

Running head: Intentional Collectives

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

Contact information:

Laboratory for Applied Ontology
Institute for Cognitive Sciences and Technology
National Research Council (ISTC-CNR)
Via Nomentana 56
00161, Roma
Italy
Email: a.gangemi@istc.cnr.it
Office: +390644161535
Fax: +390644161513

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 three formal-ontological tools of our choice (namely, DOLCE, D&S, and DDPO) 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, hence for contributing to both philosophic and scientific research on these topics.

Keywords: collective intentionality, formal ontology, social entities, collections, collectives, agents, concepts, roles, plans.

From Collective Intentionality to Intentional Collectives: an Ontological Perspective

1. Introduction

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

Reasoning along these lines, in previous work we have distinguished at least two senses in which an entity can be said to be ‘social’ (cf. Masolo, Vieu, Bottazzi, Catenacci, Ferrario, Gangemi, 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 consumer, as well as International Monetary Fund, Ethical Committee, and FIAT. This second sense of ‘social’ pivots on the idea that the social world is not only plural but also organized, i.e. it involves institutions and groups characterized by internal structures and roles. Typically, each individual agent simultaneously belongs to and acts within and across a multiplicity of such groups or collectives, ranging from family to professional, cultural, economical or political groups and organizations¹. 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. Such a vocabulary will only be successful if the chunks of

knowledge contained in the overall structure can be easily isolated, tested on their own, and updated. This is exactly what can be done in the structure presented in the following sections.

The rest of the paper is structured as follows. Section 2 presents, in an informal manner, some of the classical examples discussed in the philosophical literature, along with some less classical examples which have sprung to our minds while thinking about collective intentionality. In the same informal way, this section also anticipates our main theses on the topic. Section 3 describes the methodological backbone of our investigation. As a matter of fact, our reference scientific community is that of Applied Formal Ontology – a ‘joint venture’ of Artificial Intelligence and Philosophy, which provides formal accounts of large chunks of human knowledge for use in software applications. We begin with providing indications to the (unacquainted) reader about both the sense in which the term ‘ontology’ is used in our field and the specific ontologies adopted here to conduct our investigation: the Descriptive Ontology for Linguistic and Cognitive Engineering (Masolo, Gangemi, Guarino, Oltramari, & Schneider, 2003), Descriptions and Situations (Gangemi & Mika, 2003), and some extensions of these two which have been developed in the context of the DOLCE plus D&S Plan Ontology (Gangemi, Catenacci, Lehmann, & Borgo, 2004). Section 4 provides a formal-ontological account of the notion of collection in terms of what defined in Section 3, along with a typology of collections. Section 5 provides a treatment of intentional collectives and of their typology in terms of the formally specified notions of collection, agent and plan. Finally, Section 6 draws some conclusions.

The reported work is part of our Laboratory’s research program dedicated to social ontologies. The modules (DOLCE, D&S, DDPO) 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.

2. The notions at stake

As stated in the Introduction, our main objective is to provide a treatment of the notion of intentional collective and use it to present a general formal framework for an ontology of social reality. Consequently, the focus of the whole paper is on collections and collectives considered as social entities. In this Section, we list some classical examples of candidate intentional collectives, along with other, less classical examples, and anticipate the way in which we have chosen to reconstruct the notions involved in their analysis. We do not claim that such reconstructions are the only possible ones. The emphasis is on the results obtainable through the methodology we follow, which consists in making explicit what is often taken for granted, thus providing the means for solidly founded discussions and (dis)agreements.

2.1 Our reference examples

In his well-known book The Construction of Social Reality, Searle discusses, among others, the following putative examples of collective action (Searle, 1990): 1) a group of people in a park running to a common shelter because it has suddenly started to rain; and 2) an outdoor ballet where the choreography calls for the entire corps de ballet to converge on a common shelter because, in the fiction, it has suddenly started to rain; and 3) businessmen having the same goal (i.e. each one pursuing his own selfish interests) as well as mutual beliefs about their respective intentions, but not cooperating or acting together; and 4) a football team trying to execute a pass play.

Although a discussion of Searle's stance is outside the scope of the present paper², it is worthwhile to note that, according to his analysis, only 2 and 4 of the above-listed examples are genuine cases of collectively intentional behaviour. Since Searle seems to take collective intentionality as a fundamental feature of social facts, examples 1 and 3 would depict cases which, apparently, do not even pertain to social reality at all. From our point of view, however, the main point is that each of the examples portrays a situation which partly overlaps with, and partly differs from, any of the others.

In order to enlarge the basis of our reasoning, we have added to this first set the following sublist of candidate cases of collective social action: 5) Nazi Germans as possessed by a self-destructive unconscious desire (according to a subsequent psycho-historical reconstruction); and 6) CIA agents executing orders into a setting about which they are informed only 'on a strictly need-to-know basis'; and 7) the actors of an organization (e.g. an oil company) which, in addition to its 'constitutive' plan, plays a role in further, less transparent plans (e.g. fuelling civil wars in oil areas like African countries)³; and 8) fans in a stadium performing the so-called 'ola' (wave); and 9) the human agent seen as a collection of co-existing self systems (a case implying sub-agentive collectives); and 10) the employees/workers in a SAP workflow, or a 'Ford-style' production line.

In each of the listed examples there are both a plurality of agents and (at least) one goal. We believe that agents and goals are held together by plans. Intuitively, a plan is something that coordinates means to ends. Usually, plans are conceived by at least one of the involved agents. For us, however, a plan is such even if the conceiving agent is not involved in it. In some of the above-listed cases, the plan has not been conceived by, nor communicated to, any of the individual agents it involves, but eventually inferred or reconstructed a posteriori by another, not

necessarily involved agent. In other cases, the plan has been conceived by an involved agent, but the other involved agents are only partially, or even totally, unaware of it. In still other cases, finally, the plan is equally known to all the agents involved, but as a result of different mechanisms or reasons, ranging from spontaneous convergence (common motivation without previous agreement) to explicit agreement proper.

2.2 Our thesis: collectives are collections of agents

Provided that the notion of agent, as well as all the other notions introduced in this anticipation, will be discussed at further length in the following sections, let us state that we consider collectives to be collections of agents. Hence, each of the groups of people in the above-listed examples, as well as – let’s say – a pack of hunting wolves, will be considered here as a collective.

Collections, in turn, are considered to be social objects that (generically) depend on their members; consider, for instance, a collection of books in a library, which remains the same entity even if some books are lost and others acquired over time. Collections depend also (specifically) on the role(s) played by their members. Consider, for example, the constellation of Orion. Should the role ‘being a member of Orion’ cease to exist, the Orion constellation would disappear as well.

Another crucial feature of collections is that they must be covered by at least one role; consider, for instance, a collection of (not further specified) bones, where ‘being a bone’ is the one and only role played by the members. Collections, however, can also be (and usually are) characterized by further roles; consider, for instance, the collection of different (cutting, pasting, etc.) machines in a factory.

Collections, finally, are unified by ‘theory-like’ entities that we call descriptions, which contain and specify said roles. Intentional collectives, for instance, are unified by the kind of descriptions that we call plans (with characterizing roles), as in the case of the staff of a publishing house – including, e.g., a project manager, some assistants, some consultants, and several editors and authors – which is working at the production of a textbook.

3. Background concepts

Before getting to the heart of the matter, we introduce in this section the formal ontological apparatus used in our treatment of collections, collectives and related typologies. We employ three ontologies: the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE, in the following) (Masolo et al., 2003), the Descriptions and Situations ontology (D&S) (Gangemi & Mika, 2003), and the DOLCE plus D&S Plan Ontology (DDPO) (Gangemi et al., 2004). 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.

3.1 DOLCE

DOLCE is a foundational ontology of particular (as opposed to universal) entities. It is a formalized structure of very basic categories⁴, conceived as conceptual containers and applied in 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 comprises possibilia, i.e. possible, not only actual particulars, so that we are allowed to talk of particulars that are postulated by existentially

quantified variables, even if said particulars are not explicitly introduced in the model (Masolo et al., 2003).

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

Endurants 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 objects⁵. Social objects are a kind of non-physical objects which generically depend also on communication events and can be either agentive or non-agentive⁶.

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. 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 Figure 1, we show DOLCE by means of a UML class diagram, assuming a description logic-like semantics (Baader & 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..*).

3.2 D&S

The second formal ontological tool that we employ in this paper is D&S (Gangemi et al., 2004; Gangemi & Mika, 2003; Masolo et al., 2004), an apparatus conceived with the purpose of extending other (possibly, but not exclusively, foundational) ontologies. For instance, when using 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. 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 system⁷.

