

Collaborative Hybrid Human AI Learning through Conceptual Exploration

Bernhard Ganter^{1,2,†}, Tom Hanika^{3,*,†}, Johannes Hirth^{4,†} and Sergei Obiedkov^{5,†}

¹Ernst-Schröder-Zentrum, Karolinenpl. 5, 64289 Darmstadt, Germany

²TU Dresden, 01062 Dresden, Germany

³Intelligent Information Systems, University of Hildesheim, Universitätsplatz 1, 31141 Hildesheim, Germany

⁴Knowledge & Data Engineering Group, University of Kassel, Wilhelmshöher Allee 73, 34121 Kassel, Germany

⁵Faculty of Computer Science / cfaed / ScaDS.AI, TU Dresden, 01062, Dresden, Germany

Abstract

Conceptual Exploration is a sophisticated method for the interactive and structured acquisition of knowledge from experts. It is therefore particularly suitable for the use in hybrid settings where both humans and AIs act as experts. This article provides a brief summary of how conceptual exploration can be used in the context of Hybrid Human AI systems, as discussed within a tutorial during the third HHAI conference in Malmö, Sweden. We will recapitulate two small experiments that were carried out with the participants of the tutorial and their results. Finally, we give some pointers on how this promising link can be further researched.

Keywords

interactive AI, hybrid learning, hybrid AI, conceptual knowledge, collaborative learning

1. Introduction

Conceptual exploration is a knowledge acquisition method, which takes the approach that one would naturally take if one wanted to fully explore a defined area of knowledge. It underpins this approach with detailed algorithms.

The “expert” consulted in such an acquisition need not be human; it may also be an AI or a human-AI combination. That is one reason why we present this methodology to the HHAI community during the third conference on Hybrid Human-Artificial Intelligence [1]. A second is that the results of such explorations are documented in a human-readable format, even if they are computer-generated.

HHAI-WS 2024: Workshops at the Third International Conference on Hybrid Human-Artificial Intelligence (HHAI), June 10–14, 2024, Malmö, Sweden

*Corresponding author.

†These authors contributed equally.

✉ bernhard.ganter@tu-dresden.de (B. Ganter); tom.hanika@uni-hildesheim.de (T. Hanika); hirth@cs.uni-kassel.de (J. Hirth); sergei.obiedkov@tu-dresden.de (S. Obiedkov)

🌐 tomhanika.de (T. Hanika); <https://www.kde.cs.uni-kassel.de/hirth> (J. Hirth);

https://iccl.inf.tu-dresden.de/web/Sergei_Obiedkov (S. Obiedkov)

🆔 0000-0003-0767-1379 (B. Ganter); 0000-0002-4918-6374 (T. Hanika); 0000-0001-9034-0321 (J. Hirth);

0000-0003-1497-4001 (S. Obiedkov)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

The precise theoretical foundation allows for many, sometimes surprising, variants and generalizations. This is documented in a comprehensive monograph [2]. However, the presentation at the workshop for the HHAI conference only outlined simple basic principles of the approach; the goal was to raise awareness of the usefulness of this approach for a collaborative AI.

We explain these principles in the following section and provide some pointers to more expressive extensions. In Section 3 we present the two collaborative conceptual explorations we conducted with the about 30 tutorial participants at HHAI 2024.

2. Conceptual exploration

This method of acquisition originated in the research area Formal Concept Analysis, which sees its main task in using mathematical methods (more precisely: methods of mathematical lattice theory) to structure given data conceptually and thus make it more accessible to human understanding. The basics of the procedure described briefly in the following section are intuitive and can usually be applied without prior knowledge.


