

Norms and plans as unification criteria for social collectives

Aldo Gangemi, Jos Lehmann, Carola Catenacci

Laboratory for Applied Ontology
Institute of Cognitive Science and Technology
Italian National Research Council, Rome, Italy
{aldo.gangemi, jos.lehmann, carola.catenacci}@istc.cnr.it

Abstract. Based on the formal-ontological paradigm of *Constructive Descriptions and Situations*, we propose a definition of social collectives that includes social agents, plans, norms, and the conceptual relations between them. We also propose a typology of social collectives, including *collection of agents*, *knowledge community*, *intentional collective*, and *intentional normative collective*. Our ontology, represented as a first-order theory, provides the expressivity to talk about the contexts (social, informational, circumstantial, and epistemic), in which collectives *make* and *produce* sense.

Keywords. Formal Ontology, Constructivism, Social Entities, Semantic Web

1 Introduction

In this article we lay down the basis for an integrated ontology of the mutual dependencies between agents, collectives, concepts, information, plans, and norms. The ontology has a constructive approach, and is represented as a first-order theory, as well as an OWL(DL) ([1]) ontology, for applications in the semantic web¹, and semantic web services domains (cf. [2]).

In previous work [3], we have treated some problems of *collective intentionality* by introducing a formal-ontological definition of the notion of *intentional collective*. Our approach pivoted on two general ideas. On the one hand, we investigated and formalized the grounds based on which we define a set of items as a *collection*, and collected items as *members* of a collection². On the other hand, we proposed a way of relating collections and their members to intentional and agentive notions³. According to our reconstruction, collections can be seen as social objects (as defined in [20]) that (generically) depend on their members at a certain time. This entails, for instance, that a collection of books in a library remains the same entity even if some books are lost and others acquired over time. Collections depend also (specifically) on the *role(s)* played by their members. Consider, for example, the constellation of Orion. Should the role ‘being a member of Orion’

¹ See e.g. the EU NeOn project site: <http://www.neon-project.org>

² There is a large and heterogeneous literature on collections and plural entities; in our work, we considered in particular [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14]

³ On this topic, we made reference to classical works such as [15], [16], [17], [18], [19]

cease to exist, the relative constellation would disappear too. Collections 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 the covering or characterizing roles of the collection.

Following this notion of collection, collectives in our proposal are collections of agents which are unified by the kind of descriptions that we call *plans*. The members of a collective are ‘held together’ by one plan, which specifies a goal and (one or more) covering or characterizing role(s). In order for a collective to be intentional, there must be a plan, and the agentive members of a collection must play the covering or characterizing roles of that plan. For instance, in our view, both a group of people running towards a common shelter because of a sudden storm [21], and a pack of hunting wolves are to be considered as examples of intentional collectives.

In this article, the proposal presented in [3] is updated and enriched under two respects. Firstly, the very foundations of the original proposal are profoundly restructured by a new paradigm, called Constructive Descriptions and Situations. Secondly, our definition of intentional collectives is extended with *normative elements*, as well as with the conceptual relations between such new normative elements and the original planning elements. This move provides us with the conceptual means to define collective entities like *collection of agents*, *knowledge community*, *intentional collective*, and *intentional normative collective*. Within this framework, an issue we address is how to *represent* in our formal ontological framework the influence that norms may have on plans. We are not concerned here with providing a full-fledged theory of these interactions. We rather want to set a formal-ontological basis for modeling such theories. In other words, we are aware of the fact that relevant work on norms, and possibly on their interactions with plans, may be found both in the legal-philosophical literature (e.g. [22], [23]), in the sociological literature (e.g. [1]), as well as in the multi-agent systems literature (e.g. [24]). Here though we introduce a minimal setup of formal distinctions between the types of interactions that norms may have with plans. Future work will refine such distinctions, possibly modelling existing proposals on that basis.

Section 2 provides a brief informal overview on our previous work, including Constructive Descriptions and Situations, and the ontologies of Plans, Norms, and Collectives. Section 3 presents a formalization of our proposal. Finally, section 4 draws some conclusions.

2 Constructive DnS and its extensions at a glance

In this section we informally introduce the ontological apparatus, on which our treatment of collectives is based. We start with a brief presentation of Descriptions and Situations (DnS), an ontology developed in [25], [26], [3]. We then present Constructive DnS, a restructured version of DnS proposed in [27] and in this paper. Finally, we provide a schematic introduction to our ontologies of Plans, Norms and Collectives.

2.1 Relations to previous work

In [3], we have provided a formal-ontological definition of the notion of *intentional collective*. Our approach there pivoted on two general ideas: on the one hand, we investigated and formalized 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 proposed a way of relating collections and their members to intentional notions. The work presented in [3] was based on three ontologies: the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) [28], the ontology of Descriptions and Situations (DnS) [25] and the ontology of Plans defined in [26].

In this article, the work described above is updated and enriched under two respects. Firstly, the very foundations of our original proposal are profoundly re-structured. The definition of collectives given here is not based anymore on a combination of DOLCE and DnS, but on a brand-new version of DnS, called Constructive DnS (hereafter, *c.DnS*). Secondly, our definition of intentional collectives is extended with *normative elements*. This provides us with the conceptual means to define typologies of normed intentional collectives.

In its original version, DnS is a formal tool that allows to extend other (possibly, but not exclusively, foundational) ontologies with a number of reified concepts and relations, thus making the extended ontology more expressive without making it computationally more complex. Suppose, for instance, that you want to use DnS to extend DOLCE⁴. The final result of this extension would be DOLCE+. DOLCE+ would consist, on the one hand, of DOLCE – which would play the role of ground ontology, i.e. the ontology that specifies the entities of a given domain irrespective of any possible epistemological status or concern. On the other hand, DOLCE+ would also consist of the DnS extension, that provides the formal means to specify the epistemological perspective from which the entities of the domain are considered. By way of example, suppose that such perspective is of legal nature. The DnS extension makes it possible to express the legal constraints imposed by norms and regulations on the domain of the ground ontology, i.e. to describe the entities of DOLCE (in particular, entities pertaining to social reality) under a legal perspective. In other words, in DOLCE+ it would become possible to describe a legal view on the behaviour of DOLCE's (social) entities according to a given legal system.

Two are the key-elements of DnS:

The distinction between descriptions and situations which allows to separate, within the same model of the legal domain, relations between ‘conceptual elements’ like laws, norms, regulations, crime types, etc. (which are all descriptions) from relations between ‘observable elements’ like legal facts, cases, states of affairs, etc. (which are all situations).

A reification mechanism which allows to have descriptions and situations in the same domain of quantification (i.e. at the same logical level) and to relate them by means of a reified relation of satisfaction. For instance, according to a DnS-based legal extension of a ground ontology like DOLCE, a case (a situation) satisfies a norm (a description). This means that the norm (i.e. the description and the concepts devised in it) classifies the entities of DOLCE,

⁴ The OWL-DL version of DnS combined with DOLCE can be loaded from <http://www.loa-cnr.it/ontologies/ExtendedDnS.owl>.

the ground ontology. This very classification gives rise to the case (i.e. the situation), which is a setting for the entities of the ground ontology that satisfy (i.e. are individually classified by) the concepts devised in the description. Both the case and the norm are part of the same domain of quantification. In other words, the DnS extension makes it possible to enhance the expressivity of the language of the ground ontology while keeping its complexity under control – the usual advantage of reification.

While DOLCE+ is being effectively used in several projects, specially in its version encoded in the Web Ontology Language (OWL), we have realized that the expressive power of the reification vocabulary in DnS can be axiomatized and reused beyond the scope of a specific foundational ontology (DOLCE or any other). Based on this working hypothesis, we have created a new ontology that applies the constructivist paradigm, and remains agnostic with respect to which foundational, core, or domain ontology should be used as a primary modelling framework for a knowledge base. We have named this ontology *Constructive Descriptions and Situations (c.DnS)*.

2.2 The constructive stance

Constructivism is the epistemological stance according to which reality and its structure are not given ‘as such’ for our minds to passively discover, but are rather actively constructed by cognitive agents in specific contexts and for specific purposes. This implies a rejection of naive interpretations of Aristotelian notion of ‘truth as correspondence’ (between constructs and chunks of ‘reality’), and a deep awareness of the historical and social nature of all kinds of knowledge. It does not imply, however, that we have to reject the idea that there are (physical, biological and cultural) constraints on the way the mind builds and manages its constructs, nor that the whole scientific enterprise is devoid of meaning. Rather, constructivism promotes a view according to which every scientific theory or model should be seen as a ‘tool’, which is the product of a specific ‘knowledge collective’⁵ and whose adequacy in representing and handling specific aspects of our interaction with the world has to be tested against actual usage and effectiveness, and always be open to revision (cf. [30]).

In cognitive sciences, in particular, this has led to see our mental representations as context-dependent (or ‘situated’) and action-oriented views on the world, relating only to those aspects of the environment which are salient for the perceiver/cognizer [31]. Moreover, focusing on the non-abstract nature of cognition has led to put a new emphasis on the ‘gestaltic’ aspects of representations, i.e. the need of taking into account “the interconnected whole that gives meaning to the parts” [32].

In current ontology research and engineering, however, epistemology is usually left out of the picture. Viewpoints on, and theories of, the represented entities are assumed not to play any relevant role inside an ontology, since the latter reflects a static, ‘frozen’, and widely shared portion of knowledge in a given field. So, although even a common-sense concept, like *sun*, refers to an aspect of reality that is ‘seen’ and understood in the terms set by a culturally determined conceptualization, there seems to be little or no point in introducing this whole

⁵ The term is borrowed from Ludwik Fleck’s epistemological observations on ‘thought-collectives’ (*Denkkollektiv*) and ‘thought-styles’ (*Denkstil*); cf. [29].

conceptualization explicitly in, e.g., an ontology of weather conditions.

The intuition underlying this practice, however, comes to odds when an ontology needs to be extended with social entities, such as social institutions, organizations, plans, regulations, narratives, schedules, parameters, diagnoses, etc. Important fields of investigation have denied an ontological primitiveness to social objects, since the latter are taken to have meaning only in combination with some other entity, i.e. it is assumed that their intended meaning results from a statement (see e.g. [33]).

In that view, for example, a norm, a plan, or a social role should be better represented as a (set of) logical statement(s), not as logical individuals. This position is documented by the almost exclusive attention dedicated by many relevant frameworks (such as BDI agent model, theory of trust, situation calculus, and formal context analysis) to states of affairs, facts, beliefs, and contexts, whose logical representation is set at the level of theories or models, not at the level of concepts or relations.

In *c.DnS*, we take seriously the attempt to build a constructive formal ontology that assumes social entities as first-class citizens in a logical theory's domain of quantification.

2.3 Informal description of *c.DnS*

The core structure of *c.DnS* [27] is the following:

$$\langle D, S, C, E, A, K, I, T \rangle$$

where *D* is the class of *Descriptions*, *S* is the class of *Situations*, *C* is the class of *Concepts*, *E* is the class of *Entities*, *A* is the class of (social) *Agents*, *K* is the class of *Collections*, *I* is the class of *Information Objects*, and *T* is the class of *Time intervals*.

In intuitive terms, these classes allow to model how a social agent, as a member of a certain community, singles out a situation at a certain time, by using a descriptive relation that assigns concepts to entities within that situation. In other words, these classes are meant to formalize the constructivist assumption according to which, in order to contextualize a concept, we need to take into account the viewpoint/s or description/s inside which the concept is defined or used, the situation/s this viewpoint 'carves out' the perceived environment, the entities which are in the setting of said situation/s, the social agents who share the viewpoint, the collective, or community, of which these agents are members, the information object/s by which the viewpoint is expressed, and, finally, the time or time-span characterizing the viewpoint.

c.DnS' classes are substructured as follows: *E* is the class of everything that is assumed to exist in some domain of interest, for any possible world. *E* is partitioned in the class of 'schematic entities', i.e. entities which are axiomatized in *c.DnS* (*D, S, C, A, K, I*), and the class of 'non-schematic entities', which are not characterized in *c.DnS* (*T*). Other non-primitive social entities may be added as subtypes of *E*, depending on the needs of the modeled domain. For instance in the application of *c.DnS* presented throughout section 3 the following additional entities are also considered: *physical agents* (*PA*), *internal representations* (*R*), *physical realizations* (*PR*), *objects* (*Object*), *actions* (*Action*), and *regions* (*Region*).

The main purpose of *c.DnS* is to *redescribe* entities that are (or are assumed to be) existing. For example, an existing situation including humans, cars, roads and signs can be redescribed as a driving situation, as a racing situation, as well as a speed-limit-violation situation, depending on the circumstances and on the intention of the interpreter of that situation.

In the field of developmental psychology, this ability has been described in terms of *Representational Redescription*, “a process by which (implicit) information that is in a cognitive system becomes progressively explicit to that system” [34], allowing for greater flexibility.

This ‘redescription game’ is played in terms of a number of projections of the general *c.DnS* relation, which allows to relate schematic and/or non-schematic entities to one another. We provide here a brief overview of such projections per class.

