

Non-monotonic Extensions to Formal Concept Analysis via Object Preferences

Lucas Carr¹, Nicholas Leisegang¹, Thomas Meyer¹, and Sebastian Rudolph^{2,3}

¹ University of Cape Town and CAIR, South Africa

{crrluc003, lsgnic001}@myuct.ac.za, tommie.meyer@uct.ac.za

² Technische Universität Dresden, Germany

³ Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig
sebastian.rudolph@tu-dresden.de

Abstract. Formal Concept Analysis (FCA) is an approach to creating a conceptual hierarchy in which a *concept lattice* is generated from a *formal context*. That is, a triple consisting of a set of objects, G , a set of attributes, M , and an incidence relation I on $G \times M$. A *concept* is then modelled as a pair consisting of a set of objects (the *extent*), and a set of shared attributes (the *intent*). Implications in FCA describe how one set of attributes follows from another. The semantics of these implications closely resemble that of logical consequence in classical logic. In that sense, it describes a monotonic conditional. The contributions of this paper are two-fold. First, we introduce a non-monotonic conditional between sets of attributes, which assumes a preference over the set of objects. We show that this conditional gives rise to a consequence relation that is consistent with the postulates for non-monotonicity proposed by Kraus, Lehmann, and Magidor (commonly referred to as the KLM postulates). We argue that our contribution establishes a strong characterisation of non-monotonicity in FCA. To our knowledge, this is a novel view of FCA as a formalism which supports non-monotonic reasoning. We then extend the influence of KLM in FCA by introducing the notion of *typical* concepts through a restriction placed on what constitutes an acceptable preference over the objects. Typical concepts represent concepts where the intent aligns with expectations from the extent, allowing for an exception-tolerant view of concepts. To this end, we show that the set of all typical concepts is a meet semi-lattice of the original concept lattice. This notion of typical concepts is a further introduction of KLM-style typicality into FCA, and is foundational towards developing an algebraic structure representing a concept lattice of prototypical concepts.

Keywords: Formal Concept Analysis · Non-monotonic reasoning · Rational consequence relations · Defeasible reasoning

1 Introduction

Formal Concept Analysis (FCA) is a lattice-theoretic approach to representing and reasoning about concepts and hierarchies in data. The view of concepts adopted in FCA has clear philosophical underpinnings, describing a concept as a pair representing the dualism between extension, *what a concept refers to*, and intension, *what a concept means* [8,9,19]. As such, the setting of FCA works with data in the form of a formal context, describing a relationship between objects and attributes.

An important component of FCA involves discovering implications in the data which describe a complete correspondence between sets of attributes [19]. A concern might be that complete correspondence is too strict a requirement, and implications may fail to capture relationships that, while useful, are only partial correspondences between attribute sets. To illustrate this point, we introduce a variation of the customary example of penguins, where we look at species of animals and some of their attributes (see Example 1 for a more detailed description) [10,13]. In our example we have three objects which have the attribute `bird`: duck, robin, and penguin. Of these, duck and robin also have the attribute `flies`, and obviously penguin does not. Consequently, the implication `bird` \rightarrow `flies` is not valid in our context. One may argue that this behaviour is completely correct. Clearly, not all birds fly. A compelling response is that birds that do not fly are exceptions to the rule, and that the ability to express that birds usually do fly is useful.

Association rules offer an existing approach to address this concern. These rules introduce a notion of confidence, which quantifies the proportion of data that conforms to a given rule, if the confidence is above a user-defined threshold are then accepted [8]. While association rules are capable of capturing partial relationships, this approach may be considered somewhat blunt, relying on something analogous to *majority rules*. The reliance on a threshold for rule acceptance limits the expressivity of these rules.

We propose an alternate approach which adopts the view that we can define a preference relation over the objects in our data. While we are – at present – agnostic about the origin of this preference relation, from an applied perspective it might represent an external sentiment that one object is more *typical* than another. Partial implications between attribute sets can then be captured through the idea that the correspondence holds for *preferred* objects [10,13].

Once we have a preference relation over the set of objects, and if we regard the relation as representative of the typicality of objects, we can introduce a notion of a *typical concept*. The motivation for typical concepts, and what they mean, rests on the notion that, given some set of attributes, we want a coherent way to derive a concept which has the attributes we *typically* expect. Continuing with the example from before, we might expect that, while the intent of the “bird concept” derived from the attribute `bird` does not include the attribute `flies`, it may be useful to have a notion of a typical concept which includes the attributes of prototypical birds.

The remainder of the paper is structured as follows: Section 2 provides an account of some basic notions in FCA. Then, Section 3 provides an introduction to classical notions of consequence, non-monotonic reasoning, and an intuition for how these translate to FCA. We introduce the announced extensions to FCA in Section 4 and discuss the semantics and properties of the consequence relation they allow. In Section 5 we develop the notion of typical concepts, derived from the preferential view of objects. In Section 6 we discuss some attempts to bring KLM-style defeasibility into other logical systems. We also discuss existing work within FCA which proposes to introduce more expressive notions of concepts.

2 Formal Concept Analysis

Two fundamental notions in FCA are *formal context* and *formal concept*. A formal context has a set-theoretic definition, where three sets of objects, attributes, and a binary relation are defined. However, in reasonably sized instances it can be represented as a cross-table, where rows represent the objects, and columns are the attributes. Naturally, the presence of an ‘×’ at a position in the table indicates that the respective object has the attribute from the respective column.

Definition 1. A formal context is a triple (G, M, I) where G refers to a finite, non-empty set of objects, M to a finite, non-empty set of attributes, and $I \subseteq G \times M$ is an incidence relation. For $(g, m) \in I$, we say object g has attribute m , otherwise g does not have attribute m .

Example 1. Figure 1 shows a cross-table representation of a formal context, describing different animal species and some of their attributes.

	northern	southern	flies	antarctic	bird
orca	×	×		×	
duck	×	×	×		×
robin	×	×	×		×
penguin		×		×	×

Fig. 1: A context of animal species and some of their attributes. For example, Orcas are found in the northern hemisphere, southern hemisphere, and the antarctic.

