

Progressive Ontology Alignment for Meaning Coordination: An Information-Theoretic Foundation

Marco Schorlemmer

Institut d'Investigació en Intel·ligència Artificial
Consell Superior d'Investigacions Científiques
08193 Bellaterra (Barcelona), Catalonia, Spain
marco@iiaa.csic.es

Yannis Kalfoglou

Advanced Knowledge Technologies (AKT)
University of Southampton
Southampton SO17 1BJ, England, UK
y.kalfoglou@ecs.soton.ac.uk

Abstract

We elaborate on the mathematical foundations of the meaning coordination problem that agents face in open environments. We investigate to which extent the Barwise-Seligman theory of information flow provides a faithful theoretical description of the partial semantic integration that two agents achieve as they progressively align their underlying ontologies through the sharing of tokens, such as instances. We also discuss the insights and practical implications of the Barwise-Seligman theory with respect to the general meaning coordination problem.

1. Introduction

For two agents to interoperate, exchanging vocabulary and syntax is insufficient, because agents also need to agree upon the meaning of the communicated syntactic constructs. Separate agents, though, are most often engineered assuming different, sometimes even incompatible, conceptualisations. Ontologies have been advocated as a solution to this semantic heterogeneity: separate agents would need to match their own conceptualisations against a common ontology of the application domain, so that all communication is done according to the constraints derived from the ontology.

Although the use of ontologies may indeed favour semantic interoperability, it relies on the existence of agreed domain ontologies in the first place. Furthermore, these ontologies will have to be as complete and as stable for a domain as possible, because different versions only introduce more semantic heterogeneity. Thus, semantic-integration approaches based on *a priori* common domain ontologies may be useful for clearly delimited and stable domains, but they are

untenable and even undesirable in highly distributed, open, and dynamic environments such as those encountered in multi-agent systems. In such environments, it is more realistic to progressively achieve certain levels of semantic interoperability by coordinating and negotiating the meaning attached to syntactic constructs on the fly, as done, for instance, in approaches by Bailin and Truszkowski [1] or by Wang and Gasser [10]. Although we are skeptical that *meaning* as such can ever be coordinated or negotiated in a way such that all agents share the understanding of a communicated concept, we do argue that communication between separate agents will hardly ever be achieved if we lack the necessary commodity for meaning to be coordinated and negotiated in the first place: information.

This puts us within the philosophical tradition put forth by Dretske [4], which sees information as prior to meaning, namely as an interpretation-independent objective commodity that can be studied by its own right. Consequently, we believe that any satisfactory formalisation of semantic interoperability needs to be built upon a mathematical theory capable of describing under which circumstances information flow occurs. We shall use Barwise and Seligman's channel theory for this purpose [2]. It constitutes a general mathematical theory that aims at describing the information flow in any kind of distributed system.

Previously, Kalfoglou and Schorlemmer have been starting from the Barwise-Seligman theory of information flow in order to formalise and automate semantic interoperability [6, 7]. In this paper, though, we investigate the ways in which the Barwise-Seligman theory applies to the problem of meaning coordination. We do not present a fully-fledged theory for meaning coordination, nor do we provide a meaning coordination methodology or procedure. Instead, our aim here is to explore how the insights about information and its flow provided by the Barwise-Seligman theory translate to

the meaning coordination problem.

2. Meaning Coordination

Before applying all the channel-theoretic machinery to the meaning coordination problem, we first need to delimit the problem and state the assumptions upon which we build the theoretical framework.

We assume a scenario in which two agents A_1 and A_2 want to interoperate, but in which each agent A_i has its knowledge represented according to its own conceptualisation, which we assume is explicitly specified according to its own ontology O_i . By this we mean a concept of O_1 will always be considered semantically distinct *a priori* from any concept of O_2 , even if they happen to be syntactically equal, unless the meaning coordination process unveils sufficient semantic evidence that it means the same to A_1 as it does to A_2 . Furthermore, we assume that the agents' ontologies are not open to other agents for inspection, so that semantic heterogeneity cannot be solved by "looking into each agents' head." Hence, an agent may learn about the ontology of another agent only through interaction. Thus, following an approach similar to that of Wang and Gasser described in [10], if A_1 wants to explain A_2 the meaning of a concept, it can use a *token* of this concept, such as an instance classified under this concept, as a representation of it.

