A Categorical Framework for Translating between Conceptual Hierarchies

Rainer Osswald

Department of Computer Science, FernUniversität in Hagen, Germany

Abstract. Any classification of objects by attributes naturally determines a conceptual (or ontological) hierarchy consisting of the attribute sets that are closed under all implications holding in the classification. In this paper, we present a thorough categorical treatment of translations between implications over different sets of attributes and study in detail how such a translation is reflected on the side of the associated conceptual hierarchies.

1 Introduction

If the elements of a given domain of discourse are classified with respect to a certain set Σ of attributes then these data naturally give rise to a hierarchy of formal entities that can be interpreted both conceptually and ontologically. A well-known example of this sort of construction is provided by Formal Concept Analysis (FCA) [10], where classifications are called formal contexts and the resulting hierarchies are lattices of (formal) concepts. The formal concepts of FCA can be identified with those subsets of Σ that are closed under all conjunctive implicational statements holding in the given context or classification. This kind of systematic relationship between statements holding in a classification and conceptual hierarchies consisting of closed attribute sets can be generalized from purely conjunctive implications to those also involving disjunction, truth, and falsity, that is, to implications where arbitrary *positive* or *affirmative terms* may occur as premise and conclusion [9, 25]. The resulting hierarchies include all finite partial orders and have once been dubbed *information domains* [8].

The present paper is concerned with translations between different sets of attributes. We present a categorical framework for theories, i.e., sets of statements, which makes use of a very natural and explicit notion of translation, and study in detail how translations of theories are reflected on the side of the corresponding information domains.[1] In particular, we give equivalence criteria for theories with identical information domain and we present general construction schemes both for theories and information domains.

[1] Our exposition presumes some background knowledge of the standard definitions of category, lattice, order, and domain theory, all of which can be found in [2, 18], [6], and [1], respectively.

Overview. Section 2 introduces basic notions like interpretation, theory, and model, as well as the specialization relation given by an interpretation. We then define the information domain of a theory as the ordered universe of a universal model of that theory and review an alternative construction via the Lindenbaum algebra of positive terms determined by the theory, which is a bounded distributive lattice. Section 3 explores the relation between theories over different base vocabulary; in particular, we are interested in criteria for equivalence. To this end, we define an appropriate notion of translation between positive terms and study the effect such a translation has on the respective information domains. Categorically speaking, we study the information domain functor from the category of theories to that of (directed-complete) ordered sets. Section 4 is concerned with various ways of constructing theories and with the corresponding constructions of information domains. The most general construction we consider is the quasi-colimit of theories, which is a slight variant of the standard category-theoretic notion of colimit. The information domain functor is shown to take quasi-colimits of theories to limits of information domains. This gives us a characterization of information domains as well as possible ways to construct them.

Related Approaches. There is a significant overlap with the extensive work on *sequent structures* (*consequence systems*, *information systems*, *entailment relations*) [8, 29, 3, 7, 5], which are essentially normal forms of the theories considered in this paper. The focus of the present paper differs insofar as sequent structures are typically seen as suitable representations for domains, whereas for us a given theory is the primary object of interest. It seems that the category of theories introduced in this paper has not been studied in depth before. For instance, in [7], only primitive translations, i.e., translations of attributes by (primitive) attributes, are taken into account, whereas [5] makes use of the more general notion of an approximable relation. Categorical treatments of conjunctive theories and their relation to FCA can be found in [14, 13]. It is furthermore worth mentioning that [20, 19] take the Lindenbaum algebra of positive terms instead of the information domain as an appropriate representation of the conceptual hierarchy determined by a theory. As mentioned above, the Lindenbaum algebra canonically determines the information domain of a theory (and vice versa in the finite case).

2 Information Domains

2.1 Theories, Models, and Positive Terms

Let Σ be a set of (primitive) attributes. We introduce two special attributes V and Λ which respectively hold of everything and nothing in any universe of discourse. To build compound attributes, we employ the usual Boolean connectives

$\wedge$, $\vee$, and $\neg$. Let $B[\Sigma]$ be the term algebra of (compound) Boolean attributes inductively defined that way. The meaning of these compound attributes, when applied to classify the elements of a certain domain of discourse U, is the obvious one: Let $\models$ be a binary relation from U to Σ that expresses the *satisfaction* of (primitive) attributes by elements of U. Then $x \in U$ satisfies $\phi \wedge \psi$ iff x satisfies ϕ and ψ; similarly, x satisfies $\neg\phi$ iff x does not satisfy ϕ; etc. The satisfaction relation $\models$ can thus be inductively extended to a relation from U to $B[\Sigma]$. As usual, we write $\phi \rightarrow \psi$ and $\phi \leftrightarrow \psi$ for $\neg\phi \vee \psi$ and $\phi \rightarrow \psi \wedge \psi \rightarrow \phi$, respectively.

A natural way to reformulate the notions introduced so far within a standard logical setting is to regard attributes as *monadic predicates* (see also [25]). A satisfaction relation $\models$ from U to Σ is then essentially the same as a *(set valued) interpretation function* M from Σ to $\wp(U)$, with $M(p) = \{x \in U \mid x \models p\}$, since within (first-order) predicate logic, monadic predicates are interpreted by subsets of a universe U. Moreover, the interpretation function M can be inductively extended to a function $\hat{M}$ from $B[\Sigma]$ to $\wp(U)$ such that $\hat{M}(\phi) = \{x \in U \,|\, x \models \phi\}$. We refer to $\hat{M}(\phi)$ as the *extent* of ϕ.

In order to formulate statements that *hold* or are *true* with respect to a given interpretation (or classification), we need to quantify over the terms in $B[\Sigma]$. In this paper, only universal quantification is taken into account, which covers the approaches mentioned in the introduction. That is, we consider *universal statements* of the form $\forall x(\phi x)$, abbreviated by $\forall\phi$, with $\phi \in B[\Sigma]$, and take a *theory* over Σ to be a set of universal statements of this form. The following notions of truth and model are those of standard first-order logic: A statement $\forall\phi$ is *true* with respect to an interpretation if ϕ is satisfied by all elements of the universe. A *model* of a theory Γ is an interpretation with respect to which all statements of Γ are true.

A Boolean term $\phi \in B[\Sigma]$ is called *positive* or *affirmative* if it is free of $\neg$. Let $T[\Sigma]$ be the term algebra of positive terms over Σ. The special role of the positive terms reveals itself in connection with the *specialization relation* an interpretation defines on its universe:

Definition 1 (Specialization). *Given an interpretation of Σ with universe U and $x, y \in U$, then x* is specialized by *y, in symbols, $x \sqsubseteq y$, if y satisfies every element of Σ that is satisfied by x.*

The specialization relation $\sqsubseteq$ is reflexive and transitive, i.e. a preorder. By the definition of $\sqsubseteq$ and straightforward term induction over $T[\Sigma]$, it follows that $x \sqsubseteq y$ iff $\forall\phi \in T[\Sigma]\,(x \models \phi \rightarrow y \models \phi)$. That is, positive terms are *persistent* with respect to specialization, whereas Boolean terms in general are not (see e.g. [24]).

Given two theories Γ and Γ' over Σ, we say that Γ *entails* Γ', in symbols, $\Gamma \vdash \Gamma'$, if every model of Γ is also a model of Γ'. The theories Γ and Γ' are said to be *equivalent* if they entail each other. A theory over Σ has *conditional (or* *biconditional) form* if its statements are of the form $\forall(\phi \to \psi)$ (or $\forall(\phi \leftrightarrow \psi)$), with ϕ and ψ positive. The conditional form is *normal*, if ϕ is purely conjunctive (or V) and ψ is purely disjunctive (or Λ). It is a standard exercise in elementary logic to verify that every theory is equivalent to a theory in conditional (normal) form and to one in biconditional form. In the following, this fact will be frequently used to the effect that a theory is implicitly assumed to have conditional or biconditional normal form when appropriate. For convenience, let us introduce two binary term operators $\preceq$ and $\equiv$ such that $\phi \preceq \psi$ and $\phi \equiv \psi$ are $\forall(\phi \to \psi)$ and $\forall(\phi \leftrightarrow \psi)$, respectively.

