

Building a Pragmatic Methodology for KR Tool Research and Development

Mary A. Keeler¹ and Heather D. Pfeiffer²

¹ Center for Advanced Research Technology in the Arts and Humanities
University of Washington, Seattle, Washington 98117, USA
mkeeler@u.washington.edu

² Department of Computer Science
New Mexico State University, Las Cruces, New Mexico 88003-8001, USA
hdp@cs.nmsu.edu

Abstract. Reviewing the evidence of CG-tool research and development in conference papers, we find little attention devoted to the issues and institution of scientific methodology, and only vague awareness of how that deficiency impairs progress. To focus attention and improve awareness, we briefly delineate the evolution of C.S. Peirce's theory of inquiry toward Eduard Hovy's general methodology for research, tracing from Peirce's early pragmatism to his "conditional idealism." We claim that methodological theory suggests a pragmatic method for KR research and tool advancement, in the form of an open-ended game somewhat like a child's game of building blocks, in which the forms of the "blocks" would be propositional rather than physical, with conditional propositions establishing the "dimensions," in place of the physical dimensions of blocks. The constraints would be logical and evidential (factual) rather than geometrical and gravitational (forceful). We challenge the entire Conceptual Structures community to help build a truly pragmatic methodology.

1 Introduction

Over the 14 years of ICCS meetings, some participants have wondered why work done in the Conceptual Graphs (CG) community exhibits so little scientific conduct, and what Peirce himself might recommend to improve its "pragmatic progress." Eduard Hovy informally expressed his disappointment after hearing CG papers presented, when he was invited to speak on methodology at the 2005 meeting [see 1]. He complained that the CG researchers presented insular work, making no comparisons with other techniques in similar applications outside the CG community or evaluation of tools within. His paper critiques the progress of several research communities, including the CG, and urges that their poor record of ontology-building success could most readily improve if at least two conditions are met: good methodologies for building and evaluating ontologies are developed, and those ontologies prove their utility in real applications [see 1: 91]. He identifies the problem for knowledge representation (KR), in general, after assessing its accomplishments: "KR work has been excellent in developing formalisms for representation, and for investigating the properties and requirements of various classes of deductive systems. But for practical applications such formalisms need *content*; the deductive systems need to work *on something*"

[1: 92]. His summary observation is that KR researchers are not yet able to build large, general-purpose semantic theories or semantic resources required for practical use on the scale of NLP: "such semantic data and theories as do exist are almost always limited to small-scale (or toy) applications ... [and] no accepted standard set of relations exists either" [1: 92].

Hovy stresses that the most troublesome aspect of work underway (such as in CGs) is "the near-complete absence of methodological discussion and emerging 'methodological theory' that would provide to the general enterprise of ontology building and relation creation the necessary rigor, systematicity, and eventually, methods for verification that would turn this work from an art to a science. ... Without at least some ideas about how to validate semantic resources, both the semantics builder and the eventual semantics user are in trouble" [1: 92-93].

In this paper, we review attempts to consider methodology within the CG community. We then investigate (with the aid of a host of Peirce scholars who have blazed productive paths through the maze of Peirce's writings) what methodological insights and guidance can be derived from Peirce's theory of inquiry. We conclude that both pragmatic methods and methodology might be instituted in the creation of an open-ended game.

2 Looking for Methodology in the CG Community

In the history of the CG community, the Peirce Workbench (also called "Peirce project") was an early attempt (1992) to establish some sort of context for at least comparing CG tools in some application domains. It was launched at the 7th workshop on Conceptual Graphs as an attempt to build cooperation among some 40 researchers, by coordinating their 27 tools into one test environment of 11 types [see 2]. Each participant group used different formatting and storing operations for graphs, and there was no well-defined Application Program Interface (API) for communication among all tools, which were built on different platforms in different programming languages. Nevertheless, the project marked the beginning of work on the Conceptual Graphs Interchange Format (CGIF), with the objective of making it possible to translate tools into exchangeable languages such as Java. Although the project itself did not continue, it demonstrated some need for developing collaborative methods among those in the community, but no apparent recognition of the need for methodological discussion of the sort Hovy describes.

The first obvious attempt came at the 1998 conference, when Tepfenhart outlined some of the fundamental technical ideas that form the basis of research efforts across the conceptual structures community. He concluded, "The variety of approaches, processing styles, and assumptions make it difficult for one author to apply the results of another. The same problem, framed in different language, is being solved many times by researchers who do not realize that it is the same problem" [see 3: 345].

2.1 Promising Efforts in 1999

Then came the Sisyphus I (SCG-1) initiative in 1999, devised by researchers at the Knowledge Science Institute, which challenged CG tool developers to solve a

room-allocation task [see LNAI 1640, pp. 272-355]. Many hoped this testbed would "pour life into the CG community" by providing the opportunity: to compare various CG-tool approaches and determine the "CG state of the art," and to distinguish CG from other knowledge representation formalisms according to pertinent criteria [4: 213]. And yet, only the group testing the CoGITaNT tool even mentioned using an "experimental methodology": beginning with a theoretically formal model, then building software tools to implement the model, followed by a real-world application (presumably a prototype), and finally evaluating the system built—and reiterating this four-step process [see 5: 374]. Unfortunately, these researchers apparently attenuated that methodological plan: "This prototype is not yet a really usable tool. Further developments are needed in order to enable communication with the user in a friendly way, give the user the possibility to intervene at several stages of the solving process, and improve computational complexity of the solving process" [5: 375].