Descriptions are entities which represents a conceptualization, it is generically dependent on some (physical) agent and communicable ([20]). Examples of descriptions are regulations, plans, laws, diagnoses, projects, plots, techniques, etc. Descriptions have typical components, called concepts, and are related to other entities in *c.DnS* by means of the following projections: *defines*, *uses* (which hold between descriptions and concepts); *involves*, *individuallyConstructedAs* (compositions of relations holding between descriptions and entities); *unifies* (holding between a descriptions and a collections).

Situations are entities 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. Examples of situations (corresponding to the examples of descriptions above) are: facts, plan executions, legal cases, diagnostic cases, attempted projects, performances, technical actions, etc. Situations are related to other entities in *c.DnS* by means of the following projections: *satisfies* (holding between situations and descriptions); *hasInScope* (holding between situations); *settingFor* (holding between situations and entities).

Concepts are defined by a description and can be used in other descriptions. Concepts are related to other entities in *c.DnS* by means of the following projections: *classifies* (holding between concepts and entities); *covers*, *characterizes* (holding between concepts and collections).

Entities are anything that is assumed to exist in some domain of interest, for any possible world. Main subtypes of entities are ‘schematic’ and ‘non-schematic’. Both subtypes may have a *memberOf* relation with collections, while non-schematic entities are related to other entities in *c.DnS* by means of the following projections: *constructs* (holding between non-schematic entities); *actsFor* (holding between non-schematic entities and agents); *realizes* (holding between non-schematic entities and information objects).

Social agents may be a person or an organization, but never a bio-physical system that plays an agentive role⁶. Social agents are related to other entities in *c.DnS* by means of the following projections: *shares* (holding between social agents and descriptions); *re-describes* (holding between social agents and situations); *deputes* (holding between social agents and concepts).

⁶ Agents of the latter kind are introduced as non-schematic entities.

Collections are a naturalization in space-time of non-empty proper classes with (at least one) basic properties for membership. This seems to capture the common sense intuition underlying groups, teams, collections, collectives, associations, etc.. For an extensive treatment of similarities and dissimilarities between this notion of collection and the notions of (naturalized) set or class refer to [3].

Information Objects are units of information which are related to other entities in *c.DnS* by means of the following projections: *expresses* (between information objects and descriptions) and *about* (between information objects and entities). Collection are inversely related to other entities in *c.DnS* by means of the following projections: *unifies* (holding between a descriptions and a collections); *covers*, *characterizes* (holding between concepts and collections); *memberOf* (holding between entities and collections).

Time intervals, which are not characterized in DnS (i.e. they are non-schematic entities), are used for tagging descriptions, situations and projections. Time intervals should be added to the ontologies, that do not include them in their domain, when aligned to Constructive DnS.

c.DnS	Description (D)	Situation (S)	Concept (C)	Entity (E)	Social agent (A)	Collection (K)	Information object (I)
Description (D)	na	-satisfies	defines [1a], uses [1a]	involves [1c], individuallyConstructedAs [2b]	na	unifies [1h]	-expresses
Situation (S)	satisfies [1b]	hasInScope [1f]	na	settingFor [1d]	na	na	na
Concept (C)	-defines, -uses	na	na	classifies [1e]	na	covers [1h], characterizes [1h]	na
Entity (E)	-individuallyConstructedAs	-settingFor	-classifies	constructs [2b], -involves	acts for [2a]	memberOf [1h]	realizes [2c], -about
Social agent (A)	shares [1e]	re-describes [1f]	deputes [1e]	-acts for		na	na
Collection (K)	-unifies	na	-covers, -characterizes	-memberOf	na	na	na
Information object (I)	expresses [1g]	na	na	about [1g] -realizes	na	na	na

Table 1: c.DnS's classes, projections and principles, inverse projections

All projections mentioned above are based on two main constructive principles: the social construction principle and the grounded construction principle. On their turn, these principles are based on a larger set of other principles, listed below. Moreover, Table 1 shows the classes of *c.DnS* and their projections with a reference to the corresponding principles.

1. The *social construction principle* is based on:
 - (a) Relationality principle: concepts are always defined in a relational context (i.e. a *description* or a *gestalt*).
 - (b) Interpretability principle: situations are always emerging/interpreted in a relational context according to some expected configuration.
 - (c) Classification principle: entities are always internally represented with reference to a concept.

- (d) Situatedness principle: entities are always internally represented within a context according to some expected or reconstructed configuration.
 - (e) Sharing principle: descriptions are always dependent on social agents.
 - (f) Epistemological layering principle: given a description d_1 that involves another description d_2 , a situation s_1 that satisfies d_1 has in its scope a situation s_2 that satisfies d_2 .
 - (g) Formedness principle: descriptions are always expressed by information objects that provide a form to them.
 - (h) Containment principle: there exists a collection for all entities classified by a concept.
 - (i) Interaction principle: any social agent must be member of a knowledge collective in order to share a description.
2. The *grounded construction principle* is based on:
- (a) Agent efficacy principle: social agents should always be acted by some entity.
 - (b) Cognitive counterpart principle: for any description there is a social agent who shares it and who is acted by a physical agent that constructs an internal representation which is the individual construction of that description.
 - (c) Information grounding principle: any information object must have a physical support.

2.4 Plans, Norms and Collectives

Based on *c.DnS*, we define the notions of Plan, Norm and Collective and exploit their relations for defining typologies of normed collective entities.

For what concerns the notion of plan, we stick here to [26] where a plan is a description that represents an *action schema* that is *shared* by a social agent but *constructed* by a physical agent. In addition, a plan *defines* or *uses* or has as proper parts tasks, roles, goals, where:

Task is a concept that classifies *actions* (or similar non-schematic action-like entities)

Role is a concept that classifies *objects* (or similar non-schematic object-like entities)

Goal is the proper part of the plan that is desired by an agent

For what concerns the notion of norm, we take here the stance of [35] where a norm is a description, i.e. norms are treated there in their social sense (which includes but is not limited to the legal sense). This view on norms takes also into account Searle's distinction [21] between regulative and constitutive norms: regulative norms provide codes of conduct (i.e. regulations), while constitutive norms create social individuals and possibly contain no regulations at all. Here we mainly concentrate on regulative norms, on their social *usage* and on their influence on agents' plans and collectives. If the entities of a situation are classified by the concepts (roles, tasks, parameters) used in the norm, then the situation *falls under* the norm. This is an important difference from plans, which are *executed* in novel situations: norms are satisfied in a more complex and indirect way, because a situation that falls under a norm does not necessarily satisfy it.

Based on this simple model of norms, we address in this paper the issue of how to represent in *c.DnS* the influence that norms may have on plans. We are not concerned here with providing a full-fledged theory of this interactions. We rather

want to set a formal-ontological basis for modeling such theories. In other words, we are aware of the fact that relevant work on norms, and possibly on their interactions with plans, may be found both in the legal-philosophical literature [22], [23] and in the multi-agent systems literature [24]. Here though we base our analysis on an a set of intuitive distinctions between the types of interactions that norms may have with plans. Future will refine such distinctions, possibly based on other existing proposals. So, for the moment, in *c.DnS* norms may be seen as interacting with plans in one of three ways.

Norms as conventions that emerge from existing practices or plans. A norm, either social or legal, usually reflects an existing practice within a community. Typically, social and legal systems are the main way to maintain a stability among the members and the resources of a community, population, or country.

Norms as compliance checking protocols over social behavior or legal cases. Once a norm lifecycle is established, norms can be enforced by using them as filters for social behavior. Typically, the initiative for compliance checking is limited to the interest of other parties with respect to the behavior of one party.

Norms as constraints within plans. Once a norm lifecycle is established, and appropriate enforcing and compliance checking practices emerge, they can be used by social agents as constraints within their own plans. In this sense, norms are akin to *behavioral principles*. If taken as principles for social behaviour, norms can be executed, similarly to plans, and in fact they are executed as subplans.

Finally, *c.DnS*, together with the theories of plans and norms described above, provide us with the formal means to define the notion of collective and a typology for these:

Collection of agents is a collection unified by some rationale that is extrinsic with respect to the knowledge shared by the member agents.

Knowledge community is a collection of agents unified by descriptions that are shared by the member agents.

Intentional collective is a knowledge community that is unified by a plan shared by member agents.

Intentional normative collective is a knowledge community that is unified by a plan that, in turn, is entrenched with norms according to of the possible interaction between norms and plans described above.

3 Formal apparatus

3.1 The *c.DnS* relation

The *c.DnS* relation is given in (1). Each element of the core structure is encoded as a domain in a relation with arity=8:

$$c.DnS(d, s, c^*, e^*, a^*, k^*, i^*, t^*) \rightarrow D(d) \wedge S(s) \wedge C(c^*) \wedge E(e^*) \wedge A(a^*) \wedge K(k^*) \wedge I(i^*) \wedge T(t^*) \quad (1)$$

D can be read as *Description*, *S* as *Situation*, *C* as *Concept*, *E* as *Entity*, *A* as *social Agent*, *K* as *Collection*, *I* as *Information object*, and *T* as *Time interval*.

Intuitively, the *c.DnS* relation says that *a social agent, as a member of a given knowledge community, singles out a situation at a certain time, by using a descriptive relation that assigns concepts to entities within that situation.*

The ‘*’ variables are *ordered-list variables*, i.e. they can occur more than one time in an orderly way (ordered lists are paired, based on the projections described in section 3.3). Without list variables, *c.DnS* relation would formalize only ‘atomic’ situations, based e.g. on only one concept, one entity, one time interval, etc.). In real modelling (see 3.3), several occurrences of argument types are possible, for example admitting different agent’s and situation’s times, and several entities within a same situation, as when a detective singles out an event occurred days before, for the sake of interpreting a killer’s *modus operandi*. Such a case is exemplified in the statement 2, which contains four entities, four concepts, and two time intervals.

$$\begin{aligned} c.DnS(KnowledgeOfPreviousCases\#1, KillingSituation\#1, \\ \{Precedent, Killer, Tool, HypotheticalIntention\}, \\ \{Event\#1, PhysicalAgent\#1, PhysicalTool\#1, Plan\#1\}, \\ Detective\#1, InvestigationTeam\#1, PreviousCaseReport\#1, \\ \{TimeOfEvent\#1, TimeOfInterpretation\#1\}) \quad (2) \end{aligned}$$

In the following, list variables are not used, because *c.DnS* relation is best explained through its projections, which make it useful in most real world projects, where computational languages do not allow (or make it too complex) representing list variables and >2-ary relations.

3.2 Characterization of classes

E is the class of everything that is assumed to exist in some domain of interest for any possible world. (3):

$$\Box \forall x (E(x)) \quad (3)$$

D, *S*, *C*, *A*, *K*, *I*, and *T* are subclasses of *E* (1):

$$(D(x) \vee S(x) \vee C(x) \vee A(x) \vee K(x) \vee I(x) \vee T(x)) \rightarrow E(x) \quad (4)$$

D, *S*, *C*, *A*, *K*, and *I* are all mutually disjoint, and constitute the class *SE* of *schematic entities* (5).

$$SE(x) =_{df} D(x) \vee S(x) \vee C(x) \vee A(x) \vee K(x) \vee I(x) \quad (5)$$

All instances of *E* that are not instances of *SE* are *non-schematic* entities. *SE* and non-schematic entities cover the class *E*.

$$E(e) \equiv SE(e) \vee \neg SE(e) \quad (6)$$

Since time intervals are not in *SE*, they are non-schematic entities. Time intervals are important in *c.DnS* because we need to add a temporal indexing to some constructivist relations. In practice, when using *c.DnS* jointly with an existing ontology that does not reflect any commitment to time intervals, we need adding time intervals to it.

$$T(t) \rightarrow \neg SE(t) \quad (7)$$

The main application of *c.DnS* is *redescribing* existing entities, independent of such existence being derived from other (formal or informal) ontologies or assumed. For example, an existing situation including humans, cars, roads and signs can be redescribed as a driving situation, as a racing situation, as well as a speed-limit-violation situation, depending on other circumstances and on the intention of the interpreter of that situation. We define a *redescription* relation as a partial projection of *c.DnS* as follows:

$$\text{redescription}(a, e, s, t) \rightarrow A(a) \wedge E(e) \wedge S(s) \wedge T(t) \quad (8)$$

Axiom 8 states that redescription holds for agents and entities within a situation, at some time. For example, the sentence *the Italian road police has fined Manuel Fangio for a speed-limit violation on Thursday, January 18th 2007* can be modelled by using the redescription relation (9, 10).

$$\begin{aligned} &\text{redescription}(\text{ItalianRoadPolice}, \text{ManuelFangioDriving}, \\ &\text{SpeedLimitViolation18010732}, \text{ThursdayJanuary18th2007}) \quad (9) \end{aligned}$$

$$\begin{aligned} &\text{redescription}(\text{ItalianRoadPolice}, \text{ManuelFangio}, \\ &\text{SpeedLimitViolation18010732}, \text{ThursdayJanuary18th2007}) \quad (10) \end{aligned}$$

Each redescription concerns one of the entities that get redescribed in that situation. Statements 9 and 10 exemplify two such entities: Manuel Fangio, and his driving.

