

# Cognitive Robotics

Eric Chown<sup>1</sup> and Wai K. Yeap<sup>2(✉)</sup>

<sup>1</sup> Bowdoin College, Brunswick, ME, USA  
echown@bowdoin.edu

<sup>2</sup> Auckland University of Technology, Auckland, New Zealand  
wai.yeap@aut.ac.nz

**Abstract.** Cognitive robotics presents unique research challenges as it straddles the boundaries between two very different fields. Its name promises robots that exhibit cognitive behaviour and suggests that the research would be beneficial and interesting to researchers in both robotics and cognitive science. However, this is not generally the case. Many roboticists are disinterested in the developments of cognitive robots unless these robots are proven to be superior in terms of speed, efficiency and accuracy. Similarly, cognitive scientists are disinterested in the work of roboticists unless the robot's behavior is fit to some known empirical data and the models implemented are cognitively plausible. These different requirements, each from their own field, are often at odds, and has hampered the development of cognitive robotics, causing researchers to use robots only as a platform for simulating cognitive ideas or to use cognitive science as a weak source of ideas for robot mapping. In this article, we argue that a better synthesis of ideas from both fields must be encouraged and that cognitive robotics should move beyond its artificial limitations and in the process better serve robotics and cognitive science.

**Keywords:** Cognitive mapping · Robotics · Cognitive science · Cognitive robotics

## 1 Introduction

Academia is structured in such a way as to make interdisciplinary research difficult. Every field and subfield has unique methodology, terminology, and systems of evaluation. When stepping from one discipline to another one must quickly learn all of the differences between fields or success will be extremely difficult. In recent years computational scientists have bucked this general trend because they have brought new tools to nearly every field from the Humanities to the Sciences. Even so, within other disciplines there is always strong resistance (e.g. see Drucker 2012; McPherson 2012). Often these new methods from other fields are seen as a threat to existing methodologies and to old ways of doing things. By contrast, cognitive science has long since embraced computational methods. Computer models of psychological theories have a long history going back to Newell, Shaw and Simon's work on GPS (Newell et al. 1959). However, as with other fields, the intersection of computational methodologies with cognitive science is not always smooth. In the case of cognitive

robotics it is not so much a resistance to computation, as the fact that cognitive science and robotics as fields have such different methods of evaluation. Here we will broadly characterize the two fields while at the same time recognizing that the truth is far more nuanced.

Broadly speaking the goal of Robotics is to build ever more capable robots. This means robots that can do more, are smarter, and are more efficient. As a subfield of computer science this is reflective of computer science as a whole. In its memo on tenure the Computer Research Association had this to say, “When one discovers a fact about nature, it is a contribution per se, no matter how small. Since anyone can create something new [in a synthetic field], that alone does not establish a contribution. Rather, one must show that the creation is better.” (Patterson et al. 1999) The memo goes on to say that contributions will lead to “better results.” This is perfectly sensible for computer science as a discipline. So in publishing in computer science, and by extension robotics, it is generally necessary to prove that one’s work is “better” than what came before. In robotics, meanwhile, “better” can often be directly quantified. It might mean a more accurate map, a faster path, or more efficient computation. It should be no surprise that these are standard means of evaluating robot models.

The situation is quite different in cognitive science. The field has many goals, but most of them revolve around understanding human cognition. In many cases such understanding comes through building models. The goal of such models is not to be “better” in the same sense as computer science, but to be a more accurate reflection of human cognition and therefore be a potentially more useful tool in understanding how cognition works. When models are judged by characteristics of speed, accuracy and efficiency, it is not to be certain that such quantities are maximized, but rather that they closely reflect data on their human counterparts.