2.2 Assessment in 2000

No further obvious effort was attempted to achieve methodological procedures or standards until 2000, when Chein and Genest surmised that perhaps only a handful of operational CG applications remained. At conferences, they observed: "A program may be presented in a session [called] "applications" ... but it does not become, *ipso facto*, an AI or CGs application. ... External characteristics, those which consider the software from outside, like a black box, are essential for a software which claims to be an application" [6: 128]. In their paper ("CG Applications: Where Are We 7 Years After the First ICCS?"), they also discuss "the essential Internal characteristics" for an application: "[it] must have precise specifications, its conditions of use must be described, it must be reliable and robust, it must be validated, documentation must be available, its evolution must be anticipated, and so on" [6: 128]. And they identify four basic additional components specifically required for AI applications [see 6: 131-32].

In summary, Chein and Genest stressed: "In order to build applications, efficient tools are needed, and, what is rather distressing, we could deliver exactly the same discourse that Bob Levinson and Gerard Ellis had done when they launched the Peirce project in 1992!" They concluded that the CG community has difficulty "analyzing, synthesizing, and accumulating its knowledge," and further that "we have made numerous experiments, most of the time without drawing serious conclusions" [6: 138]. Nowhere in this paper is the question of methodology explicitly raised, but that is the question thoroughly begged throughout.

Guy Mineau's invited talk that year offered twenty-two recommendations for good practice in CG-based systems engineering, conceived in a three-layer architecture, as a "first set of guidelines" toward the common goal of a widely-supported, large-scale CG-based platform [see 7: 154]. But not until 2004, was a whole conference session devoted to frameworks for applications [see section "Conceptual Frameworks for Applications" in LNAI 3127, pp. 242-332], only two of which explicitly addressed the need for KR tool-development methodology [see 8 and 9], although it had been previously urged in 1999 and 2000 [see 10 and 11].

2.3 Recent Appraisalment

Without solid methodological requirements, Hovy explains, "the reality is that most people build ontologies to support their knowledge representation needs in some practical application, and ... are more concerned about the computational effectiveness and correctness of their application than about the formal completeness, correctness, or consistency of the ontology *per se*." While theoretically, completeness, consistency, etc., would ensure that ontologies will avoid unwelcome surprises in system behavior, he observes that in development the strictures these requirements introduce are usually "so onerous that they make adopting the requirements tantamount to placing a severe limitation on the eventual scope of the ontology and hence of the whole practical enterprise" [1: 93]. Consequently, he explains, ontologies are built as relatively simple-term taxonomies with some inheritance inference, and do not enforce stricter logical requirements.

Hovy concludes that KR (among other techniques) lacks "a systematic and theoretically motivated methodology that guides the builder and facilitates consistency and accuracy, at all levels." In fact, he finds no evidence of an adequate theory on which to base such a methodology, and furthermore: "*It is not even clear how one would begin to approach the problem of designing theoretically motivated procedures from a suitably general point of view*" (our emphasis). He finds not one of the current builders of ontologies able "to provide a set of operationalizable tests that could be applied to every concept and relation to inform in which cases his or her choices were wrong" [1: 94].

Indicating how to approach such a methodology, Hovy considers what ontology builders actually *do*, in core operation(s) and in justifying their actions. He stresses that builders perform an *act of creation*, every time they generate a new (candidate) ontology item: "[they] decide *whether* to create a term, and if so, *how to place it* with regard to the other existing terms" [1: 94]. This portion of the act of defining the term begins the process of additional specification and definition. Hovy then traces this decision process as it plays out for the five "personality types" of builders he identifies, including "the philosophers," where he classifies those using the CG technique [1: 95]. He then describes a general methodology for building domain models [see 1: 97-98].

We eventually identify Hovy's approach to a general methodology as implicitly *pragmatic*, by first explicating (or clarifying) the motivation for such a methodology in terms of Peirce's theory of inquiry, then indicating how a particular sort of game can institute Peirce's pragmatic methodology systematically in KR tool development.

3 Pursuing a General Methodology in Peirce's Theory of Inquiry

When Peirce's *Collected Papers* was published in the early twentieth century, scholars were discouraged to find that Peirce never coherently stated his theory of inquiry. Thankfully, their persistent investigations provide substantial advantage in deciphering Peirce's writings, and the hope of addressing the concerns Hovy raises. The studies of early scholars such as Bronstein, Chisholm, Weiss, and Feibleman agree that Peirce considered his pragmatism to be a maxim (or general rule) for

inquiry, and that his theory explained the course of inquiry in three elementary stages: abduction, deduction, and induction¹. Abduction, the creative foundation of inquiry (which is presupposed by all induction), he came to identify as the essence of the pragmatic method [see 13: 166].

The puzzle that drove Peirce's pursuit of a theory of inquiry was the difficulty of explaining *how man could have managed to guess so many fruitful hypotheses in such a short historical time*. He argued that there must be an affinity between our minds and the course of the universe giving us insight into the general elements of nature. But he knew this alone could not enable the progress humans have made; it could not reduce the number of possible hypotheses to manageable numbers — and would not explain the countless useless and perverse ones tried in the history of science. His "better theory," says Weiss, is that our efforts to guess are *self-corrective*, "no matter where we start we arrive at results which we can modify as the circumstances demand, until they have the shape which satisfies the entire community" [13: 173]. Still, Peirce was not convinced that this was adequate to explain how science could progress so quickly.

We all know that in any development, some starting points are better and some results are more significant, some techniques are more effective than others and some means of production more efficient. How can we explain the selection of good enough guesses, throughout this process, that humans could progress so swiftly to civilization, and their science even more swiftly? According to Peirce, any such explanation is and always will be hypothetical (in fact, history is entirely hypothetical [see CP 2.511 fn.]). Any such hypothesis then must explain how, from the countless possibilities in any situation, do we formulate and select good hypotheses to test (or good hunches to try)? All of Peirce's philosophical work comes together in this challenge, which he eventually tried to explain as "the economics of research" [see CP 7.83-90].

