This paper has been accepted for publication in *Conferences in Research and Practice in Information Technology*, Vol. 37. J. Weckert and Y. Al-Saggaf, Eds.

# Informational Realism

**Luciano Floridi**
Dipartimento di Scienze Filosofiche, Università degli Studi di Bari;
Faculty of Philosophy; Sub-Faculty of Computation; and Information Ethics Group, Oxford University.
Wolfson College, OX2 6UD, Oxford University, UK
luciano.floridi@philosophy.oxford.ac.uk

## Abstract

What is the ultimate nature of reality? This paper defends an answer in terms of *informational realism* (IR). It does so in three stages. First, it is shown that, within the debate about *structural realism* (SR), *epistemic* (ESR) and *ontic* (OSR) structural realism are *reconcilable* by using the methodology of the levels of abstractions. It follows that OSR is *defensible* from a structuralist-friendly position. Second, it is argued that OSR is also *plausible*, because not all related objects are logically prior to all relational structures. The relation of difference is at least as fundamental as (because constitutive of) any relata. Third, it is suggested that an ontology of structural objects for OSR can reasonably be developed in terms of informational objects, and that Object Oriented Programming provides a flexible and powerful methodology with which to clarify and make precise the concept of "informational object". The outcome is *informational realism*, the view that the world is the totality of informational objects dynamically interacting with each other.

*Keywords*: Structural realism, epistemic structural realism, ontic structural realism, levels of abstraction, informational ontology, object oriented programming.

## 1 Introduction: checkmate in three moves

This paper defends a metaphysical position – *informational realism* (IR) – as a solution to a problem. The problem is deceptively elementary: what is the ultimate nature of reality? The answer is misleadingly simple: it is informational. The cumulative steps from the problem to the solution are the difficult part of the process. To explain them, the analogy of a chess problem comes handy: white checkmates in three moves.

Here is the outline of the strategy. On the board we find the debate about *structural realism* (SR; Maxwell [1970]; Worrall [1989]; Ladyman [1998]). This is the view that knowledge of the world is knowledge of its structural properties. *Epistemic structural realism* (ESR; Morganti [forthcoming]) holds that objects are what remain in principle unknowable once the knowable structures of reality have been factored out. *Ontic structural realism* (OSR; Votsis [forthcoming]) holds that objects are themselves structures.

The first move shows that ESR and OSR are reconcilable, and hence that OSR is *defensible* from a structuralist-friendly perspective.

The second move shows that OSR is also *plausible* because not all related objects are logically prior to all relational structures: no relata are logically prior to the relation of *difference*.

The third move shows that an ontology of structural objects (Chakravartty [2002]) can reasonably be developed in terms of informational objects, and that *Object Oriented Programming* (OOP) provides a flexible and powerful methodology with which to clarify and make precise the concept of "informational object".

The outcome of the three moves is informational realism, the view that the world is the totality of informational objects dynamically interacting with each other. Checkmate.

## 2 The pieces on the board: structural realism and its ontological challenge

Broadly constructed, *Structural Realism* argues that:

SR) instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) can be, in the best circumstances, increasingly informative about the *relations* that obtain between the (possibly unobservable) objects that constitute the system under investigation (through the observable phenomena).

SR leaves unspecified the nature of the *relata* in the structures. This problem has resuscitated the classic question about the knowability of the ontological status of objects. Consider the following formalism. The scheme of a structure $S$ comprises four sets:

1) a non-empty set $O$ of objects (the domain of $S$),

2) a non-empty set $P$ of first-order, one-place properties of the objects in $O$,

3) a non-empty set of relations $R$ on $O$, and

4) a potentially empty set $T$ of transitions rules (operations) on $O$.

