

Coordination, Semantics and Ontologies

Francesco A. Zaccarini¹, Claudio Masolo²

¹University of Bologna (DAR), Bologna, Italy

²CNR (ISTC-LOA), Trento, Italy

Abstract

In the literature, interoperability is often informally defined in terms of the achievement of coordination among different parties in the accomplishment of some goals, ranging from the effective transmission of knowledge, usually about the world, to the execution of practical tasks. However, it is not clear whether adopting a shared ontology, or, in general, reaching an agreement on semantics, is either a necessary, or sufficient, condition to that end. In this paper, we set out to explore this topic by outlining an approach resting on minimal epistemological assumptions, questioning to which degree coordination capabilities usually associated with high-level cognitive functions (e.g., language) can be recovered in the proposed setting. It is suggested that our approach might also provide fruitful grounds for machine-learning-enhanced methodologies.

Keywords

Coordination, Semantics, Interoperability, Ontological Commitments, Machine Learning

1. Introduction

Semantic technologies are widely recognized as foundational to interoperability, as they provide standardized, machine-readable representations of data and knowledge, facilitating shared understanding, integration, and reuse. However, the notions of *interoperability* and *semantics* are employed and understood in markedly different ways in the literature. As such, despite the assumed centrality of their connection in knowledge representation and applied ontology, the relationship between the two remains opaque at best. This conceptual variability motivates a critical examination of the presumed connection, both to clarify the foundations of orthodox approaches and to explore the viability of alternatives.

In this paper, we undertake this endeavor by raising the question of whether agreement on a theory's semantics, or the endorsement of a shared worldview, is sufficient, and/or necessary, to achieve coordination among a plurality of agents. Following a brief survey of how the core notions are employed in the literature and an outline of key concerns regarding their presumed interrelation (Sect. 2), we address these questions by outlining an alternative approach grounded in minimal epistemological assumptions, looking back to philosophical tenets at the heart of the so-called “upward path to structural realism” [1]. Specifically, we explore the extent to which coordination can be achieved within a framework resting on those ideas, offering a preliminary investigation and pointing toward potential directions for further development.

The underlying theoretical foundations and the formalization of the framework are presented in Sect. 3 and Sect. 3.1, respectively. Sect. 4 focuses on illustrative examples, which are followed by an exploratory discussion (Sect. 4.1). Here, we touch on broad epistemological issues related to the capability of situated epistemic agents to navigate and make sense of the world. While considerations pertaining to cognition and communication are certainly relevant, our main focus lies in the acquisition, transmission, and aggregation of data for the purposes of coordination, prediction, and intervention. We suggest that data acquired adhering to the epistemological stance at the core of our approach may offer a fruitful basis for machine learning. Conversely, the integration of machine-learning-enhanced methodologies within this framework represents a promising avenue for future research.

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8-9, 2025, Catania, Italy

✉ francesco.zaccarini@unibo.it (F. A. Zaccarini); claudio.masolo@cnr.it (C. Masolo)

ORCID 0009-0008-8009-5009 (F. A. Zaccarini); 0000-0002-0764-8630 (C. Masolo)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. Interoperability, Semantics, Ontology

In the literature, *interoperability* is frequently understood, implicitly, if not explicitly, in terms of the ability to enable *information or meaning exchange* and to support *coordination* among systems, that is, the achievement of shared goals or the cooperative execution of tasks [2, 3, 4, 5]. Some authors use the term “*semantic interoperability*” very broadly, along these lines [6, 7, 8]. However, Euzenat [9, 10], among others, advocates for a more narrow understanding of semantic interoperability, distinguished from *semiotic interoperability* (as well as from encoding, lexical, and syntactic interoperability—categories which are not relevant to our present aims). The distinction is based on the difference between (i) *logical semantics*, i.e., “a variety of representation theorem” [11] connecting the non-logical signature of a theory to another mathematical structure (a set theoretical structure, in the case of standard Tarskian model-theoretic semantics for first-order logic), and (ii) *real-world semantics* [12] which fully fix the conditions of use, i.e., how people interpret a theory (in given contexts).

The gap between these two notions becomes evident when one adopts a theory of meaning according to which real-world semantics cannot be fully captured intra-linguistically. From a referentialist perspective, scholars who agree on the truth values of sentences (of a logical language), may associate these sentences with radically different states of affairs in the world, possibly adopting different domains of interpretation. Logical semantics constrains the interpretation of theories, but it always leaves room for different real-world interpretations. Therefore, outside of a strict, intra-linguistic theory of meaning (such as inferentialism), sharing a theory—e.g., a computational ontology including its factual and terminological axioms—does not guarantee semiotic interoperability and may lead to coordination failures. The situation worsens when the interpreters are not granted direct access to the structure of the world or to the very same conceptualization of it. This is not to say that ontologies do not facilitate semiotic interoperability. Rather, it is a way to emphasize the importance of tacit agreements and alignments among users. Meaning negotiation and calibration are just as important in semantic technologies as they are in ordinary communication. Presupposing alignment due to the sharing of ontologies might lead to costly (and hard to detect) mistakes.

Indeed, coordinating or exchanging information among agents does not seem to require the sharing of worldviews. Consider, for instance, two people, *A* and *B*, who speak entirely different languages and have radically divergent worldviews: *A* believes in Greek gods, while *B* considers themselves to be in a skeptical scenario and posits nothing but a series of states of a matrix. Now suppose *A* systematically utters the sentence “Zeus is in bad mood” before it starts raining. In time, *B* might learn to associate *A*’s utterance with the onset of rain—which *B* understands as a transformation in the matrix’s state—despite the fact that *B* does not share *A*’s worldview and would not accept the posited entities and states of affairs. In fact, it is irrelevant to *B* whether *A* is a human being uttering sentences according to a complex theory of the world or a barometer: insofar as *A* reliably tracks states of the world, and *B* is able to form a consistent association between certain perceived outputs of *A* and the latter, information about the world is effectively transmitted, and made available for planning and intervention.

The extent to which coordination and communication can be achieved without a shared language, semantics, or ontology remains an open question. In what follows, we introduce a formal framework to study this question without presupposing any kind of semantic sharedness.

3. Framing the Proposal

At the dawn of the 20th Century, several prominent philosophers sought a way to provide a foundation for both scientific inquiry and metaphysical investigation within a realist (as opposed to idealist, then prevalent) framework: a foothold in the external world withstanding Kant’s challenge. In this context, Poincaré [13] and Russell [14]—the latter drawing on the work of Meinong [15, 16]—courted the idea of there being a structural similarity between the noumenal, how things are in themselves, and the phenomenal, postulated to explain the pragmatic effectiveness of our commonsensical worldview and

our scientific theories.¹ If granted, this structural similarity provides us with an access to the world in itself, albeit in a limited and indirect fashion: relations holding among the objects in one's experience are the mirror of relations among external world entities—which *ground, and constrain, all the possible representations* having them as focus. Differences among entities are thus not created via comparison, but are merely apprehended by means of it. This guiding intuition lies at the heart of the so-called Helmholtz-Weyl Principle (hereafter *HWP*), which Psillos and Votsis regard as a cornerstone in the “upward path” to epistemic structural realism [17, 1]. The *HWP* simply states the following:

P1 Different effects (i.e., percepts) imply different causes (i.e., stimuli/physical objects).

