Concept Hierarchies from a Logical Point of View

Rainer Osswald

Intelligente Informations- und Kommunikationssysteme
Fachbereich Informatik, FernUniversität in Hagen
http://pi7.fernuni-hagen.de/osswald

Abstract. This paper addresses the infinitary nature of *Contextual Attribute Logic* and contrasts it with *finitary* attribute logic. Specifically, we characterize the concept hierarchies, also called *information domains*, that arise from various sorts of attribute logics, which includes a characterization of the information domains of finitary Boolean theories as locally closed subset systems. In passing, we discuss reformulations of Contextual Attribute Logic in terms of classical logic.

1 Introduction

Contextual Attribute Logic as introduced in [7] allows for infinite conjunction and disjunction. The infinitary nature of Contextual Logic has been criticized to the effect that items of knowledge should be finitely representable or at least approximable by finitely representable ones [9, 16]. While this line of thinking is directed towards imposing certain "approximation properties" on conceptual hierarchies, the present paper emphasizes the finitary nature of the "logical relationships" between attributes, if this sort of knowledge is to be held and processed by man or machine.

The paper consists of two main parts, the first of which, Section 2, reviews Contextual Attribute Logic, addresses its relation to classical logic, introduces some useful notations and techniques, and characterizes the relationship between various types of (infinitary) logics (Boolean, implicational, Horn, etc.) and the structural properties of the concept hierarchies determined by these logics. In the second part, Section 3, the focus is on finitary logic. Again, the relationship between the types of logics and corresponding concept hierarchies is determined. In particular, we characterize the hierarchies determined by finitary Boolean theories as the locally closed subset systems, whereas the infinitary case gives rise to arbitrary subset systems.

2 Contextual Attribute Logic

2.1 Review of Definitions and Facts

This section is essentially adapted from Ganter and Wille [7], with minor notational modifications and additions. Ganter and Wille introduce *Contextual Attribute Logic* as "a contextual version of the Boolean Logic of Signs and Classes", which is based on the notion of a *formal context*. A formal context $\langle U, \Sigma, \models \rangle$ consists of a set U of *(formal) objects*, a set Σ of *(formal) attributes*, and an *incidence relation* $\models$ from U to Σ. The purpose of Contextual Attribute Logic is characterized as follows:

> The central task of Contextual Attribute Logic is the investigation of the "logical relationships" between the attributes of formal contexts and, more generally, between combinations of attributes, such as implications and incompatibilities. [...] The logical relationships between formal attributes will be expressed via their extent. [7, p. 380]

The "combinations of attributes" are negation and (possibly infinite) conjunction and disjunction. More precisely, the class of *(compound) attributes over* Σ is the smallest class $B_\infty[\Sigma]$ such that $\Sigma \subseteq B_\infty[\Sigma]$, $\neg\phi \in B_\infty[\Sigma]$ if $\phi \in B_\infty[\Sigma]$, and $\bigwedge \Phi, \bigvee \Phi \in B_\infty[\Sigma]$ if $\Phi \subseteq B_\infty[\Sigma]$. We write $\phi \wedge \psi$ for $\bigwedge\{\phi, \psi\}$ and use a similar convention for $\vee$. Moreover, we write Λ for $\bigvee \varnothing$ and and V for $\bigwedge \varnothing$. Finally, $\phi \to \psi$ and $\phi \leftrightarrow \psi$ are defined in the usual way as $\neg\phi \vee \psi$ and $\phi \to \psi \wedge \psi \to \phi$, respectively.

The *extent* $p^\triangleleft$ of an atomic attribute $p \in \Sigma$ is $\{x \in U \mid x \models p\}$. The "extensional semantics" of the compound attributes is that of classical logic:

$$(\bigwedge \Phi)^\triangleleft = \bigcap\{\phi^\triangleleft \mid \phi \in \Phi\}, \quad (\bigvee \Phi)^\triangleleft = \bigcup\{\phi^\triangleleft \mid \phi \in \Phi\}, \quad \text{and} \quad (\neg\phi)^\triangleleft = U \setminus p^\triangleleft.$$

It is convenient to extend the incidence relation $\models$ to one from U to $B_\infty[\Sigma]$ in the obvious way: $x \models \phi$ iff $x \in \phi^\triangleleft$. The *dual* ϕ^{d} of a compound attribute ϕ is obtained by replacing all $\bigwedge$'s with $\bigvee$'s and all $\bigvee$'s with $\bigwedge$'s, i.e., $p^{\mathrm{d}} = p$ if $p \in \Sigma$,

$$(\bigwedge \Phi)^{\mathrm{d}} = \bigvee\{\phi^{\mathrm{d}} \mid \phi \in \Phi\}, \quad (\bigvee \Phi)^{\mathrm{d}} = \bigwedge\{\phi^{\mathrm{d}} \mid \phi \in \Phi\}, \quad \text{and} \quad (\neg\phi)^{\mathrm{d}} = \neg\phi^{\mathrm{d}}.$$

The so-called *complementary context* of a formal context $\langle U, \Sigma, \models \rangle$ is defined as $\langle U, \Sigma, \models^{\mathrm{c}} \rangle$, with $x \models^{\mathrm{c}} p$ iff $x \not\models p$. It is easy to see that $x \models^{\mathrm{c}} \phi$ iff $x \not\models \phi^{\mathrm{d}}$, which will be referred to as the *principle of duality*.

As for the "logical relationships" between (compound) attributes, an attribute ϕ is said to *(extensionally) imply* an attribute ψ in a given context if $\phi^\triangleleft \subseteq \psi^\triangleleft$; the attributes ϕ and ψ are called *extensionally equivalent* if they extensionally imply each other in the context. In case of extensional equivalence (or implication) in all contexts (with the same attribute set Σ), Ganter and Wille speak of *global* equivalence (or implication).

Global equivalence (and implication) can be checked in the so-called *test context* $\langle \wp(\Sigma), \Sigma, \models_\ni \rangle$ of Σ, with $P \models_\ni p$ iff $p \in P \subseteq \Sigma$.

Proposition 1. *Two attributes over Σ are globally equivalent iff they are extensionally equivalent in the test context of Σ.*

Observe that $P \models_\ni^{\mathrm{c}} \phi$ iff $(\Sigma \setminus P) \models_\ni \phi$, and thus $P \models_\ni \phi$ iff $(\Sigma \setminus P) \not\models_\ni \phi^{\mathrm{d}}$.

