

Concepts in Action: Introduction



Lucas Bechberger  and Mingya Liu 

It is impossible to talk about human cognition without talking about concepts—there simply *is* no human cognition without concepts. Concepts form an abstraction of reality that is central to the functioning of the human mind. Conceptual knowledge (of e.g., APPLE, LOVE and BEFORE) is crucial for us to categorize, understand, and reason about the world. Only equipped with concepts and words for them can we successfully communicate and carry out actions. *But what exactly are concepts? How are concepts acquired? How does the human mind use concepts?* Such questions have been a subject of discussion since antiquity and remain highly relevant in multiple fields (e.g., Murphy 2002; Margolis and Laurence 2015).

Recent decades have seen fruitful results and methodological advances on concept research in disciplines such as linguistics, philosophy, psychology, artificial intelligence, and computer science. For instance, cognitive psychologists use empirical experiments to validate formal models of concept representation and learning such as the prototype theory (Rosch et al. 1976), the exemplar theory (Murphy 2016) or other alternative theories (Rogers and McClelland 2004; Blouw et al. 2016). Linguists pursue the goal of assigning more precise meaning to natural language expressions by mainly applying logic-based formalisms (Asher 2011). In machine learning, decision boundaries in high-dimensional feature spaces are used to define membership to a concept (Mitchell 1997). Moreover, researchers in the semantic web area have

L. Bechberger (✉)

Institute of Cognitive Science, Osnabrück University, Osnabrück, Germany
e-mail: lucas.bechberger@uni-osnabrueck.de

M. Liu (✉)

Department of English and American Studies, Humboldt University of Berlin,
Berlin, Germany
e-mail: mingya.liu@hu-berlin.de

© The Author(s) 2021

L. Bechberger et al. (eds.), *Concepts in Action*, Language, Cognition, and Mind 9,
https://doi.org/10.1007/978-3-030-69823-2_1

created large ontologies (Gómez-Pérez et al. 2004) containing hierarchies of concepts formulated in description logics. Google’s “Knowledge Graph” illustrates how such ontologies can be used in industrial applications.

Despite of this plethora of research, there remain many open questions, unsolved debates and methodological challenges. For instance, the ontologies of the semantic web have been challenged as being unable to represent information about conceptual similarity and thus as being ill-suited for representing conceptual knowledge (Gärdenfors 2004). And although deep learning models are often said to acquire concepts when learning to classify pictures of dogs, umbrellas, and other objects, they can be easily fooled by slightly manipulated input images (Szegedy et al. 2013)—which highlights that they only learn patterns, but no conceptual knowledge.

One major obstacle for a better and more holistic understanding of concepts is that research on concepts has usually been carried out in different disciplines individually—with different approaches, different goals, and different results. The multi-disciplinary research efforts usually run in parallel without enough interaction; existing interdisciplinary research projects usually do not involve more than two disciplines, for example linguistics and computer science in the WordNet project (Fellbaum and Vossen 2016) or psychology and artificial intelligence in cognitive architectures like ACT-R (Anderson 2009) or SOAR (Laird 2012). In order to move the scientific understanding of concepts forward, we need a truly interdisciplinary perspective on concepts, involving a mutual understanding of the different approaches from different disciplines, a lively exchange of ideas, and synergies arising from the combination of different research perspectives and methods. Thus, our volume will focus on selected recent issues, approaches and results that are not only central to the highly interdisciplinary field of concept research, but that are also particularly important to newly emergent paradigms and challenges.

This volume focuses on three topics (i.e., three distinct points of view) that lie at the core of concept research: *representation*, *learning*, and *application*. In the following, we will first present research questions related to the three topics (Sect. 1), and then, we will provide an overview of the contributions (Sect. 2).

1 Research Questions

In order to structure an interdisciplinary discussion and exchange about concept research, we found it useful to put a focus on three essential questions that need to be answered: How can conceptual knowledge be represented (Sect. 1.1)? How are concepts acquired (Sect. 1.2)? How is conceptual knowledge applied in cognitive tasks (Sect. 1.3)?