2.1. Basic principles

Formal Concept Analysis uses one basic data type, called a *formal context*. Other data forms must first be translated into this form. A formal context defines the objects under investigation and a selection of their possible attributes, noting which objects have which of these attributes. In conceptual exploration, the set of objects for a fixed set of attributes is kept dynamic, growing as the acquisition progresses. The aim of the exploration is to determine all permissible combinations of attributes. To do this, the questions which are repeatedly asked are whether the presence of certain attributes forces further attributes. Depending on the answer, a new object or an if-then rule is noted. The (surprisingly effective) algorithmic support is provided by calculating and then asking a “simplest unanswered question” at any given time during such an acquisition. When there are no more such unanswered questions, then the acquisition is provably complete. The closure system of attribute combinations (and thus the *concept lattice* of the explored knowledge area) then has been completely acquired.

The result of such a knowledge acquisition process consists of two lists: a list of (attribute) implications of the form “every object with the attributes a, b, c, ... necessarily also has the attributes x, y, z, ...” and a second list of counterexamples of the form “this object has the attributes a, b, c, ..., but not the attribute x”. These lists are kept as short as possible using algorithms. At the end of an exploration, an if-then rule follows from the implications list if and only if it is not refuted by an element of the counterexample list.

2.2. Conceptual Exploration

Based on the above, the most important components of conceptual exploration are a *domain* to be explored, most simply described by a set of attributes M , a *query generator* (called query engine), an *exploration base* and an *expert*. The exploration base is represented by a pair $(\mathcal{E}, \mathcal{L})$, where \mathcal{E} is the set of (so far) encountered *counter examples* and \mathcal{L} the set of implications that are approved by the expert. One step in the conceptual exploration algorithm goes as follows.



\mathbb{K}	m_1 \dots m_n	$\mathcal{L} = \{ A_1 \rightarrow B_1,$
g_1	\dots	\dots
\vdots		$A_\ell \rightarrow B_\ell$
g_k	\dots	$X \rightarrow Y \}$
g_{k+1}	$\text{--- } C \text{ ---}$	

Figure 1: The conceptual exploration method in a nutshell: Exploration base (bottom), which consists of a formal context of counter examples \mathcal{E} (left) and a set of valid implications \mathcal{L} (right). The algorithm (top) computes the “simplest undecided implication” $X \rightarrow Y$ and asks the Expert/AI if this implication is valid. If the expert approves, the implication is added to the \mathcal{L} . If the expert refutes the implication, a counter example has to be provided, which is then added to the formal context of counter examples.

The query engine computes the next “best” question based on the current state of $(\mathcal{E}, \mathcal{L})$. In our (basic) setting, the question will be in the form: does every object having the attributes $X \subseteq M$ also have the attributes $Y \subseteq M$? Or in short: is $X \rightarrow Y$ valid? The expert then has two options. Either, the expert accepts $X \rightarrow Y$ as a valid implication in the domain, or, he refutes the implication. In the latter case the expert is obliged to present a counter example in the “language” of the domain, i.e., an object described by the attributes from M . We illustrate this process in Figure 1. The unique feature of the exploration algorithm is that it provably requires the minimum number of questions to elicit the complete implication knowledge of a domain [2]. Moreover, the result is a set of counter example and *the* minimal base for the set of all valid implications in the domain [3].

2.3. Extensions, interfaces, and limitations

The extensive theoretical basis of the procedure allows for many extensions, for example, the inclusion of several contradicting and unreliable experts. Therefore, this methodology is also suitable for human-machine cooperation. Because the results are fully summarized in the two lists, they can be checked in detail. Implausible consequences from an acquisition result can be traced back to the lists using pinpointing techniques, and it can be verified whether the individual entry is correct or erroneous.

In its (already quite useful) basic version, the method is limited to simple object-has-attribute data and simple if-then rules (propositional horn). However, Formal Concept Analysis also provides methods that allow the use of more expressive data and richer logics. The price for this is usually higher algorithmic complexity. To give the reader an impression we want to mention a few extensions in the following.

Most simple extension is to incorporate background knowledge using a formal context [4].

