

Tracking Agency

Ángeles Eraña, Benjamin Q. Sylvand, Dario Taraborelli*

January 10, 2005

Abstract

The aim of this paper is to propose an extension of the *object file* notion to the study of nonconceptual individuation of *agents*. Robust evidence in both developmental and empirical psychology supports the hypothesis that dedicated perceptual mechanisms mediate sensitivity to *objecthood* as opposed to *animacy*. Object individuation and perceptual animacy have been largely studied in recent literature, whereas little is known about mechanisms mediating individuation and tracking of perceptual entities endowed with agency (or *proto-agents*). By introducing a notion of an *agent file*, we aim to provide a theoretical framework for more constrained empirical investigations into the ability to perceptually track agency (prior to explicit identification/categorization) and into asymmetries with regard to ordinary object tracking. We insist, in particular, on the need to understand those mechanisms that underlie the ability to represent agents' *persistence* beyond mere agency cue detection.

Keywords: objecthood; proto-objects; agency; animacy; object files; multiple object tracking; individuation; indexicality; nonconceptual content.

* *Acknowledgments:* We wish to thank participants in the 2004 Oleron Workshop on Objects and Reference as well as Dick Carter, Roberto Casati and Erika Marchetto for valuable suggestions on earlier versions of this paper. Ángeles Eraña's work was supported by the research project CONACYT number 41196, "Filosofía de las prácticas científicas". Dario Taraborelli's work was partly supported by the European Commission IST-2002-002114 "Enactive" Network of Excellence. The *playing footsie* picture (p.9) is artwork by Stefano Mandracchia.

Address for correspondence: Institut Jean-Nicod, 1bis avenue de Lowendal, F-75007 Paris.

Email: taraborelli@ens.fr

Contents

1 Individuation: Creating and Maintaining Reference	3
1.1 Reference to Objects <i>vs.</i> Reference to Agents in Infants	3
2 Perceptual Tracking of Individuals	5
2.1 Object Files and Their Dynamics	6
2.2 From Object Files to Infants' Representation of Objects	7
2.2.1 Proto-objects <i>vs.</i> commonsense objects	8
3 Perceptual Tracking of Agency	8
3.1 Extending the File Notion to Agents	10
3.1.1 Agent Files <i>vs.</i> Agency Cues	11
4 Object <i>vs.</i> Agent Tracking: Empirical Directions	12
4.1 Object <i>vs.</i> Agent File Fixation Cues	12
4.2 Object <i>vs.</i> Agent File Preserving Properties	14
4.2.1 An Objectual Bias in Agent Files Preservation?	15
4.3 Object <i>vs.</i> Agent File Content	16
5 Are There Really Agent Tracking Mechanisms?	17
6 Conclusions	18
References	20

List of Figures

1	Playing 'footsie' and agent individuation	9
2	The Asteroids Game: objects <i>vs.</i> agents	13

This is a preliminary draft

Please do not quote/circulate without the permission of the authors

1 Individuation: Creating and Maintaining Reference

A large number of cognitive skills rely on the perceptual ability to single out individuals. In order to interact with, ascribe properties to, or reason about particular entities, we need to be able to *pick out* individuals, *establish* a referential link with them and *maintain* it over time. Many authors have pointed out that such ability must be grounded in the deployment of some nonconceptual skills: for picking out and maintaining reference to an individual, simply detecting some of its perceptual properties is not sufficient. We need to postulate a referential mechanism that provides a direct (nonconceptual and unmediated) link to this individual in order to be able to ascribe perceptual properties to it or to make perceptual judgments about it.

1.1 Reference to Objects *vs.* Reference to Agents in Infants

There is large evidence in the developmental literature indicating that, from early on, infants are able to deal with two distinct classes of individuals.

On the one hand, they are able to interact with and reason about *objects*, i.e. individuals that behave according to physical constraints. Understanding the observable behavior of objects requires — at least *prima facie* — a grasp of some of the principles underlying physical phenomena.

On the other hand, infants can interact with and reason about *agents*, i.e. entities endowed with intentionality, whose observable behavior cannot be reduced to those constraints governing objecthood. Being able to detect agency requires being sensitive to some specific observable cues of intentionality, such as purposefulness.

