Rational Concept Analysis

introducing the KLM framework to FCA

Lucas Carr
Department of Computer Science
University of Cape Town

March 10, 2025

Overview

1. Formal Concept Analysis

2. The KLM framework

3. Rational Concept Analysis

Introduction

- A lattice-theoretic framework for reasoning about concepts, objects, and properties.
- Its used for data-mining, knowledge discovery, ontologies, etc.
- Concepts are extracted from a data-structure called a formal context, a triple $\mathbb{K} = (G, M, I)$ of objects G, attributes M, and a relation $I \subseteq G \times M$ indicating when an object 'has' and attribute

$\mathbb{K}$	bird	mammal	reptile	aerial	aquatic	terrestrial	oviparous	viviparous	limbs
Chameleon			×			×		×	×
Crow	×			×			×		×
Ostrich	×					×	×		×
Penguin	×				×	×	×		×
Platypus		×			×	×	×		×
Snake			×			×	×		
Swallow	×			×			×		×
Whale		×			×			×	

Table: A formal context of animals and some of their properties

Derivation Operators

Definition (Derivation Operators)

In a formal context (G, M, I) the derivation operators are two maps $(\cdot)^{\uparrow} : G \mapsto M$ and $(\cdot)^{\downarrow} : M \mapsto G$ such that if $X \subseteq G$ and $Y \subseteq M$

$$X^{\uparrow} := \{m \in M \mid \forall g \in X : (g, m) \in I\}$$
$$Y^{\downarrow} := \{g \in G \mid \forall m \in Y : (g, m) \in I\}$$

- The derivation operators $(\uparrow, \downarrow)$ allow us to move from sets of objects to attributes, and vice verse
- Given some set of objects X, its derivation is the set of attributes which are common to all objects in X

Concepts

- A **concept** in a formal context is defined by its extension, and intension
- It is a pair (X, Y) of sets $X \subseteq G$ and $Y \subseteq M$ such that
  - Y consists of the attributes which all objects in X share
  - X consists of those objects which have all attributes in Y

Definition (Formal Concept)

In a formal context (G, M, I) a concept is an object-attribute pair (X, Y) such that $X^{\uparrow} = Y$ and $Y^{\downarrow} = X$.

- Two concepts $C_1 = (X_1, Y_1)$ and $C_2 = (X_2, Y_2)$ can be ordered such that $C_1 \leq C_2$ iff. $X_1 \subseteq X_2$ iff. $Y_2 \subseteq Y_1$. Then C_1 is a **subconcept** of C_2.
- This is a partial order, which corresponds to a complete lattice we call the **concept lattice**.

Concept Lattices

- We can represent the data in a context as a **concept lattice**
- Each node represents a concept
- Concepts inherit attributes from above, and contain objects from below

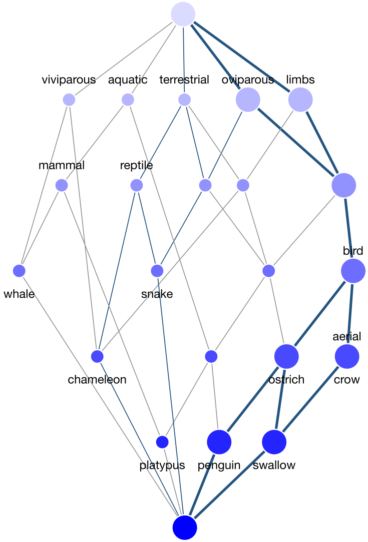

Implications I

Derivation operators also enable us to find correspondence between (sets of) attributes

1. $\mathbb{K} \models \texttt{bird} \rightarrow \texttt{oviparous}$
2. $\mathbb{K} \models \{\texttt{aquatic, terrestrial}\} \rightarrow \texttt{oviparous}$
3. $\mathbb{K} \not\models \texttt{bird} \rightarrow \texttt{aerial}$

Why?

1. All objects that are birds are oviparous, $\texttt{bird}^{\downarrow} \subseteq \texttt{oviparous}^{\downarrow}$
2. $\texttt{oviparous} \subseteq \{\texttt{aquatic, terrestrial}\}^{\downarrow\uparrow}$
3. penguin is a bird but not `aerial`

$\mathbb{K}$	bird	mammal	aerial	aquatic	terrestrial	oviparous	viviparous
Crow	×		×			×	
Ostrich	×				×	×	
Penguin	×			×	×	×	
Platypus		×		×	×	×	
Snake					×	×	
Swallow	×		×			×	
Whale		×		×			×

Table: Slightly reduced context

Implications II

- But implications which do not hold in a context might still represent useful information
- We might be dealing with error-prone data, or want to tolerate exceptions
- Existing work tackles this problem through use of association rules (with *support, confidence*)
- But these metrics are unintuitive to settle on, and difficult to explain
- We look for an approach which has a clearer pattern of reasoning, i.e., the KLM framework