Mapping provides a simple and concrete example of how the two ways of working are at odds and throughout this article, we will use this example when discussing the problems cognitive robotics face. In robotics, the mapping literature is currently dominated by SLAM (Simultaneous Localization and Mapping) methods (e.g. see Thrun 2008). In SLAM, a robot’s goal is to navigate an environment and simultaneously build an accurate map of the environment. SLAM models strive to build as precise a map of their environment as is possible. The better the map the better the robot will be able to use it later to navigate. Meanwhile human navigation is a completely different story. While humans are capable of amazing navigational feats, it has been repeatedly shown that their internal maps are actually quite distorted and sketchy. Roboticians interested in cognition might find themselves with a typical robot setup involving a small wheeled-robot with a laser rangefinder. They could then use such a robot to build human-like maps of its environment. From the perspective of mainstream robotics such maps are going to look poor compared to the best SLAM methods. From the perspective of cognitive science such maps may not be interesting simply because they were created using inputs from lasers; the latter is so unhuman.

Due to these differences and the problems inherent in evaluating such models, the development of cognitive robots has been hampered. Cognitive robotics researchers often end up with models that interest cognitive science or robotics but not both. Consequently, cognitive researchers have shown little interest in robotics beyond as a platform

for implementing and testing their ideas and roboticists have shown little interest in cognitive science beyond borrowing some simple ideas to improve upon their models. For example, the idea that SLAM models might have anything to teach cognitive scientists is seen as absurd within much of cognitive science. Indeed we are unaware of any mainstream models of human navigation or cognitive mapping that have been significantly influenced by SLAM methods. Again, on the face of it this seems reasonable, after all robots are very different than humans and most SLAM researchers would not claim that their methods are reflective of human navigation. We believe, however, that such a view is shortsighted. To get a glimpse of why this is the case one need look no further than the work on animal navigation. Cognitive research has been greatly impacted by work on many different species, from rats whose brains share a significant amount of structure with humans to ants whose brains have little in common with human brains. The common element of all such species is that they all face challenges in navigation that they have evolved to overcome. By studying a wide range of species it is possible to start finding principles that are common to all of them. Such principles may be implemented very differently for a given animal, but if the principle is powerful enough, it is likely that many different systems will have “discovered” it through the process of evolution. The very fact that diverse animals are using the same principles is evidence of the power and utility of such principles.

We are proposing that by viewing robots and humans as different species solving the same problem (an idea first proposed in Yeap 2011b), there is much to be learned from each other. From this perspective, cognitive robotics is not restricted to the use of robots for testing cognitive ideas but also to find solutions to problems that baffled cognitive scientists. The latter leads to discovering new navigational principles and new algorithms that would benefit research in robotics and provide new insights into spatial cognition. Different models (i.e. species) can then be built to evaluate the applicability of the principles identified. Indeed, the different capabilities of robots – having wheels, sonar, lasers, and others – provides a test of the generality of the methods and principles that are being proposed.

In the rest of this article we discuss the nature of interactions between cognitive science and robotics through cognitive robotics, both in the successes and the challenges that have emerged over the years but also in suggesting that much more is possible, especially with regard to the contributions that both sides have to gain from each other.

## 2 What Do Cognitive Models Offer Robotics?

There are many reasons why robotics researchers might not be interested in cognitive models of navigation. For instance, it would be easy to conclude that the human brain is just too complex and little understood. While cognitive models offer significant insights, they often lack the detail necessary for an implementation, and in particular, the details most lacking are normally computational. Further, human navigation makes use of an unparalleled object recognition system that dominates even the most sophisticated machine vision system. Meanwhile, even when navigating humans are constantly doing other things and have many other cognitive processes that impact navigation

performance. These range from emotional factors to things that take attention away from navigation. From the point of view of robotics it would appear that building systems that do not have such distractions would ultimately lead to better performance. It therefore stands to reason that doing things the way humans do them may not be the best path for robots. Further, the history of Artificial Intelligence (AI) suggests that eventually robots will be better than humans at navigation, just as with chess and with many other problems once thought to be “hard.” And just as with chess and other problems, the solutions that AI finds may have little or nothing to do with the way that humans perform them. So what does cognition have to offer robotics? Why not simply pursue a path of ever more refined mathematical models?

