EPSRC
Polaris House, North Star Avenue, Swindon, Wiltshire,
United Kingdom SN2 1ET
Telephone +44 (0) 1793 444000
Web http://www.epsrc.ac.uk/

Standard
PROPOSAL

Document Status: With Council
EPSRC Reference: EP/R030650/1

Human-like Computing: Call for feasibility studies

Organisation where the Grant would be held

Organisation	University of Cambridge	Research Organisation Reference:	RG93161
Division or Department	Computer Laboratory		

Project Title [up to 150 chars]

How to (re)represent it?

Start Date and Duration

a. Proposed start date	(	b. Duration of the grant (months)	18

Applicants

Role	Name	Organisation	Division or Department	How many hours a week will the investigator work on the project?
Principal Investigator	Dr Mateja Jamnik	University of Cambridge	Computer Laboratory	4.65

Joint Proposals

Complete this section if more than one organisation is submitting a EPSRC proposal form for this project.

Is this part of a joint proposal ?	✔Yes No
Are you the lead RO ?	✔Yes No
Joint reference	U2618202
Total number of proposals	2

Objectives

List the main objectives of the proposed research in order of priority [up to 4000 chars]

In this feasibility project, we are interested to find out:

* What are the cognitive processes that humans use to select a representation of a problem?

* What are the criteria that we can use to assess the effectiveness of a representation?

* And how do we model this computationally?

Our hypothesis is that when humans choose a representation of a problem they use cognitive and formal properties of the problem to make their choice.

This project will uniquely combine the use of empirical cognitive studies for foundational approaches to automating the choice of representation appropriate for the problem, the user and their preferences and expertise. Research contributions fall into four objectives (O1-O4) within two main streams:

Stream 1: computational studies led by Jamnik
Stream 2: cognitive studies led by Cheng

The objectives O1-O4 and the proposed work plan are as follows:

O1. Scoping of representations [Streams 1 & 2]:
Here we will build a corpus of problems and their solutions which use as many different representations as possible. We will use this corpus to identify a comprehensive list of cognitive and formal properties of representations, and then analyse them to devise a taxonomy of properties.

O2. Characterise properties of representations through cognitive studies [Stream 2]:
We will devise coding schemes that will identify cognitive and formal properties from graphical protocols captured from participants' solutions in two series of cognitive experiments. We will analyse them and develop cognitive theories about human choice and use of representations.

O3. Computationally model changes in representations [Stream 1]:
We will design and implement a system within which users can express a problem with a range of representations. We will use the taxonomy of properties from cognitive studies in devising algorithms that rank appropriate recommended representations. The result will be a system that automatically suggests a representation based on the problem and the user.

O4. Evaluate and generalise [Streams 1 & 2]:
After evaluating the utility of our system, we will generalise this work from the domain of mathematics to other domains, as well as investigate how to apply it for societal benefit, for example, in education in the form of an adviser engine in AI tutoring systems that adapt to users.

We expect that the following deliverables will result from our work:

1. Online corpus of problems with multi-representation solutions. (O1)
2. A systematic taxonomy of formal and cognitive properties. (O1)
3. A set of problems for studies of representation change. (O1)
4. Methods for conducting experiments about representation. (O2)
5. Verbal and graphical protocol analysis method for studying representation change. (O2)
6. An account of the role of formal and cognitive properties in representation change. (O2)
7. Algorithms that use a taxonomy of properties to evaluate and rank appropriate representations.(O3)
8. Heuristics encoding user preferences and level of expertise to further rank representations. (O3)
9. A system that automatically suggests a representation based on the problem and the user. (O3)
10. Extensive further project proposals that are general across problem domains. (O4)

Summary

Describe the proposed research in simple terms in a way that could be publicised to a general audience [up to 4000 chars].

Note that this summary will be automatically published on EPSRC's website in the event that a grant is awarded.

AI engines are ubiquitous in our lives: we talk to our mobile phones, ask directions from our sat-navs, and learn new facts and skills with our digital personal assistants. But most of the time, we need to learn how these systems work first: we have to adapt to them as they are not aware of our level of experience, expertise and preferences. AI engines are fast and can deal with a deluge of data much better than us, but they do so in machine-oriented ways, which are often inaccessible and unintelligible to humans.