Take, for example, the issues one has to take into account when we need to align government ministries and departments across different countries. This is a realistic scenario set out in the domain of e-governments. Our agents will have to align different conceptualisations of governmental structures as they reflect different ways of allocating responsibilities to ministries and departments. For the sake of brevity and space reasons, we only focus on four ministries—The UK Foreign and Commonwealth Office, the UK Home Office, the US Department of State, the US Department of Justice (hereafter, FCO, HO, DoS and DoJ, respectively)—and on a subset of their responsibilities as gathered from their web sites (accessible from www.homeoffice.gov.uk, www.fco.gov.uk, www.state.gov and www.usdoj.gov) and shown in Table 1.

Given these different conceptualisations, though, a UK-centred agent A_1 may explain to a US-centred agent A_2 what Home Office means by informing A_2 that to "regulate entry and settlement in the UK" is among its responsibilities. Here "regulate entry and settlement in the UK" acts as a *token* of Home Office. In principle, agents may well express responsibilities differently, since A_1 is situated in the context of the UK government, while A_2 is situated in the con-

ID	UK responsibilities
r_1	issues passports
r_2	regulate entry and settlement in the UK
r_3	executive services of the HO
r_4	promote productive relations
r_5	responsible for the work of FCO
ID	US responsibilities
s_1	passport services and information
s_2	promotes government interests in the region
s_3	heading the DoS
s_4	facilitate entry to the US
s_5	supervise and direct the DoJ

Table 1. Government responsibilities

text of the US one. But, for any successful explanation of foreign concepts by exchanging their tokens, it is sensible to assume that A_2 will be able to classify any new token coming from A_1 —a responsibility assertion in our example scenario—according to its own ontology, and vice versa. Thus, theoretically speaking, all tokens belong to a *domain of discourse*, which we denote D and may well be infinite. In our example D would include responsibility assertions (our tokens $r_1, r_2, r_3, r_4, r_5, \dots, s_1, s_2, s_3, s_4, s_5, \dots$). Although each agent manages its own finite subset of D , we assume that it is capable of processing any new token of D it may get, and classify it according to its own ontology. The focus of this paper, though, is not on how this classification is done.¹

In fact, by lacking any *a priori* domain ontology about government ministries and departments, it is hard to see how agents A_1 and A_2 could coordinate meaning as made explicit in their ontologies O_1 and O_2 in another way. It is the assumption that A_1 and A_2 are capable of classifying tokens with respect to their own conceptualisation which makes our approach to meaning coordination possible. Meaning coordination is then the progressive sharing of tokens of this domain of discourse and the subsequent mutual communication about how they are classified according to each ontology.

3. Channel-Theoretic Preliminaries

We introduce briefly the main channel-theoretic constructs needed for our foundation for meaning coordination. As we proceed, we shall hint at the intuitions

1 In this example scenario this could be done, for instance, by performing text processing on the responsibility assertions to identify relevant keywords and exploring their synonyms in public available thesauri such as WordNet®.

lying behind them, but a proper in-depth understanding of the theory is outside the scope of this paper, and we refer the interested reader to [2]. In the remainder of the paper we use the prefix ‘IF’ (information flow) in front of some of the channel-theoretic terminology to distinguish it from their usual meaning.

3.1. IF Classification, Infomorphism, and Channel

In channel theory, each component (or context) of a distributed system is modelled by means of an *IF classification*. The system itself is described by the way IF classifications are connected with each other through *infomorphisms*.

Definition 1 An *IF classification* $\mathbf{A} = \langle tok(\mathbf{A}), typ(\mathbf{A}), \models_{\mathbf{A}} \rangle$, consists of a set of *tokens* $tok(\mathbf{A})$, a set of *types* $typ(\mathbf{A})$ and a *classification relation* $\models_{\mathbf{A}} \subseteq tok(\mathbf{A}) \times typ(\mathbf{A})$ that classifies tokens to types.