Given an interpretation M of Σ with universe U, the set $\{\hat{M}(\phi) \mid \phi \in T[\Sigma]\}$ of extents of positive terms forms a *distributive lattice* with respect to $\cap$ and $\cup$, which is *bounded* by $\hat{M}(V) = U$ and $\hat{M}(\Lambda) = \varnothing$. More generally, we can consider interpretations and models in arbitrary bounded distributive lattices. We speak of such lattices briefly as *algebras*, taking them as algebras of type $\langle 2, 2, 0, 0\rangle$.

Definition 2 (Algebraic Interpretation/Model). *An* interpretation m *of* Σ *in an algebra* A*, or* A-valued interpretation*, is a function from* Σ *to* A*;* m *is an* A-valued model *of a theory* Γ *over* Σ *iff* $\hat{m}(\phi) = \hat{m}(\psi)$ *whenever* Γ *entails* $\phi \equiv \psi$*. Let* $\mathrm{Mod}(\Gamma, A)$ *(or* $\mathrm{Mod}_A(\Gamma)$*) be the set of* A*-valued models of* Γ*.*

2.2 Information Domains

There is a standard way to associate with each theory Γ over Σ a *canonical model*. Let the *canonical interpretation* of Σ in $\wp(\Sigma)$ take $p \in \Sigma$ to $\{X \subseteq \Sigma \mid p \in X\}$, that is, $X \vDash p$ iff $p \in X$. Now let $C(\Gamma)$ be the set of all $X \subseteq \Sigma$ which, under the canonical interpretation, satisfy ϕ for every statement $\forall\phi$ of Γ. The *canonical model* of Γ takes $p \in \Sigma$ to $\{X \in C(\Gamma) \mid p \in X\}$. The canonical model of Γ has the universal property that the statements true in it are precisely those entailed by Γ.[2]

Specialization on $C(\Gamma)$ is set inclusion and hence a partial order. We refer to the elements of $C(\Gamma)$ as the *consistently* Γ*-closed subsets* of Σ. Adapting the terminology of [8], we say that any ordered set order-isomorphic to $C(\Gamma)$ "is" or represents the *information domain* of Γ; notation: $D(\Gamma)$. It is not difficult to see that the information domain of a theory is *directed-complete*, i.e., closed with respect to suprema of (upwards) directed subsets.

Example 1. Let Γ be the theory over $\{a, b, c, d\}$ with statements $a \wedge b \preceq c \vee d$, $c \wedge d \preceq \Lambda$, $c \preceq a$, $d \preceq a \wedge b$. Then $C(\Gamma)$ is as depicted on the left of Figure 1.

The representation of the information domain by consistently closed sets is especially useful for practical purposes because of its concreteness. For theoret-

[2] In [9], $C(\Gamma)$ is referred to as the *free extent*; in the terminology of [4], the canonical model is essentially the same as the *generated classification*; see [25] for a comparison of terminologies.

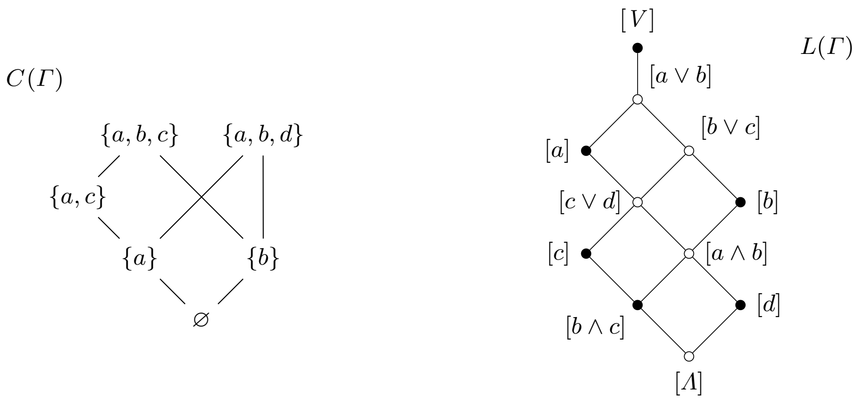

Fig. 1. Information domain and Lindenbaum algebra of Γ

ical purposes, on the other hand, a more abstract representation of the information domain as the prime spectrum of the Lindenbaum algebra is often more appropriate. The idea that underlies the Lindenbaum construction is to abstract away from syntactical differences between positive terms that are equivalent with respect to a given theory:

Definition 3 (Lindenbaum Algebra). *The* Lindenbaum algebra $L(\Gamma)$ *of a theory* Γ *over* Σ *is the quotient* $T[\Sigma]/{\simeq_\Gamma}$*, with* $\phi \simeq_\Gamma \psi$ *iff* $\Gamma \vdash \phi \equiv \psi$.

Clearly the Lindenbaum algebra of positive terms is an algebra in the sense introduced at the close of Section 2.1, i.e., a bounded distributive lattice, with $0 = [\Lambda]$, $[\phi] \wedge [\psi] = [\phi \wedge \psi]$, etc. If Γ is the theory of Example 1 then $L(\Gamma)$ is as depicted on the right of Figure 1. (The shaded circles in the diagram correspond to the $\vee$-*irreducible* elements of $L(\Gamma)$, which stand in an order-reversing one-to-one correspondence to the elements of $C(\Gamma)$.) Notice that every algebra A is isomorphic to the Lindenbaum algebra of a theory over (the carrier set of) A. For let $Th(A)$ be the theory over A consisting of all statements $\phi \equiv \psi$ such that $\hat{id}_A(\phi) = \hat{id}_A(\psi)$. Then $L(Th(A)) \simeq A$.

Let $\mathbb{2}$ be the algebra $\{0, 1\}$. There is a one-to-one correspondence between $C(\Gamma)$ and $\mathrm{Mod}_{\mathbb{2}}(\Gamma)$ (cf. Definition 2), where $X \in C(\Gamma)$ corresponds to the characteristic function χ_X of X. Consequently:

Proposition 1. $\mathrm{Mod}_{\mathbb{2}}(\Gamma)$ *represents the information domain of* Γ.

The Lindenbaum construction provides a *universal* model in the sense that the function m_Γ from Σ to $L(\Gamma)$ that takes p to $[p]_{\equiv_\Gamma}$ has the following universal property:

Proposition 2. *Every model of* Γ *in an algebra* A *factors uniquely through* m_Γ *by a homomorphism from* $L(\Gamma)$ *to* A. *Hence* $\mathrm{Mod}(\Gamma, A) \simeq \mathrm{Hom}(L(\Gamma), A)$.

Definition 4 (Spectrum). *The* spectrum $P(A)$ *of an algebra* A *is the set of all prime filters of* A *ordered by set inclusion.*

Since $h \in \mathrm{Hom}(A, \mathbb{2})$ iff $h^{-1}(1) \in P(A)$, and $F \in P(A)$ iff $\chi_F \in \mathrm{Hom}(A, \mathbb{2})$:

Proposition 3. $P(A)$ *is order-isomorphic to* $\mathrm{Hom}(A, \mathbb{2})$ *(ordered pointwise).*

Together with Propositions 1 and 2, this gives us the promised result that the information domain of a theory Γ can be represented by the spectrum of its Lindenbaum algebra $L(\Gamma)$.

3 Translations and Equivalences

3.1 Translations of Theories

Let Γ and Γ' be theories over Σ and Σ', respectively. Roughly speaking, a translation of theories is a truth-preserving translation of positive terms. That is, if a statement is true in all models of Γ, its translation is required to be true in all models of Γ':

Definition 5 (Theory Translation). *A* translation of theories *from* $\langle \Sigma, \Gamma \rangle$ *to* $\langle \Sigma', \Gamma' \rangle$ *is a function* μ *from* Σ *to* $T[\Sigma']$ *such that* $\Gamma' \vdash \hat{\mu}(\Gamma)$. *The translation* μ *is* primitive *if* $\mu(\Sigma) \subseteq \Sigma'$.

Example 2 (Theory extensions). Suppose $\Sigma \subseteq \Sigma'$ and $\Gamma \subseteq \Gamma'$. The inclusion function ε from Σ to Σ' is a primitive translation of theories from Γ to Γ', called an *extension*. We then also say that Γ' is an *extension of* Γ.