In the build-up to defining formal concepts, one introduces two operators which define an order-reversing Galois connection between the power sets $\mathcal{P}(G)$ and $\mathcal{P}(M)$. These two operators, which share the same notation, describe the derivation from a set of objects to the shared attributes, and vice versa [8,9].

Definition 2. In a formal context, $\mathbb{K} = (G, M, I)$, the derivation operator $(\cdot)'$ is defined for sets $A \subseteq G$ and $B \subseteq M$ as:

$$A' := \{m \in M \mid \forall g \in A, (g, m) \in I\}$$

$$B' := \{g \in G \mid \forall m \in B, (g, m) \in I\}$$

For a set $A \subseteq G$ of objects (resp. attributes), A' is just a set of attributes (resp. objects), and so A'' would be a set of objects (resp. attributes). The double application of derivation operators is in fact a closure operator, meaning it is extensive, idempotent, and monotonic.

Concepts are represented as a pairs of sets of objects and attributes; as one might expect, these sets are not arbitrary. In fact, they completely determine one another. Conventionally, we might denote a concept as (A, B) with $A \subseteq G, B \subseteq M$. This may be a useful notation to express the demarcation between the extent and intent, but it is not strictly necessary, for, (A, A') and (B', B) would refer to the same concept.

Definition 3. Let $\mathbb{K} = (G, M, I)$ be a formal context. A pair (A, B) , where $A \subseteq G$ and $B \subseteq M$, is a formal concept iff $A' = B$ and $B' = A$.

In fact, given an arbitrary set of objects, $A \subseteq G$, one can entirely describe a concept by (A'', A') – the same principle holds for sets of attributes [9]. If we refer back to Example 1, starting from two sets of objects and attributes, respectively, $\{\text{duck}, \text{penguin}\}$ and $\{\text{antarctic}\}$, we find the following concepts:

$$\begin{aligned} (\{\text{duck}, \text{penguin}\}'', \{\text{duck}, \text{penguin}\}') &= (\{\text{duck}, \text{robin}, \text{penguin}\}, \{\text{southern}, \text{bird}\}) \\ (\{\text{antarctic}\}', \{\text{antarctic}\}'') &= (\{\text{penguin}, \text{orca}\}, \{\text{southern}, \text{antarctic}\}) \end{aligned}$$

These two concepts suggest that one can introduce a partial ordering \leq to concepts (induced by the subset relationship on their extents), establishing the notion of sub- and super-concepts. Equipped with this ordering, the set of all concepts of a formal context \mathbb{K} , denoted by $\mathfrak{B}(\mathbb{K})$, forms a complete lattice $\underline{\mathfrak{B}}(\mathbb{K}) := (\mathfrak{B}(\mathbb{K}), \leq)$, referred to as a concept lattice [8,9]. Figure 4 provides an example of a concept lattice. Attributes connected to a concept from above (inclusive) are part of that concept's intent, while objects connected from below (inclusive) are in the extent. For a set of concepts, $\{(A_t, B_t) \mid t \in T\} \subseteq \mathfrak{B}(\mathbb{K})$, the supremum (*smallest super-concept*), and infimum (*greatest sub-concept*) are given by

$$\begin{aligned} \bigvee_{t \in T} (A_t, B_t) &:= ((\bigcup_{t \in T} A_t)'', \bigcap_{t \in T} B_t) \text{ as well as} \\ \bigwedge_{t \in T} (A_t, B_t) &:= (\bigcap_{t \in T} A_t, (\bigcup_{t \in T} B_t)''). \end{aligned}$$

2.1 Implications

A significant aspect of FCA involves implications between sets of attributes. These express the notion that certain attributes indicate the presence of others, in all objects of a formal context [8,19].

Definition 4. Let $\mathbb{K} = (G, M, I)$ be a formal context. An implication over M is an expression of the form $A \rightarrow B$ with $A, B \subseteq M$. We say that the implication is respected by another set $C \subseteq M$, iff $A \not\subseteq C$ or $B \subseteq C$. \mathbb{K} respects the implication, written $\mathbb{K} \models A \rightarrow B$, iff for every $g \in G$, g' respects $A \rightarrow B$.

It is useful to note that a formal context $\mathbb{K} = (G, M, I)$ respects an implication $A \rightarrow B$ over M exactly if $B \subseteq A''$, which, in turn, is equivalent to $A' \subseteq B'$.

3 Logical Consequence and Nonmonotonic Reasoning

3.1 Logical Consequence

A classical notion of logical consequence describes the circumstances under which one sentence is said to follow (logically) from another [6]. In a model-theoretic view, it is given as follows.

Definition 5. For two sentences α, β in the language \mathcal{L} , we say that β is a logical consequence of α , expressed as $\alpha \models \beta$, iff for every valuation $u \in \mathcal{U}$ where $u \models \alpha$ then also $u \models \beta$.

This definition can easily be generalised to the notion of *entailment*: given a set of facts it would be useful to know *what else* we can know, or, what is entailed [10].

Definition 6. Given a set of sentences \mathcal{KB} , another sentence α is entailed by \mathcal{KB} iff for every valuation $u \in \mathcal{U}$ with $u \models \mathcal{KB}$ also $u \models \alpha$ holds. This is expressed as $\mathcal{KB} \models \alpha$.

