Case Based Reasoning as a Model for Cognitive Artificial Intelligence

Susan Craw[1] and Agnar Aamodt[2]
0000-0003-1870-0323 0000-0003-0114-8931

[1] School of Computing Science & Digital Media
Robert Gordon University, Aberdeen, UK
s.craw@rgu.ac.uk www.rgu.ac.uk/dmstaff/craw-susan
[2] Department of Computer Science
Norwegian University of Science and Technology, Trondheim, Norway
agnar@ntnu.no www.ntnu.edu/employees/agnar

Abstract. Cognitive Systems understand the world through learning and experience. Case Based Reasoning (CBR) systems naturally capture knowledge as experiences in memory and they are able to learn new experiences to retain in their memory. CBR's retrieve and reuse reasoning is also knowledge-rich because of its nearest neighbour retrieval and analogy-based adaptation of retrieved solutions. CBR is particularly suited to domains where there is no well-defined theory, because they have a memory of experiences of *what* happened, rather than *why/how* it happened. CBR's assumption that '*similar problems have similar solutions*' enables it to understand the contexts for its experiences and the 'bigger picture' from clusters of cases, but also where its similarity assumption is challenged. Here we explore cognition and meta-cognition for CBR through self-reflection and introspection of both memory and retrieve and reuse reasoning. Our idea is to embed and exploit cognitive functionality such as insight, intuition and curiosity within CBR to drive robust, and even explainable, intelligence that will achieve problem-solving in challenging, complex, dynamic domains.

1 Introduction

Cognition is human-like understanding[3] of the world, context, etc., and cognitive systems aim to understand the world in a way similar to what humans do, through senses, learning, and experience [1]. Langley's 2012 article in the inaugural issue of Advances in Cognitive Systems [2] highlighted the need for AI to refocus on the human intelligence aspect from its early days. Recently there has been an upsurge of interest in cognition in AI with special issues of the AAAI AI Magazine and IEEE Intelligent Systems [3–5]. These papers highlight the impor-

[3] We use the term '*understanding*' in the sense of '*interpret in order to give meaning*' for the system involved.

tance of understanding, context and analogy for Future Intelligent Technologies (FIT)[4].

In early 2017 Launchbury published DARPA's perspective on AI [6] where he reflected on 3 waves of AI: the early approaches relying on handcrafted knowledge, and later statistical learning, and the need now for new advances in contextual adaptation. In contrast to Launchbury's historical timeline, Domingos proposes 5 'tribes' (i.e. classes) of learning algorithms: symbolic, connectionist, evolutionary, bayesian and analogical [7]. He suggests that Machine Learning in around 10 years will be dominated by deep analogy, and envisions a '*Master Algorithm*' that combines nearest neighbour, Support Vector Machines (SVMs) and analogical reasoning [8]. Forbus and Hinrichs' *Companion Cognitive Architecture* [9] also highlights the key role of analogical reasoning and the utility of qualitative representations. Case Based Reasoning (CBR) already takes advantage of many of these ideas: cases capture the context in which the experience occurs, retrieval uses nearest neighbour, adaptation is a key part of reuse, and analogy is the basis of reasoning in CBR.

Kahneman [10] proposes a classification of reasoning as '*Fast Thinking*' with intuitive, quick, stereotypical decisions, or '*Slow Thinking*' with deliberative, calculating, logical reasoning. 'Fast Thinking' may easily lead to errors, and Kahnemann gives many examples of that, while 'Slow Thinking', being more deep and elaborate, can act as a censor and make necessary corrections. We may think of CBR with a simple retrieve & reuse reasoning as replicating 'Fast Thinking' because of its assumption of intuition that similar problems will have similar solutions. In contrast, 'Slow Thinking' CBR is when similarity knowledge is complex, retrieved cases are conflicting, when adaptation is complicated or computationally demanding, etc. A CBR system is able to do both 'Fast Thinking' or 'Slow Thinking' depending on the complexity of retrieval and reuse. Its 'Fast Thinking' errors are when similar problems do NOT have similar solutions!

Gartner's dimensions of machine smartness [11] highlight the need for cognitive intelligence: handling complexity; making confidence-based predictions; learning actively/passively; acting autonomously; appearing to understand; and reflecting a well-scoped purpose. These match the three 'Ls' of cognitive computing, *Language, Learning* and (confidence) *Levels* [12]. Case-based systems go some way towards Gartner's complexity, confidence and passive learning criteria [11]. More ambitious self-reflection, introspection, and curiosity is needed to advance towards Gartner's active learning criteria and the demands of complex/changing contexts.

This paper considers the explicit knowledge in cases, but also the similarity and adaptation knowledge, in order to explore the system's *understanding* of its knowledge. This approach also broadens the CBR system's understanding to implicit knowledge from *collections* of cases, and *interactions* between case, retrieval and reuse knowledge. This will enable the system to exploit 'Fast Thinking' when possible and to streamline 'Slow Thinking' when necessary. For

[4] EPSRC FIT Priority www.epsrc.ac.uk/research/ourportfolio/themes/ict/introduction/crossictpriorities/futureintelligenttechnologies/

this we take advantage of both cognition and meta-cognition, using Cox's interpretation of Minsky's World and Self models [13]. The rest of this paper is organised as follows. Section 2 considers the knowledge sources of a CBR system and its knowledge of itself. Sections 3 and 4 explore the system's understanding of its knowledge and of itself, through cognition and meta-cognition, and how this understanding can be used for self-adaptation. The impact of cognitive CBR through its knowledge and understanding is discussed in Section 5. Section 6 highlights related work from a variety of angles, before we draw some conclusions about cognitive CBR as a model for cognitive AI in Section 7.