Let the *composite* of two theory translations μ from Γ to Γ' and μ' from Γ' to Γ'' be the composite function $\widehat{\mu'} \circ \mu$ from Σ to $T[\Sigma'']$, which is easily seen to be a theory translation from Γ to Γ'', and let the *identity translation* $\imath_\Gamma$ of Γ be the canonical inclusion function of Σ into $T[\Sigma]$. One straightforwardly verifies that these data define a *category of theories* **Th**.

The notion of isomorphism in **Th** is too restrictive to adequately characterize the equivalence of theories. To give a simple example, let Γ be the empty theory over $\Sigma = \{a\}$ and let Γ' be the theory $\{b \equiv c\}$ over $\Sigma' = \{b, c\}$. That is, within $\langle \Sigma', \Gamma' \rangle$, we have precisely two primitive terms or attributes, which are synonymous in the sense that they have the same extent, and there are no further constraints. Within $\langle \Sigma, \Gamma \rangle$, on the other hand, we have only one primitive attribute and no constraints. Clearly, these two theories should count as equivalent under any sensible notion of equivalence. But if $\hat{\mu}(\mu'(b)) = b$ then $\mu'(b) = a$, because $\hat{\mu}$ takes non-primitive terms to non-primitive terms. Similarly $\mu'(c) = a$ and thus $\hat{\mu}(\mu'(c)) = b$. So there is no isomorphism between these theories in **Th**.

In order to get an appropriate definition of equivalence between theories we need to relax the notion of an inverse morphism to that of a *quasi-inverse*, which means that the composite morphisms are not required to be identical but only to be equivalent to identity in the following sense:

Definition 6 (Equivalence of Translations). *Let μ and ν be theory translations from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$. Then μ is* equivalent *to ν, notation: $\mu \sim \nu$, iff $\forall p \in \Sigma \, (\Gamma' \vdash \mu(p) \equiv \nu(p))$.*

Definition 7 (Quasi-Inverse). *Let μ be a theory translation from Γ to Γ'. A translation ν from Γ' to Γ is* quasi-inverse *to μ iff $\nu \circ \mu \sim \imath_\Gamma$ and $\mu \circ \nu \sim \imath_{\Gamma'}$.*

We say that Γ and Γ' are *equivalent*, notation: $\Gamma \sim \Gamma'$, if there is a translation from Γ and Γ' that has a quasi-inverse. The existence of a quasi-inverse turns out be equivalent to the following two conditions:

Definition 8 (Conservative/Essentially Surjective). *A translation μ from Γ to Γ' is* conservative *iff, for all statements α over Σ, $\Gamma' \vdash \hat{\mu}(\alpha)$ only if $\Gamma \vdash \alpha$; μ is* essentially surjective *iff $\forall p \in \Sigma' \, \exists \phi \in T[\Sigma] \, (\Gamma' \vdash p \equiv \hat{\mu}(\phi))$.*

Proposition 4. *A theory translation μ has a quasi-inverse iff μ is conservative and essentially surjective.*

Proof. Suppose ν is quasi-inverse to μ. Then $\Gamma' \vdash \phi' \equiv \hat{\mu}(\hat{\nu}(\phi'))$, for every $\phi' \in T[\Sigma']$, which shows that μ is essentially surjective. If $\Gamma' \vdash \hat{\mu}(\phi) \equiv \hat{\mu}(\psi)$ then $\Gamma \vdash \hat{\nu}(\hat{\mu}(\phi)) \equiv \hat{\nu}(\hat{\mu}(\psi))$, because ν is a morphism. Moreover, Γ entails $\phi \equiv \hat{\nu}(\hat{\mu}(\phi))$ and $\psi \equiv \hat{\nu}(\hat{\mu}(\psi))$. Hence $\Gamma \vdash \phi \equiv \psi$; so, μ is conservative. Conversely, assume μ is conservative and essentially surjective. Then, by the axiom of choice, one can choose $\nu(p') \in T[\Sigma]$, for every $p' \in \Sigma'$, such that $\Gamma' \vdash p' \equiv \hat{\mu}(\nu(p'))$. Thus $\mu \circ \nu \sim \imath_{\Gamma'}$. In particular, it holds that $\Gamma' \vdash \mu(p) \equiv \hat{\mu}(\hat{\nu}(\mu(p)))$, for every $p \in \Sigma$. Hence $\Gamma \vdash p \equiv \hat{\nu}(\mu(p))$; that is, $\nu \circ \mu \sim \imath_\Gamma$.

Example 3. Let Σ be $\{a_0, a_1, \dots a_k\} \cup \{b_0, b_1, \dots b_k\}$, with k finite, and let Γ be the theory over Σ with statements $a_n \equiv a_{n+1} \vee b_{n+1}$ and $a_n \wedge b_n \equiv \Lambda$ $(0 \leqslant n < k)$. Then $C(\Gamma) = \{\varnothing, \{a_0, a_1, \dots, a_k\}\} \cup \{\{a_0, a_1, \dots, a_{n-1}, b_n\} \mid n \leqslant k\}$. So the information domain of Γ is flat and thus order-isomorphic to the information domain of the theory $\Gamma' = \{c_m \wedge c_n \equiv \Lambda \mid m \neq n\}$ over $\Sigma' = \{c_0, c_1, \dots, c_{k+1}\}$. Now consider the function μ from Σ to $T[\Sigma']$ with $\mu(a_n) = c_{n+1} \vee \dots \vee c_{k+1}$ and $\mu(b_n) = c_n$ $(0 \leqslant n \leqslant k)$. Then $\Gamma' \vdash \hat{\mu}(\Gamma)$, that is, μ is a theory translation from Γ to Γ'. Moreover, it is not difficult to see that the function ν from Σ' to $T[\Sigma]$, with $\nu(c_{k+1}) = a_k$ and $\nu(c_n) = b_n$ for every $n \leqslant k$, is a translation from Γ' to Γ which is quasi-inverse to μ. The case $k = 1$ is illustrated by Figure 2 in terms of Lindenbaum algebras and information domains (decorated with extents of primitives).

In the foregoing example, the equivalence translation μ from Γ to Γ' induces an isomorphism from $L(\Gamma)$ to $L(\Gamma')$. We shall see in the next section that this holds in general. Moreover, finite theories turn out to be equivalent whenever they have order-isomorphic information domains.

Remark 1 (Theory Morphisms in Institution Theory). It is tempting to subsume our category **Th** under the framework of *Institution Theory* (see [12, 11] for background). The sets Σ of primitives can be taken as the *signatures* of the underlying institution, the *sentences* associated with Σ are our statements

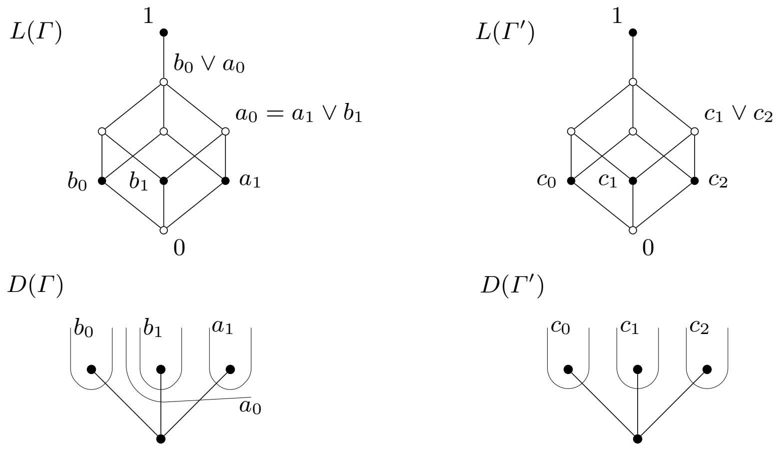

Fig. 2. The case $k = 1$ for the theory of Example 3

over Σ. A *theory* is then again a pair $\langle \Sigma, \Gamma \rangle$, where Γ is a set of statements over Σ. In order to get theory morphisms, we still need to specify the signature morphisms of the institution. But here a problem arises insofar as the standard application of institution theory to logic takes signature morphisms simply as functions from signature to signature. This would mean to restrict ourselves to *primitive* translations, which is not adequate as it has been argued above. A possible way out would be to define a signature morphism from Σ to Σ' as a function from Σ to $T[\Sigma']$. On the other hand, it seems to violate the basic division of labor in Institution Theory if term forming operations are already involved at the level of signature morphisms.