A consequence operator, \mathcal{Cn} , provides a general way to derive all sentences that should follow from a set of sentences \mathcal{KB} , under *some* notion of logical consequence [10]. Using classical entailment from Definition 6, we could have $\mathcal{Cn}(\mathcal{KB}) := \{\alpha \mid \mathcal{KB} \models \alpha\}$ [5,10]. A Tarskian consequence operator satisfies the following properties (which describe a closure operator).

$$\text{Monotonocity:} \quad \text{if } \Gamma \subseteq \Gamma' \text{ then } \mathcal{Cn}(\Gamma) \subseteq \mathcal{Cn}(\Gamma') \quad (1)$$

$$\text{Idempotence:} \quad \mathcal{Cn}(\Gamma) = \mathcal{Cn}(\mathcal{Cn}(\Gamma)) \quad (2)$$

$$\text{Inclusion:} \quad \Gamma \subseteq \mathcal{Cn}(\Gamma) \quad (3)$$

3.2 Consequence Relations

A more abstract notion of consequence is a *consequence relation*. This is a set of pairs, $\{(\Gamma_1, \gamma_1), \dots, (\Gamma_n, \gamma_n), \dots\}$, where it is typical to allow Γ_i to represent a set of formulae, and γ_i to represent a single formula in the language. The inclusion of a pair (Γ_i, γ_i) in the consequence relation—denoted as $\Gamma_i \vdash \gamma_i$ —means that γ_i can be inferred from Γ_i [7,10].

Consequence relations may be characterised by the properties they satisfy and, as such, correspond to a certain kind of reasoning—an example of this is found in Definition 7 [7]. Conversely, an algorithmic description of a kind of reasoning may in turn give rise to a consequence relation, the properties of which provide strong intuition for the implicit pattern of reasoning.

3.3 Non-monotonic Reasoning

Non-monotonic reasoning is concerned with developing formal reasoning process in which a conclusion drawn under a premise can be withdrawn under the addition of another premise [13]. The justifications for why we may want to reason non-monotonically are easy to accept – it is quite obvious that when humans reason we do so under the implicit assumption that, upon receiving new information we can change our mind. Without this assumption, navigating life would be very difficult; we would hesitate to come to any conclusion out of fear it may be the wrong one. Moreover, we frequently make statements for which we know exist exceptions, under the assumption that upon encountering an exception it would be treated as such [16].

In the setting of propositional logic, the problem is often introduced in the following way [10,13,21]: we want to accept that "penguins are birds", "birds usually fly" and "penguins do not fly". However, when we translate this into classical logic with the

propositions $\text{penguin} \rightarrow \text{bird}$, $\text{bird} \rightarrow \text{fly}$ and $\text{penguin} \rightarrow \neg\text{fly}$, we are forced to conclude $\neg\text{penguin}$. With respect to the attribute logic of FCA, we find a similar issue. Consider the context in Example 1. We do not have $\{\text{bird}\} \rightarrow \{\text{flies}\}$ as a valid implication, since penguins are an object with bird but not flies . This illustrates the final point made in the last paragraph – that monotonicity prevents us from making statements with known exceptions.

In the realm of concepts, we encounter another instance of this problem: the concept determined by bird is not a sub-concept of the concept determined by flies . Again, we lack the expressivity for the idea that *typical* birds are flying animals. In Section 4 and Section 5 we propose a solution to this lack of expressivity drawing inspiration from [12,13].

3.4 Rational Consequence Relations

The style of non-monotonic reasoning we aim to develop in FCA is that of rational consequence relations. A rational consequence relation, \sim_R , is based on preferential logics, where an order (or ranking) over valuations conveys the notion that certain valuations are preferred to others [12,14,20]. The semantics describe a notion of consequence where, if α, β are formulae in the language then β is a rational consequence of α , $\alpha \sim_R \beta$, iff β is true in all the *most preferred* models of α .

The KLM postulates, introduced in [12,14], with the inclusion of *Rational Monotonicity* are a set of inference rules that characterise a rational consequence relation.

Definition 7. A consequence relation, \sim , constitutes a rational consequence relation iff it satisfies Reflexivity, Left Logical Equivalence (LLE), Right Weakening (RW), Cut, And, Or, Cautious Monotonicity and Rational Monotonicity.

Reflexivity $A \sim A$ Reflexivity is a somewhat basic notion of any notion of consequence; it essentially prevents a self-defeating pattern of reasoning where given something, you conclude as a consequence not that thing.

Left Logical Equivalence (LLE) $\frac{\models A \equiv B, A \sim C}{B \sim C}$ LLE enforces the notion that two things that are equivalent (under a coherent notion of equivalence) should have the exact same consequences.

Right Weakening (RW) $\frac{\models A \rightarrow B, C \sim A}{C \sim B}$ RW states that the consequence of a classical implication, which has a defeasible consequence as a premise, can itself be derived from the original defeasible implication.

Cut $\frac{A \wedge B \sim C, A \sim B}{A \sim C}$ Cut allows us to use existing defeasible consequences in the premise of a new defeasible implication. However, the new implication is subject to failure should the defeasible conclusion in its premise be retracted.

Or $\frac{A \sim C, B \sim C}{A \vee B \sim C}$ Or is that given distinct premises with a common defeasible consequence, we should be able to draw this conclusion from the disjunction of the premises – that is, we need not explicitly know which one is true.

And $\frac{A \sim B, A \sim C}{A \sim B \wedge C}$ And tells us that two consequences can be concluded at the same time.

Cautious Monotonicity (CM) $\frac{A \sim B, A \sim C}{A \wedge B \sim C}$ CM ensures that adding a defeasible consequence that could already have been derived from our premises should never invalidate another defeasible consequence that could be derived from our original premises.

Rational Monotonicity (RM) $\frac{A \sim B, A \not\sim \neg C}{A \wedge C \sim B}$ RM expresses the notion, similar to conflicts in Default Logic, that only when a premise that was expected to be false is added, should we retract a rational consequence. In another sense, we can assume that when new information that does not explicitly contradict existing knowledge, we can retain existing conclusions [14,10].

4 Introducing Nonmonotonicity in FCA

4.1 Extended Formal Context

As a precursor to the non-monotonic conditionals in Subsection 4.3, we introduce the notion of an extended formal context.

Definition 8. Let $\mathbb{K} = (G, M, I)$ be a formal context. An extended formal context simply adds a partial order over G . We denote this as a quadruple, $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$.