2 What a CBR System Knows

A CBR system has knowledge of the world it models, but it also has knowledge of itself. Its world knowledge comprises its cases, and in addition reasoning knowledge through similarities and adaptations. CBR contains explicit knowledge represented as symbolic structures in some knowledge representation format. The reasoning knowledge may be explicit, allowing for meta-level reasoning, or implicit in the underlying procedures. CBR's world knowledge is well-understood. Richter's notion of knowledge containers [14] identifies four different types of knowledge (vocabulary, cases, similarities and adaptations). Richter also highlights the interactions between containers, and the possibility of moving knowledge between these, as shown in Figure 1. We shall take advantage of interactions between knowledge containers in the following sections.

CBR's self knowledge is that its memory of cases contains things it believes to be true, its similarity assumption that similar problems will have similar solutions, and that differences between query and retrieved cases may/should be reflected in the new solution.

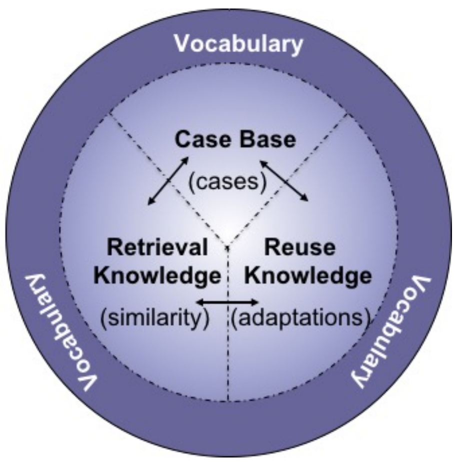

Fig. 1. Knowledge Containers (adapted from [14])

The notions of world and self knowledge fit well with Donald Rumsfeld's (in)famous statement on 'Knowns/Unknowns' [15]:

> "*. . . there are no 'knowns'. There are things we know that we know* [Known Knowns]. *There are* Known Unknowns. *That is to say there are things that we now know we don't know. But there are also* Unknown Unknowns. *There are things we don't know we don't know.*"

We add the missing combination *Unknown Knowns* for 'things we don't know we know' to Rumsfeld's list.

Figure 2 shows a problem-solving space where the horizontal axis is the system's knowledge about the *world*: what it knows towards the right, and what it does not know on the left. Similarly what the system knows about *itself* is on the vertical axis, with what it knows in the lower half, and what it does not know above. The shaded areas place the four Known/UnKnown combinations in the appropriate quadrant. For CBR the lower right Known Knowns quadrant contains things the system knows it knows; e.g. the cases in memory. The upper Unknown Knowns quadrant is information that the system does not know it knows; e.g. these may be problems that are not in the case base and CBR does not know if nearest neighbour retrieval will generate the right solution. The lower left Known Unknowns quadrant contains things the system knows it does not know; e.g. these may be problems where similar cases contain very different solutions so CBR is not confident of the solution. The upper Unknown Unknowns quadrant is information that the system does not know it does not know; e.g. these may be problems that are outliers or not similar enough to cases in the case base.

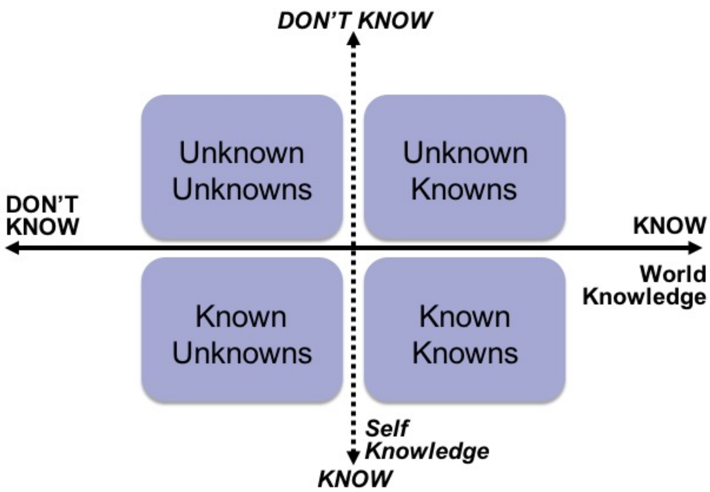

Fig. 2. Self vs World Knowledge from 'Known Knowns' to 'Unknown Unknowns'

This paper will explore a CBR system's understanding of the world knowledge 'it knows it knows' to enable better cognition in the other quadrants. For this we shall view a CBR system as a case based memory of experiences and its retrieve and reuse reasoning. Both memory and reasoning are essentially knowledge based representations [14], and so cognition in one area may enable refinements in this area or others.

We shall take advantage of Cox's work on cognition and meta-cognition [13]. Cox associates cognition with self-reflection as a system understands or makes sense of what it knows. Metacognition is cognition about cognition, or making sense of understanding, and is knowledge about when and how to use particular strategies for problem-solving. Cox associates meta-cognition with introspection. Figure 3 illustrates these ideas. Cognition is shown as 'making sense of' or 'understanding', and Meta-Cognition as 'selection strategies' or 'understanding errors'. Cognition for a model M is self-reflections or understanding of M (denoted M*), and meta-cognition is introspection or understanding of M* (M**). This diagram will be used to underpin the following 2 sections to explore cognition and meta-cognition for a CBR system composed of a case based memory and retrieve & reuse reasoning.