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 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:

According to the first sense of ‘social entity’ presented in the Introduction, we can minimally characterize a **Social Object** as a non-physical object that generically depends on agents and communication:

$$(A1) \quad \text{SocialObject}(x) \rightarrow \text{NonPhysicalObject}(x) \wedge \exists y, z. \text{GenericallyDependsOn}(x, y) \wedge \text{GenericallyDependsOn}(x, z) \wedge \text{AgentivePhysicalObject}(y) \wedge \text{CommunicationEvent}(z)$$

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

$$(A2) \quad \text{Description}(x) \rightarrow \text{NonAgentiveSocialObject}(x)$$

$$(A3) \quad \text{Description}(x) \rightarrow \exists y. \text{AgentivePhysicalObject}(y) \wedge \text{GenericallyDependsOn}(x, y)$$

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 Section 5.2 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.:

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

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

$$(A4) \quad \text{SettingFor}(x, y) \rightarrow \text{Situation}(x) \wedge \text{Particular}(y) \wedge \neg \text{Situation}(y)$$

$$(A5) \quad \text{SettingFor}(x, y) \rightarrow \exists z. \text{Perdurant}(z) \wedge \text{SettingFor}(x, z)$$

The **time and space of a situation** are the time and space of the particulars in the setting⁸:

$$(A6) \quad \forall p, s, t1, t2. (\text{Perdurant}(p) \wedge \text{TimeInterval}(t1) \wedge \text{TimeInterval}(t2) \wedge \text{TemporalLocation}(p, t1) \wedge \text{TemporalLocation}(s, t2) \wedge \text{SettingFor}(s, p)) \leftrightarrow \text{Part}(t2, t1)$$

$$(A7) \quad \forall e, s, r1, r2. (\text{Endurant}(e) \wedge \text{SpaceRegion}(r1) \wedge \text{SpaceRegion}(r2) \wedge \text{SpatialLocation}(e, r1) \wedge \text{SpatialLocation}(s, r2) \wedge \text{SettingFor}(s, e)) \leftrightarrow \text{Part}(r2, r1)$$

(A6) and (A7) state that a situation has a temporal – respectively, spatial – location that is the mereological sum of the locations of the particulars in the setting. For example, the time of World War II might span from the German invasion of Poland in 1939 to the Yalta conference in 1945; its space might include most of the Earth surface. Hence, the setting relation is not temporalized, because the time of $\text{Setting}(x, y)$ can be inferred from the previous axioms.

A **Concept**, like a description, is a social object, which is **defined by** a description. Once defined, a concept can be **used in** other descriptions. The **classifies** relation relates concepts to particulars (and possibly even concepts to concepts) at some time. There are several kinds of concepts reified in D&S, the primary ones (**role**⁹, **course**, and **parameter**) being distinguished by the categories of particulars they classify in DOLCE:

- (A8) $Uses(x,y) \rightarrow Concept(x) \wedge Description(y)$
- (A9) $Defines(x,y) \rightarrow Uses(x,y)$
- (A10) $Classifies(x,y,t) \rightarrow Concept(x) \wedge Particular(y) \wedge TimeInterval(t)$
- (A11) $Concept(x) \rightarrow NonAgentiveSocialObject(x) \wedge \exists y. Defines(y,x) \wedge Description(y)$
- (D2) $Role(x) =_{df} Concept(x) \wedge \exists y,t. Classifies(x,y,t) \wedge \forall y,t. Classifies(x,y,t) \rightarrow Endurant(y)$
- (D3) $Course(x) =_{df} Concept(x) \wedge \exists y,t. Classifies(x,y,t) \wedge \forall y,t. Classifies(x,y,t) \rightarrow Perdurant(y)$
- (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:

- (A12) $Specializes(x,y) \rightarrow Role(x) \wedge Role(y)$
- (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 are other social objects defined by descriptions. Differently from concepts, however, figures are social **individuals** (either agentive or not) and do not classify particulars. Intuitively, the difference between social concepts and social individuals is given by a comparison between the role banker and an organization like The Bank of Italy:

- (A13) $Figure(x) \rightarrow SocialObject(x)$
- (A14) $Figure(x) \rightarrow \exists y. Description(y) \wedge Defines(y,x)$
- (A15) $Figure(x) \rightarrow \neg \exists y,t. Classifies(x,y,t)$

Examples of figures are organizations (like the FIAT Company), but also political-geographic objects (like Italy) and sacred symbols (like the Holy Host).

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 Section 5. 2 below:

- (D5) $\text{AgentiveFigure}(x) =_{df} \text{Figure}(x) \wedge \text{AgentiveSocialObject}(x) \wedge \exists y. \text{Description}(y) \wedge \text{Conceives}(x,y,t)$
 (A16) $(\text{AgentiveFigure}(x) \wedge \text{Conceives}(x,y,t)) \rightarrow \exists z,t. \text{AgentivePhysicalObject}(z) \wedge \text{Conceives}(z,y,t)$

Agentive figures are established by a society or community; hence, they can act like a physical agent, can play roles, and so on. 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:

- (A17) $\text{AgentiveFigure}(x) \rightarrow \text{Figure}(x) \wedge \exists y,z,w,t. \text{Description}(y) \wedge \text{Role}(z) \wedge \text{Description}(w) \wedge y \neq w \wedge \text{Defines}(y,z) \wedge \text{Defines}(w,x) \wedge \text{Classifies}(z,x,t)$

Typical agentive figures are companies, organizations, institutions, and in general many socially constructed personas¹⁰. The notion of agentivity is taken here in the sense of being – directly or indirectly – able to conceive a description.

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):

- (A18) $\text{Deputes}(f,r,t) \rightarrow \text{Figure}(f) \wedge \text{Role}(r) \wedge \exists c,d,t. \text{Course}(c) \wedge \text{Description}(d) \wedge \text{Uses}(d,r) \wedge \text{Uses}(d,f) \wedge \text{Uses}(d,c) \wedge \text{ModalTarget}(r,c,t)$
 (A19) $\text{Deputes}(f,r,t) \rightarrow \exists r1,t1. \text{Role}(r1) \wedge \text{Classifies}(r1,f,t1)$

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

- (A20) $\text{ActsFor}(e,f,t) \rightarrow \exists t1. \text{Deputes}(f,r,t1)$
 (A21) $(\text{ParticipatesIn}(f,p,t) \wedge \text{AgentiveFigure}(f)) \rightarrow \exists e. \text{ActsFor}(e,f,t) \wedge \text{ParticipatesIn}(e,p,t)$

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 turner¹¹.

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 figure¹². Since figures are social objects, it can happen to find figures that act for other figures¹³.

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: **plays**, **sequences**, and **value for**, which apply to three different categories in DOLCE (Endurant, Perdurant, and Region, from (D2-4))¹⁴:

- (D6) $\text{Plays}(y,x,t) =df \text{Role}(x) \wedge \text{Classifies}(x,y,t)$
 (D7) $\text{Sequences}(x,y,t) =df \text{Course}(x) \wedge \text{Classifies}(x,y,t)$

$$(D8) \quad \text{ValueFor}(x,y,t) =_{df} \text{Parameter}(x) \wedge \text{Classifies}(x,y,t)$$

Roles or figures and courses are connected by relations expressing the **modalities** that, in given descriptions, (players of) roles and (representatives of) 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):

$$(A22) \quad \text{ModalTarget}(x,y,t) \rightarrow (\text{Role}(x) \vee \text{Figure}(x)) \wedge \text{Course}(y) \wedge \text{TimeInterval}(t)$$

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 agents¹⁵:

$$(A23) \quad \text{AttitudeTowards}(x,y,t) \rightarrow \text{ModalTarget}(x,y,t) \wedge \text{Task}(y) \wedge \forall e. \text{Classifies}(x,e,t) \\ \rightarrow \text{CognitiveAgent}(e)$$

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

$$(A24) \quad \text{RequisiteFor}(x,y,t) \rightarrow \text{Parameter}(x) \wedge \text{TimeInterval}(t) \wedge (\text{Role}(y) \vee \text{Course}(y))$$

Requisites are constraints over the values of the qualities of particulars. When a situation satisfies a description that uses parameters, durants 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 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 (see Figure 2):

(A25) $\text{InformationObject}(x) \rightarrow \text{SocialObject}(x)$

(A26) $\text{InformationObject}(x) \rightarrow \exists y,t. \text{Particular}(y) \wedge \text{RealizedBy}(x,y,t)$ ¹⁶

(A27) $\text{InformationObject}(x) \rightarrow \forall y,t. \text{ExpressedBy}(x,y,t) \rightarrow \text{Description}(y)$

(A28) $\text{Description}(x) \rightarrow \exists y,t. \text{InformationObject}(y) \wedge \text{ExpressedBy}(x,y,t)$

The relevance of IOs for collectives will emerge in Section 5, when considering that some collectives can exist only because their unity criterion can be communicated, shared, discussed, etc.

3.3 DDPO: DOLCE+D&S Plan Ontology

DDPO specializes the concepts and relations defined in DOLCE and in D&S. Like the latter, DDPO has a very liberal domain, which includes physical and non-physical objects (social entities, mental objects and states, conceptualizations, information objects, constraints), events,

states, regions, qualities, and ‘constructivist’ situations. The intended use of DDPO is to specify plans (which are a type of descriptions) at a very abstract level and independently from existing resources. DDPO allows also for the specification of so-called tasks, namely the types of actions of which plans are made of (Gangemi et al., 2004).

The main predicates and axioms of DDPO will be given in Section 5.3, after a discussion of the notion of agentivity.

4. 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 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 & Whitehead, 1910; Zeman, 1982), mathematics (Dauben, 1979; Dugac, 1976), and - more recently - linguistics and formal semantics (Link, 1983; Marcus, 1993). Since our focus here is not on comparing our contribute to each and any of the existing proposals, we simply introduce our treatment of this issue by excluding some of the entities we do not aim at describing.

