(x,q) ∧ Conflicts(p,q)

For example, an agreement between a service provider and a client typically contains a negotiated plan to execute the service; a disagreement between two parties about how to carry out a task is a typical conflicting plan.

Conflict can be analyzed with respect to a dedicated description for different contexts. For example, in legal conflicting norms, a *compatibility scenario* description defines the concepts used to superordinate one norm to the other (Gangemi, Prisco et al. 2003). We take conflict and **superordination** as primitives here:

(A66) Conflicts(x,y) → Description(x) ∧ Description(y)

(A67) Superordinates(x,y) → Conflicts(x,y)

In a dispositional sense, we can then introduce **stable** and **unstable** collectives:

(D48) StableIntentionalCollective(x) → IntentionalCollective(x) ∧ ∃y. NegotiatedPlan(y) ∧ Unifies(y,x)

(D49) $\text{UnstableIntentionalCollective}(x) \rightarrow \text{IntentionalCollective}(x) \wedge \exists y.\ \text{ConflictingPlan}(y) \wedge \text{Unifies}(y,x)$

A further criterion for a typology of collectives is the behavior of the unifying plan with respect to the collective. In fact, the plan can be either underlying (i.e. **devised**) or **emerging**. An emerging collective temporally follows an *extensionally equivalent* collection (of agents) that is not unified either by any or the same plan. For instance, people from a collective (otherwise unified) can suddenly adopt a new plan, starting a new collective, as in the case of a group of drivers all stopping at a same service area because of a violent storm:

(D50) $\text{EmergingCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \exists y.\ \text{Collective}(y) \wedge y \neq x \wedge \forall e,t.\ (\text{Membership}(e,y,t) \rightarrow \text{Membership}(e,x,t)) \wedge \exists p,t1,t2.\ \text{Plan}(p) \wedge \text{Unifies}(p,x) \wedge \neg\text{Unifies}(p,y) \wedge >(t2,t1) \wedge \text{PresentAt}(x,t2) \wedge \text{PresentAt}(y,t1)$

Emergence of collectives is a case of the need to distinguish between extensionally and intensionally equivalent collections.

Emerging collectives can be **casual**, *or* **spontaneous.** *Casual collectives* are unified by a plan that has at least two subplans conceived by different agents who neither conceive their respective plans, nor the unifying plan. For example, some friends meet in a bar without having planned it:

(D51) $\text{CasualCollective}(x) =df\ \text{EmergingCollective}(x) \wedge \exists y.\ \text{Plan}(y) \wedge \text{Unifies}(y,x) \wedge \exists p,q,t1.\ \text{Plan}(p) \wedge \text{Plan}(q) \wedge \text{PropertPart}(y,p,t1) \wedge \text{PropertPart}(y,q,t1) \wedge \forall a,b,t.\ a \neq b \wedge \text{Conceives}(a,p,t) \leftrightarrow \text{Conceives}(b,q,t) \wedge \neg\text{Conceives}(a,q,t) \leftrightarrow \text{Conceives}(b,p,t) \wedge \neg\text{Conceives}(a,y,t) \wedge \neg\text{Conceives}(b,y,t)$

Of course, there exists at least one agent that conceives the unifying plan, but the time of conception is usually posterior to the beginning time of the collective's life (the unifying plan is 'reconstructed', and its subplans are not necessarily dependent on each other).