(Helmholtz-Weyl Principle)

Much ink has been spilled on Russell's understanding of the *HWP*. For our purposes, we will assume that (1) from an ontological point of view, the principle comes down to an injective mapping $f : P \rightarrow \mathcal{P}(W) \setminus \{\emptyset\}$ between the set P of perceptions and *sets of possible states* of the world (W is the set of all the possible states), such that the sets in the co-domain are *disjoint*, i.e., states of the world univocally determine perceptions.² (2) From an epistemological point of view, a situated epistemological agent is warranted to infer, from the fact that they had a certain perception, rather than others, that that which is perceived (the world) is in a certain set of states, rather than others. It should be stressed that, in this case, the agent can only discriminate *contrastively* between said sets, being completely blind to the states of the world themselves, as well as to which states fall within which set. Perceptions can be thus seen as partitioning the logical space in a way reminiscent of Rayo's characterization of propositions [18].

So understood, the principle arguably has considerable intuitive plausibility. Consider a simple case involving a human being having a perception of redness: for such a perception to occur, the world has to be in a state that falls within a certain range; conversely, for the subject to perceive a different color, the world would have to be in a different range of states, one that does not include the actual state. Intuitively, perceptions allow us to navigate the world by tracking differences which are relevant to us; if something like the *HWP* did not hold, our capability of doing so, as well as the differences among perceptions, would be in need of alternative explanations. It should be noticed that in this framework a perception is naturally compatible with the world being in states differing from the actual one, insofar as the perceiver is not capable of discriminating between them:³ this can be either due to a standard lack of sensibility/precision (e.g., the human's perceptual apparatus reacting in the same way to slight differences in electromagnetic spectrum), or because the differences are not “relevant” (e.g., causally disconnected worldly differences), yet skeptical scenarios are also relevant (e.g., brain-in-a-vat scenarios in which the human being is stimulated so that they perceive redness). Indeed, these distinctions make sense only *within a certain theoretical framework* or *within a certain worldview* positing things which go way beyond what is warranted by the *HWP*, as interpreted by us.

It is no coincidence that the *HWP* emerged during a period strongly influenced by Mach/Duhem-style instrumentalism, holding scientific theories as symbolic tools for efficiently organizing sensory experience—instruments for making predictions and guiding action rather than attempts to uncover the supposed metaphysical structure underlying phenomena or to attain truth in that sense. In fact, the *HWP*, taken at face value, hints at an even more radical epistemological approach: one that does not commit to any particular theory of the world (not even as a way to “keep the score”) and resting on minimal, revisable assumptions concerning how perceptions are connected—operating entirely at that level. This marks a departure from the traditional Galilean method, hinged on idealization and

¹It is worth noting that, in his time, Kant faced criticism for positing a dependence of the phenomenal on how things are in themselves: critics argued that this move involved a misapplication of the category of causation beyond its legitimate domain, namely, the realm of possible experience. Arguably, this “unwarranted assumption” is key in the cited authors' proposals.

²Focusing on world states allows to sidestep issues concerning the identification of entities, and structure in general, in the world, among other things.

³Only given a perceiver capable of discriminating between each and every different state of the world there can thus be an *isomorphism* between phenomena and noumena (and, thus, representations carrying maximal information—i.e., singling out a single possible state of the world in contrast with all others). That said, the converse of the *HWP* gets problematically close to idealism, positing a systematic dependence of the world on perceptions.

controlled observation, insofar as matters of saliency (i.e., which entities/characteristics of a system are counter-factually relevant to answer a certain experimental question) and granularity in description presuppose a (somewhat) rigid theoretical framework. The upside of accepting these challenging restrictions is a solid metaphysical foothold hinged on a minimal principle playing an explanatory role: information about *the world* (being in a certain range of ways rather than others), contra certain strongly anti-realist versions of instrumentalism. Operating at such a level putatively sidesteps issues concerning the underdetermination of theories, and the need to establish connections across representational frameworks—this being extremely relevant for knowledge representation and interoperability.

While another movement inspired by instrumentalism, Logical Positivism, failed as a descriptive project targeting natural languages and scientific theories, one can envision a system taking some of its core principles as *prescriptive* guidelines in order to improve the transmission of information about the world across a network. The reasoning underlying the proposal is simple. Going back to points discussed in the previous sections: (1) if, necessarily, symbolic systems leave room for different interpretations, it might be worth shifting the focus towards mechanisms which connect pairs of perceptions in a way which robustly allows the second perceiver to produce actions whose success relies on the world being in a certain range of ways rather than others, as tracked by the first perception; (2) if our epistemological position prevents us from adjudicating between competing theories—assuming that one (or more) is correct—yet the theoretical assumptions in question are not critical (and sometimes an obstacle, due to unwarranted assumptions of standardization and rigidity in adjustments) for interoperability, then we should operate at a level of abstraction where data is as theory-neutral as possible.

The idea can be made more precise by comparison with instrumentalism. Point (1) invites us to assume a physicalist standpoint, shifting the focus from perceptions to the state of a generic portion of the world functioning as an observational apparatus and capable of being influenced, and influencing, states of the world. This also allows us to sidestep issues related to perception which clearly go beyond the scope of this brief paper. Indeed, assuming a physicalist standpoint, *HWP* should be even less contentious, as it can be read as simply stating that one is warranted to posit a correlation between two states of the world when one causes the other, with cascading consequences concerning the state of *portions* of the world—which can be taken to be “observational apparatuses”, since they are influenced by the world, and “agents”, since they influence it. It is now worth pointing out that, in order for the connections to be “informative”, the state of the observational apparatus has to be non-trivially *determined* by factors which do not include the observational apparatus’ state itself. This is implicit in the causal formulation of the *HWP*, given the assumption of a temporal asymmetry of causes and effects. However, to avoid the metaphysical commitments and conceptual baggage associated with causality, the principle can be more neutrally reformulated in terms of dependence. For our purposes—and endorsing the simplification with respect to time for the sake of simplicity—the core intuitions underlying the *HWP* are thus better expressed as follows:

P2 Let time t_1 precede time t_2 . The state of the world at t_2 depends on, and is determined by, its state at t_1 , i.e., different states of the world at t_2 imply different states of the world at t_1 .

(Revised Helmholtz-Weyl Principle)

According to (P2), the state of an observational apparatus P is correlated with a range of possible states of the world, which determine it (if they hold). Different states of P stand for, and “reveal”, disjoint ranges.