The partial ordering is intended to convey a preference relation between objects. That is, for two objects, $g, h \in G$, if $g \preceq h$, then we say that g is more preferred, or more typical, in comparison to h . For example, $\text{robin} \preceq \text{penguin}$ tells us a robin is “more typical” than a penguin in our context.

4.2 Minimisation

The preference relation between objects provided by an extended formal context enables us to define how the derivation operators might behave when restricted to only care about minimal objects. We formalise this notion as a *minimised derivation*.

Definition 9. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context, and let $B \subseteq M$. Then the minimised derivation, \underline{B}' , of B , is the set

$$\underline{B}' := \{g \in B' \mid \nexists h \in B' \text{ such that } h \prec g\}$$

As a reminder, for a set of attributes, $B \subseteq M$, B' provides us with all the objects which have all attributes of B . Following this, B'' extends B by including all other attributes which are shared by those objects which share attributes B .

If \underline{B}' describes the process of going from a set of attributes to the minimal objects which share B , a reasonable question might concern the procedure for returning to a “closed” set of attributes: presumably, the set of attributes common to those objects described by \underline{B}' . This is, in fact, just the composition of Definition 9 and Definition 2, taking the form $(\underline{B}')'$ - henceforth referred to as a *minimised-return operation*.

Theorem 1. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context, the minimised-return operator, $(\cdot)'$, applied to a set of attributes is nonmonotonic, extensive, and idempotent.

Proof. If the minimised-return operator were monotonic, for the sets $A, B \subseteq M$, $A \subseteq B$ would imply $(A)' \subseteq (B)'$. Refer back to the formal context in Example 1 extended with the partial order $\text{robin} \preceq \text{penguin}$. Let $A := \{\text{bird}\}$ and $B := \{\text{bird}, \text{antarctic}\}$. Clearly, $A \subseteq B$, however, application of the minimised-return operator gives $(A)' = \{\text{bird}, \text{flies}, \text{southern}, \text{northern}\}$ and $(B)' = \{\text{bird}, \text{antarctic}, \text{southern}\}$, and so $(A)' \not\subseteq (B)'$.

To show that the minimised-return operator is extensive, it should hold that for a set $A \subseteq M$, $A \subseteq (A)'$. We already know that in the classical case, $A \subseteq A''$. So, for all objects $g \in A'$, it is the case that $A \subseteq g'$. It is also clear that the minimised derivation of a set of attributes is a subset of the classical derivation – that is, $\underline{A}' \subseteq A'$. This means that for all objects $h \in \underline{A}'$, $A \subseteq h'$. It should then be straightforward to see that A is a subset of the set given by $\overline{\bigcap \{g' \mid g \in \underline{A}'\}}$, which is equivalent to $A \subseteq (\underline{A})'$.

The final task is to show that the minimised-return operator is idempotent. Given a set $A \subseteq M$, let $C := (\underline{A})'$. It is enough to show that $\underline{A}' = \underline{C}'$. C is the set of attributes such that all objects, $g \in \underline{A}'$, have C in their intent. It is also clear, since all objects in \underline{A}' have A in their intent, that $A \subseteq C$. Then, for $h \in \underline{C}'$ it follows that h has attributes C , and thus A , in its intent. Since h is minimal w.r.t. \underline{C} , and $g \in \underline{A}'$, it cannot be the case that $g \preceq h$. Conversely, since g is minimal w.r.t. A , and $h \in \underline{A}'$, it cannot be that $h \preceq g$. Then, since both g and h are elements of A' and \underline{C}' , and are arbitrary, the sets \underline{A}' and \underline{C}' refer to the same objects. Then, applying a derivation to the same sets would of course yield the same result, and so the minimised-return operator is idempotent. \square

4.3 Non-monotonic Conditional through Minimisation

From the discussion in Subsection 2.1 it is clear that the semantics of classical implications in FCA can be defined through applications of closure operators on sets of attributes. With the introduction of the minimised-return operator, one might expect to find a different notion of consequence.

The minimised-return operator still refers to the same kind of process as the closure-operator, going from a set of attributes to a set of objects, and then back to a set of attributes. The similarity enable us to define a semantics which is only a minor departure from the original implication. For now we will use *non-monotonic conditional* to refer to the minimised-derivation based implications, and denote this by \rightsquigarrow .

Definition 10. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context, with $A, B \subseteq M$. \mathbb{K}^{\preceq} respects a non-monotonic conditional $\mathbb{K}^{\preceq} \models A \rightsquigarrow B$ iff $B \subseteq (\underline{A})'$ which is equivalent to $\underline{A}' \subseteq B'$.

Although we have yet to prove the characterisation, the non-monotonic conditional from above describes a notion of consequence with strong semblance to the preferential consequence relations discussed in Subsection 3.4.

	warrior	demigod	hero
Jason	×		×
Achilles	×	×	×
Minos		×	

Fig. 2: Extended formal context of figures in Greek mythology

Example 2. Let the cross table below represent an extended formal context, where the partial order over objects is: $\text{Achilles} \preceq \text{Jason}, \text{Minos}$.

In Greek mythology there is a strong correspondence between figures who are heroes, and demigods. So much so that there is debate as to whether *hero* and *demigod* actually have distinct meanings. Thus, with some creative licence, we may wish to express the notion that hero *usually implies* demigod. The classical implication, of course, fails to express this - as Jason is a counter-example. Instead, we express this notion as the non-monotonic conditional $\text{hero} \rightsquigarrow \text{demigod}$.

In our definition of classical implications, see Definition 4, we first introduced \rightarrow as an object-level operator, saying that a set $C \subseteq M$ respects an implication $A \rightarrow B$ iff $A \not\subseteq C$ or $B \subseteq C$ (we can think of C as an object intent). However, most of the time when we talk about implications in FCA, we really refer to the meta level notion analogous to entailment that $K \models A \rightarrow B$.

For non-monotonic conditionals it makes even less sense to speak about the object level, we are interested in what our ordering allows us to conclude – and the ordering is explicitly on the meta-level. Consequently, whenever we speak about a non-monotonic conditional, we are speaking about the implication over the entire formal context.

