Associative and Formal Structures in AI

Uta Priss
Indiana University Bloomington, Email: upriss@indiana.edu

Abstract

This paper identifies a dichotomy between associative and formal structures as a common phenomenon that occurs in many disciplines. With respect to AI, arguments are provided for combining both associative and formal structures in multi-representational systems. This paper lists distinguishing features of associative and formal structures and provides some insights into how these can be formally described. The last section mentions strategies for combinations.

1 Introduction

This paper was inspired by Blank's (2001) presentation at last year's MAICS conference and the discussions it stimulated among the conference participants. Blank named his presentation "Radical Artificial Intelligence: A Postmodern Approach", which was chosen as a variation on Russell & Norvig's (1995) title "Artificial Intelligence: A Modern Approach". In his presentation, Blank distinguished two types of approaches in AI: subsymbolic ones and symbolic ones. Subsymbolic approaches employ biologically-inspired metaphors, such as "neural networks", "evolutionary computing" and "genetic algorithms". They are usually decentralized (eg. agent-based), dynamic, emergent, bottom-up and their internal representations are usually incomprehensible to humans. On the other hand, symbolic approaches employ symbolic logic. They are usually explicitly designed by humans - often in a top-down manner, as opposed to being emergent. Their internal representations are also explicitly designed and thus comprehensible to humans. Symbolic structures are evaluated with respect to correctness.

Blank's conclusion was that a "postmodern" AI that relies on subsymbolic instead of symbolic approaches appears to be more promising. This may, however, be somewhat unsettling to humans because subsymbolic approaches do not generate explanations. These approaches can implement intelligent behavior but it is opaque to humans how the internals of the system work. A "general unified theory" of the principles of AI is according to Blank thus utopia. During the subsequent discussion, I suggested that instead of favoring one approach over the other, why not combine both. There already exist a few systems that attempt to combine both approaches in so-called "hybrid systems". But maybe because the precise mechanisms of how both approaches can work together are not yet very well understood, hybrid systems are not yet very popular.

This paper presents a preliminary analysis of the issues involved in combining the two approaches. The terms "subsymbolic" and "symbolic" can be misunderstood because on some level all representations in a computer are symbolic. Therefore, in this paper the terms "associative" and "formal" are used instead of "subsymbolic" and "symbolic". Section 2 of this paper provides further examples of the divide between associative and formal structures, which occurs in many disciplines in many disguises. Section 3 lists arguments for why a combination of both approaches might be useful. Section 4 elaborates further on the difference between associative and formal structures. Blank's list of differentiating features is amended by Sloman's (1996) list and some additional features. Section 5 defines concepts. Section 6, finally proposes some strategies of how to actually combine the two approaches.

2 The divide between formal and associative structures in different disciplines

The divide between associative and formal structures occurs in many other disciplines apart from AI. In psychology, there is a divide between associative and rule-based (or formal) explanations of cognitive abilities (Sloman, 1996). In cognitive science and linguistics, traditional approaches to classification (which are often called "Aristotelian" approaches although that can be misleading) compete with fuzzy, prototype-based models that can be traced back to Wittgenstein (1953) and Rosch (1973). In information science, manually designed static and thus formal information access systems (such as Yahoo!) compete with statistical, dynamic systems or systems that are based on the network character of the Web (such as Google). In philosophical logic, Wille (2001) distinguishes "mathematical" (formal) and "logical" approaches. In this case "logical" is not used in the AI sense of symbolic logic but more in the sense of human common sense reasoning and thus more similar to "associative".

In physics, there is quantum mechanics and relativity theory - but physics is probably the only discipline where it is nowadays widely accepted that both models must co-exist! In history of science, there are externalists for whom science is a complex system within a society and internalists for whom science is a more formal collection of abstract ideas (Bauer, 1992, p. 111). And last but not least, social scientists distinguish between positivist and post-modernist theories, which, respectively, employ more formal and more associative means of argumentation.

Not all of these are precise examples of what are called associative and formal structures in this paper but overall the distinctions are similar in nature. They also all share in common the fact that proponents of the different sides often fiercely disagree with each other, occasionally, even to the extent of ridiculing each other. For example, compare the positivist versus post-modernism episode of the Sokal (1996) affair in which a physicist published a fake paper in a social science journal for the sole purpose of discrediting post-modernism. Of course, in the process of doing so he unintentionally paints a dubious picture of extreme positivism as well.

3 Arguments for combining associative and formal structures

Apart from the pragmatic but not scientific argument that by accepting both approaches as two co-existing representational facets the sometimes unfruitful controversy in some fields could be turned into a more productive atmosphere, the following additional arguments can be made for combining associative and formal approaches at least in AI.