Weiss gives us an account of the scope of deficiencies that Peirce suggested any *economic methodology* might address: "All men do not make signalized discoveries because they do not all adopt promising initial positions, are not familiar with new provocative items, do not have the requisite technical skill to express [themselves] to the satisfaction of technical arbiters whatever discoveries they might happen to make, and do not have the patience or time to add inference to inference so as to move forward steadily" [13: 178]. These are all possible errors or inefficiencies that any method at its core must help investigators routinely self-correct — by forming self-corrective *habits*.

Judged on this self-corrective requirement, all of Peirce's well-known, early formulations of a theory fail to be convincing, according to the scholarly literature. Clearly he does not maintain in his later writings the "guess" proposed in his early work, "The Fixation of Belief": that inquiry is a struggle to escape the irritation of doubt by attaining belief, true or false, and that the settlement of opinion is its sole aim [see, for best evidence, CP 6.485 (1908)]. The consequence, he realized, would be that *any* inquiry which settles belief is as good as any other, with no way of showing that some beliefs more likely conform to the facts than others.

¹ Note: Sowa and others have used these terms to discuss *types* of reasoning [12], but Peirce related these *ideal types* in his theory to account for the evolution of human knowledge.

3.1 Belief, Error, and Continuing Hypothetical Inference

Peirce's mature theory of inquiry abandons his early conclusions, on good grounds: If the settlement of opinion, by establishing belief that relieves the irritation of doubt, is the sole object of inquiry, and if belief is of the nature of habit, why should we not attain the desired end by taking anything we may fancy as answer to a question, and constantly reiterating it to ourselves, dwelling on all which conduce to that belief, and learning to turn with contempt and hatred from anything that might disturb it? [see *CP* 5.337]. Bronstein thinks Peirce revised his theory of inquiry because he realized that what we should demand of a method of attaining belief is that it should minimize *error*, not doubt. The urge to remove the irritation of doubt, is a biological fact which may cause us to adopt one or another method of fixing or re-establishing belief, but that fact tells us nothing about the validity of any method. It will not tell us *how any self-correction is determined* in the sustained inquiry that Peirce says is required to achieve true knowledge [see 13: 40-41].

In his later work Peirce proposes a more "fundamental hypothesis" to explain the successful method of science by inquiry that must lead to true knowledge, rather than merely conveniently settled opinion. He explicitly guesses: "there are Real things whose characters are entirely independent of our opinions about them; these Reals affect our senses according to regular laws" [*CP* 5.384]. We initiate inquiry when we confront a puzzling situation, then attempt to resolve the puzzle in the construction of a hypothesis that will enable us to anticipate the course of our experience (or not be surprised). A hypothesis is a *tentative* belief which may become firmer if it improves our anticipation. Theoretical as well as practical beliefs must involve expectation: when certain conditions are fulfilled, we obtain the consequence we expect. The general form of any hypothesis answers the pragmatic question: "What would be the consequences in my conduct (remembering that *thinking* is a form of conduct), if my conception of some observed phenomenon turned out to be true?" In other words, "what would I *do* (or *think*) differently if that conception explained what I observe of that phenomenon?" [see *CP* 5.534, 8.209] Any answer to this question is a hypothesis, whether implicit or explicitly expressed.

3.2 Beliefs Asserted as Conditional Propositions

Any belief can serve as an explicit hypothesis, if it is formulated in a conditional proposition whose antecedent specifies a course of action to be performed and whose consequent describes certain consequences to be expected. The hypothesis can be judged correct when we have perceived a correspondence between the *description* of these consequences and their *occurrence*. This, stresses Bronstein, is not a test of truth: "Rather, it is an attempt to explain what we mean when we say that a statement or belief is *true*" [13: 42]. In his 1903 "Lectures on Pragmatism," Peirce first clearly claims there is an "antecedent reality" and dismisses his earlier view that truth should be defined as the end of inquiry [see *CP* 5.211]. "What an inquiry presupposes," as Bronstein interprets Peirce, "is that there are phenomena subject to *law*, over which our thinking has no control," and which "we perceive directly in the course of inquiry," but which are "not introduced into the operation of knowing" [13: 43; our italics and punctuation].

Bronstein concludes that the principal weakness in Peirce's early theory was his failure to distinguish between something being true and our *knowing* that it is true. What results from inquiry is not that the belief becomes true, but that we gain some additional knowledge that we didn't have before the inquiry: that the belief is *justified* [13: 44]. Peirce clearly identified this confusion between truth and justification in his later work, but other pragmatists have perpetuated it. The confusion reaches to the depths of metaphysics, as Peirce discovered in reading Aristotle's works: "the first thing that strikes [the reader] is the author's [Aristotle's] unconsciousness of any distinction between grammar and metaphysics, between modes of signifying and modes of being" [CP 2.384; see 14]. The problem comes to focus in the question of identification and classification: "But identity ... is not a relation between two things, but between two representamens of the same thing" [CP 4.464].