Spontaneous collectives are similar to casual ones, but the subplans conceived by the agents typically 'fit together', so that the agents start conceiving the unifying plan at the time of the emergence of the collective. For example, a group of drivers all stopping at a same service area because of a violent storm, and distributing into it in a way that makes them comfortable:

(D52) $\text{SpontaneousCollective}(x) =df\ \text{EmergingCollective}(x) \wedge \exists y.\ \text{Plan}(y) \wedge \text{Unifies}(y,x) \wedge \exists p,q,t1.\ \text{Plan}(p) \wedge \text{Plan}(q) \wedge \text{PropertPart}(y,p,t1) \wedge \text{PropertPart}(y,q,t1) \wedge \forall a,b,t.\ \text{Conceives}(a,p,t) \leftrightarrow \text{Conceives}(b,q,t) \wedge \neg\text{Conceives}(a,q,t) \leftrightarrow \text{Conceives}(b,p,t) \wedge \text{Conceives}(a,y,t1) \wedge \text{Conceives}(b,y,t1) \wedge >(t1,t)$

Another type of *intentional collectives* are those unified by plans that involve agentive figures. Based on (A20-22), we know that roles of members can be deputed by a figure, and that said members can *act for* that figure. Hence, a collective can be conceived as the (reification of the) maximal set of agents that act for the figure.

Collectives unified by such means have a special status, since they are **maximal agency collectives**:

(D53) $\text{MaximalAgencyCollective}(c) =_{df} \text{IntentionalCollective}(c) \wedge \forall x,t.\ \text{Membership}(x,c,t) \rightarrow \exists f.\ \text{AgentiveFigure}(f) \wedge \text{ActsFor}(x,f,t)$

This definition says that a maximal agency collective (for example, the maximal set of Apple employees) is a collective that has only members that act for the same figure, and that at least two of them exist[25].

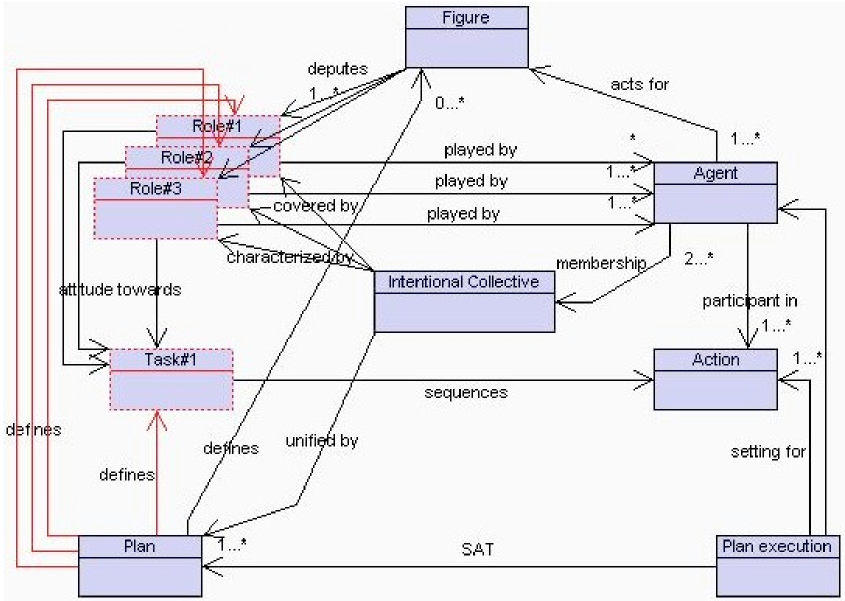

Fig. 3. A UML class diagram showing a definition of a class of (possible) intentional collectives unified by a plan that defines three roles and one task.

Given (D54) and (D18), it holds that:

(T8) $\text{MaximalAgencyCollective}(c) \rightarrow \exists f.\ \text{Characterizes}(f,c)$

Another criterion is based on the way a plan is *brought about* or *controlled* with respect to a collective. We can now add an axiom for emerging collectives: they cannot be unified by plans *brought about* by an agent:

(A68) $\text{EmergingCollective}(x) \rightarrow \forall p.\ (\text{Plan}(p) \wedge \text{Unifies}(p,x)) \rightarrow \neg\exists y.\ \text{CognitiveAgent}(y) \wedge \text{BringsAbout}(y,x)$

Contrary to emerging ones, **devised collectives** are unified by brought-about plans:

(D54) $\text{DevisedCollective}(x) =df\ \text{IntentionalCollective}(x) \wedge \forall p.\ (\text{Plan}(p) \wedge \text{Unifies}(p,x)) \rightarrow \exists y.\ \text{CognitiveAgent}(y) \wedge \text{BringsAbout}(y,x)$

[25] Of course, (D54) does not impose that actors must be classified by a single role deputed by the figure; cf. (A20) and (A21).