In the early days, AI researchers were eager to learn from nature for two reasons. First, nature provides a rich source of ideas and second, nature’s solutions are both interesting and tested by the need to survive. For mapping, roboticists have already borrowed many ideas from cognitive mapping albeit mainly at the structural level. For example, topological maps, inspired by human “route maps” have long been a staple of robotics (e.g. see Thrun and Bücken 1996). “Gateways”, first developed as part of a theory of human cognitive mapping (Chown et al. 1995) found their way into numerous robot systems (e.g. Beeson et al. 2010). In both of these cases roboticists took an initial idea from cognitive models and used it as a starting point to end up with something new.

To the extent that cognitive ideas have had a positive impact on robotics, it is because these ideas can easily be translated onto a robot whatever its capabilities. For example, a topological map consists of a network of “landmarks.” Such representations are powerful and useful regardless of whether the landmarks were learned by using vision, sonar or lasers. The navigational principles that these representations encapsulate are powerful enough that they can be implemented in an almost unlimited number of ways. The Gateway notion has a similar history. In humans Gateways occur where there is a visual occlusion followed by an opening. In other words Gateways are places where new information becomes available. A robot does not need the powerful human visual system to take advantage of the principle that locations in the environment where new information becomes available are important. One of the first uses of gateways (Kortenkamp and Weymouth 1992), for example, used sonar, a sensor that is notoriously noisy. For a robot moving down a corridor, however, sonar makes it extremely easy to identify gateways. As the robot is moving past walls it should get relatively constant, if noisy, readings. However, when a new corridor opens up, or there is an open door, the readings will jump dramatically, far beyond the magnitude of normal noise. Thus the principle that a sudden large environmental change should drive new representations is easily exploited by robots with a wide variety of sensory capabilities. This principle is powerful enough that it can be used to organize a map even for a perceptually weak robot.

The landmark example is instructive for the larger point that we are making. Topological maps have been a staple of cognitive theories of navigation going back to Piaget (Piaget and Inhelder 1967). These theories then made their way into robotics. Once robotics got ahold of them there were essentially two camps working on topological models. One consisted of cognitive roboticists who tended to slavishly try to mimic the exact details of human navigation. The second camp consisted of roboticists who wanted to make their robots navigate as effectively as possible. Both of these camps have narrow

goals that naturally limited the scope of their work. The robot camp, for example rarely considers the full complexity of landmarks. Since their only goal is to build robots that navigate as effectively as possible they focus squarely on using perception to identify landmarks. A researcher, for example, using lasers for input, will work to engineer solutions optimized for lasers. A similar story will be true for robots with cameras or sonar. In each case landmarks are taken to be any object in view that has a unique perceptual signature, and consequently landmarks are perceived almost everywhere by robots because there is almost always something in any given view that stands out.

By contrast, humans remember a much smaller number of objects as landmarks, apparently using a much different, more global, criteria, and consequently each landmark is more important and more memorable. It is much more difficult to implement this on a robot because it is still not well understood exactly how humans select landmarks, and of course the human vision system is a crucial part of the process.

In this way the two points of view represent a kind of continuum of strategies that a cognitive being might use in navigation. Different points along this continuum might have different strengths and weaknesses with regard to different environments. For example, some environments, such as a desert or prairie, are known to be landmark poor for humans. A robot might have better success in creating landmarks on the fly in such a place. On the other hand, with so many landmarks in their maps, robots may be more susceptible to becoming confused when changes in the environment occur as happens so often in so many environments. Robotic solutions along these lines are often accused of being brittle. This is one of the oldest criticisms of AI models. For example, consider the application of probabilistic solutions in robot soccer competitions such as RoboCup (Chown and Lagoudakis 2015). While these algorithms perform well in it, even small changes to the domain can cause them to completely break. This would appear to be in stark contrast to human intelligence. Humans appear to effortlessly adapt to even large changes in domain rules. For example, not only can a group of five year olds instantly adapt to the outdoor conditions of a new soccer field, they can even create a field on the fly using sticks, trees, and bushes and then localize beautifully on the ad hoc pitch. It would be more than a stretch to suggest that such children are creating highly precise internal maps to accomplish such feats. Indeed it may well be the case that the lack of such precise models is a key to such adaptability.

