Michael K. Bergman

A Knowledge Representation Practionary

Guidelines Based on Charles Sanders Peirce

Springer

A Knowledge Representation Practionary

Michael K. Bergman

A Knowledge Representation Practionary

Guidelines Based on Charles Sanders Peirce

Springer

Michael K. Bergman
Cognonto Corporation
Coralville, IA, USA

ISBN 978-3-319-98091-1 ISBN 978-3-319-98092-8 (eBook)
https://doi.org/10.1007/978-3-319-98092-8

Library of Congress Control Number: 2018954477

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To
Wendy

Preface

Human language is not the starting point for knowledge representation. Our utterances or our symbols are not the basis for what we desire to convey; they are only representations. Knowledge, the actionable side of information, is rooted in something more fundamental than language. What that something may be is what this book is about.

Competing factions have claimed truth since at least the beginning of communication. Who knows, maybe bees, whales, dingoes, and apes also have communities believing different things as true, perhaps even leading to conflict. As humans, we know from wars, missed opportunities, and personal misunderstandings the tragedy that different premises of truth may bring. We have to admit if we want to represent human knowledge to computers that we humans have not done such a hot job representing knowledge to ourselves. Since we are starting out on a journey here to explore knowledge representation (KR) for knowledge management, artificial intelligence, and other purposes, more than a bit of humility seems in order.

Information, by no means a uniformly understood concept, arises from a broader context than gestures, symbols, or sounds. For some, information is energy or when missing is entropy, the nuts-and-bits of messages. For some, information is meaning. That we continue to use 'information' in these senses, and more, in fact, tells us these senses are properly within the boundaries of the concept. Still, even if we can clear the hurdle of grokking information, we have the next obstacle of deciphering what is knowledge, that which next lies directly on our path. Further, of course, we then need to record somehow and convey all of this if we are to represent the knowledge we have gained to others. Like I say, if we have a hard time communicating all of this to other humans, what can we say about our ability to do so to machines and AI?

But maybe I overthink this. Any tasks us humans do using information that we can automate with acceptable performance may lead to more efficiency and perhaps more job satisfaction for the workers involved, maybe even more wealth. Conversely, maybe this automation leads to loss of jobs for the workers. I do know, however, if we are ever to rely upon machines to work on our behalf, requiring little or no oversight, then we need to figure out what this knowledge is and how to represent it to the machine. Such is the task of KR. What I try to provide in this book is a way to think and a practical guidebook of sorts for how to approach the questions of computers and knowledge.

The world is real. It exists independent of us or how we may think about it, though our thoughts are also part of our reality. Human history fills but a small thimble yet through the application of reason and truth-testing, including, since the Enlightenment, the scientific method, we humans have increasingly unveiled the truths of Nature, in the process creating wealth and comfort never before seen. Artificial intelligence (AI) will undoubtedly accelerate this trend. How fast that acceleration occurs is, in part, a function of how good we get at representing our knowledge. These representations are the encodings by which intelligent machines will work on our behalf. My quest in this treatise is to help promote this trend. I believe this quest to be noble and, in any case, inevitable. I believe there is something in our nature that compels us to pursue the path of useful information leading to knowledge.

The past decade was a golden one in advances in AI. We can now voice commands and requests to our phones and devices acting as virtual assistants. We are on the verge of self-driving vehicles and automation of routine knowledge worker tasks. Still, the deep learning that underlies many of these advances is an opaque, black box of indecipherable inferences. We don't know why some of this magic works or what the representations are upon which machines draw these inferences. For further advances to occur, for general AI or cognition to arise in silico, I believe we will need better ways to represent knowledge, reflective of the nature of information and its integral role in the real world.

I have had a passion for the nature and role of information throughout my professional life. I originally trained as an evolutionary biologist and population geneticist. Since my graduate days, I have replaced my focus on biological information with one based on digital information and computers. My passion has been on the role of information—biological or cultural—to confer adaptive advantage to deal with an uncertain future and as a means of generating economic wealth. My intuition—really, my underlying belief—is that there are commonalities between biological and cultural information. I have been seeking insights into this intuition for decades.

One of my first forays into information technology was a data warehousing venture, where the idea was to find ways to connect structured databases that, in native form, were stand-alone and unconnected. This venture coincided with the explosive growth of the initial Internet. To support the exploding content, we observed that large content suppliers were populating their web sites with searchable, dynamic databases, hidden from the search engines of that time (before Google's inception). We named this phenomenon the 'deep web' and did much to define its huge extent and figure out ways to mine it. We saw that, in aggregate, the web was becoming a giant, global data warehouse, though largely populated by text content and less so by structured data. We shifted our venture emphasis to text and discovery. This shift raised the perplexing question of how to place information in text onto a common, equal basis to the information in a database, such as a structured record. (Yeah, I know, kind of a weird question.)

Tim Berners-Lee, inventor of the World Wide Web, and colleagues put forward a vision of the Semantic Web in a *Scientific American* article in 2000.[1] The article painted a picture of globally interconnected data leveraged by agents or bots designed to make our lives easier and more automated. The late Douglas Adams, of *Doctor Who* and *The Hitchhiker's Guide to the Galaxy* fame, had presciently produced a fascinating and entertaining TV program on the same topic for BBC2 about 10 years earlier. Called *Hyperland*, you can see this self-labeled 'fantasy documentary' from 1990 in its entirety on YouTube. The 50-min presentation, written by and starring Adams as the protagonist having a fantasy dream, features Tom, the semantic simulacrum (actually, Tom Baker from *Doctor Who*). Tom is the "obsequious, and fully customizable" personal software agent who introduces, anticipates, and guides Adams through what is a Semantic Web of interconnected information. Laptops (actually an early Apple), pointing devices, icons, and avatars sprinkle this *tour de force* in an uncanny glimpse into the (now) future.

One of the premises of the Semantic Web is to place what we now call unstructured, semi-structured, and structured information onto a common footing. The approach uses the RDF (Resource Description Framework) data model. RDF provided an answer to my question of how to combine data with text. I am sure there were other data models out there at the time that could have perhaps given me the way forward, but I did not discover them. It took RDF and its basic *subject-predicate-object* (*s-p-o*) 'triple' assertion to show me the way ahead. It was not only a light going on once I understood but the opening of a door to a whole new world of thinking about knowledge representation.

The usefulness of ideas behind the Semantic Web and the semantic technologies supporting it lured me to switch emphasis again. I founded a new company with Frédérick Giasson, and we proceeded to provide semantic technology solutions to enterprises over the next 10 years. The Web today is almost unrecognizable from the Web of 15 years ago. If one assumes that Web technologies tend to have a 5-year or so period of turnover, we have gone through three to four generations of change on the Web since the initial vision for the Semantic Web.