3.2 Functors and Equivalences

Our primary goal is to study the *information domain functor* that takes theories to their information domains. Recall from Section 2.2 that the information domain can be identified with the spectrum of the Lindenbaum algebra, which allows us to factor the information domain functor into the *Lindenbaum functor* and the *spectrum functor*. The gain is that we can employ results from universal algebra and lattice theory.

The Lindenbaum Functor. Let μ be a theory translation from Γ to Γ'. Then $\hat{m}_{\Gamma'} \circ \mu$ is a model of Γ in $L(\Gamma')$. Since $\mathrm{Mod}(\Gamma, L(\Gamma')) \simeq \mathrm{Hom}(L(\Gamma), L(\Gamma'))$, by Proposition 2, μ gives rise to a homomorphism $L(\mu)$ from $L(\Gamma)$ to $L(\Gamma')$, and this assignment is functorial.

Proposition 5. *Two theory translations μ and ν from Γ to Γ' are equivalent iff $L(\mu) = L(\nu)$. In particular, $L(\Gamma) \simeq L(\Gamma')$ iff $\Gamma \sim \Gamma'$.*

Proof. By definition, $\mu \sim \nu$ iff, for every $\phi \in T[\Sigma]$, $\hat{\mu}(\phi) \cong_{\Gamma'} \hat{\nu}(\phi)$, that is, iff $L(\mu) \circ m_\Gamma = L(\nu) \circ m_\Gamma$. By Proposition 2, the latter condition implies that $L(\mu) = L(\nu)$.

Proposition 6. *A theory translation μ is conservative iff $L(\mu)$ is one-to-one; μ is essentially surjective iff $L(\mu)$ is onto.*

Proof. (i) $L(\mu)$ is one-to-one iff $\cong_\Gamma$ is the congruence kernel of $\hat{m}_{\Gamma'} \circ \hat{\mu}$ iff $\forall \phi, \psi \in T[\Sigma](\hat{\mu}(\phi) \cong_{\Gamma'} \hat{\mu}(\psi) \to \phi \cong_\Gamma \psi)$ iff μ is conservative. (ii) $L(\mu)$ is onto iff $\widehat{m}_{\Gamma'} \circ \hat{\mu}$ is onto iff $\forall \phi' \in T[\Sigma'] \, \exists \phi \in T[\Sigma](\phi' \cong_{\Gamma'} \hat{\mu}(\phi))$ iff μ is essentially surjective.

Let $\mathbf{L}$ be the category of bounded distributive lattices viewed as algebras of type $\langle 2, 2, 0, 0 \rangle$.

Proposition 7. *The functor L from $\mathbf{Th}$ to $\mathbf{L}$ is full and every object A of $\mathbf{L}$ is of the form $L(\Gamma)$ for some object Γ of $\mathbf{Th}$.*

Proof. Let h be a homomorphism from $L(\Gamma)$ to $L(\Gamma')$. For every $p \in \Sigma$ choose $\mu(p) \in T[\Sigma']$ such that $[\mu(p)]_{\cong_{\Gamma'}} = h([p]_{\cong_\Gamma})$. One easily shows that $\Gamma' \vdash \hat{\mu}(\Gamma)$. Hence μ is a translation with $h = \Gamma(\mu)$; so L is full. Moreover, $A \simeq L(Th(A))$.

By Proposition 5, $\sim$ is a *congruence relation* with respect to composition. Hence we can switch to the *quotient category* $\mathbf{Th}/\!\sim$ of $\mathbf{Th}$ by $\sim$, which has the same objects as $\mathbf{Th}$ whereas its morphisms are the equivalence classes of morphisms of $\mathbf{Th}$ modulo $\sim$. The functor L factors uniquely by the quotient functor from $\mathbf{Th}$ to $\mathbf{Th}/\!\sim$ and a faithful functor from $\mathbf{Th}/\!\sim$ to $\mathbf{L}$ (see [18, Sect. II.8]). Together with Proposition 7, we have:

Theorem 1. *The categories $\mathbf{Th}/\!\sim$ and $\mathbf{L}$ are equivalent.*

Let $\mathbf{Th_p}$ be the subcategory of $\mathbf{Th}$ whose objects are those of $\mathbf{Th}$ and whose morphisms are the *primitive* theory translations (cf. Definition 5). $\mathbf{Th_p}$ corresponds to the standard category of *presentations by generators and relations*. Notice that Th (see Section 2.2) can be naturally extended to a functor from $\mathbf{L}$ to $\mathbf{Th_p}$. The following fact is folklore (e.g. [20, pp. 182f]):

Proposition 8. *The functor L from $\mathbf{Th_p}$ to $\mathbf{L}$ is left adjoint to Th.*

Remark 2 (Kleisli Construction). Another method to define an appropriate notion of theory morphism makes use of the so-called *Kleisli construction*:[3] The adjunction $\langle L, Th \rangle$ from $\mathbf{Th_p}$ to $\mathbf{L}$ gives rise to the *monad* $T = L \circ Th$ on the category $\mathbf{Th_p}$, which in turn has the associated *Kleisli category* $\mathbf{K}(T)$ whose objects are those of $\mathbf{Th_p}$ and whose morphisms from Γ to Γ' are the $\mathbf{Th_p}$-morphisms from Γ to $Th(L(\Gamma'))$. Since there is a one-to-one correspondence between morphisms in $\mathbf{K}(T)$ and equivalence classes of morphisms in $\mathbf{Th}$, it follows that the categories $\mathbf{K}(T)$ and $\mathbf{Th}/\!\sim$ are equivalent. In particular, the *Kleisli comparison functor* from $\mathbf{K}(T)$ to $\mathbf{L}$ is an equivalence of categories (cf. [18, p. 144]).

[3] See [15, p. 130f] for a similar use of the Kleisli category in the context of categorical logic.

The Spectrum Functor. Every algebra homomorphism h from A to B gives rise to a function $P(h)$ from $P(B)$ to $P(A)$ such that $P(h)(F) = h^{-1}(F)$. The function $P(h)$ is *Scott-continuous*, i.e., $P(h)$ is order-preserving and preserves suprema of directed sets. Thus P is a contravariant functor from $\mathbf{L}$ to the category $\mathbf{Dcpo}$ of directed-complete ordered sets (dcpos) and Scott-continuous functions. Notice that P is naturally isomorphic to the contravariant Hom-functor Hom_2; cf. Proposition 3. The following proposition reformulates a well-know result from the theory of distributive lattices (cf. [6, p. 265]).

Proposition 9. *A homomorphism h of algebras is one-to-one iff $P(h)$ is onto; h is onto iff $P(h)$ is an order embedding.*

The Information Domain Functor. By Proposition 1, the information domain of a theory Γ is represented by $\mathrm{Mod}_2(\Gamma)$ (ordered pointwise). For each translation μ from Γ to Γ', let $\mathrm{Mod}_2(\mu)$ be the function from $\mathrm{Mod}_2(\Gamma')$ to $\mathrm{Mod}_2(\Gamma)$ that takes m to $\widehat{m} \circ \mu$. We define "the" *information domain functor D* to be any contravariant functor from $\mathbf{Th}$ to $\mathbf{Dcpo}$ that is naturally isomorphic to Mod_2. In particular, $\mathrm{Hom}_2 \circ L$ and $P \circ L$ are information domain functors. Moreover, C can be extended to an information domain functor as well. Its effect on theory translations can be described as follows:

Proposition 10. *Let μ be a morphism of theories from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$ and suppose $Y \in C(\Gamma')$. Then $C(\mu)(Y) = \{p \in \Sigma \mid Y \vDash \mu(p)\}$. In case μ is an extension of theories then $C(\mu)(Y) = Y \cap \Sigma$.*

According to Birkhoff's representation theorem for finite distributive lattices ([6]), the functor P is a dual equivalence between the category of finite algebras and the category of finite ordered sets. Hence, by Theorem 1:

Proposition 11. *The functor D induces a dual equivalence between the quotient category of finite theories modulo $\sim$ and the category of finite ordered sets.*

In other words, the information domain of a finite theory (i.e., a theory with finite Σ) represents the theory up to equivalence. For infinite Σ, however, this is not necessarily the case:

Example 4. Consider the theory $\langle \Sigma, \Gamma \rangle$ with $\Sigma = \{a_0, a_1, \ldots\} \cup \{b_0, b_1, \ldots\}$ and $\Gamma = \{a_n \wedge b_n \equiv \Lambda,\ a_n \equiv a_{n+1} \vee b_{n+1} \mid n \geqslant 0\}$. Its information domain is shown on the left of Figure 3, with extents of primitives added. Since $D(\Gamma)$ is flat, it is isomorphic to the information domain of the theory $\langle \Sigma', \Gamma' \rangle$, with $\Sigma' = \{c_0, c_1, \ldots\} \cup \{c_\omega\}$ and $\Gamma' = \{c_m \wedge c_n \equiv \Lambda \mid m \neq n\}$; see Figure 3. However, it can be shown that $L(\Gamma)$ has a non-principal prime filter whereas all prime filters of $L(\Gamma')$ are principal. So $L(\Gamma)$ is not isomorphic to $L(\Gamma')$ and thus Γ is not equivalent to Γ', by Proposition 5. Compare this result with Example 3.