Definition 2 An infomorphism $f = \langle f^{\wedge}, f^{\vee} \rangle : \mathbf{A} \rightarrow \mathbf{B}$ from IF classifications \mathbf{A} to \mathbf{B} is a contravariant pair of functions $f^{\wedge} : typ(\mathbf{A}) \rightarrow typ(\mathbf{B})$ and $f^{\vee} : tok(\mathbf{B}) \rightarrow tok(\mathbf{A})$ satisfying the following fundamental property, for each type $\alpha \in typ(\mathbf{A})$ and token $b \in tok(\mathbf{B})$:

$$\begin{array}{ccc} \alpha \vdash_{\mathbf{A}} & \xrightarrow{f^{\wedge}} & f^{\wedge}(\alpha) \\ \vdash_{\mathbf{A}} \downarrow & & \downarrow \vdash_{\mathbf{B}} \\ f^{\vee}(b) \vdash_{\mathbf{A}} & \xleftarrow{f^{\vee}} & b \end{array}$$


$$f^{\vee}(b) \models_{\mathbf{A}} \alpha \quad \text{iff} \quad b \models_{\mathbf{B}} f^{\wedge}(\alpha)$$

Definition 3 A *distributed IF system* \mathcal{A} consists of an indexed family $cla(\mathcal{A}) = \{\mathbf{A}_i\}_{i \in I}$ of IF classifications together with a set $inf(\mathcal{A})$ of infomorphisms all having both domain and codomain in $cla(\mathcal{A})$.

The basic construct of channel theory is that of an *IF channel* between two IF classifications. It models the information flow between components:

Definition 4 An *IF channel* consists of two IF classifications \mathbf{A}_1 and \mathbf{A}_2 connected through a core IF clas-

sification \mathbf{C} via two infomorphisms f_1 and f_2 :



3.2. IF Theory and Logic

Channel theory is based on the understanding that the flow of information is a result from the regularities of a distributed system. These regularities are implicit in the representation of the system as a distributed IF system of connected IF classifications, but we can make them explicit in a logical fashion by means of IF theories and IF logics:

Definition 5 An *IF theory* $T = \langle typ(T), \vdash \rangle$ consists of a set $typ(T)$ of types, and a binary relation \vdash between subsets of $typ(T)$. Pairs $\langle \Gamma, \Delta \rangle$ of subsets of $typ(T)$ are called *sequents*. If $\Gamma \vdash \Delta$, for $\Gamma, \Delta \subseteq typ(T)$, then the sequent $\Gamma \vdash \Delta$ is called a *constraint*. T is *regular* if for all $\alpha \in typ(T)$ and all sets $\Gamma, \Gamma', \Delta, \Delta', \Sigma', \Sigma_0, \Sigma_1$ of types:

1. *Identity*: $\alpha \vdash \alpha$
2. *Weakening*: If $\Gamma \vdash \Delta$, then $\Gamma, \Gamma' \vdash \Delta, \Delta'$
3. *Global Cut*: If $\Gamma, \Sigma_0 \vdash \Delta, \Sigma_1$ for each partition $\langle \Sigma_0, \Sigma_1 \rangle$ of Σ' , then $\Gamma \vdash \Delta$.²

Definition 6 An *IF logic* $\mathcal{L} = \langle tok(\mathcal{L}), typ(\mathcal{L}), \models_{\mathcal{L}}, \vdash_{\mathcal{L}}, N_{\mathcal{L}} \rangle$ consists of an IF classification $cla(\mathcal{L}) = \langle tok(\mathcal{L}), typ(\mathcal{L}), \models_{\mathcal{L}} \rangle$, a regular IF theory $th(\mathcal{L}) = \langle typ(\mathcal{L}), \vdash_{\mathcal{L}} \rangle$ and a subset of $N_{\mathcal{L}} \subseteq tok(\mathcal{L})$ of *normal tokens*, which satisfy all the constraints of $th(\mathcal{L})$; a token $a \in tok(\mathcal{L})$ satisfies a constraint $\Gamma \vdash \Delta$ of $th(\mathcal{L})$ if, when a is of all types in Γ , a is of some type in Δ . An IF logic \mathcal{L} is *sound* if $N_{\mathcal{L}} = tok(\mathcal{L})$.

Regularity arises from the observation that, given any classification of tokens to types, the set of all sequents that are satisfied by all tokens always fulfill Identity, Weakening, and Global Cut.


Every classification determines a *natural IF logic*, which captures the regularities of the classification in a logical fashion.

² A partition of Σ' is a pair $\langle \Sigma_0, \Sigma_1 \rangle$ of subsets of Σ' , such that $\Sigma_0 \cup \Sigma_1 = \Sigma'$ and $\Sigma_0 \cap \Sigma_1 = \emptyset$; Σ_0 and Σ_1 may themselves be empty (hence it is actually a quasi-partition).