In focusing on mechanisms which might, or might not, involve humans, and denying them a special place in the system, our proposal radicalizes tendencies for externalization which are shared by approaches focusing on sensors (for data acquisition) and machine learning (for data aggregation), which are usually characterized by the attempt to fit in with human conceptualization. In general, the externalization and standardization of data acquisition is a defining feature of the scientific enterprise and arguably one of the key reasons for its pragmatic and predictive success. As widely acknowledged in the literature, new theories and new instruments often co-evolve, and the mediation of instruments is not only essential for detecting world-states beyond the reach of human senses, but also to decide between different competing theories [19, 20].

Another comparison, this time with the operationalist derivations of instrumentalism, can shed light on (2). To improve information transmission and data quality, operationalist approaches ground meaning on *procedures of data acquisition*. Since no two measurements are alike, one might ideally want to include as much contextual information as possible to define these procedures precisely. However, including all the contextual information is not feasible and doing so would also prevent generalization. Notably, similar tendencies are echoed in the increasing focus on context regarding data (metadata and provenance *in primis*). Unlike these approaches, our framework does not distinguish between the target and the context when determining the state of an observational apparatus; both contribute. Thus, the issue shifts to “transmitting” and “aggregating” the states of the individuals in the world in a way that allows for effective predictions. Accordingly, our focus turns (*i*) to observational apparatuses that discriminate between different (sets of) states of the world by entering distinct internal states in response to them (possibly keeping track of past observations); and (*ii*) to mechanisms that reliably correlate the internal states of different observational apparatuses producing interventions in the world that yield expected (observable) results. Notice that our approach, which we will now illustrate with concrete examples, has precedents in this regard and shares similarities with the proposal advanced by Berto, Rossi & Tagliabue [21], as well as with previous work in cybernetics and robotics.

3.1. Introducing The Framework

To concretize the ideas outlined in the previous section, we introduce a framework that lays the foundation for an exploratory discussion on the potential developments of this approach. For our purposes, we start by outlining a framework that generalizes basic cellular automata (CA) [21], and subsequently examine simple toy models constructed within this generalized setting. Compared to CA, our framework aims to provide a more detailed and dynamic model of a world whose inhabitants (CAs’ cells) are not necessarily eternal, can change their configuration through time, and can be governed by specific laws. The following defines our general framework:

- We adopt a discrete set T of *times* and a set S of (internal) *states* that describe the states of a set D (called the domain) of *individuals*. To simplify the notation, we assume that T is a subset of the set of integers, writing $t+1$ for the successor of t and $t-1$ for its predecessor.
- The (total) function $\tau : D \rightarrow \mathcal{P}(T)$ identifies the temporal extension of an individual, i.e., the set of times at which an individual exists. We assume that for all $x \in D$, $\tau(x) \neq \emptyset$, i.e., all individuals exist at least at one time.
- The (partial) function $\sigma : D \times T \rightarrow S$ identifies the state of an individual at a time at which it exists. This function enables tracking of entities’ internal state through time.⁴ We assume that for all $x \in D$, then $t \in \tau(x)$ if and only if there exists $s \in S$ such that $\sigma(x, t) = s$.
- The (total) function $\beta : D \times T \rightarrow \mathcal{P}(D)$ identifies the *base* of an individual x at a time t at which it exists, i.e., the set of individuals (called *bases*) on which x depends at t . This function establishes (directed) connections between individuals, i.e., it determines the “shape” of the world. Since individuals can change their inputs over time, the shape of the world is not necessarily static.
- The world is regulated by discrete dynamics: the state and the bases of an individual x at time t are fully determined by (*i*) the states of the bases of x at time $t-1$ and (*ii*) the deterministic updating rules that regulate the world. These updating rules can impact both the states and the bases. This means that the base of x represent the part of the domain that influences x ’s state. The manner in which this influence is implemented depends on the specific rules considered. The framework allows for independent individuals who do not have bases and whose behavior is regulated by internal rules that do not depend on other individuals.

Notably, our framework also allows for scenarios in which individuals have “private” states—i.e., if $x \neq y$, then for all $t \in \tau(x)$ and $t' \in \tau(y)$, we have that $\sigma(x, t) \neq \sigma(y, t')$ —and are regulated by “private” updating rules, i.e., each individual has its own specific way of functioning. It is thus possible to tune the framework to accommodate both “reductionist” assumptions implicit in classic CA, where individuals

⁴The function can be made total by introducing a state that characterizes all the individuals that do not exist at a given time.

resemble Aristotelian matter and laws are universal, as well as scenarios populated by haecceities, with no room for ontologically grounded abstractions. Intermediate configurations are also possible: for instance, the domain D could be partitioned into n subdomains (types of individuals) D_1, \dots, D_n , with updating rules parameterized relative to these types. This flexibility is significant not only because inhomogeneous worlds cannot be dismissed a priori, but also to investigate the conditions under which a situated individual can effectively navigate and acquire knowledge about the world it belongs to. Showing that, under certain (epistemological) assumptions, an individual’s capacity to navigate its environment depends on some form of Millian “uniformity in nature” would arguably be a noteworthy result in itself. Naturally, the framework can be adjusted further by introducing additional constraints. For example, one could stipulate that all individuals are interconnected or have nonempty bases.

In order to better describe the examples, we distinguish the following types of individuals, each represented with a distinct graphical notation as detailed below.

- **Memory** (diamond nodes m_i). Individual m is a memory of individual x at time t if the state of m at time t depends only on the states of x and of m at $t-1$ ($\beta(m, t-1) = \{x, m\}$) and the state of x at $t+1$ depends on the states of m at t ($m \in \beta(x, t)$). The state of a memory depends on its previous state, creating a cycle (see m in Fig. 1).
- **Agent** (circle nodes a_i). Individual a is an *agent* at time t if and only if a is not the memory of any individual at t and there is a single m that is a memory for a at t . We assume that memories cannot be shared by several agents. The memory m of agent a aggregates the current states of a and m , providing the result of the aggregation as input for the agent a at the next time step. This captures the “introspective” nature of agents: their state depends not only on the states of “external” individuals, but also on their own previous internal state. In our models, memories act as externalizations of agents’ internal states and dynamics. This design enables the modeling of more complex processes with varied temporal extensions.
- **Sensor** (rectangle nodes s_i). Individual s is a *sensor* of agent a at time t if and only if (i) s is in the base of a at t ($s \in \beta(a, t)$); (ii) s does not have a memory at t ; (iii) s is not the memory of any agent at t . Sensors do not have memory, so their state depends only on the states of “external” individuals (directly or indirectly) connected to them. However, note that the agents do not have direct access to the states of the sensors (they are only influenced by them), nor do they know which portions of the world the sensors are affected by, or how reliably they track certain (sets of) world configurations. The whole battery of sensors of an agent is its *console*. Agents can share sensors, or have sensors that are connected to the sensors of other agents.

