Available Article

Author's final: This draft is prior to submission for publication, and the subsequent edits in the published version. If quoting or citing, please refer to the proper citation of the published version below to check accuracy and pagination.

Cite as: Bergman, M. K. Conclusion. in *A Knowledge Representation Practionary: Guidelines Based on Charles Sanders Peirce* (ed. Bergman, M. K.) 371–380 (Springer International Publishing, 2018). doi:10.1007/978-3-319-98092-8_17

Official site: https://link.springer.com/book/10.1007/978-3-319-98092-8

Full-text: http://www.mkbergman.com/publications/akrp/chapter-17.pdf

Abstract: Peirce posited a "third grade of clearness of apprehension" to better understand a topic at hand, a part of his pragmatic maxim. This book has attempted to adhere to this 'practionary' form, the first attempt to so apply Peirce to a single concept. In Part I, we set the context by discussing the nature of information, knowledge, and representation, as well as the challenges and opportunities facing KR. In Part II we provided a speculative grammar for KR, including the structural role of the universal categories of Firstness, Secondness, and Thirdness, and the terminology, languages, logic, and models for knowledge representation.

That foundation lets us discuss existing frameworks and KR constructs in typologies and knowledge graphs in Part III. Throughout, we provided evidence where Peirce's ideas may offer unique and valuable insights into semantic technologies, knowledge representation, and information science. There are enticing connections to very topical fields in computer science. We need to better understand the nature of signs and representation in the use of semantic technologies.

17

CONCLUSION

Peirce posited a "third-grade of clearness of apprehension" to better understand a topic at hand, what he claimed as the ultimate expression of his *pragmatic maxim*. One of the favorite quotes I have used in this book is Peirce's first formulation of this maxim:

> "Consider what effects, which might conceivably have practical bearings, we conceive the object of our conceptions to have. Then, our conception of these effects is the whole of our conception of the object." (1874, CP 5.402, EP 1:132, W 3:266)

Peirce came to believe that this initial formulation did not capture his exact intent. Here is how Kelly Parker summarized it:[1]

> "In the proposal for Memoir 32, Peirce expressed his discomfort with this formulation of the pragmatic maxim. He wrote that the paper 'was imperfect in tacitly leaving it to appear that the maxim of pragmatism led to the last stage of clearness' (NEM 4:30). Indeed, the phrasing of the maxim is potentially misleading. One might read this statement as providing guidelines for an alternative means of *defining* concepts. If we think of standard dictionaries as giving the 'second-grade' linguistic definitions of concepts, we might take the pragmatic maxim as a guide to producing a super-dictionary of 'third-grade' definitions. Such a book (a '**practionary**'?) might endeavor to list all the practical effects a thing could have in experience, and thus furnish the reader with a better conception of the object." (p. 182) (bold added)

Throughout this book, I have attempted to adhere strictly to this form, the first such attempt to apply Peirce to the interpretation of a single concept, which, in our case, is *knowledge representation*. This book is the first attempt to produce a *practionary*.

As I stated in the beginning, knowledge representation is a field of artificial intelligence dedicated to representing information about the world in a form that a computer system can utilize to solve complex tasks. We have explored this topic from background to practice and then on to implications. In *Part I,* we set the stage for the context of the concept by discussing the nature of information, knowledge, and representation, as well as the challenges and opportunities facing KR. In *Part II* we provided a speculative grammar for KR, including the structural role of the universal categories of Firstness, Secondness, and Thirdness, and the terminology, languages,

logic, and models for knowledge representation. That foundation lets us discuss existing frameworks and KR constructs in typologies and knowledge graphs in *Part III*. In our next part, we used these components to build KR and knowledge management systems, including what to construct and what to test and best practices. With a working system in hand, we were then able in *Part V* to discuss fifteen possible application areas of a Peircean approach to KR, covering informative examples in both breadth and depth. These fifteen cases were in addition to the three main application thrusts for this book in knowledge management, data interoperability, and knowledge-based artificial intelligence. As we wrap up our survey of Peirce and KR, we conclude by teasing out some cross-cutting threads and implications from our journey.