Based on the redescription relation, we define a class of *GroundEntities* as those entities that get redescribed:

$$G(e) =_{df} E(e) \wedge \exists a, s, t(\text{redescription}(a, e, s, t)) \quad (11)$$

Definition in 11 introduces ground entities as entities that are redescribed by an agent that ‘frames’ them within a situation at some time.

For the sake of intuition, the generic relation of *interpretation* (12) between an agent and an entity whatsoever at some time can be considered as a poorer projection of *c.DnS*:

$$\text{interprets}(a, e, t) \rightarrow A(a) \wedge E(e) \wedge T(t) \quad (12)$$

In practice, the constructive assumption in our ontology makes interpretation of entities by some agent at some time logically dependent on descriptions, situations, concepts, collectives, and information objects. How this assumption is unfolded, both formally and intuitively, is the theme of the next subsections. In the remainder of *c.DnS* presentation, we show how the *c.DnS* relation is projected and axiomatized so that the *redescription* relation can be actually used in as many domains as possible.

Additional entity types While a purely constructivist theory can live without postulating types of entities (besides time intervals), some common-sense type distinctions are of obvious practical advantage in the management of many social domains. In particular, we will present both a purely *constructivist* and a *grounded* (or ‘common-sense’) versions of the construction principle. In the grounded version, we make use of some types of non-schematic entities: *PA* (*physical agents*), *R* (*internal representations*), and *PR* (*physical realizations*). In the plan ontology (see below), we also use the classes *Object*, *Action*, and *Region*.

3.3 Projections of *c.DnS*: the *social construction principle*

Several projections of the *c.DnS* relation can be defined by means of binary or ternary relations, and axioms. Most projections we consider are irreflexive, asymmetric, and intransitive. Besides the specific projection signature, their axiomatization requires a temporalized *properPartOf* relation.

Here we list, mostly in an informal way, the projections that we deem necessary in order to lay a foundation for *c.DnS*. Each projection implies the full *c.DnS* relation according to the axiom schema in 13; these implication axioms should be assumed, and are not included in the axiomatization.

$$[projection](x_1 \dots x_n, x_i \in \{d, s, c, e, a, k, i, t\}) \rightarrow c.DnS(d, s, c, e, a, k, i, t) \quad (13)$$

Each projection is introduced as a *principle* for the logical representation of the construction of the social realm. The principles proposed will be finally composed into a *social construction principle*. Altogether, they constitute an extended account of the social constraints acting when an agent's ontological commitment is formed. On the other hand, one or more principles could be dropped if considered unnecessary or too strong for a particular ontology project (this implying that the construction principle is lost).

For the complete axiomatization, and for technical details on how *c.DnS* is applied in domain ontology projects, we refer to the technical reports from our ontology portal ⁷.

The following is the signature of the (basic) *c.DnS* projections:

(specializes, defines, uses, satisfies, classifies, involves, settingFor, shares, deposes, hasInScope, redescibes, expresses, about, memberOf, covers, characterizes, unifies, gUnifies)

The relationality principle The *defines* relation (14) is the projection of *c.DnS* over descriptions and concepts (cf. [20]).

$$defines(d, c) \rightarrow D(d) \wedge C(c) \quad (14)$$

Defines formalizes the intuition of a *gestalt* [36], or 'context' [32], that gives meaning to the parts. Some examples are modelled in 15, 16, 17.

$$defines(ItalianConstitution, Minister) \quad (15)$$

$$defines(LinneanTaxonomy, Species) \quad (16)$$

$$defines(CNRRegulation, SeniorResearcher) \quad (17)$$

If we assume that a *defines* relation is required for concepts, i.e. that concepts are always defined in a relational context - a *description* -, that assumption can be called the *relationality principle* (18).

$$C(c) \rightarrow \exists d(D(d) \wedge defines(d, c)) \quad (18)$$

⁷ <http://www.loa-cnr.it/ontologies/index.html>

The *uses* relation (axiom 19, exemplified in Statement 20) reflects the fact that, besides defining concepts, descriptions can also use concepts defined by some other description.

$$uses(d, c) \rightarrow D(d) \wedge C(c) \wedge \exists d' (d \neq d' \wedge defines(d', c)) \quad (19)$$

$$uses(ChiefOfStateVisitEtiquette, MasterOfCeremonies) \quad (20)$$

Descriptions can also *introduce* social agents, which are here entities such as persons, organizations, institutional figures, etc. (see 21, with some examples in 22, 23, 24)).

$$introduces(d, a) \rightarrow D(d) \wedge A(a) \quad (21)$$

$$introduces(ItalianConstitution, ItalianGovernment) \quad (22)$$

$$introduces(FIATLegalConstitution, FIAT_SpA) \quad (23)$$

$$introduces(ItalianLawBirthDeclaration, PhysicalLegalPerson) \quad (24)$$

Although introduction of agents falls under the relationality principle, like definition and usage, it has a different intuition from definition and usage, because concepts and agents are disjoint classes, where the differences are:

- agents can *share* descriptions (see section 3.3)
- agents (specially *organizations*) typically *depute* concepts (see axiom 41)
- in the grounded version of the construction principle (see section 3.4), social agents are *acted by* (axiom 81) some physical agent (axiom 84) that is classified by (axiom 29) some concept deputed by (axiom 41) a social agent

The interpretability principle The *satisfies* relation is the projection of $c.DnS$ over situations and descriptions.

$$satisfies(s, d) \rightarrow S(s) \wedge D(d) \quad (25)$$

It formalizes the intuition of an *instantiation* of a gestalt, i.e. the application of gestalts to actually occurring contexts in the life of a cognitive agent. For example:

$$satisfies(MandateForGovernmentToProdi, LawForGovernmentFormation) \quad (26)$$

If we assume that a *satisfies* relation is required for situations, i.e. that situations are always emerging/interpreted in a relational context according to some expected configuration, that assumption can be called the *interpretability* principle.

$$S(s) \rightarrow \exists d (D(d) \wedge satisfies(s, d)) \quad (27)$$

Each description generates a *situation class*, which contains all the situations that satisfy that description. For example,

$$satisfies(x, LawForGovernmentFormation) \rightarrow LegalGovernmentFormation(x) \quad (28)$$

A situation class can be empty however, since there may be descriptions that are never satisfied by any situation.

The classification principle The *classifies* relation is the projection of *c.DnS* over concepts and entities.

$$classifies(c, e, t) \rightarrow C(c) \wedge E(e) \wedge T(t) \quad (29)$$

It formalizes the intuition of a *redescription* of an entity, i.e. the application of a (new) gestaltic concept to something which is already provided with an available identity in actually occurring contexts in the life of a cognitive agent. For example, the statement 30 has the consequence that the social agent Napolitano is provided with the additional identity of ItalianPresidentRole for 2007:

$$classifies(ItalianPresidentRole, Napolitano, 2007) \quad (30)$$

Note that *ItalianPresident* is a social agent, since it has the properties given in section 3.3; anyway, that social agent also needs to *depute* (41) the concept *ItalianPresidentRole*⁸ that can classify different entities at different times, but only one at a time, while other concepts admit to classify different entities at the same time, e.g. the concept *Senator* (31). The different ways of classifying entities usually depend on the type of agent that deputes the concept, see section 3.4.

$$classifies(Senator, Napolitano, 2005) \quad (31)$$

$$classifies(Senator, LeviMontalcini, 2005) \quad (32)$$

If we assume that a *classifies* relation is required for ground entities to be considered in *c.DnS*, i.e. ground entities are always internally represented with reference to a concept, that assumption can be called the *classification* principle.

$$G(x) \rightarrow \exists c(C(c) \wedge classifies(c, x, t)) \quad (33)$$

Compositional projections can be defined from primitive ones. The projection *involves* is compositionally defined, and states that a description involves a ground entity when the latter is classified by a concept defined or used by the description.

$$\begin{aligned} involves(d, e, t) =_{df} & D(d) \wedge E(e) \wedge T(t) \wedge \\ \exists c((defines(d, c) \vee uses(d, c)) \wedge & classifies(c, e, t)) \end{aligned} \quad (34)$$

The situatedness principle The *settingFor* relation is the projection of *c.DnS* over situations and entities.

$$settingFor(s, e) \rightarrow S(s) \wedge E(e) \quad (35)$$

It formalizes the intuition of *contextualization* of an entity, i.e. the application of gestalts to actually occurring contexts in the life of a cognitive agent. For example:

$$settingFor(MandateForGovernmentToProdi, Napolitano, 2007) \quad (36)$$

If we assume that a *settingFor* relation is required for ground entities, i.e. ground entities are always internally represented within a context according to some expected or reconstructed configuration, such assumption can be called the *situatedness* principle (37).

$$G(x) \rightarrow \exists s(S(s) \wedge settingFor(s, x)) \quad (37)$$

⁸ Differently from the legal one, the common sense notion of Italian president is usually that of a concept, not of a social agent

The sharing principle The *shares* relation is the projection of *c.DnS* over social agents and descriptions.

$$shares(a, d, t) \rightarrow A(a) \wedge D(d) \wedge T(t) \quad (38)$$

It formalizes the intuition of the *social nature* of a description, i.e. the mapping of descriptions on social agents that are acted by one or more physical agents. Note that by ‘social nature’ we do not mean that a description should actually be shared by a community (although this is typically what happens), but that a description must be communicable among social agents. For example:

$$shares(Napolitano, LawForGovernmentFormation, 2006) \quad (39)$$

If we assume that a *shares* relation is required for descriptions, i.e. descriptions are always dependent on social agents, that assumption can be called the *sharing principle*.

Notice that social agents include both persons and organizations, but not physical systems that can play agentive roles (these are introduced in the grounded version of DnS, see section 3.4).

The sharing principle states that descriptions must be shared by at least one agent (40).

$$D(x) \rightarrow \exists(a, t)(A(a) \wedge T(t) \wedge shares(a, x, t)) \quad (40)$$

Besides sharing descriptions, social agents can *depute* (41) concepts (e.g. *roles*) that are supposed to enact the agent’s actions (see section 3.4).

$$deputes(a, c, t) \rightarrow A(s) \wedge C(c) \wedge T(t) \quad (41)$$

For example, a telecom company can depute the role ‘engineer’ that can classify certain entities (typically, persons with appropriate curricula) to act for the company. Back to our legal example:

$$deputes(ItalianState, ItalianPresident, 2006) \quad (42)$$

The epistemological layering principle The *hasInScope* relation reflects the intuition that situations can be *epistemologically layered*, when a description d_1 involves another description d_2 , and a situation s_1 that *satisfies* d_1 has another situation s_2 that *satisfies* d_2 in its scope. (43).

$$\begin{aligned} hasInScope(s_1, s_2) =_{df} & S(s_1) \wedge S(s_2) \wedge s_1 \neq s_2 \wedge \\ & \exists(d_1, d_2, a, t)(D(d_1) \wedge D(d_2) \wedge A(a) \wedge T(t) \wedge \\ & d_1 \neq d_2 \wedge satisfies(s_1, d_1) \wedge satisfies(s_2, d_2) \wedge \\ & involves(d_1, d_2, t) \wedge shares(a, d_1, t) \wedge shares(a, d_2, t)) \end{aligned} \quad (43)$$

An example is in 44.

$$hasInScope(MurderCase_{128}, CaesarStabbedByBrutus) \quad (44)$$

Epistemological layering is a principle in the *c.DnS* core, corresponding to the *figure-ground shifting* cognitive schema from Gestalt psychology and, more recently, cognitive linguistics [36], [37].

The application of epistemological layering is fundamental in *c.DnS*, since it accounts for the role of agents in the application of a description to some situation, i.e., in order to include the ontological commitment within an ontology's domain of discourse. In practice, ontological commitment [38] postulates the action of some agent that has the capability and the intention to (*re*)describe (=interpret) a situation.