Computer science and AI have championed the idea that problems can be viewed as search. It is certainly possible to view “navigation” as such a problem. What we are seeing is that a very small part of that space is being explored right now. There are researchers who are exploring the space that most closely resembles human navigation and there are researchers who are looking to optimize what robots with specific sets of capabilities can do. In search terms both groups are essentially looking for local maxima using different criteria for defining their maxima. The “better” required in robotics is a kind of hill climbing as is the drive for ever more realistic models in cognitive science. Cognitive robotics affords the chance to pursue more global strategies. Strategies that might provide a deeper understanding of the space and could ultimately lead to solutions that combine the flexibility of human navigation with the precision of robot navigation.

### 3 What Does Robotics Offer Cognition?

Traditionally the major, if not the only, reason for a researcher interested in cognition to use robots has been to test ideas. The benefit of using robots is concreteness. If a model cannot be implemented on a robot, or if it simply does not work when implemented, then these are strong evidence of a model's deficiencies. In turn such deficiencies are indicative of places where existing theories need to be updated or discarded. Ironically then, success, in terms of lessons learned through robotics, typically has come through failures in implementation. However, when a model is successfully implemented, it is easy to dismiss that success. Some models are dismissed out of hand on the grounds of differences in hardware, with others it is because the model used is too general. In fact it is difficult to prove anything by building a model. Generally the best that can be done is to use the model to make novel predictions that can later be checked experimentally (e.g. by testing human subjects). Despite this apparently pessimistic view where knowledge only comes through failure, the act of making theories concrete via the use of robot is invaluable.

However, while the importance of making theories concrete is hard to overstate, an important lesson learned in developing early AI models of cognition is that the theory must be formulated and tested at the appropriate level. For example, in the early history of AI and cognitive science, many models of cognition were developed at a high level of thinking (Langley 2007). Conscious thinking at that time was deemed to be the most interesting and important part of cognition. Perhaps the greatest lesson of robotics and related fields such as machine vision, is that they provided stark and unequivocal evidence of just how difficult and important perception is. This is evidenced by the famous story of how in the early days of AI Marvin Minsky assigned some of his students to solve the problem of using a camera to identify objects as a good "summer project." More than 50 years later and it is still nowhere close to being solved. Robotic provides an excellent platform for developing concrete ideas about the perceptual process that cognitive theories often lack but require. It is surprising then that this aspect of the advancements from robotics has largely been ignored. Research on SLAM provides an instructive example.

Many cognitive theories of navigation have converged around the idea that people learn a series of views as one form of navigation (Yeap 1998; Chown et al. 1995; Franz et al. 1998). All of these theories, even the ones implemented on robots, have not fully tested the implications of this idea. Meanwhile SLAM models have been implemented and tested on enormous scale at levels far beyond what the cognitively inspired models have managed. What SLAM researchers have found is that the process of building a global map of an environment based on integrating successive views of the environment has a fundamental problem – the accumulation of error. As has long been known in robotics, when an agent moves through an environment little errors tend to accumulate. Over time and space these little errors turn into large errors. Imagine, for example, that you want to head north, but your heading is off by a small amount. The further you go on that heading the more you stray from true north. Over large distances you will end up a long way away from your goal. SLAM researchers have had to learn to cope with this problem and have come up with techniques that allow them to correct for the errors that cannot help but occur when building a global map.