I will let you, dear reader, judge whether this practionary achieved its objective of attaining a "third grade of clearness of apprehension" covering "all of the conceivable practical effects" of a Peircean interpretation of knowledge representation. For me, the author struggling to understand a lone genius working in isolation more than a century ago, I have found Peirce's guidance invaluable. Now, as we wrap up our discussion, I would like to stand back from this framework of a practionary and offer some thoughts as to where this journey has led us.

THE SIGN AND INFORMATION THEORETICS

Peirce's understanding of semiosis and signs connects intimately with his views and understanding of logic. Both, I have argued, are themselves prescinded from Peirce's universal categories of Firstness, Secondness, and Thirdness. Indeed, I have argued that the universal categories provide the overarching framework for how we need to organize and categorize our world. The reality of the universal categories is that fundamental.

In Peirce's descriptions of *prescission*, which we introduced in *Chapter 7*:

> "Now, the categories cannot be dissociated in imagination from each other, nor from other ideas. The category of first can be prescinded from second and third, and second can be prescinded from third. But second cannot be prescinded from first, nor third from second. The categories may, I believe, be prescinded from any other one conception, but they cannot be prescinded from some one and indeed many elements. You cannot suppose a first unless that first be something definite and more or less definitely supposed. Finally, though it is easy to distinguish the three categories from one another, it is extremely difficult accurately and sharply to distinguish each from other conceptions so as to hold it in its purity and yet in its full meaning." (1880, CP 1.353)

By this understanding, we prescind Firstness and Secondness from Thirdness, which Peirce reaffirms many places in other ways, often using the term 'degenerate.' Thirdness, too, as we have seen, is where *meaning* resides and is also perhaps best characterized as 'continuity,' the force which Peirce calls *synechism*.*

We certainly need to place knowledge representation in Thirdness. Knowledge

* This topic is more fully discussed in *Appendix A.*

representation is ultimately symbolic because we need to relate it to a computer. KR also is tied intimately to meaning, since we have come to understand that knowledge is information we believe and upon which we may act.

At the same time, we have ample evidence for information, the basis for knowledge, being energetic and physical. Shannon and following researchers have provided quantitative ways and relationships for understanding the nature of information, messages, transmission losses, and entropy. We have seen how structure is also intimately related to this theory, providing the substrate by which free energy gets dissipated in a high-energy input non-equilibrium system, characteristic of life here on Earth. These ideas of structure and canonical forms further help us to think about the architectural designs of our information systems.

Peirce's ideas about information being a limit function, comprised of the full breadth and depth of what we can know about any given thing, approximate, for the totality of things, what may be the limit of information in the absolute. It affirms that much over which we may reason is best expressed as statistics or probabilities. The absolute limit of information, though unknowable, should perhaps be estimable on information theoretic bases.

As we first diagrammed in *Figure 2-1*, I think a deep relationship exists between Shannon's information theoretics and Peircean semiosis. We have the building blocks to tie together absolute information, messages and losses, recipient response, and meaning and entropy. My intuition, still to be tested, is that the absolute limit of what we have come to understand as information is energetic and physical. Nadin, also from a Peircean but different perspective, sees a similar complementarity between information processes and semiotic processes.[2]

In inspecting these relationships, we have seen the advantages of the simple over the complex in our structures, and how recursion and automata make simpler patterns act like engines. The combination of logic (broadly defined to include abduction as Peirce did) and mathematics and entropy, informed by the guidelines of the universal categories and semiosis, should prove a fruitful playground for musing about knowledge representation and our tools to work with it.

PEIRCE: THE PHILOSOPHER OF KR

I discuss Peirce the person and some of his unrelated aspects of philosophy in *Appendix A*. But, as our constant companion through this book, it is now apparent that Peirce is something like a patron saint of knowledge representation. There has not been a single topic within KR for which Peirce does not offer trenchant insights. This illumination is not limited to the direct items of information, knowledge, and representation. Most importantly, Peirce's insights relate to how we think about and conduct knowledge representation, and how we choose practically amongst alternatives moving forward.