This is formalized by means of the relation *redescribes* (45), which is the projection of *c.DnS* over social agents and situations.

$$\begin{aligned} \text{redescribes}(a, s_2, t) =_{df} & A(a) \wedge S(s_2) \wedge T(t) \wedge \\ & \exists s_1 (S(s_1) \wedge s_1 \neq s_2 \wedge \text{shares}(a, d_1, t) \wedge \\ & \text{satisfies}(s_1, d_1) \wedge \text{hasInScope}(s_1, s_2)) \end{aligned} \quad (45)$$

An example of application of *redescribes* is in 46.

$$\text{redescribes}(\text{SherlockHolmes}, \text{HoundOfBaskervilleFact}, 1890) \quad (46)$$

The formedness principle The *expresses* relation is the projection of *c.DnS* over information objects and descriptions. It formalizes the intuition of the intrinsic *communicability* of every description (cf. [34]).

$$\text{expresses}(i, d, t) \rightarrow I(i) \wedge D(d) \wedge T(t) \quad (47)$$

For example:

$$\text{expresses}(\text{ItalianConstitutionText}, \text{ItalianConstitution}, 1946) \quad (48)$$

We call *formedness* principle (49) the assumption that an *expresses* relation is required for descriptions, i.e. descriptions are always expressed by information objects that provide a form to them.

$$D(x) \rightarrow \exists(i, t)(I(i) \wedge (T(t) \wedge \text{expresses}(i, x, t))) \quad (49)$$

Information objects appear in other projections of *c.DnS*, which can be defined compositionally. For example, the *aboutness* of information objects can be defined through composing the *expresses*, *satisfies*, and *settingFor* relations (50).

$$\begin{aligned} \text{about}(i, e, t) =_{df} & I(i) \wedge E(e) \wedge T(t) \wedge \exists(d, s)(D(d) \wedge S(s) \wedge \\ & \text{expresses}(i, d, t) \wedge \text{satisfies}(s, d) \wedge \text{settingFor}(s, e)) \end{aligned} \quad (50)$$

Aboutness states that, if the description expressed by an information object is satisfied by a situation, the information object can be *about* any entity that is in the setting of said situation. For example, in 51, the Italian Constitution is (also) about Italy.

$$\text{about}(\text{ItalianConstitutionText}, \text{Italy}, 2006) \quad (51)$$

The containment principle The entities that are classified by a same concept or a same set of concepts, either defined by the same description or not, are easier to compare, and can be put in a same collection (K). An appropriate *memberOf* relation (52) holds for sets of said entities⁹

$$memberOf(e, k, t) \rightarrow E(e) \wedge K(k) \wedge T(t) \quad (52)$$

For example, in 54, D'Alema is a member of the Italian Government collective in 2007. Italian Government collective is intended here as the collection of all members from any particular Italian Government.

$$memberOf(D'Alema, ItalianGovernmentCollective, 2007) \quad (53)$$

Note that *ItalianGovernmentCollective* is not the same entity as *ItalianGovernment*, which is a social agent. The difference is not purely academic, because Italian Government identity depends on the current Italian Constitution, and Italian Government collective is bound to it; but there are collectives of Italian governments that have different identities based on the way they have been elected, nominated, or behaved.

For example, from the legal viewpoint, *Prodi2Collective* has a different identity from the general Italian Government collective, although the two collectives are co-extensional (have the same members) for a limited time period (see example ??). Of course, also Prodi2 is a social agent, which *specializes* ItalianGovernment (cf. 55).

$$memberOf(D'Alema, Prodi2Collective, 2007) \quad (54)$$

$$specializes(Prodi2, ItalianGovernment) \quad (55)$$

If we postulate a collection comprising all entities classified by a concept, for each concept, the resulting axiom represents the *containment* principle (56).

$$C(c) \rightarrow \forall(e, t)((E(e) \wedge T(t) \wedge classifies(c, e, t)) \rightarrow \exists x(K(x) \wedge memberOf(e, x, t)) \quad (56)$$

The concept(s) that classify all the members of a collection are said to *cover* a collection (57):

$$covers(c, k) =_{df} C(c) \wedge K(k) \wedge \forall(e, t)((E(e) \wedge memberOf(e, k, t)) \rightarrow classifies(c, e, t)) \quad (57)$$

Statement 58 is about the fact that the collective *ItalianMinisterCouncil* has all members that are classified by the concept *Minister*.

$$covers(Minister, ItalianMinisterCouncil) \quad (58)$$

Many collections can have subcollections covered by different concepts. In that case, we say that those concepts *characterize* the collection (59). Since subcollections can change without affecting the identity of a collection, *characterizes* is temporalized.

$$characterizes(c, k, t) =_{df} C(c) \wedge K(k) \wedge T(t) \wedge \exists(k_1)(covers(c, k_1) \wedge properPartOf(k_1, k, t)) \quad (59)$$

⁹ Cf. [26] and [3] for a different axiomatization that also assumes DOLCE).

Statement 60 is about the fact that (from a socio-political viewpoint), the collective *ItalianMinisterCouncil* has some members that are classified by the concept *Reformer*.

$$characterizes(Reformer, ItalianMinisterCouncil, 2006) \quad (60)$$

A complex concept, whose component concepts collectively characterize all members of a collection, results to *cover* it (cf. 61).

$$\begin{aligned} \forall(c, k)(C(c) \wedge K(k) \wedge \exists(c_1, c_2, t)(c_1 \neq c_2 \wedge \\ characterizes(c_1, k, t) \wedge characterizes(c_2, k, t) \wedge \\ properPartOf(c_1, c, t) \wedge properPartOf(c_2, c, t) \wedge \\ \neg \exists(c_3)(characterizes(c_3, k, t) \wedge c_1 \neq c_3 \wedge c_2 \neq c_3)) \rightarrow \\ covers(c, k, t)) \end{aligned} \quad (61)$$

The descriptions that define the concept(s) or concept collections that *cover* a collection are said to *unify* it (62).

$$\begin{aligned} unifies(d, k) =_{df} D(d) \wedge K(k) \wedge \\ \exists(c)(defines(d, c) \wedge covers(c, k)) \end{aligned} \quad (62)$$

Statement 63 is inferred from 62, 15, and 58: since *unifies* composes the relations *defines* and *covers*, the description *ItalianConstitution* defines the concept *Minister*, and *Minister* covers the collection *ItalianMinisterCouncil*, then *ItalianConstitution* unifies *ItalianMinisterCouncil*.

$$unifies(ItalianConstitution, ItalianMinisterCouncil) \quad (63)$$

When unification is applied to the parts of an entity, so that the unifying description defines concepts that characterize the configuration aspects of that entity, the collection is called *configuration* (*Cfg*, definition 64).

$$\begin{aligned} Cfg(k) =_{df} K(k) \wedge \\ \forall(e, t)(memberOf(e, k, t) \rightarrow \\ \exists(e_1, d, t)(properPartOf(e, e_1, t) \wedge \\ unifies(d, k) \wedge involves(d, e_1, t))) \end{aligned} \quad (64)$$

For example, the collection of all parts of a car, when the unifying description is its functional design, is a configuration.

In case a collection is covered or characterized by more than one concept defined in different descriptions, so that all entities in the collection are classified by characterizing concepts, then the collection results to be unified by a *bundle* of descriptions, in which the characterizing concepts are defined (definition 65).

$$\begin{aligned} Bundle(d) =_{df} D(d) \wedge \exists(d_1, d_2, t)(properPartOf(d_1, d, t) \wedge \\ properPartOf(d_2, d, t) \wedge (\exists(k, c_1, c_2)(defines(d_1, c_1) \wedge \\ defines(d_2, c_2) \wedge characterizes(c_1, k) \wedge \\ characterizes(c_2, k) \wedge unifies(d, k)) \vee \\ (\exists(s)(satisfies(s, d_1) \wedge satisfies(s, d_2) \wedge satisfies(s, d)))))) \end{aligned} \quad (65)$$

The notion of bundle of descriptions is defined in 65: a bundle is a (mereological) sum of (at least two) descriptions that are either all satisfied by a situation, or all define concepts that characterize a same collection.

The taxonomy principle The *specializes* relation (66) is the projection of *c.DnS* between schematic entities (in [20] this relation is limited to concepts only). It conveys the intuition of a taxonomic schema across schematic entities, for example in 67, the social agent *Prodi2 Government* specializes *Italian Government*.

$$specializes(se_1, se) \rightarrow SE(se_1) \wedge SE(se) \quad (66)$$

$$specializes(Prodi2Government, ItalianGovernment) \quad (67)$$

The difference between *specializes* and the traditional *subClassOf* and *instanceOf* relations is subtle. Firstly, *specializes* can be considered as a *reification* of *subClassOf*, since the latter holds for logical classes, while *specializes* holds for schematic entities.¹⁰

Secondly, since we are using first-order logic with a model-theoretic semantics, the *subClassOf* and *instanceOf* relations can also be used with schematic entities, and the choice between *specializes* and *instanceOf* often results to be a matter of good practice. For example, we may want to consider *Government* as a class instead of a social agent, if there is no given description that *introduces* (cf. axiom 21) government as a social agent. On the contrary, *ItalianGovernment* is introduced by the description *ItalianConstitution*, therefore it can be suitably modeled as a social agent. As a consequence, *Government* is *subClassOf* A (Social Agent), *Prodi2Government* *specializes* *ItalianGovernment*, and both *Prodi2Government* and *ItalianGovernment* are *instanceOf* *Government*.

The interaction principle A constructivist ontology should be able to contextualize agents' knowledge within their communities, called *thought collectives* in (cf. [39]). To this purpose, we introduce a class *KC* of knowledge communities, as a collection whose members share at least one description (68).

$$KC(k) =_{df} K(k) \wedge \exists(d)(D(d) \wedge \forall(a, t)(memberOf(a, k, t) \rightarrow shares(a, d, t))) \quad (68)$$

A knowledge community is different from a simple agent collection, as in cases like biological species, epidemiological groups, etc., since the members of a simple agent collection do not necessarily share any description, and it is even doubtful if agents like members of a biological or clinical group could be considered as agents at all, in the sense proposed here.

When the *membership* relation is considered necessary for descriptions to exist, the resulting axiom corresponds to the *interaction* principle (69): any social agent must be member of a knowledge community in order to share a description.

$$shares(a, d, t) \rightarrow \exists kc(memberOf(a, kc, t) \wedge KC(kc) \wedge unifies(d, kc) \wedge \exists d_1(\forall a_1(memberOf(a_1, kc, t) \rightarrow shares(a_1, d_1, t)))) \quad (69)$$

¹⁰ In a similar vein, descriptions can be considered as reifications of intensional relations, concepts as reifications of intensional classes, situations as reifications of extensional relations, and collections as reifications of extensional classes.

The statements in 70 through 75 exemplify how sharing is independent from *assuming* a description: *FlogistonTheory* was shared by both Stahl and Lavoisier, but only Stahl assumed it; on the other hand, it is not sure that Stahl actually shared *OxygenTheory*, because in the original debate there is no proof that he understood it, therefore, we are only allowed to state that Lavoisier shared (and assumed) it.

$$\text{shares}(\text{Stahl}, \text{FlogistonTheory}) \quad (70)$$

$$\text{shares}(\text{Lavoisier}, \text{FlogistonTheory}) \quad (71)$$

$$\text{shares}(\text{Lavoisier}, \text{OxygenTheory}) \quad (72)$$

Assumes is here proposed (73) as a more specific way of sharing a description, but without defining it. Defining assumption would require much more, e.g., we should axiomatize the relation between assumptions of descriptions, and beliefs about situations: while sharing a description is certainly required to an agent in order to believe a situation that satisfies that description, it is not sufficient to conclude that sharing is sufficient to that agent to actually believe it. The issue is even subtler, because we cannot either conclude that assuming that description is sufficient to believe that situation, since there can be additional constraints that make a situation unbelievable. Conversely, there can be cases in which a situation is believed without assuming the description it satisfies. We do not attempt an axiomatization of these epistemological issues here.

$$\text{assumes}(a, d, t) \rightarrow \text{shares}(a, d, t) \quad (73)$$

$$\text{assumes}(\text{Stahl}, \text{FlogistonTheory}) \quad (74)$$

$$\text{assumes}(\text{Lavoisier}, \text{OxygenTheory}) \quad (75)$$

Although we stay neutral with reference to how assumptions and beliefs are intertwined, we can use *assumes* as a primitive to introduce the notion of *paradigm* (76), which is important for a constructive ontology and to characterize collectives. Paradigms are defined here as bundle-based configurations of descriptions that are assumed by the members of a knowledge community. Those knowledge communities (common either in the commonsense or the scientific domains) result to be unified by paradigms.

$$\begin{aligned} \text{Paradigm}(p) =_{af} & \text{Bundle}(p) \wedge \forall(d, t)(\text{properPartOf}(d, p, t) \rightarrow \\ & D(d) \wedge \exists(kc)(KC(kc) \wedge \forall(a)(\text{memberOf}(a, kc, t) \rightarrow \\ & \text{assumes}(a, d, t)))) \quad (76) \end{aligned}$$

The social construction principle The unification relation holding for collections can also be used for ground entities, and we generalize a temporal version of it for all ground entities (*gUnifies*, or generalized unification, 80).