An attribute ϕ over Σ is said to be *all-extensional* in a context $\langle U, \Sigma, \models \rangle$, or to *hold* in the context, if $\phi^\triangleleft = U$. The *(infinitary) Boolean attribute logic* of $\langle U, \Sigma, \models \rangle$ is the subclass of those elements of $B_\infty[\Sigma]$ that are all-extensional in the context. The attribute logic of a formal context is invariant under *object-clarification*, i.e., under identification of objects that are indistinguishable by attributes in that they have the same intent, where the *intent* $x^{\rhd}$ of an object $x \in U$ is $\{p \in \Sigma \mid x \vDash p\}$. The standard method for object-clarification is to employ a quotient construction on U. A more concrete representation is to take the object intents as objects of the object-clarified context: Let us call $\langle \mathcal{U}, \Sigma, \vDash_{\ni} \rangle$, with $\mathcal{U} = \{x^{\rhd} \mid x \in U\}$, the *canonical object-clarification* of $\langle U, \Sigma, \vDash \rangle$.

Clause Logic Clauses provide a convenient normal form for theoretical as well as for practical purposes. A *sequent*, or *clause*, over Σ is an attribute over Σ of the form $\bigvee(\{\neg p \mid p \in A\} \cup B)$, with $A, B \subseteq \Sigma$, which is also written as $\langle A, B \rangle$. The sequent is said to be *finite* if $A \cup B$ is finite, *disjoint* if $A \cap B = \varnothing$, and *full* if $A \cup B = \Sigma$. Notice that a sequent $\langle A, B \rangle$ is (globally) equivalent to the *conditional normal form* $\bigwedge A \to \bigvee B$. Notice further that $P \subseteq \Sigma$ belongs to the extent of $\langle A, B \rangle$ in the test context of Σ just in case $A \nsubseteq P$ or $B \cap P \neq \varnothing$.

Proposition 2. *Every attribute is globally equivalent to a conjunction of full disjoint sequents.*

One can therefore restrict oneself to the *clause logic* of a formal context, which is the set of all sequents that hold in the context. In characterizing clause sets that arise as clause logics of contexts, Ganter and Wille essentially follow [2, Chap. 9]. The clause logics are precisely the *regular* clause sets, where regular means being closed under certain rules like *Weakening* and *Partitioning*. Without going into details, let us note for further use that the clause logic of a context is the *regular closure* of its full sequents. Notice, however, that this fact hinges on the infinitary nature of the logical language involved here, whereas within finitary logic over an infinite attribute set, there are no full sequents at all (cf. Section 3 below).

Free Extents The *free extent* of a set Γ of attributes over Σ is the extent of $\bigwedge \Gamma$ in the test context of Σ. If Γ is a regular clause set then $P \subseteq \Sigma$ is an element of the free extent of Γ just in case the sequent $\langle P, \Sigma \setminus P \rangle$ does not belong to Γ. Suppose now that Γ is the clause logic of a formal context $\langle U, \Sigma, \vDash \rangle$. Then Γ is also the clause logic of the associated canonical object-clarification, which means that $P \subseteq \Sigma$ is an object intent iff $\langle P, \Sigma \setminus P \rangle$ is not in Γ. Consequently:

Proposition 3. *The elements of the free extent of the (infinitary) Boolean attribute logic of a formal context are precisely the object intents of that context.*

Given a formal context $\langle U, \Sigma, \vDash \rangle$, the *intent* $V^{\rhd}$ of a set $V \subseteq U$ of objects and the *extent* $P^{\lhd}$ of a set $P \subseteq \Sigma$ of attributes are defined as follows:[1]

$$V^{\rhd} = \{p \in \Sigma \mid x \vDash p \text{ for every } x \in V\} = \bigcap\{x^{\rhd} \mid x \in V\},$$
$$P^{\lhd} = \{x \in U \mid x \vDash p \text{ for every } p \in P\} = \bigcap\{p^{\lhd} \mid p \in P\}.$$

[1] Notice that $P^{\lhd} = (\bigwedge P)^{\lhd}$.

Then $\langle V, P\rangle$ is a *formal concept* if $V^{\triangleright} = P$ and $P^{\triangleleft} = V$. Moreover, $\langle (V^{\triangleright})^{\triangleleft}, V^{\triangleright}\rangle$ is a formal concept for each $V \subseteq U$, and every formal concept is of this form. By definition, the system $\{V^{\triangleright} \mid V \subseteq U\}$ of concept intents is the closure of the system of object intents with respect to intersection.

Sequents of the form $\langle A, \{p\}\rangle$ are called *implications*. The *(infinitary) implicational attribute logic* of a formal context is the class of all implications over Σ that hold in the context. The following well-known result says that the formal concepts of a context are determined by its implicational logic (cf. [8, Sect. 2.3]):

Proposition 4. *The elements of the free extent of the (infinitary) implicational attribute logic of a formal context are precisely the concept intents of that context.*

The characterization of the free extent for other restricted logics and a reformulation of Propositions 3 and 4 by means of closure operators will be given in Section 2.3 below.

2.2 Reformulations within Classical Logic

The terminology introduced in the previous section differs considerably from that you find in standard textbooks on formal logic. According to Ganter and Wille this is intended "since the aim of Contextual Logic differs from that of Mathematical Logic" [7, p. 380]. Nevertheless, they point out that there is a close connection to *propositional logic* (see also [6]).

Propositional Logic To get a formulation of Contextual Attribute Logic in terms of propositional logic one can simply regard the elements of the attribute set Σ as *atomic propositions* (or *propositional variables*). Then $B_\infty[\Sigma]$ is the class of all (possibly infinitary) *propositional formulas* over Σ. A *(propositional) theory* over Σ is a set of formulas over Σ.

Let $\mathbb{2}$ be a set $\{0, 1\}$ of two elements equipped with the standard Boolean operations. A $\mathbb{2}$-valued (or truth-valued or Boolean) *interpretation* (or valuation) of Σ is a function m from Σ to $\mathbb{2}$; every such interpretation can be uniquely extended to a function from $B_\infty[\Sigma]$ to $\mathbb{2}$ such that $m(\bigwedge \Phi) = \bigwedge\{m(\phi) \mid \phi \in \Phi\}$, etc. A propositional formula ϕ over Σ is called *satisfiable* if there is an interpretation m of Σ such that $m(\phi) = 1$; the interpretation m is then called a *satisfier* or *model* of ϕ; m is a model of a theory Γ if m is a model of every formula of Γ.