1.1 Representation: How Can We Formally Describe and Model Concepts?

One of the major challenges in concept research is to find a formal representation of concepts that is on the one hand able to explain a wide range of empirical observations and experimental results and that can on the other hand be easily applied in practice. Exemplar and prototype theories from psychology focus on the crucial role of representative instances, whereas knowledge-based theories (Murphy and Medin 1985) emphasize that concepts do not occur in isolation, but always stand in relations to other concepts. Ontologies (Gómez-Pérez et al. 2004) from the semantic web area provide a formal way of describing such networks of concepts. The logical-formal approaches from linguistics aim at accounts of the context-independent and context-dependent aspects of meaning and can be related to logic-based representations in artificial intelligence (Russell and Norvig 2002). Finally, the feature spaces commonly used in the field of machine learning (Mitchell 1997) (for example in nearest-neighbor classifiers) can be linked to prototype and exemplar approaches from psychology. When analyzing formal representations of concepts, the following questions should be considered:

- What are the underlying assumptions of different representation approaches? How are they motivated and to what extent are they compatible with one another?
- How can different representation formalisms be compared with one another? What are useful and meaningful criteria for making such a comparison?
- How transferable are the different approaches to other domains? For instance, are there any benefits in using a prototype approach in linguistic analyses of natural language semantics?
- How can different representational approaches be augmented or combined with one another in order to arrive at a more holistic model?

1.2 Learning: Where Do Concepts Come from and How Are They Acquired?

Another major issue in concept research is concerned with concept acquisition, which is not only important per se but also essential for evaluating whether a specific theory of human concepts is psychologically plausible (Carey 2015). While there are well-established assumptions about children's acquisition of core concepts such as the basic-level bias and the taxonomic assumption, the exact nature of the underlying processes remains controversial. On a larger time scale, the evolution of concepts in human societies (Hull 1920) and similar processes in groups of robots (Spranger 2012) can give insights into learning processes. Moreover, studying concept learning across languages and cultures enables a better understanding of universality and diversity in concepts (Imai et al. 2010). Furthermore, to successfully coin and transfer new concepts, it is crucial to understand differences between everyday concepts and

expert concepts, e.g., in mathematics (Rips et al. 2008). These and other related issues (e.g., innateness, groundedness and embodiment) require researchers to not only strive for advances in their own field (such as in terms of improved machine learning algorithms in artificial intelligence), but also to start in-depth exchanges with neighboring disciplines. The following questions can provide useful guidelines when approaching concept learning:

- Which kinds of learning mechanisms (e.g., supervised vs. unsupervised, multi-modal vs. unimodal) are in the focus of the different research disciplines? Can multiple learning mechanisms be combined with one another?
- Which representational assumptions are made by the different learning mechanisms?
- How does concept learning interact with the development of low-level and other high-level cognitive (e.g., motorsensory, reasoning) abilities? What are their underlying mechanisms?
- What are the differences between learning concrete (e.g., APPLE) and abstract concepts (e.g., LOVE) and between learning expert and everyday concepts?

1.3 Application: How Are Concepts Used in Cognitive Tasks?

The last decade has witnessed an exploding utilization of conceptual knowledge bases, unprecedented both in scale and range of applications. The conceptual core of the semantic web (Berners-Lee et al. 2001) and artificial agents like IBM's Watson (Ferrucci et al. 2010) is largely based on AI technologies dating back to the last millennium (e.g., description logics). The new development has clearly shown the potential but also the limits of such approaches. The questions that arise here obviously link to other fields: The combination of a multitude of potential resources asks for modern AI methods to reason over heterogeneous and inconsistent data (Potyka and Thimm 2017). The application of conceptual knowledge in communication, including conceptual combination and application of conceptual knowledge in context, are classical problems in linguistics. And the problem of generating new concepts may find answers in recent psychological theories on creativity (Schorlemmer et al. 2014). The following questions are important with regard to the application of concepts:

- Which aspects of applying concepts are analyzed in the different disciplines? Is there any considerable overlap?
- How do the different representational formalisms constrain specific application scenarios and vice versa?
- Which contextual effects occur in the application of concepts? How are these effects handled in different frameworks?
- Which mechanisms exist for performing conceptual combination? Which constraints apply? How are conflicts that emerge from conflicting conceptualizations detected and resolved?