A composition of *c.DnS* projections leads to the *social construction principle* (77):

$$\begin{aligned}
 G(x) \leftrightarrow \\
 \exists(d, s, c, a, i, kc, t, c_1, d_1, s_1, t_1) & (D(d) \wedge S(s) \wedge C(c) \wedge A(a) \wedge \\
 & I(i) \wedge KC(kc) \wedge T(t) \wedge C(c_1) \wedge D(d_1) \wedge S(s_1) \wedge \\
 T(t_1) \wedge \text{classifies}(c, x, t) \wedge \text{settingFor}(s, x) \wedge \text{defines}(d, c) \wedge \\
 & \text{satisfies}(s, d) \wedge \text{shares}(a, d, t) \wedge \text{unifies}(d, kc) \wedge \\
 \text{memberOf}(a, kc, t) \wedge \text{deputes}(a, c_1, t) \wedge \text{expresses}(i, d, t) \wedge \\
 & \text{settingFor}(s, t) \wedge \text{re-describes}(a, s, t_1) \wedge \text{shares}(a, d_1, t_1) \wedge \\
 & gUnifies(d_1, d, t_1) \wedge \text{satisfies}(s_1, d_1) \wedge \text{hasInScope}(s_1, s) \quad (77)
 \end{aligned}$$

According to the social construction principle, when re-described by *c.DnS*, a ground entity *x* gets characterized as follows:

- *x* is always *classified* at some time by at least one concept that is *defined* in a description that is *satisfied* by a situation that is a *setting* for *x*
- *x* description has to be *shared* by a social agent that is a *member* of a knowledge community
- the description has to be *expressed* by an information object
- the social agent has to *depute* concepts that classify entities from a situation
- the social agent that shares *x*'s description *re-describes* *x*'s situation by means of *involving* *x*'s description into another description. This is equivalent to having *x*'s situation *in the scope* of the redescription situation.¹¹

The social construction principle can be interpreted as a *unity criterion* (cf. also [40] for a non-reified account of unity criteria), which becomes available to *x*. In other words, its redescription allows *x* to be unified (80). The unification relation holding for collections (62) can be used for any ground entity, and we generalize a temporal version of it for all ground entities (*gUnifies*, or generalized unification, 78).

$$\begin{aligned}
 gUnifies(d, g, t) =_{df} & D(d) \wedge G(g) \wedge T(t) \wedge \\
 \exists(c) & (C(c) \wedge \text{defines}(d, c) \wedge \text{classifies}(c, g, t) \quad (78)
 \end{aligned}$$

The intuition of generalized unification is that we can imply a “singleton” collection that is covered by the concept *c* (axiom 79), and whose unique member is the ground entity *g*. Since membership is temporalized, we need a temporal index in 78.

$$\begin{aligned}
 gUnifies(d, g, t) \rightarrow \exists(k, c) & (K(k) \wedge \text{memberOf}(g, k, t) \wedge \\
 & \text{unifies}(d, k) \wedge \text{defines}(d, c) \wedge \text{covers}(c, k) \wedge \\
 & \neg \exists(e) (g \neq e \wedge \text{memberOf}(e, k, t)) \quad (79)
 \end{aligned}$$

Another simpler way to explain the notion of ground entity is therefore based on *gUnifies* (axiom 80).

$$G(x) \leftrightarrow \exists(d, t) (gUnifies(d, x, t)) \quad (80)$$

¹¹ Note that there is no room for infinite regression, because the construction principle does not apply to ground entities that are also SE, therefore the redescription situation is not itself in the scope of a further redescription situation, unless requested by the case.

3.4 Projections of *c.DnS*: the *grounded construction principle*

So far, we have only concentrated on the core DnS relation, which focuses of schematic entities and how they can be used to provide ground entities with a unity criterion. On the other hand, we have also anticipated that a more practical version of DnS can include a mild commitment to certain types of (non-schematic) entities. In the following, we introduce new relations that will eventually allow the specification of a grounded construction principle.

The following is the signature of the additional projections for a grounded *c.DnS*:

$$\langle actsFor, constructs, individuallyConstructedAs, realizes \rangle$$

The agentEfficacy principle In 3.3 we have assumed that social agents can *depute* (41) concepts (e.g. *roles*) that are supposed to classify the entities that can act for the agent. For example, a telecom company can depute the role engineer that can classify certain entities (typically, natural persons with appropriate curricula) to act for the company.

The *actsFor* (81) relation holds for entities and social agents. It formalizes the intuition of *acting for* a social agent, i.e. the mapping of entities as actors that are classified by concepts that are *deputed by* a social agent.

$$actsFor(e, a, t) =_{df} E(e) \wedge A(a) \wedge T(t) \wedge \exists c(classifies(c, e, t) \wedge deputes(a, c, t)) \quad (81)$$

An example is provided in 42.

$$actsFor(Napolitano, ItalianState, 2007) \quad (82)$$

If we assume that an *actsFor* relation is required for social agents, i.e. social agents should always be acted by some entity, that assumption can be called the *agentEfficacy* principle (83).

$$A(x) \rightarrow \exists(e, t)(E(e) \wedge actsFor(e, x, t)) \quad (83)$$

Typically, social agents are acted by physical organisms, but actors can also be natural and legal persons, animals, robots, or even viruses. However, agent efficacy could be supported with a stronger claim, i.e. the *rationalAgentEfficacy* principle, stating that social agents must be acted by entities that can have internal meta-representations, hence only by (a subclass of) cognitive systems.

The cognitiveCounterpart principle We introduce a class for entities that can ground the action of social agents, and call them *physical agents*, or *PA* (84).

$$PA(pa) \rightarrow E(pa) \wedge \neg SE(pa) \quad (84)$$

Similarly, we ground descriptions in entities that can be localized into individual physical agents, and call them *internal representations* (85).

$$R(r) \rightarrow E(r) \wedge \neg SE(r) \quad (85)$$

The *constructs* relation holds for physical agents (or cognitive systems) and internal representations. Physical agents are considered non-schematic entities,

so we have to identify them in the domain of existing ontologies, and possibly add them to an ontology when missing (similarly to time intervals).

$$\text{constructs}(pa, r, t) \rightarrow PA(pa) \wedge R(r) \wedge T(t) \quad (86)$$

Statement 87 exemplifies grounded construction.

$$\begin{aligned} &\text{constructs}(\text{NapolitanoAsOrganism}, \\ &N.'sRepresentationOfItalianConstitution, 2007) \end{aligned} \quad (87)$$

The *individuallyConstructedAs* relation is the projection of *c.DnS* over descriptions and internal representations. It formalizes the correlate of descriptions as internal representations in a physical agent (intended as a cognitive system) that *actsFor* a social agent that *shares* the description. A *cognitiveCounterpart* principle states that for any description there is a social agent that shares it, and which is acted by a physical agent that constructs an internal representation that is the individual construction of that description.

$$\begin{aligned} &\text{individuallyConstructedAs}(d, r, t) \rightarrow \exists(pa)(PA(pa) \wedge \\ &\text{shares}(a, d, t) \wedge \text{constructs}(pa, r, t) \wedge \text{actsFor}(pa, a, t)) \end{aligned} \quad (88)$$

Statement 89 exemplifies individual grounded construction.

$$\begin{aligned} &\text{individuallyConstructedAs}(\text{ItalianConstitution}, \\ &N.'sRepresentationOfItalianConstitution, 2007) \end{aligned} \quad (89)$$

The agentEfficacy and the cognitiveCounterpart principles can be composed, in order to create a dependency of schematic entities on non-schematic ones, i.e. assuming that sharing descriptions requires constructing internal representations.

$$\begin{aligned} &\text{shares}(a, d, t) \rightarrow \exists(pa, r)(PA(pa) \wedge \text{constructs}(pa, r, t) \\ &\wedge \text{actsFor}(pa, a, t)) \end{aligned} \quad (90)$$

The informationGrounding and groundedConstruction principles An important projection concerning the way descriptions are substantially shaped and communicated is the *realizes* relation (93), holding between information objects and (physical) ground entities, which we call *physical realizations*, or *PR* (91).

$$PR(pr) \rightarrow E(pr) \wedge \neg SE(pr) \quad (91)$$

$$\text{realizes}(pr, i, t) \rightarrow PR(pr) \wedge I(i) \wedge T(t) \quad (92)$$

For example, the original paper document of the Italian Constitution realizes the Italian Constitution text in 1946 (93)

$$\begin{aligned} &\text{realizes}(\text{OriginalPaperDocumentOfItalianConstitution}, \\ &\text{ItalianConstitutionText}, 1946) \end{aligned} \quad (93)$$

This is related to information *grounding*, which is an obvious precondition for communication to happen: any information object must have a physical ‘support’.

Similarly to time intervals and physical agents, physical realizations of information must be present in the ground domain, or need to be added to it.

The agentEfficacy and the cognitiveCounterpart principles can be composed with the informationGrounding principle, in order to create a dependency of description sharing (and their internal construction) on realizing information objects that, as physical realizations, express the descriptions (94).

$$\begin{aligned} \text{shares}(a, d, t) \rightarrow \exists(pa, r, i, pr) & (PA(pa) \wedge \text{constructs}(pa, r, t) \wedge \\ & \text{actsFor}(pa, a, t) \wedge \text{expresses}(i, d) \wedge \text{realizes}(pr, i)) \end{aligned} \quad (94)$$

Axiom 94 enables a stronger interpretation of the *redescription relation* (8), since it now requires a social agent to be grounded in a physical agent, a description to be grounded in an internal representation constructed by the physical agent, and an information object to be grounded in a physical realization that realizes it. Stronger redescription results into *grounded scoping* (96), and the *grounded construction principle* (97). Grounded scoping allows us to distinguish two different times for a situation s : the time of its *setting*, and the time of its *redescription*. Only the first is the ‘real’ time of s , while the second one is actually the time of the redescription situation $s' \neq s$, so that $\text{hasGroundedScope}(s', s)$.

Notice that not all applications of *c.DnS* need the specification of a redescription situation. Such a situation is postulated by the theory, but its explicit naming and specification are useful only when the epistemological decision concerning interpretation is of some concern. E.g., when assessing witnesses in a legal case, or when selecting between judgments that may have different authoritativeness, trust, or contextual bindings.

$$\begin{aligned} \text{hasGroundedScope}(s_1, s_2) =_{df} \text{hasInScope}(s_1, s_2) \wedge \\ \exists(a, t, pa, r) & (A(a) \wedge T(t) \wedge PA(pa) \wedge R(r) \wedge \\ & \text{redescribes}(a, s_2, t) \wedge \text{actsFor}(pa, a) \wedge \text{constructs}(pa, r)) \end{aligned} \quad (95)$$

An example of grounded scoping is 96.

$$\begin{aligned} \text{hasGroundedScope}(\text{SherlockHInterpretation}, \\ \text{HoundOfBaskervilleFact}) \end{aligned} \quad (96)$$

Finally, we present the enriched construction principle with grounding in the axiom 97.

$$\begin{aligned} G(x) \leftrightarrow \\ \exists(d, s, c, a, r, i, kc, t, c_1, pa, pr, d_1, s_1, t_1) & (D(d) \wedge S(s) \wedge C(c) \wedge \\ & A(a) \wedge R(r) \wedge I(i) \wedge KC(kc) \wedge T(t) \wedge C(c_1) \wedge PA(pa) \wedge \\ & PR(pr) \wedge D(d_1) \wedge S(s_1) \wedge T(t_1) \wedge \text{classifies}(c, x, t) \wedge \\ & \text{settingFor}(s, x) \wedge \text{defines}(d, c) \wedge \text{satisfies}(s, d) \wedge \\ & \text{shares}(a, d, t) \wedge \text{unifies}(d, kc) \wedge \text{memberOf}(a, kc, t) \wedge \\ & \text{constructs}(pa, r, t) \wedge \text{actsFor}(pa, a, t) \wedge \text{deputes}(a, c_1, t) \wedge \\ & \text{classifies}(c_1, pa, t) \wedge \text{expresses}(i, d, t) \wedge \text{realizes}(pr, i, t) \wedge \\ & \text{settingFor}(s, t) \wedge \text{settingFor}(s, pa) \wedge \text{redescribes}(a, s, t_1) \wedge \\ & \text{settingFor}(s_1, pa) \wedge \text{settingFor}(s_1, pr) \wedge \text{settingFor}(s_1, r) \wedge \\ & \text{shares}(a, d_1, t_1) \wedge \text{gUnifies}(d_1, d, t_1) \wedge \text{satisfies}(s_1, d_1) \wedge \\ & \text{hasInScope}(s_1, s)) \end{aligned} \quad (97)$$

According to the construction principle, when redescribed by $c.DnS$, a ground entity x gets now an additional characterization:

- the description of x has to be shared by a social agent that is acted by at least one physical agent capable of constructing an internal representation
- the description has to be expressed by an information object that is realized by a physical realization
- the social agent has to depute a concept that classifies the physical agent acting for it.

3.5 Plans

Before discussing our typology of collectives, we introduce here some axioms for *plans* [26], which have the following properties (103):

- A **plan** is a description that represents an *action schema*
- Coherently with $c.DnS$, we assume that a plan is *shared* by a social agent, provided that it is *constructed* by a physical agent (90).
- A plan *defines* or *uses* at least one *task* (101) and one *role* (100), which are two kinds of concepts.
- A plan has at least one *goal* (108 below) as a proper part (*properPart* is assumed with its usual mereological semantics).