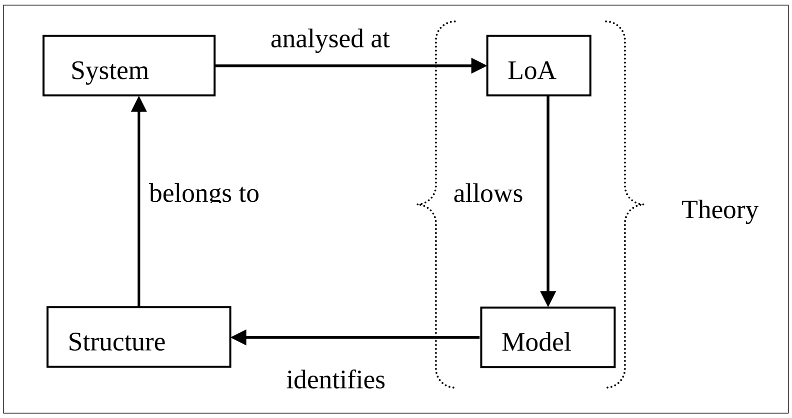


**Table 1: the SLMS scheme**

What are the objects in *O*? According to *Epistemic Structural Realism*, objects can only be posited as *ontic residua*, what remains in principle unknowable once the knowable structures of reality have been factored out:

ESR) models can be increasingly informative *only* about the relations that obtain between the objects, but not about the first-order, one-place predicates (the intrinsic nature) qualifying the objects in themselves.

According to *Ontic Structural Realism*, however, the closing clause needs not to be granted:

OSR) ultimately, objects are themselves structures and in the best cases they can be indirectly captured in our models, at least in principle.

## 3 First move: ESR and OSR are not incompatible

The first move is to try to combine the virtues of both brands of SR without incurring their costs.

### 3.1 Indirect knowledge

Structural realism seeks to combine some ontological commitment with a degree of epistemic resilience to theory change. It achieves this flexibility by decoupling not knowledge from reality, like instrumentalism does, but, within knowledge itself, the descriptions of the knowable structural characteristics of the system from the explanations of its intrinsic properties. Both can still crash, but the structural descriptions are far more resilient to change than the ontological explanations.

All this comes at a price. SR gains epistemic resilience only by endorsing a view of knowledge as an *indirect* relation between the epistemic (possibly multi-) agent and the system under analysis, and at the expense of a somewhat weakened ontological commitment. This can be achieved by relying on the *method of levels of abstraction* (LoAs) (Floridi and Sanders [2004]), in the following way.

### 3.2 Levels of abstraction

The method of LoA comes from modelling in science and is influenced by an area of Computer Science, called *Formal Methods*, in which discrete mathematics is used to specify and analyse the behaviour of information systems.

A LoA consists of a collection of *observables*. An *observable* is an *interpreted typed variable*, that is, a variable with a well-defined possible set of values together with a statement of the properties of the system under consideration for which it stands. The target of a LoA is called the *system*. A system may be accessed and described at a range of LoAs and so can have a range of models.

An explicit reference to the LoA clarifies that

i) the model of a system is a function of the available observables;

ii) different LoA may be fairly ranked depending on how well they satisfy modelling specifications (e.g. informativeness, coherence, elegance, explanatory power, consistency with the data, predictive power, etc.); and

iii) different analyses can be fairly compared provided that they share the same LoA.

The method of LoA is also an efficient way of making explicit and managing the ontological commitment of a theory. Let us see how.

### 3.3 Ontological commitments and levels of abstractions

A theory comprises at least a LoA and a model. The LoA allows the theory to analyse the system and to elaborate a model that identifies a structure of the system at the given LoA. So the structuralist position becomes:

$SR^{LOA}$) instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) *at a given LoA* can be, in the best circumstances,