Individual approaches are unsuccessful

Both associative and formal approaches by themselves have not yet succeeded in producing human-like artificial intelligence. While associative approaches have been very successful in building perceptive interfaces (such as face recognition), generating and controlling movement (such as in robots or in compensating for movement in hand-held cameras) and controlling variables in dynamic systems (such as controlling tire pressure in a car), so far it has not been possible to generate natural language of considerably quality in such a manner. On the other hand, while symbolic logic has been useful in some expert systems, for creating chess computers and other software applications, even more than 15 years of manual labor invested in the CYC ontology (Lenat, 2000) have also not yet led to completely satisfactory language processing capacities. Therefore, if both approaches individually have so far not been entirely successful in creating human-like artificial intelligence, maybe they have failed because intelligence requires a combination of several approaches.

Multiple levels of representation

There are a variety of phenomena, such as feedback loops, emergence and what Clark (1997) calls "mangrove effects", which intrinsically require several components of systems or levels of representations. Feedback loops involve an interaction between different components of a system. Maybe associative and formal representations together can invoke feedback loops which initiate a growth of representational capacities that are significantly better than any individual representation can produce (compare section 6).

Many emergent phenomena can be explained as a shift of representational viewpoints. For example, Clark describes the movement of a traffic jam as an emergent phenomenon. Cars follow local rules that describe their behavior. They all move forward. If the conditions are right – sufficient numbers of cars that travel at a sufficient speed – a traffic jam can emerge. In contrast to the forward movement of the cars, the traffic jam moves backward. But to recognize this emergent phenomenon, the viewpoint has to shift from the micro-level of individual cars and their local behavior to the macro-level of the behavior of the traffic jam as a whole. Thus to recognize, describe or use emergent phenomena, at least two levels of representation are required.

Another phenomenon described by Clark (1997) are "mangrove effects". These are effects that arise from shifting between different representations similarly to how mangroves create land from the sea which itself provides an opportunity for more mangroves to grow. For example, humans often employ external modes of representation, such as drawing graphs or scribbling notes, when solving abstract problems, which could also be entirely formally described.

These three phenomena, feedback loops, emergent structures and mangrove effects, all facilitate growth of representational capacities but require the presence of multiple representational levels. Therefore, maybe associative and formal structures are such two representational levels which can facilitate growth of expressibility once they are combined.

Formal structures are grounded in associative structures

Results from research in animal cognition demonstrate that there is a stepwise development of intelligent behavior. Priss (2002) relates these steps to a stepwise development of conceptualizations. This research supports the idea that behavior that requires very little conceptualization or in other words very small brain masses, can be easily modeled with artificial neural networks and other associative structures. For example, artificial life simulations of insect behavior are easily obtained. Similarly, brain functions that have developed at an early evolutionary stage, such as reflexive behavior, are easily modeled with associative representations.

Formal concepts, however, require the existence of a natural or formal language and thus occur at a later evolutionary stage. Formal concepts do not replace associative concepts but instead build on them and receive a grounding from them. This is because evolution occurs gradually and older structures are usually integrated, adapted or exapted but not usually discarded. The idea of grounding language (and thus formal concepts) in associative structures is supported by the research of the embodiment of cognition as expressed, for example, by Lakoff & Johnson (1999).

Thus associative structures by themselves are sufficient to model some evolutionary older cognitive behaviors but formal structures by themselves are necessary but not sufficient to model evolutionary younger cognitive behaviors. Therefore it is not surprising that associative structures in AI are more successful if their tasks are restricted to associative tasks, whereas formal structures by themselves can never be as successful because they require associative structures. But an assertion that associative structures by themselves might be used even to model formal structures, such as human language, is pure speculation and not supported by any evidence. For example, Regier et al. (2001) claim that they have built a model that explains how words (formal concepts) can emerge directly from associative structures without a qualitative change, i.e. without explicitly invoking formal structures. But their reasoning is misleading in this aspect because their network receives two types of input: potential forms and potential referents! Thus a formal versus associative distinction is built into the system from the start.

4 Features of formal and associative structures

There are several characteristic features that distinguish associative and formal structures. Table 1 summarizes the features and is partly inspired by Blank's (2001) and Sloman's (1996) lists. Both formal and associative structures are dependent on the contexts in which they are viewed. The same object can be both formal or associative in different contexts. For example, unicorns are initially formal because they are invented. But humans can integrate them into a virtual world and think about them in the same associative manner and level of detail (e.g. think about their smell, texture, shape, movement) as they do about horses. Devlin (2000) explains that mathematicians think about formal mathematical objects in an associative manner. Formal and associative modes of thinking are thus highly intertwined in human cognition.