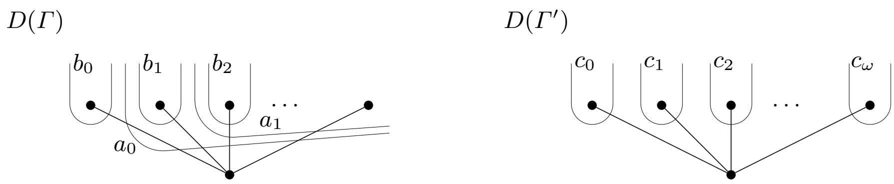

Fig. 3. Non-equivalent theories with isomorphic information domains

Remark 3. The dual equivalence of Proposition 11 can be generalized to the category of *finitistic theories* ([23, 22]) and the category of coherent algebraic domains with continuous functions restricted appropriately; see also [28, 1].

We conclude this section by combining Propositions 6 and 9:

Proposition 12. *A translation μ of theories is conservative iff $D(\mu)$ is onto; μ is essentially surjective iff $D(\mu)$ is an order embedding.*

4 Constructions

4.1 Quasi-Colimits

Many universal constructions in category theory can be seen as limits or colimits. Colimits, for instance, cover coproducts, pushouts, and inductive limits, besides others. Since in **Th**, we are more interested in characterizing theories *up to equivalence* than up to isomorphism, the notion of a *quasi-colimit* proves more fruitful than that of a colimit.

Let G be directed graph with vertex set I, edge set E, and two functions s and t from E to I, where $s(e)$ and $t(e)$ are respectively source and target of $e \in E$. A *diagram* $\mathcal{D}$ of shape G in a category **C** is a pair consisting of a family $\langle D_i \rangle_{i \in I}$ of **C**-objects and a family $\langle f_e \rangle_{e \in E}$ of **C**-morphisms such that f_e is a morphism from $D_{s(e)}$ to $D_{t(e)}$.

Definition 9 (Quasi-Cocone/-Colimit). *A* quasi-cocone $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ *of a diagram $\langle \langle \Gamma_i \rangle_{i \in I}, \langle \mu_e \rangle_{e \in E} \rangle$ of theories is a theory Γ and a family of translations μ_i from Γ_i to Γ such that $\mu_{s(e)} \sim \mu_{t(e)} \circ \mu_e$. A quasi-cocone $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ is a* quasi-colimit *if for every other quasi-cocone $\langle \Gamma', \langle \mu'_i \rangle_i \rangle$ of the diagram there is a translation μ from Γ to Γ', which is unique up to equivalence, such that $\mu'_i \sim \mu \circ \mu_i$.*

Quasi-Colimits can be shown to exist for all diagrams in **Th** by a fairly standard construction. In the following, let ι_i be the canonical injection of Σ_i into the disjoint union $\biguplus_i \Sigma_i$ of a family $\langle \Sigma_i \rangle_{i \in I}$ of sets.

Proposition 13. *If $\langle \langle \langle \Sigma_i, \Gamma_i \rangle \rangle_{i \in I}, \langle \mu_e \rangle_{e \in E} \rangle$ is a diagram of theories, the theory*

$$\bigcup_i \hat{\iota}_i(\Gamma_i) \cup \{\iota_{s(e)}(p) \equiv \hat{\iota}_{t(e)}(\mu_e(p)) \mid e \in E \wedge p \in \Sigma_{s(e)}\}$$

over $\biguplus_i \Sigma_i$, together with the family $\langle \iota_i \rangle_i$ of injections, is a quasi-colimit of the diagram.

Proof. Let $\langle \Sigma, \Gamma \rangle$ be the theory defined in the proposition. Clearly $\langle \Gamma, \langle \iota_i \rangle_i \rangle$ is a quasi-cocone since, by definition, $\iota_{s(e)} \sim \iota_{t(e)} \circ \mu_e$. Suppose $\langle \langle \Sigma', \Gamma' \rangle, \langle \mu_i \rangle_i \rangle$ is a quasi-cocone of the given diagram, i.e. $\mu_{s(e)} \sim \mu_{t(e)} \circ \mu_e$ for all $e \in E$. We claim that there is a translation μ from Γ to Γ' with $\mu_i = \mu \circ \iota_i$. Since every element of $\biguplus_i \Sigma_i$ is of the form $\iota_i(p)$, where i and $p \in \Sigma_i$ are uniquely determined, the only function μ from $\biguplus_i \Sigma_i$ to $T[\Sigma]$ satisfying the desired condition takes $\iota_i(p)$ to $\mu_i(p)$. Then μ is a translation, since Γ' entails $\hat{\mu}_i(\alpha)$ and thus $\hat{\mu}(\hat{\iota}_i(\alpha))$, for all $\alpha \in \Gamma_i$. This leaves us to check that Γ' entails $\hat{\mu}(\iota_{s(e)}(p)) \equiv \hat{\mu}(\hat{\iota}_{t(e)}(\mu_e(p)))$, i.e. $\mu_{s(e)}(p) \equiv \hat{\mu}_{t(e)}(\mu_e(p))$, for every $p \in \Sigma_{s(e)}$. But this is just the assumption that $\mu_{s(e)} \sim \mu_{t(e)} \circ \mu_e$; so μ is a translation from Γ to Γ'. It remains to show that if ν is another translation from Γ to Γ', with $\mu_i \sim \nu \circ \iota_i$, then $\mu \sim \nu$. Since $\mu \circ \iota_i = \mu_i \sim \nu \circ \iota_i$, it follows that $\Gamma' \vdash \hat{\mu}(\iota_i(p)) \equiv \hat{\nu}(\iota_i(p))$ for all $p \in \Sigma_i$. Hence $\mu \sim \nu$, by definition.

Example 5 (Coproducts). Let $\langle \langle \Sigma_i, \Gamma_i \rangle \rangle_{i \in I}$ be a family of theories and let $\biguplus_i \Gamma_i$ be $\langle \biguplus_i \Sigma_i, \bigcup_i \hat{\iota}_i(\Gamma_i) \rangle$. Then $\langle \biguplus_i \Gamma_i, \langle \iota_i \rangle_i \rangle$ is a coproduct of $\langle \langle \Sigma_i, \Gamma_i \rangle \rangle_{i \in I}$ in **Th**.

Since, by definition, the quotient functor from **Th** to **Th**/$\sim$ takes quasi-colimits to colimits, we have by Theorem 1:

Proposition 14. *The functor L takes quasi-colimits in* **Th** *to colimits in* **L**.

By Proposition 8, the restriction of L to a functor from $\mathbf{Th_p}$ to **L** has a right adjoint and hence preserves colimits ([18, Sect. V.5]). Colimits need not exist for all diagrams in $\mathbf{Th_p}$ (because $\mathbf{Th_p}$ has not enough morphisms). If, however, the colimit of a diagram in $\mathbf{Th_p}$ exists then it is also a quasi-colimit of that diagram in **Th**.

Proposition 15. *Colimits in* $\mathbf{Th_p}$ *are quasi-colimits in* **Th**.

Proof. Suppose a diagram $\mathcal{D}$ in $\mathbf{Th_p}$ has a colimit $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ in $\mathbf{Th_p}$. Then $\langle L(\Gamma), \langle L(\mu_i) \rangle_i \rangle$ is a colimit of the diagram $L(\mathcal{D})$ in **L**. By Proposition 13, $\mathcal{D}$ has a quasi-colimit $\langle \Gamma', \langle \mu'_i \rangle_i \rangle$ in **Th**, which, by Proposition 14, is taken to another colimit of $L(\mathcal{D})$. Hence $L(\Gamma) \simeq L(\Gamma')$ and thus $\Gamma \sim \Gamma'$, by Proposition 5. So $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ is a quasi-colimit of $\mathcal{D}$ in **Th**.