Definition 7 The *natural IF logic* is the IF logic $Log(\mathbf{C})$ generated from an IF classification \mathbf{C} , and has as classification \mathbf{C} , as regular theory the theory whose constraints are the sequents satisfied by all tokens, and whose tokens are all normal.

3.3. Distributed IF Logic

The key channel-theoretic construct we shall use in order model the semantic interoperability between agents with different ontologies is that of a *distributed IF logic*, which is the logic that represents the flow of information occurring in a distributed system. Semantic interoperability between agents A_1 and A_2 is then described by the IF theory of the distributed IF logic of IF channel



representing the information flow between A_1 and A_2 , and which describes how the different types from A_1 and A_2 are logically related to each other, both respecting the local IF classification systems of each agent and interrelating types whenever there is a similar semantic pattern (i.e., a similar way communities classify related tokens). The distributed IF logic is defined by *moving* an IF logic on the core \mathbf{C} of the channel to the sum of components $A_1 + A_2$.

Definition 8 Given an infomorphism $f : \mathbf{A} \rightarrow \mathbf{B}$ and an IF logic \mathcal{L} on \mathbf{B} , the *inverse image* $f^{-1}[\mathcal{L}]$ of \mathcal{L} under f is the IF logic on \mathbf{A} , whose theory is such that $\Gamma \vdash \Delta$ is a constraint of $th(f^{-1}[\mathcal{L}])$ iff $f^{\sim}[\Gamma] \vdash f^{\sim}[\Delta]$ is a constraint of $th(\mathcal{L})$, and whose normal tokens are $N_{f^{-1}[\mathcal{L}]} = \{a \in tok(\mathbf{A}) \mid a = f^{\sim}(b) \text{ for some } b \in N_{\mathcal{L}}\}$. If f^{\sim} is surjective on tokens and \mathcal{L} is sound, then $f^{-1}[\mathcal{L}]$ is sound.

Definition 9 Given an IF channel $\mathcal{C} = \{f_{1,2} : \mathbf{A}_{1,2} \rightarrow \mathbf{C}\}$ and an IF logic \mathcal{L} on its core \mathbf{C} , the *distributed IF logic* $DLog_{\mathcal{C}}(\mathcal{L})$ is the inverse image of \mathcal{L} under the sum infomorphisms $f_1 + f_2 : \mathbf{A}_1 + \mathbf{A}_2 \rightarrow \mathbf{C}$. This sum is defined as follows: $\mathbf{A}_1 + \mathbf{A}_2$ has as set of tokens the Cartesian product of $tok(\mathbf{A}_1)$ and $tok(\mathbf{A}_2)$ and as set of types the disjoint union of $typ(\mathbf{A}_1)$ and $typ(\mathbf{A}_2)$, such that for $\alpha \in typ(\mathbf{A}_1)$ and $\beta \in typ(\mathbf{A}_2)$, $\langle a, b \rangle \models_{\mathbf{A}_1 + \mathbf{A}_2} \alpha$ iff $a \models_{\mathbf{A}_1} \alpha$, and $\langle a, b \rangle \models_{\mathbf{A}_1 + \mathbf{A}_2} \beta$ iff $b \models_{\mathbf{A}_2} \beta$. Given two infomorphisms $f_{1,2} : \mathbf{A}_{1,2} \rightarrow \mathbf{C}$, the sum $f_1 + f_2 : \mathbf{A}_1 + \mathbf{A}_2 \rightarrow \mathbf{C}$ is defined by $(f_1 + f_2)^{\sim}(\alpha) = f_i^{\sim}(\alpha)$ if $\alpha \in \mathbf{A}_i$ and $(f_1 + f_2)^{\sim}(c) = \langle f_1^{\sim}(c), f_2^{\sim}(c) \rangle$, for $c \in tok(\mathbf{C})$.

3.4. Ontologies in Channel Theory

For the purposes of meaning coordination described in this paper, we adopt a definition of ontology that includes some of its core components: *Concepts*, organized in an *is-a hierarchy*, and notions of *disjointness* of two concepts—when no instance can be considered of both concepts—and *coverage* of two concepts—when all instances are covered by two concepts.³ Disjointness and coverage are typically specified by means of ontological axioms. In this paper we take these kind of axioms into account including disjointness and coverage into the hierarchy of concepts by means of two binary relations ‘ \perp ’ and ‘ $|$ ’, respectively. In [6], Kalfoglou and Schorlemmer included also *relations* over concepts in their core treatment of ontologies. We have left them out here for the ease of presentation.