The aim of this project is to identify and study how humans \textit{represent} information that they want to work with and from which they will obtain new knowledge. Humans have the capability to choose the representation that is just right for them to enable them to solve a new problem, and moreover, if the representation needs to be changed, they can spot this and change it. Unlike humans, machines in general have fixed representations and do not have the understanding of the user. For example, sat-nav systems will only give directions with elementary spatial commands or route planning functions, whereas humans give directions in many forms, for instance in terms of landmarks or other geographic features that are based on shared knowledge.

We want to model in computational systems this inherently human ability to choose or change appropriate representations, and make machines do the same. We want to find out what are the cognitive processes that humans use to select representations, what criteria they use to choose them, and how we can model this ability on machines. Our hypothesis is that when humans choose a representation of a problem, they use cognitive and formal properties of the problem and its representation to make their choice. In this project, we will test this hypothesis by achieving the following goals:

1. Collect a corpus of problems and candidate representations to study and categorise their cognitive and formal properties.

2. Devise coding schemes and conduct cognitive studies to identify cognitive and formal properties that people use in choosing representations. Develop cognitive theories based on these experiments.

3. Design and implement computational algorithms that allow users to choose alternative representations. Build a ranking and recommendation system based on the taxonomy from cognitive studies to suggest appropriate representation given a particular problem and user.

4. Evaluate the utility of the system and generalise the approach to other domains outside of mathematics. Investigate how to apply our cognitive and computational models in education in the form of AI tutors that are adaptable to users.

Our work is novel in that it will address the problem of appropriate representation choice. Moreover, we will build novel cognitive theories and computational models that will allow AI systems to operate in more human-like ways and adapt to the requirements of the problem and the needs of the user. Thus, the potential impact will span numerous domains where systems interact with humans to represent information and use it for extracting new knowledge.

Academic Beneficiaries

Describe who will benefit from the research [up to 4000 chars].

The results of this project will have numerous academic beneficiaries. The artificial intelligence community will benefit from systems that operate in more human-like ways: they can choose and change representations as is most suitable for the user. This will make AI systems better able to understand the human and interact with the human -- which is ultimately one of the main aims of AI systems.

Our algorithms and architectures will benefit computer scientists, for example, the automated reasoning community, as it will provide them with ways to recommend, rank and choose representations that are suitable for the problem, the system or the user. As Polya alluded: representing the problem the right way will make finding its solution trivial.

Our taxonomy of cognitive and formal properties, as well as the coding scheme for assessing them in a representation will benefit cognitive scientists and cognitive psychologists. They will provide new cognitive theories about human preferences in representing and reasoning with information. Ultimately, our repository of classified problems and their solutions, accompanied with cognitive theories and coding schemes for choosing representations will provide a richer and more realistic problems for the study of creativity and insight.

Human-computer interaction (HCI) researchers are interested in interfaces between systems and their users: these have the choice of how to represent the material that is interacted with in accessible way - hence our insights and algorithms about choosing representations should be of great benefit to them.

Similarly, information and data visualisation communities will be able to use our repository and classification of representations and their properties for enriching their repertoire of potential visualisation representations. We envisage a shift in this area from information visualization to knowledge visualization.

Ultimately, AI engines are ever more ubiquitous in our society, and thus any area that uses them and wants to extract new information from given data, will benefit from representing this information and interacting with humans about this information in ways that are tailored to the user, their expertise and the problem. We will take education and intelligent tutoring systems as our initial focus in this project, but the landscape of potential domains is substantial.

Impact Summary

Impact Summary (please refer to the help for guidance on what to consider when completing this section) [up to 4000 chars]

"How to (re)represent it?" is an ambitious project where we want to find out how humans choose and also change representation during problem solving. We want to develop engines that will give machines the same capability. Such machines will accrue many of the benefits that humans obtain from changing representation, including: more effective communication with a human by selecting a representation that is well-suited to their level of familiarity with the target topic; better understanding of a human by identifying and adopting the representation that the human favours; and greater flexibility in adapting to the needs of the user. As such, the project will have economic, societal and knowledge impact.

The impact of this work will be wide as AI systems are becoming ubiquitous. When interacting and exchanging information with humans, developers of such products need to represent information in human understandable ways. We will devise techniques that will help developers of AI systems to build products that choose representations appropriate for their users -- and thus have potential economic benefit to them.

Making machines more human accessible will also benefit the society at large, in particular, as it will potentially bring technical tools to those that are less technically versed.

We have already outlined how our work will contribute to the expansion of knowledge in Academic Beneficiaries. The techniques we will develop are novel and will provide a scientific advance in all areas (in academia and in industry) that must represent information for solving problems, reasoning, etc.