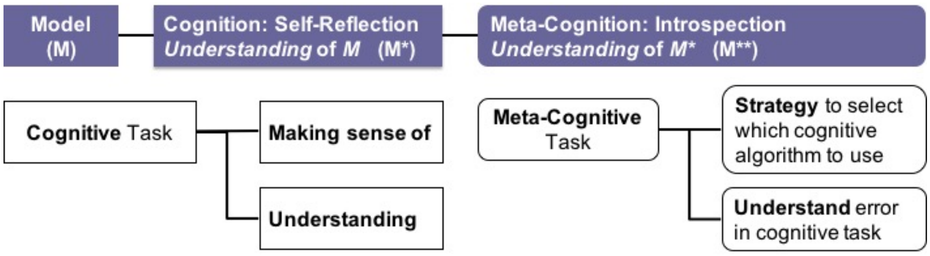

Fig. 3. Cognition and Meta-cognition

3 Cognition from Self-Reflection

In this section we explore cognition for a CBR system by looking at the system's understanding of its memory of cases and its reasoning, and how it makes sense of its case and reasoning knowledge. Figure 4 applies cognition in Figure 3 to a CBR system, describing cognition as context for each case, insight of the domain from collections of cases, intuitive reasoning from nearest neighbour, and analogy to take account of problem differences. Reflections on memory provide understanding of each experience as a whole and the relationships among its various facets, and the landscape captured by the collection of experiences. This will enable an understanding of implicit knowledge corresponding to 'what you don't know you know'. Understanding the reasoning will exploit the fundamental

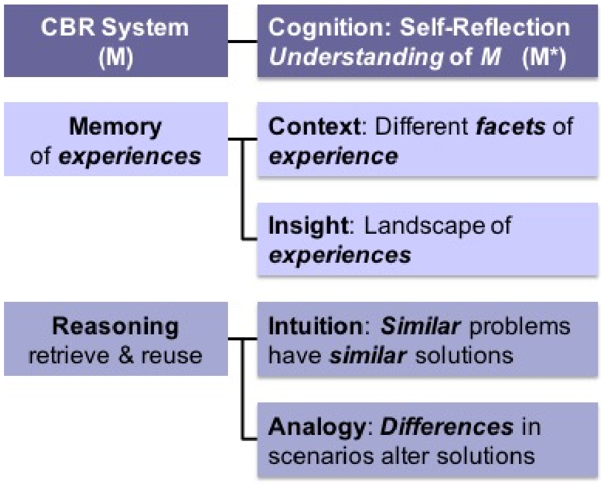

Fig. 4. Cognition with CBR

assumption of case based systems that 'similar problems have similar solutions', but that some differences are significant and alter the solution.

A case captures a collection of related facets for an experience, and different combinations of facets can be used as a specification or scenario in order to retrieve solutions or suggestions contained in the other facets. Learning relationships between sets of facets within clusters or neighbourhoods of the case base allows the identification of important concepts and relationships between them, and so an understanding of the different contexts in which each case is relevant.

The collection of cases offers an opportunity to understand the landscape for the domain. Areas where there are many similar cases could validate the contents of individual experiences but it also shows where reuse of similar experiences is less risky. In contrast, areas where the problem and solution spaces are not well aligned mean the landscape is complex and more reasoning is needed to reuse these cases. Competence and complexity models for CBR maintenance [16, 17] use a similar approach, but these identify redundant or noisy cases whereas here we are interested in areas where 'slow reasoning' may be needed.

Case based systems assume that 'similar problems have similar solutions' and so understanding the reasoning becomes understanding the alignment of cases in similarity space. In areas of regularity, where similar problems do indeed have similar solutions, an intuitive reasoning that reuses similar cases is appropriate, but in complex areas a more sophisticated reasoning is needed. A more complex, finer-grained local similarity can be learned [18], or an uncertainty-based reuse of multiple similar cases is needed by mining neighbourhoods.

The use of analogy to exploit cases beyond their areas of intuitive reasoning may require reuse that includes significant adaptation to take account of differences across the neighbourhood. Understanding when differences between cases in similarity space become significant allows adaptations that reflect the differences in scenarios as alterations to solutions. This is based on learning adaptation knowledge by understanding relationships among cases; e.g. ensemble learning for adaptation [19], or gradient learning for adaptation [20].

4 Meta-cognition

Meta-cognition is 'cognition about cognition' and so is making sense of cognition. Cognition in Section 3 has developed the 'bigger picture' of what the system knows; here we explore how the system should know things that it currently solves wrongly. Meta-cognition is the system understanding when and how to use particular knowledge and strategies for problem-solving. Figure 5 shows meta-cognition from Figure 3 applied to a CBR system. Here we explore how introspective models capture an understanding of how the CBR system should know something, by making sense of different contexts and insights within memory, and by understanding retrieve and reuse reasoning failures. Under meta-cognition we also include curiosity in which an extrospective curiosity builds understanding from external sources.

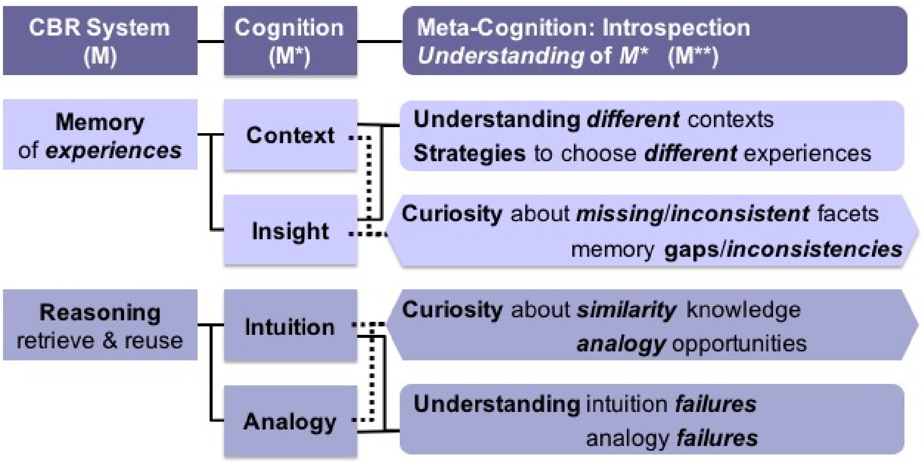

Fig. 5. Meta-cognition with CBR

4.1 Meta-cognition from Introspection

Introspection for CBR memory focuses attention on understanding different contexts in cases. Clusters of similar cases in the problem space allows different selections of facets or features to be identified as key features for similarity matching. Areas of redundancy in the case base enables rich alternative views of