In ontology, collections are usually considered as plural entities, an open and debated topic in the field (cf. Simons, 1987). We face the problem 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, namely Zermelo-Fraenkel plus the axiom of Choice, and von Neumann-Bernays-Gödel; see Devlin, 1993; von Neumann, 1967): 1) a set is uniquely determined by its members, i.e. it changes when its members or its cardinality change (axiom of extension), while a collection is not, unless explicitly specified; and 2) any two sets can be summed forming a union (axiom of union), while this is not tenable for any two collections; and 3) sets do not need an identity criterion for members (axiom of specification does not apply to all sets), while collections do (there is at least one property P that is true for all members); and 4) sets can be empty or singletons, but no empty or singleton collections are allowed; and 5) (hyper)sets can be members of themselves (anti-foundation axiom), while collections do not; and 6) the elements of a set have no part or constitution relation to the set, while collection elements have; and 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); and 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 evolution¹⁷. It follows 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; and ii) we assume that

members have an explicit identity criterion; and 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.

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 object(s) applies (in explicit, varied ways) to a plurality of entities, a collection appears. When such complex ceases to be conceived by any agent, or stops being applied to a plurality, then the collection dies. General criteria on the lifecycle of social objects are given in (Masolo et al., 2004).

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

Since collections are considered here as cognitive or social objects, but they also depend on their members, their space-time behavior is peculiar. Collections can participate in actions or processes either ‘on a member basis’ or ‘on a whole basis’. For example, some cows step on a guy, and the guy recognizes a moving herd ‘stepping on him’: the herd steps on the guy ‘on a member basis’. An opposite example: in 1914, 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 members¹⁸.

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

4.1 Definition of 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.

$$(D9) \quad \text{Collection}(x) =_{df} \text{SocialObject}(x) \wedge \forall t. \text{PresentAt}(x,t) \rightarrow \exists y,z. \\ \text{GenericConstituentOf}(y,x,t) \wedge \text{GenericConstituentOf}(z,x,t) \wedge y \neq z \wedge \exists c. \text{Role}(c) \\ \wedge \forall w,t. \text{GenericConstituentOf}(w,x,t) \rightarrow \text{Classifies}(c,w,t)$$

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

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

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:

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

In order to comply with the common meaning of the notion, we provisionally (and contextually) assume that members of collections cannot be parts of the same endurant:

$$(A29) \quad \text{Collection}(x) \rightarrow \neg \exists z. \text{Endurant}(z) \wedge \neg \text{ArbitrarySum}(z) \wedge \forall y,t. \text{Membership}(y,x,t) \\ \rightarrow \text{ProperPartOf}(y,z,t)$$

This rules out, for instance, that a sum of the corners of one (and the same) room, or a set of shady places in one (and the same) garden, could count as collections.

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):

$$(D12) \quad \text{Covers}(r,c) =_{df} \text{Role}(r) \wedge \text{Collection}(c) \wedge \forall w,t. \text{Membership}(w,c,t) \rightarrow \\ \text{Classifies}(r,w,t)$$

Notice that, due to the necessary presence of (at least) one covering role, all collections are homogeneous in a broad and trivial sense. However, we can refine and restrict the property of **homogeneity** on purely logical grounds, by saying that a collection is homogeneous (at a time t) if all members are classified by one covering role under the same leaf type in the current ground ontology:

$$(A30) \quad (\text{Homogeneous}(x,t) \wedge \text{Collection}(x)) \leftrightarrow \forall y,t. \text{Membership}(y,x,t) \rightarrow (\varphi(y) \wedge \neg \exists \psi \forall z. \psi(z) \rightarrow \varphi(z))$$

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 $\text{Saxophone}(x)$, $\text{Drum}(x)$, or $\text{Guitar}(x)$; all the endurants that constitute our collection are simply $\text{NonAgentivePhysicalObjects}$ (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.

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 & Nunez, 2000), which we represent as a description.¹⁹

- (S1) Description(ContainmentSchema)
- (S2) Role(Container)
- (S3) Role(Containee)

- (S4) Defines(ContainmentSchema, Container)
 (S5) Defines(ContainmentSchema, Containee)
 (A31) $\forall c,t. \text{Classifies}(\text{Container},c,t) \rightarrow \text{Collection}(c)$
 (A32) $\forall e,t. \text{Classifies}(\text{Containee},e,t) \rightarrow \text{Endurant}(e)$
 (A33) $\text{Collection}(c) \rightarrow \text{Covers}(\text{Containee},c)$

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.

4.2 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** – i.e., for instance, a collection of saxophones, or a mass of lymphocytes²⁰ – is a collection having only covering roles (for the definition of Characterizes, see (D18) below):

- (D13) $\text{SimpleCollection}(c) =_{df} \text{Collection}(c) \wedge \exists r. \text{Covers}(r,c) \wedge \neg \exists s. \text{Characterizes}(s,c)$

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) $\text{MaximallyGenericCollection}(c) =_{df} \text{SimpleCollection}(c) \wedge \neg \exists r. \text{Covers}(r,c) \wedge r \neq \text{Containee}$

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)
 (A34) $\forall e,t. \text{Classifies}(\text{Containee\#DB},e) \rightarrow \text{DinosaurBone}(e)$
 (D15) $\text{DinosaurBonesCollection}(c) =_{df} \text{Collection}(c) \wedge \text{Covers}(\text{Containee\#DB},c)$

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:

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

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):

$$(D17) \text{ ParametrizedCollection}(c) =_{df} \text{SimpleCollection}(c) \wedge \exists r, p, v, t. \text{Covers}(r, c) \wedge \\ \text{RequisiteFor}(p, r, t) \wedge \text{ValueFor}(v, p) \wedge \forall e. \text{Membership}(e, c, t) \rightarrow \exists q. \\ \text{InheresIn}(q, e, t) \wedge \text{Q-Location}(q, v)^{21}$$

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)²².

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:

$$(D18) \text{ Characterizes}(r, c) =_{df} \text{Role}(r) \wedge \text{Collection}(c) \wedge \exists e, f, o, s, t. (\text{Figure}(o) \vee \\ \text{Description}(o) \vee \text{Collection}(o)) \wedge \text{Role}(s) \wedge e \neq f \wedge r \neq s \wedge \text{Membership}(e, c) \wedge \\ \text{Membership}(f, c) \wedge (\text{Uses}(o, r) \vee \text{Deputes}(o, r) \vee \text{Covers}(r, o)) \wedge (\text{Uses}(o, s) \vee \\ \text{Deputes}(o, s) \vee \text{Covers}(s, o)) \wedge \text{Classifies}(r, e, t) \wedge \text{Classifies}(s, f, t) \\ (T2) \text{ Characterizes}(r, c) \rightarrow \exists s. \text{Role}(s) \wedge r \neq s \wedge \text{Characterizes}(s, c) \\ (D19) \text{ OrganizedCollection}(c) =_{df} \text{Collection}(c) \wedge \exists r, s. \text{Characterizes}(r, c) \wedge \\ \text{Characterizes}(s, c) \wedge r \neq s$$

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

$$(A35) \text{ Collection}(c) \rightarrow \exists d. \text{Description}(d) \wedge \text{SpecificallyDependsOn}(c, d)$$

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 (A37-39), and is not temporalized, since changing the description (differently from changing some members) creates a new collection:

- (A36) $\text{Unifies}(x,y) \rightarrow \text{Description}(x) \wedge \text{Collection}(y)$
 (A37) $\text{Covers}(x,y) \rightarrow \exists d. \text{Description}(x) \wedge \text{Defines}(d,x) \wedge \text{Unifies}(d,y)$
 (A38) $\text{Characterizes}(x,y) \rightarrow \exists d. \text{Description}(x) \wedge \text{Defines}(d,x) \wedge \text{Unifies}(d,y)$
 (A39) $(\text{Characterizes}(x,y) \wedge \exists f. \text{Deputes}(f,x)) \rightarrow \exists d. \text{Description}(x) \wedge \text{Uses}(d,f) \wedge \text{Unifies}(d,y)$

From (A11), (D2), (D9), (D12) and (A36) we can derive that a collection must be unified by at least one description, which provides to said collection its unity criterion:

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

We can imagine roles that are used by, deputed by, or that cover more than one description, figure, or collection²³. 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 survive change of some members. This behaviour ‘corresponds to’ the extensional/intensional nature of classes. For instance, if my collection of saxophones loses one member, it is still my collection of saxophones, because it respects the intensional criterion of being saxophones (possibly of a certain type).

On the other hand, collections can change identity while preserving the same members. What about, for instance, your collection of saxophones, which is covered or characterized by the same roles as mine? They result to be identical under the sole intensional criterion suggested by the solution above. In this case, we must postulate additional constraints, like exact restrictions on role playing. For instance, my/your collection will require that a role owner be played 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’.

Let us recapitulate here the network of relations among the main notions introduced thus far. A concept is defined by a description and can classify some particular (a role being a concept that classifies only endurants), while a figure is defined by a description, but cannot classify any particular, and must act by means of something else. A collection, on the other hand, is not directly defined by a description, and cannot classify any particular, but has members that are classified by at least one and the same role. 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 directly defined by a description.

5. Towards collectives

Now that we have an explicit notion of collection, we want to use it for making a step towards the definition of collectives as social entities. Despite the fact that traditional literature on collective intentionality is usually not committed to this kind of entities – preferring, instead, the notion of ‘social group’ – our move finds some theoretical support in the literature from various other fields. For instance, Sartre (1982) considers collectives as intermediate entities

between collections and social groups. A similar view has been entertained, too, both in sociology (French, 1984) and linguistics (Borschev & 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 (Bratman, 1992; Gilbert, 1992; Searle, 1990, 1995; Tuomela, 2003b). Moreover, the classical literature is also strongly committed to the notion of ‘we-mode’: in a social group, agents think or speak about themselves in terms of ‘we’. Some of the authors (Gilbert, 1992; Searle, 1990, 1995) consider this ‘we’ as a primitive; others (Bratman, 1992; Castelfranchi, 2003; Tuomela, 2003b), on the contrary, believe in the reducibility of this we-modality to a complex composition of I-modalities. Our position, however, is that this sort of psychological notions are not strictly necessary in order to define collectives.