Many of our engagements were proprietary, though we did provide three notable open source projects. We developed a general semantic platform for ontology (knowledge graph) and data management, the still-active Open Semantic Framework project. To help information interoperate, we created UMBEL, a subset of Cyc and a contributor to our current efforts, as a set of reference concepts that users can share across different Web datasets. Based on that experience, we designed a successor reference knowledge structure, KBpedia, a combination of upper knowledge graph and leading public knowledge bases. We talk much about KBpedia throughout since it is this book's reference knowledge structure.

The marrying of electronic Web knowledge bases—such as Wikipedia or internal ones like the Google search index or its Knowledge Graph—with improvements in machine learning algorithms is systematically mowing down what used to be called

[1] Berners-Lee, T., Lassila, O., and Hendler, J., "The Semantic Web," *Scientific American Magazine*, 2001.

the Grand Challenges of computing, such as machine translation or language understanding. Sensors are also now entering the picture, from our phones to our homes and our cars, that exposes the higher-order requirement for data integration combined with semantics. Natural language processing (NLP) kits have improved in accuracy and execution speed; many semantic tasks such as tagging or categorizing or questioning already perform at acceptable levels for most projects. We naturally call the marriage of these knowledge sources with AI 'knowledge-based artificial intelligence.' KBAI is one of the potential payoffs that would arise from better ways to represent knowledge and thus is a common theme throughout the book.

Combining information goes beyond the technical challenges of matching forms and formats. We need to tackle the question of meaning, inextricably entwined with context and perspective. Cinemaphiles will readily recognize Akira Kurosawa's *Rashomon* film of 1951. In the 1960s, one of the most popular book series was Lawrence Durrell's *The Alexandria Quartet*. Both, each in its way, tried to get at the question of what is the truth by telling the same story from the perspective of different protagonists. Whether you saw Kurosawa's movie or read Durrell's books, you know the punchline: truth is very different depending on the point of view and experience—including self-interest and delusion—of each protagonist.

All of us recognize this phenomenon of the blind man's view of the elephant. The problem we are trying to solve is how to connect information meaningfully. For that, we need to somehow capture the ideas of perspective and context, as well as the usual vagaries of imprecise semantics. Root cause analysis for what it takes to achieve meaningful, interoperable information suggests one pivotal factor is to describe source content adequately in context to its use. Capturing and reflecting context is essential if we are to get information sources to work together, a capability we give the fancy label of 'interoperability.' We also need to assemble and represent this information such that we can reason over it and test new knowledge against it, a structural form we call a 'knowledge graph.' All of this requires a logical and coherent theory—a *grounding*—for how to represent knowledge.

Our client efforts over the past decade were converging on design thoughts about the nature of information and how to signify and communicate it. The bases of an overall philosophy regarding our work emerged around the teachings of Charles Sanders Peirce and Claude Shannon, each explicating one of the boundary senses of information. Shannon emphasized the message and mechanical aspects of information; Peirce emphasized meaning in both breadth and depth. In the combination, we see semantics and groundings as essential to convey accurate messages. Simple forms, so long as they are correct, are always preferred over complex ones because message transmittal is more efficient and less subject to losses (inaccuracies). How we could represent these structures in graphs affirmed the structural correctness of our design approach. The now visible reawakening of artificial intelligence helps to put the Semantic Web in its proper place: a key subpart, but still a subset, of AI.

I first encountered Charles S. Peirce from the writings of John Sowa about a decade ago. Sowa's writings are an excellent starting point for learning about logic and ontologies, especially his articles on Peirce and signs.[2] Early on it was clear to me that knowledge modeling needed to focus on the inherent meaning of things and concepts, not their surface forms and labels. Sowa helped pique my interest that Peirce's theory of semiotics was perhaps the right basis for getting at these ideas.

In the decade since that first encounter, I have based some writings on Peirce's insights. I have developed a fascination with his life and teachings and thoughts on many topics. I have become convinced that Peirce—an American philosopher, logician, scientist, and mathematician—was possibly one of the greatest thinkers ever. While the current renaissance in artificial intelligence can certainly point to the seminal contributions of George Boole, Shannon, Alan Turing, and John von Neumann in computing and information theory (among many others), my view, not alone, is that C.S. Peirce belongs in those ranks from the perspective of knowledge representation, the *meaning* of information, and hewing to reality.

The importance of studying Peirce for me has been to tease out those principles, design bases, and mindsets that can apply Peircean thinking to the modern challenge of knowledge representation. This knowledge representation is like Peirce's categorization of science or signs but is broader still in needing to capture the nature of relations and attributes and how they become building blocks to predicates and assertions. In turn, we need to subject these constructs to logical tests to provide a defensible basis for what is knowledge and truth given the current information. Then, all of these representations need to be put forward in a manner (symbolic representation) that is machine readable and computable.

In reading and studying Peirce for more than a decade, it has become clear that he had insights and guidance on every single aspect of this broader KR problem. My objective has been to take these piece parts (Peirce parts?) and recombine them into a whole consistent with Peirce's *architectonic*. How can Peirce's thinking be decomposed into its most primitive assumptions to build up a new KR representation? These are the points I argue in the book while also sharing the experience of how we may integrate these viewpoints into working knowledge management systems.

I have no intent for balance in this exposition. There are wonderful textbooks and handbooks available if you are seeking a neutral presentation on knowledge representation in computer and information science. The lens I use is strictly that of Peirce and his views that contribute to an understanding of knowledge representation, at least how I read and understand those views. Peirce further guides the scope and organization of this book. One of Peirce's signal contributions was the philosophy of pragmatism, according to a specific maxim and a recommended methodology to follow, what the Peirce scholar Kelly Parker calls a 'practionary.' To my knowledge, this book employs this Peircean methodology for the first time. Given this emphasis, we will by necessity need to tackle many Peircean concepts, some with arcane or jaw-breaking labels. That is a small price to pay to gain entry into Peirce's brilliant insights.

I also minimize math and equations in the book. I provide many salient references for exploring topics further. I try to emphasize how to think and organize.

[2] Use https://www.google.com/search?as_q="peirce"&as_sitesearch=jfsowa.com for a listing.

I avoid cookbook steps or prescriptive techniques or methods. I do not recommend specific tools. Rather, because of the coherence of Peirce's views, I use how I understand him and his writings, including interpretations by others, to bring a consistent approach, logic, and mindset to the question of knowledge representation. By straddling today's two separate worlds of Peirce scholarship and knowledge representation, I perhaps risk disappointing both camps. One of my points, though, is that the camps should be separated no longer.

I would first like to thank my colleague and business partner, Frédérick Giasson, for his creativity and effort in our commercial ventures over the past decade. He was not only the implementer of the many systems we developed, and a constant fount of ideas and innovation, but a great friend and a calm and cool influence during those engagements. Though I am the recorder of the results in this book, he deserves co-billing for why and how this book came into being.