Peirce then began to distinguish his theory of inquiry from the many contemporary theories known as "pragmatism" by calling it *conditional idealism*, expressed in a conditional proposition that aligned his theory with experimentalism: If a certain experiment were performed, then an experience of a given description would ensue. [CP 5.494]

3.3 What Will Be the Facts?

When we make a statement of fact, according to Peirce, we are asserting a *real possibility*: an imaginable event that would be realized under certain describable conditions, which must be specified in an explicit hypothesis. Bronstein reminds us that, on Peirce's view, "only individuals have actual existence (Firstness), ... [but] the world would be a mere chaos if these individuals were not related to each other and also subject to laws which govern their behavior" [13: 48]. The facts of science, then, are the *conditional* discoveries of those relations and behaviors (or habits). There is no such thing as an isolated fact, and the relations we discover and call facts are all conditionally dependent on how we perceive and conceive them. Any statement of fact (or assertion of something's reality) relies on the *assumed truth* of some *general conditional proposition* [see CP 5.457]. An assertion of scientific fact *means* that under certain conditions something *would be* true, whether the assertion explicitly (using the subjunctive conditional form) includes reference to those conditions or not. (This view underlies Peirce's theory of induction and distinguishes his philosophy, fundamentally, from some claimed followers, such as Popper, whose falsificationism is strictly deductive [see Haack (1995) *Evidence and Inquiry*, p. 131]. Putnam tells us that Popper even rejected the very idea of inductive logic [see Putman (2002) *The Collapse of the Fact/Value Dichotomy*, p. 141]).

Peirce's later theory distinguishes between a *proposition* and the *assertion* of that proposition, and insists that when you assert a proposition, you become responsible for it, as though you had placed a wager on it. "Now, if in the future it shall be ascertained that your assertion is correct, you win the wager. The rational meaning, then, or intellectual purport of your assertion lies in the future" [CP 5.543].

4 Genuine Inquiry and the Growth of Knowledge

Bronstein contends that Peirce's "signal contribution," in his revision of what we will continue to refer to as "pragmatism," was to realize "the importance of subjunctive conditionals for science and their irreducibility to material implications" [see 13: 49; *CP* 4.580]. Along with Bronstein, Chisholm concludes that Peirce's account of inquiry still relies on the account of belief as habit, but adds what might be called "outlook," expectations that anticipate imposed real constraints. A belief-habit is a law connecting behavior and possible experience, and a symbol (such as a statement or a sentence) is meaningful only if it can express a belief. But if it purports to concern matters that "transcend the limits of possible experience," then it cannot mean anything: any *meaningful* hypothesis must be *verifiable* (see expanded explanation, below, in section 4.3) [see 13: 93-4; *CP* 5.536; 5.597]. His theory of inquiry then becomes also a theory of meaning [see 13: 50].

Peirce's remark, "That belief gradually tends to fix itself under the influence of inquiry is, indeed, one of facts with which logic sets out," encouraged Chisholm to attempt a formulation of Peirce's basic theory of inquiry ("independently of the complications of his account of probability and induction") [*CP* 2.693; see also 3.161]. Chisholm compiled ten distinct tenets, which he entitled "The Conduct of Inquiry," from "the bewildering collection of Peirce's statements concerning doubt, surprise, habits of action, judgments of perception, common-sense, indubitability, fallibility, and truth" in the *Collected Papers* [see 13: 93]. We have modified Chisholm's list to serve in tracing the evolution of inquiry through Peirce's three stages (abduction, deduction, and induction) [see 13: 95-99]. We include abbreviated relevant discussion and comments from Peirce after each tenet; but we found it necessary to enrich the tenets we list under "Deduction," since Chisholm almost entirely neglects what Peirce describes must occur in that stage.

4.1 Abduction: Belief, Surprise, and Conjecture

"Abduction merely suggests that something may be" [*CP* 5.171].

1) The inquirer must have some beliefs to begin with; for inquiry does not begin until experience, by shock or surprise, breaks in upon some belief-habit. An empty mind cannot be surprised [see *CP* 5.512].

2) The inquirer should be guided by those personal beliefs which have survived the shock, many of which are indubitable. There is no genuine "complete doubt," as Cartesian theory presumes [see *CP* 5.416, 5.265].

3) As the inquirer ponders the surprising phenomena in relation to beliefs already held, conjectures or hypotheses instinctively suggest themselves. Each conjecture or hypothesis furnishes a possible explanation: "a syllogism exhibiting the surprising fact as necessarily consequent upon the circumstances of its occurrence together with the truth of the credible conjecture, as premisses" [*CP* 6.469]; each comes in a flash of *insight* [see *CP* 5.173].

4) The relative plausibility of these hypotheses is considered in terms of what the inquirer happens to believe already. The hypothesis ought first, to be "entertained interrogatively," for there are many hypotheses "in regard to which knowledge already in our possession may, at once, quite justifiably either raise them to the rank of opinions, or even positive beliefs, or cause their immediate rejection" [CP 6.524]. "Accepting the conclusion that an explanation is needed when facts contrary to what we should expect emerge, it follows that the explanation must be such a proposition as would lead to the prediction of the observed facts, either as necessary consequences or at least as very probable under the circumstances" [CP 7.202].

4.2 Deduction: Conditional Prediction of Possible Consequences

"Deduction proves that something must be" (under what we can imagine would be the ideal conditions) [CP 5.171].

5) Most of the hypotheses which any inquirer is thus led to adopt will be false, but many of them will be true, at least under ideal conditions. Unfortunately, the faculty which originates conjectures supplies us with more false conjectures than with true ones. (The terms "true" and "false" are explained below, under "Induction."). Our insight is "strong enough not to be overwhelmingly more often wrong than right" [CP 5.173]. But in order to prevent future surprises and the ensuing misfortunes and irritable states, we need a way of sifting out the bad guesses. "Deduction ... relates exclusively to an ideal state of things. A hypothesis presents such an ideal state of things, and asserts that it is the icon, or analogue of an experience" [CP 7.205]. Experimentation requires preliminary "logical analysis," which renders the hypothesis as distinct as possible and *deduces* experimental predictions from the conjunction of the hypothesis and our other beliefs [see CP 6.472; 6.527]. The hypothesis has two relations to experience, one to the facts [induction] but the other to the hypothesis, and what effect that hypothesis, if embraced, must have in modifying our expectations in regard to future experience. Peirce clarifies in a footnote, "we infer by Deduction that if the hypothesis be true, any future phenomena of certain descriptions must present such and such characters" [CP 7.115 fn. and see 7.115].