Our own definition of collectives is built around a ‘descriptive’ interpretation of the notions of intentionality (sec. 5.1), agent (sec. 5.2), and – especially – plan (sec. 5.3). This commitment to the notion of plan gives our approach a teleological flavour. In this sense, we comply with some of the classical work in Philosophy and AI (Bratman, 1992; Cohen & Levesque, 1990), which assumes that the kind of rationality usually guiding actors in a society is a means-end rationality, and that the latter plays a crucial role in many contexts of (contemporary) productive societies²⁴. 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. 5.4 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.

5.1 Our approach to intentionality

Intentionality is still a debated notion in philosophy. From a historical point of view, the first modern account of intentionality is due to Brentano (Brentano, 1924), who gave new life and meaning to the medieval notion of intentio²⁵ 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”²⁶. This ordinary use of ‘intention’, on the other hand, seems closer to that typically made in Belief-Desire-Intention approaches (BDI in the following), where ‘intentions’ (as representations of the goals an agent is committed to achieve) are considered to be the third type of mental states which, together with beliefs and desires, plays a crucial role in the modeling of agent behavior (Ferrario & Oltramari, 2004).

Providing an ontology of mental states, however, is definitely beyond the aim of this paper. Although a first move towards such objective has actually been done in our laboratory (Ferrario & Oltramari, 2004), and although our ontology of descriptions should ideally ‘correspond’ to an

ontology of mind, there is still not enough agreement either in Cognitive Sciences or Philosophy of Mind on the nature of mental entities, and the currently available primitives are not sufficient for developing typologies of mental states and their mutual relations – hence, to handle intentionality and intentions in formal ontological terms. In the DOLCE+ framework, we will consider the descriptive equivalent of a type of mental states which undoubtedly include intentions, namely plans. Due to the above-mentioned lack of a sufficiently developed ontology of mind, however, we cannot provide a one-to-one correspondence between the two ontologies as yet. Therefore, we will characterize intentional agents and collectives within our ontology of descriptions alone.

5.2 Agents

Our characterization of agentivity, however, takes into account some of the main features attributed to (intentional) agents in the philosophical, AI and (mostly) BDI literature. As reported in (Ferrario & 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 descriptions²⁷. 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:

$$(A40) \text{ Conceives}(x,y,t) \rightarrow \text{GenericallyDependsOn}(y,x) \wedge \text{Agent}(x) \wedge \text{Description}(y)$$

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:

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

In simple words, this first level of agentivity is defined in (D20) in a wide sense as implying conception (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:

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

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

$$(A41) \quad (\text{Conceives}(x,y,t) \wedge \text{AgentiveSocialObject}(x)) \rightarrow \exists z,t. \\ \text{CognitiveAgentivePhysicalObject}(z) \wedge \text{GenericallyDependsOn}(x,z) \wedge \\ \text{Conceives}(z,y,t)$$

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:

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

We also introduce a restricted class for cognitive agents:

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

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

$$(A42) \quad \text{Creates}(x,y) \rightarrow \text{CognitiveAgent}(x) \wedge \text{Description}(y) \wedge \\ \text{SpecificallyDependsOn}(y,x) \wedge \exists t. \text{Conceives}(x,y,t) \wedge \neg \exists x',t'. t' < t \wedge \\ \text{Conceives}(x',y,t')$$

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

$$(A43) \quad \text{Adopts}(x,y,t) \rightarrow \text{Conceives}(x,y,t) \wedge \text{CognitiveAgent}(x) \wedge \text{Description}(y) \wedge \exists z. \\ \text{CognitiveAgent}(z) \wedge \text{Creates}(z,y)$$

$$(A44) \quad \text{Adopts}(x,y,t) \rightarrow \exists t1. >(t1,t) \wedge \text{Conceives}(x,y,t1)$$

5.3 Plans

Before discussing our typology of collectives, we introduce here some axioms for plans²⁸.

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:

- (A45) $\text{Plan}(x) \rightarrow \text{Description}(x)$
- (A46) $\text{Plan}(x) \rightarrow \exists y, t. \text{Conceives}(y, x, t) \wedge \text{CognitiveAgent}(y)$
- (A47) $\text{Plan}(x) \rightarrow \exists y. \text{Task}(y) \wedge \text{Uses}(x, y)$
- (A48) $\text{Plan}(x) \rightarrow \exists c. (\text{Role}(c) \wedge \forall a, t. \text{Classifies}(c, a, t) \rightarrow \text{Agent}(a)) \wedge \text{Uses}(x, c)$
- (A49) $\text{Plan}(x) \rightarrow \exists g. \text{Goal}(g) \wedge \text{ProperPart}(x, g)$

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 (possibly a different agent from the one initiating the task) implies a 'desire attitude' towards them:

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

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:

- (A50) $\text{ConstitutiveDescription}(x) \rightarrow \text{Description}(x)$
- (T4) $\forall x, f. (\text{Plan}(x) \wedge \text{Figure}(f) \wedge \text{Uses}(x, f) \wedge \neg \text{Defines}(x, f)) \rightarrow \exists y. \text{ConstitutiveDescription}(y) \wedge \text{Defines}(y, f)$
- (T5) $\forall x, f. (\text{Plan}(x) \wedge \text{Figure}(f) \wedge \text{Defines}(x, f)) \rightarrow \exists y. \text{ConstitutiveDescription}(y) \wedge \text{ProperPart}(x, y) \wedge \text{Defines}(y, f)$

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:

(A51) $\text{Desire}(x) \rightarrow \text{Description}(x)$

(A52) $\text{Desire}(x) \rightarrow \exists y, t. \text{Conceives}(y, x, t) \wedge \text{CognitiveAgent}(y)$

(A53) $\text{Desire}(x) \rightarrow \exists y, z. (\text{Role}(y) \wedge \forall a, t. \text{Classifies}(y, a, t) \rightarrow \text{Agent}(a)) \wedge \text{Course}(z) \wedge \text{Uses}(x, y) \wedge \text{Uses}(x, z) \wedge \text{DesireTowards}(y, z, t)$

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

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

(D26) $\text{Goal}(x) =_{df} \text{Desire}(x) \wedge \exists p. \text{Plan}(p) \wedge \text{ProperPart}(p, x)$

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):

(D27) $\text{MainGoal}(p1, x) =_{df} \text{ProperPart}(p1, x) \wedge \text{Plan}(p1) \wedge \text{Goal}(x) \wedge \neg \exists p2. \text{Plan}(p2) \wedge \text{ProperPart}(p1, p2) \wedge \text{ProperPart}(p2, x)$

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

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

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

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

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

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

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

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

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

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

$$(D31) \text{ PlanExecution}(x) =_{df} \text{Situation}(x) \wedge \exists y. \text{Plan}(y) \wedge \text{Satisfies}(x,y) \wedge \exists t. \text{PresentAt}(y,t) \wedge \neg \text{PresentAt}(x,t)$$

Subplan executions are parts of the whole plan execution:

$$(A54) \quad \forall p1,p2,s1,s2. (\text{Plan}(p1) \wedge \text{Plan}(p2) \wedge \text{ProperPart}(p1,p2) \wedge \text{Satisfies}(p1,s1) \wedge \text{Satisfies}(p2,s2)) \rightarrow \text{ProperPart}(s1,s2)$$

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

$$(D32) \quad \text{GoalSituation}(x) =_{\text{df}} \text{Situation}(x) \wedge \exists y. \text{Goal}(y) \wedge \text{Satisfies}(x,y)$$

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

$$(A55) \quad \text{GoalSituation}(x) \rightarrow \forall y,p,s. (\text{Goal}(y) \wedge \text{Satisfies}(x,y) \wedge \text{Plan}(p) \wedge \text{ProperPart}(p,y) \wedge \text{Satisfies}(s,p)) \rightarrow \neg \text{ProperPart}(s,x)$$

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:

$$(T7) \quad \forall p1,p2,s1 \neg \forall s2. (\text{Plan}(p1) \wedge \text{Plan}(p2) \wedge \text{ProperPart}(p1,p2) \wedge \text{Satisfies}(p1,s1) \wedge \text{Satisfies}(p2,s2)) \rightarrow \text{ProperPart}(s1,s2)$$

This makes it possible to account for the execution of plans containing abort or suspension conditions (the plan would be satisfied even if the goal has not been reached, see below), and for cases of incidental satisfaction, when a situation satisfies a goal without being intentionally planned (but anyway desired).

5.4 Collectives

As anticipated in Section 2.2, we define a collective as a collection of agents:

$$(D33) \quad \text{Collective}(c) =_{\text{df}} \text{Collection}(c) \wedge \forall x,t. \text{Membership}(x,c,t) \rightarrow \text{Agent}(x)$$

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

Similarly to what we have done for building our typology of collections, we distinguish simple and organized collectives. On the other hand, we need a finer-grained set of criteria for figuring out where collective action comes from. To this purpose, we make use of plans.

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 emerges 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 action²⁹.