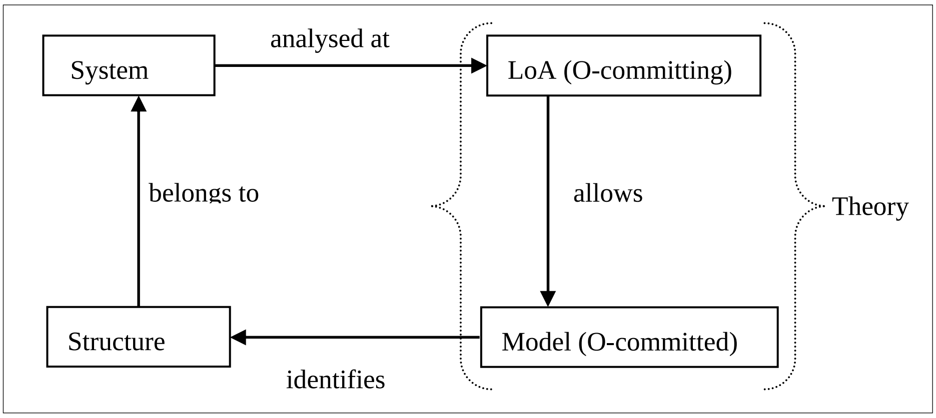


**Table 2: the SLMS scheme with ontological commitment**

increasingly informative about the relations that obtain between the (possibly unobservable) objects that constitute the system under investigation (through the observable phenomena).

This is the system-level-model-structure (SLMS) scheme (Table 1).

The ontological commitment of a theory can now be understood by distinguishing between a *committing* and a *committed* component.

A theory *commits* itself ontologically by adopting a (set of) LoA. Compare this to the case in which one has chosen a car model but has not bought one yet. A theory *shows* its ontological commitment in full through its model, which is therefore the bearer of the actual commitment. The analogy here is with the specific car one has actually bought. So LoAs commit a theory to types while models commit it to the corresponding tokens (Table 2).

The distinction just introduced can now be used to reconcile ESR and OSR.

### 3.4 How to reconcile ESR and OSR

ESR endorses a *minimalist* approach. It argues that a theory is justified in adopting a LoA that commits it (the theory itself) ontologically to a realist interpretation of only the structural properties of the system identified by the model produced by the theory at the chosen LoA. This minimal ontological commitment is *primary* or *first-order*, that is, it concerns first-order knowledge of the structural properties of the system under investigation.

ESR is arguably correct. On the one hand, the adoption of any LoA supporting a degree of first-order ontological commitment higher than the epistemic-structuralist one is metaphysically risky and suspicious. It is certainly unnecessary and not backed up by a general conception of knowledge as an indirect relation with the world. In short, Ockham's razor applies. Now, the current debate on SR has been developing on the assumption that OSR and ESR are incompatible, one of the two must go, and it is not ESR.

This argument against OSR is misguided because ESR and OSR work at separate LoA. As far as a first-order analysis is concerned, SR justifies only one kind of ontological commitment, the one endorsed by ESR. However, at a derivative or *second-order* level of analysis, OSR correctly argues in favour of an economical view of objects.

Consider again the SLMS scheme. The assumption is that there is no *direct* (i.e. LoA-free) knowledge of the intrinsic nature of the objects constituting the system under investigation. However, once a theory has ontologically committed itself to the structural properties of the system, one is entitled to infer *indirectly* that, whatever the system and its components (i.e., the objects or relata) are in themselves, they must be such as to allow the theory to capture their structural/relational properties. But what are the conditions of possibility of knowledge (the knowledge offered by the theory) of the structural properties of the system? The question is whether there might be a justifiable LoA that commits the theory to some kind of description of the type of *relata* that support the structural properties of the system identified by the model produced by the theory. ESR simply has nothing to offer here, for its concern is with a first-order commitment. OSR, on the contrary, can argue in favour of a minimalist approach. And this is a very reasonable step to take. The LoA justifiably adoptable at this second level is one according to which the relata are themselves understood in structural terms.