Complexity and representation

The objects of associative structures are embedded in an external world or in an externalized virtual world. Agents do not have direct access to such objects, but instead their perceptions of objects are determined by the agent's perceptive apparatus and by the knowledge structures that the agent obtains through socialization. Clark (1997) explains that humans do not have a complete model of an external world in their minds. Instead world models are continuously updated and completed by perceptual input. Chaitin's (1999) research, which provides a continuation of Gödel's ideas, shows that even mathematics is quasi-empirical. Chaitin's claim is that most of the objects of mathematics are complex, random and can be observed but cannot be explained by axiomatic theory. Complexity, randomness and depth of detail are features of associative structures and their embedding in a complex external system. Lenat (2000) refers to such objects as "rich objects". But the complexity only persists as long as the objects are viewed in their embedding in a local, temporal-spatial and cultural context. From a more abstract, formal viewpoint or context, which is independent of time, space and culture, the objects must be represented in a more shallow, less detailed manner. In mathematics, such a formal context can be represented by an axiom system. Chaitin's work shows that each (finite) axiom system can only provide explanations for a finite set of objects which must not be significantly more complex than the axiom system itself. Thus formal contexts are "globally" valid because they are independent of time, space and culture but they only contain limited, shallow content.

Formal structures are entirely symbolical and completely defined by their relationships to other symbols. Although Lakoff & Johnson's (1999) claim that all human philosophy has been shaped by our bodily experiences suggests that even formal thinking is always motivated by an external world, its motivation applies to a system of formal structures as a whole. But individual elements of a system of formal structures do not need a direct relationship with an external world. Individual formal structures can be represented entirely via relationships to other abstract entities. Formal structures can furthermore represent hypothetical situations, nested and recursive structures.

Both formal and associative structures need to be represented in some manner to be communicated. Formal concepts require to be represented in a symbolic language. But they can be renamed without affecting their meanings as long as their relationships to other formal concepts remain. They are thus entirely syntactic and do not require semantics. Associative concepts can be represented in non-linguistic formats (for example, as images or photographs) but they are much more dependent on their representations and their connotations. For example, highly associative forms of representation such as poetry can often not be translated; music or images can usually not be adequately described in words. Finding an appropriate representation for an associative concept is an essential part of associative reasoning.

Cognitive activities and styles of argumentation

The cognitive activities involved in associative and formal reasoning are different. Formal concepts and contexts are evaluated with respect to logical correctness, consistency and completion. Formal arguments do not require grounding but instead rely on logical inference. Associative cognitive activities, however, are based on filtering, co-occurrence, resonance and analogy. Filtering is necessary because of the potential infinite complexity and detail of external objects. Gestalt perception is an example of filtering. The formal equivalent of filtering is defining. Observation of co-ocurrent objects is the associative equivalent of formal explanations based on presumed causation between objects. Resonance refers to structural alignment. Shastri (1997) provides a low-level neuropsychological model of how cognitive structures, such as language comprehension and subtyping, can be represented as synchronous firing of cells, which is a form of resonance. Last but not least, analogy is the associative strategy that facilitates expanding knowledge into new areas. Analogy also involves structural alignment.

Associative concepts and contexts are evaluated with re-

	formal	associative
metaphors	information processing	biologically inspired (neural networks, evolution)
design	top-down	bottom-up
objects:	symbolic, shallow abstract	complex, grounded “rich”
concepts:	precise, defined designed	fuzzy, emergent filtered, observed gestalt-based
conceptual structures:	type lattices	semantic networks diagrams
contexts:	intensional axiomatic, shallow global	extensional associative, deep temporal, spatial, cultural
representations:	symbolic no connotations syntactic	subsymbolic or symbolic connotations semantic
systems:	simple, organized correct, complete macro-level discrete, static	complex, decentralized probabilistic, dynamic micro-level continuous, “chaotic”
extensions: intensions:	set of “objects” formal definitions formal logic rules	prototypes, exemplars stereotypes, image schemata family resemblances
reasoning:	logical, formal causal explanation hierarchical inference, deduction classification correctness	similarity, contiguity co-occurrence observation associative analogy clustering resonance
implementations	formal logic algorithmic, modular	fuzzy logic learning-based

Figure 1: Features of associative and formal structures

spect to their grounding. They are convincing to an agent if they resonate with observations and prior knowledge. Of course, most human arguments contain both formal and associative elements. An example for this difference between associative and formal arguments is the fact that it is possible to follow the logical structure of a mathematical proof and to be convinced of its correctness but to still not understand its meaning if the proof does not resonate with prior knowledge. Another example of this difference are political discussions about scientific topics. These discussions can be expected to be more associative in nature, i.e. referring to social, ethical and psychological aspects, than formal, i.e. referring to the formal properties of the topic. Confusion arises, if the participants are not aware of these two levels and inappropriately apply formal arguments to associative problems and vice versa.