It is not surprising that most perhaps best know Peirce is the founder of pragmatism, despite the depth and breadth of his contributions in other areas. The logical endpoints of his inquiries most often lead to the practical aspects of how to act. Para-

doxically, most still treat Peirce as a subject of theoretical discussion and rarely put his guidance into practice. In computer science, for example, no working Peircean semiotic systems exist to my knowledge, and the field has effectively ignored abductive reasoning. The lack of applying Peirce's ideas of pragmatism to real problems feels disappointing. We are overlooking manifest opportunities. It is time to square this circle.

Knowledge and Peirce

Peirce wrote a series of papers arguing against a Cartesian view of the world, a view that places truth solely in the mind and refuses to accept the primacy of external reality.[3] The world is not exclusively one of deductive logic. Objective truth can be approximated by the scientific method. This approximation of truth can always be the subject of inquiry based on different perspectives, or new facts or insights. Beliefs imposed from without or driven by social pressures alone are dead ends to knowledge and understanding. What is real is mostly external to us that we collectively adjudicate through reason and consensus. How we think about, organize and define our problem spaces is central to that process. In the words of Qiwei Chen:[4]

> "Peirce teaches us that the human capacity for knowledge is both unlimited and limited. It is limited in the sense that perfect knowledge cannot be fulfilled in any one individual person and any one particular moment, but as the history of science has shown, every presumed limit has been proved to break down and to be overcome by the progress of knowledge from generation to generation. If it is considered as a process realized in all human beings both past and future, human knowledge is constantly increasing and 'may increase beyond any assignable point,' that is, there is no absolute limit that might restrict it. Indeed, 'an absolute termination of all increase of knowledge is absolutely incognizable, and therefore does not exist.' (CP 5.330)" (p. 47)

Peirce insists that probabilities and chance amidst continuity also direct us to use inductive and abductive logic to anticipate the future. Peirce provides clear guidance on what is information, with meaning defined ultimately upon what we believe and act. Information is a product function of what is intensional that characterizes something with what are extensional connections to external things. Through habit or repeated observations, we may come to believe this information sufficient to act, at which point we are responding to knowledge. This knowledge is not immutable, though it does require a 'surprising fact' or loss of belief to stimulate new inquiry. Abductive reasoning and then the choice of working hypotheses to test follows.

Generating new ideas and testing the truth of them is a logical process that we can formalize. Critical to this process is the proper bounding, definition, and vocabulary upon which to conduct the inquiries. As Peirce argued, we need to express the potentials central to the inquiries for a given topic through a suitable speculative grammar. The guiding lens for how we do this thinking comes from the purpose or nature of the inquiries at hand. In the case of machine learning applied to knowledge bases, this lens, I have argued, should be grounded in Peirce's categories of Firstness,

Secondness, and Thirdness, all geared to feature generation upon which machine learners may operate. The structure of the system should also be oriented to enable (relatively quick and cheap) creation of positive and negative training sets upon which to train the learners. In the end, the nature of how to structure and define knowledge bases depends upon the uses we intend them to fulfill.[5]

We also see, however, that knowledge representation, while symbolic, is not limited to the realm of the symbolic. Some of our knowledge is unconscious or instinctual and may be triggered by the dyadic kinesthetic or by the sudden alarm or alert. The nature of the stimulus (or predicate) giving rise to these signals helps direct what kind of signal and action-response might get triggered. Looking to embed our efforts to understand human language and communication in robotic testbeds should help continue to guide our understanding of these factors.

These strands of argument point to Peircean insights about the nature of knowledge. Peirce's contributions extend to the representational as well. The general ideas of signs and sign-making are the first level of contribution. We also gain much from Peirce's concepts of denotations and indexicality. The rationale for splitting our predicates into the broad groupings of attributes, external relations, and representations is a significant advance over conventional upper ontologies. The fact we have a working knowledge artifact, KBpedia, available for free to use in semantic technology and knowledge representation instantiations is a crucial basis for testing and extending Peirce's ideas about knowledge further.