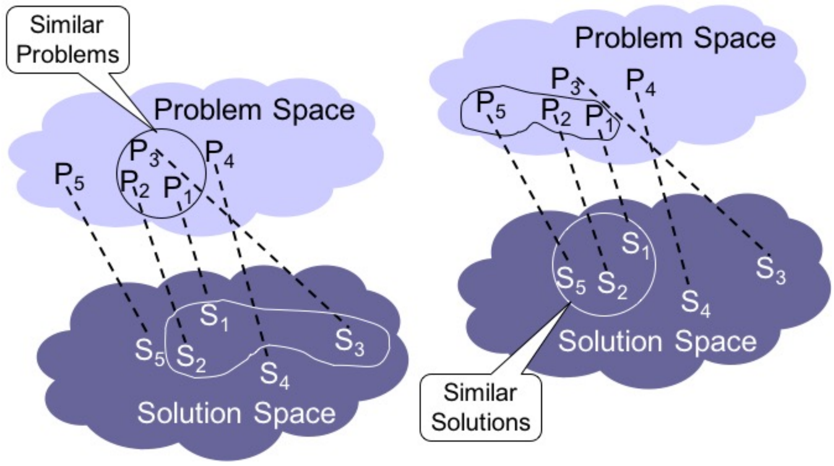

Fig. 6. Problem-Solution Alignment

the context of an experience. Clusters of similar cases in the solution space can identify dimensions in the problem space where similarity is found. By taking advantage of similarity in the problem or solution spaces, we can define feature selection strategies to create different contexts from these cases. Figure 6 shows two different views of 5 cases comprising problem-solution pairs (P_i,S_i). The left diagram shows a circle of neighbouring problems P_1, P_2, P_3 in the problem space, and a less regular cluster of the corresponding solutions S_1, S_2, S_3. P_4 is a neighbouring problem whose solution S_4 is closer to the others than S_3. P_4 can give important pointers about features and similarity in the problem space in relation to P_1 and $P2_2$ and in comparison with P_3. The diagram on the right shows the same cases but now focuses on the neighbourhood S_1, S_2, S_5 in the solution space. This highlights P_5 as a potentially useful neighbour of P_1 and P_2. In a similar way as previously, P_5 can help to highlight important features and similarity in the P_1, P_2, P_3 region of the case base.

By understanding how different facets or experiences are relevant, new selection strategies can be learned. Richter's knowledge containers allow knowledge to be shifted between containers. So it is natural that a given selection strategy can be implemented within different knowledge containers; e.g. a different memory selection can be achieved by altering the representation (different facets) or the retrieval knowledge. For areas of complexity in the case base, cases with similar problems do not have similar solutions. We might take advantage of the contexts learned from other areas of the case base to reduce this complexity by feature/context selection. Alternatively we could identify this region as needing more focused search or more deliberative reasoning. Memory introspection takes advantage of self-reflection for reasoning to associate faulty solutions with the need to learn selection strategies that use alternative contexts or more narrowly focused regions within complex regions of the landscape of cases.

Introspection for retrieve and reuse reasoning in Figure 5 highlights understanding the failures to identify similar cases or to use analogy to adapt retrieved cases. So here we explore the system's understanding of faulty solutions where the reasoning, rather than memory, is to blame. These methods should go beyond explaining the failure, to understanding what may have caused the failure. The similarity based retrieval may be the cause of the failure, and understanding will involve repairing faulty retrieval knowledge or refining it by adding new similarity knowledge. If the reuse of the retrieved cases is to blame then the adaptation knowledge should be repaired or refined. As before, the interaction between different knowledge containers means that equivalent refinements or repairs can be achieved in similarity or reuse knowledge. An important aspect of understanding reasoning failures is exploring which options are available and where do changes have least potential impact and are most natural for future understanding issues. Introspection for reasoning also takes advantage of self-reflection of reasoning to understand the limits of intuition and the need for adaptation in more deliberative reuse.

4.2 Curiosity Towards Unknown Unknowns

Faulty solutions and reasoning failures that trigger introspection may also act as cues to instigate extrospection so that curiosity discovers new facets or experiences to expand the memory, or new similarities or adaptations that alter retrieval or reuse from memory, as shown in Figure 5. Whereas Section 4.1 focuses on how the current memory and existing reasoning can resolve failures, here we consider a proactive outward facing understanding, where curiosity searches to find external information that will alter memory and/or reasoning. Exploring external sources identifies relevant problem-solving knowledge that fills some of the knowledge gaps described previously as *known unknowns* and the particularly elusive *unknown unknowns*.