Definition 10 An *ontology* is a tuple $\mathcal{O} = (C, \leq, \perp, |)$ where

1. C is a finite set of concept symbols;
2. \leq is a reflexive, transitive and anti-symmetric relation on C (a partial order);
3. \perp is a symmetric and irreflexive relation on C (disjointness);
4. $|$ is a symmetric relation on C (coverage); and

When an ontology $\mathcal{O} = (C, \leq, \perp, |)$ is used in some particular application domain, we need to populate it with instances. First, we will have to classify objects of a set X according to the concept symbols in C by defining a binary classification relation $\models_{\mathbf{C}}$. This determines an IF classification $\mathbf{C} = (X, C, \models_{\mathbf{C}})$, where $X = tok(\mathbf{C})$ and $C = typ(\mathbf{C})$. The classification relation $\models_{\mathbf{C}}$ will have to be defined in such a way that the partial order \leq , the disjointness \perp , and the coverage $|$ are respected:

Definition 11 A *populated ontology* is a tuple $\tilde{\mathcal{O}} = (\mathbf{C}, \leq, \perp, |)$ such that $\mathbf{C} = (X, C, \models_{\mathbf{C}})$ is an IF classification, and $\mathcal{O} = (C, \leq, \perp, |)$ is an ontology, and for all $x \in X$ and $c, d \in C$,

1. if $x \models_{\mathbf{C}} c$ and $c \leq d$, then $x \models_{\mathbf{C}} d$;
2. if $x \models_{\mathbf{C}} c$ and $c \perp d$, then $x \not\models_{\mathbf{C}} d$;
3. if $c | d$, then $x \models_{\mathbf{C}} c$ or $x \models_{\mathbf{C}} d$.

Our approach to meaning coordination uses the fact that, in the context of channel theory, a populated

³ Both disjointness and coverage can easily be extended to more than two concepts.

ontology $\tilde{\mathcal{O}} = (\mathbf{C}, \leq, \perp, |)$ —with $\mathbf{C} = (X, C, \models_{\mathbf{C}})$ —determines a local logic $\mathcal{L} = (X, C, \models_{\mathbf{C}}, \vdash)$ whose theory (C, \vdash) is given by the smallest regular consequence relation (i.e., the smallest relation closed under Identity, Weakening, and Global Cut) such that, for all $c, d \in C$:

$$c \vdash d \text{ iff } c \leq d \quad c, d \vdash \text{ iff } c \perp d \quad \vdash c, d \text{ iff } c | d$$

4. Progressive Semantic Integration


In order to formalise the semantic integration of a collection of agents via the precise mathematical construct of an IF channel, Kalfoglou and Schorlemmer articulated in [7] the following four steps:

1. Modelling the populated ontologies of agents by means of IF classifications.
2. Defining an IF channel—its core and infomorphisms—connecting the agents' IF classifications.
3. Defining an IF logic on the core of the IF channel representing the information flow between agents.
4. Distributing the IF logic to the sum of agent IF classifications to obtain the IF theory that describes the desired semantic interoperability.

They pointed out that these steps had to be understood in the context of a theoretical exercise and would hardly be implemented directly as engineering steps in actual interoperability scenarios. Indeed, the definition of an IF channel and an IF logic on the core of this channel representing the information flow between agents (steps 2 and 3) requires a global view of all involved parties, which we seldom will possess in general. On the contrary, in this paper we started from the assumption that the agents' ontologies are not open to other agents for inspection, and that an agent learns about the ontology of another agent only through interaction.

4.1. The Global Ontology

The four steps above determine what we call the *global ontology* of two semantically integrated agents A_1 and A_2 . It is the IF theory of the distributed IF logic of an IF channel \mathcal{C} connecting IF classifications \mathbf{A}_1 and \mathbf{A}_2 modelling the agents' populated ontologies $\tilde{\mathcal{O}}_1$ and $\tilde{\mathcal{O}}_2$ respectively:



At the core of IF channel \mathcal{C} , $typ(\mathbf{C})$ covers $typ(\mathbf{A}_1)$ and $typ(\mathbf{A}_2)$, while the elements of $tok(\mathbf{C})$ connect tokens from $tok(\mathbf{A}_1)$ with tokens from $tok(\mathbf{A}_2)$. By defining an IF logic on the core of the channel and distributing it to the sum of IF classifications $\mathbf{A}_1 + \mathbf{A}_2$ we get the *global ontology* that captures the overall semantic integration of the scenario.