Enticing connections occur between Peirce's ideas and very topical fields in computer science beyond machine learning, natural language understanding, and robotics. Two of these are possible bridges between description logics and category theory[6] [7] and the emerging field of homotopy type theory. We also have the links to the many promising approaches to computational linguistics as discussed in *Chapter 16*.

Time to Move from Theory to Practice

The semantic Web needs to play a central role in data integration and interoperability. Fortunately, as we have seen in other areas, semantic technologies lend themselves to generic functional software that can be designed for re-use in most any knowledge domain, chiefly by changing the data and ontologies guiding them. This design means that we can build reference libraries of groundings, mappings, and transformations over time and reuse them across enterprises and projects. Functional programming languages align well with the data and schema in knowledge management functions and ontologies and DSLs, domain-specific languages. These prospects parallel the emergence of knowledge-based AI (KBAI), which marries electronic Web knowledge bases with improvements in machine-learning algorithms.

We have ample evidence of the possible areas for which Peirce's ideas may offer unique and valuable insights to all areas of semantic technologies, knowledge representation, and information science. It is time, after a hundred years and many books and learned papers, for how we learn from and use Peirce to move from the theoreti-

cal to the practical.

I hope that this practionary and KBpedia stimulate more practical use and testing of Peirce's insights. Whether KBpedia, an outgrowth of it, or something entirely different, seeing a reference standard emerge for interoperating across multiple datasets and communities would be a potent seed to nucleate still further insights and understanding. We have not yet seen the catalyst that will trigger the cascade of emergent properties one would see from the network effect.

I think one of the reasons we have seen theory prevail over practice with Peirce is the fear of failing, the intimidation of trying to encapsulate a working system that captures the breadth and depth of C.S. Peirce's genius. However, Peirce himself had a pretty sanguine view of his limitations, as he stated in 1906 in "Pragmatism in Retrospect: A Last Formulation":*

> "I here owe my patient reader a confession. It is that when I said that those signs that have a logical interpretant are either general or closely connected with generals, this was not a scientific result, but only a strong impression due to a life-long study of the nature of signs. My excuse for not answering the question scientifically is that I am, as far as I know, a pioneer, or rather a backwoodsman, in the work of clearing and opening up what I call semiotic, that is, the doctrine of the essential nature and fundamental varieties of possible semiosis; and I find the field too vast, the labor too great, for a first-comer. I am, accordingly, obliged to confine myself to the most important questions. The questions of the same particular type as the one I answer on the basis of an impression, which are of about the same importance, exceed four hundred in number; and they are all delicate and difficult, each requiring much search and much caution. At the same time, they are very far from being among the most important of the questions of semiotic. Even if my answer is not exactly correct, it can lead to no great misconception as to the nature of the logical interpretant. There is my apology, such as it may be deemed." (CP 5.488)

Besides espousing 'fallibility,' Peirce took fallibility to heart. We have surely made many mistakes in our efforts to apply Peirce's guidance to a working knowledge representation system in KBpedia. I have perhaps misunderstood what Peirce had to say in multiple areas. Likely, some areas where we have accurately followed Peirce's guidance may simply be wrong. We provide facilities on the KBpedia Web site to communicate those mistakes to us and to participate in KBpedia's ongoing improvement. Charles Sanders Peirce, the philosopher of knowledge representation, would undoubtedly prefer to see us struggle, fail, and improve upon his insights in making our knowledge representations practical, than not try at all.

REASONS TO QUESTION PREMISES

One often finds at the end of a journey that what one thought they would discover or experience on the journey did not prove out. We learn things while on the journey that may cause us to change our initial premises. We encounter new things

* See also Buchler, p. 284,[8]

and take forks in the road. These shifting directions are the idea of fallibility in action, and it is also useful to look at why we got some of our premises wrong and what we have learned.