When domain attributes are defined in some formal way (backed through some formal logic), new possibilities for the exploration method occur. For mathematical reasons, a connection to description logics [5] is an obvious link that has been extensively researched. In this context, exploration was also adapted for various description logics [6, 7]. Contrary approaches for more expressive domain languages can be found using triadic (i.e., ternary) context representations [8] as well as fuzzy values [9]. In its original form, conceptual exploration was defined as a method for acquiring the knowledge of exactly one expert. Given a set of multiple experts, various problems arise, which were investigated using different approaches [10, 11]. For a more extensive overview on extension we refer the reader to Ganter and Obiedkov [2].

3. Tutorial and outcomes

We conducted during the interactive part of the tutorial two experiments with the participants. In the first setting the audience explored properties from the free collaborative knowledge graph Wikidata [12, 13]. In the second setting the participants were challenged individually with a set of questions about the meaning of words. Afterwards, the answers were used in a collaborative analysis of the corresponding semantic field.

3.1. Exploring Wikidata

In Wikidata¹ (WD), knowledge is represented using statements. These link entities to values via properties, which in turn can be entities. For example, the statement John von Neumann was a computer scientist is represented by a connection from item Q17455 (John von Neumann) to item Q82594 (computer scientist) using property P106 (occupation). In the tutorial experiment the participants cooperatively explored the WD knowledge graph, more specifically, a chosen subset of WD properties. For this we employed *The Exploration Game*² (TEG) [14], an interactive tool that implements conceptual exploration and interfaces with WD.

In every step of the exploration TEG asks for the validity of an implication $X \rightarrow Y$ and provides information about how many statements support and how many statements refute it. After conducting the collaborative conceptual exploration, the tutorial participants discussed the results. They pointed out the advantages of the conceptual approach for learning collaboratively within a domain. Also, they pointed out limitations, such as the expressivity of WD statements and, moreover, of the used Horn rules.

3.2. Exploring Semantic Fields

Lexical typology is a relatively young area of linguistics that focuses on a comparative analysis of word meanings in different languages and the search for regularities in lexical categorization of reality by humans [15, 16]. Somewhat simplifying, the main question of the area can be summarized as follows:

How do words in different languages cover a conceptual space of related meanings?

¹<https://wikidata.org/>

²<https://teg.toolforge.org/>




Figure 2: Semantic map for the semantic field ‘sharp’ [18].

The question is usually asked with respect to a specific *semantic field*, and an answer typically comes in the form of a *semantic map*, which shows the interrelations between elementary meanings within the semantic field based on their colexification across various languages. These elementary meanings are sometimes referred to as *semantic frames* and are assumed to correspond to prototypical situations relevant to the semantic field.

Several types of semantic maps have been considered. *Classical semantic maps*, as they are called [15], are undirected graphs where nodes are meanings and an edge connecting two nodes suggests that the corresponding meanings can be colexified, i.e., covered by a single lexical item, in at least one language. Such a map is expected to be consistent with the *connectivity hypothesis*: Any relevant language-specific and/or construction-specific category should map onto a connected region in conceptual space [17], which, in practical terms, requires every subset of meanings that can be colexified to induce a connected subgraph of the semantic map.

This requirement is, of course, not sufficient to produce useful semantic maps, since it is trivially satisfied by a complete graph. Thus, in addition, the *economy principle* is usually adopted: No edge is needed between frames *A* and *C* if linguistic items expressing *A* and *C* always express *B* [15].

Figure 2 shows a simple semantic map for adjectives in the semantic field ‘sharp’ from [18]. It suggests, in particular, that, if the same word can be used to describe the sharpness of a knife and a thorn, then it can also be used to describe the sharpness of an arrow. However, it remains unclear whether such a word would necessarily be appropriate for describing a sharp nose. The map does not enforce it, but the problem is that, with a classical map, it cannot be enforced without prohibiting some other combinations observed in data. This is a well-known problem with the classical semantic map: by trying to accommodate all combinations of meanings attested in the data, “it predicts much more than is actually found” [19].