4.4 Non-monotonic Conditionals and the KLM Postulates

We now present an argument that the consequence relation given by the non-monotonic conditional described in the previous sub-section satisfies all the KLM postulates required to be a characterisation of rational consequence relations. We should remind ourselves that the postulates are usually described in a truth-theoretic language, such as propositional logic. A consequence being that some initial work needs to be translated into the attribute-logic formalism considered here.

For *Reflexivity* and *RW*, the translation is obvious and doesn't require additional intuition. All that is required for *LLE* is to describe what a notion of equivalence between sets of attributes means. To this end, we say that in a formal context two attribute sets are equivalent, $A \equiv B$, iff $A' = B'$. *Cut* says that if $\alpha \wedge \beta \vdash \gamma$ and $\alpha \sim \beta$ then $\alpha \vdash \gamma$. The conjunction of two formulae in a truth-theoretic logic requires satisfaction of each formula. In attribute logic, the equivalent notion is given by the union of two attribute sets, satisfied by objects with *all* attributes from both sets. The conjunction in *CM* and *And* is re-phrased in the same way.

It is not immediately clear what *RM* might mean, as it is uncommon to talk about implications with a negation. We use the following definitions to make this notion explicit.

Definition 11. Let $\mathbb{K} = (G, M, I)$ be a formal context and $A \rightarrow \neg B$ an implication over M . We say that a set $C \subseteq M$ respects the implication $A \rightarrow \neg B$ iff $A \not\subseteq C$ or $B \not\subseteq C$. \mathbb{K} respects the implication, $\mathbb{K} \models A \rightarrow \neg B$, iff for every object g , g' respects $A \rightarrow \neg B$. Equivalently, $A' \cap B' = \emptyset$.

Definition 12. An extended formal context, \mathbb{K}^{\preceq} , respects the non-monotonic conditional $A \rightsquigarrow \neg B$ iff, for every object g in \underline{A}' , $B \not\subseteq g'$. Equivalently, $\mathbb{K}^{\preceq} \models A \rightsquigarrow \neg B$ iff $\underline{A}' \cap B' = \emptyset$.

The intuition is that one set of attributes corresponds to the absence of another if there are no objects which have both sets in their intent. On the object level of a formal context, for an object $g \in G$, if $g \not\models A \rightarrow B$ then $g \models A \rightarrow \neg B$. We do not, however, have that if $\mathbb{K} \not\models A \rightarrow B$ then $\mathbb{K} \models A \rightarrow \neg B$. Translating the intuition of $A \rightarrow \neg B$ to $A \rightsquigarrow \neg B$ is straightforward. The non-monotonic conditional, $A \rightsquigarrow \neg B$, concerns only objects $g \in \underline{A}'$, which, by definition of the minimised derivation operator, must have A in their intent. Consequently, we need only show that g does not have B in its intent.

We have yet to address the *Or* postulate. The difficulty here is that attribute logic does not have (nor is it intuitive to introduce) a notion of disjunction [15]. As such, we recognise a departure from the KLM framework, and omit the *Or* postulate entirely. Consequence relations that satisfy only *Reflexivity*, *LLE*, *RW*, *Cut*, and *CM* are considered *cumulative consequence relations* [13]. However, while we do not have *Or*, we show that we still have *Rational Monotonicity*, and continue to regard our relation as a characterisation of *rational consequence**.⁴

Theorem 2. *The consequence relation derived from \rightsquigarrow satisfies Reflexivity, Left Logical Equivalence, Right Weakening, Cut, And, Cautious Monotonicity, and Rational Monotonicity, and thus, constitutes a rational consequence relation*.*⁵

This gives us a notion of non-monotonic consequence that corresponds to what we may think *logical* non-monotonic reasoning should look like, as described by rational consequence. We are able to use intermediary results as the basis for more conclusions through the use of *Cut* and *CM* [10,21]. *RM* is a useful property in the sense that it reduces the work one has to do - unless we have some consequence that is contradictory to what we expect, we can retain existing conclusions. The notion of consequence also preserves classical implications in FCA, and so it is strictly more expressive than regular implications in FCA [21].

Now is a good place to remind ourselves of what we gain from introducing non-monotonic conditionals to FCA. If we accept rational consequence relations as a good account of non-monotonic reasoning, and that we can impose a preference relation on objects in a formal context, non-monotonic conditionals provide FCA with the expressivity to reason about what may be the *typical* case.

⁴ The absence of the *Or* postulate is indicated by an asterisk

⁵ The proofs for the non-monotonic conditional satisfying the KLM-postulates* can be found in the Arxiv submission with a matching title

5 Typical Concepts

The partial ordering over objects was introduced as a means to developing a non-monotonic consequence relation. We now show how this ordering, under some constraint, gives rise to a notion of typical concepts.

5.1 Naive Typical Concepts

Definition 13. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context. Then, for a set $A \subseteq M$, we define a typical concept as a concept of the form

$$\left((\underline{A}')'', (\underline{A}')' \right).$$

The set of all typical concepts of an extended formal context is denoted $\mathfrak{T}(\mathbb{K}^{\preceq})$.

In this case, the intent of a typical concept is the set of attributes common to the minimal objects which have A in their intent. The extent is then all objects which have this extended set of attributes. Referring back to Example 1, we add the partial order $\text{duck} \preceq \text{penguin}$ and $\text{robin} \preceq \text{penguin}$, expressing the sentiment that ducks and robins are more *typical* than penguins, but incomparable to orcas. Given only the attribute bird , $(\{\text{duck}, \text{robin}\}, \{\text{northern}, \text{southern}, \text{flies}, \text{bird}\})$ is the derived typical concept. Without this notion of a typical concept, the derived concept would be $(\{\text{duck}, \text{robin}, \text{penguin}\}, \{\text{southern}, \text{bird}\})$. Of course, both of these concepts exist in the classical concept lattice; however, typical concepts provide an instrument to, given only the condition of bird , to arrive at a concept which we might consider a more natural characterisation of things that are birds. We should remind ourselves that, while this example is picked to match real-world expectations, we should think of “expectations” as being an expression of the preference relation on objects.