Human reasoning usually involves both associative and formal methods in combination because without associative concepts, reasoning would have no ground and no relationship to an external world. Without formal concepts, reasoning would be limited to objects within the actual physical environment and their evident features and relations. Without formal structures, no broader consequences, abstract structures or possibilities could be considered. Devlin's (2000) main conclusion about what differentiates people with and without a mathematical ability is that mathematicians are capable of thinking about mathematical objects in the same gossip-like (and thus associative) manner as other people think about soap operas. Bauer (1992) explains that science itself does not follow the formal model of the "scientific method" but instead also heavily depends on certain other social and instrumental (and thus associative) factors. It is thus important that formal models about reasoning include both the formal and the associative aspects of reasoning.

5 A definition of concepts

If formal and associative structures are to be combined, it is essential to have a solid understanding of how they interact. Models are needed that shed some light on the features, which were mentioned in the previous section, and their relationships. Priss (2002) describes such a model in form of a ten-fold classification of concepts. Without presenting too much detail from that paper, it should be mentioned that the model is derived by analyzing concepts with respect to a triadic relationship between objects, internal representations and signs.

Objects are elements of an external world or of an externalized virtual world filtered by an agent based on perception and influenced by the agent's socialization. Internal representations are employed by agents in their cognitive processes. In most cases, internal representations are opaque to an observer and even to the agents themselves. But some psychological tests provide insights into internal representations used by humans and animals. Signs are external representations (or behaviors) that according to the viewpoint of an observer are produced by agents in response to encountering objects. Because internal representations are opaque and because observers are usually different from producers of signs, there is always some ambiguity involved in "reading signs", i.e., in interpreting the relationship between objects and signs.

Concepts are defined in this paper as external models for internal representations. They consist of an extension, which contains the objects to which they refer, and an intension, which contains features, attributes and types of the objects. Because internal representations are opaque and signs are ambiguous, concepts are models or attempts at explaining the relationship between objects and signs from the viewpoint of an observer. There are several levels of contextuality involved in each concept: the context of the object, the context of the sign producer and the context of the observer. If the produced sign is an element of a sign system, such as a natural language, the sign itself may also invoke a further, linguistic context.

The main purpose for this definition of concepts is to provide a means for explaining the relationship between objects and signs in a dynamic, context-dependent manner. Only in very primitive organisms, sign production is a simple input-output function. For example, reflexive behavior can be modeled as simple, deterministic stimulus-response behavior and does not really require concepts as a model. But this is an external view: internally even simple stimulus-response behavior can be complicated. Complex behavior is usually context-dependent, dynamic and dependent on the choices an agent makes. This behavior cannot be modeled as input-output functions. But it can be modeled as a result of conceptual processing in a network of concepts.

The classification in Priss (2002) describes a stepwise (but ordinal, not linear) progression from stages without conceptualization, to stages with associative concepts and a finally to a stage of formal concepts. The formal difference between associative and formal concepts in this classification is that the objects of associative concepts are filtered external objects whereas the objects of formal concepts are symbols. Some of these symbols refer to associative concepts, others refer to other formal concepts. Formal structures derive meaning from those symbols that refer to associative concepts. Meaningful formal structures thus receive some grounding from associative structures. But this grounding does not affect the consistency of a formal structure because consistency is entirely formally defined.

As mentiond in section 3, an argument for a distinction between associative and formal structures are the feedback loops that arise from the interaction between these two modes of representation. The classification in Priss (2002) facilitates the identification of precise mechanisms that are involved in such feedback loops. The next section explains how these dynamic structures can be exploited.

6 Combining associative and formal structures

A first associative stage

A system that successfully exploits both associative and formal structures has to start with an associative perceptive interface. The system learns to recognize certain patterns and gestalts from continuous input data. The identification of gestalts can be unsupervised. But the filtering process, which forces the system to focus on those gestalts that are more relevant to the system, requires either scaffolding based on spatial interactions with the environment and with other agents or it requires supervision. The filtering process is intensified by associations between gestalts and symbols which the system learns from its interactions with other agents. These symbols are not yet part of a language, but simply labels for associative concepts, such as occurring in animal calls or in the linguistic development of a one to two year old child. The existence of labels, without any syntax, presents a strong guidance for the agent's focus of attention. The agent must assume that labeled concepts are of higher importance than unlabeled ones. The driving force at this stage that leads to a maximum amount of learning is the interaction with the environment and the other agents. The more uses for objects an agent learns, the more concepts it develops. Regier et al.'s (2001) paper about the "emergence of words" describes an artificial neural network that accomplishes most of the tasks required at this stage.