I want to thank those who have encouraged me over many years to write this book, including from many commenters on my *AI3:::Adaptive Information* blog. I especially thank Fred, Steve Ardire, Alianna Maren, Alan Morrison, Amit Sheth, and Peter Yim for their encouragement. I further thank Amit for his kind efforts to help me find and secure a publisher.

I thank my former colleague, Jacquie Bokow, for early editorial assistance and advice. I much appreciate the complete and detailed reviews I got on the first draft from Michael Buckland, Scott David, Rob Hillard, John Huntley, and Jack Park. I am grateful for the commentary and errors found in my readings of Peirce from Jon Alan Schmidt and Edwina Taborsky, as well as insights I have gained from the Peirce-L discussion group. I further thank William Anderson, Andreas Blumauer, Fred Giasson, Alan Morrison, Amit Sheth, Aleksander Smywiński-Pohl, Bobbin Teagarden, and Tom Tiahrt for their reviews and commentary. Despite their best efforts to find and correct my errors and to make great suggestions, I am sure that errors remain, which are entirely my responsibility. I ask your forbearance for any errors or oversights. I lastly thank Susan Lagerstrom-Fife and Caroline Flanagan for helping to shepherd the manuscript through the publication process.

I find it wondrous that the human species has come to learn and master symbols. That mastery, in turn, has broken the shackles of organic evolution and has put into our hands and minds the very means and structure of information itself. The *lingua franca* for doing so is knowledge representation, best done, I believe, following the guidelines of Charles Sanders Peirce.

Coralville, IA
October 2018

Michael K. Bergman

In-line Citations

Here are the conventions and sources used for quotations for Peirce's writings as used in the book. Items separated by a period or colon are page or clause number.

Abbrv.	Example	Source
CP	CP 1.343	Peirce, C. S., *The Collected Papers of Charles Sanders Peirce, 8 Volumes*, Cambridge, MA: Harvard University Press, 1931.
EP	EP 2:43	Peirce, C. S., *The Essential Peirce: Selected Philosophical Writings, Vol 1 (1867–1893)*, Bloomington: Indiana University Press, 1992. Peirce, C. S., *The Essential Peirce: Selected Philosophical Writings, Vol 2 (1893–1913)*, Bloomington: Indiana University Press, 1998.
MS	MS 32	Robin, R. S., *Annotated Catalogue of the Papers of Charles S. Peirce*, Amherst, Massachusetts: The University of Massachusetts Press, 1967. Robin, R. S., "The Peirce Papers: A Supplementary Catalogue," *Transactions of the Charles S. Peirce Society*, 1971, pp. 37–57.
NEM	NEM 4:83	Peirce, C. S., *The New Elements of Mathematics by Charles S. Peirce*, Hague, Netherlands: The, Mouton Publishers, 1976 (four volumes).
W	W 3:266	Peirce, C. S., *The Writings of Charles S. Peirce—Chronological Editions Vo1 8,* compiler, Peirce Edition Project, Bloomington: Indiana University Press, 1982–2009 (by volume).

I have attempted to date each Peirce quote, given the current tendency in scholarship and its usefulness to place his views into a chronology. Unlike most practice, I list the year first and then the citation.

Permission

Significant portions of the material in this book were first published on the author's *AI3:::Adaptive Information* blog, at http://mkbergman.com. We thank the author for the permission to use this copyrighted material.

Contents

Chapter 1
Introduction

Knowledge representation, of course, deals with knowledge, itself based on information. Knowledge representation is shorthand for how to represent human symbolic information and knowledge to computers, preferably in the most effective manner. Formally, and the working definition for this book, *knowledge representation*[1] 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. KR applications range from semantic technologies and machine learning and artificial intelligence to information integration, data interoperability, and natural language understanding.

I am not even-handed in this book. My explicit purpose is to offer a fresh viewpoint on knowledge representation and ontology engineering, informed by a variety of projects over the past dozen years, and guided by the principles of Charles Sanders Peirce, as I best understand them. Many others have different perspectives on knowledge representation. For more balance and to understand this diversity, I recommend the excellent KR reference texts by van Harmelan [1] or Brachman and Levesque [2].

C.S. Peirce (1839–1914), pronounced 'purse,' was an American logician, scientist, mathematician, and philosopher of the first rank. His profound insights and writings spanned a half-century, and cover topics ranging from the nature of knowledge and epistemology to metaphysics and cosmology.[2] His universal categories of Firstness, Secondness, and Thirdness provide the mindset and theories that guide this book. Peirce, along with Gottlob Frege, is acknowledged as a founder of predicate calculus, to which Peirce provided a notation system, and which formed the basis of first-order logic. Peirce's theory of signs and sign-making, semiosis, is a seminal understanding of icons, indexes, and symbols, and the way we perceive and

[1] Many of the italicized terms in this book are defined when first used and listed in the *Glossary*.

[2] Appendix A, from which I borrow these two sentences, is a summary biography and reading suggestions for Charles Sanders Peirce. He is also referenced in the literature as Peirce, Charles Peirce, C.S. Peirce, or CSP.

M. K. Bergman, *A Knowledge Representation Practionary*,
https://doi.org/10.1007/978-3-319-98092-8_1

understand objects. Peirce's semiosis (*semeiosis*, his preferred spelling) and approach arguably provide the logical basis for description logics and other aspects underlying the semantic Web building blocks of the RDF data model and, eventually, the OWL language. Peirce is the acknowledged founder of pragmatism, the philosophy of linking practice and theory in a process akin to the scientific method. He was also the first formulator of existential graphs, a basis to the field now known as model theory [3], and the basis for conceptual graphs, a KR formalism. No aspect of knowledge representation exceeded his grasp.

This book also weaves the open-source knowledge artifact, KBpedia, through its later chapters and observations. KBpedia combines the information from multiple public knowledge bases, prominently including Wikipedia and Wikidata, under the conceptual structure of the KBpedia Knowledge Ontology (KKO), a *knowledge graph* organized according to the Peircean universal categories. KBpedia's 55,000 reference concepts, classified into 85 mostly separate *typologies*, and with access to millions of notable *entities* and *events*, is a modular resource that may be leveraged or expanded for particular domain purposes. However, the confederation between this book and KBpedia is loose. Each stands on its own without reliance on the other.

We have witnessed enormous and mind-boggling strides over the past decade in *artificial intelligence*. *Machine learning* has leveraged massive *knowledge bases* to deliver breakthrough capabilities in automated question answering and intelligent virtual assistants. Deep learning, with its mostly indecipherable black-box layers, has enabled automatic recognition of voice, images, and patterns at speeds and accuracies often exceeding that of humans.

Still, we struggle to integrate information, get data to interoperate, or discover or manage knowledge. Our current AI techniques appear close to reaching limits, including whether we even understand what those techniques are doing. Peircean ideas hold the tantalizing prospect to unlock better ways to represent knowledge. KR is the foundation upon which, I believe, next breakthroughs will come. I believe Peircean ideas provide the way to better represent human knowledge such that AI-powered computers can organize, index, reference, and cross-check information in any digital form. This prospect will obliterate current boundaries to information sharing. If the past is a guide, innovation, transformation, and wealth will follow.