Perhaps un-intuitively, this formulation allows for a typical concept to contain non-typical objects in its extent, as long as they share all the attributes of their typical counterparts. We should recognise, however, that they do not contribute to describing the intent, they just happen to be consistent with it. Although we ascribe specific meaning to \underline{A}' , it can be considered to be an arbitrary set of objects. Then, per Definition 3, an arbitrary typical concept, $((\underline{X}')'', (\underline{X}')')$ where $X \subseteq M$, is always equivalent to some formal concept. As such, we can define a map from the set of concepts to the set of typical concepts.

Definition 14. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context, then $\mathfrak{B}(\mathbb{K}^{\preceq})$ denote the set of concepts of \mathbb{K}^{\preceq} , and $\mathfrak{T}(\mathbb{K}^{\preceq})$ denote the set of typical concepts of \mathbb{K}^{\preceq} . We define a map:

$$\varphi : \mathfrak{B}(\mathbb{K}) \mapsto \mathfrak{T}(\mathbb{K}^{\preceq})$$

For a concept $(A, B) \in \mathfrak{B}(\mathbb{K}^{\preceq})$, φ is defined as:

$$\varphi(A, B) := ((\underline{B}')'', (\underline{B}')')$$




Fig. 3: Classical concept lattices, with the typical concepts derived from the respective ordering marked as grey

The mapping function inherits idempotency from the minimised-return operator. For a typical concept $(A, B) \in \mathfrak{T}(\mathbb{K}^{\preceq})$, there must exist some set $C \subseteq M$ such that $(\underline{C}')' = B$. It follows that $((\underline{C}')')' = (\underline{C}')' = B$. Consequently, $\varphi(A, B) = (A, B)$, or, typical concepts map to themselves. Differently put, $\varphi(A, B) = \varphi(\varphi(A, B))$. We also note that by the extensive property of minimized-derivation, for any concept (A, B) , we have $(\underline{B}')' \supseteq B$ and therefore $\varphi(A, B) = ((\underline{B}')'', (\underline{B}')') \leq (A, B)$. That is, φ is *anti-extensive*.

Theorem 3. *In an extended formal context, $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$, let $(A, B) \in \mathfrak{B}(\mathbb{K}^{\preceq})$, then $(A, B) \in \mathfrak{T}(\mathbb{K}^{\preceq})$ iff $(\underline{B}')'' = A$.*

Proof. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context with a concept, $(A, B) \in \mathfrak{B}(\mathbb{K}^{\preceq})$. Suppose (A, B) were a typical concept, then there exists a set $C \subseteq M$ such that $A = (\underline{C}')''$ and $B = (\underline{C}')'$, per the definition of a typical concept. Through the properties of the minimised-return operator in Theorem 1 we have that $\underline{C}' = (\underline{C}')'$, equivalently, $\underline{C}' = \underline{B}'$. Then, obviously, $(\underline{C}')'' = (\underline{B}')''$. \square

A next point of departure might be to investigate the structure the set of typical concepts. Certainly, we have that $\mathfrak{T}(\mathbb{K}^{\preceq}) \subseteq \mathfrak{B}(\mathbb{K}^{\preceq})$. Then in particular, we want to investigate whether the subset of “typical” objects forms a sub-lattice, if we can form a sub-lattice of only those “typical” concepts, we are able to use all the same algebraic tools to analyze the data in our context “prototypically”, as opposed to the classical analysis which relies on strict implications and concepts. Using the same approach to ordering typical concepts as we did for formal concepts (see Section 2), two examples make it clear that the current formulation unfortunately does not guarantee that the pair $(\mathfrak{T}(\mathbb{K}^{\preceq}), \leq)$, henceforth denoted $\underline{\mathfrak{T}}(\mathbb{K}^{\preceq})$, will form any type of lattice. Figure 3a is the concept lattice of some formal context, \mathbb{K} . If one were to extend \mathbb{K} with the partial order $c \preceq b$ and $a, d \preceq c$, the resultant set of typical concepts would be the four concepts marked with grey

nodes. We pay special attention to the concepts

$$\begin{aligned} C_1 &:= \left((\underline{X}')'', (\underline{X}')' \right) = \left((a)'', (a)' \right) = \left(\{a, b, c\}, \{X\} \right) \\ C_2 &:= \left((\underline{Y}')'', (\underline{Y}')' \right) = \left((d)'', (d)' \right) = \left(\{b, c, d\}, \{Y\} \right) \end{aligned}$$

The greatest common sub-concept of these two concepts is given by

$$C_1 \wedge C_2 = \left(\{a, b, c\} \cap \{b, c, d\}, (\{X\} \cup \{Y\})'' \right) = \left(\{b, c\}, \{X, Y\} \right)$$

However, from Theorem 3, it is clear that this is not a typical concept, as $(\{X, Y\})'' = \{c\} \neq \{b, c\}$. This result tells us that the subset of typical concepts is not closed under meets. If we shift attention to Figure 3b, it becomes clear that it is not closed under joins either, nor does it guarantee a lattice with respect to the concept-lattice ordering. For the typical concepts

$$\begin{aligned} C_1 &:= \left((\underline{X}')'', (\underline{X}')' \right) = \left((a)'', (a)' \right) = \left(\{a, b\}, \{X\} \right) \\ C_2 &:= \left((\underline{\{Y, Z\}})'', (\underline{\{Y, Z\}})' \right) = \left((b, c)'', (b, c)' \right) = \left(\{b, c\}, \{Y, Z\} \right) \end{aligned}$$

the least common super-concept is given by

$$C_1 \vee C_2 = \left((\{a, c\} \cup \{b, c\})'', \{X, Z\} \cap \{Y, Z\} \right) = \left(\{b, c\}, \{Z\} \right).$$

For $C_1 \vee C_2$, the closure of the minimised-derivation of the intent, $(Z)'' = \{a, c\}$, does not equal the concept extent. So, we do not have that typical concepts are closed under joins. To make matters worse, the top element of the lattice is the concept $(\{a, b, c, d\}, \emptyset)$, but not a typical concept, as $(\emptyset)'' = (a)'' = \{a, c\}$, and so we do not have a lattice structure of typical concepts. In fact, we do not even have an upper bound to the set of typical concepts in Figure 3b.