KLM Framework

- Kraus, Lehmann, and Magidor (KLM) argue that a logic for non-monotonic should be able to express something like *birds usually fly (even though some may not)*

$$b \mid\!\sim f$$

- Expressed as a consequence relation satisfying certain properties (so called Rationality Postulates)

1. **Reflexivity:** $\dfrac{\alpha \rightarrow \alpha}{}$
2. **LLE:** $\dfrac{\alpha \equiv \varphi, \alpha \mid\!\sim \gamma}{\varphi \mid\!\sim \gamma}$
3. **RW:** $\dfrac{\alpha \rightarrow \varphi, \gamma \mid\!\sim \alpha}{\gamma \mid\!\sim \varphi}$
4. **And:** $\dfrac{\alpha \mid\!\sim \varphi, \alpha \mid\!\sim \gamma}{\alpha \mid\!\sim \varphi \wedge \gamma}$
5. **Or:** $\dfrac{\alpha \mid\!\sim \varphi, \gamma \mid\!\sim \varphi}{\alpha \vee \gamma \mid\!\sim \varphi}$
6. **Cut:** $\dfrac{\alpha \wedge \gamma \mid\!\sim \varphi, \alpha \mid\!\sim \gamma}{\alpha \mid\!\sim \varphi}$
7. **CM:** $\dfrac{\alpha \mid\!\sim \gamma, \alpha \mid\!\sim \varphi}{\alpha \wedge \varphi \mid\!\sim \gamma}$
8. **RM:** $\dfrac{\alpha \mid\!\sim \varphi, \alpha \not\mid\!\sim \neg\gamma}{\alpha \wedge \gamma \mid\!\sim \varphi}$

Semantics

- The semantics for a statement like $b \mid\!\sim f$ rest on the idea of ordering valuations $\mathcal{V}$ by preference (*typicality*)[1]

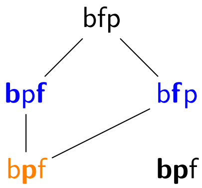

- Any *minimal* $v \in \mathcal{V}$ where birds is true, $v \Vdash f$
- So, also $p \mid\!\sim \neg f$

[1] **boldface** indicates false

How do we develop a rational consequence relation in FCA?

Extending the Language

- In FCA, implications are defined over sets of attributes
- We do not naturally have a way of expressing $\vee$ or $\neg$
- Even if we want to restrict ourselves to these inexpressive implications, we need to show that the pattern of non-monotonic reasoning corresponds to a rational consequence relation
- So we extend the language to **compound attributes**

Definition (The Language $\mathcal{L}_M$)

$$\phi := m \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid \phi_1 \vee \phi_2$$

$$m^{\downarrow} = \{g \mid (g, m) \in I\}$$

$$(\phi_1 \wedge \phi_2)^{\downarrow} = \phi_1^{\downarrow} \cap \phi_2^{\downarrow}$$

$$\neg\phi^{\downarrow} = G \setminus \phi^{\downarrow}$$

$$(\phi_1 \vee \phi_2)^{\downarrow} = \phi_1^{\downarrow} \cup \phi_2^{\downarrow}$$

Compound Attributes: Example

$(G, M, I) =$

	Republican	Quaker	Pacifist
Nixon	×	×	
Bush	×		
Penn		×	×

- $R^{\downarrow} = \{\texttt{Nixon, Bush}\}, Q^{\downarrow} = \{\texttt{Nixon, Penn}\}, P^{\downarrow} = \{\texttt{Penn}\}$
- *Those objects that are not republicans, or pacifists*
- $(\neg\mathtt{R} \vee \mathtt{P})^{\downarrow} = (G \setminus \mathtt{R}^{\downarrow}) \cup \mathtt{P}^{\downarrow} = \{\texttt{Penn}\}$

Cont.

$(G, M \cup \neg\mathtt{R} \vee \mathtt{P}, I) =$

	Republican	Quaker	Pacifist	$\neg\mathtt{R} \vee \mathtt{P}$
Nixon	$\times$	$\times$		
Bush	$\times$			
Penn		$\times$	$\times$	$\times$

Preferential Context

- An analogue to the view of a partial order over valuations representing preference is to order the objects in $\mathbb{K}$
- A preferential context $\mathbb{P} = (G, M, I, \preceq)$ where $\preceq$ is a partial-order over G.