If humans do indeed build maps out of successive views, than this process is necessarily going to run into the same problems of accumulating errors. It is also possible, and may even be likely, that the solutions being found by SLAM researchers have also been “found” by evolution. At the very least SLAM is leading to a more thorough understanding of the problems and issues inherent in such processes. The question is what if any of this knowledge has found its way back into psychological models and testing? The answer, as far as we can tell, is “little or none.” Cognitive scientists by and large are not interested in SLAM because it isn’t a cognitive model. Worse, from their point of view, it is often implemented on robots with very different capabilities than humans. A cognitive scientist might say that SLAM isn’t relevant because humans can resolve their errors through the use of their superior vision systems or in some other way that is different than SLAM. It is possible that this is even true but what would such a model be? Cognitive science has not proposed any as of yet. It is also possible that the general principles used by SLAM systems may be similar, or even the same, as the principles used in human cognition. If this is the case then it is clearly a good idea to identify those principles. A different cognitive scientist might say that it doesn’t matter anyway since human cognitive maps are known to be distorted and sketchy. Indeed one of the co-authors of this article has leveled both of these criticisms at SLAM in the past. Saying a cognitive map is distorted, however, does not mean that it can be distorted without limit. Indeed this suggests an interesting line of research on the nature of distortions in internal maps and just how distorted they can become while still being functional. More to the point these are things that are not known. There is still a great deal to be learned about the nature of cognitive maps, about the nature of dealing with errors in maps, and about how navigation works in general. Robotics provides a tool for exploring these questions and the answers that roboticists are finding may provide critical information for better understanding how humans navigate. Cognitive roboticists, and by extension cognitive scientists, would be wise to pay more attention to these explorations and to use them as springboards for their own work.

#### **4 Cognitive Robotics – on the Edge of Discovering New Ideas**

Early work on cognitive robotics, rightly so, has emphasized on the use of robots as test beds for evaluating cognitive theories. However, the fact that robots are so different from humans has meant that cognitive scientists are unlikely to view the work of cognitive roboticists as making a significant contribution to their understanding of how the mind works. Similarly, the fact that cognitive ideas are so semantically laden and so specific to humans has meant that roboticists are unlikely to view cognitive robotics as practically significant. What we are proposing in this article is that cognitive robotics should pay more attention to the questions cognitive scientists raise and attempt to find answers to those questions using robots often borrowing developments in robotics to do so. Initially, these questions should center on perceptual problems since this is where cognitive theories tend to be weakest and where robotics must necessarily find solutions. In this section, we provide two examples of exploring such questions.



## 4.1 Contributions to Cognitive Science

While cognitive scientists in general and psychologists in particular have made significant discoveries concerning cognitive mapping, their theorizing often lacks important computational details concerning the underlying process itself. Such omissions are not a matter of mere details since without those details the correctness of any such theory is impossible to verify. This is because any representation inferred from behavioral results could be challenged with an alternative representation. For example, even the most fundamental idea of cognitive mapping – that humans and animals compute a map of their environment, first proposed by Tolman (1948) and later given significant support from two prominent pieces of work: Lynch (1960) and O’Keefe and Nadel (1978) – is controversial precisely because proponents of the theory have failed to demonstrate exactly the kind of map computed and how it is learned. Consequently, by providing alternative explanations to account for the behavior observed, many have challenged the very idea that a map is computed at all. For example, for rats searching for food next in a radial arm maze, Brown (1992) argue that they could have considered only which alley not visited and thus would not need to use a spatial map. In the water maze problem, Benhamou (1996) argue that rats could use some orientation mechanisms and not a spatial map to locate the platform in the water maze (for further discussions, see Yeap 2011b). Even for humans, the idea that we compute a map in the head has been challenged recently. Wang and Spelke (2002) argue that humans maintain only a transient egocentric map that allows awareness of their immediate surroundings but not an enduring non-egocentric map. The latter is too empowering (Yeap 2014) and based on their observation that other lower animals do not compute such a map and their belief that all animals navigation abilities should build on a common set of mechanisms, they conclude that no such map is computed at the perceptual level.