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

A preliminary typology of collectives is introduced (sec. 5.5) 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 two notions are introduced: **bringing about** and **control**.

Agents bring about³⁰ 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, Sagri, Steve, & Tiscornia, 2003):

$$(A56) \text{ BringsAbout}(x,y) \rightarrow \text{CognitiveAgent}(x) \wedge \text{Collective}(y) \wedge \exists z. \text{Plan}(z) \wedge \text{Creates}(x,z) \wedge \text{Unifies}(z,y)$$

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

(A57) $\text{Controls}(x,y,t) \rightarrow \text{Agent}(x) \wedge \text{Collective}(y) \wedge \exists z,w,r. \text{Plan}(z) \wedge \text{Unifies}(z,y) \wedge \text{Plan}(w) \wedge \text{Conceives}(x,w,t) \wedge \text{Uses}(w,r) \wedge \text{Plays}(z,r,t)$

Trust is a further notion that could be used for refining our typology of collectives. In the framework we are proposing, trust could be directed at members, at those who brought about the collective, at controllers, or at plans. Trust could be about truth, validity, or plausibility of a description, as well as about known reliability, disposition to follow norms and plans, etc. At a very general level, we could preliminarily see it as another form of conceivability over descriptions. At a different – and more specific – level, trust could be treated as a special kind of social relationship (see below), linked to the notion of communities of trust. However, since the formal treatment of this complex issue is still at a very preliminary stage (cf. the large literature, e.g. Castelfranchi, 2000, 2001; Castelfranchi & Falcone, 2000), we postpone the inclusion of this notion to further developments of our Laboratory's work on social entities.

5.5 Typology of collectives

Collectives can be classified according to different property kinds. As for all collections, the first property 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 distinguished from ecological and social collectives, based on the (biological or social) properties ascribed to members.³¹

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. Obviously, such kind of collectives do not require the presence of a plan; they rather are unified by other kinds of descriptions like, e.g., scientific theories.

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, it is precisely the competence of members to conceive plans (in the generic sense outlined here of action schemas) that 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 can plans unify those collectives? These collectives 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:

(D34) $\text{SimpleCollective}(x) =_{\text{df}} \text{Collective}(x) \wedge \neg \exists y, r. \text{Description}(y) \wedge \text{Unifies}(y, x) \wedge \text{Role}(r) \wedge \text{Uses}(y, r) \wedge \text{Characterizes}(r, x)$

(D35) $\text{OrganizedCollective}(x) =_{\text{df}} \text{Collective}(x) \wedge \exists y, r. \text{Description}(y) \wedge \text{Unifies}(y, x) \wedge \text{Role}(r) \wedge \text{Uses}(y, r) \wedge \text{Characterizes}(r, x)$

However, differently from what applied to collections, we will use 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:

$$(D36) \text{ SimplePlannedCollective}(x) =_{df} \text{SimpleCollective}(x) \wedge \text{AgentiveSocialObject}(x) \\ \wedge \exists y, r. \text{Plan}(y) \wedge \text{Unifies}(y, x) \wedge \text{Role}(r) \wedge \text{Uses}(y, r) \wedge \text{Covers}(r, x)$$

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

$$(D37) \text{ IntentionalCollective}(x) =_{df} \text{Collective}(x) \wedge \text{AgentiveSocialObject}(x) \wedge \exists y, r. \\ \text{Plan}(y) \wedge \text{Unifies}(y, x) \wedge \text{Role}(r) \wedge \text{Uses}(y, r) \wedge \text{Characterizes}(r, x)$$

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:

$$(D38) \text{ Redescribes}(x, y) =_{df} \text{IntentionalCollective}(x) \wedge \text{SimplePlannedCollective}(y) \wedge \\ \text{ExtensionallyEquivalent}(x, y) \wedge \exists t1, t2. \text{PresentAt}(x, t1) \wedge \text{PresentAt}(y, t2) \wedge \\ >(t1, t2)$$

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 there being 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:

- (A58) $\text{IntentionalCollective}(x) \rightarrow \exists y. \text{Plan}(y) \wedge \text{Unifies}(y,x) \wedge \forall z. (z \neq y \wedge \text{Plan}(z) \wedge \text{Unifies}(z,x)) \rightarrow \text{Part}(y,z)$
- (D39) $\text{MaximalPlan}(x) =_{df} \text{Plan}(x) \wedge \exists y. \text{IntentionalCollective}(y) \wedge \text{Unifies}(x,y) \wedge \forall z. (z \neq x \wedge \text{Plan}(z) \wedge \text{Unifies}(z,y)) \rightarrow \text{ProperPart}(x,z)$

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

- (D40) $\text{NegotiatedPlan}(x) =_{df} \text{MaximalPlan}(x) \wedge \neg \exists p,q,y. \text{Plan}(p) \wedge \text{Plan}(q) \wedge \text{ProperPart}(x,p) \wedge \text{ProperPart}(x,q) \wedge \text{Conflicts}(p,q)$
- (D41) $\text{ConflictingPlan}(x) =_{df} \text{MaximalPlan}(x) \wedge \exists p,q,y. \text{Plan}(p) \wedge \text{Plan}(q) \wedge \text{ProperPart}(x,p) \wedge \text{ProperPart}(x,q) \wedge \text{Conflicts}(p,q)$

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.

Conflicts 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 et al., 2003). We take conflict and **superordination** as primitives here:

- (A59) $\text{Conflicts}(x,y) \rightarrow \text{Description}(x) \wedge \text{Description}(y)$
- (A60) $\text{Superordinates}(x,y) \rightarrow \text{Conflicts}(x,y)$

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

$$(D42) \text{ StableIntentionalCollective}(x) \rightarrow \text{IntentionalCollective}(x) \wedge \exists y. \text{NegotiatedPlan}(y) \wedge \text{Unifies}(y,x)$$

$$(D43) \text{ UnstableIntentionalCollective}(x) \rightarrow \text{IntentionalCollective}(x) \wedge \exists y. \text{ConflictingPlan}(y) \wedge \text{Unifies}(y,x)$$

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:

$$(D44) \text{ EmergingCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \exists y. \text{Collective}(y) \wedge y \neq x \wedge \forall e,t. (\text{Membership}(e,y,t) \rightarrow \text{Membership}(e,x,t)) \wedge \exists p,t1,t2. \text{Plan}(p) \wedge \text{Unifies}(p,x) \wedge \neg \text{Unifies}(p,y) \wedge >(t2,t1) \wedge \text{PresentAt}(x,t2) \wedge \text{PresentAt}(y,t1)$$

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

Emerging collectives can either 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:

$$(D45) \text{ CasualCollective}(x) =_{df} \text{EmergingCollective}(x) \wedge \exists y. \text{Plan}(y) \wedge \text{Unifies}(y,x) \wedge \exists p,q,t1. \text{Plan}(p) \wedge \text{Plan}(q) \wedge \text{PropertPart}(y,p,t1) \wedge \text{PropertPart}(y,q,t1) \wedge \forall a,b,t. a \neq b \wedge \text{Conceives}(a,p,t) \leftrightarrow \text{Conceives}(b,q,t) \wedge \neg \text{Conceives}(a,q,t) \leftrightarrow \text{Conceives}(b,p,t) \wedge \neg \text{Conceives}(a,y,t) \wedge \neg \text{Conceives}(b,y,t)$$

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:

$$(D46) \text{ SpontaneousCollective}(x) =_{df} \text{EmergingCollective}(x) \wedge \exists y. \text{Plan}(y) \wedge \\ \text{Unifies}(y,x) \wedge \exists p,q,t1. \text{Plan}(p) \wedge \text{Plan}(q) \wedge \text{PropertPart}(y,p,t1) \wedge \\ \text{PropertPart}(y,q,t1) \wedge \forall a,b,t. \text{Conceives}(a,p,t) \leftrightarrow \text{Conceives}(b,q,t) \wedge \\ \neg \text{Conceives}(a,q,t) \leftrightarrow \text{Conceives}(b,p,t) \wedge \text{Conceives}(a,y,t1) \wedge \text{Conceives}(b,y,t1) \\ \wedge >(t1,t)$$

Another type of intentional collectives are those unified by plans that involve agentive figures. Based on (A18-20), 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 (see Figure 3).

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

$$(D47) \text{ MaximalAgencyCollective}(c) =_{df} \text{IntentionalCollective}(c) \wedge \forall x,t. \\ \text{Membership}(x,c,t) \rightarrow \exists f. \text{AgentiveFigure}(f) \wedge \text{ActsFor}(x,f,t)$$

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.