Structure of the Book

This book is structured into parts and chapters. The central portion of the book (Parts II–IV) reflects C.S. Peirce's universal categories of Firstness, Secondness, and Thirdness. Across nearly five decades of writings, Peirce likens the universal categories to more than 60 different expressions (Table 6.2). The expression used for this central portion of the book is Peirce's logic triad of *grammar* (1ns), logics and tools (or *critic*) (2ns), and methods (or *methodeutic*) (3ns).[3] We use this triadic organization to explain the what and how for a working knowledge representation system, with frequent reference to KBpedia.

Parts I and V are bookends around this central portion. Part I, the opening bookend, provides the context for why one should be interested in the topic of knowledge representation and what kind of functions KR should fulfill. Part V, the closing bookend, provides practical speculation for what kinds of benefits and applications may result from a working KR system built according to Peircean principles. A couple of chapters tee up this structure.

The structural approach of this book is consistent with Peirce's *pragmatic maxim* to achieve the "third grade of clearness of apprehension" (1878, W 3:266)[4] covering "all of the conceivable practical effects," regarding an understanding of something. If a *dictionary* is for the definition of terms, a *practionary* is for the definition of methods and potential applications resulting from an explication of a domain. In the case of this book, that domain is *knowledge representation*.[5]

To my knowledge, this is the only Peirce book dedicated solely to knowledge representation, and the only KR book exclusively devoted to Peirce.[6] Some reviewers of drafts of this book have suggested splitting the book into multiple parts. I admit there is some logic to that suggestion. Early chapters discuss contexts of information theory, economics, and social circumstances. Middle parts of the book are theoretical, even philosophical, that evolve into how-to and practice. The latter parts of the book are speculative and span potential applications in breadth and depth. My answer in keeping these parts together is to try to be faithful to this overall ideal of a Peircean *practionary*. I welcome you to a soup-to-nuts banquet of Peircean perspectives on the challenge of knowledge representation.

Overview of Contents

Before we start the formal structure of the book, we begin with this chapter and then Chap. 2 discussing the core concepts of *information*, *knowledge*, and *representation*. Gregory Bateson defined information as the "difference that makes a difference." Claude Shannon, the founder of information theory, emphasized the

[3] 1ns, 2ns, and 3ns are shorthand for Firstness, Secondness, and Thirdness, respectively.

[4] See the note on Abbreviations after the *Preface* for the citations format for the Peirce quotations used throughout.

[5] The term of *practionary* comes from Kelly Parker based on his study of Peirce [3]; I thank him for graciously allowing me to use the term.

[6] John Sowa's 1999 book, *Knowledge Representation: Logical, Philosophical, and Computational Foundations* (Brooks Cole Publishing Co., Pacific Grove, CA, 2000), was much influenced by Peirce. Sowa in his work on conceptual graphs builds directly from Peirce's existential graphs. However, Sowa's book is not based exclusively on Peirce, nor is his ontology (see http://www.jfsowa.com/ontology/toplevel.htm). Still, Sowa's is the closest Peirce-KR treatment to my knowledge without being solely based on him.

engineering aspect of information, defining it as a message or sequence of messages communicated over a channel; he specifically excluded meaning. Peirce emphasized meaning and related it to the triadic relationship between immediate object, representation, and interpretation. We associate knowledge and its discovery with terms such as open, dynamic, belief, judgment, interpretation, logic, coherence, context, reality, and truth. Peirce's pragmatic view is that knowledge is fallible information that we believe sufficiently upon which to act. I argue in this book, consistent with Peirce, that knowledge representation is a complete triadic sign, with the meaning of the information conveyed by its symbolic representation and context, as understood and acted upon by the interpreting agent. A challenge of knowledge representation is to find structured representations of information—including meaning—that can be simply expressed and efficiently conveyed.

We then begin the structural portions of the book. Part I and its three chapters attempt to place knowledge representation, as practiced today, in context. Chapter 3 describes the situation and importance of information to enterprises and society. Knowledge representation is a primary driver for using computers as a means to improve the economic well-being of all peoples. Solow, a student of Schumpeter, had the insight into two papers in the 1950s, for which he won a Nobel Prize, that technological change is the 'residual' leftover from empirical growth once we remove the traditional inputs of labor and capital. This residual is what we now call total-factor productivity. Romer's subsequent work internalized this factor as a function of information and knowledge, which in contrast became the endogenous growth model. Innovation and its grounding in knowledge had finally assumed its central, internal role in economists' understanding of economic growth. Unlike the historical and traditional ways of measuring assets—based on the tangible factors of labor, capital, land, and equipment—information is an intangible asset. If we are to improve our management and use of information, we need to understand how much value we routinely throw away.

Once we survey the situation, Chap. 4 begins to surface some of the opportunities. The path to knowledge-based artificial intelligence (KBAI) directly coincides with a framework to aid data interoperability and responsive knowledge management (KM). A knowledge graph (or ontology) provides the overall schema, and semantic technologies give us a basis to make logical inferences across the knowledge structure and to enable tie-ins to new information sources. We support this graph structure with a platform of search, disambiguation, mapping, and transformation functions, all of which work together to help achieve data interoperability. KBAI is the use of large statistical or knowledge bases to inform feature selection for machine-based learning algorithms. We can apply these same techniques to the infrastructural foundations of KBAI systems in such areas as data integration, mapping to new external structure and information, hypothesis testing, diagnostics and predictions, and myriad other uses to which researchers for decades hoped AI would contribute. We apply natural language processing to these knowledge bases informed by semantic technologies.

To complete the context, we discuss other vital precepts (or premises) in Chap. 5. Knowledge should express a coherent reality, to reflect a logical consistency and structure that comports with our observations about the world. How we represent reality has syntactic variation and ambiguities of a semantic nature that can only be resolved by context. A hub-and-spoke design with a canonical data model is a superior way to organize, manipulate, and manage input information. By understanding the sources of semantic heterogeneity, we set the basis for extracting meaning and resolving ambiguities. Once we resolve ('disambiguate') the source information, we need to organize it into 'natural' classes and relate those classes coherently and consistently to one another. This organization takes the form of a knowledge graph. Traditional relational databases do not; they are inflexible and fragile when the nature (schema) of the world changes, and require expensive re-architecting in the face of new knowledge or new relationships.