Due to the one-to-one correspondence between subsets of Σ and their characteristic functions, one can identify interpretations of Σ with subsets of Σ. Under this identification, a subset P of Σ is a model of a formula ϕ iff the characteristic function of P takes ϕ to 1. The inductive definition of satisfaction then has the following formulation: $P \subseteq \Sigma$ satisfies $p \in \Sigma$ iff $p \in P$, P satisfies $\bigwedge \Phi$ iff P satisfies ϕ for all $\phi \in \Phi$, and so on. In other words, P is a $\mathbb{2}$-valued model of ϕ just in case P belongs to the extent of ϕ in the test context of Σ. Correspondingly, *the free extent of a theory Γ is the class of all $\mathbb{2}$-valued models of Γ.*

The connection of Contextual Attribute Logic to propositional logic proposed in [7, p. 381] for arbitrary formal contexts $\langle U, \Sigma, \models\rangle$ is to regard U as a set of situations, where $x \vDash p$ indicates that the proposition p is true in the situation x. Strictly speaking, however, this view transcends plain propositional logic in that evaluation with respect to certain situations or "worlds" comes into play. A much more natural view is that of *monadic predicate logic*, where the elements of Σ, the attributes, are predicated of the elements of U, the objects.

Attributes as Predicates The framework of propositional logic adopted in the foregoing sections has the advantage of familiarity. Moreover, propositional logic is commonly seen as the most basic sort of logic – witness any textbook on logic. Conceptually, however, it seems rather awkward to regard attributes as propositions. If attributes are formalized within a logical language at all then the most natural way to do so is to represent them as monadic predicates. For assigning attributes to objects is essentially a predication.

The propositional viewpoint can be replaced by a predicational one without much effort. The elements of Σ, the attributes, are now regarded as *atomic (monadic) predicates*. We can still apply the Boolean connectives to members of Σ by using a variable-free notation, i.e., $(\neg\phi)x$ means $\neg\phi x$, $(\bigwedge\Phi)x$ means $\bigwedge\{\phi x \mid \phi \in \Phi\}$, and so on.[2] In addition, we write $\forall\phi$ for the *universal statement* $\forall x(\phi x)$ and assume *theories* over Σ to consist of universal statements of this form.

Recall the standard definition of an *interpretation* within the framework of predicate logic: an interpretation of Σ consists of a universe U and a function that takes each monadic predicate $p \in \Sigma$ to a subset of U. Now observe that a formal context $\langle U, \Sigma, \vDash\rangle$ uniquely corresponds to an interpretation M of Σ, and vice versa: simply define $M(p) = p^{\lhd} = \{x \in U \mid x \vDash p\}$. The notion of an interpretation gives us the notion of truth and model as well: a statement $\forall\phi$ is *true* with respect to the interpretation M if $M(\phi) = U$ (with M extended to $B_\infty[\Sigma]$); the interpretation M is a *model* of a theory Γ if every statement of Γ is true with respect to M. Instead of saying that a (compound) predicate (or attribute) ϕ is "all-extensional" in a formal context (cf. Section 2.1) we can now say that $\forall\phi$ is true in (the interpretation corresponding to) the context.

Under this perspective, the test context of Σ provides a *canonical interpretation* of Σ. Moreover, if $\mathcal{U}$ is the free extent of a theory Γ over Σ, then the context $\langle \mathcal{U}, \Sigma, \vDash_\ni\rangle$ determines a *canonical model* of Γ. It is not difficult to see that this model is the "largest object-clarified" model of Γ, in the sense that every other such model is embeddable into it, and that the canonical model is *universal* in the sense that a statement is true in that model if it is true in all models of Γ.

The predicational viewpoint also allows a clear distinction between a background theory consisting of universal statements on the one hand and assertions about a specific object on the other. Within Contextual Attribute Logic as presented in [7], this distinction is only implicit in that compound attributes serve "not only as "generalized attributes", but also as logical rules that may or may

[2] Formally, this can be realized by *predicate abstraction*.

Ganter & Wille [7]	Barwise & Seligman [2]	Present paper [10, 12]
formal context	classification	satisfaction relation, interpretation
test context	powerset classification	canonical interpretation
object-clarified	separated	identity of indiscernibles
consistent set	consistent partition	consistently closed set
free extent	(tokens of) generated classification	canonical model, information domain
Boolean attribute logic of formal context	generated theory (only clauses)	canonical (Boolean) theory

Table 1. Comparison of terminologies

not hold in the given context" [7, p. 382]. It is this implicit universal quantification of rules that becomes explicit under the present reformulation.

Fixing Terminology and Notation In the rest of this paper, we employ the following terminological and notational conventions, which more or less resemble those of [10, 12]. Table 1 provides an overview of how these conventions are related to that of Ganter and Wille and of Barwise and Seligman, respectively.

We take up the predicational view presented above, that is, we speak of the class $B_\infty[\Sigma]$ of compound predicates over a set Σ of atomic (monadic) predicates and of theories as sets of statements that are universally quantified compound predicates.[3] Moreover, we speak of interpretations and their associated satisfaction relation, of the canonical interpretation of Σ, and of the canonical model of a theory. We use $C(\Gamma)$ to refer to the universe of the canonical model of Γ, that is, $C(\Gamma)$ is the free extent of Γ. Adapting the terminology of [5], we call $C(\Gamma)$ the *information domain* of Γ. Notice that two theories over Σ are equivalent, i.e., entail each other (in the sense of having the same models), if and only if they have identical information domains.

Instead of $\forall(\phi \rightarrow \psi)$ and $\forall(\phi \leftrightarrow \psi)$ we also write $\phi \preceq \psi$ and $\phi \equiv \psi$, respectively. A (compound) predicate is called *positive* or *affirmative* if it is free of $\neg$. A statement of the form $\phi \preceq \psi$ (or $\phi \equiv \psi$), with ϕ and ψ positive, is said to have *conditional* (or *biconditional*) *form*. The conditional form is called *normal*, if ϕ is purely conjunctive and ψ is purely disjunctive. A conditional normal form is thus the same as a *clausal* form.