Tasks are concepts that classify action-like entities, which we assume here as **Actions** (98) without a specific characterization, while roles are concepts that classify object-like entities, which are also assumed generically as **Objects** (99)¹². Finally, roles can have tasks as *targets* (102).

$$Action(e) \rightarrow E(e) \quad (98)$$

$$Object(e) \rightarrow E(e) \quad (99)$$

$$Role(c) =_{df} C(c) \wedge \forall(e, t)(classifies(c, e, t) \rightarrow Object(e)) \quad (100)$$

$$Task(c) =_{df} C(c) \wedge \forall(e, t)(classifies(c, e, t) \rightarrow Action(e)) \wedge \exists(r)(Role(r) \wedge targets(r, c)) \quad (101)$$

$$targets(x, y) \rightarrow Role(x) \wedge Task(y) \quad (102)$$

In [26], roles are explicitly defined as concepts that classify DOLCE objects, while tasks are defined as concepts that classify DOLCE actions, but here we do not make any commitment on how action-like or object-like entities should

¹² The choice of introducing actions and objects as pure primitives follows our practice of avoiding overcommitment, i.e. the attempt to provide axiomatic constraints without a specific need coming from a domain or problem to be represented or solved. We have followed the same practice in general $c.DnS$ when introducing time intervals, physical agents and physical realizations with no characterization, except being *entities*

be represented in a (legacy) ontology. Based on previous axioms, a *Plan* class is characterized in 103.

$$\begin{aligned} Plan(d) \rightarrow & Description(d) \wedge \exists(a, t, c_1, c_2, g)(shares(a, d, t) \wedge \\ & A(y) \wedge T(t) \wedge Task(c_1) \wedge uses(d, c_1) \wedge Role(c_2) \wedge \\ & uses(d, c_2) \wedge Goal(g) \wedge properPartOf(g, d, t)) \end{aligned} \quad (103)$$

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.

Plans can have a rich internal structure, because they can have *subplans*, main and intermediate goals, roles that target more than one task, tasks that are targeted by more than one role, hierarchical tasks and roles, parameters on attributes of entities classified by tasks or roles, etc. A rich axiomatization of plan structures and task types is provided in [26]; here we only concentrate on goals.

Parts of plans A plan can have several *proper parts* (regulations, goals, norms), including other plans. For example, social agents are introduced by *constitutive descriptions* (104); if a plan *introduces* (21) a social agent, the related *constitutive description* is a *proper part* of the plan (105):

$$\begin{aligned} ConstitutiveDescription(x) =_{df} & D(x) \wedge \exists(a)(introduces(x, a)) \quad (104) \\ Plan(x) \rightarrow & introduces(x, a) \leftrightarrow \\ \exists(y)(ConstitutiveDescription(y) \wedge & defines(y, a) \wedge \\ & properPartOf(y, x, t)) \end{aligned} \quad (105)$$

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

Plans can have subplans (106).

$$\begin{aligned} subPlan(x, y, t) \rightarrow & properPartOf(y, x, t) \wedge Plan(x) \wedge \\ & Plan(y) \wedge T(t) \end{aligned} \quad (106)$$

Goals are necessary proper parts of plans, and are considered here as desires (another kind of description) that are proper parts of a plan.

$$Desire(x) \rightarrow Description(x) \quad (107)$$

For example, a *desire to start a relationship* can become a *goal to start a relationship* if someone assumes a plan in order to *take action* - or to let someone else take action on her behalf - with the purpose of starting that relationship. We propose 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 anyway to be a proper part of a plan:

$$\begin{aligned} Goal(x) =_{df} & Desire(x) \wedge \exists(p, t)(Plan(p) \wedge \\ & properPartOf(x, p, t)) \end{aligned} \quad (108)$$

Goal dependencies A *main goal* (109) is 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):

$$\begin{aligned} \text{mainGoal}(p_1, x, t) =_{df} & \text{properPartOf}(x, p_1, t) \wedge \\ & \text{Plan}(p_1) \wedge \text{Goal}(x) \wedge T(t) \wedge \\ & \neg \exists (p_2) (\text{Plan}(p_2) \wedge \text{properPartOf}(p_2, p_1, t) \\ & \wedge \text{properPartOf}(x, p_2, t)) \end{aligned} \quad (109)$$

$$\begin{aligned} \text{subGoal}(p_1, x, t) =_{df} & \text{properPartOf}(x, p_1, t) \wedge \\ & \text{Plan}(p_1) \wedge \text{Goal}(x) \wedge T(t) \wedge \\ & \exists (p_2) (\text{Plan}(p_2) \wedge \text{properPartOf}(p_2, p_1, t) \\ & \wedge \text{properPartOf}(x, p_2, t)) \end{aligned} \quad (110)$$

It is not necessarily for a subgoal of a plan to be a part of the main goal of that plan. E.g. consider the main goal: *being satiated*; *eating food* can be a subgoal, but it is not a part of *being satiated*. Nonetheless, we can also conceive of an *influence* relation between a subgoal and the main goal of the plan the first goal is a subgoal of (111).

$$\begin{aligned} \text{influenceOn}(x, y) =_{df} & \text{Goal}(x) \wedge \text{Goal}(y) \wedge \exists (z, t) \\ & (\text{Plan}(z) \wedge \text{subGoal}(z, x, t) \wedge \text{mainGoal}(z, y, t)) \end{aligned} \quad (111)$$

InfluenceOn can be used to talk of expected causal dependencies between goals, either within a same or different plans.

By using the previous definitions, we can also define a *disposition* relation (112) between the roles used in a plan having a main goal, and the influenced goal.

$$\begin{aligned} \text{dispositionTo}(x, y) =_{df} & (\text{Role}(x) \wedge \text{Goal}(y) \wedge \\ & \exists (p, g, t) (\text{Plan}(p) \wedge \text{Goal}(g) \wedge \text{mainGoal}(p, g, t) \wedge \\ & \text{uses}(p, x) \wedge \text{influenceOn}(g, y)) \end{aligned} \quad (112)$$

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

In interesting cases, supergoals can be created in order to support the adoption of a subgoal. In order to describe these cases, we need an adoption relation for either plans (114) or goals (113).

$$\begin{aligned} \text{adoptsGoal}(a, g, t) =_{df} & \text{shares}(a, g, t) \wedge A(a) \wedge \text{Goal}(g) \\ & \wedge \forall (p, z) (\text{Plan}(p) \wedge \text{Task}(z) \wedge \text{uses}(p, z) \wedge \\ & \text{properPartOf}(g, p)) \rightarrow \text{adoptsPlan}(a, z, t) \end{aligned} \quad (113)$$

$$\text{adoptsPlan}(a, p, t) \rightarrow \text{shares}(a, p, t) \wedge A(x) \wedge \text{Plan}(p) \quad (114)$$

Adoption is a kind of sharing, but not a kind of assuming: assuming concerns beliefs, and not executions. From that viewpoint, the BDI (Belief-Desire-Intention) paradigm is not distinctive enough: when some agent adopts a plan, that agent might believe the (meta-fact) that execution will be appropriate if complying to

the adopted plan.

But even in that case, adoption is different from assumption, because in the latter case the assumed description is supposed to be directly satisfied by the believed situation. This is not the case in adoption: an agent does not ‘believe’ an executed situation or its resulting goal situation, but (maybe) the possibility of its execution or outcome.

In 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*). These cases can be represented by interlacing adopted goals with influences between them.

Executions Plan executions (115) are situations that proactively satisfy a plan, meaning that plan sharing time precedes (anticipates) its execution time:

$$\begin{aligned} PlanExecution(s) =_{df} & S(s) \wedge \exists(p)(Plan(p) \wedge satisfies(s, p) \wedge \\ & \exists(t_1, t_2)(successor(t_1, t_2) \wedge shares(a, p, t_1) \wedge settingFor(s, t_2)) \end{aligned} \quad (115)$$

Axiom 116 formalizes that subplan executions are parts of the whole plan execution.

$$\begin{aligned} \forall(p_1, p_2, s_1, s_2)((Plan(p_1) \wedge Plan(p_2) \wedge properPartOf(p_2, p_1) \wedge \\ PlanExecution(s_1) \wedge PlanExecution(s_2) \wedge satisfies(s_1, p_1) \wedge \\ satisfies(s_2, p_2)) \rightarrow properPartOf(s_2, s_1) \end{aligned} \quad (116)$$

A goal situation is a situation that satisfies a goal:

$$GoalSituation(x) =_{df} S(x) \wedge \exists(y)(Goal(y) \wedge satisfies(x, y)) \quad (117)$$

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

$$\begin{aligned} GoalSituation(x) \rightarrow \neg \forall(y, p, s, t)((Goal(y) \wedge Plan(p) \wedge \\ PlanExecution(s) \wedge satisfies(x, y) \wedge properPartOf(y, p) \wedge \\ satisfies(s, p)) \rightarrow properPartOf(x, s, t) \end{aligned} \quad (118)$$

In other words, it is not true in general that any situation satisfying a goal is also part of the situation that satisfies the overall plan. This can account for the following cases:

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

3.6 Norms

Norms are treated here in their social (including legal as a special case) sense, as some specification of a conceptualization whose objective is regulatory. On the other hand, the very idea of a regulation is far from clearcut, and we prefer to

delimit our area of interest to the relations between aspects of social norm *usage*, agents' plans, and agent collectives.

We follow the major distinction proposed by Searle [21] between regulative and constitutive norms. Constitutive norms create social individuals, and can contain very few or no regulation. Here we deal only with regulative norms.

In [35], a norm is defined as a description which a case can fall under, if the entities in that case are properly classified by the concepts (roles, tasks, parameters) used in the norm. Differently from plans, which are *executed* in novel situations, norms are satisfied in a more complex and indirect way, which we formalize in the *fallsUnder* relation (129). Norm execution is limited to the cases represented in 130, and is anyway dependent on plan execution.

Norm aspects Firstly, we must consider at least three different aspects, in which norms relate to plans:

1. Norms as conventions out of existing practices or plans (see axiom 121). A norm, either social or legal, usually reflects an existing good practice within a community. Typically, social and legal systems are the main way to maintain a stability within the members and resources of a community, population, or country, and that stability is dynamically addressed by evolving practices (shared plans), either in a positive form (norm creation), or in a negative one (norm deletion or update). This aspect of norms makes them contributions to social engineering, i.e. to the creation of social reality as reflecting either ideology or existing practices.

For example, a legal speed limit is supposed to encode the social practice of avoiding excessive speed on vehicles that can be dangerous for people. Therefore, a legal speed limit provides constraints to any plan execution that requires driving at a certain speed. A plan, on its turn, is supposed to encode the way a desire can be realized. Such ways can be limited by existing norms. For example, the execution of my intention to drive fast for my pleasure, or to arrive somewhere in a short time, has constraints coming from a legal speed limit.

Norms as conventionalized practices depart from plans because they are not executed, and do not have an inherent goal. When (specially in legal domains) there is a talk about the goal or objective of this aspect of a norm, we assume that that goal is actually the goal of social regulation politics, aimed at enforcing established practices that are negotiated appropriately to a community dynamics.

Alternatively, politics or authoritative social ruling may enforce norms that are not established practices: in that case, the norm is imposed, and the relationships with the other aspects are affected (see section).

2. Norms as compliance checking protocols over social behavior or legal cases (see axiom 123). Once a norm lifecycle is established, norms can be enforced by using them as filters for social behavior. Typically, the initiative for compliance checking is limited to the interest of other parties with respect to the behavior of one party.

Norms as case descriptors depart from plans because they are not executed, and do not have an inherent goal. On the contrary, the goal involved in this aspect of norms is the goal of an *interpretation* plan that is aimed at finding a social (or specifically legal) framework to a case or behavior. Typical examples of this aspect of norms include e.g. investigations about the (social or

legal) responsibility for a certain event that caused damage to someone else. When checking the compliance of visible behavior, in many cases it is necessary to attribute plans to agents' behavior. For this reason, also the plans that can be assumed as being executed in those cases are also involved in this aspect of norms. Moreover, sometimes we must attribute norms as parts of agents' plans, and therefore to assess if and how agents' plans differ from the norm they were expected to follow within those plans.

3. Norms as behavioral rules (constraints) within plans (see axiom 124). Once a norm lifecycle is established, and appropriate enforcing and compliance checking practices emerge, they can be used by social agents as constraints within their own plans. In this sense, norms are akin to *behavioral principles*. If taken as principles for social behaviour, norms can be executed, similarly to plans, and in fact they are executed as subplans.

Norms as constraints within plans depart from plans only because their goal is dependent on the main goal of an agent that shares the plan they are part of. Typical examples of this aspect of norms include e.g. the assumption of agent's knowledgeability of a norm within a certain community. These aspects of norms evidence the mutual dependency between plans and norms. Plans are constrained by norms, norms encode conventional plans, and are supposed to constrain agents' plans.