Whilst we exposed education as our initial concrete application target, our contributions will be general and could apply across different domains. They will lead to better understanding of representations and their relation to human expertise and preferences. As intelligent tools make part of our every day life, we will all benefit from having more human-like systems that adapt to us, rather than us adapting to them.

To achieve maximum impact, we will follow a comprehensive dissemination strategy. We will publish our results in all relevant communities (e.g., artificial intelligence, cognitive science, computer science, automated reasoning, diagrams, knowledge representation, human-computer interaction, information visualisation) to achieve wide dissemination. We will demonstrate our work by organising a workshop and preparing tutoring material aimed at academic as well as more specific industrial communities. This will also provide a community building opportunity for people interested in representations and more generally, human-like computing. We will continue to participate at outreach activities to raise the awareness of the importance of human-like computing. We will create a web repository to enable free and public access to our papers, corpus of problems and their solutions, software and tutorials. Finally, we have enlisted an advisory board of experts spanning all areas relevant to this project, and coming from academia and industry - their advice will help us stay focused on relevant problems, influence diverse communities that they lead, and transfer our technology widely.

How to (re)represent it?

Mateja Jamnik
Computer Laboratory, University of Cambridge

Peter Cheng
Department of Informatics, University of Sussex

1 Research Track Record

This interdisciplinary proposal brings together unique expertise from the Universities of Cambridge in logic, diagrammatic and automated reasoning, and from Sussex, in cognitive science, empirical studies and representational systems. The Cambridge team pioneered advances in mechanising reasoning with diagrammatic and heterogeneous representations. The Sussex team are leaders in empirically studying human mental processes and devising representational systems that aid human cognition.

The two Investigators, Dr Mateja Jamnik (PI) and Prof Peter Cheng (Co-I), will be joined by two RAs, one based at Cambridge and one based at Sussex. The project will be supported by the advisory board comprising: Dave Barker-Plummer (Stanford University, USA) who is an expert on computationally modelling the use of heterogeneous representations in tutoring systems; Irvin Katz, Senior Research Director at the Cognitive Accessibility and Technology Sciences Center of ETS (Educational Testing Service, Princeton, USA); and Gem Stapleton (Brighton University, UK) an expert on diagrammatic and logic notations, and empirical testing of their use.

University of Cambridge: Dr Mateja Jamnik (PI) is a Reader in Artificial Intelligence at the University of Cambridge Computer Laboratory. She is also appointed as the Specialist Adviser to the House of Lords Select Committee on AI. Mateja holds a Leverhulme Trust Grant [RPG-2016-082] for an interdisciplinary project with cognitive scientists about *ARD: Accessible Reasoning with Diagrams*. Previously, she held the EPSRC Advanced Research Fellowship (ARF) *Automating Informal Human Mathematical Reasoning* [GR/R76783/01]. Her PhD thesis at the University of Edinburgh broke new ground in automated reasoning [12]. She was invited by CSLI Press to write a book about this work [10].

Mateja explores how people solve problems in mathematics [1, 11] and computationally models this on machines enabling them to reason in a similar way to humans [13]. Her research on diagrammatic reasoning is internationally recognised as pioneering. It investigates the relationship between formal, logical proofs and diagrammatic proofs. In the systems Diamond [12] and Speedith [22] users can interactively prove theorems of mathematics purely by applying operations on diagrams.

The ambition of Mateja's work was recognised by the EPSRC ARF, for which she combined diagrammatic and existing symbolic reasoning techniques in a framework that more closely resembles human reasoning: it enables switching between any representation, at any point in the same proof, but still ends up with a formally correct proof [20, 21].

Mateja is capitalising on these advancements to bring accessible representations and systems to a wide range of stake-holders [19, 17, 18]. The Leverhulme Trust recognises the novelty of this work that is empirically establishing which diagrammatic representations are accessible to humans, and is building reasoning systems based on these findings.

In addition, Mateja was awarded a number of other relevant grants: NSF/1036113, EP/G020523/1, CASE/CNA/07/66, GR/S87676/01. She serves in many research roles, most notably as: a member of the CPHC/BCS Distinguished Dissertation committee; a member of the EPSRC Peer Review College (since 2002); one of the founders of the international interdisciplinary conference series "Diagrams" on the theory and application of diagrams (program chair in 2016, chair in 2010, and on its Steering Committee since 2004).