Governed collectives are distinguished from **ungoverned** ones based on the presence of an agent, who *controls* the collective by means of a plan or metaplan. An example of governed collective is the crew of a vessel. An example of an ungoverned collective is a rioting crowd (without any underlying manipulation):

(D55) GovernedCollective(x) =df IntentionalCollective(x) ∧ ∀p. (Plan(p) ∧ Unifies(p,x)) → ∃y,t. CognitiveAgent(y) ∧ Controls(y,x,t)

(D56) UngovernedCollective(x) =df IntentionalCollective(x) ∧ ∀p. (Plan(p) ∧ Unifies(p,x)) → ¬∃y,t. CognitiveAgent(y) ∧ Controls(y,x,t)

How many members of a collective share the conception of its unifying plan? On this basis, we can distinguish the following types:

Transparency: all members conceive the whole (maximal) plan; for instance, when a group of friends agrees on the destination of a trip:

(D57) TransparentCollective(x) $=_{df}$ IntentionalCollective(x) ∧ ∀e. Membership(e,x,t) → ∃p,t. Plan(p) ∧ Unifies(p,x) ∧ Conceives(e,p,t)

Opaqueness: not all members conceive the whole (maximal) plan; for instance, when a group of friends organizes a surprise party for one of them:

(D58) PartlyTransparentCollective(x) $=_{df}$ IntentionalCollective(x) ∧ ¬∀e. Membership(e,x,t) → ∃p,t. Plan(p) ∧ Unifies(p,x) ∧ Conceives(e,p,t)

Obscurity: no member conceives the whole (maximal) plan, while conceiving of a proper part of it; for instance, a collective of agents in a security network:

(D59) OpaqueCollective(x) $=_{df}$ IntentionalCollective(x) ∧ ∀e,t. Membership(e,x,t) → ¬∃p. Plan(p) ∧ Unifies(p,x) ∧ Conceives(e,p,t) ∧ ∃d. Plan(d) ∧ ProperPartOf(d,p,t) ∧ Conceives(e,d,t)

The degree of sharing of a unifying plan across members depends on the information objects that express the plan; in other words, *communication* plays a major role in intentional collectives. The way information is conveyed and spread out is another criterion to distinguish between collectives (not investigated here) .

A similar typology of sharing can be created by substituting *conception* with other 'modes' of sharing plans: *goal sharing, adoption sharing* or *trust sharing.*

For example, a **transparently embracing collective** is an intentional collective whose members have all adopted the conceived (maximal) plan, e.g. a group of friends decides to leave to a common destination:

(D60) TransparentlyEmbracingCollective(x) $=_{df}$ IntentionalCollective(x) ∧ ∀e,t. Membership(e,x,t) → ∃p. Plan(p) ∧ Unifies(p,x) ∧ AdoptsPlan(e,p,t)

As a second example, a **transparently trustful collective** is an intentional collective whose members all trust the conceived (maximal) plan; for instance, when a group of friends is confident to leave to a common destination with sufficient resources, good directions, competent drivers, etc.:

(D61) TransparentlyTrustfulCollective(x) $=_{df}$ IntentionalCollective(x) ∧ ∀e,t. Membership(e,x,t) → ∃p. Plan(p) ∧ Unifies(p,x) ∧ Trusts(e,p,t)

Shared adoption and trust, too, are influenced by communication. The dynamics of identification, imitation, deception, and leadership are additional dimensions for describing collectives.