Curiosity-inspired learning may be triggered in response to faulty reasoning highlighted during problem-solving. However curiosity about gaps or inconsistencies in memory and reasoning knowledge can also come naturally from self-reflection in Section 3. Proactive learning strategies may be applied based on the system's awareness of its own competencies; e.g. in identifying relevant trending stories in social media. Introspective processes may help identify the type of information needed, but curiosity-driven learning will provide autonomous reasoning that interrogates web based memories to refine existing knowledge and assemble latent cases. Trust and provenance will play an important role in selecting knowledge sources that range from trusted, well established, domain relevant ontologies, through to unstructured, uncorroborated content on social media. Mixed strategies based on provenance, previous performance and extent of verification is needed to select and verify suitable sources.

5 Understanding in Cognitive CBR

The previous sections have explored cognitive extensions to CBR to enable understanding of CBR's memory and reasoning at different levels: self-reflection for cognition, introspection for meta-cognition and curiosity for exploration beyond the CBR system itself to discover relevant new knowledge and understanding. These cognitive enhancements have built on the multiple, and interacting, sources of knowledge in a CBR system, the knowledge containers.

Cognitive CBR can have *insights* from the collection of cases in its memory. Relationships between cases can uncover different facets that offer alternative scenarios for retrieval. Collections of cases offer a problem-solving landscape where localised generalisation makes sense. This enables it to develop *intuition* by knowing which contexts are relevant and where similar problems have similar solutions. However it also has an understanding of when and why a more *deliberative* reasoning is needed and how to apply relevant similarity based retrieval and analogy based reuse knowledge. *Curiosity* stems from an understanding that the memory should explore relevant external knowledge or that the reasoning needs to discover similarity or analogy knowledge that is not already available in the CBR system.

Figure 7 demonstrates our ideas of CBR and its cognitive enhancements superimposed on the (Un)Known (Un)Knowns diagram in Figure 2. The CBR system's memory contains the things *it knows it knows* about the world, the Known Knowns. Adding Self-Reflection offers an understanding of *what* the CBR system knows to discover what the system did not know it knew; i.e. its insights and intuition to uncover Unknown Knowns. Introspection provides an understanding of its understanding and so an understanding of *how* it *should* know things, the Known Unknowns. Curiosity takes steps towards Unknown Unknowns by understanding *how* it *should* know *what* it *should* know!

6 Related Work

IBM Watson demonstrates cognitive behaviour in the way it reasons about the facts that it has learned from the Web to 'flesh out' its concept model for a domain. It was able to reason about some '*Unknown Knowns*' when winning the Jeopardy! game show in the US [21]. Watson's Jeopardy! success depended on its DeepQA question-answering cognitive knowledge engine [22]. DeepQA combines Natural Language Processing, Machine Learning and Evidence-based Experimentation to reason about the meaning of queries, to discover relevant information from its memory of extracted facts, and to gather evidence to rank the candidate answers [23, 9:1–12,10:1–14,14:1–12]. IBM's vision for Watson is to exploit DeepQA to underpin decision support in specialised domains. However, priming Watson to understand a new domain is a significant challenge, as found with Healthcare Watson [24]. It must be able to extract meaning from new text content, to understand new questions/scenarios, and to reason about new con-

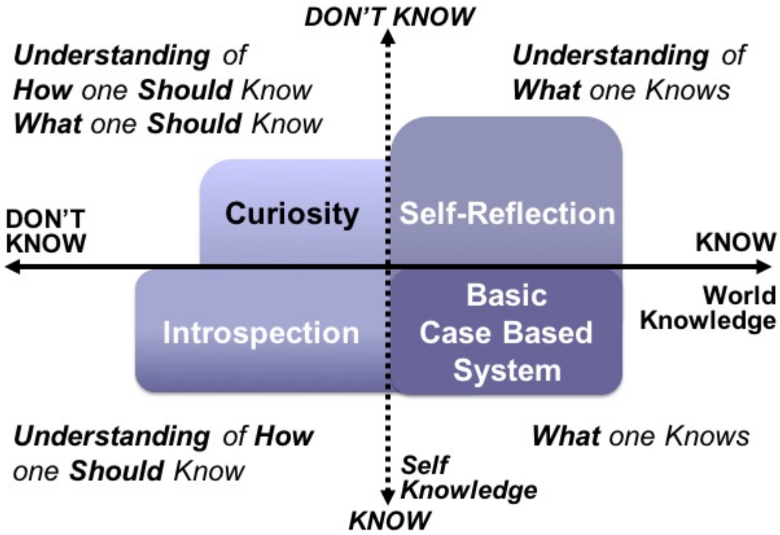

Fig. 7. Cognitive CBR

cepts [25]. Although IBM Watson Knowledge Studio[5] is designed to allow experts to teach Watson about a new domain, this instruction is quite knowledge poor – annotating texts to highlight domain entities and relationships. Nevertheless, Goel's application *Jill Watson*, the virtual teaching assistant, has been highly successful for supporting students during learning because its knowledge source of previous years student queries and answers is well matched to its task [26].

Case-based systems are a different sort of cognitive system; they are already knowledge-rich. Their knowledge is based on experiences in memory, but also explanations from their retrieve and reuse reasoning. Case-based systems, and cognitive systems more generally, apply knowledge-driven, localised, just-in-time search at run-time. This '*lazy learning*' contrasts with other learning approaches that create a generalised model of their data; e.g. Bayesian Networks, Neural Networks and Deep Learning. Watson DeepQA's cognitive reasoning about its knowledge contrasts sharply with Google DeepMind AlphaGo's deep learning of inscrutable 'value' and 'policy' networks [27], and CMU's poker-playing Libratus' efficient pruning of game trees [28]. Rather than capturing expertise in network models, Rubin & Watson's case-based poker-playing system captures decisions of expert players and its knowledge-driven approach reasons with, adapts, and learns from the play of experts [29, 30].