4.2 The Information Domain in the Limit

We now show that the information domain functor takes quasi-colimits in **Th** to limits in **Dcpo**. In other words, the information domain of the quasi-colimit of a diagram of theories is the limit of the corresponding diagram of information domains. Limits in **Dcpo** can be constructed as canonical limits in **Set**, i.e., as subsets of Cartesian products, with elements ordered coordinatewise (cf. [1, p. 45]). The following lemma gives us a criterion in which cases a limit of a diagram of dcpos (under the forgetful functor) in **Set** is actually a limit in **Dcpo**.[4]

Lemma 1. *Suppose $\langle D, \langle f_i \rangle_{i \in I} \rangle$ is a cone of a diagram $\mathcal{D}$ in* **Dcpo** *which is taken to a limit cone of $\mathcal{D}$ in* **Set** *by the forgetful functor. Then the cone is a limit of $\mathcal{D}$ in* **Dcpo** *iff, for every two $x, y \in D$, it holds that $x \leqslant y$ whenever $f_i(x) \leqslant f_i(y)$ for all $i \in I$.*

[4] The argument is essentially the same as that used in [16, p. 249].

Proof. Let $\langle D', \langle p_i \rangle_i \rangle$ be the canonical limit cone in **Dcpo** of $\mathcal{D}$. Then there is a unique Scott-continuous function f from D to D' such that $f_i = p_i \circ f$. By assumption, f is one-to-one because $\langle D', \langle p_i \rangle_i \rangle$, under the forgetful functor, is a limit cone of $\mathcal{D}$ in **Set**. Since f is Scott-continuous and one-to-one, it is enough to check that f is an order embedding to make sure that f is an isomorphism in **Dcpo**. Suppose $f(x) \leqslant f(y)$, for $x, y \in D$. Then $f_i(x) = p_i(f(x)) \leqslant p_i(f(y)) = f_i(y)$ for every i, since f and p_i are order preserving. Hence $x \leqslant y$, by assumption.

Theorem 2. *The functor* $\mathrm{Hom}_{\mathbb{2}}$ *takes colimits in* **L** *to limits in* **Dcpo**.

Proof. Let $\langle A, \langle h_i \rangle_i \rangle$ be a colimit of a diagram in **L**. $\mathrm{Hom}_{\mathbb{2}}$, as a contravariant Hom-functor, takes colimits in **L** to limits in **Set**. Suppose $v, w \in \mathrm{Hom}_{\mathbb{2}}(A)$ with $\mathrm{Hom}_{\mathbb{2}}(h_i)(v) \leqslant \mathrm{Hom}_{\mathbb{2}}(h_i)(w)$ for all i. Then $v(h_i(a)) \leqslant w(h_i(a))$ for all i and $a \in A_i$. Since A is generated by $\bigcup\{h_i(A_i) \mid i \in I\}$, it follows inductively that $v(a) \leqslant w(a)$ for every $a \in A$; so $v \leqslant w$. Now apply Lemma 1.

Corollary 1. *The functor D takes quasi-colimits in* **Th** *to limits in* **Dcpo**.

Inductive Limits of Extensions. Let I be a directed ordered set and suppose $\langle\langle \Sigma_i, \Gamma_i \rangle\rangle_{i \in I}$ is an inductive system of theory extensions over I with extensions ε_{ij}, for $i \leqslant j$. Then the theory $\bigcup_i \Gamma_i$ over $\bigcup_i \Sigma_i$ together with the extensions ε_i from Γ_i to $\bigcup_i \Gamma_i$ is a quasi-cocone of the inductive system.

Proposition 16. *Let $\langle \Gamma_i \rangle_i$ be an inductive system of theory extensions. Then $\langle \bigcup_i \Gamma_i, \langle \varepsilon_i \rangle_i \rangle$ is an inductive quasi-limit of that system.*

Proof. According to Proposition 15, it suffices to show that $\langle \bigcup_i \Gamma_i, \langle \varepsilon_i \rangle_i \rangle$ is an inductive limit of $\langle \Gamma_i \rangle_i$ in $\mathbf{Th_p}$. Suppose there are primitive translations μ_i from $\langle \Sigma_i, \Gamma_i \rangle$ to a theory $\langle \Sigma, \Gamma \rangle$ such that $\mu_i = \mu_j \circ \varepsilon_{ij}$ for $i \leqslant j$. Let μ be the function from $\bigcup_i \Sigma_i$ to Σ that takes $p \in \Sigma_i$ to $\mu_i(p)$; μ is well defined because if $p \in \Sigma_i \cap \Sigma_j$, there is a k with $i, j \leqslant k$ and thus $p \in \Sigma_k$; so $\mu_i(p) = \mu_k(p) = \mu_j(p)$. Notice that μ is the only function satisfying $\mu_i = \mu \circ \varepsilon_i$ for all i. It remains to check that μ is a translation, i.e. that $\Gamma \vdash \hat{\mu}(\alpha)$ for every $\alpha \in \bigcup_i \Gamma_i$. But if $\alpha \in \Gamma_i$ then $\hat{\mu}(\alpha) = \hat{\mu}_i(\alpha)$, and μ_i is a translation from Γ_i to Γ.

This result has the following straightforward but useful application. Suppose $\langle \Sigma, \Gamma \rangle$ is a theory. For $S \subseteq \Sigma$ let $\Gamma|_S$ be the set of all statements of Γ whose primitives belong to S. Let $\mathcal{F}$ be the directed set of finite subsets of Σ. Then the family $\langle\langle F, \Gamma|_F \rangle\rangle_{F \in \mathcal{F}}$ together with the extensions from $\Gamma|_F$ to $\Gamma|_{F'}$, whenever $F \subseteq F'$, is an inductive system, whose inductive limit is $\langle \Sigma, \Gamma \rangle$, by Proposition 16. Therefore:

Corollary 2. *Every theory is an inductive quasi-limit of finite theories.*

Suppose Σ is countable, i.e. $\Sigma = \{p_0, p_1, p_2, \ldots\}$. Let Σ_i be $\{p_0, p_1, \ldots, p_i\}$ and Γ_i be $\Gamma|_{\Sigma_i}$. Then $\langle \Sigma, \Gamma \rangle$ is the inductive limit of the inductive system $\langle\langle \Sigma_i, \Gamma_i \rangle\rangle_{i \in \omega}$, with extensions from Γ_i to Γ_j for all $i \leqslant j$. Hence:

Corollary 3. *Every theory over a countable set of primitives is an inductive quasi-limit of a sequence of finite theories.*

Another consequence, in combination with Proposition 11, is Speed's [27] characterization of spectra as *profinite* ordered sets, where profinite means to be the projective limit of a projective system of *finite* ordered sets:

Corollary 4. *Every information domain is profinite and every profinite ordered set is an information domain.*