Given (D47) 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:

(A61) $\text{EmergingCollective}(x) \rightarrow \forall p. (\text{Plan}(p) \wedge \text{Unifies}(p,x)) \rightarrow \neg \exists y. \text{CognitiveAgent}(y) \wedge \text{BringsAbout}(y,x)$

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

(D48) $\text{DevisedCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \forall p. (\text{Plan}(p) \wedge \text{Unifies}(p,x)) \rightarrow \exists y. \text{CognitiveAgent}(y) \wedge \text{BringsAbout}(y,x)$

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):

(D49) $\text{GovernedCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \forall p. (\text{Plan}(p) \wedge \text{Unifies}(p,x)) \rightarrow \exists y,t. \text{CognitiveAgent}(y) \wedge \text{Controls}(y,x,t)$

(D50) $\text{UngovernedCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \forall p. (\text{Plan}(p) \wedge \text{Unifies}(p,x)) \rightarrow \neg \exists y,t. \text{CognitiveAgent}(y) \wedge \text{Controls}(y,x,t)$

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:

(D51) $\text{TransparentCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \forall e. \text{Membership}(e,x,t) \rightarrow \exists p,t. \text{Plan}(p) \wedge \text{Unifies}(p,x) \wedge \text{Conceives}(e,p,t)$

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

(D52) $\text{PartlyTransparentCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \neg \forall e. \text{Membership}(e,x,t) \rightarrow \exists p,t. \text{Plan}(p) \wedge \text{Unifies}(p,x) \wedge \text{Conceives}(e,p,t)$

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:

$$(D53) \text{ OpaqueCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \forall e,t. \text{Membership}(e,x,t) \rightarrow \\ \neg \exists p. \text{Plan}(p) \wedge \text{Unifies}(p,x) \wedge \text{Conceives}(e,p,t) \wedge \exists d. \text{Plan}(d) \wedge \text{ProperPartOf}(d,p,t) \\ \wedge \text{Conceives}(e,d,t)$$

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: e.g., goal sharing or adoption 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:

$$(D54) \text{ TransparentlyEmbracingCollective}(x) =_{df} \text{IntentionalCollective}(x) \wedge \forall e,t. \\ \text{Membership}(e,x,t) \rightarrow \exists p. \text{Plan}(p) \wedge \text{Unifies}(p,x) \wedge \text{AdoptsPlan}(e,p,t)$$

Shared adoption, too, is 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 and adoption, 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. theories, scripts, or rules.

Other types of collectives may be singled out based on whether a goal is conceived or adopted 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 the master role has rights towards certain tasks, 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**:

$$(D55) \text{ TemporaryCollective}(x) =_{df} \text{ Collective}(x) \wedge \exists y,z. \text{ Plan}(y) \wedge \text{ Schedule}(z) \wedge \\ \text{ Unifies}(y,x) \wedge \text{ Defines}(y,z)$$

Notice that, according to the typology – not a partition! –drafted so far for collectives (see Table 1 and Figure 4 for a summary), only the first of the classical examples listed in Section 2.1 – the group of people running to a common shelter because it has suddenly started to rain – does not involve an intentional collective. Due to the lack of characterizing roles in its unifying plan, it ranks rather as a case of simple-planned (and temporary) collective.

A table of the remaining classified examples that we are using to check the applicability of our typology of intentional collectives is proposed in Table 2.

6. 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 DDPO), 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 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? And 2) More specifically, how much development would be needed for the treatment of the notion of organization? Our taxonomy of plans and collectives allows, in principle, for the characterization of different types of organizations, from very simple to very complex ones. What else is needed for making this a viable option? And 3) How can our proposal contribute to handling unambiguous sharing of plans and negotiation of meaning? This is a hot topic in Semantic Web circles. And 4) Our modular approach to intentional collectives addresses a problem which is often overlooked in the theoretical literature: each different exemplar from a variety of collective entities relates differently to intentions. Could this line of research result in any relevant contribution to, for instance, distributed AI or to the methodology of sociological research? And 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.

References

- Baader, F., & Nutt, W. (2003). Basic description logics. In P. F. Patel-Schneider (ed.), The Description Logic Handbook: Theory, Implementation and Applications (pp. 43-95). Cambridge: Cambridge University Press.
- Biddle, B. J. (1979). Role Theory. Expectations, Identities, and Behaviors. New York: Academic Press.
- Borschev, V., & 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., & Pacheco, O. (2003). A Role Based Model for the Normative Specification of Organized Collective Agency and Agents Interaction. Autonomous Agents and Multi-Agent Systems, 6(2), 145-148.
- Castelfranchi, C. (1998). Modelling social action for AI agents. Artificial Intelligence, 103(1), 157-182.
- Castelfranchi, C. (2000). Founding Agent's 'Autonomy' On Dependence Theory. Paper presented at the ECAI-2000: 14th European Conference on Artificial Intelligence, Berlin, Germany.

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

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

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

Cocchiarella, N. (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., & Levesque, H. J. (1990). Intention is choice with commitment. Artificial Intelligence, 42, 213-261.

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

Ferrario, R., & Oltramari, A. (2004). Towards a Computational Ontology of Mind. Paper presented at the 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., Catenacci, C., Lehmann, J., & Borgo, S. (2004). Task taxonomies for knowledge content: EU 6FP METOKIS Project D07, <http://metokis.salzburgresearch.at>; updated version available at <http://www.loa-cnr.it/Publications.html>.

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

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

Gilbert, M. (1992). Social Facts. Princeton, New Jersey: Princeton University Press.

Gilbert, M. (1996). Living together : rationality, sociality, and obligation. Lanham, Md.; London: Rowman & Littlefield Publishers.

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

Jacob, P. (2003). Intentionality. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Fall 2003 Edition). <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. In K. Guilfooy (Ed.), The Cambridge companion to Abelard. 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., & Nunez, R. E. (2000). Where mathematics comes from : how the embodied mind brings mathematics into being. New York: Basic Books.

Link, G. (1983). The Logical Analysis of Plurals and Mass Terms: A Lattice-theoretical Approach. In A. von Stechow (Ed.), Meaning, Use and Interpretation of Language (pp. 302–323). Berlin: Walter de Gruyter.

Marcus, R. B. (1993). Classes, Collections, Assortments, and Individuals. In R. B. Marcus (ed.), Modalities: Philosophical Essays (pp. 90-100). New York: Oxford University Press.

Masolo, C., Gangemi, A., Guarino, N., Oltramari, A., & Schneider, L. (2003). WonderWeb Deliverable D18: The WonderWeb Library of Foundational Ontologies.

Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Ferrario, R., Gangemi, A., et al. (2004). Social Roles and their Descriptions. Paper presented at the Proceedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning, Whistler, Canada.

Pacherie, E. (2005). Is collective Intentionality really primitive?
<http://pacherie.free.fr/papers/Coll-intentionality-PDF.pdf>

Perler, D. (2003). Théories de l'intentionnalité au moyen âge. Paris: Vrin.

Pettit, P. (2003). Groups with Minds of Their Own. In F. F. Schmitt (Ed.), Socializing metaphysics: the nature of social reality. Lanham, Md. ; Oxford: Rowman & Littlefield Publishers.

- Rorty, A. O. (1986). Self-deception, akrasia and irrationality. In J. Elster (Ed.), The Multiple Self. Cambridge, UK: Cambridge University Press.
- Russell, B., & Whitehead, A. N. (1910). Principia Mathematica. Cambridge, UK: Cambridge University Press.
- Sartre, J. P. (1982). Critique of dialectical reason (Verso pbk. ed.). London ; New York, N.Y.: Verso/NLB.
- Scherer, A. G. (2003). Modes of Explanation in Organization Theory. In C. Knudsen (Ed.), The Oxford handbook of organization theory (pp. 310-344). Oxford: Oxford University Press.
- Searle, J. R. (1983). Intentionality. An Essay in Philosophy of Mind. Cambridge, UK: Cambridge University Press.
- Searle, J. R. (1990). Collective Intentions and Actions. In M. E. Pollak (Ed.), Intentions in Communication (pp. 401-415). Cambridge, Massachusetts: MIT Press.
- Searle, J. R. (1995). The Construction of Social Reality. New York: The Free Press.
- Searle, J. R. (2001). Rationality in Action. Cambridge, MA: MIT Press.
- Simons, P. (1987). Parts: a Study in Ontology. Oxford: Clarendon Press.
- Tuomela, R. (1995). The importance of us : a philosophical study of basic social notions. Stanford, 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. In F. Schmitt (Ed.), Socializing Metaphysics: The Nature of Social Reality. Lanham, Md: Rowman and Littlefield.

von Neumann, J. (1967). An Axiomatization of Set Theory. In J. Van Heijenoort (Ed.), From Frege to Gödel; a source book in mathematical logic, 1879-1931. Cambridge: Harvard University Press.

Weber, M. (1968). Economy and society; an outline of interpretive sociology. New York,: Bedminster Press.

Wooldridge, M. J. (2000). Reasoning about rational agents. Cambridge, Mass.; London: MIT Press.

Zeman, J. J. (1982). Peirce on Abstraction. The Monist, 65, 211-222.

Acknowledgments

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

Table 1. Our preliminary formal typology of collectives (schema)

COLLECTIVE

Simple (covered by a concept which is not defined in a plan with concepts characterizing the collective)

Type-based (covered by a concept which is not defined in a plan)

Genetic | Taxonomic | Epidemiological

Simple-planned (unified by a plan which defines covering, but no characterizing concepts for the members)

Organized

Intentional (unified by a plan which defines characterizing concepts for the members)

Stable vs. Unstable (based on negotiated vs. conflicting plans)

Devised vs. Emerging (based on s.o. bringing about the collective)

[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)

By internal structure of plans and/or related descriptions

Temporary (scheduled)

Table 2. Classified examples of intentional collectives. Transparent/opaque/obscure on goal (g), conception (c) or adoption (a).