A prevailing working hypothesis in developmental studies claims that, when dealing with entities belonging to each of these domains (objects *vs.* agents), different individuation mechanisms are at work. Paul Bloom (Bloom, 2004) asserts that infants are *commonsense dualists*: ‘who have two ways of looking at the world: in terms of bodies and in terms of souls’ (p.191). Kuhlmeier et al. (2004) have designed and carried out a number of experiments to support this idea: infants display precocious abilities to distinguish inanimate objects (entities that behave according only to physical constraints) from animate entities (e.g. humans). These abilities suggest that infants’ early understanding of animate entities does not rely on the typical individuation principles and constraints on which the individuation of objects depends (p.7).

Most studies have focused so far on the ability to distinguish animacy cues from objecthood cues, assuming that this distinction - resulting in significantly different observable patterns - is one of the main sources of infants’ early knowledge

about animate beings as opposed to inanimate entities. Yet little is known about how the animate vs. inanimate distinction relates to another more fundamental one, that of agents vs. objects. As the authors of this study (Kuhlmeier et al. 2004) themselves explicitly acknowledge, it is debatable how empirical data should be interpreted with respect to this distinction:

It [...] remains an open question whether the results of the present study are due to a distinction between animate versus inanimate entities, intentional agents versus non-intentional objects, or humans versus other entities. (*ibid.*, p.7)

If on the one hand, infants' perceptual representation of objects has been largely studied, little is known, on the other hand, about infants' perceptual representation of *agents*. Spelke (Spelke, 1994) has suggested a number of principles circumscribing the notion of perceptual objecthood in infants. These principles can be formulated, following Bloom (2004, p.12), as:

1. *Cohesion*. Objects are connected masses of stuff that move as a whole. If you want to know where the boundaries of an object are, an easy test is to grab some portion of stuff and pull—what comes with what you are pulling belongs to the same object; what remains does not.
2. *Solidity*. Objects are not easily permeable by other objects; if you tap at an object with your finger, your finger does not penetrate.
3. *Continuity*. Objects move in continuous paths; they travel through space without gaps. An object would violate this rule if it disappeared from one location and reappeared in another.
4. *Contact*. Objects move through contact. A ball on a pool table is not going to move unless something contacts it; it will not run from the cue or come when it is called. The exceptions to this rule are animate creatures, like people and dogs, and also certain complex artifacts, such as robots and cars.

Taken together these principles define what counts as an 'object' for children. More precisely, they characterize a cluster of properties that children systematically privilege in picking out a certain kind of entities in the environment, prior to any conceptual identification. For this reason such clusters (often referred to as 'Spelke Objects') have been described as *proto-objects*.

An analogous characterization of what counts as a perceptual 'agent' for infants has not yet been clearly proposed. We argue that a principled inquiry into infants'

ability to distinguish objects from agents cannot be based on mere sensitivity to animacy vs. inanimacy cues, since the ability to perceptually individuate and track agents relies on a much more fundamental distinction: the ability to individuate and track *bearers of intentionality* as opposed to objects.¹

Are there, beyond mere animacy sensitivity, perceptual mechanisms for tracking agency? To which extent can agents (as opposed to objects) be perceptually individuated, parsed and tracked by children?

We submit that understanding the infant's perceptual representation of agents might benefit from introducing the notion of a *proto-agent*, i.e. a cluster of properties that determine the way in which infants single out intentional entities prior to conceptual identification.

In order to achieve this aim, we will propose an extension of the object individuation paradigm that relies on Pylyshyn's proposal for visual object tracking. To construe a notion of *proto-agent*, we will assume:

- (a) following Pylyshyn (2000) and Pylyshyn (2001), that human beings individuate and track single entities via a nonconceptual mechanism which can be operationalized by appealing to the notion of an *object file*;
- (b) following Carey and Xu (2001), that the same nonconceptual mechanisms underly infants' early understanding of objects and adults' ability to keep track of perceptual entities;
- (c) according to the evidence mentioned above, that there are plausibly different processes that allow children to individuate agents as opposed to objects.

Our main contention is that while tracking perceptual entities endowed with agency (or *proto-agents*) and their persistence over time, the human cognitive system opens a special sort of file, which we will call an *agent file*. Furthermore, we will suggest that the same mechanisms underlie infants' individuation of agents and adults' ability to perceptually keep track of agents before explicit identification.