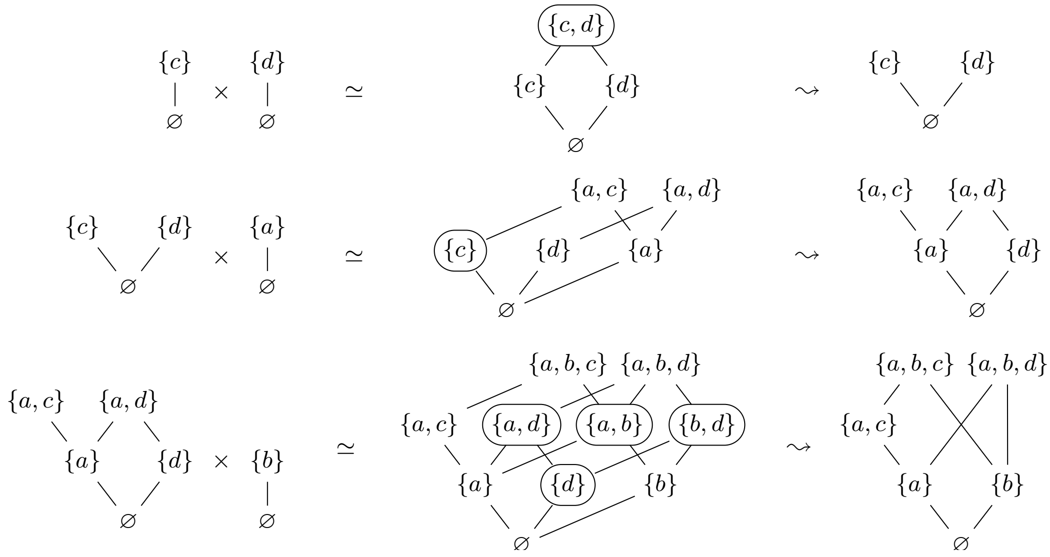

Fig. 4. Step-by-step construction of the information domain of Example 1

4.3 Applications

A Generic Construction Scheme. Suppose $\langle \Gamma, \Sigma \rangle$ is a finite theory. Let $\Sigma_0, \Sigma_1, \ldots, \Sigma_n$ be a strictly increasing sequence of sets, with $\Sigma_0 = \varnothing$ and $\Sigma_n = \Sigma$, and let Γ_i be $\Gamma|_{\Sigma_i}$. The extension from $\langle \Sigma_i, \Gamma_i \rangle$ to $\langle \Sigma_{i+1}, \Gamma_{i+1} \rangle$ factors through an extension from $\langle \Sigma_i, \Gamma_i \rangle$ to $\langle \Sigma_{i+1}, \Gamma_i \rangle$ and one from $\langle \Sigma_{i+1}, \Gamma_i \rangle$ to $\langle \Sigma_{i+1}, \Gamma_{i+1} \rangle$. Since $\langle \Sigma_{i+1}, \Gamma_i \rangle = \langle \Sigma_i, \Gamma_i \rangle \uplus \langle \Sigma_{i+1} \setminus \Sigma_i, \varnothing \rangle$, we get $C(\Gamma_{i+1})$ by deleting from $C(\Gamma_i) \times \wp(\Sigma_{i+1} \setminus \Sigma_i)$ all elements that are not consistently closed with respect to $\Gamma_{i+1} \setminus \Gamma_i$.

Example 6. Consider the theory Γ over $\Sigma = \{a, b, c, d\}$ introduced in Example 1. Let $\Sigma_1 = \{c\}$, $\Sigma_2 = \{c, d\}$, $\Sigma_3 = \{a, c, d\}$, and $\Sigma_4 = \Sigma$. Then $\Gamma_1 = \varnothing$, $\Gamma_2 \setminus \Gamma_1 = \{c \wedge d \preceq \Lambda\}$, $\Gamma_3 \setminus \Gamma_2 = \{c \preceq a\}$, and $\Gamma_4 \setminus \Gamma_3 = \{a \wedge b \preceq c \vee d,\ d \preceq a \wedge b\}$. The construction of $C(\Gamma_{i+1})$ from $C(\Gamma_i)$ is depicted by the i-th row of Figure 4, where the framed elements of $C(\Gamma_i) \times \wp(\Sigma_{i+1} \setminus \Sigma_i)$ are not consistently closed with respect to $\Gamma_{i+1} \setminus \Gamma_i$ and thus subject to deletion.

Figure 5 presents an algorithmic formulation of this generic construction scheme, with F and Σ' for $\Sigma_{i+1} \setminus \Sigma_i$ and Σ_{i+1}, respectively. Notice that the algorithm says nothing about how to choose F. The proper choice of the partition of Σ into the sets $\Sigma_{i+1} \setminus \Sigma_i$ is of course essential for keeping $|C(\Gamma_i)|$ small during the construction process, since calculating $C(\Gamma_{i+1})$ requires to check $|C(\Gamma_i)| \cdot 2^{|\Sigma_{i+1} \setminus \Sigma_i|}$ sets against $|\Gamma_{i+1} \setminus \Gamma_i|$ statements. To develop good heuristics for choosing such a partition is part of future research.

function $C\,(\Sigma$: set; Γ: theory): system of sets

if not $cc?\,(\varnothing,\, \Gamma|_{\varnothing})$ **then** $C := \varnothing$
 else
     $C := \{\varnothing\};\ \Sigma' := \varnothing$
     while $\Sigma \neq \varnothing$ **and** $C \neq \varnothing$ **do**
         F := any nonempty subset of Σ
         $\Sigma' := \Sigma' \cup F;\ \Gamma' := \Gamma|_{\Sigma'};\ C' := \varnothing$
         foreach $X \in C,\ Y \subseteq F$ **do**
             $X' := X \cup Y$
             if $cc?\,(X',\, \Gamma')$ **then** $C' := C' \cup \{X'\}$
         $\Sigma := \Sigma \setminus F;\ \Gamma := \Gamma \setminus \Gamma';\ C := C'$

function $cc?\,(X$: set; Γ: theory): boolean   { true if X is consistently Γ-closed, false otherwise }
foreach $(\phi \preceq \psi) \in \Gamma$ **do**   **if** $X \vDash \phi$ **and** $X \nvDash \psi$ **then return** (false)   **return** (true)

Fig. 5. A generic algorithm for constructing information domains

Merging. Suppose a given domain of discourse is described by two theories $\langle \Gamma_1, \Sigma_1 \rangle$ and $\langle \Gamma_2, \Sigma_2 \rangle$ whose vocabulary turns out to be partly equivalent in the sense that a set Δ of statements of the form $\phi_1 \equiv \phi_2$ is held to be true, with $\phi_1 \in T[\Sigma_1]$ and $\phi_2 \in T[\Sigma_2]$. For ease of exposition, let us assume that Σ_1 and Σ_2 are disjoint. Merging these two theories in a way that respects the equivalences in Δ then simply means to take the theory $\Gamma_1 \amalg_\Delta \Gamma_2 = \langle \Sigma_1 \cup \Sigma_2, \Gamma_1 \cup \Gamma_2 \cup \Delta \rangle$. Now let Σ_Δ be the set of all pairs $\langle \phi_1, \phi_2 \rangle$ such that $\phi_1 \equiv \phi_2$ belongs to Δ and let μ_1, μ_2 be the projections from Σ_Δ to $T[\Sigma_1]$ and $T[\Sigma_2]$, respectively. Then $\Gamma_1 \amalg_\Delta \Gamma_2$ is the *(quasi-)pushout* of μ_1, μ_2 seen as theory morphisms from $\langle \Sigma_\Delta, \varnothing \rangle$ to $\langle \Sigma_1, \Gamma_1 \rangle$ and $\langle \Sigma_2, \Gamma_2 \rangle$, respectively. By Corollary 1, it follows that the information domain of $\Gamma_1 \amalg_\Delta \Gamma_2$ is the *pullback* of $D(\mu_1)$, $D(\mu_2)$.

Example 7. Let Γ be the theory $\{a \wedge b \preceq \Lambda,\ a \preceq c\}$ over $\Sigma_1 = \{a, b, c\}$, let Γ_2 be the theory $\{d \wedge e \preceq \Lambda\}$ over $\Sigma_2 = \{d, e\}$, and suppose $\Delta = \{b \wedge c \equiv d,\ b \equiv d \vee e\}$. Figure 6 shows the information domain of $\Gamma_1 \amalg_\Delta \Gamma_2$ as the pullback of $C(\mu_1)$, $C(\mu_2)$, where μ_1 and μ_2 are as introduced before and ι_1, ι_2 are the canonical extensions. In order to determine the effect of $C(\mu_i)$ observe that $C(\mu_i)(X) = \{\langle \phi_1, \phi_2 \rangle \in \Sigma_\Delta \mid X \vDash \phi_i\}$, by definition of μ_i and Proposition 10. So $C(\mu_1)$ takes $\varnothing$, $\{c\}$, and $\{a, c\}$ to $\varnothing$, whereas $\{b\}$ is taken to $\{\langle b, d \vee e \rangle\}$ and $\{b, c\}$ to $\{\langle b \wedge c, e \rangle, \langle b, d \vee e \rangle\}$. And $C(\mu_2)$ takes $\varnothing$ to $\varnothing$, $\{d\}$ to $\{\langle b \wedge c, e \rangle, \langle b, d \vee e \rangle\}$, and $\{e\}$ to $\{\langle b, d \vee e \rangle\}$.