[3] Notice that we do not require a theory to be closed under entailment.

2.3 Theory Types and Closure Operators

Let us now return to the issue addressed at the close of Section 2.1. In addition to the general class $B_\infty[\Sigma]$ of Boolean predicates over Σ, we make use of the following notations for special types of predicates (with $p, q \in \Sigma$ and $P, Q \subseteq \Sigma$):

$H_\infty[\Sigma]$	Horn predicates	$\bigwedge P \to \Lambda$, $\bigwedge P \to q$ (or $\bigwedge P \to \bigwedge Q$)
$I_\infty[\Sigma]$	implications	$\bigwedge P \to q$ (or $\bigwedge P \to \bigwedge Q$)
$O_\infty[\Sigma]$	contradictions	$\bigwedge P \to \Lambda$ (or $\neg \bigwedge P$)
$S[\Sigma]$	simple implications	$p \to q$

By an *(possibly infinitary) Horn statement* or *H_∞-statement* over Σ we mean a statement of the form $\forall\phi$ with $\phi \in H_\infty[\Sigma]$. An *(infinitary) Horn theory* or *H_∞-theory* is a set of Horn statements. A corresponding terminology is employed for the other cases.

Recall from Section 2.1 that in order to determine the Boolean attribute logic of a formal context we can work with the canonical object-clarification instead. This comes down to determining the Boolean attribute logic of contexts of the form $\langle \mathcal{U}, \Sigma, \vDash_\ni \rangle$, where $\mathcal{U}$ is a subset system over Σ. Adapted to the present terminology, this means to ask for the class of all Boolean statements over Σ that are true in the canonical interpretation of Σ restricted to $\mathcal{U}$. We speak of the *canonical Boolean theory* (or *B_∞-theory*) $T_{B_\infty}(\mathcal{U})$ associated with $\mathcal{U}$. In the same vein, we define the *canonical τ-theory* $T_\tau(\mathcal{U})$ of $\mathcal{U}$ with τ ranging over the other types of theories introduced above. By a *complete τ-theory* of $\mathcal{U}$ we mean a τ-theory Γ over Σ that is equivalent to $T_\tau(\mathcal{U})$, that is, $\Gamma \subseteq T_\tau(\mathcal{U})$ and Γ entails $T_\tau(\mathcal{U})$.

It is easy to see that $C \circ T_\tau$ is a closure operator on $\wp(\Sigma)$. For let $\Pi_\tau(\mathcal{U})$ be the set $\{\phi \,|\, P \vDash_\ni \phi$ for every $P \in \mathcal{U}\}$ and let $Q_\forall(\Phi)$ be $\{\forall\phi \,|\, \phi \in \Phi\}$, for $\Phi \subseteq \tau[\Sigma]$. Then $T_\tau = Q_\forall \circ \Pi_\tau$ and the pair $\langle \Pi_\tau, C \circ Q_\forall \rangle$ is the Galois connection between $\wp(\wp(\Sigma))$ and $\wp(\tau[\Sigma])$ induced by the satisfaction relation $\vDash_\ni$.[4] In particular, it follows that if $\Gamma \subseteq \Gamma'$ then $C(\Gamma') \subseteq C(\Gamma)$.

Suppose now that the subset systems that arise as information domains (i.e., as free extents) of τ-theories are characterized by certain closure properties. Then the closure operator $C \circ T_\tau$ takes an arbitrary subset system $\mathcal{U}$ to the closure of $\mathcal{U}$ with respect to these properties. The information domains of implicational theories, for instance, are known to be closed with respect to the intersection of arbitrary subsets. Hence $C(T_{I_\infty}(\mathcal{U}))$ is the closure of $\mathcal{U}$ with respect to intersection, which essentially resembles Proposition 4. Proposition 3, on the other hand, tells us that $C(T_{B_\infty}(\mathcal{U})) = \mathcal{U}$, which means that there are no closure conditions for the information domains of B_∞-theories. See Table 2 for the respective closure conditions for the other theory types. ('Nonempty intersection' is short for 'intersection of nonempty subsets'.) Notice that the closure conditions given in the table are not only necessary but also sufficient for a subset system to be the information domain of a theory of the respective type.

[4] This technicality concerning the universal quantifier is surely one the less attractive features of the predicational viewpoint.

τ	Closure properties of $C(\Gamma)$	Order-theoretic characterization
B_∞	none	poset
H_∞	nonempty intersection	bounded-complete poset
I_∞	intersection	complete lattice
O_∞	subsets	bounded-complete atomic poset with completely coprime atoms
S	intersection + union	completely distributive complete lattice
S, O_∞	nonempty intersection + bounded union	

Table 2. Relationship between theories and information domains for infinitary logic

The following example, which is adapted from [12], gives an illustration of the various closure operations for a context over a finite set of attributes (where we drop the index ∞).

Example 1. Let $\langle U, \Sigma, \models\rangle$ be the formal context depicted by the cross table on the left of Figure 1. The diagram on the right of the figure shows the system $\mathcal{U}$ of object intents determined by the context. Let Γ be the theory over $\{a, b, c, d, e\}$ consisting of the statements

$$d \preceq c, \quad c \preceq a, \quad e \preceq a \wedge b, \quad b \wedge c \preceq d,$$
$$c \wedge e \preceq \Lambda, \quad a \wedge b \preceq c \vee e, \quad V \preceq a \vee b.$$

The information domain $C(\Gamma)$ of Γ coincides with the system of object intents $\mathcal{U}$, as the reader will easily check. In other words, Γ is a complete Boolean theory of $\mathcal{U}$ (and thus of the given context). Figure 2 provides an overview

	a	b	c	d	e
x_1	×	×			×
x_2	×				
x_3		×			
x_4	×	×	×	×	
x_5	×		×		
x_6	×	×	×	×	
x_7	×		×	×	