2 Perceptual Tracking of Individuals

Nonconceptual mechanisms for individuating and tracking objects have been largely studied in the perceptual literature. In particular, humans' abilities to visually track individual entities have been extensively studied in recent years by

¹We use hereafter the term 'tracking' in a technical sense to refer to the perceptual ability to parse and maintain reference to individual entities (see Pylyshyn, 2000). Issues related to *reidentification* of individuals (as mediated by conceptual representations) are beyond the scope of the present analysis and should be kept distinct from the narrow notion of tracking that we adopt in this article.

Pylyshyn and collaborators [Pylyshyn \(2000\)](#); [Blaser et al. \(2000\)](#); [Pylyshyn \(2001\)](#); [Scholl et al. \(1999, 2001\)](#).

There is robust evidence that a referential link to a particular entity, an individual, is established in virtue of low-level indexing mechanisms (which Pylyshyn dubs ‘FINSTs’) that:

- precede the deployment of focal attention;
- operate before and independently from the perceptual identification of the tracked item;
- are pretty insensitive to major featural modifications in the visual appearance of the tracked items;

These three aspects suggest the *nonconceptual nature* of such referential mechanisms: they allow a subject to keep track of an item without necessarily encoding specific features or attributes of that item. The nonconceptual nature of this referential link has to be considered, according to defendants of the FINSTs theory, as a precondition to any form of identification or perceptual judgment about objects. In this respect, the visual index theory can be considered as an extension and partial amendment of the hypothesis according to which humans access and store information about perceptual objects in terms of ‘files’.

2.1 Object Files and Their Dynamics

The notion of an object file is due to the work of [Kahneman and Treisman \(1992\)](#). They proposed that object perception is mediated by the opening of episodic ‘files’ within which object tokens are constructed. Information about particular perceptual items is thus selected from the sensory array, integrated over time, and stored in such files.

Pylyshyn has suggested that object files can account for the way in which perceptual information is *stored*, but not for the *creation* and *maintenance* of the referential link to the object. Pylyshyn’s model enriches and extends the previous theory at two different levels: first, it introduces a mechanism of nonconceptual reference as a requirement for any theory of perceptual reference to individuals; second, it embeds the traditional object file notion in a larger framework that accounts for the whole dynamics of perceptual items. Following Pylyshyn, hence, we should distinguish three independent classes of properties that are relevant for understanding the dynamics of object tracking, namely:

- (1) INDEX-GRABBING FEATURES
properties that cause the assignment of an index (and enable an object file to be opened for the indexed item);
- (2) INDEX-PRESERVATION FEATURES
properties that allow the indexed item to be tracked (and the file to persist over time);
- (3) ENCODED FEATURES
properties that can be ascribed to the indexed item (information that can be stored in the file).

Taken together, (1)-(3) characterize how ‘files’ work: the *content* of a file, i.e. the information attached to the indexed individual, is defined by properties of type (3); the *dynamics* of the tracking is defined by properties of type (1) and (2). It should be noted that properties belonging to (1) and (2) need not be encoded in the file: they need not be used for identifying entities *as* objects, i.e. they do not provide *per se* any basis for identifying the object or making perceptual judgments about it. Properties belonging to (3), on the contrary, are properties stored in the file, and they can be used for further qualification (categorization/identification) of the object at a conceptual level.

2.2 From Object Files to Infants’ Representation of Objects

Up to now we have been referring to perceptual objects, i.e. individuals picked out by our perceptual systems without benefit from concepts or prior knowledge. [Carey and Xu \(2001\)](#) have argued that adults’ nonconceptual representation of perceptual items and infants’ object representation share a number of characteristics:

- they privilege spatiotemporal information in decisions of individuation;
- they are subject to the same set of size limitations for parallel individuation;
- they survive occlusion and are sensitive to the distinction between cessation of existence and temporary loss of visual contact (*ibid.*, p.186).

In particular, the individuation mechanisms described in developmental psychology literature and those described by the object tracking literature both seem to privilege entities that are bounded, cohesive and that persist over time.

According to these authors, these strong similarities suggest that the account provided by studies regarding how human adults track perceptual objects and the widely discussed results in developmental psychology regarding how infants represent objects are two descriptions of the *same* system. If their proposal is correct,

then it is plausible to assume that the same mechanisms are put to work when adults track perceptual objects and when infants track objectual entities: nonconceptual representations of individual entities involved in both domains can hence be described as the same kind, which we refer to using the notion of a *proto-object* (Pylyshyn, 2004, p.23).