In general, insofar as all individuals receive inputs and produce outputs, they can all be considered as “reasoners” or “performers of computations”, given a broader understanding of these notions.⁵ Without aiming to fully capture the notion of agentivity, the previous “structural” characterizations of agents, memories, and sensors are sufficient for our illustrative purposes and they avoid committing to complex cognitive architectures or specific ontological analyses.

Let’s now consider some simple examples. For simplicity, we will assume that all the individuals exist and have the same bases at all times (for all $x \in D$, $\tau(x) = T$ and $\beta : D \rightarrow \mathcal{P}(D)$).

Example 1 (redness sensation of human a). To familiarize with the framework, we consider the simple graph in Fig. 1 where the nodes stands for *individuals*, i.e., $D = \{i_1, i_2, s_1, s_2, a, m\}$, and the arrows indicate the *bases* of individuals, i.e., $\beta(i_1) = \emptyset$, $\beta(i_2) = \emptyset$, $\beta(s_1) = \{i_1\}$, $\beta(s_2) = \{i_2\}$, $\beta(a) = \{s_1, s_2, m\}$, $\beta(m) = \{a, m\}$. According to the above definitions of types of individuals, at all times, except the first and last ones, a is an agent; m is a memory of a ; s_1 and s_2 are sensors of a . For simplicity, we consider only 5 times: $T = \{0, 1, 2, 3, 4\}$ and we mainly focus on times 1, 2, 3, where the previous definitions are effective. We introduce states in S as needed. Fig. 1 shows the subset of S relevant for the following discussion. In this figure, node x is labeled $t:\alpha$ when $\sigma(x, t) = \alpha$.

We informally describe the example as: the human a has memory m , visual apparatus s_1 which is affected only by i_1 , a part of the world not further specified, and auditive apparatus s_2 which is affected by i_2 , a part of the world different from i_1 . State α_v assures a is in visual rest while state α_r assures

⁵This option is for instance favored by [21].

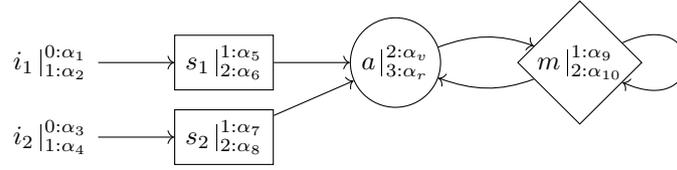


Figure 1: Illustration of the model related to Ex.1.

a has a redness sensation. The interpretation of states α_1 - α_{10} is left unspecified. With the intuitive interpretation in mind, let us follow the evolution of the system in our example. The stimulation of the visual apparatus is modeled via the transition of s_1 from α_5 (at 1) to α_6 (at 2). Since only the part of the world i_1 is in the base of s_1 , the transition of s_1 is fully determined by the transition of i_1 from α_1 (at 0) to α_2 (at 1). Similarly for the stimulation of the auditive apparatus s_2 . The state of the human being a is affected by the states of their perceptive apparatuses s_1 and s_2 as well as memory m ($\beta(a) = \{s_1, s_2, m\}$). The contemporaneous transitions of s_1 from state α_5 (at 1) to state α_6 (at 2), of s_2 from α_7 to α_8 , and of m_1 from α_9 to α_{10} determine the transition of a from visual rest α_v (at 2) to redness sensation α_r (at 3). \square

Ex. 1 allows us to make some general considerations about our framework. First note that our framework is compatible with two perspectives: (i) a descriptive one, in which one describes the state and the shape of the world (i.e., the states and bases of all individuals) at all the times, but the updating rules remain implicit (we adopted this perspective in Ex. 1); and (ii) a constructive or simulative one, in which the updating rules are explicit, and the state and shape of the world is generated step by step via such rules, starting from an initial configuration. Perspective (i) is clearly weaker than perspective (ii). Epistemological approaches typically exclude exact knowledge of the rules that govern the world, which can only be inferred inductively a posteriori.

Second, both perspectives presuppose knowledge of the world's structure, i.e., the dependencies between individuals. This knowledge enables the study of how states propagate within a given world. For instance, in Ex. 1, the state of the visual apparatus s_1 is only affected by the state of i_1 . Hence, s_1 tracks the state of i_1 at a previous time without interference, i.e., the state of s_1 can be systematically calculated from that of i_1 , preserving information about the range of configurations of the world to the degree that s_1 can discriminate between different states of i_1 . For example, s_1 could enter the same state despite i_1 being in several different states. In this case, s_1 clusters different states of the world, but it does that in a systematic way. The same applies to the agent a when considering their visual apparatuses and memory. That is, given the resolution a disposes of, a systematically tracks the configurations of the states of s_1 , s_2 , and m . *Ideal channels* (and *fibers*), as defined below, generalize the idea of reliable tracking by considering chains of tracking steps.

- **Ideal channel and fiber.** At time t , $\langle D_1, \dots, D_n \rangle$ with $D_i \subseteq D$, $D_n = \{e\}$, and $n \geq 2$, is a (transmission) *ideal channel* of length $n-1$ ending in e if and only if, for $0 \leq i \leq n-2$, (i) $D_{n-i-1} \subseteq \bigcup_{x \in D_{n-i}} \beta(x, t-i)$ and (ii) for all $x \in D_{n-i}$, $\beta(x, t-i) \subseteq D_{n-i-1}$ or $\beta(x, t-i) \cap D_{n-i-1} = \emptyset$. A *fiber* is an ideal channel with the form $\langle \{x_1\}, \dots, \{x_n\} \rangle$.

First of all, D_{n-1} contains all the direct “inputs” (the bases) of the ending individual e at t . Thus, ideal channels of length 1, which have the form $\langle \beta(e, t), \{e\} \rangle$, formally capture the examples discussed above: $\langle \{i_1\}, \{s_1\} \rangle$ (a fiber), $\langle \{i_2\}, \{s_2\} \rangle$ (a fiber), $\langle \{s_1, s_2, m\}, \{a\} \rangle$, and $\langle \{a, m\}, \{m\} \rangle$ are the only ideal channels of length 1 (at all times except 0) in Ex. 1. Secondly, D_{n-2} collects *some* of the inputs of e “mediated” by at least one of its bases, i.e., as stated in clause (i), D_{n-2} is a subset of the union of the bases (at $t-1$) of the bases of e (at t). Clause (ii) assures that when an input mediated by $x \in \beta(e, t)$ is included in D_{n-2} , all of x 's other bases (at $t-1$) are also included in D_{n-2} , i.e., all x bases are included in the channel as mediated inputs of e . It is however possible to exclude all the bases of an $x \in \beta(e, t)$ from D_{n-2} . In this case, x is considered a “starting point” of the channel. Similarly, when x has no base. This procedure is iterated until D_1 is reached, which allows the channel to collect inputs

mediated by $n-2$ layers of individuals. Formally, the set of the starting points of channel $\langle D_1, \dots, D_n \rangle$ is $D_1 \cup \bigcup_{1 < i < n} \{j \mid j \in D_i \text{ and } \beta(j) \cap D_{i-1} = \emptyset\}$. Intuitively, the ending individual e of an ideal channel tracks (with a given resolution) the states of the starting points of the channel.