This gap in cognitive science research is what cognitive roboticists could fill via experimentations using robots. To do so, cognitive roboticists need to address how the key ideas/representations identified in cognitive science are physically realized and in ways that match their characterization by cognitive scientists. For example, cognitive roboticists, unlike traditional roboticists, must not only show how a map of the environment is computed, but the map must also bear many of the characteristics of a cognitive map. For example, one distinctive characteristic of a cognitive map is that it is fragmented and inexact. Successfully implementing such a process on a mobile robot, even with sensors that differ from cognitive agents, would provide insights into the nature of the process. Recently, Yeap (2011a) describes one such process implemented on a mobile robot equipped with a laser sensor and an odometer (see also Yeap et al. 2011). Unlike SLAM, Yeap’s process eschews error corrections, continuous updating, and continuous self-localization. Instead it takes a snapshot of the environment as a kind of map of a local environment that it is about to explore, and as it moves out of its current bounded space, it takes a new snapshot corresponding to the next local environment that it will explore and so on. What is computed as it explores is a trace of the individual local maps. When moving in each of the local areas, objects within the area are tracked in each subsequent view by the moving robot and these tracked objects enable the robot to recover its pose (its position and orientation in space) in these maps, thereby allowing



the robot to generate a global metric map. This metric map ends up being incomplete and inexact, but it is not as distorted as a similar map generated via integrating successive views. The map produced is accurate enough to allow the robot to orient itself in the environment.

While Yeap's process is not an exact analog of how human and/or animal cognitive mapping works, it nonetheless bears many interesting commonalities with human cognitive mapping, especially when compared with the SLAM approach. For example, the central idea is that tracking objects as the robot moves can compensate for the normal errors that might build up during the mapping process. The objects provide a natural source of error correction as the robot moves. Global correctness is not necessarily important. This is as opposed to the SLAM-based approach whereby a robot has to continuously localize and correct its position in the map. In SLAM, robots are constantly trying to place themselves in the correct position in the global map, whereas in Yeap's approach the robot is merely trying to solve the simpler problem of determining where it is relative to nearby landmarks. As already noted, human navigation makes use of an unparalleled object recognition system and thus it would be natural that any process proposed for human mapping would take maximum advantage of such a system. Imagine our hypothetical traveller heading to the north. As they head north they will naturally accumulate error. But then if they see a known landmark along the way the landmark will naturally and automatically correct the errors that they have accumulated. Such a traveller need not have a precise map in their head, they need only have one good enough to get to the next landmark. As has been noted before in robotics, "the world is its own best model" (Brooks 1991). Given this, it is hard to imagine that the humans would compute an exact and complete map, and indeed it is well known that they do not.

## 4.2 Contributions to Robotics

Within its own space, roboticists are proud of finding a solution to the principle problem that they have identified – namely how to correct the sensor errors and produce a correct map while simultaneously exploring the environment – and rightly so. Their confidence in this approach has led some to predict that the future key challenges lie in developing ever larger, more persuasive demonstrations of the approach, such as mapping a city, or massive structures such as the Barrier Reef or the surface of Mars (Bailey and Durrant-Whyte 2006; Durrant-Whyte and Bailey 2006). Not surprisingly then, these algorithms have been extended for handling dynamic environments (e.g. Fox et al. 1999; Hahnel et al. 2003), for creating maps in large outdoor environments (Thrun and Montemerlo 2006; Folkesson and Christensen 2007), for creating 3D maps (Nüchter et al. 2007; Pathak et al. 2010), for creating sub-maps as places in a topological map (Konolige et al. 2011; Ranganathan and Dellaert 2011), and for use with vision (Ho and Newman 2007; Schleicher et al. 2010). More recently SLAM-based approaches have become a popular choice for use with drones.

Despite all of this work and all of these successes with SLAM, and despite the fact that the basic processes involved in SLAM are similar to what cognitive scientists are interested in with humans, cognitive science has not paid attention to the results of SLAM research, even though some of these results indicate problems that must be resolved by cognitive models going forward. Humans may not resolve the accumulation

of error in the same way that SLAM does, for example, but SLAM research has very effectively shown that it must be resolved in some way, it cannot simply be ignored. Conversely, nature must have discovered an alternative approach to the one being studied and adopted by roboticists because there are so many examples of successful animal navigators. As discussed above, one such alternative, as illustrated in Yeap (2011a), could provide an alternative paradigm for robot mapping. As we have already witnessed in the development of SLAM-based approaches, if the robotics community take interests in such alternative approaches, they will help accelerate the development and understanding of the alternative paradigms and thus more effectively help to search the space for solving the larger problem of general purpose navigation, both for humans and robotics.