2.2.1 Proto-objects vs. commonsense objects

It should be noted that the notion of nonconceptual representation of objectual entities (proto-objects) need not be relevant for our commonsense understanding of *physical objects*. While Pylyshyn (2004) acknowledges that trackable individual items are typically the proximal counterpart of commonsense physical objects, it is debatable whether this is always the case. We should hence be prudent not to conflate the theoretical notion of proto-object which is relevant for perceptual and developmental psychology with our shared intuitions about objects, which can be shaped by conceptual, linguistic and cultural factors (Casati, 2004).

The legitimacy of the distinction between a theoretical notion of (proto)objecthood and the commonsense notion of an object motivates our claim that the study of *agency* and agent individuation should respect a similar prescription: what we intend to suggest is that - no matter what our shared intuitions are about agents in everyday life, how we individuate them or identify them - there might be a theoretical notion, comparable to the notion of a proto-object, (we might call it the notion of a 'proto-agent') that is relevant for the understanding of nonconceptual individuation and tracking of entities endowed with agency and intentionality. Such a notion should, then, not be conflated with its commonsense counterpart: from now on we will refer to 'agents' *tout court* for this theoretical notion of a 'proto-agent'.

3 Perceptual Tracking of Agency

Picking out and maintaining a referential link to agents seems *prima facie* to raise the very same problems involved in referring to objects. We argue that merely being sensitive to agency cues is not sufficient to account for the way agents are individuated, tracked and referred to. What is needed, much as in the case of object tracking, is to understand how a subject is able not only to detect agency, but to maintain reference to an individual which:

- is unique in spite of multiple agency cues;
- can persist over time;

- can survive to changes in some of its features;
- can cease to exist, split or merge with other entities;
- can be tracked in parallel and independently of other entities of the same kind.

The need for such a mechanism that allows agent tracking can be illustrated through a case like the following:



Figure 1: Playing 'footsie' and agent individuation

PLAYING 'FOOTSIE'. John is invited for dinner at a friend's place. He is seated across from a number of guests. At a certain point in the dinner he notices that something is going on under the table: he detects an agency cue, like someone willing to play 'footsie' with him. After a while, a second agency cue is detected. John has a problem understanding whether there actually is an agent behind the detected cues, and if it there is, whether one and the same agent is responsible for both cues or more than one agent is involved.

Thus, keeping track of agents seems to require some sort of mechanism for the selection of individuals, the creation of a referential link and its maintenance over time. It is unclear, though, whether and how keeping track of agents could be done using the mechanism used for ordinary object tracking. The mechanism involved in object tracking is not triggered by properties such as purposefulness which could plausibly be a necessary condition for parsing an agent as persistent. If this is so, then such mechanism would not detect the agent's willingness

of playing footsie with John and, thus, would not allow him to keep track of the agent behind the agency cue (not to say that it would be helpless for John's understanding of whether one and the same agent is responsible for both, the first and the second, cues). Furthermore, there may be cases in which objectual cues conflict with agency cues and it is hard to establish whether individual entities are tracked in virtue of the former or the latter. The literature on infants' perception of intentionality in self-propelling shapes (Heider and Simmel, 1944; Premack, 1990; Scholl and Tremoulet, 2000) represents a good case study for investigating such issues. Heider and Simmel (1944), for instance, showed that people tend to associate particular intentions or intentionality (doing some good to someone, being good, etc.) to certain systematic ways in which some geometrical figures (circles, triangles, etc.) move while they "tell a story" in a movie. One might ask under which conditions criteria of individuation and perceptual tracking are based on objectual cues rather than agency cues. It is an interesting empirical issue to study whether:

- an item preserving its agency cues and changing its objectual features can still be tracked as the same agent (*e.g.*, an 'evil' triangle turning into an 'evil' square without violation of other constraints or an 'evil' triangle disappearing and reappearing at another place with a violation of spatio-temporal continuity);
- an item preserving its objectual features and changing its agency features can be still tracked as the same object (*e.g.* an 'evil' triangle suddenly turning into a 'good' triangle).

Evidence from such experiments supports the claim that there *are* specific individuation processes that depend on agency which can be dissociated from individuation of the same entities as objects (Bloom, 1996). In what follows we will argue that the question of whether children are able to individuate and refer to agents by appealing to mechanisms similar to those of object tracking might benefit from an extension of the notion of 'file' to the case of intentional agents.

3.1 Extending the File Notion to Agents

An interesting way to operationalize perceptual individuation and the tracking of agents involves an appeal to the notion of file described above, extending it from the domain of (proto-)objects to the domain of (proto-)agents by postulating the notion of an *agent file*.