Further discussions can clarify the relations between conception, adoption, and trust, but the complexity involved here makes us move to the more general perspective of *social relationships*, namely to the *internal structure* of unifying plans.

Criteria based on **social relationships** take into account the following elements: the *goals* that are proper parts of the unifying plans, the *relations between the concepts* defined or used within the plan, *other types of descriptions* intertwined with the maximal plan, e.g. *scripts* or *rules*.

Other types of collectives may be singled out based on whether a goal is *conceived, adopted,* or *trusted* without, though, conceiving the whole plan. For example, a group of people could enthusiastically adopt the goal of going to a restaurant that is 3 km far, as suggested by one of them or by an external agent. Such adoption does not require to conceive the actual plan adopted by the proposing agent(s), which may involve, for instance, to reach the place on foot, rather than by car.

As concerns relations between concepts, an example can be a group of two physical agents constituting a *master-slave* collective, unified by a plan in which, *towards* certain tasks, the *master* role has *rights*, while the *slave* role has only *duties*.

As a more general example, maximal agency collectives from a company can be analyzed into complex collectives interacting according to subtle hierarchies of roles, statuses, functions, tasks, etc. The relations among roles based on reciprocal influence, responsibility, obligations, expectations, or even trust, create an extremely varied typology of intentional collectives. We leave this area – traditionally investigated by sociology – untouched for the moment, although very interesting suggestions for further ontological analysis can be derived from new ideas about communities and acquaintance arising, for instance, in *web-based social networks* (Mika and Gangemi, 2004). The contribution of this paper to the analysis of complex collectives is limited to the presented formal framework, and to the basic tools it provides to an understanding of organized collectives; a framework which is plugged into a foundational ontology that can be reused as a component in philosophical inquiries, organizational studies, and information science.

The criteria presented so far have used properties of the members of collectives. Other criteria may obtain by looking at properties of collectives as *wholes, i.e. as interacting with other collectives or objects of any kind*. For example, collectives unified by a plan that defines a schedule (i.e. having an explicit execution time), with a limited lifecycle, can be considered to be **temporary**:

(D62) TemporaryCollective(x) =df Collective(x) ∧ ∃y,z. Plan(y) ∧ Schedule(z) ∧ Unifies(y,x) ∧ Defines(y,z)

We wrap up here the typology (not a partition!) drafted so far for collectives (Fig. 4):

Collective
- Simple (covered by roles, and not unified by plans with *characterizing* roles)
  - Type-based
    - Genetic
    - Taxonomic
    - Epidemiological

Simple-planned
Organized
Intentional (unified by plans with *characterizing* roles)
Stable vs. Unstable (based on negotiated vs. conflicting plans)
Devised vs. Emerging (based on presence of bringing about)