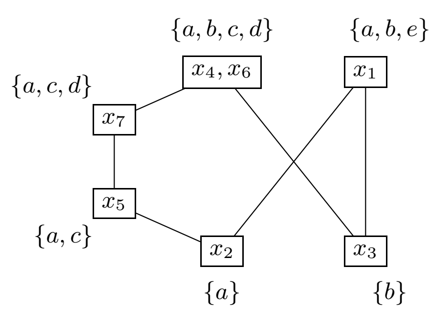

Fig. 1. Formal context and induced system of object intents

of the information domains of several complete τ-theories of $\mathcal{U}$, with varying τ. The top of the figure shows the information domain of a complete simple implication theory of $\mathcal{U}$; it is the closure of $\mathcal{U}$ with respect to intersection and union. A (nonredundant) complete simple implication theory of $\mathcal{U}$ is given by the statements $d \preceq c$, $c \preceq a$, $e \preceq a$, and $e \preceq b$. The diagram below the top on the left depicts the closure of $\mathcal{U}$ with respect to intersection of nonempty subsets and union of bounded subsets. It is the information domain of the extension of the above simple implication theory by the contradiction statement $c \wedge e \preceq \Lambda$. Addition of the Horn statement $b \wedge c \preceq d$ further weakens the closure properties of the associated information domain. If the statement $b \wedge c \preceq d$ is added to the simple implication theory before the contradiction statement $c \wedge e \preceq \Lambda$, the resulting effect on the respective information domains is as depicted by the right branch of Figure 2. Finally, adding the statements $V \preceq a \vee b$ and $a \wedge b \preceq c \vee e$ leads to a Γ and hence to $\mathcal{U}$. □

The last column of Table 2 gives an order-theoretic characterization of the information domains of theories of the respective type. Consider, for instance, the case of (infinitary) Horn theories. An ordered set (poset) is bounded-complete, i.e., has suprema for all upwards-bounded subsets, just in case it has infima for all nonempty subsets (see e.g. [4]). The down-set representation of this ordered set as an ordered set of sets is therefore closed with respect to the intersection of nonempty subsets. To give another example, take the implicational theories. Their information domains coincide with the complete lattices of concept intents, and every complete lattice is known to arise that way.

3 Finitary Logic

As mentioned in the introduction, it seems reasonable to assume that an adequate representational framework for knowledge and information in mind and machine is restricted to finitary "logical relationships". Ganter and Wille, however, take infinitary compound attributes into account. The same is true of the framework of Barwise and Seligman, who treat finitary theories only briefly in passing [2, Sect. 9.2]. In contrast, [1, 15] and especially [5] are devoted to the study of the information domains of finitary theories, where theories are typically represented as sets of sequents, also called *sequent structures* or *nondeterministic information systems*.

In what follows, we apply the techniques developed in Section 2 to the study of finitary attribute logic. In particular, we give a characterization of the information domains of finitary Boolean theories.

3.1 Informations Domains of Finitary Theories

Given a set Σ of atomic predicates (or attributes) let $B_\omega[\Sigma]$ be the set of compound predicates over Σ that are inductively constructed by applying $\neg$, $\bigwedge$, and $\bigvee$, where the operators $\bigwedge$ and $\bigvee$ take only *finite* sets. A *(finitary) Boolean*