If the ability to track agents is not reducible to mere sensitivity to agency cues, in virtue of the agent's continuity and persistency over time, it seems reasonable to assume that in order to track such entities and ascribe properties to them a subject might use files. Their dynamics can be then described by three classes of

properties analogous to those we introduced for object tracking.

1. AGENCY-GRABBING PROPERTIES

properties that enable individuation of an agent (i.e., properties that are necessary to open an agent file);

2. AGENCY-PRESERVING PROPERTIES

properties that allow an individual agent to be tracked (and the file to persist over time);

3. ENCODED PROPERTIES

properties that can be stored and retrieved from the agent file.

In the example introduced earlier, such properties could be: John's detection of an animacy cue on his leg (an agency-grabbing property); the spatio-temporal coherence of this stimulation (for instance its spatial orientation or its frequency: agency-preservation features); any other information that can be ascribed to the agent (including non-perceptual information such as beliefs about the agent's explicit intention of playing 'footsie': encoded properties).

The main idea is that files are domain-specific so that the kind of information and properties that an agent file can store is different from that which is stored in an object file: the encoded information in the former type of file can include intentions, which can be retrieved in reasoning or while making judgments about the agent.

3.1.1 Agent Files vs. Agency Cues

It should be noted that by positing the existence of a class of properties that allow an individual agent to be picked out we are not endorsing the idea that as soon as such properties are *detected* an agent file is automatically opened.

In the case of objects, it has been shown that detection of objectual cues is not per se sufficient for tracking an entity: well-formed perceptual objects might still lack conditions for being tracked over time (Scholl et al., 2001). The mere presence of objectual cues does not entail the presence of reference fixation cues.

Similarly, there can be cases in which merely detecting an animacy or agency cue, although sufficient to respond to its presence, does not imply that an agent is individuated and susceptible of perceptual tracking. John can for example withdraw his leg as a reaction to the stimulation, without necessarily individuating an agent. Since agent *tracking* is not required in such cases (either because there is

no need to individuate the source of the cue or because there is no need to represent this entity as persistent over time), we assume that opening an agent file is not needed. Actually, we claim that agency-grabbing properties are necessary conditions for agent tracking but they are not *per se* sufficient for the creation of an agent file. Distinguishing between simple detection of properties and full-blooded perceptual tracking - as in the case of detection of objectual cues vs. object tracking - is a natural consequence of our approach.

4 Object vs. Agent Tracking: Empirical Directions

Introducing the notion of an agent file raises the problem of understanding whether and how such files are related to files that apply to tracking of objects. If we accept the hypothesis that there are two distinct mechanisms for agent vs. object tracking, it is reasonable to ask how such mechanisms can interact or be mutually related in specific experimental conditions. Broadly speaking, there can be two general options:

(A) INDEPENDENCE VIEW

At each level of description, object files and agent files *share no* common features. Properties relevant for describing the two kinds of files and their dynamics (file fixation, preservation and content) are *necessary and sufficient* to account for the two distinct kinds of tracking.

(B) DEPENDENCE VIEW

At some level of description, object files and agent files *might share* common features. Properties relevant for describing the two kinds of files and their dynamics are *sufficient but not necessary* to account for the two distinct kinds of tracking. Some properties of one kind of file can be exploited at a given level by the other kind of file.

We will outline in what follows some cases in which the relation between the two kinds of files can be empirically studied at each level.

4.1 Object vs. Agent File Fixation Cues

Once we acknowledge that object vs. agent tracking can be articulated at three distinct levels, we can raise the question of whether, in the case of tracking², different classes of properties allow *fixing of* reference to an item in the case of agents and in

²We do not intend to discuss here cases of mere *sensitivity* to objectual cues vs. agency cues. See the paragraph in section 3.1.1 on the distinction between detecting and tracking agency

the case of objects. Let us consider an example drawn from a classic arcade game of the '80.

THE ASTEROIDS GAME. A player must drive a space vessel in order to avoid and destroy both asteroids and enemy vessels. Asteroids are characterized by passive physical movements, while enemy vessels are characterized by motor patterns revealing purposeful behavior (like avoiding asteroids, actively changing speed and direction, and shooting other vessels). The player must be able to track both asteroids and enemy vessels and react appropriately to their movements in order to destroy them, survive and win the game.