I remain convinced that enormous opportunities exist for applying Peircean semiotics to knowledge representation. I started with that premise, and end with that premise. With half of the modern U.S. economy based on information, with a rapidly growing percentage globally doing the same, figuring out how to turn that information into knowledge and then to leverage that for economic benefit would be a Rosetta Stone. I also began with the premise that the failures to adopt working knowledge management systems were a combination of technology and culture. That premise, too, remains unchanged, but I also do not have a better idea as to which of culture or technology is more operative. What is clear is that a change in perspective is required to unleash new growth, one which demands energy and management attention.

AI is a Field of KR

I have found Peirce's idea of prescission powerful and subtle. It is powerful because it is an entirely logical, non-psychological way to decide a subsumption relationship.* Prescission, or its verbs *prescind* or *prescinded from*, is the process of comparing two items and seeing if either may exist independent of the other. If so, we say the independent one is prescinded from the dependent one; it is one way to determine a subsumption relationship. The idea of prescission is subtle because, personally, I find getting the direction of the predicate correct is sometimes difficult, and some cases require much thought to discern. In Peirce's terms, 'prescission' is not yet so general for me as being habitual.

When I began this book, I blithely assumed that knowledge representation was a subfield of artificial intelligence. Every taxonomy that I have seen about AI subfields and that included consideration of knowledge representation shows KR as a subsidiary field. I frankly had never questioned the relationship.

However, when considered, mainly using prescission, it becomes clear that KR can exist without artificial intelligence, but AI requires knowledge representation. We can only pursue artificial intelligence via symbolic means, and KR is the translation of information into a symbolic form to instruct a computer. Even if the computer learns on its own, we represent that information in symbolic KR form. This changed premise for the role of KR now enables us to think, perhaps, in broader terms, such as including the ideas of instinct and kinesthetics in the concept. This kind of re-consideration alters the speculative grammar we have for both KR and AI, helpful as we move the fields forward.

So, rather than the definition at the beginning of this book as repeated a few pages prior, we should now state *knowledge representation* is dedicated to symbolizing information about the world in a form that a computer system can utilize to solve

* Or a sibling relationship where precission works in both directions, as for *red* and *blue*, or *squares* and *triangles*.

complex tasks and useful to sub-fields such as artificial intelligence.

Hurdles to be Overcome

Unfortunately, many of the lessons learned deal with the impediments to effective knowledge use and management. Most of the critical obstacles to overcome are not technological, rather social or attitudinal. We need to break away from dichotomous or Cartesian thinking. We need to inculcate a better appreciation for information and knowledge as assets, including the value of purposeful discovery and management. We need to understand the nature of signs and representation and commit to the use of semantic technologies to bridge differences and capture meaning. These are skills that can be learned. However, without the commitment of top-level managers, incentives and processes will not be put in place to encourage their adoption.

Besides these failures of attitude and management, the manner in which we promulgate knowledge management in the organization fails for a further two reasons. One failure is to view knowledge management as its own 'application,' somehow separate and independent of standard work tasks. As we have argued, we need to include distributed, specific functions within current applications, coordinated as services to some form of governing workflow engine and ontologies. At the same time, this realization also opens up opportunities across the board in business process improvements. Knowledge management is itself a leading candidate for these improvements.

The second further failure is in not driving the KM function directly to the knowledge workers and users. Knowledge nurturing, discovery, definition, and use should be directly in the hands of those we pay for those responsibilities. KM, let alone the questions of KR, should not be the responsibility of IT. (And RTFM while you are at it.) Information technology has rightful responsibility for the security, operations, and maintenance of the information infrastructure, and should hold sway on those aspects for KM as well. Hegemony should stop there.

I noted before the advances shown in manufacturing in many of these areas. We are also now witnessing how product and distribution fulfillment centers are starting to see the fruits of automation and robotics. The next frontier is in the white collar, knowledge-oriented portions of the economy. Here is where the next innovation wave is due. Peircean approaches to knowledge representation combined with semantic technologies are the bright path to follow moving forward.