Mateja co-founded the pioneering women@CL national network for promoting women in computing, for which she won the 2016 Royal Society Athena Award. She is an advocate for bringing science to the general public and frequently participates in outreach engagements (TEDx talk, write for BBC Focus magazine, film for BBC Tomorrow's World).

University of Sussex: Prof Peter Cheng (Co-I) is Professor of Cognitive Science, in the Department of Informatics, and leads the Representational Systems Lab. He was a member of the governing board of the Cognitive Science Society (CSS) and a chair of the society (2009-11). He was PI on six research awards from the EPSRC and ESRC (e.g., RES-328-25-0011, L328253012, GR/K74258; all 3 yr X 2 postdoc) and served as deputy director of the ESRC CREDIT Centre, at the University of Nottingham, which ran from 1992 to 2002. His PhD at the Human Cognition Research Lab at the Open University developed computation models of scientific discovery. He then held a SERC Postdoc Fellowship at Carnegie Mellon University, where he collaborated with Herbert Simon on modelling the role of diagrams in scientific discovery.

Peter studies cognitive science of representational systems; how notations and diagrams encode knowledge and what impacts they have on higher

Case for Support

forms of cognition [6, 4]. He studies the mental processes that underpin our interaction with external representations using verbal protocol analysis, eye-movement recording [14] and graphical protocol analysis [16]. He invents novel representational systems for conceptually challenging domains and information intensive problem solving, ranging from topics such as circuit electricity [5], probability theory [2] and logic [3] through to production planning and scheduling [7]. His empirical studies show that these novel representations substantially improve users comprehension, problem solving and learning [5, 2, 7].

Peter's previous collaborations with computer scientists and computational modellers included projects on: representations for automated scheduling [15]; models of graphical perceptual chunking in memory [9]; cognitive modelling of reasoning with graphs [14], which won the J. H. Ely prize for the best article in volume 45 of the Human Factors Journal.

In addition to his service to the CSS, Peter is a member of the EPSRC Peer Review College and has been a leading figure in the Diagrams research community since co-chairing the first Diagrams conference and founding the series.

Grecia Garcia Garcia will be employed as the RA at Sussex. She is ideally skilled to design, conduct and analyse experiments on representation change in problem solving given her experiences developing innovative methods in empirical cognitive studies [8].

Project Team's Selected Publications

[1] A. Bundy, M. Jamnik, and A. Fugard. What is a proof? *Philosophical Transactions of Royal Society A: The Nature of Mathematical Proof*, 363(1835):2377–2391, 2005.

[2] P. C.-H. Cheng. Probably good diagrams for learning: Representational epistemic re-codification of probability theory. *Topics in Cognitive Science*, 3(3):475–498, 2011.

[3] P. C.-H. Cheng. Graphical notations for syllogisms: How alternative representations impact the accessibility of concepts. *Journal of Visual Languages and Computing*, 25:170–185, 2014.

[4] P. C.-H. Cheng. What constitutes an effective representation? In J. Mateja, U. Yuri, and S. Stephanie, editors, *Diagrammatic Representation and Inference: 9th International Conference, Diagrams 2016*, Lecture Notes in Artificial Intelligence, pages 17–31. Springer, Heidelberg, Germany, 9781 edition, 2016.

[5] P.-H. Cheng. Electrifying diagrams for learning: principles for effective representational systems. *Cognitive Science*, 26(6):685–736, 2002.

[6] P.-H. Cheng, R. Lowe, and M. Scaife. Cognitive science approaches to diagrammatic representations. *Artificial Intelligence Review*, 15(1-2):79–94, 2001.

[7] P. C.-H. Cheng, R. Barone, P. I. Cowling, and S. Ahmadi. Opening the information bottleneck in complex scheduling problems with a novel representation: Stark diagrams. In M. Hegarty, B. Meyer, and N. H. Narayanan, editors, *Diagrammatic representations and inference: Second International Conference, Diagrams 2002*, pages 264–278. Springer-Verlag, Berlin, 2002.

[8] G. Garcia-Garcia and R. Cox. An infra-red finger tracking system used in the assessment and remediation of graph-as-picture misconceptions. In R. Cox and J. S. Diego, editors, *Proceedings of the Workshop on Technology-Enhanced Diagrams Research.*, pages 42–51. CEUR Workshop Proceedings, New York, 2012.