In Ex. 1, both $\langle \{i_1, i_2, a, m\}, \{s_1, s_2, m\}, \{a\} \rangle$ and $\langle \{i_1\}, \{s_1, s_2, m\}, \{a\} \rangle$ are ideal channels of length 2 ending in a (at all times except 0). As we have seen, in general, the state of a at $t+1$ depends on the states of all its bases at t (which are included in D_{n-1}). Thus, a 's redness sensation is, in principle, affected also by the state of a 's auditory apparatus and by past sensations stored in memory—i.e., all the context is included. In our framework, memories can thus provide a way to model environmental adaptation, that is the production of varied outputs in response to otherwise identical external sensory inputs. This is a desirable feature when human and artificial agents, *qua* agents, are involved, as per our example. The framework is likewise capable of modeling *robustness*: measurement devices are typically designed to be more sensible to differences in the state of certain (posited) bases and, conversely, to minimize the impact of others, usually including memories. Memories can thus be used to model the internal states of devices due to repeated excitations or permanent deformations (e.g., wear). A robust device either controls the states of interfering bases (the noise) or compensate for their impact on the states of its inputs. In the first case, the (update function of the) device may be insensitive to the states of interfering inputs within certain environmental conditions (in controlled environments, the device can control them, at least partially). In the second case, the (update function of the) device may take interferences into account, adjusting its output state based on this information.

The concepts of ideal channel and fiber can be exploited to model more complex correlations among the states of two individuals at different times, also taking into account the world's initial/possible states and the rules of update. These correlations are especially interesting when shifting from an external view of the system to a situated epistemological perspective, as they appear necessary to explain how an agent could perform effectively with limited access to the external world and restrictive epistemological constraints. As mentioned in the previous sections, adhering strictly to *HWP* means that when, in Ex. 1, a has a sensation of redness at t , they do not immediately gain information about the actual state of their sensors (and memories) or of the individual(s) that affected them. They also do not gain information about the number of the individuals in their base and in the world overall. They can only “know” that the world is in a certain range of ways rather than others, insofar as they have access to their (previous) internal states (e.g., via the memory) and assume that they are robust with respect to a subclass of their direct sensors. A little help from the world thus appears *prima facie* necessary to increase the position's plausibility. Again, one should be careful not to be misled by the model's apparent simplicity on this point. The range of the possible states for an agent can be extremely wide, and the agent may aggregate a huge number of inputs to produce a relevant output. Similarly, memories might encode complex structures, becoming progressively more complex via feedback loops, further enriching the agent's behavioral repertoire. Additionally, Berto, Rossi & Tagliabue make a case for positing *strongly reversible* “worldly rules of transition”. According to the trio, different inputs cannot result in the same output (when it comes to the rules actually governing the world, though this might not hold for observational apparatuses which are actually composite mechanisms) and, their Turing-Machine-like world is time-reversal invariant [21]. As such, the ideas underlying the framework remain a promising avenue for exploration. Moreover, contrary to assumptions commonly endorsed alongside *HWP* on the upward path to structural realism, it does not seem necessary to postulate that agents have magically access to their ontological internal state for their behavior to be understood as an expression of knowledge about the world: memories show that an agent's behavior can be affected by their knowledge about the world, even if they don't have direct access to it. Conversely, as long as there is a mechanism that produces variable outputs by aggregating *external* inputs (without considering the previous states of a given agent or their memory), all the relevant functions appear to be expressed through behavior alone. Although our first example involved a human being's perception of redness, the framework can be used to describe a non-sentient individual responding to environmental stimuli, such as a thermometer. In this respect, the model treats brains and hard disks as functionally analogous. One could even extend the analogy to include anything that keeps track of an individual's previous states (previous ranges of states of the world—the states of an individual being especially salient).

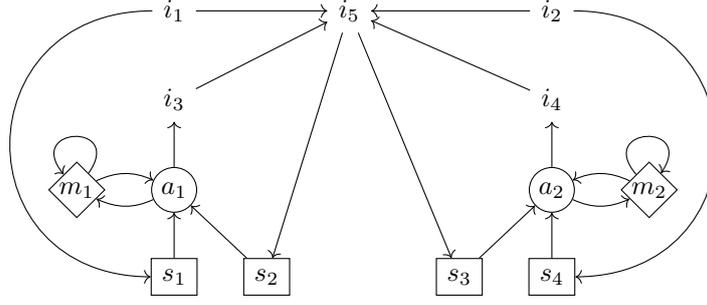


Figure 2: Illustration of the model related to Ex.2.

4. Additional Illustrative Examples

In this section, we present two further examples of increasing complexity. Ex. 2 explores interactions involving inter-agent connections where one agent updates its memory and/or takes action based on sensors directly accessible only to the other. Ex. 3 highlights the expressive power of our framework and the proposed approach by addressing a more nuanced coordination scenario that incorporates dynamic updates and goal-driven behavior. As for Ex. 1, the intuitive descriptions for Ex. 2 and Ex. 3 are provided to help readers follow the examples and situate them in a concrete scenario. However, it is always possible to take a neutral and formal reading that does not go beyond the labels of the individuals and the dependencies depicted in the figure.

Example 2 (interaction between weather responsive windows a_1 and a_2 via light i_5). This example illustrates how mechanistic connections between sensors (of different agents) can model one-way tracking and capture basic, yet meaningful, forms of “communication”. The notion of robustness is also further refined in the process. We consider the overall system in Fig. 2 designed to maximize home safety against adverse weather conditions by having one window anticipate adverse weather conditions based on the behavior of the other window. Informally, individuals a_1 and a_2 are two weather-responsive windows. Through sensor s_1 (s_4), window a_1 (a_2) indirectly accesses the weather condition of the portion of the sky above it, i_1 (i_2). The closure of window a_1 (a_2) turns on switch i_3 (i_4). The physical light i_5 , partially under i_1 and partially under i_2 , is designed to be on when one of the switches i_3 and i_4 is on and the other is off but its intended functionality is affected by the weather conditions of the portions of the sky i_1 and i_2 . Window a_1 (a_2) indirectly accesses the state of light i_5 through sensor s_2 (s_3). Memories m_1 and m_2 regulate the behavior of windows a_1 and a_2 , respectively, through a feedback loop that keeps track of previously detected weather conditions. Windows a_1 and a_2 are in open or closed states (noted \square/\square); i_1 and i_2 , as well as s_1 and s_4 , are in good or bad weather states (noted G/B), while i_3 , i_4 and i_5 , as well as s_2 and s_3 , are in on or off states (noted \oplus/\ominus). \square

Table 1 reports the states of the individuals in the system from time 0 to time 10 disregarding the dependencies of i_5 on i_1 and i_2 . The data in the table satisfy the informal update rules described above. Furthermore, for the moment, we do not consider the memories m_1 and m_2 . We can then “interpret” the data in Table 1 using the intuitive reading of Fig. 2.