To summarise, at the first-order level (see Table 3), we adopt a transcendental approach to knowledge – as a LoA-mediated relation with reality – and use Ockham's razor as a methodologically safety measure to limit the *number* of (types of) components to which a theory should be ontologically committed. In the best circumstances, LoAs should ontologically commit a theory at most to the existence in the world of the (type of) structures identified by its models. At a second-order level, we re-adopt a transcendental approach to what makes possible the previous first-order knowledge of structures, and apply a qualitatively modified version of Ockham's razor, which now suggests to keep the *nature* of objects as simple as possible: *entia non sunt complicanda praeter necessitatem*, as it were. According to this new safety principle, it is reasonable to assume that relata are structural objects, for this is all the theory

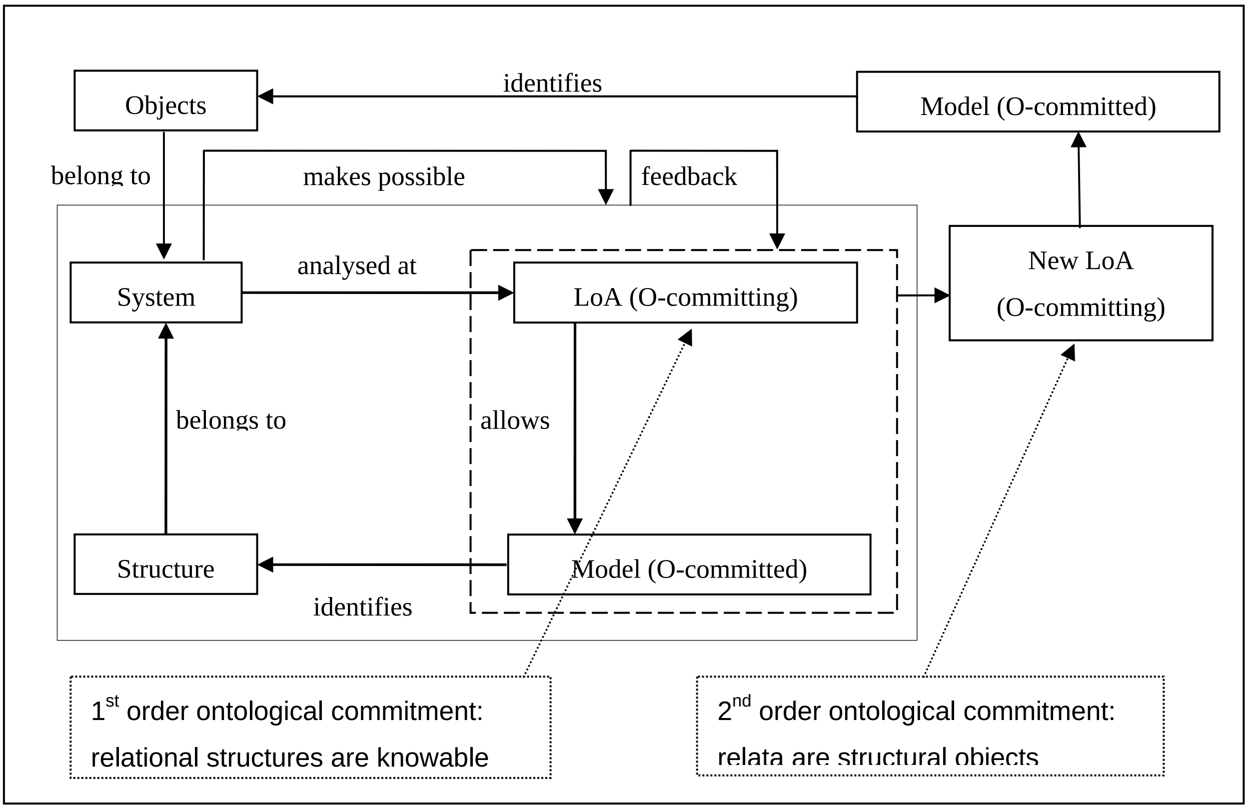


**Table 3: the SLMS scheme with ordered ontological commitments**

needs to justify its first-order commitment. Since the two ontological commitments occur at different levels there is no incompatibility and hence no objection. End of the first move.