5.2 Restriction on the Partial Order

One cause of the aforementioned problem lies in the fact that we allow non-minimal objects to be included in typical concepts. This allows for rankings where non-minimal objects, which should be in the meet or join of two concepts, are *lost*. As a solution to this, we may restrict the partial orders to prevent non-minimal objects from being included in typical concepts.

Definition 15. Let $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$ be an extended formal context. A valid partial order over objects is one which for all $A \subseteq M$, $\underline{A}' = (\underline{A}')''$.

Consider a set of attributes $X \subseteq M$ and then take $Y \subseteq M$ to be the set of all attributes common to the minimal objects which satisfy X – of course, Y would be the set X with additional attributes. The restriction to orders \preceq that are valid ensures that any object that has all the attributes in Y must itself be minimal w.r.t. X . Put differently, Y serves as a total characterisation of the minimal objects satisfying attributes X , if an object matches this characterisation, then it should have been one of these minimal objects.

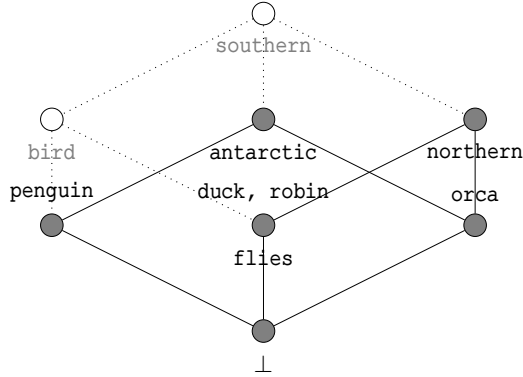


Fig. 4: The concept lattice for Example 1. Given the preference order $\text{duck} \preceq \text{penguin}$ and $\text{robin} \preceq \text{penguin}$, the dark grey concepts show the structure of typical concepts

If we accept this restriction on the partial order, the set of typical concepts preserves meets, that is, the greatest common sub-concept of two typical concepts is itself a typical concept. Before this is shown, Definition 13, Definition 14, and Theorem 3 might be amended to reflect that the closure of a set \underline{A}' , where $A \subseteq M$, is superfluous.

Theorem 4. *For an extended formal context, $\mathbb{K}^{\preceq} = (G, M, I, \preceq)$, with a restriction on the partial-order such that for any $A \subseteq M$, $\underline{A}' = (\underline{A}')''$, the greatest common sub-concept of two typical concepts is itself a typical concept. Hence, the subset of typical concepts form a \wedge -subsemilattice of the concept lattice.*

Proof. We note that it is a sufficient condition to show that $\underline{A}' \cap \underline{B}' = \underline{C}'$ for some $C \subseteq M$. We claim that we can choose $C = (\underline{A}')' \cup (\underline{B}')'$. Note

$$\underline{A}' \cap \underline{B}' = (\underline{A}')'' \cap (\underline{B}')'' = ((\underline{A}')' \cup (\underline{B}')')',$$

where the first equality follows from our assumed condition, and the second from our preliminary lemma for Galois connections. Then $\underline{A}' \cap \underline{B}' = \underline{((\underline{A}')' \cup (\underline{B}')')}$. Note that

$$\underline{A}' \cap \underline{B}' = \{g \in \underline{A}' \cap \underline{B}' \mid \nexists h \in \underline{A}' \cap \underline{B}' \text{ such that } h < g\}.$$

Therefore, if $g \in \underline{A}' \cap \underline{B}'$, since g is a minimal member of \underline{A}' and of \underline{B}' there is no h in \underline{A}' or in \underline{B}' with $h < g$ by definition and so $g \in \underline{A}' \cap \underline{B}'$. Therefore $\underline{A}' \cap \underline{B}' \subseteq \underline{A}' \cap \underline{B}'$, and clearly $\underline{A}' \cap \underline{B}' \subseteq \underline{A}' \cap \underline{B}'$. If the intersection has no elements $\underline{A}' \cap \underline{B}' = \emptyset = \emptyset$. Hence $\underline{A}' \cap \underline{B}' = \underline{A}' \cap \underline{B}' = \underline{((\underline{A}')' \cup (\underline{B}')')}$ and we are done. \square

It is also worth noting here that the bottom element of the concept lattice is always a typical concept since $\varphi(\perp) \leq \perp$, by anti-extensivity, and so $\varphi(\perp) = \perp$. This result allows us to begin characterizing the algebraic structure of those concepts considered “typical” in an extended formal context. However, there is still much structure that this restriction on the order does not account for. In particular, we do not have that $\mathfrak{T}(\mathbb{K}^{\preceq})$ preserves least common super-concepts – Figure 3b is an example of a concept lattice

whose order satisfies the condition in Definition 15, but does not have a \top -concept. That being said, we are still able to preserve some structure using a somewhat intuitive restriction on the order in our extended formal context.

6 Related Work

There is a large corpus of work on developing KLM-style notions of non-monotonic reasoning outside of the standard of propositional logic. [4], and more extensively [18], introduce KLM-style defeasible reasoning to datalog, which can be considered a fragment of first-order logic used for database queries. [2,1] are recent efforts to introduce a notion of defeasible subsumption to description logics (DL) that follows the KLM properties. Remaining in the realm of DL, [17] is a general characterisation of to KLM in DL \mathcal{ALC} . [3] investigates the KLM framework in the context of defeasible modalities, introducing new modal operators for *defeasible necessity* and *distinct possibility*.