Concept lattices offer an alternative to classical semantic maps that does not have this problem [20], with the additional advantage of being constructible automatically from data. The formal context is built in a straightforward way by letting the words of the semantic field be its objects and frames be its attributes. The concept lattice of the semantic field ‘sharp’ is shown in Figure 3. It suggests the same relations between the *knife*, *arrow*, and *thorn* frames as the semantic map in Figure 2 (the infimum of *knife* and *thorn*, labeled *inf*, is below *arrow*), but, in addition, it makes it clear that there are words that colexify these three frames without expressing the meaning ‘object with a sharp form’ (*inf* is not below *nose*).

Collecting a representative dataset for building semantic maps is itself a challenge. There are two main approaches [15]. In the *onomasiological approach*, the first step is to identify core




Figure 3: Concept lattice for the semantic field ‘sharp’; words are not shown [18]. The node labeled as *inf* is the infimum of the attribute concepts of ‘instrument with a sharp functional edge (knife)’ and ‘object/surface that pricks (thorn)’.

meanings in the semantic field and then search for individual forms that express these meanings in different languages. This is a perfect case for attribute exploration: fix a few meanings and collect words as counterexamples to implications over these meanings. Questions to be answered in the process are of the following form: If a word colexifies a set A of meanings, must it express a meaning b too?

An alternative is the *semasiological approach*: one starts by choosing a single meaning as a pivot and then lists the other meanings of the linguistic items expressing the pivot meaning. This can be supported by object exploration, where one would fix a set of words sharing a meaning and identify their other meanings as counterexamples to implications over the words. Such implications look as follows: If a meaning is shared by a set C of words, it is also shared by the word d . Alternating between object and attribute exploration, it is possible to collect a representative set of lexical items and to identify elementary meanings of the semantic field.

Thus, concept lattices provide an interesting alternative to classical semantic maps. They can be built automatically if data is already collected, and, if not, attribute and object exploration can help organize data collection. However, better software is needed for linguists to be able to use these methods.

The participants in the tutorial were native speakers of different languages, including German, English, Dutch, Italian, and Swedish. We presented questions involving the semantic field of the verb “falling” to them and the resulting concept lattice is depicted in Figure 4. Based on this, we have demonstrated how a conceptual exploration can also be carried out on already collected data. This led to some exciting discussions among the various native speakers.

Interested readers are welcome to contact the authors of this article if they would like a more detailed description of the methods and results, in particular the data recorded.

Acknowledgments

This work is partly supported by DFG in project 389792660 (TRR 248, Center for Perspicuous Systems), by BMBF in ScaDS.AI, and by BMBF and DAAD in project 57616814 (SECAI).




Figure 4: The concept lattice for the collaboratively explored semantic field of the verb “falling”.

References

- [1] F. Lorig, J. Tucker, A. D. Lindström, F. Dignum, P. K. Murukannaiah, A. Theodorou, P. Yolum (Eds.), HHAI: Proceedings of the 3rd International Conference on Hybrid Human-Artificial Intelligence, volume 386 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2024. URL: <https://doi.org/10.3233/FAIA386>. doi:10.3233/FAIA386.
- [2] B. Ganter, S. Obiedkov, *Conceptual Exploration*, Springer, 2016. doi:10.1007/978-3-662-49291-8.
- [3] J.-L. Guigues, V. Duquenne, Familles minimales d’implications informatives résultant d’un tableau de données binaires, *Mathématiques et Sciences humaines* 95 (1986) 5–18.
- [4] B. Ganter, Attribute exploration with background knowledge, *Theor. Comput. Sci.* 217 (1999) 215–233. doi:10.1016/S0304-3975(98)00271-0.
- [5] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, P. F. Patel-Schneider (Eds.), *The*

Description Logic Handbook: Theory, Implementation, and Applications, Cambridge University Press, 2003.