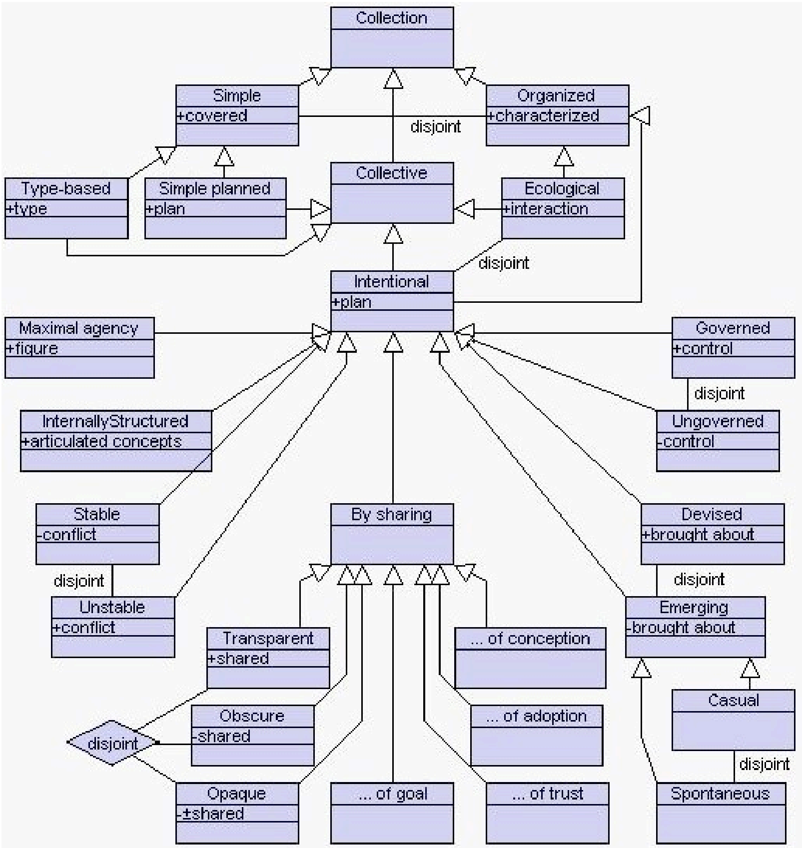

Fig. 4. Our preliminary formal typology of collectives.

[Emerging]: Casual vs. Spontaneous (based on time of plan conception)
Maximal agency collective (based on figure)
Governed vs. Ungoverned (based on control)
Transparent, Opaque, Obscure (based on degree of plan sharing across members)
By modes of plan sharing
(of goal)
(of conception)

(of adoption)
(of trust)
By internal structure of plans and/or related descriptions
Temporary (scheduled)

An excerpt of the classical and new examples that we are using to check the applicability of our typology of intentional collectives is proposed here:

1. A group of people running to a common shelter because it has suddenly started to rain (Searle 1990). *Simple: planned, temporary.*
2. An outdoor ballet where the choreography calls for the entire *corps de ballet* to converge on a common shelter (Searle 1990). *Intentional: stable, devised, transparent (on goal, conception, adoption, trust), internally structured.*
3. Businessmen having the same goal (i.e. pursuing their own selfish interests) as well as mutual beliefs about their respective intentions, but not cooperating or acting together (Searle 1990). *Intentional: stable*, *emerging*, *transparent (on goal), ungoverned, internally unstructured*.
4. A football team trying to execute a pass play (Searle 1990). *Intentional: stable, devised, transparent (on conception, goal and adoption), governed, internally structured.*
5. Nazi Germans as possessed by a self-distructive desire (according to a subsequent psycho-historical reconstruction). *Intentional: emerging, spontaneous, obscure (on goal and adoption)*.
6. CIA agents executing orders into a setting about which they are informed "on a strictly need-to-know basis". *Intentional: stable, devised, obscure (on goal, conception, adoption), governed, internally structured, temporary.*
7. The actors of an organization (e.g. an oil company) which, in addition to its "constitutive" plan, plays a role in further plans (e.g. fuelling civil wars in oil areas like African countries)[26]. *Intentional: maximal agency, stable, devised, opaque (on goal, conception, adoption), governed, internally structured.*
8. Fans in a stadium performing the so-called "ola" (wave). *Intentional: stable, emerging, spontaneous, transparent (on goal, adoption, trust), ungoverned, temporary.*
9. The human agent seen as a collection of temporal parts of herself, or as a collection of co-existing self systems (sub-agentive collectives). *Intentional: maximal agency, unstable, emerging, opaque (on goal, conception, adoption, trust), (un)governed (depending on possible pathologies and/or different neuropsychological theories).*
10. The employees/workers in a SAP workflow, or a "Ford-style" production line. *Intentional: maximal agency, stable, devised, obscure (on conception, adoption), governed.*

5. Conclusions 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 D&S), and various other foundational notions (adopted from DOLCE). Moreover, by applying a reification mechanism (D&S), 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

[26] Maybe these are not properly 'further' plans; consider, for instance, an organization whose 'constitutive' overall plan, or *mission*, does not explicitly specify *sub-plans,* but whose representatives actually implement some of them (like fuelling civil wars in oil areas).

relationships. This constitutes an important step towards computational tractability.