6) Experience can eliminate a false hypothesis by virtually predicting the results of possible experiment. As soon as a hypothesis is adopted, deduction must trace out its necessary and probable experiential consequences [see CP 7.202]. Deduction draws virtual predictions. A virtual prediction is "an experiential consequence deduced from the hypothesis, and selected from among possible consequences independently of whether it is known, or believed, to be true, or not; so that at the time it is selected as a test of the hypothesis, we are either ignorant of whether it will support or refute the hypothesis, or, at least, do not select a test which we should not have selected if we had been so ignorant" [CP 2.96]. "The Deductions which we base upon the hypothesis produce conditional predictions concerning our future experience" [CP 7.115 fn.]. "'Conditional' is the

right appellation, and not 'hypothetical,' if the rules of [my] Ethics of Philosophical Terminology are to be followed" [CP 2.316 fn.]. Experience, if given the opportunity, will perform a sifting or "pruning" function. In fact, this is the main service which experience renders — "to precipitate and filter off the false ideas, eliminating them" [CP 5.50].

4.3 Induction: Conditional Truth, Verifiability, Essential Fallibility

"Induction shows that something actually is operative" [CP 5.171].

7) If experience causes surprise, the new surprise will be accompanied by a series of events similar to those which accompanied the first surprise. "We now institute a course of quasi-experimentation in order to bring these predictions to the test, and thus to form our final estimate of the value of the hypothesis, and this whole proceeding I term Induction. ... By quasi-experimentation I mean the entire operation either of producing or of searching out a state of things to which the conditional predictions deduced from hypothesis shall be applicable and of noting how far the prediction is fulfilled" [CP 7.115 fn.]. If experience surprises, there will be new doubt leading to a new struggle to recover belief with a new hypothesis. If experience confirms, it will not inspire further inquiry, which is particularly unfortunate when the confirmed hypothesis happens to be false. Even when the confirmed hypothesis is true, complacency may result, which can impair the ultimate success of inquiry [see CP 5.168-69]. In so far as they greatly modify our former expectations of experience and in so far as we find them, nevertheless, to be fulfilled, we accord to the hypothesis a due weight in determining all our future conduct and thought. Even if the observed conformity of the facts to the requirements of the hypothesis has been fortuitous, we have only to persist in this same method of research and we shall gradually be brought around to the truth [see CP 7.115].

8) If given sufficient opportunity, experience would eliminate all false beliefs and leave us with none but true beliefs; this follows from Peirce's definitions of "true" and "false." In the long run "if inquiry were sufficiently persisted in" and experience given every opportunity to prune out the unstable beliefs, "the community of inquirers" would reach an agreement and all would share the same perfectly stable beliefs; some beliefs would continued to be re-affirmed and some would be denied [see CP 5.384, 5.311]. And the more we thus persist, particularly if we work together in *community*, the closer we come to this ideal. Peirce defines a *true* belief as a belief which would thus be an "ultimate conclusion" of inquiry. "The truth" is that to which belief would tend if it were to tend indefinitely toward absolute fixity [see CP 5.416]. These definitions become the foundation of his "conditional idealism," in which the concept of *truth* becomes the ideal (or hope) that motivates inquiry to persist indefinitely [see CP 5.494, 5.358 fn.]. It is much more important to frustrate the false hypothesis, than to confirm true ones, since we make more false ones than true ones. Hence, although it is not *solely* by surprise, it is *primarily* by surprise that experience teaches us [CP 5.51].

9) In order for experience to perform this function most efficiently, the inquirer should endeavor to submit all of his hypotheses and belief-habits to constant *experimental* test. We should make efforts to expose our hypotheses to the effects of experience, and reject those methods of fixing belief which attempt to shield our beliefs from those effects. The best effort is to conduct planned experimentation. An explanatory hypothesis must be *verifiable*; which is to say that it is "little more than a ligament of numberless possible predictions concerning future experience, so that if they fail, it fails" [CP 5.597]. We then submit these predictions to active test, by means of these actions: first, the act of choice by which the experimenter singles out certain identifiable objects to be operated upon; next, the external (or quasi-external) *act* by which the experimenter modifies those objects; and next, the subsequent *reaction* of the world upon the experimenter in a perception; and finally, the experimenter's recognition of the teaching of the experiment. "While the two chief parts of the event itself are the action and the reaction, yet the unity of essence of the experiment lies in its purpose and plan" [CP 5.424].

4.4 Self-correction as the Inquirer's Habit of Mind

10) This endeavor requires, in turn, that all inquirers have a "will to learn" and a constant dissatisfaction with their state of opinion at any time. Finally, "inquiry of every type, fully carried out, has the vital power of self-correction and of growth," but only if the inquirer has an intense desire to learn: a "dissatisfaction with one's present state of opinion," and a "sense that we do not know something" [CP 5.584; see 5.582; 5.583; also see 6.428]. Here can be the danger of the complacency when we get confirmation rather than surprise. An "indispensable ingredient" in any experiment is "sincere doubt in the experimenter's mind" concerning the truth of any hypothesis [see CP 5.424]. We should desire to know and be willing to work to find out [see CP 5.584]. While opinions that naturally suggest themselves are more often wrong than right, they can be corrected provided that we have a genuine dissatisfaction with them. Consequently we should not regard any as finally settled or certain, which is the basis of Fallibilism. "The first condition of learning is to know that we are ignorant. A man begins to inquire and to reason with himself as soon as he really questions anything and when he is convinced he reasons no more" [CP 7.322].

