Object Permanence in Human Development and Robotics - A Survey

Tooba Ahsen, Meghan O'Brien, Elizabeth Lanzilla

October 27, 2017

Abstract

The field of developmental robotics is built upon a concrete understanding of children's cognitive development. One of the central components of such development is object permanence, which is the knowledge that an object continues to exist despite its disappearance from view. There are practical advantages to the development of object permanence for concepts involving hidden objects, people leaving a field of vision, and object tracking. This survey paper aims to consolidate extensive research on object permanence in robots, infants, and animals. Topics explored include the relation of object permanence to language acquisition, object motion, and search behavior. Analysis is conducted between the research in robots, infants, and animals to identify gaps in the field of developmental robotics. As an effect, future studies in developmental robotics are proposed as a way to close such gaps in research and benefit the field.

1 Introduction

"Object permanence" describes the perceptual understanding that objects continue to exist even when they cannot be seen, heard, or otherwise observed. Jean Piaget proposed that infants gain knowledge of object permanence at the end of the sensorimotor stage, which ends at around two years of age [11]. However, recent developmental research suggests that infants may achieve this milestone as early as three months old [1]. Object permanence has been studied extensively in psychology, as it is considered a crucial phenomenon in the advancement of human cognitive development. Comparative studies have also investigated cross-cultural and cross-species differences in order to establish the biological and environmental influences associated with object permanence. More recently, continued advances in the field of developmental robotics have inspired research into the ways in which perceptual abilities relating to object permanence can be tested and applied in robots. Robotics research regarding object permanence has many applications, including the augmentation of robots' searching and tracking abilities.

Many testing paradigms, including reaching tasks and impossible-event tasks, have been developed to measure infants' perception of objects that are unobservable. Research in robotics seeks to adapt these paradigms for testing in robots, in order to quantify the extent to which robots can understand the continued existence of hidden objects. Additionally, researchers in the field of developmental robotics seek to understand how object permanence contributes to infants' perception and understanding of other psychological concepts, and extend these findings to applications of robotics in human society. Object permanence has been linked to language acquisition, perception of motion, and searching behaviors in infants, all of which are interesting fields to investigate in robotic development. Paralleling psychological and comparative studies in infants and animals with studies involving robots will allow researchers to design artificial machines that can gain the same developmental traits as humans, with the added bonus of superior processing and computing power.

2 Related Work

In our survey paper we will explore different applications of object permanence across infants, animals, and robots. We have started to conduct initial research in areas such as language acquisition, object motion, and searching behavior. Gaps will be identified thoroughly after more extensive research in order to propose future studies.

2.1 Object Permanence in Robots - A brief glimpse

There are various mechanisms and implementation strategies that can be used when teaching concepts related to object permanence to robots. These implementation strategies largely depend upon the type of robots being used (sensors and hardware etc), the kind of learning being done (supervised, unsupervised, reinforcement etc), and the kind of problem that is being addressed. The research, therefore, has been grouped ahead based on the over-arching theme of the work and the implementation strategies and authors' approaches to tackling object permanence in that particular grouping have been elaborated. While there may be some commonalities, such as the use of machine learning techniques in learning approaches, the use of 3D graphing rendering approaches for forming a robot's representation of the real world and so on, these approaches have been fine-tuned for the relevant case study being conducted [12, 3]. One study, for example, makes use of exploratory behavior and machine learning strategies in order to detect when a robot's movement of its arm (that is holding a shield) causes objects in its visual field to be occluded [3]. Perhaps an interesting approach to this survey paper would be to also identify the kinds of high-level algorithmic approaches (machine learning, 3D graphics, graph theory, visualizations) being used for different studies, if there any general algorithms that are used across the board, and which algorithms are the most beneficial.

2.2 Object Permanence Testing Paradigms in Infants, Robots, and Animals

One of the most important aspects of research in robotics involves finding methods of testing developmental traits in robots. Adapting paradigms that have been tested on humans to accommodate robots allows researchers to evaluate the ability of a robot to perform tasks that are paramount in human existence. One of these such paradigms is the drawbridge experiment [4], in which infants view a screen that oscillates back and forth 180 degrees, similar to a drawbridge. After the infants have habituated to the screen, they see a box placed behind the screen, and are then shown a "possible event" and an "impossible event." In the impossible event, the screen appears to move through the box as if it were not there, while the possible event showed the screen stopping when it reached where the box was known to be. Results of studies utilizing this paradigm have observed that infants look longer at events perceived to be impossible, indicating that the screen defied the infants' expectations. Weng, et al. tested this same paradigm on their SAIL robot using novelty measurements to determine that the robot also perceived the impossible events as surprising [14].