[9] F. Gobet, P. Lane, S. Croker, P.-H. Cheng, G. Jones, I. Oliver, and J. M. Pine. Chunking mechanisms in human learning. *Trends in Cognitive Science*, 5(6):1236–243, 2001.

[10] M. Jamnik. *Mathematical Reasoning with Diagrams: From Intuition to Automation*. CSLI Press, Stanford, CA, 2001.

[11] M. Jamnik and A. Bundy. Psychological validity of schematic proofs. In D. Hutter and W. Stephan, editors, *Mechanizing Mathematical Reasoning: Essays in Honor of Jörg H. Siekmann on the Occasion of His 60th Birthday*, volume 2605 of *Lecture Notes in Artificial Intelligence*, pages 321–341. Springer, 2005.

[12] M. Jamnik, A. Bundy, and I. Green. On automating diagrammatic proofs of arithmetic arguments. *Journal of Logic, Language and Information*, 8(3):297–321, 1999.

[13] M. Jamnik, M. Kerber, and M. Pollet. Automatic learning in proof planning. In *Proceedings of 15th ECAI*, pages 282–286. European Conference on Artificial Intelligence, IOS Press, 2002.

[14] D. Peebles and P. C.-H. Cheng. Modelling the effect of task and graphical representations on response latencies in a graph-reading task. *Human factors*, 45(1):28–45, 2003.

[15] D. Ranson and P.-H. Cheng. Graphical tools for heuristic visualization. In G.Kendall, L.Lei, and M.Pinedo, editors, *Proceedings of the 2nd Multidisciplinary International Conference on Scheduling: Theory and Applications.*, volume II, pages 658–667. New York University, New York, 2005.

[16] R. Roller and P. C.-H. Cheng. Observed strategies in the freehand drawing of complex hierarchical diagrams. In P. Bello, M. Guarini, M. McShane, and B. Scassellati, editors, *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, pages 2020–2025. Cognitive Science Society, Austin, TX, 2014.

[17] Y. Sato, G. Stapleton, M. Jamnik, Z. Shams, and A. Blake. How network-based and set-based visualizations aid consistency checking in ontologies. In *10th International Symposium on Visual Information Communication and Interaction, VINCI-2017*. ACM.

[18] Z. Shams, M. Jamnik, G. Stapleton, and Y. Sato. Reasoning with concept diagrams about antipatterns in ontologies. In *10th International Conference on Intelligent Computer Mathematics - CICM-2017*, LNCS 10383:255–271.

[19] G. Stapleton, M. Jamnik, and A. Shimojima. What makes an effective representation of information: A formal account of observational advantages. *Journal of Logic, Language and Information*, 26(2):143–177, 2017.

[20] M. Urbas and M. Jamnik. Diabelli: A heterogeneous proof system. In *IJCAR*, volume 7364 of *Lecture Notes in Artificial Intelligence*, pages 559–566. Springer, 2012.

[21] M. Urbas and M. Jamnik. A framework for heterogeneous reasoning in formal and informal domains. In *Diagrams*, LNCS 8578:277-292, 2014.

[22] M. Urbas, M. Jamnik, and G. Stapleton. A reasoner for spider diagrams. *Journal of Logic, Language and Information*, 24(4):487–540, 2015.

Case for Support

2 Project Description

In this interdisciplinary project, we want to understand how humans choose representations of problems in sufficient depth and detail so we can build AI engines to do the same. We will investigate human cognitive abilities of finding representations that suitably match problems, and the process by which humans adapt or switch between representations. We will build computational models of this, and automate them in a new generation of adaptive AI systems.

Consider this example. What is the formula for the area of a parallelogram? In how many different ways can you demonstrate it? Here are three (amongst many):

1. You could cut it up into unit squares, and reassemble them on the base line of the parallelogram, but this is very approximate.
2. You could cut off a triangle from one end and move it to the other side, so you get a rectangle.
3. Or you could use calculus and integrate a function. This works, but is like using a sledge hammer to crack a nut.

Which of these are effective representations for the problem? It depends; the first is probably best for young school children; the last for more advanced mathematicians. Which one is best for a machine?

In this project, we are interested to find out:

- What are the cognitive processes that humans use to select a representation of a problem?
- What are the criteria that we can use to assess the effectiveness of a representation?
- And how do we model this computationally?

Our hypothesis is that when humans choose a representation of a problem they use cognitive and formal properties of the problem to make their choice.