Of Crystals and Robots

As first noted in *Chapter 11*, Peirce famously claimed thought does not necessarily occur in the brain, that we may find thought in the work of crystals and bees, inanimate matter and insects. (1906, CP 4.551) We have also talked about its applicability to robots and AI. The most important lesson to emerge from our investigations might well be that some fundamental truths underlie the universal categories. During the second great wave of artificial intelligence in 1988 Daniel Dennett wrote that:[9]

> "AI is, in large measure, philosophy. It is often directly concerned with instantly recognizable questions: What is mind? What is meaning? What is reasoning and rationality? What are the necessary conditions for the recognition of objects in perception? How are decisions made and justified?" (p. 283)

Peirce, I believe, gives us guidance on all of these questions. Still, as a voice of theory, not yet validated by practice, Peirce may point the way, yet leaves many questions tantalizingly open.

Peirce understood graph structures. His language formulations and understanding of relations are at the forefront of much current computational linguistic research. His conception of mind embraced the external world if not was dominated by it. His interest in moving algebra to geometric forms and then topology fits well with the probability landscapes that now inform much thinking in machine learning and statistical mechanics. His writing about logic machines and electrical computation indicate he was anticipating much that has come to pass.* His attempts to construct more elaborate and structured sign systems foreshadowed many aspects of ontologies and knowledge graphs. We can construct every idea that Peirce advocated from realities in the external world agreed to by the community. He was clear about the fundamental concepts of reality, existence, actuality, being, truth, chance, and continuity.

The neuroscientist Eugen Izhikevich in a recent debate with Roger Penrose said:[10] "We are at the stage of understanding consciousness as we were for information before Shannon. We lack a theory and definition for it that is agreed as likely correct." That is a fair assessment. Hopefully, we have taken some tiny steps on the path to that theory.

We want a theory grounded in reality, including quantum reality. We want a theory that embraces Shannon's information theory, yet one that extends its embrace to include meaning. We want a theory of signification and representation that can model energy fluxes that extend from inanimate matter to human symbol systems. We want a theory that captures the logic and message content of human language, one that can effectively communicate a symbolic representation to computers. We want a theory with a set of primitives that give us these capabilities while being small and straightforward. It will take many minds and much tinkering to complete the journey on this path that Charles Sanders Peirce has blazed for us.

Chapter Notes

1. Parker, K. A., *The Continuity of Peirce's Thought*, Nashville: Vanderbilt University Press, 1998.
2. Nadin, M., "Information and Semiotic Processes The Semiotics of Computation," *Cybernetics & Human Knowing*, vol. 18, 2011, pp. 153–175.
3. Peirce wrote three articles for the *Journal of Speculative Philosophy* in 1868-9: "Questions Concerning Certain Faculties Claimed for Man"(CP 5.213-63; EP 2:11-27); "Some Consequences of Four Incapacities" (CP 5.264-317; EP 1:28-55); and "Grounds of Validity of the Laws of Logic: Further Consequences of Four Incapacities"

* See discussion on this topic in *Appendix A*.

(CP 5.318-357; EP 1:56-82).

4. Chen, Q., "Some Aspects of Peirce's Theory of Knowledge," *Living Doubt*, G. Debrock and M. Hulswit, eds., Dordrecht: Springer Netherlands, 1994, pp. 43–53.
5. Bergman, M. K., "A Speculative Grammar for Knowledge Bases," *AI3:::Adaptive Information*, Jun. 2016.
6. Patterson, E., "Knowledge Representation in Bicategories of Relations," *arXiv*, vol. 1706, Jun. 2017.
7. Zalamea, F., *Synthetic Philosophy of Contemporary Mathematics*, New York, NY: Urbanomic Sequence Press, 2012.
8. Peirce, C. S., *Philosophical Writings of Peirce: Selected Writings*, New York NY: Routledge and Kegan Paul Ltd., reissued by Dover Publications, 1940.
9. Dennett, D. C., "When Philosophers Encounter Artificial Intelligence," *Daedalus*, vol. 117, 1988, pp. 283–295.
10. Qualcomm Corporation, *AI Debate: How Far Can the AI Revolution Go?*, 2018.