Another variation of the impossible event paradigm shows an object moving along a track and disappearing behind a screen that has a gap in the middle. According to physical laws, the object should be visible as it passes through the gap. However, in "impossible" situations, the object does not reappear in the window [1]. Similar to the experiments conducted by Weng, et al., researchers have found that infants as young as 2.5 months looked longer at impossible situations than at possible situations. This paradigm exists in many forms, with additional variations that investigate infants' ability to discriminate the possibility of a situation based on the height of the window in the screen relative to the height of the object.

Similar to the moving object paradigms in the impossible event tasks, a category of testing paradigms known as visible and invisible displacement tasks has been largely used to study the perceptual abilities of animals [5]. In the "visible displacement" version of these tasks, an object moves across an area and disappears behind a screen. In the "invisible displacement" version, which requires a much more sophisticated skill set, the object is first placed in a container before moving behind the screen, during which time it is then removed, so that only the empty container emerges from the other side of the screen. In both visible and invisible displacement tasks, animals are tasked with searching for the object which, in both cases, will end up hidden behind the screen.

Another category of testing paradigms is the category of reaching tasks, in which infants learn to reach for a mechanism to release a toy, and are then tested on their tendency to reach for the mechanism even when the toy in question is shielded from view. In their 1997 study, Munakata and McClelland tested this paradigm on 7 month old infants, and /'found that infants were less likely to reach for a toy-retrieving mechanism when the toy in question was hidden behind a screen [8].

2.3 Object Permanence and Language Acquisition

The development of object permanence in infants is not an isolated developmental milestone. In fact, there has been research that establishes the relationship between the development of different stages of object permanence and the use of relational words. The study titled 'Cognitive bases of lexical development: object permanence and relational words' refers to the seminal work by McCune-Nicolich that looked into the relationship between object permanence and the development of relational words, in a qualitative manner [13]. The aforementioned study argues that the use of words such as 'more, gone, up' is special because it requires an understanding of the concept of spatial and temporal object transformations such as displacement [13] . The author stresses that the 'emergence of relational words' in infants occurs alongside their entrance into the sixth stage of object permanence development; the stage in which a child is able to understand that objects that have been completely removed from his/her field of view still exist [11]. Having said that, while the work is indeed the first study that finds a positive relationship between the development of object permanence and semantic content, there are several limitations to the study, the most significant one being that the author fails to distinguish between relational words that develop at the fifth stage of object permanence ('visible displacements') from those that develop during the sixth stage ('invisible displacements'), and how the meaning of these words is modified based on the infant's evolved understanding of object permanence.

The current study, therefore, follows the lexical development of six children that are in the fifth stage of object permanence development and follows them into the sixth stage, all the while gathering qualitative data about novel relational word use, the context it was used in and whether the use of the word was productive (used in a novel situation) [13]. This is done through maintaining records of the child's words, weekly maternal interviews and observation of the child's playing sessions. The results are affirmative, in that, the authors find that 'a variety of present-relational words emerged' during stage five. When these six children transitioned into the sixth stage, an emergence of 'absent-relational' words was seen. Children would use the phrase 'allgone' when they found a cup or a box of toys empty or when food was taken out of their perceptual field of vision. The findings also indicated that quantitative metrics, such as the size of a child's vocabulary, did not relate to the development of object permanence. It therefore established that the previous studies that used these metrics and concluded that there was no relationship between semantic content and object permanence, were erroneous.

In the field of robotics, object permanence has been explored alongside a robot's understanding of human speech and an understanding of another person's perspective. In the work 'Mental Imagery for a Conversational Robot', the authors look into what is necessary for a robot to engage in face to face conversations with humans [12]. One such necessary ability is to be able to connect what the speaker is saying to what is present in the environment. The robot must also take into account that the user may be speaking from their perspective; talking about something that is to 'my left'. Therefore, the robot must have the capacity to imagine shift's between its own, and the human's perspective, and must maintain object permanence; a human may be talking about something that is occluded, and therefore cannot be seen by the robot, such as, 'the cup behind the book'. Their approach therefore allows the robot to 'understand and generate' about objects that are not in the camera's field of view. The robot maintains a 'mental model' of the real world, and by using this mental model to imagine shifts in viewpoint, the robot is able to generate 'mental imagery'.