## 5 Concluding Remarks

When two different fields intersect in a new way it is natural for the early work at the intersection to run into the problems inherent in attempting to please two different masters with two distinct sets of needs. This has certainly been the case with cognitive robotics. Ultimately the different standards of evaluations stemming from the two fields can be restrictive, stifling and even self-defeating. In this article we have proposed that cognitive robotics should move in new directions aimed less at slavishly modeling human navigation and instead focus more on processes and principles. In the end we expect that all of the fields involved will gain. Cognitive science can benefit from the work done in robotics in exploring the problems of situating real agents in the real world. This work continues to uncover new problems and alternative solutions to such problems. Meanwhile, robotics can benefit from ideas stemming from systems that are far more general and flexible than any produced by man to date.

## References

- Bailey, T., Durrant-Whyte, H.: Simultaneous localization and mapping (SLAM): Part II. *IEEE Robot. Autom. Mag.* **13**(3), 108–117 (2006)
- Beeson, P., Modayil, J., Kuipers, B.: Factoring the mapping problem: Mobile robot map-building in the hybrid spatial semantic hierarchy. *Int. J. Robot. Res.* **29**(4), 428–459 (2010)
- Benhamou, S.: No evidence for cognitive mapping in rats. *Anim. Behav.* **52**, 201–212 (1996)
- Brown, M.: Does a cognitive map guide choices in the radial-arm maze? *J. Exp. Psychol. Anim. Behav. Process.* **18**, 56–66 (1992)
- Brooks, R.A.: Intelligence without representation. *Artif. Intell.* **47**, 139–159 (1991)
- Chown, E., Kaplan, S., Kortenkamp, D.: Prototypes, location and associative networks (PLAN): towards a unified theory of cognitive mapping. *Cog. Sci.* **19**, 1–52 (1995)
- Chown, E., Lagoudakis, M.: The standard platform league. In: Bianchi, R., Akin, H., Ramamoorthy, S., Sugiura, K. (eds.) *RoboCup 2014. LNCS*, vol. 8992, pp. 636–648. Springer, Heidelberg (2015)
- Durrant-Whyte, H., Bailey, T.: Simultaneous localization and mapping (SLAM): Part I. *IEEE Robot. Autom. Mag.* **13**(2), 99–110 (2006)
- Drucker, J.: Humanistic theory and digital scholarship. In: Gold, M.K. (ed.) *Debates in the Digital Humanities*, pp. 85–95. University of Minnesota Press, Minneapolis (2012)