The resulting typologies of collections and collectives 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, in the future we would like to address the following issues concerning the applicability of our proposal to areas where collective/intentional concepts play a role:

1. Can our framework support the understanding and/or the representation of social and institutional reality? How?
2. More specifically, how much development would be needed for the treatment of the notion of organization? In theory our taxonomy of plans and collectives allows 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?
3. How can our proposal contribute to handling unambiguous sharing of plans and negotiation of meaning? This is a hot topic in Semantic Web circles.
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. Could this line of research result in any relevant contribution to, for instance, distributed AI or (with due adaptations!) to the methodology of sociological research?
5. On the technical side, how could our analysis of collectives be used as a basis to define *qualified participation*? Imagine, for instance, a relation of **organized co-participation** within a collective, i.e. a relation that exploits a path through the related tasks towards which the agents playing the relevant roles have an attitude.

Acknowledgments

We would like to thank Nicola Guarino, Claudio Masolo, Roberta Ferrario and Geri Steve for their help and precious comments on draft versions of this paper.

References

Baader, F. and W. Nutt (2003). Basic description logics. The Description Logic Handbook: Theory, Implementation and Applications. F. Baader, D. Calvanese, D. McGuinness, D. Nardi and P. F. Patel-Schneider. Cambridge, Cambridge University Press: 43-95.

Biddle, B. J. (1979). Role Theory. Expectations, Identities, and Behaviors. New York, Academic Press.

Borschev, V. and Partee B. H. (2001). "Genitive Modifiers, Sorts, and Metonymy." Nordic Journal of Linguistics 24(2): 140-160.

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. and O. Pacheco (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. ECAI-2000: 14th European Conference on Artificial Intelligence, Berlin, Germany, IOS Press.

Castelfranchi, C. (2001). Information Agents: The Social Nature of Information and the Role for Trust. Cooperative Information Agents V, 5th International Workshop, CIA 2001. M. Klusch and F. Zambonelli. Modena, Italy, Springer. 2182: 208-210.

Castelfranchi, C. (2003). Grounding We-Intention in Individual Social Attitudes: On Social Commitment Again. Realism in Action - Essays in the Philosophy of Social Sciences. M. Sintonen and K. Miller. Dordrecht, Kluwer.

Castelfranchi, C. and R. Falcone (2000). "Trust and control: a dialectic link." Applied Artificial Intelligence 14: 799-823.

Cocchiarella, N. (2004). "Denoting concepts. reference, and the logic of names, classes as many, groups and plurals (to be published)." http://www.formalontology.it/essays/plurals.pdf.

Cohen, P. R. and H. J. Levesque (1990). "Intention is choice with commitment." Artificial Intelligence 42: 213-261.

Cohen, R. S. and T. Schnelle, Eds. (1986). Cognition and Fact - Materials on Ludwik Fleck. Reidel, Dordrecht.

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 mathematiques. Paris, J. Vrin.

Eco, U. (1997). Kant e l'Ornitorinco. Milano, Bompiani.

Fales, E. (1977). "The Ontology of Social Roles." Philosophy of Social Sciences 7(2): 139-161.

Feinberg, J. (1968). "Collective Responsibility." Journal of Philosophy, 21: 222–51.

Ferrario, R. and A. Oltramari (2004). Towards a Computational Ontology of Mind. to appear in Formal Ontology in Information Systems, Proceedings of the International Conference FOIS 2004, Torino.

French, P. A. (1984). Collective and Corporate Responsibility. New York, Columbia University Press.

Gangemi, A., C. Catenacci, J. Lehmann and S. Borgo (2004). Task taxonomies for knowledge content, EU 6FP METOKIS Project D07, http://metokis.salzburgresearch.at.