A second associative stage

The next developmental step leads to what Devlin (2000) calls "protolanguage". At this stage, the agent learns to combine labels for different concepts in simple two word associations. These associations do not follow any syntactic rules but are essentially just co-occurrences or phrases. It has been shown that gorillas can learn to communicate at this level. The lack of syntax does not pose a problem because the semantics of these associative structures is provided by their grounding in an external world. For example, it is obvious what "Koko banana" means in a context in which Koko does not have a banana but someone else does.

At this stage, an agent can start to express metaphors because the agent can observe features in the external world and can learn to express these and apply these to other similar situations. Metaphor (or analogy) is the driving force that leads to a maximum of conceptual development at this stage. Formally, metaphor involves gestalt and pattern recognition on a network of associative concepts, but is guided by the agent's interactions with an external world. Metaphor development initiates a feedback loop with the previous stage because seeing objects from a metaphoric viewpoint can lead to the discovery of new objects.

A formal stage

At the formal stage, the agent needs to learn a formal syntax that facilitates combining almost any symbol with almost any other symbol. A driving force at this stage is a desire of the agents to formally establish relationships between all objects that they encountered so far. New objects can arise from the processes of defining, classifying, role filling, and establishing of logical inferences. Feedback loops to both previous stages can be initiated if new objects are formally defined and need to be established in a virtual external world. As before, all activity within the formal realm is guided by the agent's interaction with the external environment and other agents.

7 Conclusion

This paper argues that formal and associative structures need to be combined in artificial intelligence systems to achieve human-like cognitive abilities for artificial agents. While some of the arguments provided in this paper are theoretical and have not yet been tested in implementations, there are implementations (such as Shastri (1997) and Regier et al. (2001)) that support some of the major steps mentioned in this paper. Future research will determine how exactly the combination of associative and formal structures in AI can be achieved and whether the predictions made in this paper are accurate.

References

Blank, Douglas S. (2001) *Radical Artificial Intelligence: A Postmodern Approach.* In: Moore (ed.). Proceedings of the 12th Midwest Artificial Intelligence and Cognitive Science Conference.

Bauer, Henry H. (1992) *Scientific Literacy and the Myth of the Scientific Method.* University of Illinois Press.

Chaitin, G. J. (1999) *The Unknowable.* Springer Verlag.

Clark, Andy (1997). *Being There. Putting Brain, Body, and World Together.* MIT Press.

Devlin, Keith (2000) *The Maths Gene. Why everyone has it, but most people don't use it.* Weidenfeld & Nicolson, UK.

Lenat, Douglas (2000). *The Dimensions of Context Space.* On-line available at http://www.cyc.com/publications.html.

Lakoff, George; Johnson, Mark (1999). *Philosophy in the Flesh.* Basic Books, New York.

Priss, Uta (2002). *Associative and Formal Concepts.* Proceedings of ICCS'02. (to appear)

Regier, T.; Corrigan, B.; Cabasaan, R.; Woodward, A.; Gasser, M.; Smith, L. (2001). *The Emergence of Words.* Proceedings of the 23rd Annual Meeting of the Cognitive Science Society, p. 815-820.

Rosch, Eleanor (1973). *Natural categories.* Cognitive Psychology, Vol 4.

Russell, S.; Norvig, P. (1995) *Artificial Intelligence: A Modern Approach.* Prentice Hall.

Shastri, Lokendra (1997). *A Model of Rapid Memory Formation in the Hippocampal System.* Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society, p. 680-685, Stanford University, CA. August.

Sloman, S. A. (1996). *The empirical case for two systems of reasoning.* Psychological Bulletin, 119, p. 3-22.

Sokal Affair, The (1996) On-line available at http://www.drizzle.com/ jwalsh/sokal/articles/articles.html

Wille, Rudolf (2001). *Why Can Concept Lattices Support Knowledge Discovery in Databases?* In: E. Mephu Nguifo (ed.) ICCS-2001 International Workshop on Concept Lattices-based Theory, Methods and Tools for Knowledge Discovery in Databases.

Wittgenstein, Ludwig (1953) *Philosophical Investigations.* Basil Blackwell, Oxford.