Figure 2: The Asteroids Game: objects vs. agents

This example illustrates a case in which a subject needs to pick out and maintain reference to two different types of individual (objectual entities, like *asteroids*, vs. intentional entities, like *enemy vessels*). We assume that in order to do this, the subject must detect two classes of cues prior to any further identification or categorization. It seems that, to establish reference, the subject does not need to access properties that might be used to *identify* objects and agents (e.g., asteroids and enemy targets might have the same *shape*). In short, we are claiming that a subject becomes able to individuate and maintain reference to entities belonging to two distinct classes (proto-objects vs. proto-agents) in virtue of his ability to detect objectual vs. agency cues, and of the specific task requirements that force him to maintain perceptual reference to individuals over time. If it can be empirically demonstrated that in similar conditions subjects display robust capabilities to differentially detect reference-grabbing properties belonging to two mutually exclusive classes as a condition for tracking items, then we might plausibly conclude that two distinct and independent individuation processes are at work. This distinction would support the claim that object tracking and agent tracking are *independent* at the level of reference fixation properties.

The existence of two distinct classes of reference fixing properties — although sufficient to support the *independence view* at this level (*agency cues can be segregated from*

objectual cues) — is not sufficient to conclude that the two mechanisms of agent and object tracking are necessarily independent. We must also consider the relation between agents and objects at two other levels, viz. reference preservation and feature encoding.

4.2 Object vs. Agent File Preserving Properties

We have insisted on the fact that fixing perceptual reference is still not enough for *tracking* individual entities over time. Tracking implies *preserving* a referential link to a perceptual item already picked out. This raises the question of determining in virtue of which properties reference to a single individual can be maintained over time.

In the example above, in order to be able to avoid an asteroid, a subject must be able to track it as *persisting* over time. Following Pylyshyn, we assume that the properties used for fixing reference when the individual is picked out *need not* be the same as those that preserve reference. Once an item is individuated, reference fixing properties can be discarded without breaking up the referential link itself. Provided there are some *file preserving properties*, an item can undergo major changes without losing its singularity: we argue that if preservation conditions are met, an individual will not cease to be treated as a single perceptual item even if the properties initially used for its individuation have disappeared. Assuming that an asteroid was picked out as an individual object through its shape, it is not necessary that its shape be maintained over time in order for the object to persist in a perceptual tracking task.

Recent literature on *Multiple Object Tracking* has demonstrated that tracked items can survive several kinds of disruption of their features. It seems, though, that certain properties are required for an item to preserve its individuality. It is reasonable to assume, on the basis of this literature, that in our asteroids game example, although asteroids might ‘survive’ temporary occlusions which do not alter their trajectory, they would fail to maintain a perceptual link in cases of spatio-temporal incoherence, like sudden shrinking or disappearing and reappearing at a different location (Scholl and Pylyshyn, 1999). *Coherence of trajectory* as well as *cohesiveness* (Van Marle and Scholl, 2003) are hence examples of properties that seem to be required in order to maintain reference to perceptual items in tracking tasks.

We might then ask whether the properties used to *keep reference alive* are different in the case of agent or object tracking.

4.2.1 An Objectual Bias in Agent Files Preservation?

Many observable properties might in principle be recruited to preserve agency. For instance, the reiteration of animacy cues (e.g. *an enemy vessel periodically shooting at the player*) or the lack of cues of agency disruption (e.g. *the lack of observable indicators of an enemy vessel being destroyed*) are good candidates for the properties that contribute to the preservation of agent files. This might suggest that at the level of file preservation as well, agents and objects do not share any property (and hence that *object files are independent from agent files*). Nevertheless, given the fact that intentional entities are *usually* coinstantiated with objectual entities ('bodies') in our environment, it is plausible to assume that agents can be perceptually tracked via objectual preservation properties. We call this phenomenon an *objectual bias* in agent files preservation. This amounts to saying that:

- (a) there is a class of properties that are sufficient to maintain reference to an object, once an object file is opened.
- (b) there is another class of properties that are sufficient to maintain reference to an agent, once an agent file is opened;
- (c) among each of the above classes there are single properties that are more or less strongly correlated with preservation of the file: in the case of objects, there might be properties strongly or weakly correlated with object persistence;
- (d) one and the same property can be correlated with different degrees of reliability to agency or objecthood preservation;
- (e) object preservation properties that are weakly correlated with preservation of agency might nonetheless be recruited for preserving agency when other agency preservation cues are absent or not available.