Gangemi, A. and P. Mika (2003). Understanding the Semantic Web through Descriptions and Situations. International Conference ODBASE03, Catania, Sicily (Italy), Springer.

Gangemi, A., A. Prisco, M. T. Sagri, G. Steve and D. Tiscornia (2003). Some ontological tools to support legal regulatory compliance, with a case study. WORM03 Workshop at ODBASE03 Conference, Catania, Italy, Springer.

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.

Hobbes, T. (1996). Leviathan. Oxford, Oxford University Press.

Jacob, P. (2003). Intentionality. The Stanford Encyclopedia of Philosophy (Fall 2003 Edition). E. N. Zalta. URL = <http://plato.stanford.edu/archives/fall2003/entries/intentionality/>.

Johnson, M. (1987). The body in the mind : the bodily basis of meaning, imagination, and reason. Chicago ; London, University of Chicago Press.

King, P. (2004). The Metaphysics of Peter Abelard. The Cambridge companion to Abelard. J. E. Brower and K. Guilfoy. Cambridge, U.K. ; New York, Cambridge University Press.

Ladd, J. (1970). "Morality and the Ideal of Rationality in Formal Organizations." The Monist 54(4): 488-516.

Lakoff, G. and R. E. Nunez (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. Meaning, Use and Interpretation of Language. R. Bauerle, C. Schwarze and A. von Stechow. Berlin, Walter de Gruyter: 302–23.
Marcus, R. B. (1993). Classes, Collections, Assortments, and Individuals. Modalities: Philosophical Essays. R. B. Marcus. New York, Oxford University Press: 90-100.
Masolo, C., A. Gangemi, N. Guarino, A. Oltramari and L. Schneider (2003). WonderWeb Deliverable D18: The WonderWeb Library of Foundational Ontologies.
Masolo, C., L. Vieu, E. Bottazzi, C. Catenacci, R. Ferrario, A. Gangemi and N. Guarino (2004). Social Roles and their Descriptions. Procedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning, Whistler, Canada.
Mika, P. and A. Gangemi (2004). Descriptions of Social Relations. 1st Workshop on Friend of a Friend, Social Networking and the Semantic Web. D. Brickley, L. Miller, S. Decker, R.V. Guha. URL = <http://www.w3.org/2001/sw/Europe/events/foaf-galway/>.
Perler, D. (2003). Théories de l'intentionalité au moyen âge. Paris, Vrin.
Pettit, P. (2003). Groups with Minds of Their Own. Socializing metaphysics : the nature of social reality. F. F. Schmitt. Lanham, Md. ; Oxford, Rowman & Littlefield Publishers.
Rorty, A. O. (1986). Self-deception, akrasia and irrationality. The Multiple Self. J. Elster. Cambridge, UK, Cambridge University Press.
Russell, B. and A. N. Whitehead (1910). Principia Mathematica. Cambridge, UK, Cambridge University Press.
Sartre, J. P. (1982). Critique of dialectical reason. London, New York, Verso/NLB.
Scherer, A. G. (2003). Modes of Explanation in Organization Theory. The Oxford handbook of organization theory. H. Tsoukas and C. Knudsen. Oxford, Oxford University Press: 310-344.
Searle, J. R. (1983). Intentionality. An Essay in Philosophy of Mind. Cambridge, UK, Cambridge University Press.
Searle, J. R. (1990). Collective Intentions and Actions. Intentions in Communication. P. R. Cohen, J. Morgan and M. E. Pollak. Cambridge, Massachusetts, MIT Press: 401-415.
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, Calif., 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. Socializing Metaphysics: The Nature of Social Reality. F. Schmitt. Lanham, Md, Rowman and Littlefield.
von Neumann, J. (1967). An Axiomatization of Set Theory. From Frege to Gödel; a source book in mathematical logic, 1879-1931. J. Van Heijenoort. 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: 211-22.