- [6] F. Baader, Computing a minimal representation of the subsumption lattice of all conjunctions of concepts defined in a terminology, in: Proc. Intl. KRUSE Symposium, August, Citeseer, 1995, pp. 11–13.
- [7] D. Borchmann, Learning terminological knowledge with high confidence from erroneous data, Ph.D. thesis, Dresden University of Technology, 2014. URL: <https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa-152028>.
- [8] M. Felde, G. Stumme, Triadic exploration and exploration with multiple experts, in: A. Braud, A. Buzmakov, T. Hanika, F. L. Ber (Eds.), ICFCA 2021, Proceedings, volume 12733 of LNCS, Springer, 2021, pp. 175–191. doi:10.1007/978-3-030-77867-5_11.
- [9] C. V. Glodeanu, Attribute exploration with fuzzy attributes and background knowledge, in: M. Ojeda-Aciego, J. Outrata (Eds.), CLA, volume 1062 of CEUR Workshop Proceedings, CEUR-WS.org, 2013, pp. 69–80. URL: <https://ceur-ws.org/Vol-1062/paper6.pdf>.
- [10] T. Hanika, J. Zumbrägel, Towards collaborative conceptual exploration, in: P. Chapman, D. Endres, N. Pernelle (Eds.), Graph-Based Representation and Reasoning (ICCS), volume 10872 of LNCS, Springer, 2018, pp. 120–134. doi:10.1007/978-3-319-91379-7_10.
- [11] M. Felde, G. Stumme, Interactive collaborative exploration using incomplete contexts, Data Knowl. Eng. 143 (2023) 102104. doi:10.1016/J.DATAK.2022.102104.
- [12] D. Vrandečić, M. Krötzsch, Wikidata: a free collaborative knowledgebase, Commun. ACM 57 (2014) 78–85. doi:10.1145/2629489.
- [13] D. Vrandečić, Wikidata: a new platform for collaborative data collection, in: A. Mille, F. Gandon, J. Misselis, M. Rabinovich, S. Staab (Eds.), Proceedings of 21st WWW 2012 (Companion Volume), ACM, 2012, pp. 1063–1064. doi:10.1145/2187980.2188242.
- [14] T. Hanika, M. Marx, G. Stumme, Discovering implicational knowledge in wikidata, in: D. Cristea, F. L. Ber, B. Sertkaya (Eds.), 15th ICFCA 2019, Proceedings, volume 11511 of LNCS, Springer, 2019, pp. 315–323. doi:10.1007/978-3-030-21462-3_21.
- [15] T. Georgakopoulos, S. Polis, The semantic map model: State of the art and future avenues for linguistic research, Language and Linguistics Compass 12 (2018) e12270. doi:<https://doi.org/10.1111/lnc3.12270>, e12270 LNCO-0727.R1.
- [16] E. Rakhilina, D. Ryzhova, Y. Badryzlova, Lexical typology and semantic maps: Perspectives and challenges, Zeitschrift für Sprachwissenschaft 41 (2022) 231–262. URL: <https://doi.org/10.1515/zfs-2021-2046>. doi:doi:10.1515/zfs-2021-2046.
- [17] W. Croft, Radical Construction Grammar: Syntactic Theory in Typological Perspective, Oxford University Press, 2001. doi:10.1093/acprof:oso/9780198299554.001.0001.
- [18] D. Ryzhova, S. Obiedkov, Formal concept lattices as semantic maps, in: Computational Linguistics and Language Science, volume 1886 of CEUR Workshop Proceedings, 2017, pp. 78–87. URL: http://ceur-ws.org/Vol-1886/paper_10.pdf.
- [19] M. Cysouw, Martin Haspelmath, Indefinite pronouns (Oxford Studies in Typology and Linguistic Theory). Oxford: Clarendon Press, 1997. pp. xvi+364., Journal of Linguistics 37 (2001) 593–625. doi:10.1017/S0022226701231351.
- [20] T. Georgakopoulos, E. Grossman, D. Nikolaev, S. Polis, Universal and macro-areal patterns in the lexicon: A case-study in the perception-cognition domain, Linguistic Typology 26 (2022) 439–487. doi:doi:10.1515/lingty-2021-2088.