Intentional Collective	T	agreement	devisal	transparency	control	structure
Common shelter ballet	I	stable	devised	transparent (g,c,a)	governed	structured
Selfish businessmen	I	unstable	emerging (casually)	transparent (g)	ungoverned	unstructured
Football team	I	stable	devised	transparent (g,c,a)	governed	structured
Self-destructive Nazis	I	unstable	emerging (spont.)	obscure (g,a)	ungoverned	unstructured
CIA agents	I	stable	devised	obscure (g,c,a)	governed	structured
Oil company (as partly cognizant maximal agency)	I	stable	devised	opaque (g,c,a)	governed	structured
Fans in a stadium (ola)	I	stable	emerging (spont.)	transparent (g,c,a)	ungoverned	unstructured
Co-existing self systems (or maximal agency of selves)	I	unstable	emerging	opaque (g,c,a)	(un)governed	structured
SAP workers's maximal agency	I	stable	devised	obscure (c,a)	governed	structured

Footnotes

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. (R. S. Cohen & Schnelle, 1986).

² A detailed presentation and philosophical criticism of Searle's theses on collective intentionality can be found in Pacherie (2005).

³ Alternatively, 'fuelling civil wars in oil areas' could be considered as a sub-plan which, although not explicitly specified in the 'constitutive' overall plan, or mission, of the oil company, it is actually implemented by some of its representatives.

⁴ Most of the ontologies mentioned in the following are available in various formal languages and formats on <http://www.loa-cnr.it>.

⁵ 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 5.2 below.

⁶ See <http://dolce.semanticweb.org> for more details.

⁷ An extension of DOLCE which does not use D&S, but introduces descriptions in order to treat social entities is presented in Masolo et al. (2004).

⁸ All 't' variables in the formulas denote time intervals.

⁹ In this paper a role is simply a concept that classifies an endurant. This definition is different from the one provided in Masolo et al. (2004), which relies crucially on the properties of anti-rigidity and foundation.

¹⁰ 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 & Pacheco (2003).

¹¹ 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).

¹² Cases of full delegation or representation, however, are quite unusual, and even prohibited in some legal contexts.

¹³ 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.

¹⁴ 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.

¹⁵ See Section 5.2 for more on agents, and Section 5.3 for more on tasks.

¹⁶ Information objects are reifications of pure information as social objects, hence they are assumed to be in space-time, and realized by some entity; see Gangemi et al. (2004).

¹⁷ We do not enter here the complex debate on the primacy of biology vs. society.

¹⁸ ‘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.

¹⁹ 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. Currently, cognitive schemata are only informally represented in the literature. Our Laboratory has an ongoing

research on the formal representation of cognitive schemata. The choice of representing schemata as descriptions lies in the similarity between the schema metaphor in cognition (simultaneous activation of multiple intentional systems, through their physiological counterparts), and in the reification semantics of D&S (relational patterns within a reified theory). However, besides the technically seamless formalization of cognitive schemata as very general ontological descriptions, experimental validation is still unavailable.

²⁰ Including masses among collections is not straightforward, since most traditional views take masses as continuous. However, the traditional criterion for membership identification, i.e. discreteness, is (implicitly) based on the perceivability of members. Since members can be either directly or instrumentally identifiable, or even inferrable, we need to enlarge the range of possible collections. On the other hand, a mass cannot be considered a collection when no member is identifiable or epistemically relevant. For instance, if we take a piece of gold without considering its constituent atoms, it cannot be a collection, it is just ‘stuff’.

²¹ Cf. Section 3.1 and Figure 1.

²² 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.

²³ 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.

²⁴ For an account of rationality and social reality, see Searle (2001), and also Pettit (2003).

²⁵ For a treatment of intentionality in medieval philosophy cf., for instance, Perler (2003).

²⁶ For an introduction to the different senses of ‘intentionality’, cf. also Jacob (2003).

²⁷ Notice that descriptions are expressed by some information object (see (A28)), which can be a natural-language proposition as well as some other (not necessarily language-like) encoding.

²⁸ Gangemi et al. (2004). See Section 3.3 above.

²⁹ 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.).

³⁰ BringAbout is used here in a different way with respect to deontic logic; see, e.g., Carmo & Pacheco (2003).

³¹ Notice that biological, ecological and social collectives can be extensionally equivalent.

Figure Captions

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

Figure 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.

Figure 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.

Figure 4: Our preliminary formal typology of collectives (diagram).

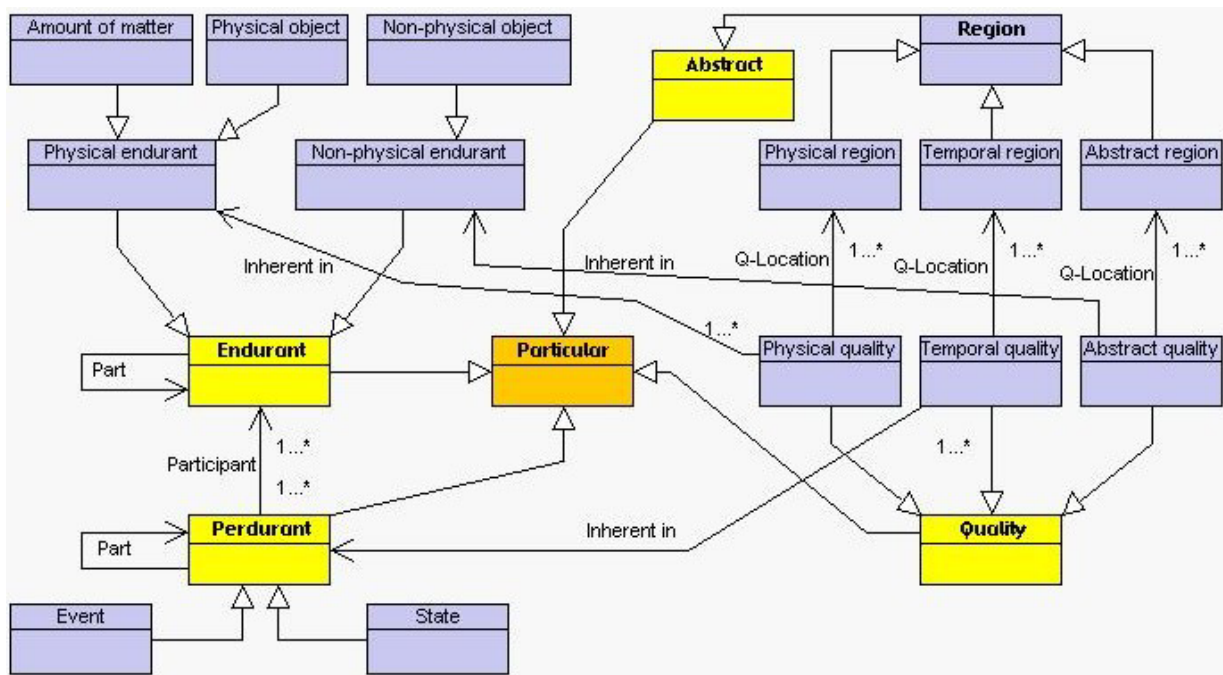


Figure 1. A UML class diagram showing the basic classes and relations of DOLCE.

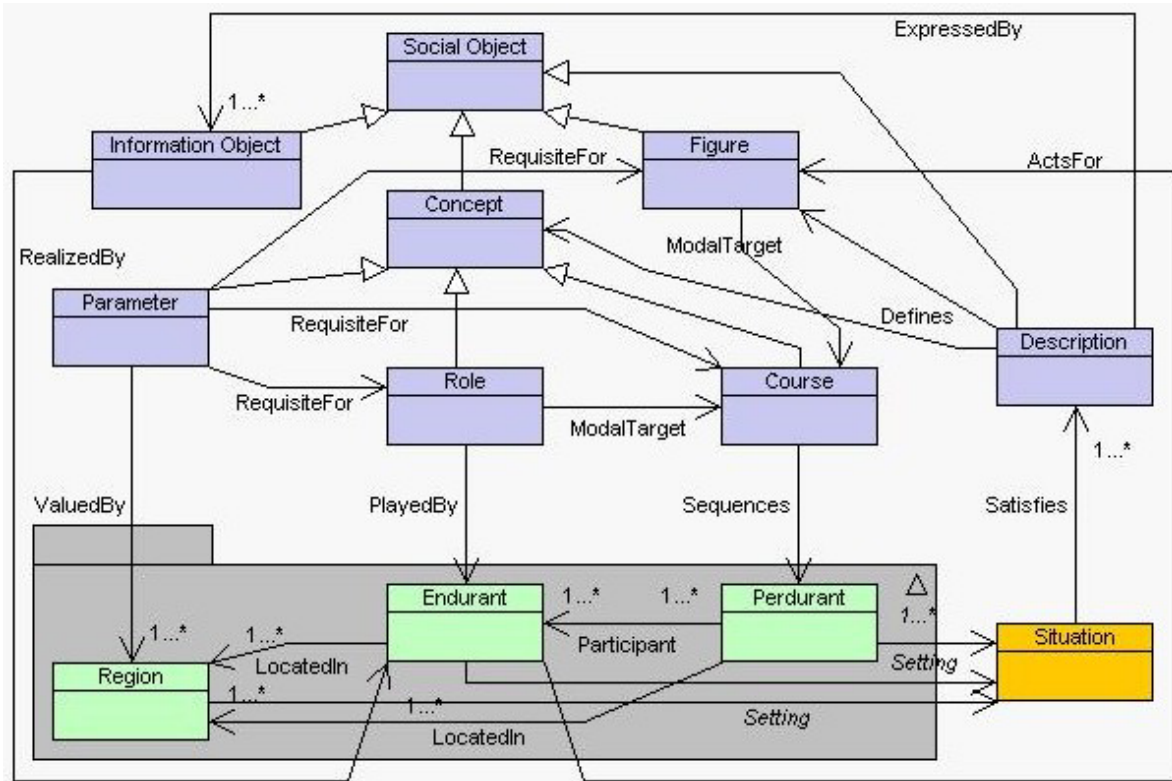


Figure 2. A UML class diagram for D&S. The lower part of the pattern (within the grey package) is called the ground ontology, the higher is called the descriptive ontology; a situation satisfies a description if the two parts match according to the axioms specified for the concepts defined by the description.