5 The Essential Open Mind

Peirce's persistent questioning took him well beyond his original pragmatism, as we have seen. And yet, his theory maintains the initial distinction that he says logic supposes in any investigation, between doubt and belief, which later became the distinction between a *question* and a *proposition* [see CP 7.313]. He explicitly relates these distinct elements in various expressions as his theory advances, beginning in his 1878 essays: "The object of reasoning is to find out, from the consideration of what we already know, something else which we do not know" [CP 5.365]. Meanwhile, his progress in developing relative logic gave him increasingly more insight into the

strategic fine points of this relationship in the process of inquiry and, as Weiss concludes in one example, the critical balances to be maintained.

The most satisfactory inference [according to Peirce's theory] is one which enables us with the least risk to make the furthest advance beyond our present position. It must evidently move as far from the initial position as possible, but only to what is relevant. If we violate the first condition, we unnecessarily limit the range of application of the result; if we violate the second, we risk losing control over whatever material or truth we originally had. The combination of these two considerations enables us to preserve the past in the future, to extend the range, enrich the meaning while preserving the achievements of our inherited position. [13: 179]

By 1905, after he changed "pragmatism" to "pragmaticism," Peirce's maxim explicitly incorporated his *realist* insight, to balance the earlier *idealist* insight that true ideas are revealed by an "*a priori* evolutionism of the mind" (or because mind has evolved to be attuned to nature). According to pragmaticism, for the ultimate meaning of concepts to be true, they must represent their objects in the form of conditional resolutions consisting of conditional propositions and their antecedents, which he concludes amounts to saying "that possibility is sometimes of a real kind" [CP 5.453]. This realist *outlook* was the basis of his prolonged harangue against *nominalism*. From his early studies in the history of philosophy, he understood that nominalism had historical alliance with idealism and, by 1893, he had formulated his realist criticism of the nominalists' subtle problem: "they merely restate the fact to be explained under another aspect; or, if they add anything to it, add only something from which no definite consequences can be deduced. A scientific explanation ought to consist in the assertion of some positive matter of fact, other than the fact to be explained, but from which this fact necessarily follows; and if the explanation be hypothetical, the proof of it lies in the experiential verification of predictions deduced from it as necessary consequences" [see CP 8.30; 6.273]. Accepting a nominalist explanation as *sufficient* "would be to block the road of inquiry" [CP 6.273].

By 1906, he realized: "According to the nominalistic view, the only value which an idea has is to represent the facts, and therefore the only respect in which a system of ideas [or a generalization] has more value than the sum of the values of the ideas of which it is composed is that it is compendious" [CP 4.1]. But then, while insisting that hypothetical generalizations should be "submitted to the minutest criticism before being employed as premisses," he declares, "It appears therefore that in scientific method the nominalists are entirely right. Everybody ought to be a nominalist at first, and to continue in that opinion until he is driven out of it by the *force majeure* of irreconcilable facts. Still he ought to be all the time on the lookout for these facts, considering how many other powerful minds have found themselves compelled to come over to realism" [CP 4.1]. From all this, we conclude that though a nominalist explanation is not *sufficient*, it could be at least *necessary*?

5.1 Nominalism as a Game of Inquiry

Peirce formulated a question to distinguish realists and nominalists by their opposite answers: "Do names of natural classes (such as man and horse) correspond with

anything which all men, or all horses, really have in common, independent of our thought, or are these classes constituted simply by a likeness in the way in which our minds are affected by individual objects which have in themselves no resemblance or relationship whatsoever" [CP 8.13]. Pursuing inquiry with the latter, nominalist, form of investigation would result in reducing knowledge to a sort of "map" for the "unknowable territory" of existential being, and then accepting *that* representation as all we *can* know. Inquiry then becomes a sort of "game," in which we "take the map to be the territory." The advantage of "map-making" is that we can construct coherent accounts or stories to satisfy some conceived purposes, and submit those representations to "the minutest criticism" with respect to those purposes before admitting them as serious hypotheses to be tested [see 15].

We propose that KR researchers, instead of being unconsciously nominalistic in their conduct [see 11], should adopt nominalism consciously, as a formal method of inquiry by which to continue building and submitting to minutest criticism "their maps," (or what Chein and Genest have referred to as "essential internal characteristics" of software), while remaining "on the outlook for the facts," just as Peirce recommends (which are external to what Chein and Genest call that software "black box"). The strategic refinement made possible by Peirce's logic, especially his Existential Graphs (EG), suggests to us what sort of game could be constructed for instituting that method [see details in 15]. As Peirce says, "relative logic shows that from any proposition whatever, without a second, an endless series of necessary consequences can be deduced; and it very frequently happens that a number of distinct lines of inference may be taken, none leading into another" [CP 3.641]. The object of the game would be to "prune, filter, and select" the worthy hypotheses to test (or the essential internal characteristics to validate), among all the possible ones players might "wager" in the game.

A game method would be particularly appropriate for creating test-worthy hypotheses collaboratively. First, in explicitly (but sportively) formalizing the process of inquiry; second, by encouraging participants to make strategic contributions responsive to the progress of collaboratively formulating verifiable hypotheses that can be reconciled into testable predictions. The game must be "open-ended," as Peirce's theory mandates, and as simple to play as possible, leaving the challenge of building the "fail-proof ligaments in the construction" to KR processing. The analogy of a children's game of building blocks, comes to mind. In the game of inquiry, the forms of the "building blocks" would be propositional rather than physical, with conditional propositions establishing the "dimensions," in place of the physical dimensions of blocks. Rather than geometrical and gravitational (forceful) constraints, they would be logical and evidential (factual). These, conditionally-related, building blocks would "behave" as complex systems adapting to an "environment," in which fallibility would serve as "gravity" does in physical systems [see 15].