Why is this important? AI engines that can choose representations of problems in a similar way will be an essential component of human-like computers. They will give machines a powerful ability to adapt representations so that they are better suited to the particular preferences or abilities of the human user. This is especially important when the machine must give users intelligible explanations about its reasoning (see EPSRC Call Objective CO-1). Human representational choices give us clues of their problem solving approaches. This will aid the construction of the world model (CO-4) that reflects that of a human. Consequently, the machines will be able to not only better adapt to the individual user, but also better interpret instructions or information provided by humans (CO-2). Ultimately, this will lead to machines collaborating with humans in more intuitive ways, and also tutoring humans to develop their creative problem-solving skills.

2.1 Background and Context

We distinguish between cognitive and formal properties of the representation and the domain: they are critical for choosing and changing representation. Cognitive properties are general characteristics of cognitive processes demanded of a particular representation [26] (e.g., problem state space characteristics; applicable state space search methods; attention demands of recognition; inference operators complexity) [3, 4, 5, 6, 7]. Formal properties are general characteristics of the nature of the content of the representation domain (e.g., operation types like associative or commutative, symmetries, coordinate systems, quantity or measurement scales).

From the computational perspective, systems have a fixed representation that users can solve problems with – it is typically symbolic (e.g., logic). As a more accessible alternative, the use of diagrammatic representations has been modelled [10, 22], however, diagrams are again the *only* available representation [24]. Heterogeneous reasoning refers to the human ability to switch between any type of representation. Whilst formal foundations for the design of heterogeneous reasoning systems were laid by Hammer [25], very few attempts have been made to implement such reasoning tools: e.g., the Openbox framework [23] and MixR [21] provide a platform for plugging in a suite of representations, but make no stipulation about what is the appropriate representation for a given problem or user.

2.2 Programme and Methodology

This project will uniquely combine the use of empirical cognitive studies for foundational approaches to automating the choice of representation appropriate for the problem, the user and their preferences and expertise. Research contributions fall into four objectives (O1–O4) within two main streams with the following proposed plan:

Stream 1: computational studies led by Jamnik
Stream 2: cognitive studies led by Cheng

O1: Scope representations [Streams 1 & 2]

For this objective we will develop a taxonomy of formal as well as cognitive properties of problems and their solutions with alternative representations. The examples for our initial set of properties were listed above in Background. Our aim is to devise comprehensive lists of properties and to systematically organise them. We will build an extensive corpus of problems that are solved with alternative representations. Target sources include popular puzzle-problem books, and cognitive science literature about creativity, insight, analogy, etc. We will set up a web-based repository to crowd-source (e.g., Diagrams community) with incentives (prizes) a large set of diverse examples with the most divergent representations. A sample will be selected

Case for Support

from the corpus for analysis, which will: identify and explicate how formal properties are encoded in each representation of every selected problem; assess problem solution with respect to the cognitive properties, using cognitive task analyses.

Deliverables: (1) Online corpus of problems with multi-representation solutions. (2) A systematic taxonomy of formal and cognitive properties. (3) A set of problems for studies of representation change.

O2: Cognitively study properties [Stream 2]

For this objective we will probe the role that cognitive and formal properties play in human changes of representation. Two series experiments will be conducted in which participants are asked to generate diverse solutions to problems (from deliverable 3). Prompts will be designed to encourage diverse solutions. Participants' "work scratchings" will be digitally captured as graphical protocols and with concurrent and retrospective verbal protocols. Coding schemes will be developed to identify formal and cognitive properties from the protocols. The first series of experiments will be naturalistic, whereas the second will test the benefits of training novices in the terminology of the formal and cognitive properties. We predict that such training will facilitate representation change.

Deliverables: (4) Methods for conducting experiments about representation. (5) Verbal and graphical protocol analysis method for studying representation change. (6) An account of the role of formal and cognitive properties in representation change.

O3: Computationally model, implement [Stream 1]

Here, we will design and implement a system that allows users to express a problem with a range of representations. An existing platform, such as MixR or Openproof, will be extended to examples from our problem domain (from deliverable 3). Next, we will model computationally the application of criteria identified in the cognitive studies about effectiveness of representations (deliverable 6) to evaluate and suggest an appropriate representation to the user. Finally, we will build heuristics that reflect human cognitive preferences and level of expertise with regards to properties, and implement a recommender system that automatically selects and suggests an effective representation.

Deliverables: (7) Algorithms that use a taxonomy of properties to evaluate and rank appropriate representations. (8) Heuristics encoding user preferences and level of expertise to further rank representations. (9) A system that automatically suggests a representation based on the problem and the user.