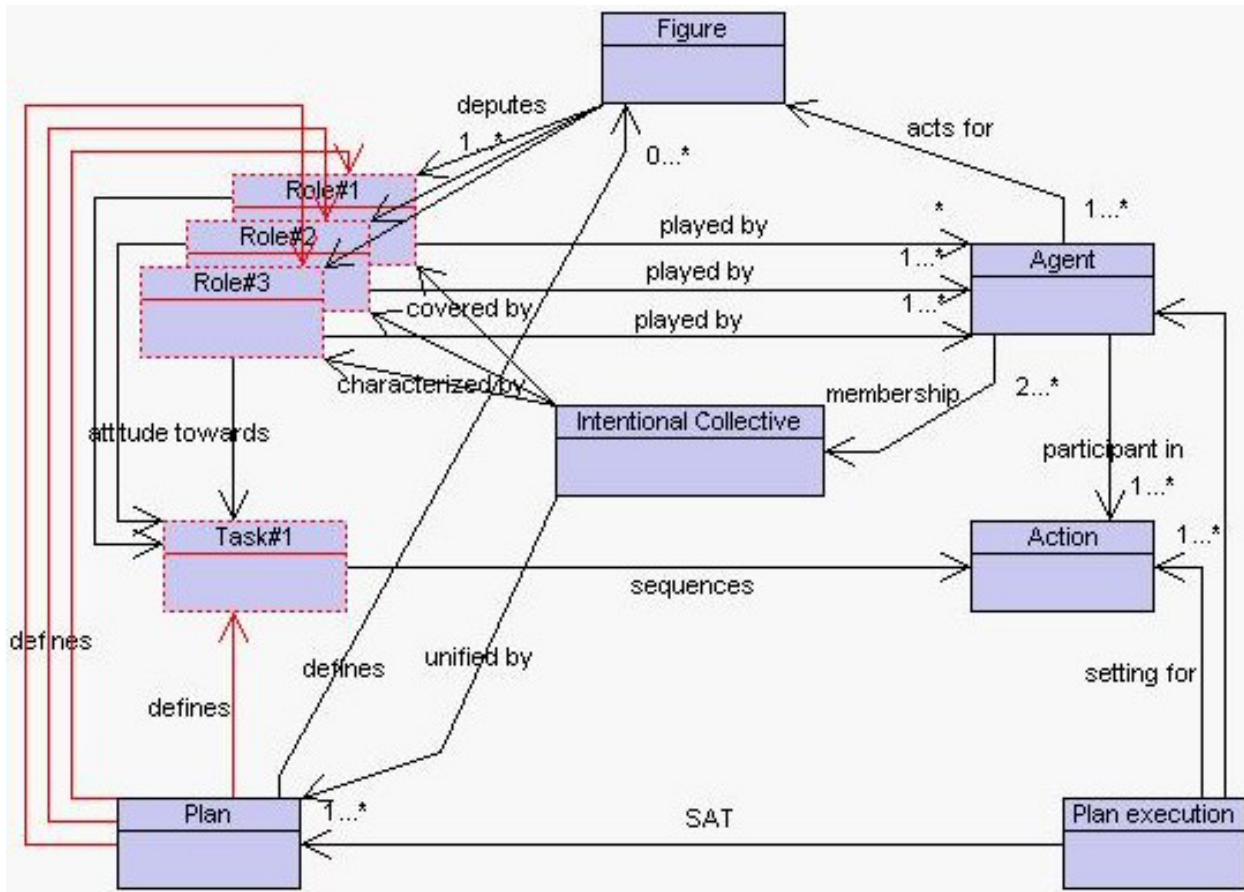


Figure 3. A UML class diagram showing a definition of a class of (possible) intentional collectives unified by a plan that defines three roles and one task.

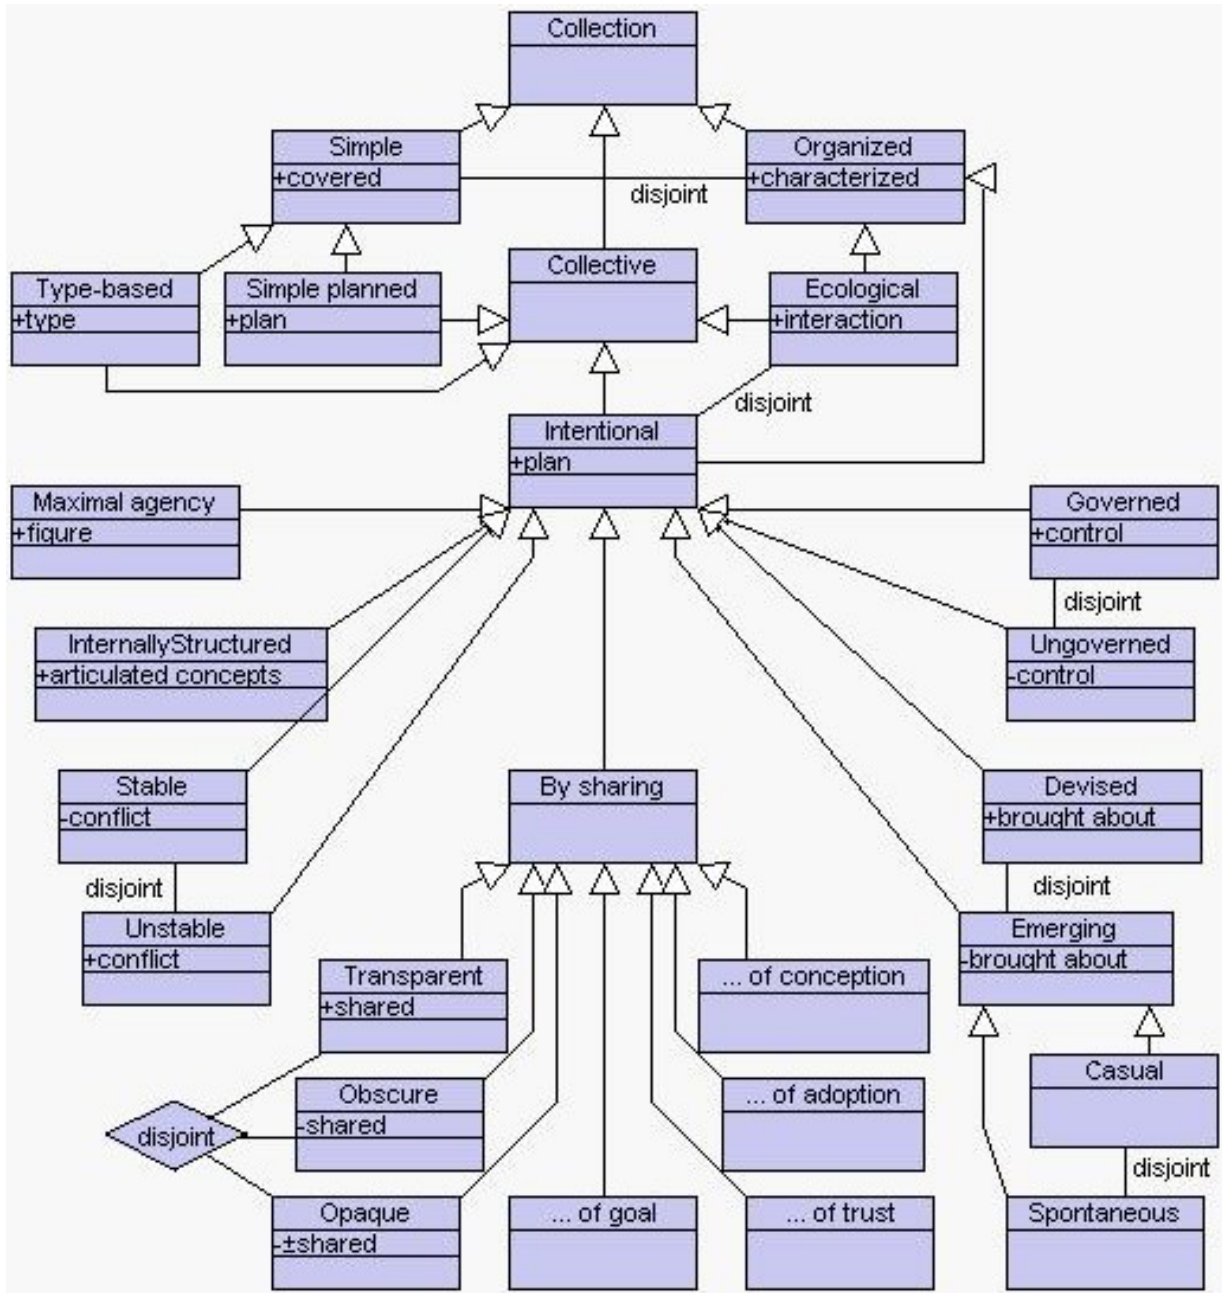


Figure 4. Our preliminary formal typology of collectives (diagram)