Peirce's fallibilism (identified with all three stages of inquiry, unlike Popper's falsificationism) reminds us that nothing is known for certain, that we should conduct inquiries so that sooner or later experience will catch up with any unstable (that is, invalid or unreliable) belief-habits and eliminate them. And since possible predictions are "numberless," fallibilism entails that no hypothesis can be completely verifiable" [see CP 1.149, 1.404, 2.663]. If inquiry's purpose is to find belief-habits that are reliable enough to serve as stable strategies in its evolution, then within the

game context for investigation, complexity theory could apply. The building blocks for evolving the stable strategies of complex adaptive systems, composed of interacting agents, are described in terms of rules. Holland explains: "a major part of the environment of any given adaptive agent consists of other adaptive agents, so that a portion of any agent's efforts at adaptation is spent adapting to other adaptive agents" [16: 10]. Agents adapt their behavior by changing their rules as experience accumulates, just as improving hypotheses must do. We are exploring this framework for the game, in future work.

If, as Peirce predicts, "belief gradually tends to fix itself under the influence of inquiry," then a game strategy in which players are "always on the outlook for the facts" should prove successful [CP 2.693]. The nominalist game of inquiry must therefore be played in context where the facts can be found.

5.2 Realism in the Context for Inquiry

Returning to Hovy's paper on methodology, we find that his analysis and examples of current methodological deficiencies pinpoint the need for both minute criticism of ideas and for expanded scope in responding to external realities. He describes an *ideal* methodological objective for ontology creation that would reconcile the identified concept-creation procedures, and even assign relative priorities to their various methods and justification criteria *a priori*. But Hovy would agree with Peirce's caution that although there are universal presuppositions operative in the thinking of scientific investigators which may be essentially valid, "the history of science illustrates many of the stumbling blocks created by *a priori* dicta" [CP 1.114]. We have argued that something like the "nominalist game" should be instituted for such *idealistic* abduction and deduction, to accomplish collaborative reconciliation, but also *as a game* to remind us that the reconciled hypotheses and their predictions must not "block the road of inquiry."

Hovy then describes a more *realistic* approach. His general methodology: *continual graduated refinement* [CGR] for building domain models proceeds in seven steps [1: 97-98]. CGR builders would begin by selecting "the principal criteria" for concept-creation and justification methods, and specify their order — which he insists *must be determined by the specific task or enterprise to be served*. He suggests basic questions to be answered in this preliminary step, with refinements to follow in a two-step process. These include: "Is this domain model to be used in a computational system? Or is it a conceptual product for domain analysis and description? Who are the intended users/readers of the ontology? What is their purpose for it? What justification criteria will they find most understandable? Do these criteria match the purpose or task in mind?" [1: 98]

These could be formulated as *pragmatic* questions, but typically in computer system development they are answered in the manner of traditional realism, rather than investigated in the manner that Peirce's critical realism demands. The traditional realist would be concerned with logical self-consistency and a "cognitive reality" of answers to these questions, rather than experientially by scientific induction. The result would be a *logical construction* of what the builder's finite intellect can apprehend as answers, rather than an outcome of the process of scientific inquiry.

Peirce identified this as the "nominalistic tendency" that distinguished traditional realism from his thorough-going realism [CP 1.29 (1869)].

We have previously described a *testbed* context [see 15, 17, 18] in which the answers to such questions could be experimentally obtained, where tool versions would be proposed as hypothetical solutions and these *predictions* tested on real tasks to solve. As Peirce says: "the entire meaning of a hypothesis lies in its conditional experiential predictions," to the extent that its predictions are true, the hypothesis is true [CP 1.29]. Why not build a context in which tools can be evaluated according to his scientific methodology? The very task of representing and managing the knowledge in such a context would be a healthy challenge for KR tool development. We think it is the appropriate challenge for any technology based in Peirce's logical theory, especially in his Existential Graphs (EG).

Peirce's EG have a central function in his philosophy, as instruments for observing experiments in reasoning. Greaves explains that they were designed to fill that role: "they would make both the logical structure and the entailments of propositions directly observable, in the same way that ... molecular diagrams make both the atomic structure and possible combinations of organic compounds observable" [19: 172]. Computer technology has given us Geographical Information Systems (GIS), we now need comparable technology for visualizing Peirce's pragmatic conduct of inquiry, and the CG community needs to pursue such an *outlook*.

6 Conclusions

At the beginning of Peirce's work on pragmatism, he briefly outlined what he conceived as the evolution of ways to establish beliefs, which he later advanced in his famous essays of 1877-8.

Men's opinions will act upon one another and the method of obstinacy will infallibly be succeeded by the method of persecution and this will yield in time to the method of public opinion and this produces no stable result. Investigation differs entirely from these methods in that the nature of the final conclusion to which it leads is in every case destined from the beginning, without reference to the initial state of belief. ... But this will not be true for any process which anybody may choose to call investigation, but only for investigation which is made in accordance with appropriate rules. Here, therefore, we find there is a distinction between good and bad investigation. This distinction is the subject of study in logic. [CP 7.318-20 (1873)]

In the end, his theory of inquiry retains the objective of explaining a method for finding stable belief, but not at the cost of settling for what we call a "nominalist game." And yet, he came to realize that what is most wonderful about the mind is the ability to create ideas for which there is yet no existing prototype and "by means of this utter fiction it should manage to predict the results of future experiments and by means of that power should during the nineteenth century have transformed the face of the globe?" [CP 7.686] Clearly, the open-ended creation of ideas is also essential.

In Feibleman's evaluation of the lessons to be learned from Peirce's philosophy for those who might carry on his work (but not be blinded by it), the final lesson is something that Peirce was first to see: "the possibility of constructing an *open*

system." All systems conceived by philosophers before Peirce had been *ipso facto* closed, he explains: "there is no real reason why there must be a limit to the size of our hypotheses. ... to maintain a single proposition tentatively should be no easier than to maintain a consistent set" [13: 325, 334].