Planning domains are particularly amenable to cognitive approaches and Muñoz-Avila & Cox et al. are embedding cognition into architectures of planning systems [31, 32]. Researchers from MIT's CSAIL are trying to improve automated planners by giving them the benefit of human intuition [33]. Gottlieb

[5] https://www.ibm.com/watson/services/knowledge-studio/

et al. highlight links between curiosity in Psychology, and exploration in Active and Reinforcement Learning, as key to information-seeking behaviours [34].

The ideas in this paper have built on existing areas of research in CBR. Cognition and meta-cognition for memory relates to case base maintenance and TCBR indexing. Memory based reflection has been useful to identify redundant or noisy cases for case base maintenance [17, 35, 36]. Memory based introspection has been used in facet learning and case indexing. Meta-level reasoning has been used in a clinical decision support system for combining reasoning methods at run-time [37]. This was later extended to an architecture for learning how to select reasoning methods dynamically during execution time, using a lazy learning approach [38]. For recommendation, Smyth et al. have used opinion mining from reviews to learn relevant features for the products to be recommended [39, 40]. Curiosity builds on previous work on case discovery and case indexing [41, 42]. Cognition and Meta-Cognition for reasoning relates to CBR research in introspective learning of retrieval knowledge in changing environments [43], self-reflection for improving retrieval and reuse [44, 45], and introspective learning of adaptation knowledge to reuse retrieved solutions [46–48]. Introspection for reasoning also builds on previous research in textual contexts through understanding failures [49].

7 Conclusions

In this paper we have explored the possibility of extending CBR to embed cognition and meta-cognition. CBR offers a suitable framework for this enhancement because CBR comprises local independent cases in its memory and a just-in-time localised generalisation at run-time. Both its memory and reasoning are driven by explicit qualitative knowledge that allows experimentation and refinement. Compared to the generalised models of other AI systems, CBR is able to understand its knowledge and reasoning, and update it as needed. In this way CBR can use its existing framework to capture self-reflection and introspection.

Cognitive CBR may also address the features of Domingos' proposed Master Algorithm: nearest neighbour, SVMs and analogical reasoning [7]. CBR already uses nearest neighbour retrieval and analogical reasoning in its R^4 Retrieve-Reuse-Revise-Retain approach [50, 51]. Self-reflection and introspection enables cognitive CBR to achieve feature/facet learning, case refinement and local similarity learning. These could be thought of as learning the efficient problem solving representation corresponding to SVM's planes.

Self-understanding through reflection and introspection offers both cognition and meta-cognition, and thus provides opportunities for adaptive self-improvement towards a cognitive system with high competence and robust intelligence. Understanding of both self- and world-knowledge will also contribute to explainability. Cognitive CBR will underpin Explainable CBR (XCBR) since the system has its (self) understanding of its knowledge and problem-solving. Thus explainability for a human is transformed into interpreting the system's understanding and explanation into understanding in the human's view. There are links be-

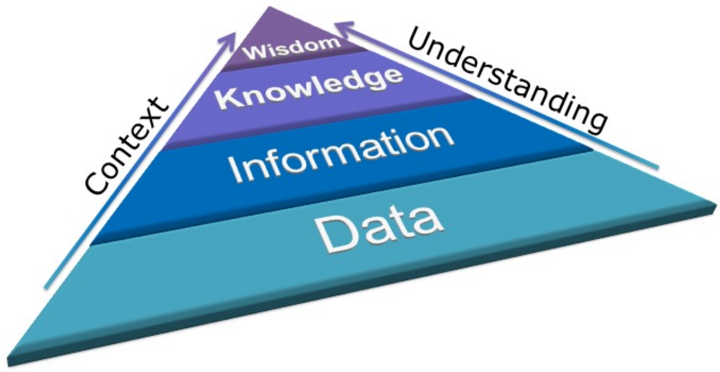

Fig. 8. Data – Information – Knowledge – Wisdom Pyramid

tween cognition and the important AI goals of explainability, competence, and robustness. As a result Cognitive CBR could make a valuable contribution to an XAI that is robust in complex and changing environments.

The well-known Data – Information – Knowledge – Wisdom Pyramid [52] shown in Figure 8 demonstrates the need for increased understanding and context as systems fit in the higher layers of Knowledge and Wisdom compared to Data and Information nearer the base. A CBR system certainly fits in the Knowledge layer through its knowledge in cases, the patterns and relationships captured by similarity based retrieval, and analogy-based adaptations in reuse. So does cognitive CBR achieve wisdom? Its understanding of the CBR system at the knowledge layer builds additional context, insight and intuition and so extends cognitive CBR beyond knowledge. We argue that cognitive CBR captures some aspects of Wisdom in its understanding and higher level reasoning, but human wisdom may include other more perceptive or emotional aspects not yet found in cognitive CBR.

Acknowledgments Susan Craw wishes to thank Stewart Massie for very helpful discussions of cognition and curiosity during the early stages of developing a research proposal on cognitive CBR that underpins this paper. Agnar Aamodt wishes to thank Enric Plaza for discussions of CBR, analogy, and cognition relevant to this paper, in preparing an invited talk at ICCBR 2017.