2 Summaries of the Contributed Chapters

This volume consists of seven individual chapters from different scientific disciplines, each of which relates to at least one of the other topics presented in Sect. 1. Figure 1 illustrates how the individual contributions relate to each other, based on their underlying disciplines, common themes, and the three focus topics from Sect. 1. Figure 1 illustrates both the strong relations between the individual contributions and the broad spectrum of this edited volume. We will now introduce the individual contributions in more detail.

Bechberger and Kühnberger’s contribution “Generalizing Psychological Similarity Spaces to Unseen Stimuli – Combining Multidimensional Scaling with Artificial Neural Networks” (Chap. 2) addresses the focus topic of *learning*. It uses a *spatial model* of concepts as regions in psychological similarity spaces based on Gärdenfors’ cognitive framework of conceptual spaces. These similarity spaces are typically obtained based on dissimilarity ratings from psychological studies and the technique of “multidimensional scaling” (MDS). This approach is however unable to generalize to unseen inputs. The authors propose to use MDS on human similarity ratings for initializing the similarity space and ANNs (artificial neural networks) to learn a mapping from raw stimuli into this similarity space. This proposal is a valuable contribution for integrating *psychology* and *artificial intelligence*. In order to

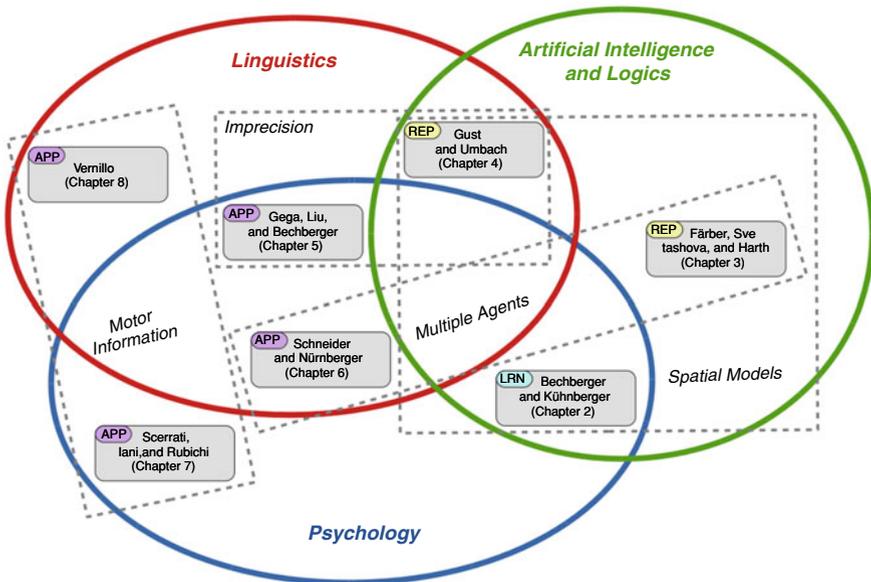


Fig. 1 Visualization of the contributed chapters based on scientific disciplines (solid ellipses), common research themes (dashed rectangles), and classification based on the three focus topics representation (REP), learning (LRN), and application (APP)

validate their hybrid approach, the authors conducted a feasibility study. Their results show that while their proposal works in principle, the generalization capabilities of the ANNs are still limited and need to be improved further.

Färber, Svetashova, and Harth's contribution "Theories of Meaning for the Internet of Things" (Chap. 3) is concerned with the *representation* of concepts in the context of the Internet of Things (IoT) from the perspective of *artificial intelligence*. They compare different representational frameworks from philosophy and computer science, taking a simple smart home setting as an application example. Overall, they consider four different approaches, namely model-theoretic semantics (which are based on first-order logic), possible world semantics (using modal logic), situation semantics, and cognitive and distributional semantics (i.e., *spatial models* of meaning). With the IoT application in mind, the authors assess whether these representational frameworks are able to represent intersubjectivity (i.e., *multiple agents*) and dynamics (i.e., changes in the state of the world) and to what extent they can be connected to perception. The authors conclude that none of the existing approaches is able to completely satisfy all three requirements. They propose to further investigate a combination between situational and distributional semantics as a promising avenue for future research.