The fact that one and the same property P might be relevant to different degrees for maintaining reference either to an object or to an agent does not threaten *per se* the very possibility that there are two distinct kinds of file. We are just suggesting that it is empirically possible that, in virtue of the robustness of the agent-body correlation in our environment, nonconceptual tracking of agents might exploit objectual properties.

Which classes of properties are actually being used for tracking a proto-agent (whether they are properly described as agency-related rather than object-related) is, thus, an empirical issue that calls for experimental work. We maintain that — until a number of explicit conditions for distinguishing between these two classes

are formulated — empirical results supporting the hypothesis that infants track individual entities as *persistent objects*, as *persistent agents* or as *agents persisting in virtue of their objectual features* might be seriously undermined. As a possible suggestion, one might test whether the fact of using objectual properties for keeping track of agents' persistence has any consequences on the number of items an infant is able to individuate and track at the same time.

4.3 Object vs. Agent File Content

Once the conditions for fixing and preserving an item are met, we have at our disposal a device (a file) to *store* information attached to this individual. The way in which a file attaches properties to an entity is radically different from the way in which the nonconceptual mechanism at work while tracking an individual establishes and preserves reference to it.

On the one hand, nonconceptual tracking provides a *direct link* to an individual, i.e. allows a particular entity to be grabbed and segregated from other individuals or other properties of the scene: tracking is thus a necessary condition for parsing an individual as persisting.

On the other hand, having a file attached to an individual is required in order to ascribe some properties to it.

In the object file literature, a file *content* typically consists of properties that can be *predicated* of an object and used in perceptual judgment, categorization and identification. Following Scholl et al. (1999, p.2):

[This kind] of property determines the object's appearance — what a particular object looks like — including its color, shape, lightness, and texture. We call these *featural properties*.

If our hypothesis on the extensibility of the object file paradigm to the case of agent tracking is valid, then we can think of the content of an agent file as a temporary structure attached to an individual entity which can store information about this intentional entity.

It is an empirical issue to understand whether and what are the particular constraints on properties that can be stored in an agent file as opposed to an object file. Yet, we submit that the extension of the file notion to the case of agency can explain how perceptual agency properties can be *ascribed* to a particular agent and used in perceptual judgment.

When an infant is asked to identify among a number of perceptually available agents *which is* the agent displaying, say, aggressive behavior, we claim that she is making use of information stored in a file to perform this task. It should be noted

that since properties stored in a file are those properties that are conceptually accessible, they need not have any relevance for understanding how the nonconceptual parsing and tracking of individual entities is done.

5 Are There Really Agent Tracking Mechanisms?

The previous paragraphs were meant to outline a number of empirical issues related to the possible distinction of object vs. agent tracking at three different levels: (a) the level of reference-fixing, (b) the level of keeping reference alive and, (c) the level of ascribing properties. It might be objected that our proposal of an extension from the domain of objects to that of perceptual agents is in fact a mere redescription of the functional role of object files and thus the idea that we are able to track agents in virtue of dedicated mechanisms would be brought into question. The proposal underlying such objection can be called the *deflationary view* on agent tracking.

Our reply to this objection can be articulated at different levels.

- The rationale for the existence of agent tracking mechanisms is that if we want to account *not only* for detection of and sensitivity to agency or animacy cues as opposed to objectual cues, but also for the ability to maintain reference to an intentional entity *persisting* over time, then we need to explain how this representation of persistence is achieved.
- From the fact that there might be significant similarities in the dynamics and nature of object vs. agent tracking (what we called the *dependence view*), it does not follow that the latter should be reduced to the former. In the previous paragraph we made some suggestions about possible empirical ways to assess the similarity/difference between the two mechanisms.
- Our proposal is consistent with a large literature in developmental psychology that has demonstrated the existence in children of two distinct domains of perceptual properties: the domain of animacy and the domain of objecthood. Our contribution can be considered as a framework for extending these investigations to the question of how individuals endowed with animacy and agency can be grabbed as entities persisting over time.
- The hypothesis of the existence of agent files and agent-related tracking abilities opens up some interesting research directions leading to the study of possible conflicts and dissociations between object and agent tracking. Empirical research might shed light on the fact that:

- intentional and objectual entities can compete for the same attentional or computational resources: it is possible that the limit on the number of items a subject can track at one time is dependent on (or independent of) the class of tracked items (*objects only*, *agents only* or *objects plus agents*);
- there may be interesting cases of dissociations, i.e. selective impairments of either of these abilities without functional consequences for the other: we might imagine cases of subjects being able to track objectual entities but not agents or viceversa.
- inattentive blindness studies might investigate whether the existence of two distinct classes of entities has any effect on their neglect: this might suggest that the traditional list of ‘styles of attention’ (object-based vs. space-based attention) should be extended to include a third kind of style (agent-based attention);
- developmental investigations might benefit from the notion of an agent tracking mechanism to establish the conditions under which infants represent the *persistence* of a perceptual item over time.