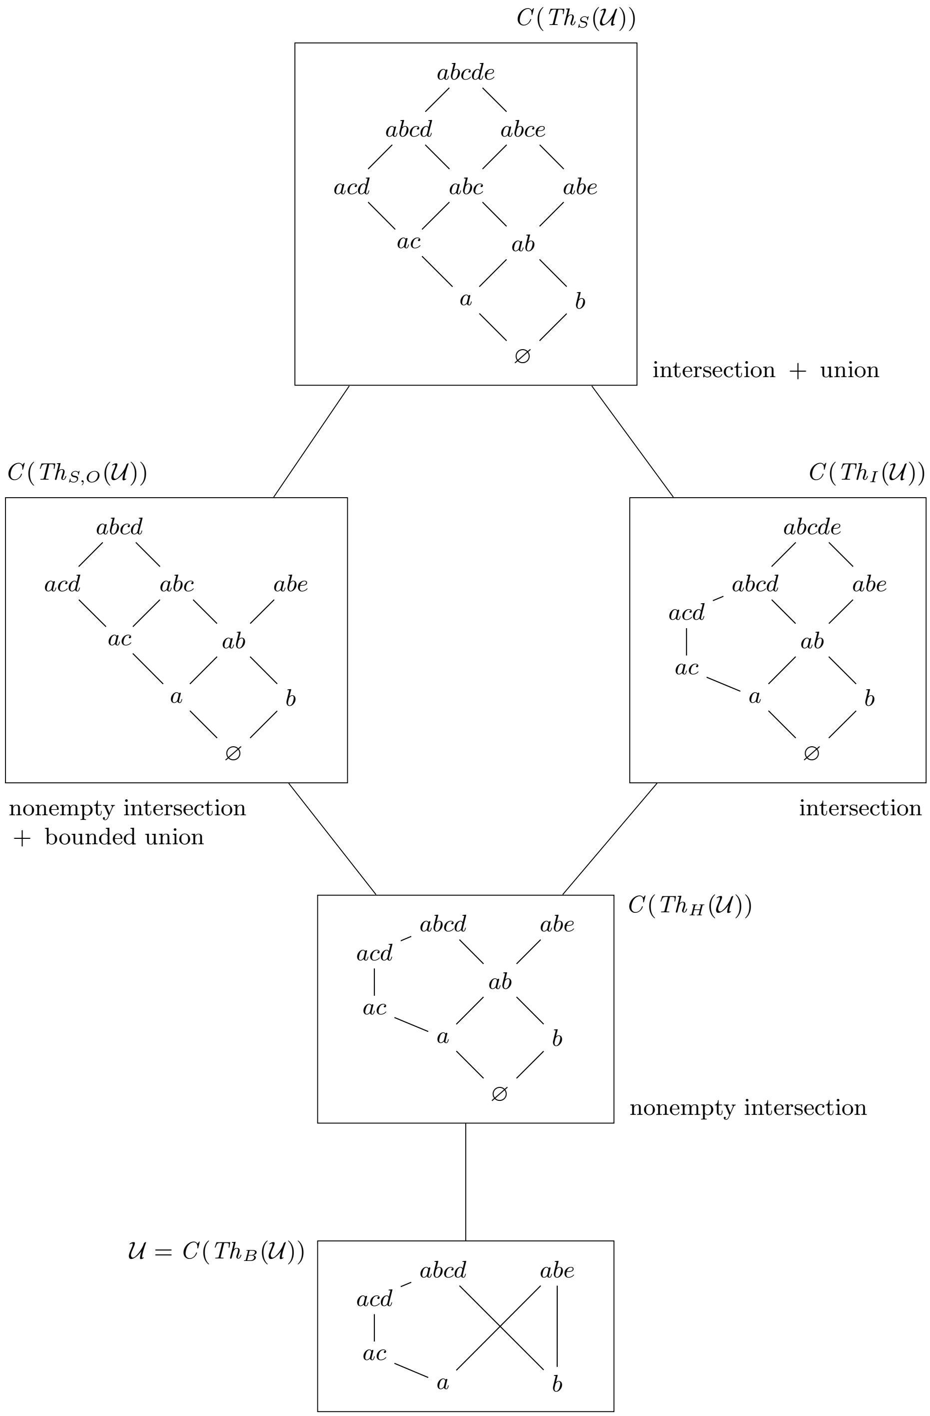

Fig. 2. Information domain (free extent) of complete τ-theory of $\mathcal{U}$ with varying τ

theory is a set of statements of the form $\forall\phi$, with $\phi \in B_\omega[\Sigma]$. All other definitions from Section 2 carry over in a similar vein. In the following, let $\vDash$ always be the canonical satisfaction relation $\vDash_\ni$, if not otherwise indicated. Recall that a predicate is called *positive* (or *affirmative*) if is free of $\neg$. (Notice that Λ is positive too.) It is easy to see that every positive predicate ϕ is *persistent* (or *increasing*) in the sense that if $P \vDash \phi$ and $P \subseteq Q$ then $Q \vDash \phi$, for all $P, Q \subseteq \Sigma$.

Information domains of finitary theories, in contrast to those of infinitary ones, turn out to be *(upwards and downwards) directed complete.* A nonempty system $\mathcal{S}$ of subsets of Σ is called *upwards directed* if for all $P, Q \in \mathcal{S}$ there is a $R \in \mathcal{S}$ with $P \cup Q \subseteq R$. A subset system $\mathcal{S}$ is called *downwards directed* if for all $P, Q \in \mathcal{S}$ there is an $R \in \mathcal{S}$ with $R \subseteq P \cap Q$.

Lemma 1. *Let $\mathcal{S}$ be subset system and ϕ a positive predicate over Σ.*

(i) If $\mathcal{S}$ is upwards directed, then $\bigcup\mathcal{S} \vDash \phi$ iff $P \vDash \phi$ for some $P \in \mathcal{S}$.

(ii) If $\mathcal{S}$ is downwards directed, then $\bigcap\mathcal{S} \vDash \phi$ iff $P \vDash \phi$ for every $P \in \mathcal{S}$.

Proof. (i) is straightforwardly proved by induction. Let us verify the induction step for conjunction: by induction hypothesis, $\bigcup\mathcal{S} \vDash \phi \wedge \psi$ iff for some elements P and Q of $\mathcal{S}$, $P \vDash \phi$ and $Q \vDash \psi$, that is, iff $\phi \wedge \psi$ is satisfied by some element of $\mathcal{S}$ since $\mathcal{S}$ is upwards directed and $\phi \wedge \psi$ is persistent. As for (ii) we can apply the principle of duality to (i): If $\mathcal{S}$ is downwards directed then $\mathcal{S}' = \{\Sigma \setminus P \mid P \in \mathcal{S}\}$ is upwards directed and $\bigcap\mathcal{S} = \Sigma \setminus (\bigcup\mathcal{S}')$. Hence $\bigcap\mathcal{S} \vDash \phi$ iff $\bigcup\mathcal{S}' \nvDash \phi^{\mathrm{d}}$ iff, for every $P \in \mathcal{S}$, $(\Sigma \setminus P) \nvDash \phi^{\mathrm{d}}$, i.e. $P \vDash \phi$. □

Proposition 5. *The information domain of a (finitary) Boolean theory is closed with respect to the union of upwards directed subsets and with respect to the intersection of downwards directed subsets.*

Proof. Let Γ be a theory over Σ and suppose $\mathcal{S} \subseteq C(\Gamma)$ is upwards directed. Without loss of generality, we may assume that Γ consists of statements of the form $\phi \preceq \psi$, with ϕ, ψ positive. We need to show that $\bigcup\mathcal{S} \vDash (\phi \preceq \psi)$. But if $\bigcup\mathcal{S} \vDash \phi$ then, by Lemma 1, $P \vDash \phi$ for some $P \in \mathcal{S}$, hence $P \vDash \psi$, and thus $\bigcup\mathcal{S} \vDash \psi$. The closure of $C(\Gamma)$ with respect to the intersection of downwards directed sets follows in the same way by applying the principle of duality. □

The foregoing result is an instance of the general fact that all order-theoretic properties of information domains of Boolean theories are invariant under (order-theoretic) duality. For if $\mathcal{U}$ is the information domain of a theory Γ over Σ then, by the principle of duality, $\{\Sigma \setminus P \mid P \in \mathcal{U}\}$ is the information domain of $\{\neg\phi^{\mathrm{d}} \mid \phi \in \Gamma\}$.

The following two examples show that information domains of (finitary) Boolean theories are not necessarily algebraic nor coherent, if algebraic (see also [5]).

Example 2. Let Γ be the theory over $\Sigma = \{a_1, a_2, \ldots\}$ with statements

$$a_{n+1} \preceq a_n \quad \text{and} \quad a_n \preceq a_1 \vee a_{n+1} \quad (n > 1).$$

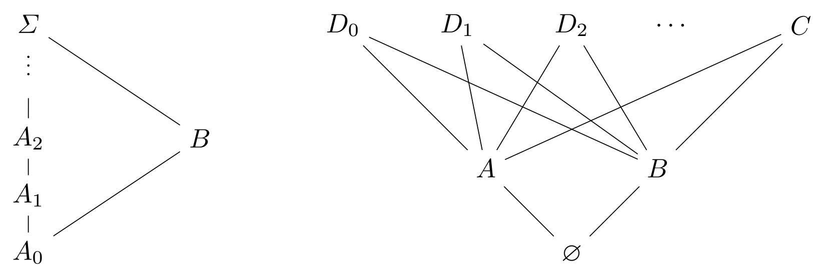

Fig. 3. Non-algebraic domain and non-coherent algebraic domain

The information domain $C(\Gamma)$ is shown on the left of Figure 3, with $A_n = \{a_m \,|\, m \leq n\}$ and $B = \{a_n \,|\, n > 1\}$. This ordered set is *not algebraic*, because B is neither compact nor the supremum of a directed set of compact elements.