We next embark on the central portion of our thesis, Parts II–IV. Part II covers the grammar of knowledge representation. I discuss in detail Peirce's universal categories of Firstness, Secondness, and Thirdness in Chap. 6. The ideas behind Peircean pragmatism are how to think about signs and representations (*semiosis*); logically reason and handle new knowledge (*abduction*) and probabilities (*induction*); make economic research choices (*pragmatic maxim*); categorize; and let the scientific method inform our inquiry. The connections of Peirce's sign theory, his threefold logic of deduction-induction-abduction, the importance of the scientific method, and his understanding about a community of inquiry have all fed my intuition that Peirce was on to some fundamental insights suitable to knowledge representation. We can summarize Firstness as unexpressed possibilities; Secondness as the particular instances that may populate our information space; and Thirdness as general types based on logical, shared attributes. Scholars of Peirce acknowledge how infused his writings on logic, semiosis, philosophy, and knowledge are with the idea of 'threes.' Understanding, inquiry, and knowledge require this irreducible structure; connections, meaning, and communication depend on all three components, standing in relation to one another and subject to interpretation by multiple agents in multiple ways.

Our next topic within the *speculative grammar* of the KR space addresses basic terminology, which we cover in Chap. 7. We begin our analysis with the relevant 'things' (nouns, which are *entities, events*, *types*, or *concepts*) that populate our world and how we organize them. We pair these things with three kinds of internal and external relations to other things. Attributes are the intensional characteristics of an *object*, event, entity, type (when viewed as an *instance*), or concept. External relations are actions or assertions between an event, entity, type, or concept and another particular or general. Representations are signs and the means by which we point to, draw attention to, or designate, denote, or describe a specific object, entity, event, type, or general. We now know that *attributes* are a Firstness in the universal categories; that Secondness captures all events, entities, and relations; and that Thirdness provides the types, context, meaning, and ways to indicate what we refer to in the world.

Chapter 8 presents the logic basis and introduces the actual vocabularies and languages to express this grammar. Knowledge graphs and knowledge bases need to be comprehensive for their applicable domains of use, populated with 'vivid' knowledge. We use deductive logic to infer hierarchical relationships, create forward and backward chains, check if domains and ranges are consistent for assertions, assemble attributes applicable to classes based on member attributes, conform with transitivity and cardinality assertions, and test virtually all statements of fact within a knowledge base. We want a knowledge representation (KR) language that can model and capture intensional and extensional relations; one that potentially embraces all three kinds of inferential logic; that is decidable; one that is compatible with a design reflective of particulars and generals; and one that is open world in keeping with the nature of knowledge. Our choice for the knowledge graph is the W3C standard of OWL 2 (the Web Ontology Language), though other choices may be just as valid.

Using this grammatical and language foundation, Part III transitions to discuss the working components of a KR system. In Chap. 9, I argue the importance of openness and keeping an open design. Open content works to promote derivative and reinforcing factors in open knowledge, education, and government. Open standards encourage collaboration and make it easier for data and programs to interoperate. Open data in public knowledge bases are a driver of recent AI advances in knowledge. Open also means we can obtain our knowledge from anywhere. Our knowledge graphs useful to a range of actors must reflect the languages and labels meaningful to those actors. We use reference concepts (RCs) to provide fixed points in the information space for linking with external content. We now introduce KBpedia to the remainder of the discussion. We use RDF as a kind of 'universal solvent' to model most any data form. We match this flexible representation with the ability to handle semantic differences using OWL 2, providing an open standard to interoperate with open (or proprietary) content.

In Chap. 10, we shift the emphasis to modular, expandable typologies. The idea of a SuperType is equivalent to the root node of a typology, wherein we relate multiple entity types with similar essences and characteristics to one another via a natural classification. Our typology design has arisen from the intersection of (1) our efforts with SuperTypes to create a computable structure that uses powerful disjoint assertions; (2) an appreciation of the importance of entity types as a focus of knowledge base terminology; and (3) our efforts to segregate entities from other constructs of knowledge bases, including attributes, relations, and annotations. Unlike more interconnected knowledge graphs (which can have many network linkages), typologies are organized strictly along these lines of shared attributes, which is both simpler and also provides an orthogonal means for investigating type-class membership. The idea of nested, hierarchical types organized into broad branches of different entity typologies also offers a flexible design for interoperating with a diversity of worldviews and degrees of specificity.

Typologies are one component of our knowledge graphs and knowledge bases, to which we shift our attention in Chap. 11. Relations between nodes, different than those of a hierarchical or subsumptive nature, provide still different structural connections across the knowledge graph. Besides graph theory, the field draws on methods including statistical mechanics from physics, data mining and information visualization from computer science, inferential modeling from statistics, and social structure from sociology. Graph theory and network science are the suitable disciplines for a variety of information structures and many additional classes of problems. We see the usefulness of graph theory to linguistics by the various knowledge bases such as WordNet (in multiple languages) and VerbNet. Domain ontologies emphasize conceptual relationships over lexicographic ones for a given knowledge domain. Furthermore, if we sufficiently populate a knowledge graph with accurate instance data, often from various knowledge bases, then ontologies can also be the guiding structures for efficient machine learning and artificial intelligence. We want knowledge sources, preferably knowledge bases, to contribute the actual instance data to populate our ontology graph structures.

We have now discussed all of the conceptual underpinnings to a knowledge representation *system*. Part IV, also spread over three chapters, presents how these components are now combined to build a working platform. In Chap. 12, we outline the basic KR platform and the accompanying knowledge management (KM) capabilities it should support. The platform should perform these tasks: insert and update concepts in the upper ontology; update and manage attributes and track specific entities as new sources of data are entered into the system; establish coherent linkages and relations between things; ensure that these updates and changes are done wholly and consistently while satisfying the logic already in place; update how we name and refer to things as we encounter variants; understand and tag our content workflows such that we can determine provenance and authority and track our content; and do these tasks using knowledge workers, who already have current duties and responsibilities. These requirements mean that use and updates of the semantic technologies portion, the organizing basis for the knowledge in the first place, must be part of daily routines and work tasking, subject to management and incentives.

Once a platform is available, it is time to build out the system, the topic of Chap. 13. Critical work tasks of any new domain installation are the creation of the domain knowledge graph and its population with relevant instance data. Most of the implementation effort is to conceptualize (in a knowledge graph) the structure of the new domain and to populate it with instances (data). In a proof-of-concept phase, the least-effort path is to leverage KBpedia or portions of it as is, make few changes to the knowledge graph, and populate and test local instance data. You may proceed to create the domain knowledge graph from pruning and additions to the base KBpedia structure, or from a more customized format. If KBpedia is the starting basis for the modified domain ontology, and if we test for logic and consistency as we make incremental changes, then we are able to evolve the domain knowledge graph in a cost-effective and coherent manner.