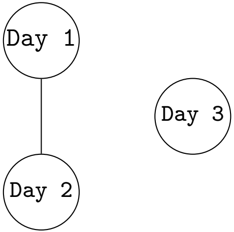

	Sun	Rain	Wind	Hot
Day 1	×	×	×	×
Day 2				×
Day 3	×		×	×

Defeasible Implications in FCA

Definition (Defeasible Implication over $\mathcal{L}_M$)

A defeasible implication $\mathtt{A} \mid\!\sim \mathtt{B} \in \mathcal{L}_M$ holds in an preferential context $\mathbb{P}$ iff. the minimal objects $\underline{\mathtt{A}}^{\downarrow} \subseteq \mathtt{B}^{\downarrow}$

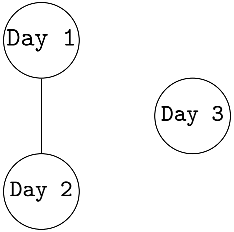

	Sun	Rain	Wind	Hot
Day 1	×	×	×	×
Day 2				×
Day 3	×		×	×

- $\mathtt{Hot} \mid\!\sim \neg\mathtt{Rain}$ then holds in $\mathbb{P}$ since $\underline{\mathtt{Hot}}^{\downarrow} \subseteq \neg\mathtt{Rain}^{\downarrow}$
- A partial order doesn't guarantee Rational Monotonicity

Rational Monotonicity

$$\textbf{Rational Monotonicity} = \frac{\alpha \mathrel{|\!\sim} \varphi, \alpha \mathrel{|\!\not\sim} \neg\gamma}{\alpha \wedge \gamma \mathrel{|\!\sim} \varphi}$$

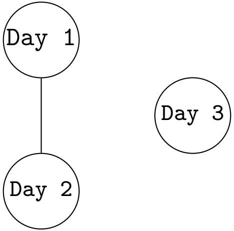

	Sun	Rain	Wind	Hot	¬Rain
Day 1	×	×	×	×	
Day 2				×	×
Day 3	×		×	×	×

1. $\mathbb{P} \models \texttt{Hot} \mathrel{|\!\sim} \neg\texttt{Rain}$
2. $\mathbb{P} \models \texttt{Hot} \mathrel{|\!\not\sim} \neg\texttt{Wind}$
3. $\mathbb{P} \not\models \texttt{Hot} \wedge \texttt{Wind} \mathrel{|\!\sim} \neg\texttt{rain}$ (because of Day 1)

Ranked Context

Definition

A ranked context $\mathbb{R} = (G, M, I, \mathrm{R}, \Delta)$ is a formal context with a ranking function $\mathrm{R} : G \mapsto \mathbb{N}$, where the partial order imposed by R is *modular*, and Δ is a set of (defeasible) constraints we place over (G, M, I).

- Object g is more "typical" than g' iff. $\mathrm{R}(g) < \mathrm{R}(g')$
- *Modularity* enforces that if two objects are incomparable then they occupy the same rank
- Δ represents some external domain knowledge about how we expect objects in a context to behave (i.e. *that mammals are usually viviparous*)
- We derive R by a translated `BaseRank` algorithm (not shown)

Ranked Context

R	$\mathbb{K}$	bird	mammal	reptile	aerial	aquatic	terrestrial	oviparous	viviparous	limbs
0	Swallow	×			×			×		×
	Crow	×			×			×		×
1	Ostrich	×					×	×		×
	Penguin	×				×	×	×		×
2	Snake			×			×	×		
	Whale		×			×			×	
3	Chameleon			×			×		×	×
	Platypus		×			×	×	×		×

Table: A ranked context of some vertebrates

- bird $\mid\!\sim$ aerial
- terrestrial $\wedge$ bird $\mid\!\sim$ $\neg$aerial

Recap

- We argued that FCA could benefit from "softer" rules, enabling it to reason with error-prone, or exceptional data
- Proposed KLM-style defeasible implications as an attempt to do this
- Requires a change from implications over attributes to implications over compound attributes
- And that we place some order, encoding a view of typicality, over the object set
- Have an algorithm to determine this order, given some background information about the context (not shown)

Future Work & Paper

- If we accept that defeasible implications contain useful information, it would be nice to have consistency w.r.t concepts
- i.e. if $\texttt{bird} \mid\!\sim \texttt{fly}$ then the concept derived from `bird` should contain `fly`
- This is the notion of a *typical concept*
- We would obviously like to then have a concept lattice for the set of typical concepts - this is the tricky part

- If you are interested, we have a paper “Non-monotonic Extensions to Formal Concept Analysis via Object Preferences”
- Also, Ding, Yiwen, et al. ”Defeasible Reasoning on Concepts.” arXiv preprint arXiv:2409.04887 (2024)

Thanks

References

Bernhard Ganter and Rudolf Wille.
Contextual attribute logic.
In *International Conference on Conceptual Structures*, pages 377–388. Springer, 1999.

Bernhard Ganter and Rudolf Wille.
Formal concept analysis: mathematical foundations.
Springer Science & Business Media, 2012.

Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato.
Semantic characterization of rational closure: From propositional logic to description logics.
Artificial Intelligence, 226:1–33, 2015.

Sarit Kraus, Daniel Lehmann, and Menachem Magidor.
Nonmonotonic reasoning, preferential models and cumulative logics.
Artificial intelligence, 44(1-2):167–207, 1990.

Daniel Lehmann and Menachem Magidor.
What does a conditional knowledge base entail?
Artificial intelligence, 55(1):1–60, 1992.