Example 3. Suppose Γ is the theory over $\{a, b\} \cup \{c_0, c_1, \ldots\} \cup \{d_0, d_1, \ldots\}$ that consists of the statements

$$a \wedge b \equiv c_0 \vee d_0, \quad c_n \equiv c_{n+1} \vee d_{n+1}, \quad c_n \wedge d_n \equiv \Lambda \quad (n \geq 0).$$

Then $C(\Gamma)$ consists of $\varnothing$, $A = \{a\}$, $B = \{b\}$, $C = \{a, b\} \cup \{c_0, c_1, c_2, \ldots\}$, and $D_n = \{a, b\} \cup \{c_0, c_1, \ldots, c_{n-1}, d_n\}$, $n \geq 0$; see the diagram on the right of Figure 3. Every element of $C(\Gamma)$ is compact, but the set of minimal upper bounds of $\{A, B\}$ is infinite; so $C(\Gamma)$ is *not coherent*.

3.2 Locally Closed Subset Systems

In this section, we present a characterization of the information domains of finitary Boolean theories as subset systems. The key definition of a locally closed system draws on an idea of Davey [3], who employs a similar concept to characterize the (ordered) sets of prime ideals of distributive lattices.[5]

Definition 1 (Local membership/locally closed). *Let $\mathcal{U}$ be a subset system over Σ. A subset Q of Σ is* locally a member of $\mathcal{U}$ *if for every* finite *subset F of Σ there is an $P \in \mathcal{U}$ such that $Q \cap F = P \cap F$. The system $\mathcal{U}$ is* locally closed *if it contains every subset of Σ which is locally a member of $\mathcal{U}$.*

Clearly, every subset system $\mathcal{U}$ over a *finite* set Σ is locally closed. For if Q is locally a member of $\mathcal{U}$ then, since Σ is finite, there is an element P of $\mathcal{U}$ such that $Q = Q \cap \Sigma = P \cap \Sigma = P$.

[5] The reason that Davey's work proves useful here is that the information domain of a finitary Boolean theory Γ over Σ can be represented as the set of prime filters of a distributive lattice, namely the quotient of the class of positive predicates over Σ modulo equivalence with respect to Γ; see e.g. [11, Chap. 6] for details.

Proposition 6. *The information domain of a finitary Boolean theory is locally closed.*

Proof. Suppose $Q \subseteq \Sigma$ is locally a member of $C(\Gamma)$ and $\phi \in \Gamma$. Since the set F of all elements of Σ occurring in ϕ is finite, there is a $P \in C(\Gamma)$ such that $P \cap F = Q \cap F$; hence $Q \models \phi$. It follows that $Q \in C(\Gamma)$. $\square$

Example 4. If Σ is infinite, there is no theory over Σ with information domain $\mathcal{U} = \{\{p\} \,|p \in \Sigma\}$. To see this, notice that $\varnothing$ is locally a member $\mathcal{U}$, since for every finite subset F of Σ there is a $p \in \Sigma$ such that $p \notin F$, that is, $\varnothing \cap F = \{p\} \cap F$.

In order to show that the property of being locally closed characterizes the information domain of finitary Boolean theories, it remains to check that any locally closed system $\mathcal{U}$ over Σ is the information domain of some theory over Σ. This comes down to proving that $\mathcal{U} = C(T(\mathcal{U}))$, with T short for T_{B_ω}. From Section 2, we know the following:

Lemma 2. *If Σ is finite then $C(T(\mathcal{U})) = \mathcal{U}$.*

Let $\mathcal{U}$ be a subset system over Σ and S a subset of Σ. The *restriction $\mathcal{U}|_S$ of $\mathcal{U}$ to S* is the subset system over S defined by $\mathcal{U}|_S = \{P \cap S \mid P \in \mathcal{U}\}$.

Lemma 3. *If $Q \in C(T(\mathcal{U}))$ then $Q \cap S \in C(T(\mathcal{U}|_S))$.*

Proof. Suppose $Q \in C(T(\mathcal{U}))$ and $\phi \in T(\mathcal{U}|_S)$. We need to show that $Q \cap S \models \phi$, that is, $Q \models \phi$, since ϕ is a predicate over S. But for the same reason ϕ belongs to $T(\mathcal{U})$. $\square$

Proposition 7. *If $\mathcal{U}$ is locally closed then $C(T(\mathcal{U})) = \mathcal{U}$.*

Proof. We need to show that $C(T(\mathcal{U})) \subseteq \mathcal{U}$. Suppose $Q \in C(T(\mathcal{U}))$ and F is a finite subset of Σ. Then, by Lemma 3, $Q \cap F$ belongs to $C(T(\mathcal{U}|_F))$, and thus to $\mathcal{U}|_F$, by Lemma 2. Hence $Q \in \mathcal{U}$, since $\mathcal{U}$ is locally closed. $\square$

Together with Proposition 6, this give us the following characterization of information domains of finitary Boolean theories:

Theorem 1. *The information domain of a finitary Boolean theory over Σ is locally closed and every locally closed subset system over Σ arises that way.*

3.3 Theory Types and Closure Operators

As in Section 2.3, we get closure operators $C \circ T_\tau$ for the finitary versions B_ω, H_ω, I_ω, and O_ω of theory types. (Notice that simple implication theories are finitary by definition since no conjunction or disjunction is involved.) Again, if the subset systems that arise as information domains of τ-theories can be characterized by closure conditions, then $C(T_\tau(\mathcal{U}))$ is the closure of $\mathcal{U}$ with respect to these conditions.

For $\tau = B_\omega$, Theorem 1 tells us that the information domain of a complete finitary Boolean theory of $\mathcal{U}$ is the closure of $\mathcal{U}$ with respect to local membership. Formulated in the language of Section 2.1, the free extent of the finitary Boolean logic of a formal context is the closure of the system of object intents with respect to local membership. (Compare this to Proposition 3.) Table 3 list the closure properties that correspond to the various theory types. The proofs are fairly straightforward and can be found, e.g., in [11, Chap. 3].