Before releasing for formal use, the system and its build-outs should be tested in various ways and developed using best practices. Chapter 14 addresses these needs. The problems we are dealing with in information retrieval (IR), natural language understanding or processing (NLP), and machine learning (ML) are all statistical classification problems, specifically in binary classification. The most common scoring method to gauge the ‘accuracy’ of these classification problems uses statistical tests based on two metrics: negatives or positives, and true or false. We discuss a variety of statistical tests using the four possible results from these two metrics (e.g., false positive). We offer best practices learned from client deployments in areas such as data treatment and dataset management, creating and using knowledge structures, and testing, analysis, and documentation. Modularity in knowledge graphs, or consistent attention to UTF-8 encoding in data structures, or the emphasis on 'semiautomatic' approaches, or the use of literate programming and notebooks to record tests and procedures are just a few of the examples where lines blur between standard and best practices.

In the concluding Part V, the last bookend in our structured organization, we tackle the "conceivable practical effects" that may result from following these pragmatic Peircean approaches. As before, three chapters comprise this part. The first two chapters present what kind of benefits and practical effects can result from following these guidelines to KR. I offer each potential use as a 'mini-story' following the same structure as the book.[7] Chapter 15 speculates on 12 potential applications in *breadth*. Four of these are near-term applications in word sense disambiguation, relation extraction, reciprocal mapping, and extreme knowledge supervision. Four are logic and representation applications in automatic hypothesis generation, encapsulating KBpedia for deep learning, measuring classifier performance, and the thermodynamics of representation itself. The last four areas in Chap. 15 include new applications and uses for knowledge graphs. Two of these, self-service business intelligence and semantic learning, have been on wish lists for years. The last two apply Peirce's ideas and guidance to nature and questions of the natural world. These examples show the benefits of organizing our knowledge structures using Peirce's universal categories and typologies. Further, with its graph structures and inherent connectedness, we also have some exciting graph learning methods that we can apply to KBpedia and its knowledge bases.

Chapter 16 discusses three potential uses in *depth*. The three application areas are workflows and business process management (BPM), semantic parsing, and robotics. Workflows are a visible gap in most knowledge management. One reason for the gap is that workflows and business processes intimately involve people. Shared communication is at the heart of workflow management, a reason why semantic technologies are essential to the task. In semantic parsing, a lexical theory needs to handle word senses, sentences and semantics, cross-language meanings, common-sense reasoning, and learning algorithms. We can map the compositional and semantic aspects of our language to the categorial perspectives of Peirce's logic and semiosis, and then convert those formalisms to distributions over broad examples provided by KBpedia's knowledge. Cognitive robots embrace the ideas of learning and planning and interacting with a dynamic world. Kinesthetic robots may also be helpful to our attempts to refine natural language understanding.

In our last Chap. 17, we are now able to draw some conclusions looking across the broad sweep of our completed *practionary*. Peirce posited a "third grade of clearness of apprehension" to better understand a topic at hand, a part of his pragmatic maxim. As was first stated, knowledge representation is a field of artificial

[7] Namely, that structure is parts organized as context and practical outcomes that are the bookends surrounding the logic triad of grammar (1ns), modes of logic (2ns), and methods (3ns).

intelligence dedicated to representing information about the world in a form that a computer system can utilize to solve complex tasks. Peirce (at least how I interpret him) offers a fresh and realistic take on the question of KR. The foundation of the universal categories and other Peircean ideas offer unique and valuable insights into semantic technologies, knowledge representation, and information science. We need to better understand the nature of signs and representation in the use of semantic technologies. More minds and more scrutiny will improve our understanding and will increase the knowledge we may derive from Peirce's ideas.

I provide supplementary material in three appendices. Appendix A is a short bio of Charles S. Peirce, a most accomplished and fascinating person. Most Peircean scholars acknowledge changes in Peirce's views over time, from his early writings in the 1860s to those at the turn of the century and up until his death in 1914. In Peirce's cosmogony, the primitives of chance (Firstness), law (Secondness), and habit (Thirdness) can explain everything from the emergence of time and space to the emergence of matter, life, and then cognition. Synechism, which Peirce equated with continuity,[8] is the notion that space, time, and law are continuous and form an essential Thirdness of reality in contrast to existing things and possibilities. Peirce made a profound contribution to mathematical logic, where he pioneered many new areas. We can also point to a second area in probability theory, then known as the Doctrine of Chances. Peirce's universal categories of Firstness, Secondness, and Thirdness provide the mindset for how to think about and organize knowledge. The appendix concludes with an annotated list of resources for learning more about Peirce.

Appendix B provides overview information on the KBpedia knowledge artifact. KBpedia is structured to enable useful splits across a myriad of dimensions from entities to relations to types that can all be selected to create positive and negative training sets, across multiple perspectives. The disjointedness of the SuperTypes that organize the 55,000 entity types in KBpedia provides a robust selection and testing mechanism. We organize KBpedia using a knowledge graph, KKO, the KBpedia Knowledge Ontology, with an upper structure based on Peircean logic. KKO sets the umbrella structure for how we relate KBpedia's six constituent knowledge bases to the system. We split the KBpedia knowledge graph into concepts and topics, entities, events, attributes, annotations, and relations and their associated natural classifications or types.

Appendix C discusses the KBpedia features suitable for use by machine learners. This systematic view, coupled with the large-scale knowledge bases such as Wikipedia and Wikidata in KBpedia, provides a basis for faster and cheaper learners across a comprehensive range of NLP tasks. For natural language, a feature may be a surface form, like terms or syntax or structure (such as hierarchy or connections);

---

[8] Peirce states, "I have proposed to make *synechism* mean the tendency to regard everything as continuous" (1893, CP 7.565). He goes on to say, "I carry the doctrine so far as to maintain that continuity governs the whole domain of experience in every element of it. Accordingly, every proposition, except so far as it relates to an unattainable limit of experience (which I call the Absolute,) is to be taken with an indefinite qualification; for a proposition which has no relation whatever to experience is devoid of all meaning" (CP 7.566).

it may be derived (such as statistical, frequency, weighted, or based on the ML model used); it may be semantic (in terms of meanings or relations); or it may be latent, as either something hidden or abstracted from feature layers below it. I present and organize an inventory of more than 200 feature types applicable to natural language. They include lexical, syntactical, structural, and other items that reflect how we express the content in the surface forms of various human languages.

Throughout the book, I try to stick with more timeless concepts and guidelines, rather than current tools or specific methods. Tools and methods change rapidly, with current ones rather easily identified at implementation time.[9] I also try to limit mathematical notations or overly technical discussions. The abundance of references and endnotes provided at the conclusion of each chapter or appendix offers further entry points into these topics. A glossary of technical and Peircean terms and an index conclude the book.

Key Themes

Some themes recur throughout this book. Sometimes how I discuss these concepts may differ by context. To help reduce confusion, let's tackle some of these concepts early.

The first theme is the concept of Peirce's universal categories of Firstness, Secondness, and Thirdness. I devote Chap. 6 to this concept due to its importance and prominence. Peirce's penchant for threes and his belief in the universal categories perfuse his writings across all eras. Peirce's terminology for these 'threes' differs in the contexts of sign-making (semiosis), logic, thought, phenomenology, evolution, protoplasm, information, and so on. As I have tallied across his writings to date, Peirce employs the idea of the universal categories across more than 60 different contexts (see Table 6.2). OK, then, so what is an absolute universal category?