Also **Gust and Umbach's** contribution "A Qualitative Similarity Framework for the Interpretation of Natural Language Similarity Expressions" (Chap. 4) explores the *representation* of concepts in the context of natural language semantics. It aims at the interpretation of expressions of similarity and sameness, such as *so/similar/same* in English or their counterparts in German. The authors argue that treating similarity as a primitive predicate is unsatisfactory because semantic differences between individual similarity expressions could not be accounted for and the role of similarity expressions in creating ad-hoc kinds, for example, by similarity demonstratives and scalar and non-scalar equatives would be obscured. The framework proposed in the paper introduces a non-metric qualitative concept of similarity which makes use of a *spatial model* called attribute spaces equipped with systems of predicates corresponding to predicates on the domain. Individuals are mapped to points in attribute spaces by generalized measure functions. Two individuals count as similar if their images in a particular attribute space given a particular predicate system cannot be distinguished. This allows representations of varying granularity and hence of different degrees of *imprecision*. The authors argue that the framework is suited for modeling the meaning of natural language similarity expressions and, moreover, account for their role in ad-hoc kind formation constructions. It thus provides a *logic*-based formalism which is able to explain *linguistic* phenomena.

Gega, Liu and Bechberger's contribution "Numerical Concepts in Context" (Chap. 5) deals with the semantics and pragmatics of numerical expressions, with a focus on their precise or *imprecise* interpretations. While the precise interpretation most prominently appears in mathematical contexts, the imprecise interpretation seems to arise when numbers (as quantities) are applied to real world contexts (e.g., *the rope is 50m long*). Earlier literature shows that the (im)precise interpretation can depend on different factors, e.g., the kind of approximators a numeral appears with (precise vs. imprecise, e.g., *exactly* vs. *roughly*) or the kind of the numeral

itself (round vs. non-round, e.g., 50 vs. 47). The authors report on a *corpus-linguistic* study and a *psycholinguistic* rating experiment of English numerical expressions. The results confirm the effects of both factors, and additionally also an effect of the kind of unit, namely, whether it refers to discrete versus continuous concepts (e.g., PEOPLE vs. METER).

Schneider and Nürnberger's contribution "Evaluating Semantic CoCreation by using a Marker as Linguistic Constraint in Cognitive Representation Models" (Chap. 6) explores the *application* of conceptual knowledge in communication between *multiple agents*. More specifically, they address semantic co-creation, i.e., the convergence of the cognitive models of the interlocutors within a conversation. The authors hypothesize that a shared marker can facilitate this coordination of representations. In order to validate this hypothesis, they conducted an experiment where groups of three participants needed to identify a target location on a given map. One participant (the describer) was given the target and had to describe it to the two other participants who needed to correctly identify this target location. One of them (the committer) was able to give feedback to the describer while the other one (the observer) had to remain passive. The authors considered four experimental conditions which differed in the availability of a shared marker (i.e., a movable point on the map) and in the complexity of the task (measured by the number of cities displayed on the map). Their results show that when task complexity was low, no real interaction between the participants was necessary to successfully solve the task. Contrary to their expectations, the shared marker was not able to improve performance in the high-complexity scenario. While their results highlight that a certain level of complexity is necessary to elicit interactions, it also casts doubt on the assumption that additional means of communication (such as a shared marker) necessarily improve the outcome of the interaction. Their work thus urges for further research both in *psychology* and *linguistics* to gain a deeper understanding of the observed effects.

Scerrati, Iani and Rubichi's contribution "Does the Activation of Motor Information Affect Semantic Processing?" (Chap. 7) considers the *application* of concepts in lexical decision tasks, focusing on the influence of pre-activated *motor information*. The authors report on a *psychological* priming experiment in which the subjects were instructed to make keypress responses depending on two factors: One factor is word type with target words being relevant/irrelevant/unrelated to action (e.g., *handle/ceramic/eyelash*) with respect to a prime object (e.g., image of a frying pan). The other factor is spatial compatibility with the related part of the prime object (e.g., *handle* for a frying pan) either on the same side or on the opposite side of the key to be pressed. The dependent measures were reading time (RT) latencies and error rates for the question whether the target word was an Italian word. The results of the RT latencies did not show any significant effects or an interaction. The results of the error rates however showed a significant main effect of word type with the lexical decision responses being more accurate with action-relevant target words than with action-irrelevant words or unrelated words. This indicates that motor activation may indeed influence semantic processing, thus complementing and enriching the literature that focuses on the reverse effect of semantic content on motor activation.