- Folkesson, J., Christensen, H.I.: Graphical SLAM for outdoor applications. *J. Field Robot.* **24**(1/2), 51–70 (2007)
- Fox, D., Burgard, W., Thrun, S.: Markov localization for mobile robots in dynamic environments. *J. Artif. Intell. Res.* **11**, 391–427 (1999)
- Franz, M.O., Schölkopf, B., Mallot, H.A., Bühlhoff, H.H.: Learning view graphs for robot navigation. *Auton. Robots* **5**, 111–125 (1998)
- Hahnel, D., Burgard, W., Fox, D., Thrun, S.: An efficient FastSLAM algorithm for generating cyclic maps of large-scale environments from raw laser range measurements. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, pp. 206–211 (2003)
- Ho, K.L., Newman, P.: Detecting loop closure with scene sequences. *Int. J. Comput. Vision* **74**(3), 261–286 (2007)
- Konolige, K., Marder-Eppstein, E., Marthi, B.: Navigation in hybrid metric-topological maps. In: *Proceedings of IEEE International Conference on Robotics and Automation (ICRA 2011)*, pp. 3041–3047 (2011)
- Kortenkamp, D., Weymouth, T.: Using gateways to build a route map. In: *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*, pp. 2209–2214 (1992)
- Langley, P.: Artificial intelligence and cognitive systems. In: Cohen, P. (ed.) *AI: The First Hundred Years*. AAAI Press, Menlo Park (2007)
- Lynch, K.: *The Image of the City*. MIT Press, Cambridge (1960)
- McPherson, T.: Why are the Digital Humanities so white? Or Thinking the histories of race and computation. In: Gold, M.K. (ed.) *Debates in the Digital Humanities*, pp. 139–160. University of Minnesota Press, Minneapolis (2012)
- Newell, A., Shaw, J.C., Simon, H.A.: Report on a general problem solving program. In: *Proceedings of the International Conference on Information Processing*, pp. 256–264 (1959)
- Nüchter, A., Lingemann, K., Hertzberg, J., Surmann, H.: 6D SLAM-3D Mapping outdoor environments. *J. Field Robot.* **24**(8–9), 699–722 (2007)
- Pathak, K., Birk, A., Vaskevicius, N., Poppinga, J.: Fast registration based on noisy planes with unknown correspondences for 3D mapping. *IEEE Trans. Rob.* **26**(3), 424–441 (2010)
- Ranganathan, A., Dellaert, F.: Online probabilistic topological mapping. *Int. J. Robot. Res.* **30**(6), 755–771 (2011)
- O’Keefe, J., Nadel, L.: *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford (1978)
- Patterson, D., Snyder, L., Ullman, J.: Best practices memo: evaluating computer scientists and engineers for promotion and tenure. Computing Research Association, Special Insert (1999)
- Piaget, J., Inhelder, B.: *The Child’s Conception of Space*. W.W. Norton, New York (1967)
- Schleicher, D., Bergasa, L.M., Ocana, M., Barea, R., Lopez, E.: Real-time hierarchical stereo visual SLAM in large-scale environments. *Robot. Auton. Sys.* **58**(8), 991–1002 (2010)
- Tolman, E.C.: Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948)
- Thrun, S.: Simultaneous localization and mapping. In: Jefferies, M.E., Yeap, W.K. (eds.) *Robotics and Cognitive Approaches to Spatial Mapping*. STAR, vol. 38, pp. 13–41. Springer, Heidelberg (2008)
- Thrun, S., Bücken, A.: Integrating grid-based and topological maps for mobile robot navigation. In: *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, OR, pp. 944–950 (1996)
- Thrun, S., Montemarlo, M.: The GraphSLAM algorithm with applications to large-scale mapping of urban structures. *Int. J. Robot. Res.* **25**(5–6), 403–429 (2006)
- Wang, R.F., Spelke, E.S.: Human spatial representation: Insights from animals. *Trends Cogn. Sci.* **6**(9), 376–382 (2002)

- Yeap, W.K.: Towards a computational theory of cognitive maps. *Artif. Intell.* **34**, 297–360 (1998)
- Yeap, W.K.: A computational theory of human perceptual mapping. In: *Proceedings of the Cognitive Science Conference, Boston, USA*, pp. 429–434 (2011a)
- Yeap, W.K.: How Albot<sub>0</sub> finds its way home: a novel approach to cognitive mapping using robots. *Top. Cogn. Sci.* **3**(4), 707–721 (2011b)
- Yeap, W., Hossain, M., Brunner, T.: On the implementation of a theory of perceptual mapping. In: Wang, D., Reynolds, M. (eds.) *AI 2011. LNCS*, vol. 7106, pp. 739–748. Springer, Heidelberg (2011)
- Yeap, W.: On egocentric and allocentric maps. In: Freksa, C., Nebel, B., Hegarty, M., Barkowsky, T. (eds.) *Spatial Cognition 2014. LNCS*, vol. 8684, pp. 62–75. Springer, Heidelberg (2014)