Basic axiomatization of norms Aspects of norms are axiomatized here as additional axioms to the class of norms, because their complementarity makes them parts of a unique ontology design pattern for normative descriptions. A norm is assumed as a description, similarly to plans (119).

$$Norm(x) \rightarrow D(x) \quad (119)$$

Norms are disjoint from plans (120), following the rationale we have given in the previous section.

$$Norm(x) \rightarrow \neg Plan(x) \quad (120)$$

The three aspects of norms are axiomatized as follows: conventionalized practices (121, 122), case descriptors (123), constraints within plans (124). From the viewpoint of a conventionalized practice, a norm uses concepts that are defined by a (usually precedent) social practice (121). In addition, norms are parts of plans that address a community, and use the norms to maintain some kind of social equilibrium (122).

$$Norm(x) \rightarrow \exists(p)(Plan(p) \wedge \forall(c)((C(c) \wedge uses(x, c)) \rightarrow defines(p, c)) \quad (121)$$

$$Norm(x) \rightarrow \exists(p, kc)(Plan(p) \wedge KC(kc) \wedge properPartOf(x, p, t) \wedge involves(p, kc, t)) \quad (122)$$

From the viewpoint of case descriptors, a norm is used by an agent to interpret an agent's behavior (as a plan execution), so providing a redescription for it.

$$Norm(x) \rightarrow \exists(pe, a_1, a_2, t_1, t_2) (PlanExecution(pe) \wedge A(a_1) \wedge A(a_2) \wedge setting(a_1, pe, t_1) \wedge shares(a_2, x, t_2) \wedge redescrives(a_2, pe, t_2) \wedge involves(x, pe, t_2)) \quad (123)$$

From the viewpoint of plan constraints, a norm is a proper part of agents' plans.

$$Norm(x) \rightarrow \exists(p, a, t)(Plan(p) \wedge A(a) \wedge shares(a, p, t) \wedge properPartOf(x, p, t)) \quad (124)$$

Norm application and falling under Having clarified the multiple aspects of social norms, now we can distinguish norm satisfaction, called *appliedIn* (126), obtained by a class of situations called *norm applications* (125), from what we call *falling under* (129). A social situation (case, behavior, etc.) falls under a norm when an agent applies that norm in order to redescribe the situation, either in positive (compliance) or negative (non-compliance) terms.

Formally, a plan execution can: a) be outside the scope of a norm application (127); b) fall under a norm (129); c) execute a norm (130). In the first case, the plan is not constrained by the norm. In the second case, the plan execution is a case for the norm. In the third case, the plan explicitly includes the norm as a proper part.

The *fallsUnder* relation holds between plan executions and norms, and is defined by composing the *hasInScope* relation, and other relations, including *appliedIn*.

$$NormApplication(x) =_{df} Situation(x) \wedge \exists(y)(Norm(y) \wedge satisfies(x, y)) \quad (125)$$

$$appliedIn(x, y) =_{df} Norm(x) \wedge NormApplication(y) \wedge satisfies(y, x) \quad (126)$$

$$outsideTheScopeOf(x, y) =_{df} \neg hasInScope(y, x) \quad (127)$$

$$outsideTheScopeOfNorm(x, y) =_{df} outsideTheScopeOf(x, y) \wedge PlanExecution(x) \wedge NormApplication(y) \quad (128)$$

$$fallsUnder(x, y) =_{df} PlanExecution(x) \wedge Norm(y) \wedge \exists(z) (hasInScope(z, x) \wedge satisfies(z, y)) \quad (129)$$

$$executesNorm(x, y) =_{df} PlanExecution(x) \wedge Norm(y) \wedge \exists(p, pe, t)(Plan(p) \wedge properPartOf(y, p, t) \wedge executes(pe, p) \wedge properPartOf(x, pe, t)) \quad (130)$$

Since norm aspects are complementary and interrelated, also norm axioms and relations are entrenched in interesting ways. For example, norm applications may include different situations, e.g. applying a legal norm to a case requiring an up-to-date interpretation based on current social practices; applying a social norm to overstate the inappropriate behavior of an agent; applying a norm as a principle in regulating an agent's friendship relations, etc.

These different examples of norm application can be made complementary in some cases; e.g. applying an ethical principle with friends may be adopted or not depending on the source of the norm, and if that source is a conventionalized practice in a community whose members share similar descriptions of the social world, that adoption can be easier for, or considered advisable by, a larger number of agents.

3.7 Collectives

We have introduced, so far, two notions of *agent collection*. The first is a simple *collection of agents* (either social or physical), which are unified by some rationale that is extrinsic with respect to the knowledge shared by the agents, e.g. the collection of all agents that use to drink beer, or have green eyes, or the collection of all mosquitoes. The second notion, more relevant for this work, is *knowledge community* (*KC*, cf. 68), which is a collection of social agents that is unified by descriptions that are shared by the members, e.g., the community of semantic web researchers.

We have also suggested that knowledge communities can be based on *paradigms* (76) whose descriptions are *assumed*, not only shared, specially in scientific communities and communities of practice. The notion of community of practice leads us to a more complex notion of collective, which is based on sharing (or assuming) ways of doing things, i.e. plans and workflows.

Therefore, in this section we augment the notion of knowledge community with more types, based on more specific descriptions that agents can share: plans and norms. We call *intentional collectives* (see definition 131) those knowledge communities that are unified by a plan; while we call *intentional normative collectives* (see definition 134) those knowledge communities which are unified by plans that, in turn, are entrenched with norms, according to the aspects described and formalized in section 3.6. Finally, we introduce *knowledge collectives* (see definition 138) as those intentional normative collectives that also share an *epistemic workflow* (see axiom 137), in order to exchange or modify knowledge.

Social and physical agents in collectives In knowledge communities, covering and characterizing concepts classify social agents. How to talk of knowledge communities and collectives whose members are physical agents? From a constructivist's viewpoint, this move is not necessary; in fact, since social agents need to be acted by physical agents (or cognitive systems), every non-empty knowledge community will be eventually enacted by physical agents. From the grounded construction principle (97), we can infer that, whenever we talk of a knowledge community, the possibility is created of having physical agents that construct internal representations of descriptions shared by social agents. Therefore, in the following we do not make any attempt to distinguish knowledge communities whose members are social from those whose members are physical, since the dual nature of the construction principle (both in the social and grounded version) guarantees expressive means as well as correct inferences.

Intentional collectives As proposed in [3], collective action can only originate from the adoption of a common action schema, i.e. from the unification of a collective by means of a plan. Intentional collectives are defined in a straightforward way in (131) as knowledge communities unified by plans.

$$\text{IntentionalCollective}(x) =_{df} KC(x) \wedge \exists(p)(\text{Plan}(p) \wedge \text{unifies}(p, x)) \quad (131)$$

Organizations, teams, task forces, governments, committees, etc. can be modeled as social agents that are acted by intentional collectives unified by shared plans.

For example, we can generalize over *organizations* as in axiom 132.

$$\begin{aligned} \text{Organization}(x) \rightarrow A(x) \wedge \forall(y, t)(\text{actsFor}(y, x, t) \rightarrow \\ \exists(k)(\text{memberOf}(y, k, t) \wedge \text{IntentionalCollective}(k))) \end{aligned} \quad (132)$$

On the other hand, organizations and most free associations are also based on rules. In the next section, we show how this aspect of complex social agents can be modeled.

Intentional normative collectives Plans can be framed in a wider descriptive context, including regulations, normative constraints, social relationship types, etc. In that case, collective action results to emerge from the ‘bundle’ of descriptions that unifies the collective. Our notion of *Bundle*, introduced in 3.3, helps us in creating another type of knowledge communities, which are defined as unified by a bundle including entrenched plans and norms.

Firstly, we introduce the notion of *NormPlanBundle*.

$$\begin{aligned} \text{NormPlanBundle}(b) =_{df} \text{Bundle}(b) \wedge \exists(p, n, t)(\text{Plan}(p) \wedge \\ \text{Norm}(n) \wedge \text{properPartOf}(p, b, t) \wedge \text{properPartOf}(n, b, t) \wedge \\ (\text{involves}(n, p, t) \vee \text{involves}(p, n, t) \vee \text{properPartOf}(n, p, t))) \end{aligned} \quad (133)$$

Definition 133 refers to entrenchment of norms and plans as three possible cases:

- A norm involves a plan, i.e. when a norm, which is supposed to rule the behavior of a community, considers some typical plans of the agents specifically. E.g. *a speed limit regulation that promotes specific counteractions for the drivers that attempt to escape the enforcement of the norm*
- A plan involves a norm, e.g. *a driving plan is built in order to avoid the consequences of not complying to a speed limit (as when decreasing speed when seeing a police vehicle)*
- A norm is a part of a plan, e.g. *a speed limit is straightforwardly considered as a parameter in a driving plan shared by an agent*

Now, an intentional normative collective (134) is a knowledge community unified by a NormPlanBundle (133):

$$\begin{aligned} \text{IntentionalNormativeCollective}(x) =_{df} \text{KC}(x) \wedge \\ \exists(y)(\text{NormPlanBundle}(y) \wedge \text{unifies}(d, x)) \end{aligned} \quad (134)$$

Whereas a NormPlanBundle is explicitly stated (‘anticipated’), like in a closed set of tasks that describe, for instance, the possible actions for a social agent, there exists a unique, communicable motivation (the plan defining the tasks) for the collective action.

On the contrary, whereas a bundle is not anticipated, collective action is an *epiphenomenon*, or something that dynamically appears out of local conditions. Here we do not attempt a formalization of epiphenomenic bundles, leaving it to further research.

Knowledge collectives Having introduced norms into the identity criteria for complex agents and their collective action, we are still left with the problem of defining the *knowledge-level* structure of those agents, and how the collectives that act for them can be characterized at that level. In order to do that, we

merge the notion of *Paradigm* (76), with that of *NormPlanBundle*. *Epistemic influences* (definition 137) are *NormPlanBundles* that govern the influence between the agents from a community, with respect to their **core** knowledge, i.e. the collection of their assumed descriptions, or *paradigm* (76)¹³. In order to define it, we need some specific concept types.

A *knowledge role* (axiom 135) is a concept that classifies only information objects that have an epistemic relevance, i.e. they express descriptions that help maintaining the identity of a community by being exchanged, enriched, or revised according to appropriate plans and norms. E.g. the term “force” plays a knowledge role in contemporary physics because the relations and axioms (the *descriptions*) that are assumed when using that term contribute to the stability of contemporary physics’ paradigm (the notion of *paradigm* has been formalized here in 76).

$$\begin{aligned} & KnowledgeRole(x) \rightarrow C(x) \wedge \\ & \forall(i, t)(classifies(x, i, t) \rightarrow I(i) \wedge \exists(d, p)(expresses(i, d, t) \wedge \\ & Paradigm(p) \wedge properPartOf(d, p, t))) \quad (135) \end{aligned}$$

An *agent role* (definition 136) is a concept that can only classify social agents.

$$AgentRole(x) =_{df} C(x) \wedge \forall(a, t)(classifies(x, a, t) \rightarrow A(a)) \quad (136)$$

Epistemic influence (definition 137) is now formalized as a *NormPlanBundle* (definition 133) that necessarily *uses* (axiom 19) at least one knowledge role (axiom 135) and at least one agent role (axiom 136).

$$\begin{aligned} & EpistemicInfluence(x) =_{df} NormPlanBundle(x) \wedge \exists(y, z) \\ & (KnowledgeRole(y) \wedge AgentRole(z) \wedge uses(x, y) \wedge uses(x, z)) \quad (137) \end{aligned}$$

The notion of epistemic influence is very flexible, since it can be used to talk about one agent that is influenced by some knowledge, as well as about two or more agents that mutually influence each other through their individual knowledge.

Based on it, we define *knowledge collectives* are intentional normative collectives whose unifying bundle is an epistemic influence that has a paradigm (76) as part.

$$\begin{aligned} & KnowledgeCollective(x) =_{df} \\ & IntentionalNormativeCollective(x) \wedge \\ & \exists(y, p, t)(EpistemicInfluence(y) \wedge unifies(d, x) \wedge \\ & Paradigm(p) \wedge properPartOf(p, y, t)) \quad (138) \end{aligned}$$

Epistemic influences found most agent interactions. In particular, the involvement in a *social relation* (definition 139) depends on the fact that involved agents are members of a same knowledge collective, i.e. on agreeing on a shared epistemic influence bundle, towards which all agents are accountable. Accountability is not treated here, but it will be axiomatized on the more basic notion of *assumption* (axiom 73), which is used here for axiom 139.

$$\begin{aligned} & SocialRelation(x) =_{df} EpistemicInfluence(x) \wedge \exists(kc) \\ & (KnowledgeCollective(kc) \wedge unifies(x, kc) \wedge \forall(a, t) \\ & (memberOf(a, kc, t) \rightarrow assumes(a, x, t))) \quad (139) \end{aligned}$$

¹³ We remark that we are not interested here in how descriptions are grounded in physical agents, but only in the fact that social agents share those descriptions

Intuitively, if the agents that participate in a social relationship do not comply to the plans and norms that they are expected to assume (according to a given epistemic influence bundle that unifies the relationship), the underlying collective cannot be “brought about” by them.