First, the closure of a_1 at 3 can be attributed to the bad weather conditions that arose at location i_1 at time 1. This dependence is tracked through the ideal channel $\langle\{i_1\}, \{s_1, s_2\}, \{a_1\}\rangle$,⁶ given that the state of s_2 did not change from 1 to 2. Second, the switch of i_5 at 5 can be attributed to the closure of a_1 at 3, as tracked by the ideal channel $\langle\{a_1\}, \{i_3, i_4\}, \{i_5\}\rangle$, given that the state of i_4 did not change from 3 to 4. Third, the closure of a_2 at 7 (to maximize safety) can be traced back to the state of i_5 at 5 through the ideal channel $\langle\{i_5\}, \{s_3, s_4\}, \{a_2\}\rangle$ where s_4 reports good weather at i_2 at both 5 and 6. The last two dependencies can be composed to say that window a_2 closes at 7 in response to the closure of window a_1 at 3, i.e., the two windows “communicated”, as tracked by the ideal channel

⁶Given that in the examples the structure of the world is static, it is not necessary to temporally qualify the ideal channels.

Table 1

Extensional description of the evolution over time of the individuals in Ex. 2 assuming $\beta(i_5) = \{i_3, i_4\}$.

| | i_1 | i_2 | i_3 | i_4 | i_5 | s_1 | s_2 | s_3 | s_4 | a_1 | a_2 | m_1 | m_2 |
|-----------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| 0 | G | G | ⊖ | ⊖ | ⊖ | G | ⊖ | ⊖ | G | 0 | 0 | 000 | 000 |
| 1 | B | G | ⊖ | ⊖ | ⊖ | G | ⊖ | ⊖ | G | 0 | 0 | 000 | 000 |
| 2 | B | G | ⊖ | ⊖ | ⊖ | B | ⊖ | ⊖ | G | 0 | 0 | 000 | 000 |
| 3 | B | G | ⊖ | ⊖ | ⊖ | B | ⊖ | ⊖ | G | C | 0 | 000 | 000 |
| 4 | B | G | ⊕ | ⊖ | ⊖ | B | ⊖ | ⊖ | G | C | 0 | CC0 | 000 |
| 5 | B | G | ⊕ | ⊖ | ⊕ | B | ⊖ | ⊖ | G | C | 0 | CC0 | 000 |
| 6 | B | G | ⊕ | ⊖ | ⊕ | B | ⊖ | ⊕ | G | C | 0 | CCC | 000 |
| 7 | B | G | ⊕ | ⊖ | ⊕ | B | ⊖ | ⊕ | G | C | C | CCC | 000 |
| 8 | B | G | ⊕ | ⊕ | ⊕ | B | ⊖ | ⊕ | G | C | C | CCC | CC0 |
| 9 | B | G | ⊕ | ⊕ | ⊖ | B | ⊖ | ⊕ | G | C | C | CCC | CC0 |
| 10 | B | G | ⊕ | ⊕ | ⊖ | B | ⊖ | ⊖ | G | C | C | CCC | CCC |

$\langle \{a_1\}, \{i_3, i_4\}, \{i_5\}, \{s_3, s_4\}, \{a_2\} \rangle$ with starting points a_1 , i_4 , and s_3 , given that the state of i_4 did not from 3 to 4 and the state of s_4 did not change from 5 to 6.

Assume now that $\sigma(a_2, 11) = 0$ given the states of s_4 and s_3 at 10. Assuming the stability of i_1 and i_2 and the previously described rules, it is easy to see that the state of window a_2 commutes every four time steps. Memories play a fundamental role in avoiding this commuting. Consider then m_1 and m_2 and assume that the state of memory m_i stores the last three states of a_i . For example, $\sigma(m_i, t) = \text{CC0}$ indicates that the previous three states of a_i were C, C, and 0, where the first state in the list is the oldest. Since $\sigma(s_4, 0) = \text{G}$, $\sigma(s_3, 0) = \ominus$ and $\sigma(m_2, 0) = 000$, there is no reason to update the state of a_2 to C. This holds for a_2 until time 5. At time 2, sensor s_1 reports bad weather above a_1 . This is sufficient to close a_1 in the next time step. Since i_1 does not change anymore, independently of the state of the memory, a_1 remains closed. At time 6 the sensor s_3 reports “disagreement” between the states of the switches i_3 and i_4 . Since $\sigma(m_2, 6) = 000$, a_2 closes in the next time step because, three states ago, it was open while a_1 closed. The situation is similar at times 7, 8, and 9. At time 10, we have that $\sigma(m_2, 10) = \text{CCC}$, $\sigma(s_3, 10) = \ominus$, and $\sigma(s_4, 10) = \text{G}$. In this case, s_3 reports that a_1 and a_2 were in the same state three time steps ago and m_2 reports that a_2 was in the state C three time steps ago. Therefore, even though $\sigma(s_3, 10) = \ominus$ and $\sigma(s_4, 10) = \text{G}$, it makes sense not to open a_2 , i.e., $\sigma(a_2, 11) = \text{C}$. Commutation of a_2 is then prevented, given that m_2 remembers the state of a_2 from three time steps ago and the states of i_4 and i_3 take three time steps to propagate to a_2 .

Since there are no dependencies between i_1 , i_2 and i_5 , $\langle \{a_1, a_2\}, \{i_3, i_4\}, \{i_5\}, \{s_2\} \rangle$ is an ideal channel with a single aggregation step (from i_3 and i_4 to i_5). Now, introduce the dependencies of i_5 on both i_1 and i_2 . One might suppose that extreme weather conditions, such as lightning, could alter the behavior of light i_5 . This could cause i_5 to be in the state \ominus even if the states of i_3 and i_4 differ. In this case, given the ideal channel $\langle \{i_1, i_2, i_3, i_4\}, \{i_5\} \rangle$, i_5 is not robust with respect to $\{i_3, i_4\}$ and this could indirectly affect the behavior of a_1 and a_2 .

This suggests refining the notion of robustness of x with respect to a given ideal channel ending in x by focusing on certain ranges of the states of the starting points. In the previous example, one could exclude extreme weather conditions, by stating that robustness with respect to $\{i_3, i_4\}$ is guaranteed only when i_1 and i_2 are not in extreme weather conditions (i.e., provided that i_1 and i_2 are not in certain states). Alternatively, one could assume that (i) i_3 and i_4 have multiple possible states that correspond to different degrees of “intensity” of being on or off; and (ii) extreme weather conditions do not impact the aggregation step provided by i_5 when the intensity of its inputs exceeds a given threshold.