The answer, I think, is it still depends. As I suggest in Chap. 6, perhaps the base definition comes from *hypostatic abstraction* applied to the ideas of First, Second, and Third. Still, all my suggestion does is to substitute one abstract First for another slightly different abstract Firstness. Labels seem to twist us up into literalness and miss the broader point, the one I often harken to in this book about mindset. If we look to the most grounded primitives from which all things, ideas, and concepts are built, according to Peirce, nothing seems as irreducible as one, two, and three. If we further take the understanding of our signs as built from more primitive signs, which combine into more complicated statements and arguments, we can bring Peirce's conception of the universal categories into clear focus. They are meant to inform a process of investigation, refinement, and community, each new concept and term building upon others that came before it. If we reduce that process to its most reduc-

[9] For example, for nearly a decade I started and maintained a listing of semantic technology tools that eventually grew to more than 1000 tools, called Sweet Tools. I ultimately gave up on trying to maintain the listing because of the rapid creating and abandonment of tools. Only a small percentage of these tools lasted for more than a few years.

tive level, it is pretty hard to get more primitive than Firstness, Secondness, and Thirdness. In other words, we can represent anything that we can describe, perceive, or understand using the universal categories for a given context. Our and Peirce's different ways to describe these categories depend on where we are in the representational hierarchy, which is just another way of saying context.

Given the context of knowledge representation, then, what might be the best way to label these categories of Firstness, Secondness, and Thirdness? Many of the optional expressions shown in Table 6.2 approximate this answer. Since the context of knowledge representation is the real world and what we can know and verify, let's take that perspective.

Figure 1.1 is a working conception for what the base context may be for the knowledge representation domain. The unexpressed possibilities or building blocks that might contribute to a given knowledge category I term potentialities, a Firstness.[10] (One could argue that Peirce preferred the idea of possibilities as a Firstness over potential, and good scholarly bases exist to support that contention. However, in this context, it makes sense to limit our possible building blocks to those likely for the category at hand. I think 'potential' better conveys this restriction that some possibilities are more likely for a given topic category than others.) Potentialities include any unexpressed attribute, such as shape, color, age, location, or any characterization that may apply to something in our current category.

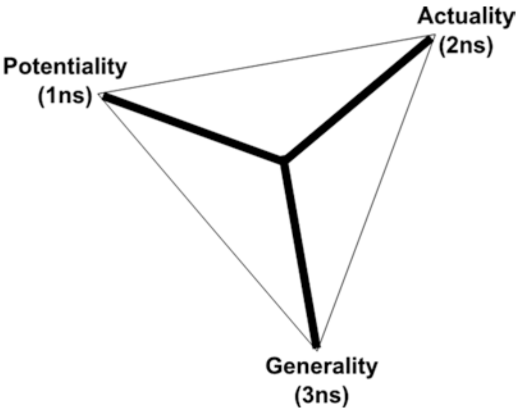

Fig. 1.1 A version of the "universal categories"

Potentialities, when expressed, are done so by the actualities of the world, a Secondness in the universal categories.[11] Actualities are the real, actual things that populate our domain, specifically including entities and events. These actual things may not have a corporal or physical existence (for example, Casper the friendly ghost, with 'fiction' being a legitimate attribute), but they can be pointed to, referred to, or described or characterized. What we find as commonalities or regularities

[10] In the KBpedia Knowledge Ontology, we term the Firstness (1ns) branch as Monads. Also, recall the earlier shorthand of 1ns, 2ns, and 3ns for the three universal categories.

[11] In the KBpedia Knowledge Ontology, we term the Secondness branch as Particulars.

across actual things we can call generalities,[12] a Thirdness in the universal categories. Generalities include types, laws, methods, and concepts that cut across many actuals or generals. Given a different context, the labeling of these universal categories may differ quite substantially, as Table 6.2 affirms. However, virtually any context invoking the universal categories would still retain some sense of these distinctions of potentiality, actuality, and generality.

(Another aspect to note in Fig. 1.1 is its central, heavier lined image, which we can describe as a three-pronged spoke or three-pointed star. Many Peirce scholars prefer this image. It is the form used by Peirce in his writings.[13] We can ascribe the lighter lined equilateral image in Fig. 1.1 to the 'meaning triangle' approach of Ogden and Richards in 1923, also apparently informed by Peirce's writings [4]. Most current Peircean practice favors the equilateral image, which I also tend to use. Though perhaps deep implications reside in the choice of image, I find either image acceptable.[14])

Given the variety of expressions for the universal categories, always ask yourself what the context is for a particular reference. As I state multiple times in the book, the universal categories are a mindset of how to decompose the signs of the world, and plumbing the use and application of the categories in different contexts is one way to better apprehend that mindset.

Another area of 'threes' in this book, but not directly related to the universal categories, is the idea of a *triple*. A triple—so named because it triply combines a *subject* to a *predicate* and to an *object* (*s*-*p*-*o*) that is the basic *statement* or assertion in the RDF and OWL languages that we use in this *practionary*. The triple is equivalent to what Peirce called a proposition. We often represent triples as barbells, with the subject and objects being the bubbles (or nodes), and the connecting predicate being the bar (or edge). Figure 1.2 is such a representation of a basic triple.

Fig. 1.2 Basic "triple"

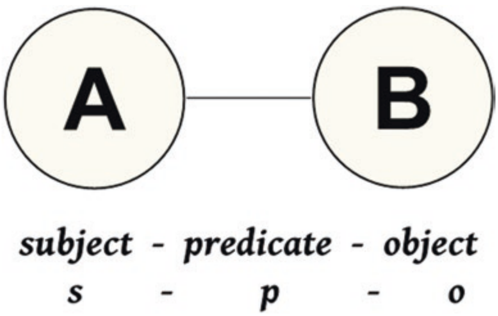

[12] In the KBpedia Knowledge Ontology, we term the Thirdness branch as Generals.

[13] Edwina Taborsky is one vocal advocate for using the "umbrella spoke triad" image as she calls it, noting that it is open and not closed (equilateral triangle), and is the form used by Peirce.

[14] I skewed Fig. 1.1 10 degrees so as to not convey a preference or order to any of the three universal categories.

The triple statements are basic assertions such as 'ball is round' or 'Mary sister of John.' Sometimes an assertion may point to a value, such as "Mary age 8," but it also may be a true object, such as 'John citizen of Sweden.' Objects, then, in one triple statement might be the subject of a different one, such as 'Sweden located Northern Hemisphere.'