Brings about (axioms 140 and 141) is a specialized projection of *c.DnS* maximal relation, holding for social agents and knowledge collectives at a certain time, and requires agents to assume an epistemic influence bundle that has to be adopted by the member agents in the collective.

$$\begin{aligned} \text{bringsAbout}(x, k, t) &\rightarrow \text{actsFor}(x, k, t) \wedge \\ &A(x) \wedge \text{KnowledgeCollective}(k) \wedge T(t) \end{aligned} \quad (140)$$

$$\begin{aligned} \text{bringsAbout}(x, k, t) &\rightarrow \exists(y)(\text{assumes}(x, y, t) \wedge \\ &\text{EpistemicInfluence}(y) \wedge \text{unifies}(y, k) \wedge \\ \forall(a, t)(\text{memberOf}(a, k, t) &\rightarrow \text{adoptsPlan}(a, y, t))) \end{aligned} \quad (141)$$

When applied to social relations in general, *brings about* requires participating agents to assume it, not just to adopt it as a plan. In other words, social relations are maintained by knowledge collectives that are brought about by their members.¹⁴

With *c.DnS*, norms, plans, epistemic influence bundles, knowledge collectives, and social relationships, we have got a rich ontology to describe the nature and the behavior of complex social agents like organizations, institutions, corporations, teams, lobbies, movements, etc.

A recap of the main classes of the schematic entities introduced here, with their taxonomy and disjointness axioms, is depicted in Figure 1.

4 Conclusions

In this article, we have presented a formal framework to represent social agents, collectives, plans, norms, and their dependencies. The framework is based on the constructive version of the ontology of Descriptions and Situations (*c.DnS*), which has been applied to the modelling of social reality and information objects in several applications for the semantic web, business interaction, healthcare informatics, digital libraries, etc.

c.DnS provides a complex pattern for social entities, axiomatized as the so-called *construction principle*, and can be extended to represent a *grounded* version of the principle. Plans and norms are represented here as extensions of *c.DnS*, and a typology of collectives is defined on top of these extensions.

Social agents are taken as primitives, independently from their embodiments as cognitive systems, organisms, robots, etc. Embodiments can be represented in

¹⁴ The notion of social relation proposed here may be perceived as having a ‘scientific’ flavor, because we are proposing that agents in a social relationship assume a common paradigm, i.e. a bundle of plans and norms. But this common paradigm should be taken as minimal as possible, and subject to continuous revision, on the basis of the dynamics that operate on the agents. Therefore, no special claim on the stability or scientific foundedness of a social relationship is made here.


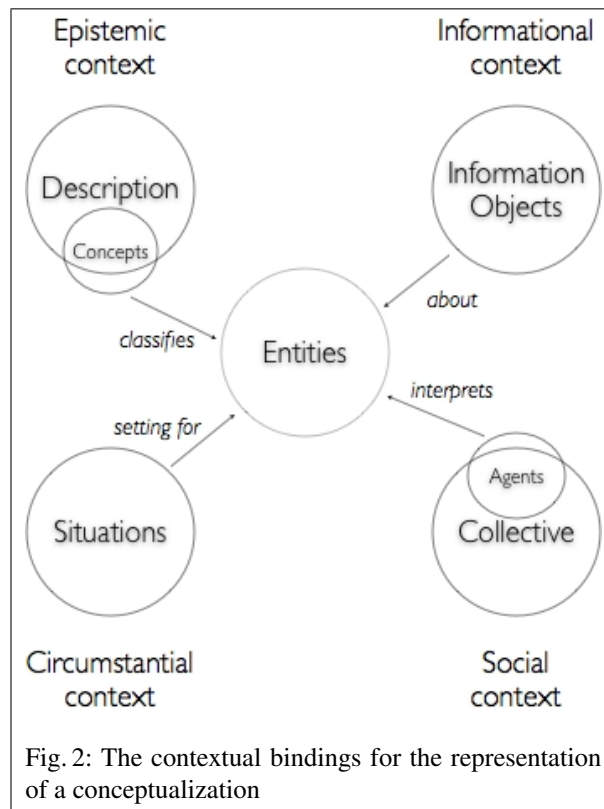


Fig. 1: The main classes of schematic entities, with their taxonomy (sub-class of, bold arrows) and disjointness axioms (red light arrows)



the grounded version of *c.DnS*. This move avoids the typical multi-hierarchies generated by classes that can be agentive or non-agentive, depending on local commitments.

Information and knowledge play a major role in *c.DnS*, and have constructive counterparts, i.e. so-called *information objects*, e.g. texts, images, etc., and *descriptions*, i.e. reified relations between entities. Constructivism (knowledge is inherent in *thought collectives* that create *knowledge paradigms*) is represented by positing collectives of agents that share descriptions expressed by information objects. These descriptions use concepts to classify entities within situations. All these notions: collective, agent, description, information object, concept, situation, are first-class citizens in our ontology.

In other words, *c.DnS* allows to represent the contextual binding of conceptualizations on the circumstantial level (via situations), on the cognitive level (via descriptions and concepts), on the social level (via collectives and agents), and on the informational level (via information objects) (see Figure 2).

Plans and norms have been introduced on top of *c.DnS*, as description types. Plans and norms are shown to be disjoint classes, but entrenched within ‘bundles’ of descriptions that provide formal unification criteria for intentional normative collectives.

Finally, we have used our formal apparatus in order to model the interdependencies of social agents, collectives, plans, and norms against the background of *c.DnS*. The main outcome can be considered a novel ability to talk -in a first-order theory- about complex agents like organizations and institutions, together with their knowledge-level structure (knowledge collectives, unifying plans and norms, paradigms), as well as arriving at a more sophisticated formal notion of social relationships as immersed into (and therefore dependent on) knowledge collectives.

Future plans include a finer-grained classification of collectives, based on more common sense distinctions, like:

- the type of their members (e.g. physical persons, boys, cows, left-handers);
- their knowledge domains (e.g. genetic, taxonomic, epidemiological);
- their related social practices (e.g. neighborhood, geographic, ethnic, linguistic, commercial, industrial, scientific, political, religious, institutional, administrative, professional, sportive, interest-based, stylistic, devotional);
- the ways members of collectives explicitly interact with the description bundles that are expected to unify their collectives (e.g. complete or partial adoption, external control and distribution of accountability, emergence, negotiation, trustfulness);
- the causal relations that characterize them, according to the set-up of the *DnS*-based treatment of causal relations presented in [41].

Another item for our research agenda is investigating the assumptions inherent in relevant theories of action, collective intentionality, plural reference, etc., and describing them according to the ontology introduced here.

Acknowledgments

Part of the research presented in this paper (in particular, the notions of epistemic influence, knowledge communities, and related ones) was developed under the

auspices of the NeOn project¹⁵ funded by the European Commission within the 6th IST Framework Programme.

References

1. : OWL Web Ontology Language Family of Specifications. <http://www.w3.org/2004/OWL> (2004)
2. Oberle, D., Mika, P., Gangemi, A., Sabou, M.: Foundations for service ontologies: Aligning OWL-S to DOLCE. In: Proceedings of the World Wide Web Conference (WWW2004). Volume Semantic Web Track. (2004)
3. Bottazzi, E., Catenacci, C., Gangemi, A., Lehmann, J.: From collective intentionality to intentional collectives: an ontological perspective. *Cognitive Systems Research - Special Issue on Cognition and Collective Intentionality* 7 (2006) 2–3
4. Cocchiarella, N.: Denoting concepts. Reference and the logic of names, classes as many, groups and plurals. <http://www.formalontology.it/essays/plurals.pdf> (2004)
5. King, P.: The Metaphysics of Peter Abelard. In K., G., ed.: *The Cambridge companion to Abelard*. Cambridge University Press, Cambridge, U.K.; New York (2004)
6. Russell, B., Whitehead, A.N.: *Principia Mathematica*. Cambridge University Press, Cambridge, UK (1910)
7. Zeman, J.J.: Peirce on Abstraction. *The Monist* 65 (1982) 211–222
8. Dauben, J.W.: *Georg Cantor: His Mathematics and Philosophy of the Infinite*. Princeton University Press, Princeton (1979)
9. Dugac, P.: *Richard Dedekind et les fondements des mathématiques*. J. Vrin, Paris (1976)
10. Link, G.: The Logical Analysis of Plurals and Mass Terms: A Lattice-theoretical Approach. In von Stechow, A., ed.: *Meaning, Use and Interpretation of Language*. Walter de Gruyter, Berlin (1983) 302–323
11. Marcus, R.B.: Classes, Collections, Assortments, and Individuals. In Marcus, R.B., ed.: *Modalities: Philosophical Essays*. Oxford University Press, New York (1993) 90–100
12. Simons, P.: *Parts: a Study in Ontology*. Clarendon Press, Oxford (1987)
13. Devlin, K.: *The Joy of Sets: Fundamentals of Contemporary Set Theory*. Springer-Verlag, New York (1993)
14. von Neumann, J.: An Axiomatization of Set Theory. In Van Heijenoort, J., ed.: *From Frege to Gödel; a source book in mathematical logic, 1879-1931*. Harvard University Press, Cambridge, MA (1967)
15. Searle, J.R.: *Intentionality. An Essay in Philosophy of Mind*. Cambridge University Press, Cambridge, UK (1983)
16. Gilbert, M.: *Social Facts*. Princeton University Press, Princeton, New Jersey (1992)
17. Bratman, M.E.: Shared Cooperative Activity. *The Philosophical Review* 101 (1992) 327–341
18. Tuomela, R.: Collective Acceptance, Social Institutions, and Social Reality. *The American Journal of Economics and Sociology* 62 (2003) 123–165

¹⁵ <http://www.neon-project.org>

19. Wooldridge, M.J.: Reasoning about rational agents. MIT Press, Cambridge, Mass.; London (2000)
20. Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Ferrario, R., Gangemi, A., Guarino, N.: Social roles and their descriptions. In: KR. (2004) 267–277
21. Searle, J.R.: The construction of social reality. Free Press, New York (1995)
22. Hart, H.: The Concept of Law. Oxford: Clarendon Press (1961)
23. Kelsen, H.: General Theory of Norms. Oxford University Press (1991)
24. Boella, G., van der Torre, L., Verhagen, H.: Introduction to normative multiagent systems. Computational and Mathematical Organization Theory **12** (2006) 71–79
25. Gangemi, A., Mika, P.: Understanding the semantic web through descriptions and situations. In: CoopIS/DOA/ODBASE. (2003) 689–706
26. Gangemi, A., Borgo, S., Catenacci, C., Lehmann, J.: Task Taxonomies for Knowledge Content. Deliverable D07 of the Metokis Project. http://www.loa-cnr.it/Papers/D07_v21a.pdf (2005)
27. Gangemi, A., Catenacci, C.: A constructive ontology of descriptions and situations. Technical report, ISTC-CNR (2006) <http://www.loa-cnr.it/TR/ConstructiveDnS.pdf>.
28. Masolo, C., Gangemi, A., Guarino, N., Oltramari, A., Schneider, L.: WonderWeb EU Project Deliverable D18: The WonderWeb Library of Foundational Ontologies. <http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf> (2004)
29. Cohen, R., Schnelle, T.E.: Cognition and Fact - Materials on Ludwik Fleck. Reidel, Dordrecht (1986)
30. Stanzione, M.: Epistemologie Naturalizzate. Bagatto, Roma (1990)
31. Churchland, P.S., Ramachandran, V.S., Sejnowski, T.J.: A critique of pure vision. In Koch, C., Davis, J., eds.: Large scale neuronal theories of the brain. MIT Press, Cambridge, Mass. (1994)
32. Light, P., Butterworth, G.e.: Context and cognition: Ways of learning and knowing. Lawrence Erlbaum Associates, Hillsdale, NJ (1992)
33. Moore, M.S.: Legal Reality: A Naturalist Approach to Legal Ontology. Law and Philosophy **21** (2002)
34. Karmiloff-Smith, A.: Précis of ‘Beyond modularity: A developmental perspective on cognitive science’. Behavioral and Brain Science **17** (1994)
35. Gangemi, A., Sagri, M.T., Tiscornia, D.: A constructive framework for legal ontologies. In: Law and the Semantic Web. Volume LNCS 3369. Springer (2005) 97–124
36. Köhler, W.: Gestalt Psychology. Liveright, New York (1947)
37. Talmy, L.: Toward a Cognitive Semantics. MIT Press, Cambridge, Mass. (2003)
38. Quine, W.: On what there is. In: From a Logical Point of View. 2nd edn. Harvard University Press, Cambridge, MA (1980)
39. Fleck, L.: The Problem of Epistemology. In Cohen, R., Schnelle, T., eds.: Cognition and Fact - Materials on Ludwik Fleck. Reidel, Dordrecht (1936/86) 79–112
40. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A.: Understanding top-level ontological distinctions. In Stuckenschmidt, H., ed.: Proceedings of the IJCAI Workshop on Ontologies and Information Sharing, 2001. (2001)
41. Lehmann, J., Gangemi, A.: An ontology of physical causation as a basis for assessing causation in fact and attributing legal responsibility. Artificial Intelligence and Law DOI **10.1007/s10506-007-9035-3** (2007)