O4: Evaluate and generalise [Streams 1 & 2]

For this objective we will evaluate the accessibility of and human performance with our system (from deliverable 9). We will explore the generalisation avenues from this feasibility study to develop larger research project proposals that will look across different domains not just mathematics. We envisage exploiting our model as an adviser engine in AI tutoring systems that adapt to users. Our future vision for this work is that we will illuminate issues of accessibility, generality of representations, the properties of representations, cognitive criteria for choices of representation, and thus ultimately make AI systems more human-like.

Deliverable: (10) Extensive further project proposals that are general across problem domains.

Success Criteria and Risk Management

Our key criteria for measuring success include: *intervention* – does improving users' awareness of the properties improve how they choose representation?; and *automation* – can a machine choose a representation that humans find appropriate? Our main risks and alternative plans for them include: (1) *People may not be capable of re-representing, perhaps it is only a skill of experts.* This is a matter of degree since expertise depends on learning. If needed, we will examine expert-novice differences and refocus our empirical studies on the effects of training about the properties. (2) *Representation change processes may be beyond any current reasoning automation platform.* We built MixR which we will use as a starting point and extend it to allow automated representation choice. If needed, we will use our expertise and build a brand new platform.

2.3 National Importance

This project contributes to the recognised priority of creating Future Intelligent Technologies that are adaptive and can make decisions without human control. Such systems must be able to interpret and make sense of information about the world around them. Our work will provide systems with means to comprehend and be comprehensible to humans in order to successfully communicate and collaborate with them. We will translate two decades of cognitive science discoveries about human representation use into a form suitable for incorporating into Artificial Intelligence, Robotics and Automation.

Additional References

[23] D. Barker-Plummer, J. Etchemendy, A. Liu, M. Murray, and N. Swoboda. Openproof: A flexible framework for heterogeneous reasoning. In *Diagrams*, volume 5223 of *LNAI*, pages 347–349. Springer, 2008.

[24] B. Chandrasekaran, J. Glasgow, and N. Narayanan, editors. *Diagrammatic Reasoning: Cognitive and Computational Perspectives*. AAAI Press/MIT Press, Cambridge, MA, 1995.

[25] E. Hammer. Reasoning with sentences and diagrams. *Notre Dame J. Formal Logic*, 35(1):73–97, 1994.

[26] A. Miyake and P. Shah. *Models of working memory.* C.U.P., Cambridge, 1999.