## 4 Second move: relata are not logical prior to all relations

Unfortunately, eliminating the apparent inconsistency between ESR and OSR does not yet show that OSR is *plausible*. A direct and quite substantial objection against OSR still needs to be neutralized. This is our second move. Relations (structures) require relata (structured objects), which therefore cannot be further identified as relations without running into some vicious circularity or infinite regress. Yet this is precisely what OSR needs to argue, if the very idea of structural objects is supposed to make sense.

Certainly *external* relations usually require relata. Distance is a good example. Yet *internal* relations constitute their relata for what they are. "Married" comes easily to one's mind: Mary and John are married only because of their mutual relation.

Unfortunately, internal relations seem to supervene on their relata and further qualify them. Mary and John did not come into existence by getting married; they only acquired a new contingent property, their marital status, after their wedding. If one wants to defend the logical priority of internal relations over their relata, then one has to show much more, namely that the essential properties of the objects in question depend upon some fundamental internal properties. This is difficult but not impossible. For there is a significant case, metaphysically more fundamental than the case in which the essence of the relata is in question. This is the (internal) relation of *difference*, which constitutes its relata as a precondition for their existence. *Relata* as merely *differentiated* entities and nothing else (at least not yet) are possible only because there is a relation of difference in the first place. For consider what a completely undifferentiated object *x* might be. It would be one unobservable at any possible LoA. This means that there would be no possible world in which *x* would exist. And this is equivalent to saying that *x* would not exist. Call this the modified Leibnizian law of *the necessary inexistence of the undifferentiable in principle*. The law still says very little about the nature of the entities/relata in question, but there is finally a clear sense in which relata are not necessarily prior to a more fundamental internal relation of difference. Difference is our *Ur*-relation, as it were. In Eddington's words: "The relations unite the *relata*; the *relata* are the meeting points of the relations. The one is unthinkable apart from the other. I do not think that a more general starting-point of structure could be conceived" (Eddington [1928], 230–231, quoted in French [2003], 233). This leads us to the last move.

## 5 Third move: the concept of a structural object is empty

Even if ESR and OSR are compatible, and even if it is not true that, ultimately, structured relata are necessarily logically prior to all structural relations, we still lack a clear grasp of what these *structural objects* might be like, even if indirectly or metatheoretically.

A straightforward way of making sense of these structural objects is as *informational objects*, that is as clusters of data, not in the alphanumeric sense of the word, but in an equally common sense of differences *de re*, i.e. mind-independent points of lack of uniformity.

In its simplest form, a datum can be reduced to just a lack of uniformity, that is, a binary difference, like the presence and the absence of a black dot, or a change of state, from there being no black dot at all to there being one. Admittedly, a datum is usually classified as the relatum exhibiting the anomaly, often because this is perceptually more conspicuous or subjectively less redundant than the other relatum, seen as a background condition. However, the relation of difference is binary and symmetric. A white sheet of paper is a constitutive part of the datum itself, together with the fundamental relation of difference that couples it with the dot. In this specific sense, nothing is a datum *per se*, without its counterpart, exactly like nobody can be a wife without there being a husband. So, ontologically, data, as still unqualified differences, are purely relational entities. Of course, from a structural perspective, they remain unknowable in themselves. One never deals with pure data but only with somewhat interpreted data. Now clusters of data as relational entities are the relata we are looking for in our modified version of OSR.

As French [2001] points out, in the context of the philosophy of physics, "if we want to continue to talk, in everyday language, about electrons as objects - because we lack the logico-linguistic resources to do otherwise - then we can do so 'only indirectly', '... not insofar as they themselves, as individuals, are given, but so far as they are describable as "points of intersection" of certain relations' (Cassirer, *ibid*.)".

So what an informational ontology could look like? *Object Oriented Programming* (OOP) provides an excellent example of a flexible and powerful methodology with which to clarify and make precise the concept of "informational object".