Also **Vernillo's** contribution "Grounding Abstract Concepts in Action: The Semantic Analysis of Four Italian Action Verbs Encoding Force Events" (Chap. 8) focuses on the *application* of conceptual knowledge, comparing the concrete and metaphorical uses of the four Italian action verbs *premere*, *spingere*, *tirare*, and *trascinare* (i.e., 'press', 'push', 'pull', and 'drag'). The underlying hypothesis is that the image schema of their literal meaning also constrains their usage in the metaphorical meaning. The *linguistic* study uses the representation of verb meanings through 3D scenes from the IMAGACT database. Based on the extracted data, the author provides a description of the semantic resemblances and differences in terms of salient image-schematic structures. The results show that the four verbs under consideration belong to the same semantic class of force (involving *motor information* and movement), and that they share commonalities in their literal and metaphorical use. At the same time, one can also observe systematic differences: For instance, while the literal meaning of *premere* focuses on the force exerted on the object, *spingere* emphasizes the resulting movement. These different connotations are also transferred to the metaphorical usage where *spingere* entails a change of state while *premere* does not. The results of this analysis support the view that metaphors are not just a linguistic phenomenon, but are grounded in embodied conceptual knowledge.

References

- Anderson, J. R. (2009). *How can the human mind occur in the physical universe?* Oxford University Press.
- Asher, N. (2011). *Lexical meaning in context: A web of words*. Cambridge University Press.
- Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web: A new form of web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American*, 284(5), 34–43.
- Blouw, P., Solodkin, E., Thagard, P., & Eliasmith, C. (2016). Concepts as semantic pointers: A framework and computational model. *Cognitive Science*, 40(5), 1128–1162.
- Carey, S. (2015). Why theories of concepts should not ignore the problem of acquisition. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 415–454). MIT Press.
- Fellbaum, C., & Vossen, P. (2016). *Proceedings of the 8th Global Wordnet Conference*.
- Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A. A., et al. (2010). Building Watson: An overview of the DeepQA project. *AI Magazine*, 31(3), 59–79.
- Gärdenfors, P. (2004). How to make the semantic web more semantic. In *Formal ontology in information systems* (pp. 19–36).
- Gómez-Pérez, A., Fernández-López, M., & Corcho, O. (2004). *Ontological engineering: With examples from the areas of knowledge management, e-commerce and the semantic web*. Springer.
- Hull, C. L. (1920). Quantitative aspects of evolution of concepts: An experimental study. *Psychological Monographs*, 28(1), i–86.
- Imai, M., Saalbach, H., & Stern, E. (2010). Are Chinese and German children taxonomic, thematic, or shape biased? Influence of classifiers and cultural contexts. *Frontier in Psychology*, 1(194).
- Laird, J. E. (2012). *The soar cognitive architecture*. MIT Press.
- Margolis, E., & Laurence, S. (Eds.). (2015). *The conceptual mind: New directions in the study of concepts*. MIT Press.
- Mitchell, T. M. (1997). *Machine learning*. McGraw Hill.

- Murphy, G. L. (2002). *The big book of concepts*. MIT Press.
- Murphy, G. L. (2016). Is there an exemplar theory of concepts? *Psychonomic Bulletin & Review*, 23(4), 1035–1042.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289–316.
- Potyka, N., & Thimm, M. (2017). Inconsistency-tolerant reasoning over linear probabilistic knowledge bases. *International Journal of Approximate Reasoning*, 88, 209–236.
- Rips, L. J., Bloomfield, A., & Asmuth, J. (2008). From numerical concepts to concepts of number. *Behavioral and Brain Sciences*, 31(6), 623–642.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. MIT Press.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382–439.
- Russell, S. J., & Norvig, P. (2002). *Artificial intelligence: A modern approach* (International Edition). Schorlemmer, M., Smaill, A., Kühnberger, K.-U., Kutz, O., Colton, S., Cambouropoulos, E., & Pease, A. (2014). COINVENT: Towards a computational concept invention theory. In *Proceeding of the ICCV*.
- Spranger, M. (2012). The co-evolution of basic spatial terms and categories. In L. Steels (Ed.), *Advances in interaction studies experiments in cultural language evolution* (Vol. 3, pp. 111–141). Amsterdam: John Benjamins.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). *Intriguing properties of neural networks*.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