Case for Support

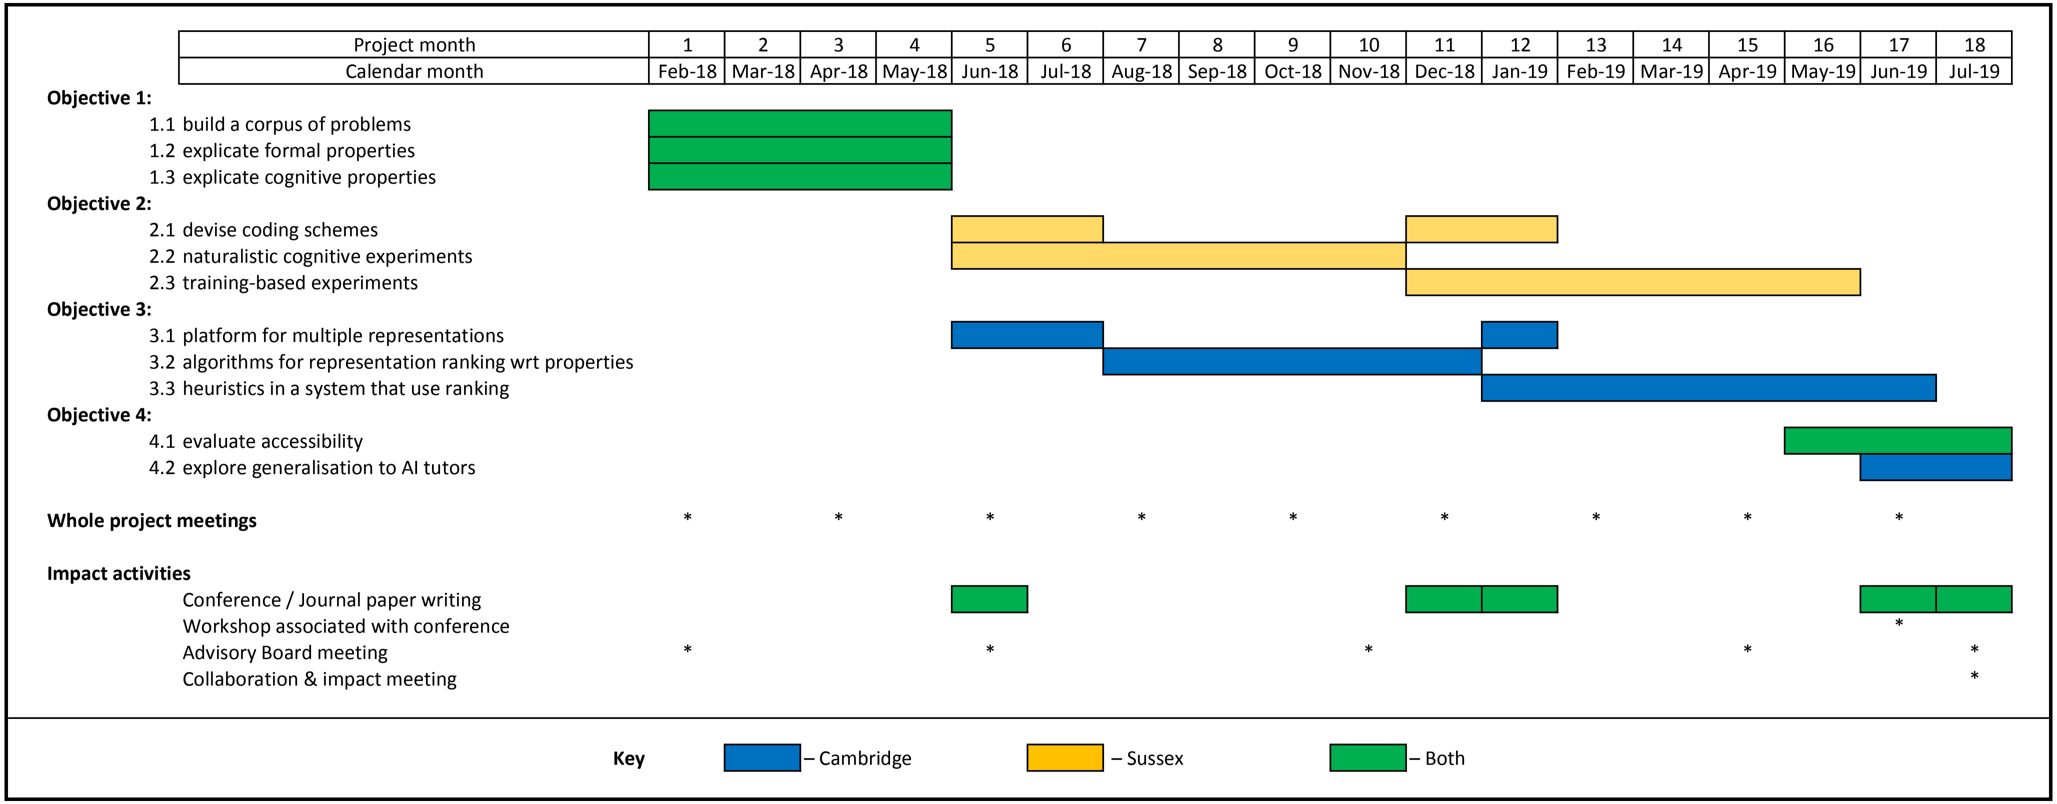
Project month
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
Calendar month
Feb-18
Mar-18
Apr-18
May-18
Jun-18
Jul-18
Aug-18
Sep-18
Oct-18
Nov-18
Dec-18
Jan-19
Feb-19
Mar-19
Apr-19
May-19
Jun-19
Jul-19
Objective 1:
1.1 build a corpus of problems
1.2 explicate formal properties
1.3 explicate cognitive properties
Objective 2:
2.1 devise coding schemes
2.2 naturalistic cognitive experiments
2.3 training-based experiments
Objective 3:
3.1 platform for multiple representations
3.2 algorithms for representation ranking wrt properties
3.3 heuristics in a system that use ranking
Objective 4:
4.1 evaluate accessibility
4.2 explore generalisation to AI tutors
Whole project meetings
Impact activities
Conference / Journal paper writing
Workshop associated with conference
Advisory Board meeting
Collaboration & impact meeting
Key
– Cambridge
– Sussex
– Both

Workplan