τ	Closure properties of $C(\Gamma)$	Order-theoretic characterization
B_ω	local membership	profinite poset
H_ω	nonempty intersection + directed union	Scott domain (bounded-complete algebraic dcpo)
I_ω	intersection + directed union	complete algebraic lattice
O_ω	subsets + finitely bounded union	bounded-complete atomic dcpo with completely coprime atoms

Table 3. Relationship between theories and information domains for finitary logic

Table 3 also contains order-theoretic characterizations for the information domains of the types of theories discussed so far. To the best of our knowledge, no satisfying, purely order-theoretic characterization of the information domains of finitary Boolean theories has been given up to now. A result worth mentioning is that of Speed [13], which implies that an ordered set can be represented as an information domain of a finitary theory just in case it is *profinite*, i.e., a projective limit of a projective system of finite ordered sets. In addition to this somewhat indirect characterization, Table 3 provides order-theoretic characterizations for the information domains of Horn theories and their subtypes (see [11, Chap. 4] for proofs and pointers to the literature).

Approximable Concepts For Formal Concept Analysis proper, the theory type I_ω is of special interest. According to Table 3, the free extent of the finitary implicational logic of a formal context is the closure of the system of object intents with respect to the intersection of arbitrary subset *and the union of directed subsets*. Furthermore, the resulting "concept lattice" is algebraic and every complete algebraic lattice arises that way.

According to [8, pp. 33f], a concept lattice is algebraic just in case the system $\mathcal{U}$ of concept intents consists precisely of those sets $P \subseteq \Sigma$ of attributes with

$$(F^{\lhd})^{\rhd} \subseteq P \quad \text{for every finite subset } F \text{ of } P. \tag{1}$$

The following fact, which is easily verified by unraveling definitions, sheds more light on this condition:

$$p \in (F^{\lhd})^{\rhd} \quad \text{iff} \quad (\bigwedge F \preceq p) \in T_{I_\omega}(\mathcal{U}).$$

So the above condition means that P is a concept intent iff P is in the free extent of the *finitary* implicational logic of the context. Put differently, *a concept lattice is algebraic just in case its implicational logic is finitary.*

Zhang and Shen [16], who study formal contexts under the name of *Chu spaces*, define an *approximable concept* as a set $P \subseteq \Sigma$ of attributes that satisfies condition (1).[6] According to what has just been said, the approximable concepts are the elements of the free extent of the finitary implicational logic of the context. This proves the following theorem of Zhang and Shen:

Proposition 8. *The information domain of a complete (finitary) implicational theory of a formal context consists of the approximable concepts of that context.*

The lattice of approximable concepts is thus obtained from the system of object intents by taking arbitrary intersections and directed unions.

4 Conclusion

The systematic approach to relate logical properties of theories to structural properties the resulting information domains (or concept hierarchies) which has been put forward in this papers may be further pursued in various directions. To name just two of them, although considerable effort has been spent to characterize "good" algebraic domains beyond Scott domains in terms of sequent structures (see e.g. [5] or the work of G. Q. Zhang), a satisfying logical characterization of the respective theories seems to be still missing. A second topic for further research is to allow infinite disjunction while adhering to finite conjunction, which leads to the so-called *observational* or *geometric* logic (e.g. [14]).

References

1. Jon Barwise. Information links in domain theory. In S. Brookes, M. Main, A. Melton, M. Mislove, and D. Schmidt, editors, *Proceedings of the 7th International Conference on the Mathematical Foundations of Programming Semantics*, Lecture Notes in Computer Science 598, pages 168–192, Berlin, 1992. Springer.
2. Jon Barwise and Jerry Seligman. *Information Flow. The Logic of Distributed Systems*. Cambridge Tracts in Theoretical Computer Science 44. Cambridge University Press, Cambridge, 1997.
3. Brian A. Davey. A note on representable posets. *Algebra Universalis*, 3:345–347, 1973.

[6] In [16], implicational theories are represented by so-called *information systems*, which are a special sort of sequent structures.

4. Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 2nd edition, 2002.
5. Manfred Droste and Rüdiger Göbel. Non-deterministic information systems and their domains. *Theoretical Computer Science*, 75:289–309, 1990.
6. Bernhard Ganter. Attribute exploration with background knowledge. *Theoretical Computer Science*, 217:215–233, 1999.
7. Bernhard Ganter and Rudolf Wille. Contextual attribute logic. In William Tepfenhart and Walling Cyre, editors, *Proceedings of ICCS 1999*, Lecture Notes in Artificial Intelligence 1640, pages 377–388, Berlin, 1999. Springer.
8. Bernhard Ganter and Rudolf Wille. *Formal Concept Analysis*. Springer, Berlin, 1999.
9. 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, *Proceedings of ICCS 2004*, Lecture Notes in Artificial Intelligence 3127, pages 170–185, Berlin, 2004. Springer.
10. Rainer Osswald. Classifying classification. In Geert-Jan Kruijff, Larry Moss, and Dick Oehrle, editors, *Proceedings of the Joint Conference on Formal Grammar and Mathematics of Language (FG/MOL-01)*, Electronic Notes in Theoretical Computer Science 53, 2001.
11. Rainer Osswald. A logic of classification – with applications to linguistic theory. Technical Report 303 – 9/2003, FernUniversität Hagen, Fachbereich Informatik, 2003.
12. Rainer Osswald and Wiebke Petersen. A logical approach to data driven classification. In *KI-2003: Advances in Artificial Intelligence*, Lecture Notes in Artificial Intelligence 2821, pages 267–281, Berlin, 2003. Springer.
13. Terence P. Speed. On the order of prime ideals. *Algebra Universalis*, 2:85–87, 1972.
14. Steven Vickers. *Topology via Logic*. Cambridge Tracts in Theoretical Computer Science 5. Cambridge University Press, Cambridge, 1989.
15. Guo-Qiang Zhang. *Logic of Domains*. Birkhäuser, Basel, 1991.
16. Guo-Qiang Zhang and Gongqin Shen. Approximable concepts, Chu spaces, and information systems. *Theory and Applications of Categories*, 2004. to appear.