The current robot, Ripley, is able to examine a set of object through touch (its gripper) and vision (its cameras), and is able to manipulate them based on instructions from a human. Ripley's mental model is made up of a model of its own body, a predefined model of the workspace table, a model of the human's body (represented by a sphere), and objects present on the workspace table. Each object in the model is identified by metrics such as its position, color, orientation and so on. In order to register and keep track of objects that it sees, Ripley has an object permanence module called the 'objector'. This module receives input from the 'input processor' unit, which in turn receives a stream of images from Ripley's cameras. When the Objector module finds an image region that exists for a specified amount of time, it creates an object in its mental model and therefore, becomes 'aware' of the existence of the object. Even of Ripley's line of sight changes, the information about the object persists in the mental model. The Objector uses the 2D images, its knowledge about the position of the workspace table, and projective geometry, to represent these objects in the mental model in 3D form. The mental model is combined with a 3D graphics renderer. By positioning a synthetic camera, the graphics environment may be rendered from different viewpoints (mental imagery), therefore allowing Ripley to have shifts in perspective. As far as the dialog system is concerned, the system has multiple components such as 'sensorimotor lexicons and speech recognition parser'. The meaning of words is defined in terms of sensorimotor representations; meaning that verbs like 'pickup' are associated with a series of perceptual conditions and actions, such as ensuring the object is in view, extending the arm, grasping the object and so on. With the help of such Sensorimotor Networks, as well as speech recognizers, and parsers, commands are translated into actions. Words such as 'your' or 'my' are evaluated first, so that the mental model can be altered to account for the shift in perspective. While the design is a lot more extensive, the key takeaway here is that in order to be able to shift between perspectives, the robot must have a concept of object permanence, which has been implemented as stated above.Therefore the system is able to understand spoken language and shifts in perspective, and based on its own mental model and a human's commands, it is able to manipulate objects on the workspace table. This is perhaps just one example of the connection between object permanence and language acquisition as applied to robotics. A more extensive literature survey may lead to more findings.

2.4 Motion

Object permanence is most traditionally thought of as the ability to realize that hidden objects continue to exist. However, object permanence can be applied to hidden objects in motion too. Objects in motion have been tested by infants using anticipatory tracking paradigms briefly discussed earlier in this paper. In past research, there have been conflicting results about an infant's ability to understand that an object in motion can transfer locations because it continues moving out of view. Baillargeon approaches this research question through an experiment involving both the existence and location of stationary and moving hidden objects [2].

Baillargeon presented 6- to-8-month-old infants with both impossible and possible events regarding object permanence. A ramp was placed in the view of the infants with a screen. The infants witnessed a toy car slide down the ramp, pass behind the screen, and emerge on the other side as part of a habituation event. The impossible event was illustrated by a solid green wooden box being placed behind the screen and the toy car appearing to slide through it. The possible event also depicted the toy car sliding down the ramp, but the wooden box had been positioned behind the car so as not to obstruct its path. An infant's interest in either event was measured by how long the infant looked at the event [2].

The experiment determined that when infants were subjected to the impossible event, they demonstrated a greater interest in it by looking longer. It is extrapolated that this result occurred because the infant understood the impossibility of the event. Therefore, the infant exhibited object permanence with the stationary box covered from view with the screen, the car's continued movement behind the screen, and the car's inability to pass through a solid object [2]. These results can be extended to the field of developmental robotics for applications such as continuous object tracking when the object disappears from view.

Lovett and Scassellati examined Baillargeon's study through a robot, Nico, which is a humanoid robot that represents a 9-month old infant at the Yale Social Robotics Lab. The ramp and screen were set up in the same manner as Baillargeon's study. Instead of a box representing an obstruction to the ramp, a green folder was used. However, to Nico this difference is nonexistent as it cannot conceptualize 3-d objects. In order to determine a loss of interest in the trials, Nico needed to look away for 3 seconds [6].

In this experiment, it was determined that the robots did look longer at the impossible event, the event when the green folder should obstruct the path of the car but does not. However, Lovett and Scassellati posit that Nico took more interest in the impossible event because of the novelty of the green object, not because of the event's contradictions with object permanence. In relation to Baillargeon's study, Lovett and Scassellati believe that the novelty of the green block may perhaps explain the results [6]. Following the inconclusive results of this study, further research into objects in motion must be conducted in both the fields of infants and robots to better understand the phenomenon.

2.5 Searching

The concept of object permanence displays contradictions in regards to searching behavior. While evidence of object permanence has been found in infants as young as 3 months old, infants do not exhibit search behavior until after 8 or 9 months despite being physically able. This phenomenon is widely known as "Paradox of Object Permanence." Research into this paradox alludes that an infant's understanding of an object's continued existence is only strong enough to produce looking behavior, not searching. However, there is a lack of evidence for why search tasks would need deeper representation [9].

O'Connor and Russell hypothesize that the paradox in object permanence arises due to the difference in looking being a passive action and searching being active. If an infant learns that object permanence can be extended to retrieving a hidden object and allowing it to reappear, it is thought that searching will ensue [9].