Figure 1. Aligning ontologies with a pair of maps

For example, an IF channel for the UK-US government alignment scenario of Section 2 is shown in Figure 1. It corresponds to the globally constructed alignment described by Kalfoglou and Schorlemmer in [7]. At the core of this channel the connections $\langle r_2, s_4 \rangle$, $\langle r_1, s_1 \rangle$, and $\langle r_4, s_2 \rangle$ link particular tokens (i.e., responsibilities among those shown in Table 1) of type HO or FCO together with particular tokens of type DoJ or DoS in such a way that their resulting classification into the four concepts HO, FCO, DoJ, and DoS, determines an IF theory about how these concepts are semantically related. This theory is given by the distributed IF logic of the natural IF logic of the core classification: $DLogc(Log(\mathbf{C}))$. It includes among its constraints:

$$\begin{array}{ll} \vdash \text{HO, DoS} & \text{DoJ} \vdash \text{HO} \\ \text{FCO} \vdash \text{DoS} & \text{DoJ, FCO} \vdash \end{array}$$

i.e., that $\text{HO} | \text{DoS}$, $\text{DoJ} \leq \text{HO}$, $\text{FCO} \leq \text{DoS}$, and $\text{DoJ} \perp \text{FCO}$. Other IF channels modelling a different semantic integration are possible in principle, although this one reflects the particular relationship linking together immigration control (r_2 and s_4), passport services (r_1 and s_1), and promotion of productive relations (r_4 and s_2), which was taken as given in [7].

In meaning coordination scenarios we cannot assume that we will be able to define a global IF channel that connects \mathbf{A}_1 and \mathbf{A}_2 directly, capturing thus their semantic integration. In the channel of Figure 1, for example, it is not clear from where we would gain the additional understanding that allowed us to link tokens

in the way we did. Nor can we assume that we ever will be able to define such a channel completely, linking all tokens and defining an IF theory on the union of all types. Therefore, the global IF channel is not appropriate as a mathematical model for describing the process of meaning coordination.


4.2. The Coordinated Channel

We shall model meaning coordination with a *coordinated channel* instead, an IF channel that captures how O_1 and O_2 are progressively coordinated, and which captures the semantic integration achieved through interaction between A_1 and A_2 . As we have described in Section 2, if A_1 wants to explain A_2 the meaning of a concept, it can do so using a token of this concept as a representation of it.

The coordinated channel is a mathematical model of this coordination that captures the *degree of participation* of an agent A_i at any stage of the coordination process. This degree is determined both, at the type and at the token level, since

- an agent A_i will have attempted to explain a subset of its concepts to other agents, and
- other agents will have exchanged with agent A_i some of its tokens, incrementing in this way the set of tokens originally available to agent A_i .


This degree of participation can be captured in a straightforward way with an infomorphism $g_i : \mathbf{A}'_i \rightarrow \mathbf{A}_i$, for which functions f_i and g_i are the inclusions $typ(\mathbf{A}'_i) \subseteq typ(\mathbf{A}_i)$ and $tok(\mathbf{A}_i) \subseteq tok(\mathbf{A}'_i)$, respectively. The coordination is then established not between the original IF classifications \mathbf{A}_i , but between the *subclassifications* \mathbf{A}'_i that result from the interaction carried out so far:



In Section 2 we argued that although agents may handle different token sets, any successful explanation of foreign concepts by exchanging tokens will need to assume that A_2 is able to identify tokens of A_1 as belonging to a theoretical domain of discourse D common to its own tokens, and that it will be able to classify, in theory, any element of D according to its own ontology. We also assumed disjoint sets of concepts among agents. These assumptions ultimately determine the coordinated channel C' ; this is mathematically captured by an IF classification \mathbf{S} with no types,

$typ(\mathbf{S}) = \emptyset$, the domain of discourse as its token set, $tok(\mathbf{S}) = D$, and empty classification relation.