With respect to work that extends our view of concepts in FCA, [11] introduces *Rough Concept Analysis*, a merging of rough set theory and FCA that uses equivalence classes on objects to define upper and lower-bound approximations of concepts. [22] investigates an expansion which enables “rough concepts” to be defined not only by objects.

7 Conclusions and Future Work

By extending a formal context with a preference relation on the objects we have introduced a non-monotonic variant of implication between attribute sets which characterise a rational consequence relation*. This strictly increases the expressivity of the attribute logic of FCA by creating a notion of non-monotonicity that corresponds to the KLM view of how a *logical* non-monotonic system should behave. In terms of FCA, we have introduced a way of discovering and representing relationships between attribute sets that tolerates exceptions, and is capable of representing what our data (formal context) shows in the *typical* case.

With a slight restriction on what constitutes a *valid* preference relation over objects, we presented a formalisation of typical concepts which has its foundations in a KLM-style typicality. We were able to show that with this notion we could create a structure of typical concepts that at least preserved the sub-concept relation, and as such is a *meet-subsemilattice* of the original concept lattice.

To our knowledge, the introduction of KLM-style typicality, and preferences over a formal context as a whole, presents a novel view on, and non-trivial extension to FCA.

This work represents the initial investigation of introducing KLM-style typicality into FCA. As such, we believe there is a considerable amount that remains to be looked at. Addressing concerns of existing work, we aim to find an approach to defining typical concepts which is closed under joins (super-concepts), essentially meaning that the structure of the set of typical concepts would be a sub-lattice of the original concept lattice. To this end, we could neatly “reduce” a concept lattice to its typical counterpart. In another branch of this work, we would like to investigate the relationship between the set of all non-monotonic conditionals and the typical concept lattice; determining if something analogous to a *canonical basis* can be found.

Acknowledgments. This work is a culmination of effort that took place when students at the University of Cape Town visited the Technical University of Dresden. We would like to express immense appreciation and gratitude to the School of Embedded Composite Artificial Intelligence (SECAI) – project 57616814 funded by BMBF (the Bundesministerium für Bildung und Forschung) and DAAD (German Academic Exchange Service) – for funding the visit, and to its members in Dresden who were so welcoming and engaging.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Britz, K., Casini, G., Meyer, T., Moodley, K., Sattler, U., Varzinczak, I.: Principles of klm-style defeasible description logics. *ACM Transactions on Computational Logic (TOCL)* **22**(1), 1–46 (2020)
2. Britz, K., Casini, G., Meyer, T., Varzinczak, I.: A klm perspective on defeasible reasoning for description logics. In: *Description Logic, Theory Combination, and All That: Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday*, pp. 147–173. Springer (2019)
3. Britz, K., Varzinczak, I.: From klm-style conditionals to defeasible modalities, and back. *Journal of Applied Non-Classical Logics* **28**(1), 92–121 (2018). <https://doi.org/10.1080/11663081.2017.1397325>, <https://doi.org/10.1080/11663081.2017.1397325>
4. Casini, G., Meyer, T., Paterson-Jones, G., Varzinczak, I.: Klm-style defeasibility for restricted first-order logic. In: *International Joint Conference on Rules and Reasoning*. pp. 81–94. Springer (2022)
5. Citkin, A., Muravitsky, A.: Consequence relations an introduction to the tarski-lindenbaum method (2021), <https://arxiv.org/abs/2106.10966>
6. Etchemendy, J.: Tarski on truth and logical consequence. *The Journal of Symbolic Logic* **53**(1), 51–79 (1988), <http://www.jstor.org/stable/2274427>
7. Gabbay, D.M.: A general theory of structured consequence relations. *Theoria: An International Journal for Theory, History and Foundations of Science* pp. 49–78 (1995)
8. Ganter, B., Obiedkov, S., Rudolph, S., Stumme, G.: *Conceptual exploration*. Springer (2016)
9. Ganter, B., Wille, R.: *Formal concept analysis: mathematical foundations*. Springer Science & Business Media (2012)
10. Kaliski, A.: An overview of KLM-style defeasible entailment. Master’s thesis, University of Cape Town (2020)
11. Kent, R.E.: Rough concept analysis: a synthesis of rough sets and formal concept analysis. *Fundamenta informaticae* **27**(2-3), 169–181 (1996)
12. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence* **44**(1-2), 167–207 (1990)
13. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic Reasoning, Preferential Models and Cumulative Logics (Feb 2002), <http://arxiv.org/abs/cs/0202021>, arXiv:cs/0202021
14. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? *Artificial Intelligence* **68**(2) (1994)
15. Lei, Y., Sui, Y., Cao, C.: The logical operations on relations, scales and concept lattices. In: *2009 Fifth International Conference on Semantics, Knowledge and Grid*. pp. 160–167 (2009). <https://doi.org/10.1109/SKG.2009.37>
16. Makinson, D.: How to go nonmonotonic. In: *Handbook of Philosophical Logic*, pp. 175–278. Springer (2005)

17. Moodley, K., Meyer, T., Sattler, U.: Practical defeasible reasoning for description logics. In: STAIRS 2014, pp. 191–200. IOS Press (2014)
18. Paterson-Jones, G.: KLM-Style Defeasible Reasoning for Datalog. Master’s thesis, University of Cape Town (2022)
19. Rudolph, S.: Relational Exploration - Combining Description Logics and Formal Concept Analysis for Knowledge Specification. Universitätsverlag Karlsruhe (December 2006)
20. Shoham, Y.: A semantical approach to nonmonotonic logics, p. 227–250. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1987)
21. Strasser, C., Antonelli, G.A.: Non-monotonic Logic. In: Zalta, E.N., Nodelman, U. (eds.) The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, Summer 2024 edn. (2024)
22. Yao, Y.: Rough-set concept analysis: interpreting rs-definable concepts based on ideas from formal concept analysis. *Information Sciences* **346**, 442–462 (2016)