The above considerations suggest that dedicated mechanisms for agent tracking are likely to have psychological reality and account for a number of capacities involved in perceptual reference to persistent entities endowed with agency. The relation of such mechanisms to those involved in perceptual reference to objects remains, however, an open empirical issue deserving further investigation. We have given arguments in favor of a moderate *dependence view* suggesting that in some cases objectual properties might be recruited to establish and maintain reference to agents.

6 Conclusions

A crucial step for understanding our nonconceptual abilities to refer to individual entities consists in explaining how such individuals are *picked out* and *tracked over time*, prior to any form of categorization or conceptualization. Such mechanisms of direct reference to individuals lay probably at the basis of both infants’ abilities to parse objects and adults’ capacity to keep track of multiple perceptual items. We have proposed an extension of the studies on object individuation and tracking to the domain of perceptual individuation and tracking of entities endowed with agency. This extension is intended to fill a gap between the study of perceptual sensitivity to agency cues and a full-fledged understanding of how perceptual items

tracked as agents (what we called *proto-agents*) can persist over time. We have argued that unless some explanation of the mechanisms underlying agent *persistence* is provided, many empirical results concerning infants abilities to track animate entities could be dramatically undermined. Our proposal of three distinct levels that might be involved in agency fixing, preservation and ascription is meant to provide the basic requirements for any explanation of perceptual capabilities to track agents. We argue that by analyzing the dynamics of agents at these three different levels, more principled answers might be given to the question of possible interferences and biases between mechanisms dedicated to agent vs. object individuation.

Institut Jean Nicod
Paris

References

- Blaser, E., Pylyshyn, Z. W., and Holcombe, A. O. (2000). Tracking an object through feature space. *Nature*, 408:196–199. 6
- Bloom, P. (1996). Possible individuals in language and cognition. *Current Directions in Psychological Science*, 50:90–94. 10
- Bloom, P. (2004). *Descartes' Baby. How The Science Of Child Development Explains What Makes Us Human*. Basic Books. 3, 4
- Carey, S. and Xu, F. (2001). Infant's knowledge of objects: beyond object files and object tracking? *Cognition*, 80:179–213. 5, 7
- Casati, R. (2004). Common sense, philosophical and theoretical notions of an object: some methodological problems. In *Barcelona ESPP-SPP Joint Conference*. 8
- Heider, F. and Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology*, 57:243–249. 10
- Kahneman, D. and Treisman, A. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24:175–219. 6
- Kuhlmeier, V., Bloom, P., and Wynn, K. (2004). Do 5-month-old infants see humans as material objects? *Cognition*, 94(1):95–103. 3
- Premack, D. (1990). The infant's theory of self-propelled objects. *Cognition*, 36(1):1–16. 10
- Pylyshyn, Z. W. (2000). Situating vision in the world. *Trends in Cognitive Sciences*, 4:197–207. 5, 6
- Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, 80:127–158. 5, 6
- Pylyshyn, Z. W. (2004). Notes for the oleron summer workshop on objects. 8
- Scholl, B. J. and Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, pages 259–290. 14
- Scholl, B. J., Pylyshyn, Z. W., and Feldman, J. (2001). What is a visual object: Evidence from target-merging in multiple-object tracking. *Cognition*, 80(1-2):159–177. 6, 11

- Scholl, B. J., Pylyshyn, Z. W., and Franconeri, S. L. (1999). The relationship between property-encoding and object based attention: Evidence from multiple object tracking. [6](#), [16](#)
- Scholl, B. J. and Tremoulet, P. (2000). Perceptual causality and animacy. *Trends Cogn Sci*, 4(8):299–309. [10](#)
- Spelke, E. (1994). Initial knowledge: Six suggestions. *Cognition*, 50:431–445. [4](#)
- Van Marle, K. and Scholl, B. J. (2003). Attentive tracking of objects versus substances. *Psychological Science*, 14(5):498–504. [14](#)