The optimal coordinated IF channel that captures the semantic integration achieved by the agents is mathematically described by the universal property of the category-theoretical *colimit* (see, e.g., [8]) $C' = colim\{\mathbf{A}'_1 \leftarrow \mathbf{S} \rightarrow \mathbf{A}'_2\}$ of the diagram linking the IF subclassifications that model each agent’s participation through the assumptions of the scenario:



4.3. Partial Semantic Integration

The diagram above is a general model of the coordinated channel between two agents, and it faithfully captures the semantic integration between them, according to the Barwise-Seligman theory of information flow. Initially, when the agents have not yet coordinated themselves, the IF classifications modelling the agents’ participation have no types since none of them have been communicated yet, and the token set of the core of the coordinated channel is empty (as no tokens have been shared yet):

$$\begin{matrix} typ(\mathbf{A}'_i) = \emptyset & typ(\mathbf{C}') = \emptyset \\ tok(\mathbf{A}'_i) = tok(\mathbf{A}_i) & tok(\mathbf{C}') = \emptyset \end{matrix}$$

After A_1 told A_2 that $r_1 \models HO$ (i.e., “issues passports” is a responsibility of the Home Office) and A_2 told A_1 that $r_1 \models DoS$ (i.e. “issues passports” would be a responsibility of the Department of State), A_1 participates in the coordinated channel with type HO and A_2 participates in the coordinated channel with type DoS. Furthermore A_2 will have extended its token set with the shared token r_1 , which yields the coordinated channel of Figure 2.




Figure 2. Partially coordinated channel

Furthermore, after A_2 told A_1 that $s_2 \models DoS$ (i.e., “promotes government interests in the region” is a re-

sponsibility of the Department of State) and A_1 told A_2 that $s_2 \models \text{FCO}$ (i.e., “promotes government interests in the region” would be a responsibility of the Foreign and Commonwealth Office), new types participate in the meaning coordination, and new tokens are shared, yielding the newly coordinated channel of Figure 3.




Figure 3. Partially coordinated channel

At each stage a new coordinated channel arises. The distributed IF logic of the natural logic determined by the core of each new channel **captures the semantic integration achieved so far**. For instance, for this last coordinated channel the theory of the distributed IF logic $DLog_C(\text{Log}((C')))$ would include among its constraints:

$$\vdash \text{DoS} \quad \vdash \text{HO, FCO} \quad \text{HO, FCO} \vdash$$

4.4. Complete Semantic Integration

In the optimal limit case, all types would be eventually communicated and all tokens shared, which would yield a situation of complete semantic integration in which the IF classifications modelling the agents’ participation in the coordination would include each agent’s types and would have the domain of discourse as their token set:

$$\begin{aligned} \text{typ}(A_i) &= \text{typ}(A_i) & \text{typ}(C') &= \bigcup_i \text{typ}(A_i) \\ \text{tok}(A_i) &= D & \text{tok}(C') &= D \end{aligned}$$

This is an ideal scenario, in which agents would have exchanged their entire IF classification (all tokens, all types, and the entire classification relation). In our example (and restricting it only to types HO, FCO, DoJ, and DoS, and to tokens r_1, r_2, r_4, s_1, s_2 , and s_4) complete semantic integration would have been achieved with the coordinated channel shown in Figure 4. The IF theory of the distributed IF logic of this channel happens to be equivalent to that of the global ontology discussed above, although the core IF classification of the channel shows a different set of tokens.




Figure 4. Completely coordinated channel

Because in practice complete semantic integration will seldom be achieved (e.g., because it would be computationally too expensive) the ontology coordination process will usually yield only a partial semantic integration involving a fraction of communicated types and shared tokens. In these cases it is important to have a faithful formalisation of the resulting situation, which we believe is achieved with its modelling as a coordinated IF channel.

5. Concluding Discussion

Channel theory emphasises that, since information is carried by particular tokens, information flow crucially involves both types and tokens. Barwise and Seligman realised the fundamental duality between types and tokens, which is central to all channel-theoretic constructions. Thus, although meaning coordination is usually thought of as a process during which concepts of separate ontologies are being aligned at the type-level, the logical relationship between concepts arises when tokens are being connected by means of an IF channel. Knowing what these connections at the token-level are is therefore fundamental for determining the semantic integration of ontologies at the type-level.