OOP is a method of programming that radically changed the approach to software development. Historically, a program was viewed as an algorithmic procedure that takes input data, processes it, and produces output data. The difficulty was then represented by the elaboration of the algorithmic process. OOP has shifted the focus from the logic procedures, required to manipulate the objects, to the objects that need to be manipulated. Consider a pawn in a chess game. Its identity is not determined by its contingent properties as a physical body, including its shape and colour. Rather, a pawn is a cluster of states (properties like white or black, and its strategic position on the board) and behavioural rules (e.g., it can move forward only), which in turn are possible only in relation to other pieces and the logical space of the board. For a good player, the pawn is only a placeholder standing for an "informational object". It is not a material thing but a set of typed variables, using the LoA terminology, a mental entity, to put it in Berkeley's terms, or an entity constituted by a bundle of properties, to use a Humean expression, whose existence and nature is determined by the differences and nomological relations that characterise the game of chess. Now in OOP, data structures (e.g. the pawn's property of being white) and their behaviour (programming code, e.g. the pawn's power to capture pieces only by moving diagonally forward) are packaged together as (informational) objects. Objects are then grouped in a hierarchy of classes (e.g. pawns), with each class inheriting characteristics from the class above it (e.g. all pieces but the king can be captured, so every pawn can be captured). A class is a named representation for an abstraction, where an abstraction is a named collection of attributes and behaviour relevant to modelling a given entity for some particular purpose at a certain LoA. The routines or logic sequences that can manipulate the objects are called *methods*. A method is a particular implementation of an operation, i.e. an action or transformation that an object performs or is subject to by a certain class. Objects communicate with each other through well-defined interfaces called *messages*.

Clearly OOP provides us with a concept of informational objects that is richer than our minimalist approach to informational realism may allow us. Recall that all we are placing in the world are informational objects as the conditions of possibility of those structures that our first-order LoAs allow us to attribute to the world in the first place. OOP is not a viable way of doing philosophical ontology, but a valuable methodology to clarify the nature of our ontological components.

Time has come to summarise the proposed solution to the ontological problem.

## 6 Informational realism

Informational realism (IR) is a version of SR. As a form of realism, it is committed to the existence of a mind-independent reality. Like ESR, it supports a first-order, minimal ontological commitment in favour of the structural properties of reality. Like OSR, it also supports a second-order, minimal ontological commitment in favour of objects understood informationally. This second-order commitment is justified by epistemic reasons. We are allowed to commit ourselves ontologically to whatever minimal conception of objects is useful to make sense of our first-order commitment. IR argues that, as far as we can tell, the ultimate nature of reality is informational, that is, it makes sense to adopt a LoA at which our mind-independent reality is constituted by relata that are neither substantial nor material (they might well be, but we have no reasons to suppose them to be so) but informational. We are ready for a definition:

IR) instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) at a given LoA can be, in the best circumstances, increasingly informative about the relations that obtain between the (possibly unobservable) informational objects that constitute the system under investigation (through the observable phenomena).

## 7 Conclusion

Philosophers often rely on spatial analogies to explain their theories. References to rooms are particularly popular. Sextus Empiricus thought that we are like men searching for gold, or shooting at a target in a dark room: no matter how long the search or the shooting proceeds, it is pointless because in principle there is no way to establish whether any of us has found a nugget or hit the mark. Turing used different rooms for his test. And Searle devised a Chinese room for his couterexample. I shall rely on their examples and suggest a double box analogy to illustrate IR. But first, a final bit of terminology.

In software engineering, *black-box* refers to a test-design method that focuses on testing functional or behavioral requirements of a program. The methodology treats the *analyisandum* as a nontransparent and closed system, avoiding using explicit knowledge of its internal nature or structure to understand the way it works. The opposite methodology is known as *white-* or *glass-box* test design. This allows one to "look inside" the system, and it focuses on using specific and detailed knowledge of the program code to guide the selection of test data. A *grey-box* approach is one that allows only a partial view of the internal properties of the system.