Finally, it is worth pointing out that, while this case might appear to be one involving predetermined compatibility, i.e., privileged one-to-one connections dependent on inherent similarities that facilitate mutual understanding, or *ad-hoc* interfaces between specific system types, similar exchanges are always possible insofar as one agent produces an output which can be reliably tracked by the other. There are thus no requirements concerning the agents themselves, but only the channels connecting them.

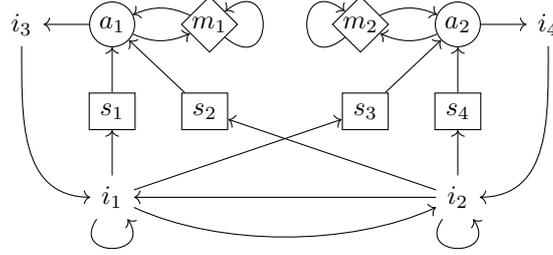


Figure 3: Illustration of the model related to Ex.3.

Example 3 (goal-directed coordination between agents a_1 and a_2). This example involves goal-directed coordination. Due to the complexity of the scenario we only provide an abstract description of the underlying processes omitting the formal details. We consider all the individuals and dependencies in Fig. 3. Informally, individuals i_1 and i_2 are concrete blocks which are in physical contact with each other. Individuals i_3 and i_4 are coolers/heaters connected to i_1 and i_2 , respectively. Sensors s_1 , s_2 , s_3 , and s_4 are analog mercury column thermometers. Agents a_1 and a_2 control coolers/heaters i_3 and i_4 , respectively. Depending on the states of memories m_1 and m_2 and of thermometers on their private consoles, they set the level of activation of coolers/heaters. \square

We start by considering a scenario in which sensors s_2 and s_3 are absent. In this case, the agents only have access to the state of the block affected by the cooler/heater that they control. Suppose also that a_1 's goal concerns only the state of i_1 , and a_2 's goal concerns only the state of i_2 . In this scenario, "coordination" is achieved only when the world behaves in a way that allows both agents to achieve their goals. For example, a situation where both the agents intend to keep the block more directly under their control at a temperature of 30° is plausibly achievable. The coolers/heaters can be activated for just enough time to reach the desired temperature, which can then be maintained with successive intermittent activations. Insofar as the temperatures of blocks i_1 and i_2 are similar, there is minimal energy exchange between them, resulting in a relatively stable system. Even when agents intend to maintain different temperatures for each block, both goals seem potentially achievable. However, the system is unstable due to energy transfer from the hottest block to the colder one.

A second scenario includes sensors s_2 and s_3 . Still agent a_1 acts directly only on the temperature of block i_1 , and a_2 acts directly only on the temperature of block i_2 but now both agents can monitor the temperature of the two blocks. In this scenario, the two agents share the goal of achieving a specific temperature configuration for the blocks, i.e., entering in respective states corresponding to the world being in a specific range of states. Here, coordination is intended in a more explicit and stronger way, i.e., as the sharing of a goal (in the previous scenario, each agent had a private goal). Suppose that the agents have the shared goal of setting the temperature of i_1 to 70° and the temperature of i_2 to 10° . This situation is similar to the one discussed in the previous scenario, but now the agents "know" the temperature difference between the two blocks. This information could support input-output relations characterizing the agentive system which more efficiently lead to the desired goals. However, it is unclear whether this increase in efficiency is solely due to the additional information or if the agents' different goals also play a role.

4.1. Discussion

Taking stock of the three examples discussed, and returning to the questions that motivated this inquiry, it appears increasingly plausible that, under reasonable mechanistic presuppositions about how world-states depend one another, coordination does not necessarily require shared language, semantics or the endorsement of a common ontology (at least given the specifications of the notions we considered), but can also emerge from agential mechanisms being influenced, and influencing, the world in accordance with the world's laws. In itself, this suggests that alternative approaches for the achievement of

coordination, or relatively to the acquisition, transmission and aggregation of information about the world—approaches less tightly coupled to human-centric models of communication and cognition, but rather reminiscent of programming—might be worth considering and exploring.

Of course, the modeling of the examples is extremely simplistic in relation to our goals, and the results rest on various assumptions, leaving open reasonable doubts about whether coordination could be achieved in more complex toy-models, or in relation to more complex tasks or goals. In general, it might be questioned whether high-level cognitive functions (such as language use, theory-building, or conceptual abstraction) are ultimately (directly, or indirectly—via systems’ design) indispensable for coordination in real-world settings. However, rather than undermining the proposed framework, these doubts highlight what we take to be a productive line of inquiry. For instance, it appears pertinent to ask whether linguistic and cognitive function can fit within a mechanistic setting of information exchange—also connecting to questions currently raised concerning large language models. In this context, even hard-limits could reveal insights into how semantic artifacts could be improved.

One could consider the possibility of building (or aligning) artifacts fully bottom up, exploiting worldly connections to ensure that they reliably track world states.⁷ In relation to this point, other questions worth asking are whether, under which conditions, and to what degree, situated agents could “reverse engineer” their world’s structure given the minimal epistemological assumptions endorsed in our exploratory investigation: our framework, once expanded and adjusted, might provide a basis for practical work in this direction, examining toy models given different sets of constraints. Particularly interesting seems the application of techniques resting on machine learning—and much work has already been done in this direction in other contexts, such as robotics and data-driven scientific inquiry. Conversely, the kind of “minimal information” about the world resting on the *HWP* appears to provide solid foundations on which to apply such techniques, being as “raw” and less “theory-laden” as one can imagine. Nevertheless, even if such a project were viable, a connection between the resulting representations and human-friendly conceptual schemas or ontologies would likely remain necessary, if explainability is to be retained, as suggested by [12].

5. Concluding Remarks

In this paper, we critically examined the widely presumed foundational link between semantic technologies and interoperability, emphasizing how this connection is contingent upon how the latter is understood. Our analysis revealed salient conceptual divergences in the literature, which not only obscure the discussion but may also lead to misguided expectations among stakeholders. We argued that, insofar as “interoperability” is understood in terms of (the achievement of) *coordination* among systems, the adoption of a common semantic artifact (and endorsement of the encoded semantics) is neither a necessary nor a sufficient condition—provided that coordination can undoubtedly be facilitated, and in some cases actually achieved, in conjunction with tacit agreements and prior alignments among agents, thus underscoring the importance of factual sharedness. What was initially taken to be a foundational link at the heart of the discipline, seemingly implicitly assumed across research programs, thus emerges as a complex interplay calling for deeper philosophical and practical scrutiny, and leaving room for the exploration of alternative approaches.