In this paper, we have been formalising an meaning coordination approach in which token connection is the result of passing “responsibility assertions” between agents. But the general formalisation based on channel theory presented here provides a wide view about what we can consider to be a *token* and a *connection between tokens*. This allows for accommodating different understandings of semantics—depending on the particularities of the interoperability scenario—whilst retaining the core aspect that will allow coordination among agents: connections through their tokens. Schorlemmer showed in [9] how the type-token duality helps to pin down some of the reasons why ontologies appear to be insufficient in certain interoperability scenarios for which a common verified ontology is not enough for knowledge sharing, as pointed out

by Corrêa da Silva et al. [3]. Depending on the scenario being analysed, the role of tokens is taken either by instances, model-theoretic structures, or even proof-theoretic derivations.

An information-theoretic analysis of meaning coordination based on channel theory highlights the fact that a coordination process can hardly be absolute. On the contrary, not only is it relative to the respective ontologies being coordinated, but also

1. to the way ontologies are actually used in the context of specific application domains (what we have been calling the populated ontologies);
2. to the way ontologies are characterised as IF logics: the particular understanding of semantics of the interoperability scenario is relative to our choice of types and tokens and its classification relation; (this is closely related to what Farrugia calls the *logical setup*, and which he claims needs to be established first before any meaning negotiation between agents can start [5];)
3. to the way ontologies are linked together via connected tokens: as discussed in [9] reliable semantic integration is only guaranteed on connected tokens, which nicely includes into the framework the unavoidable imperfections of most meaning coordination processes, unless complete semantic integration is achieved.

It would be interesting, for instance, to explore the channel-theoretical notion of *induced IF logic* in the meaning coordination context. This logic characterises how an agent extends its own ontology with the understanding it has gained of other agents' ontologies *relative to the coordinated channel*. This logic is defined by moving the distributed IF logic of the coordinated channel to its restriction to one particular agent's IF classification. It turns out that the resulting induced IF logic is only sound and complete when the infomorphisms constituting the coordinated channel are surjective on tokens (see Definition 8). Such a particular case is when we achieve complete semantic integration, but it would be desirable to find conditions for meaning coordination processes that, without obtaining complete semantic integration, lead to coordinated channels for which sound and complete induced IF logics exist.

Acknowledgments. M. Schorlemmer is supported by a 'Ramón y Cajal' Research Fellowship from Spain's Ministry of Science and Technology. Y. Kalfoglou is supported by the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under

grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

References

- [1] S. Bailin and W. Truszkowski. Ontology negotiation between intelligent information agents. *Knowledge Engineering Review*, 17(1):7–19, 2002.
- [2] J. Barwise and J. Seligman. *Information Flow: The Logic of Distributed Systems*. Cambridge University Press, 1997.
- [3] F. Corrêa da Silva, W. Vasconcelos, D. Robertson, V. Brilhante, A. de Melo, M. Finger, and J. Agustí. On the insufficiency of ontologies: Problems in knowledge sharing and alternative solutions. *Knowledge-Based Systems*, 15(3):147–167, 2002.
- [4] F. Dretske. *Knowledge and the Flow of Information*. MIT Press, 1981.
- [5] J. Farrugia. Logical systems: Towards protocols for web-based meaning negotiation. In *Meaning Negotiation, Papers from the AAAI Workshop*, pages 56–59. The AAAI Press, 2002.
- [6] Y. Kalfoglou and M. Schorlemmer. IF-Map: An ontology-mapping method based on information-flow theory. In *Journal on Data Semantics I*, volume 2800 of *Lecture Notes in Computer Science*, pages 98–127. Springer, 2003.
- [7] Y. Kalfoglou and M. Schorlemmer. Formal support for representing and automating semantic interoperability. In *The Semantic Web: Research and Applications. ESWS 2004. Proceedings*, Lecture Notes in Computer Science, pages 45–60, Heraklion, Crete, Greece, May 2004. Springer.
- [8] B. Pierce. *Basic Category Theory for Computer Scientists*. MIT Press, 1991.
- [9] M. Schorlemmer. Duality in knowledge sharing. In *7th International Symposium on Artificial Intelligence and Mathematics*, Ft. Lauderdale, Florida, USA, 2002.
- [10] J. Wang and L. Gasser. Mutual online ontology alignment. In *OAS'02 Ontologies in Agent Systems, Proceedings of the AAMAS 2002 Workshop*, volume 66 of *CEUR Workshop Proceedings*, 2002.