Note that I earlier likened the subject and object to *nodes* and predicates to *edges*. This terminology is the language of *graphs*. As one accumulates statements, where subjects of one statement may be an object in another or vice versa, we can see how these barbells grow linked together. When these accrete or accumulate as encountered, we have a bottom-up image of how graphs grow, as illustrated in Fig. 1.3, wherein a single statement grows to become a longer story:

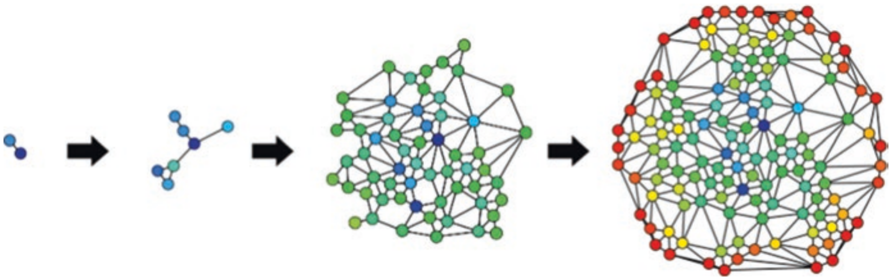

Fig. 1.3 A bottom-up view of graph growth

Of course, we can also create graphs in a top-down manner. An *upper ontology* is one example. We often intend top-down graphs to be a sort of coherent scaffolding of vetted (coherent) relationships upon which we can hang the statements for new instances. Graphs are a constant theme in this book. Chapter 11 is largely devoted to graphs and their uses. The specific kind of graph our knowledge structures assume is a DAG, a *directed acyclic graph*. This fancy term means that the edge relationships in the graph are not all transitive (both directions); one or more exhibit directionality, such as 'George father of Mary.'

Last, let me raise a crucial theme, *fallibility*. Our knowledge of the world is continually changing, and our understanding of what we believe justifies that belief may still be in error—both central tenets of Peirce. I believe that arming ourselves with how to think—and with logical methods to discover, test, select, and relate information—is the right adaptable and sustainable response to a changing world.

References

1. F. van Harmelen, V. Lifschitz, B. Porter (eds.), *The Handbook of Knowledge Representation* (Elsevier, Amsterdam, 2008)
2. R.J. Brachman, H.J. Levesque, *Knowledge Representation and Reasoning* (Morgan Kaufmann, Burlington, 2004)
3. K.A. Parker, *The Continuity of Peirce's Thought* (Vanderbilt University Press, Nashville, 1998)
4. C.K. Ogden, I.A. Richards, *The Meaning of Meaning* (Harcourt, Brace, and World, New York, 1923)

Chapter 2
Information, Knowledge, Representation

Practitioners of knowledge representation (KR) should have a shared working understanding of what the concepts of information, knowledge, and knowledge representation mean. That is the main thrust of this chapter.[1] As a symbolic species [1], we first used symbols as a way to convey the ideas of things. Simple markings, drawings, and ideograms grew into more complicated structures such as alphabets and languages. The languages came to embrace still further structure via sentences, documents, and ways to organize and categorize multiple documents, including ordered alphabets and categorization systems.

Grammar is the rules or structures that govern language. It is composed of syntax, including punctuation, traditionally understood as the sentence structure of languages, and morphology, which is the structural understanding of a language's linguistic units, such as words, affixes, parts of speech, intonation, or context. The field of linguistic typology studies and classifies languages according to their structural features. However, grammar is hardly the limit to language structure. In the past, semantics, the meaning of language, was held separate from grammar or structure. Via the advent of the thesaurus, and then linguistic databases such as WordNet and more recently concept graphs or *knowledge graphs* that relate words and terms into connected understandings, we have now come to understand that semantics also has structure. It is the marriage of the computer with language that is illuminating these understandings, enabling us to capture, characterize, codify, share, and analyze. From its roots in symbols, we are now able to extract and understand those very same symbols to derive information and knowledge from our daily discourse. We are doing this by gleaning the structure of language, which in turn enables us to relate it to all other forms of structured information.

[1] Some material in this chapter was drawn from the author's prior articles at the *AI3:::Adaptive Information* blog: "The Open World Assumption: Elephant in the Room" (Dec 2009); "Give Me a Sign: What Do Things Mean on the Semantic Web?" (Jan 2012); "The Trouble with Memes" (Apr 2012); "What is Structure?" (May 2012); "The Irreducible Truth of Threes" (Sep 2016); "The Importance of Being Peirce" (Sep 2016); "Being Informed by Peirce" (Feb 2017).

M. K. Bergman, *A Knowledge Representation Practionary*,
https://doi.org/10.1007/978-3-319-98092-8_2

What Is Information?

Many definitions of *information* may be found across the ages, often at variance because of what sense is primary. Some definitions are technical or engineering in nature; others emphasize intention, context, or meaning. Gregory Bateson offered one of the more famous definitions of information, claiming it the "difference that makes a difference" [2]. Claude Shannon, the founder of information theory, emphasized a different aspect of information, defining it as a message or sequence of messages communicated over a channel; he specifically segregated the meaning of information from this engineering aspect [3]. For Charles S. Peirce, information is equivalent to *meaning*, which is measurable as the breadth times the depth of the object. Despite this difference, I see both Shannon and Peirce talking broadly about the same underlying thing, though from different aspects of the universal categories. Shannon is addressing a Firstness of information, Peirce a Thirdness, as I will explain.[2]

Some Basics of Information

The idea of information has an ethereal quality. It is something conveyed that reflects a 'difference,' to use Bateson's phrase, from some state that preceded it. Indeed, Norbert Wiener, of cybernetics fame, stated in 1961 that "Information is information, not matter nor energy" [4]. By coincidence, that was also the same year that Rolf Landauer of IBM posited the physical law that all computing machines have irreversible logic, which implies physical irreversibility that generates heat. This principle sets theoretical limits to the number of computations per joule of energy dissipated. By 1991 Landauer was explicit that information was physical [5]. Physicists confirmed that data erasure is a dissipative heat process in 2012 [6]. The emerging consensus is that information processing does indeed generate heat [7]. By these measures, information looks to have a physical aspect.

The motivation of Shannon's 1948 paper on information theory was to understand information losses in communication systems or networks [3]. Much of the impetus for this came about because of issues in wartime communications and early ciphers and cryptography and the emerging advent of digital computers. The insights from Shannon's paper also relate closely to the issues of data patterns and data compression. In a strict sense, Shannon's paper was about the amount of information that could be theoretically and *predictably* communicated between a sender and a

[2] One might try to avoid the term "information" because it has multiple meanings or differing interpretations, substituting instead narrower definitions for each meaning. Moreover, multiple meanings do not have crisp boundaries, so that there is always a probability of misunderstanding or misassignment. However, when a term is commonly used, we need to accept the reality of its use. Thus, Peirce tried mightily, a practice to which I adhere, to define terms as understood and put forward. This consideration applies to any term or definition.