According to IR, any white-box approach to reality is excluded in principle, given the indirect nature of knowledge. We look at the world as if we were in Sextus' dark room. This is the first box. We are in it, but our target is not representationalist, nor are our interactions with the furniture in the room unidirectional, as Sextus assumed. And contrary to Sextus' it is a grey-box. In the best cases, it allows the reconstruction of the structural properties relating the furniture of the room, i.e. our informational objects. The latter are our second type of boxes. Like in Turing's test, they are black-boxes, not directly knowable, but "epistemically interactable". Sometimes we can indirectly capture their nature by observing their behaviour and mutual interactions, but we do not know their intrinsic properties.

IR takes our epistemic goal to be constructionist, not representationalist. Knowledge is not a matter of hitting or finding, but of designing and constructing, and one may construct successfully even in the dark. Since we wish to devise an intelligible conceptual environment, we do so not by trying to photocopying whatever is in the room (representationalist epistemology), but by interacting with it as a *resource* for our semantic tasks, interrogating it through experience, tests and experiments. Out of metaphor, reality is not a source but a resource for knowledge. Structural objects (clusters of data as relational entities) work like *constraining affordances*, exploitable by a theory, at a given LoA, as input of adequate queries to produce information (the model) as output. In this respect, we are more like Turing's interrogator, since the model of investigation is erotetic: we query the world as a database. But since our task is not to find out who is who, we resemble Searle in his Chinese room: we get the data on one side and produce information on the other. The difference, in this case, is that we understand the semantic rules of the game. Black-boxes in a grey-box, this is the basic idea behind IR. The last specification is that the qualifications are level-dependent, as the distinction between system and component or unit of the system is. A black-box may be opened, but opening it transforms it into a grey-box, in which more black-boxes may be found. Whether *ad infinitum* we simply cannot know.

## 8 References


Chakravartty, A. forthcoming, "The Structuralist Conception of Objects", in *18th Biennial Meeting of the Philosophy of Science Association, Nov. 7-10, 2002*, Milwaukee, WI.

Eddington, A. 1928, *The Nature of the Physical World* (Cambridge: Cambridge University Press).

Floridi, L., and Sanders, J. W. 2004, "The Method of Abstraction" in *Yearbook of the Artificial. Nature, Culture and Technology. Models in Contemporary Sciences*, edited by M. Negrotti (Bern: Peter Lang), Preprint from http://www.wolfson.ox.ac.uk/~floridi/papers.htm

French, S. 2001, "Symmetry, Structure and the Constitution of Objects", in *Symmetries in Physics, New Reflections: Oxford Workshop, January 2001*, Oxford.

French, S. 2003, "Scribbling on the Blank Sheet: Eddington's Structuralist Conception of Objects", *Studies in History and Philosophy of Modern Physics*, 34, 227–59.

Ladyman, J. 1998, "What Is Structural Realism?" *Studies in History and Philosophy of Science*, 29A(3), 409-24.

Maxwell, G. 1970, "Structural Realism and the Meaning of Theoretical Terms" in *Analyses of Theories, and Methods of Physics and Psychology*, edited by Stephen Winokur and Michael Radner (Minneapolis: University of Minnesota Press), 181-92.

Morganti, M. forthcoming, "On the Preferability of Epistemic Structural Realism", *Synthese*.

Votsis, I. forthcoming, "Is Structure Not Enough?" *Philosophy of Science, Supplement*.

Worrall, J. 1989, "Structural Realism: The Best of Both Worlds?" *Dialectica*, 43, 99-124.


## Acknowledgements


I am grateful to John Weckert and the Centre for Applied Philosophy and Public Ethics (CAPPE) for the kind invitation to deliver the keynote lecture at the Computing and Philosophy conference (the Australian National University in Canberra on 31 October – 2 November, 2003) on which this article is based. I wish to thank the colleagues attending the conference for their useful feedback. Gianluca Paronitti and Matteo Turilli read a final version of the text and made several useful comments. Of course I am the only person responsible for any remaining mistake.