References

1. Kelly III, J.E., Hamm, S.: Smart Machines: IBM's Watson and the Era of Cognitive Computing. Columbia University Press (2013)
2. Langley, P.: The cognitive systems paradigm. Advances in Cognitive Systems **1** (2012) 3–13

3. Nirenburg, S.: Cognitive systems: Towards human-level functionality. AI Magazine **38**(4) (2017) 5–12
4. Langley, P.: Interactive cognitive systems and social intelligence. IEEE Intelligent Systems **32**(4) (2017) 22–30
5. Langley, P.: Meta-algorithms in cognitive computing. IEEE Intelligent Systems **32**(4) (2017) 35–39
6. Launchbury, J.: A DARPA perspective on artificial intelligence. (2017) www.darpa.mil/attachments/AIFull.pdf http://machinelearning.technicacuriosa .com/2017/03/19/a-darpa-perspective-on-artificial-intelligence
7. Domingos, P.: The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. Basic Books (2015)
8. Vasant, D., Pedro, D.: Pedro Domingos on the Master Algorithm: A conversation with Vasant Dhar. Big Data **4**(1) (2016) 10–13 doi: 10.1089/big.2016.29003.pdo
9. Forbus, K., Hinrichs, T.: Analogy and qualitative representation in the companion cognitive architecture. AI Magazine **38**(4) (2017) 5–12
10. Kahneman, D.: Thinking, Fast and Slow. Farrar Straus Giroux (2011)
11. Austin, T., LeHong, H.: Digital business innovation with smart machines. Report G00265758, Gartner, Stamford, CT (August 2014) www.gartner.com/doc/2818821/digital-business-innovation-smart-machines
12. Spohrer, J., Banavar, G.: Cognition as a service: An industry perspective. AAAI Artificial Intelligence Magazine **Winter** (2015) 71–86
13. Cox, M.T.: Metacognition in computation: A selected research review. Artificial Intelligence **169**(2) (2005) 104–141 doi: 10.1016/j.artint.2005.10.009
14. Richter, M.M.: Knowledge containers. http://pages.cpsc.ucalgary.ca/∼mrichter/Papers/Knowledge%20Containers.pdf
15. Rumsfeld, D. Press Conference, NATO HQ, Brussels (June 2002) www.nato.int/docu/speech/2002/s020606g.htm
16. Smyth, B., McKenna, E.: Competence models and the maintenance problem. Computational Intelligence **17**(2) (2001) 235–249 doi: 10.1111/0824-7935.00142
17. Craw, S., Massie, S., Wiratunga, N.: Informed case base maintenance: A complexity profiling approach. In: Proceedings of the 22nd AAAI Conference on Artificial Intelligence, Vancouver, Canada, AAAI Press (2007) 1618–1621 www.aaai.org/Papers/AAAI/2007/AAAI07-258.pdf
18. Jarmulak, J., Craw, S., Rowe, R.: Self-optimising CBR retrieval. In: Proceedings of the 12th IEEE International Conference on Tools with AI, Vancouver, Canada, IEEE Press (2000) 376–383 doi: 10.1109/TAI.2000.889897
19. Jalali, V., Leake, D., Forouzandehmehr, N.: Learning and applying case adaptation rules for classification: an ensemble approach. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), Melbourne, Australia, AAAI Press (2017) 4874–4878 http://static.ijcai.org/proceedings-2017/0685.pdf
20. McDonnell, N., Cunningham, P.: A knowledge-light approach to regression using case-based reasoning. In Proceedings of the 8th European Conference on Case-Based Reasoning. LNAI 4106, Fethiye, Turkey, Springer (2006) 91–105 doi: 10.1007/11805816_9
21. Murdock, J.W.: Structure mapping for Jeopardy! clues. In: Proceedings of the 19th International Conference on Case-Based Reasoning. LNAI 6880, Greenwich, UK, Springer (2011) 6–10 doi: 10.1007/978-3-642-23291-6_2
22. Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J., Schlaefer, N., Welty, C.: Building Watson: An overview of the DeepQA project. AAAI Artificial Intelligence Magazine **Fall** (2010) 59–79

23. Murdock, editor, J.W.: This is Watson (Special Issue). IBM Journal of Research and Development **56**(3-4) (2012)
24. Ferrucci, D., Levas, A., Bagchi, S., Gondek, D., Mueller, E.: Watson: Beyond Jeopardy! Artificial Intelligence **199-200** (2013) 93–105 doi: 10.1016/j.artint.2012.06.009
25. Kalyanpur, A., Murdock, J.W.: Unsupervised entity-relation analysis in IBM Watson. In: Proceedings of the Third Annual Conference on Advances in Cognitive Systems. (2015) 12:2–12:12 www.cogsys.org/papers/ACS2015/article12.pdf
26. Goel, A.K., Polepeddi, L.: Jill watson: A virtual teaching assistant for online education. Presented to the Learning Engineering for Online Learning Workshop, Harvard University (2017) http://smartech.gatech.edu/handle/1853/59104
27. David Silver, Aja Huang, et al: Mastering the game of Go with deep neural networks and tree search. Nature **529**(7587) (2016) 484–489 doi: 10.1038/nature16961
28. Brown, N., Kroer, C., Sandholm, T.: Dynamic thresholding and pruning for regret minimization. In: Proceeding of the 31st AAAI Conference on Artificial Intelligence, San Francisco, CA, AAAI Press (2017) 421–429 http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14855/13793
29. Rubin, J., Watson, I.: On combining decisions from multiple expert imitators of performance. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), Barcelona, Spain, AAAI Press (2011) 344–349 http://ijcai.org/Proceedings/11/Papers/067.pdf
30. Rubin, J., Watson, I.: Decision generalisation from game logs in No Limit Texas Hold'Em. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), Beijing, China, AAAI Press (2013) 3062–3066 http://dl.acm.org/citation.cfm?id=2540128.2540578
31. Cox, M.T., Alavi, Z., Dannenhauer, D., Eyorokon, V., Muñoz-Avila, H., Perlis, D.: MIDCA: A metacognitive, integrated dual-cycle architecture for self-regulated autonomy. In: Proceedings of the 30th AAAI Conference on Artificial Intelligence, Phoenix, AZ, AAAI Press (2016) 3712–3718 www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12292/12151
32. Muñoz-Avila, H., Dannenhauer, D., , Cox, M.T.: Towards cognition-level goal reasoning for playing real-time strategy games. In: Annual Conference on Advances in Cognitive Systems: Workshop on Goal Reasoning. (2015) 120–132 http://smartech.gatech.edu/handle/1853/53646
33. Kim, J., Banks, C., Shah, J.: Collaborative planning with encoding of users' high-level strategies. In: Proceeding of the 31st AAAI Conference on Artificial Intelligence, San Francisco, CA, AAAI Press (2017) 955–961 http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14840/13867
34. Gottlieb, J., Oudeyer, P.Y., Lopes, M., Baranes, A.: Information-seeking, curiosity, and attention: computational and neural mechanisms. Trends in Cognitive Sciences **17**(11) (2013) 585 – 593 doi: 10.1016/j.tics.2013.09.001
35. Lupiani, E., Massie, S., Craw, S., Juarez, J.M., Palma, J.: Case-base maintenance with multi-objective evolutionary algorithms. Journal of Intelligent Information Systems **46**(2) (2016) 259–284 doi: 10.1007/s10844-015-0378-z
36. Jalali, V., Leake, D.: Adaptation-guided case base maintenance. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence, Quebec, Canada, AAAI Press (2014) 1875–1881 www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/download/8360/8824
37. Houeland, T.G., Aamodt, A.: An introspective component-based approach for meta-level reasoning in clinical decision-support systems. In: Proceedings of the 1st