Remark 4. Since the information domain of a theory can be viewed as its *generic ontology* ([22]), the described construction can be regarded as a case of *ontology merging* via pullbacks. In [17], in contrast, it is argued that ontology merging is best captured by pushouts. Although a thorough discussion of this issue is beyond the scope of the present paper, a few brief remarks may be in or-

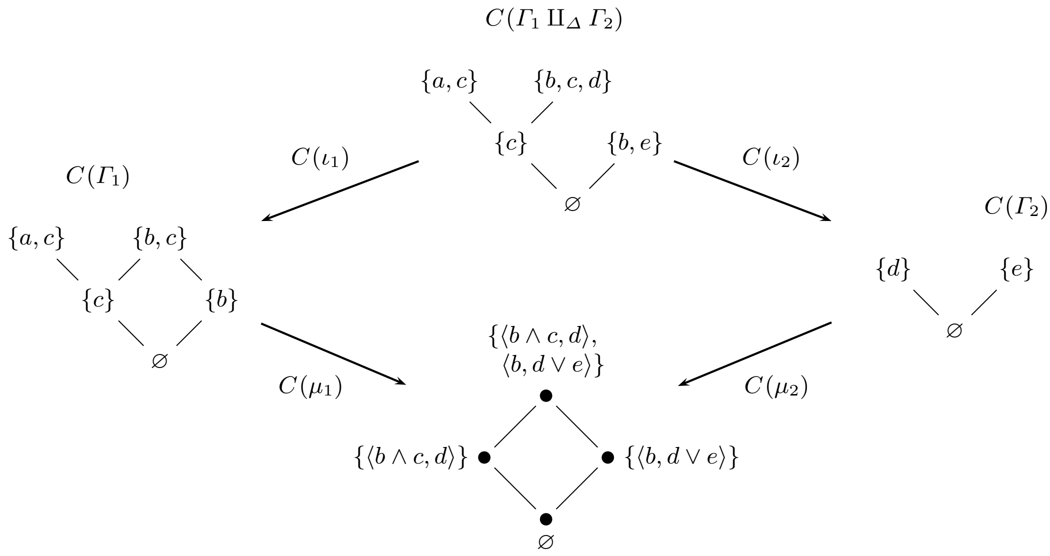

Fig. 6. Pullback of $C(\mu_1)$, $C(\mu_2)$

der. Consider the base case of merging two ontological or conceptual hierarchies without any equivalence constraints. Then the pullback construction leads to the direct product whereas pushouts are given as disjoint unions, i.e., coproducts. If, for instance, a typical zoological taxonomy of vertebrates is merged with one of birds, the disjoint union would probably be the first choice (maybe extended by an additional most general concept). Now notice that this construction implicitly takes the concepts of the two taxonomies as pairwise incompatible. If, however, concepts are conceived as sets of attributes, as in the present paper, then combining attributes from different conceptual hierarchies should be possible without necessarily requiring them to be equivalent. In fact, our merging approach can be easily extended by allowing explicit incompatibility constraints.

5 Conclusion

We have presented a natural and explicit notion of translation between classifications, seen as theories, over different base vocabularies. The resulting category of theories, together with the information domain functor, provides a nice categorical framework for studying the effect a translation of classifications has on the associated conceptual hierarchies as well as for constructing conceptual hierarchies by means of categorical constructions on the level of theories.

A major topic for future research would be an appropriate conception of translation, and thus of morphism, for more expressive logical frameworks, especially those involving *attributive* descriptions, like *description logics* and *feature logics* ([26, 21]).

References

1. Samson Abramsky and Achim Jung. Domain theory. In Samson Abramsky, Dov M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science, Vol. 3: Semantic Structures*, pages 1–168. Oxford University Press, Oxford, 1994.
2. Michael Barr and Charles Wells. *Category Theory for Computing Science*. Prentice Hall, 1995.
3. Jon Barwise. Information links in domain theory. In S. Brookes, M. Main, A. Melton, M. Mislove, and D. Schmidt, editors, *Proc. of the 7th International Conference on the Mathematical Foundations of Programming Semantics*, LNCS 598, pages 168–192, Berlin, 1992. Springer.
4. Jon Barwise and Jerry Seligman. *Information Flow. The Logic of Distributed Systems*. Cambridge University Press, Cambridge, 1997.
5. Thierry Coquand and Guo-Qiang Zhang. Sequents, frames, and completeness. In *Computer Science Logic*, LNCS 1862, pages 277–291, Berlin, 2000. Springer.
6. Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 2nd edition, 2002.
7. Georgi Dimov and Dimiter Vakarelov. On Scott consequence systems. *Fundamenta Informaticae*, 33(1):43–70, 1998.
8. Manfred Droste and Rüdiger Göbel. Non-deterministic information systems and their domains. *Theoretical Computer Science*, 75:289–309, 1990.
9. Bernhard Ganter and Rudolf Wille. Contextual attribute logic. In William Tepfenhart and Walling Cyre, editors, *Proc. of ICCS 1999*, LNAI 1640, pages 377–388, Berlin, 1999. Springer.
10. Bernhard Ganter and Rudolf Wille. *Formal Concept Analysis*. Springer, Berlin, 1999.
11. Joseph Goguen. Information integration in institutions. Manuscript, 2004.
12. Joseph Goguen and Rod Burstall. Institutions: Abstract model theory for specification and programming. *Journal of the Association for Computing Machinery*, 39(1):95–146, 1992.
13. Pascal Hitzler, Markus Krötzsch, and Guo-Qiang Zhang. A categorical view on algebraic lattices in formal concept analysis. Manuscript, 2004.
14. Pascal Hitzler and Guo-Qiang Zhang. A cartesian closed category of approximable concept structures. In Karl Erich Wolff, Heather D. Pfeiffer, and Harry S. Delugach, editors, *Proc. of ICCS 2004*, LNAI 3127, pages 170–185, Berlin, 2004. Springer.
15. Bart Jacobs. *Categorical Logic and Type Theory*. North-Holland, Amsterdam, 1999.
16. Peter T. Johnstone. *Stone Spaces*. Cambridge University Press, Cambridge, 1982.
17. Markus Krötzsch, Pascal Hitzler, Marc Ehrig, and York Sure. What is ontology merging? In *Proc. of the First International Workshop on Contexts and Ontologies*, Pittsburgh, 2005.
18. Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, New York, 1971.
19. Jørgen Fischer Nilsson. A logico-algebraic framework for ontologies. In P. Anker Jensen and P. Skadhauge, editors, *Proc. of the 1st International OntoQuery Workshop*, pages 43–56, University of Southern Denmark – Kolding, 2001.
20. Frank J. Oles. An application of lattice theory to knowledge representation. *Theoretical Computer Science*, 249:163–196, 2000.
21. Rainer Osswald. Semantics for attribute-value theories. In Paul Dekker, editor, *Proc. of the 12th Amsterdam Colloquium*, pages 199–204, Amsterdam, 1999. University of Amsterdam, ILLC.
22. Rainer Osswald. Generic ontology of linguistic classification. In Benedikt Löwe, Wolfgang Malzkorn, and Thoralf Räsch, editors, *Foundations of the Formal Sciences II*, Trends in Logic 17, pages 203–212. Kluwer, Dordrecht, 2003.
23. Rainer Osswald. A logic of classification – with applications to linguistic theory. Technical Report 303 – 9/2003, FernUniversität in Hagen, Department of Computer Science, 2003.
24. Rainer Osswald. Assertions, conditionals, and defaults. In Gabriele Kern-Isberner, Wilhelm Rödder, and Friedhelm Kulmann, editors, *Conditionals, Information, and Inference: International Workshop, WCII 2002, Revised Selected Papers*, LNAI 3301, pages 108–130, Berlin, 2005. Springer.
25. Rainer Osswald. Concept hierarchies from a logical point of view. In *Supplementary Proceedings of ICFCA-2005*, pages 15–30, Lens, France, 2005.
26. William C. Rounds. Feature logics. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*, pages 475–533. North-Holland, Amsterdam, 1997.
27. Terence P. Speed. On the order of prime ideals. *Algebra Universalis*, 2:85–87, 1972.
28. Steven Vickers. *Topology via Logic*. Cambridge University Press, Cambridge, 1989.
29. Guo-Qiang Zhang. *Logic of Domains*. Birkhäuser, Basel, 1991.