To those ends, we proposed a framework inspired by reflections central to the upward path to structural realism, focusing specifically on the *Helmholtz-Weyl Principle*, adapted to present purposes. Through (relatively) simple illustrative examples, we explored the extent to which coordination capabilities, typically associated with high-level cognitive functions, can be recovered within a purely mechanistic and deterministic setting, while emphasizing pragmatic aspects related to data acquisition, transmission, and aggregation. On the one hand, our approach moves towards demonstrating that coordination among diverse (not inherently-compatible), possibly non-sentient agents can be

⁷It is worth reminding that, from the point of view of a situated agent with no direct access to the states of a certain system, and of the world in general, this reduces to connections among perceptions at different times. Nevertheless, one is naturally free to produce revisable hypotheses concerning the underlying mechanisms.

achieved without reliance on shared semantics or a common worldview. On the other hand, we hinted at the fact that, “raw information about the world”—resting on the presupposed capability of situated agents to reliably and objectively discriminate among different ranges of states of the world—might provide grounds for the bottom-up development of knowledge systems and agentive ones, or to connect disparate worldviews, theories, and, for what concerns us specifically, ontologies. While ontologies remain pivotal for explainability, especially for epistemic agents with limited cognitive resources, both apparent and factual agreement or disagreement can emerge at the level of constructed posits.⁸ Likewise, program-like mechanistic input-output connections grounded in worldly dependencies might ensure the achievement of coordination without relying on sharedness, but instead progressively leading to it.

While the purported results of our examples may not come as a surprise to those familiar with the cybernetics tradition, and are by no means definitive given the high-level nature of our models, they may nonetheless offer thought-provoking insights for researchers in knowledge representation and applied ontology. We hope this work contributes to ongoing discussions about the epistemological and practical foundations of coordination, and encourages further exploration of alternative and complementary approaches, while at the same time improving conceptual clarity around interoperability. Needless be said, much work remains to be done to even just fix, let alone concretize, the ideas outlined in this exploratory paper, yet some of the points touched in the examples’ discussion, in relation to machine learning and the bottom-up engineering of knowledge systems, seem to us promising lines of research.

Acknowledgments

We would also like to thank three anonymous reviewers for their helpful comments to previous versions of this manuscript.

Declaration on Generative AI

During the preparation of this work, the authors used M365 CoPilot and DeepL in order to: grammar and spelling check, paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] I. Votsis, The upward path to structural realism, *Philosophy of Science* 72 (2005) 1361–1372. doi:10.1086/508974.
- [2] IEEE Standard Glossary of Software Engineering Terminology, *IEEE Std 610.12-1990* (1990) 1–84. doi:10.1109/IEEESTD.1990.101064.
- [3] H. Panetto, Towards a classification framework for interoperability of enterprise applications, *International Journal of Computer Integrated Manufacturing* 20 (2007) 727–740. doi:10.1080/09511920600996419.
- [4] M. D. Wilkinson, et al., The FAIR Guiding Principles for scientific data management and stewardship, *Scientific data* 3 (2016). doi:10.1038/sdata.2016.18.
- [5] Regulation (EU) 2023/2854 of the European Parliament on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act), https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202302854, 2023.
- [6] A. L. Szejka, O. Canciglieri, H. Panetto, E. Rocha Loures, A. Aubry, Semantic interoperability for an integrated product development process: a systematic literature review, *International Journal of Production Research* 55 (2017) 6691–6709. doi:10.1080/00207543.2017.1346314.
- [7] M. Noura, M. Atiquzzaman, M. Gaedke, Interoperability in Internet of Things Infrastructure: Classification, Challenges, and Future Work, in: Y.-B. Lin, D.-J. Deng, I. You, C.-C. Lin (Eds.),

⁸Compare with [22, 23], which also present an approach to describe the relevant phenomena ontologically.

- IoT as a Service, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Springer International Publishing, Cham, 2018, pp. 11–18. doi:10.1007/978-3-030-00410-1_2.
- [8] T. Hagelien, H. A. Preisig, J. Friis, P. Klein, N. Konchakova, A Practical Approach to Ontology-Based Data Modelling for Semantic Interoperability, 14th WCCM-ECCOMAS Congress 2020 (2021). doi:10.23967/wccm-eccomas.2020.035.
- [9] J. Euzenat, Towards a principled approach to semantic interoperability, in: A. Gómez-Pérez, M. Gruninger, H. Stuckenschmidt, M. Uschold (Eds.), Proceedings of the IJCAI-01 Workshop on Ontologies and Information Sharing Seattle, USA, August 4-5, 2001, volume 47 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2001. URL: <https://ceur-ws.org/Vol-47/euzenat.pdf>.
- [10] J. Euzenat, P. Shvaiko, *Ontology Matching*, 2nd ed., Springer, Heidelberg, 2013.
- [11] H. Halvorson, *The Logic in Philosophy of Science*, Cambridge University Press, Cambridge and New York, 2019.
- [12] G. Guizzardi, N. Guarino, Explanation, semantics, and ontology, *Data & Knowledge Engineering* 153 (2024) 102325. doi:<https://doi.org/10.1016/j.datak.2024.102325>.
- [13] H. Poincaré, *Science and Hypothesis*, Scott, New York, 1905.
- [14] B. Russell, *The Analysis of Matter*, Kegan Paul, London, 1927.
- [15] A. Meinong, *Untersuchungen zur Gegenstandstheorie und Psychologie*, J.A. Barth, Leipzig, 1904. From the collections of the New York Public Library.
- [16] K. Mulligan, Early analytic philosophy’s austrian dimensions, in: A. Coliva, P. Leonardi, S. Moruzzi (Eds.), *Eva Picardi on Language, Analysis and History*, Palgrave, 2018, pp. 7–29.
- [17] S. Psillos, Is structural realism possible?, *Philosophy of Science* 68 (2001) 13–24. doi:10.1086/392894.
- [18] A. Rayo, *The Construction of Logical Space*, Oxford University Press, Oxford, England, 2013.
- [19] M. Boon, Instruments in science and technology, in: J. K. B. O. Friis, S. A. Pedersen, V. F. Hendricks (Eds.), *A Companion to the Philosophy of Technology*, Wiley-Blackwell, 2012, pp. 78–83.
- [20] H. Chang, *Inventing Temperature: Measurement and Scientific Progress*, Oxford University Press, 2004. URL: <https://doi.org/10.1093/0195171276.001.0001>. doi:10.1093/0195171276.001.0001.
- [21] F. Berto, G. Rossi, J. Tagliabue, *The mathematics of the models of reference*, College Publications, New York, 2010.
- [22] A. Gangemi, P. Mika, Understanding the semantic web through descriptions and situations, in: R. Meersman, Z. Tari, D. C. Schmidt (Eds.), *On The Move to Meaningful Internet Systems 2003: CoopIS, DOA, and ODBASE*, Springer, Berlin, Heidelberg, 2003, pp. 689–706.
- [23] A. Gangemi, V. Presutti, *Formal Representation and Extraction of Perspectives*, *Studies in Natural Language Processing*, Cambridge University Press, 2022, p. 208–228.