We think that tools created with technology that is theoretically based in Peirce's work should be developed as open-ended experiments, relying on Peirce's minutely-critical logical instruments and his realistic outlook, for a suitably pragmatic (self-corrective) methodology that can cope with the challenges of their necessarily collaborative development.

References

General Notes: For all *CP* references, *Collected Papers of Charles Sanders Peirce*, 8 vols., ed. Arthur W. Burks, Charles Hartshorne, and Paul Weiss (Cambridge: Harvard University Press, 1931-58).

1. Hovy, E. [2005]. "Methodology for the Reliable Construction of Ontological Knowledge." In F. Dau, M-L. Mugnier, and G. Stumme (Eds.): *Lecture Notes in Artificial Intelligence*, Vol 3596. Berlin: Springer, pp. 91-106.
2. Ellis, G and Levinson, R. [1992]. "The Birth of PEIRCE: A Conceptual Graph Workbench." In H.D. Pfeiffer and T.E. Nagle (Eds), *Proceedings of the 7th Workshop on Conceptual Graphs*, LNAI 754. Berlin: Springer, pp. 219-228.
3. Tepfenhart, W. M. [1998]. "Ontologies and Conceptual Structures." In M-L. Mugnier, and M. Chein (Eds.): *Lecture Notes in Artificial Intelligence*, Vol 1453. Berlin: Springer, pp. 334-48.
4. Damianova, S. and Tpitampva, K. [1999]. "Using Conceptual Graphs to Solve a Resource Allocation Task." In W. Tepfenhar and W. Cyre, (Eds.): *Lecture Notes in Artificial Intelligence*, Vol 1640. Berlin: Springer, pp. 297-314.
5. Baget, J.F., Genest, D., and Mugnier, M.L. [1999]. "A Pure Graph-Based Solution to the SCG-1 Initiative." In W. Tepfenhar and W. Cyre, (Eds.): *Lecture Notes in Artificial Intelligence*, Vol 1640. Berlin: Springer, pp. 355-76.
6. Chein, M. and Genest, D. [2000]. "CGs Applications: Where Are We 7 Years After the First ICCS? " In B. Ganter, and G. Mineau (Eds.): *Lecture Notes in Artificial Intelligence*, Vol.1876. Berlin: Springer, pp. 127-39.
7. Mineau, G. [2000]. "The Engineering of a CG-Based System: Fundamental Issues." In B. Ganter, and G. Mineau (Eds.): *Lecture Notes in Artificial Intelligence*, Vol.1876. Berlin: Springer, pp. 140-56.
8. Keeler, M. [2004]. "Using Brandom's Framework to Do Peirce's Normative Science." In K.E. Wolff, H.D. Pfeiffer, and H.S. Delugach (Eds.): *Lecture Notes in Artificial Intelligence*, Vol. 3127, Springer-Verlag, pp. 37-53.
9. de Moor, A [2004]. "Improving the Testbed Development Process in Collaboratories." In K.E. Wolff, H.D. Pfeiffer, and H.S. Delugach (Eds.): *Lecture Notes in Artificial Intelligence*, Vol. 3127, Springer-Verlag, pp. 261-73.
10. Kayser, D. [1998]. "Ontologically Yours." In M-L. Mugnier, and M. Chein (Eds.): *Lecture Notes in Artificial Intelligence*, Vol 1453. Berlin: Springer, pp. 35-47.
11. Keeler, M. [2000]. "Pragmatically Yours." In B. Ganter, and G. Mineau (Eds.): *Lecture Notes in Artificial Intelligence*, Vol.1876. Berlin: Springer, pp. 82-99.
12. Sowa, J. [1990]. "Crystallizing Theories Out of Knowledge Soup." In Z. W. Ras and M. Zemankova (Eds.): *Intelligent Systems: State of the Art and Future Directions*. Ellis Horwood Ltd.: London, pp. 456-487. <http://www.jfsowa.com/pubs/challenge>

13. Wiener, P. P. and Young, F. H. [1952]. *Studies in the Philosophy of Charles S. Peirce*. Cambridge, MA: Harvard University Press.
14. Keeler, Mary [2003]. "Hegel in a Strange Costume: Reconsidering Normative Science in Conceptual Structures Research." In A. de Moor, W. Lex, and B. Ganter (Eds.): *Lecture Notes in Artificial Intelligence*, Vol. 2746. Berlin: Springer, pp. 37-53.
15. Keeler, M. and Pfeiffer, H.D. [2005]. "Games of Inquiry for Collaborative Concept Structuring." In F. Dau, M-L. Mugnier, and G. Stumme (Eds.): *Lecture Notes in Artificial Intelligence*, Vol 3596. Berlin: Springer, pp. 396-410.
16. Holland, J. H. [1995]. *Hidden Order: How Adaptation Builds Complexity*. Basic Books.
17. de Moor, A., Keeler, M. and Richmond, G. [2002]. "Towards a Pragmatic Web." In U. Priss, D. Corbett, and G. Angelova (Eds.): *Lecture Notes in Artificial Intelligence*, Vol. 2393, Springer-Verlag, pp. 235-249.
18. Pfeiffer, H.D. [2004]. "An Exportable CGIF Module from the CP Environment: A Pragmatic Approach." In K.E. Wolff, H.D. Pfeiffer, and H.S. Delugach (Eds.): *Lecture Notes in Artificial Intelligence*, Vol. 3127, Springer-Verlag, pp. 319-332.
19. Greaves, Mark [2002]. *The Philosophical Status of Diagrams*. Stanford University: CSLI Publications.