Norwegian Artificial Intelligence Symposium (NAIS'09), Trondheim, Tapir Forlag (2009) 121–132
38. Houeland, T.G., Aamodt, A.: A learning system based on lazy metareasoning. Progress in Artificial Intelligence (In Press, Online November 2017) doi: 10.1007/s13748-017-0138-0
39. Dong, R., O'Mahony, M.P., Schaal, M., McCarthy, K., Smyth, B.: Combining similarity and sentiment in opinion mining for product recommendation. Intelligent Information Systems **46**(2) (2016) 285–312 doi: 10.1007/s10844-015-0379-y
40. Dong, R., Schaal, M., O'Mahony, M.P., Smyth, B.: Topic extraction from online reviews for classification and recommendation. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), Beijing, China, AAAI Press (2013) 1310–1316 http://dl.acm.org/citation.cfm?id=2540128.2540317
41. Massie, S., Craw, S., Wiratunga, N.: Complexity-guided case discovery for case based reasoning. In: Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05), Pittsburgh, PA, AAAI Press (2005) 216–221 www.aaai.org/Papers/AAAI/2005/AAAI05-035.pdf
42. Recio-Garcia, J.A., Wiratunga, N.: Taxonomic semantic indexing for textual case-based reasoning. In: Proceedings of the 18th International Conference on Case-Based Reasoning. LNAI 6176, Alessandria, Italy, Springer (2010) 302–316 doi: 10.1007/978-3-642-14274-1_23
43. Craw, S., Jarmulak, J., Rowe, R.: Maintaining retrieval knowledge in a case-based reasoning system. Computational Intelligence **17**(2) (2001) 346–363 doi: 10.1111/0824-7935.00149
44. Sizov, G., Øzturk, P., Aamodt, A.: Evidence-driven retrieval in textual CBR: Bridging the gap between retrieval and reuse. In: Proceedings of the 23th International Conference on Case-Based Reasoning. LNAI 9343, Frankfurt am Main, Germany, Springer (2015) 351–365 doi: 10.1007/978-3-319-24586-7_24
45. Murdock, J.W., Goel, A.: Meta-case-based reasoning: Self-improvement through self-understanding. Journal of Experimental and Theoretical Artificial Intelligence **20**(1) (2008) 1–36 doi: 10.1080/09528130701472416
46. Craw, S., Wiratunga, N., Rowe, R.C.: Learning adaptation knowledge to improve case-based reasoning. Artificial Intelligence **170**(16-17) (2006) 1175–1192 doi: 10.1016/j.artint.2006.09.001
47. Arcos, J.L., Mulayim, O., Leake, D.B.: Using introspective reasoning to improve CBR system performance. In Metareasoning: Thinking about Thinking. MIT Press (2011) 167–182 doi: 10.7551/mitpress/9780262014809.003.0011
48. Leake, D., Powell, J.: A general introspective reasoning approach to web search for case adaptation. In: Proceedings of the 18th International Conference on Case-Based Reasoning. LNAI 6176, Alessandria, Italy, Springer (2010) 186–200
49. Jayanthi, K., Chakraborti, S., Massie, S.: Introspective knowledge revision in textual case-based reasoning. In: Proceedings of the 18th International Conference on Case-Based Reasoning. LNAI 6176, Alessandria, Italy, Springer (2010) 171–185
50. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. AI Communications **7**(1) (1994) 39–59 citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.39.1670
51. López de Mántaras, R., McSherry, D., Bridge, D., Leake, D., Smyth, B., Craw, S., Faltings, B., Maher, M.L., Cox, M.T., Forbus, K., Keane, M., Aamodt, A., Watson, I.: Retrieval, reuse, revision, and retention in case-based reasoning. Knowledge Engineering Review **20**(3) (2005) 215–240 doi: 10.1017/S0269888906000646
52. Rowley, J.: The wisdom hierarchy: representations of the DIKW hierarchy. Journal of Information Science **33**(2) (2009) 163–180 doi: 10.1177/0165551506070706