To approach this hypothesis, 44 infants between 6 and 8 months old were randomly assigned to an Agency group or a Means-Ends group. The Agency group was presented with a turntable, toys, and an opaque screen between the two sides of the table. The Mean-Ends group was contrastingly presented with the same setup, but with a transparent screen. In the training portion of the study, an infant selected a toy, watched from an immobile position as the toy was placed underneath a cloth, and was given 45 seconds to search for the toy. There were two trials of this before the turntable was presented. The researcher taught the infant how to spin the turntable by playing music after every successful turn. Four training sessions were conducted with the turntable in which the infant was taught to bring a hidden toy successfully back into view with the researcher's assistance [9]. Following these training sessions, the infant was presented the toy again and restrained while it was hidden. The infant searched for the toy for 30 seconds without any outside help. Following 30 seconds, the researcher drew the infant's attention to the toy (lifting the screen in the Agency group) and gave the infant another 30 seconds to search [9].

O'Connor and Russell determined that presenting infants with the experience of uncovering a hidden object creates stronger performances in searching behaviors of infants aged 6 to 8 months old. The Agency group improved in their searching behavior more than the Means-Ends group following two days of testing. Overall, the two groups performed very similarly. In terms of this studies application to developmental robotics, researchers can see that the object permanence is a critical precursor to the development of searching behavior. Furthermore, searching behavior can potentially be improved in robotics through experiences derived from an understanding of object permanence [9].

Object permanence is also an important concept when it comes to search strategies. The ability to scan the environment and identify object in the real world requires the understanding that objects may be partially occluded. A cup is still a cup, even if a robot is unable to see a part of the cup. The ability to correctly identify and retrieve objects that may be partially hidden is a very important quality, especially in robots that are designed to serve as assistants. In the work titled 'Curious George: An attentive semantic robot', the authors tackle this very task [7]. During the first phase, the robot must build up a representation of a list of objects using images from the internet (training phase). In the next phase, it must explore the environment and correctly identify the objects in the list by matching images it gained from the training phase, to those it gained in the real world (testing phase). Seeing as objects in the real world may look very different from those in images acquired from the internet (the object may be at a different angle, it may be hidden etc), the robot is programmed to take images of the real world objects even as it moves past them, and as it leaves them behind. Taking images from numerous viewpoints increases the likelihood of a match between the training images and the testing images. The concept of object permanence is important here as the robot must be aware of the location of the object that it leaves behind (and that is no longer in its visual field) so that it may take its image from numerous points of view.

2.6 Animal studies

Some of the research surrounding object permanence has revolved around the question of whether or not object permanence is a uniquely human characteristic. Research by de Blois et al. in 1998 found differences in the abilities of nonhuman primates regarding object permanence, after testing orangutans and squirrel monkeys [5] on a series of visible and invisible displacement tasks. The results of these studies found that orangutans were able to successfully complete the visible displacement tasks, despite showing increased difficulty for visible displacement tasks that included three displacements. The orangutans were also able to complete the invisible displacement tasks that only included one displacement, but failed to solve double invisible displacements. However, a follow-up study found that, when a cup was used as a container for the hidden objects, instead of a human hand, the orangutans were able to perform better than chance on the double invisible displacement tasks. Together, these results suggest that, while orangutans have the capacity to understand invisible displacement, their abilities are limited for more complex tasks. In the studies with squirrel monkeys, researchers observed that squirrel monkeys were similarly able to complete tasks involving visible displacements, but showed significant declines in performance for tasks involving invisible displacements.

Additionally, Irene Pepperberg and Florence Kozak expanded the comparative study of object permanence to include birds, when they investigated the perceptual capabilities of an African Grey parrot that was, at the time, being trained and assessed on a number of language tasks [10]. The parrot, Alex, was tested on a battery of tasks, which involved searching for hidden or partially hidden objects, following moving objects, noticing the disappearance of objects, and finding objects that have been displaced one or more times. The results of this study found that Alex was able to complete all of the tasks, demonstrating a capacity for object permanence in non-mammalian animals. While researchers cannot generalize that the subject was using the same mechanisms to perform these tasks as humans would, these results show that object permanence is not uniquely human, a fact that also has broader implications for the application of object permanence to machines.

3 Future Work

So far we have tried to identify some of the more interesting research problems being addressed and the different ways research into object permanence in infants and robots is grouped. However what is lacking in our current perusal of this research space is the identification of sub areas where there is a dearth of research. Once we identify these gaps, by way of doing a more extensive literature survey, we can perhaps suggest approaches to how research in those areas can be conducted, and if some concepts of object permanence development in infants may be useful here.

4 Timeline

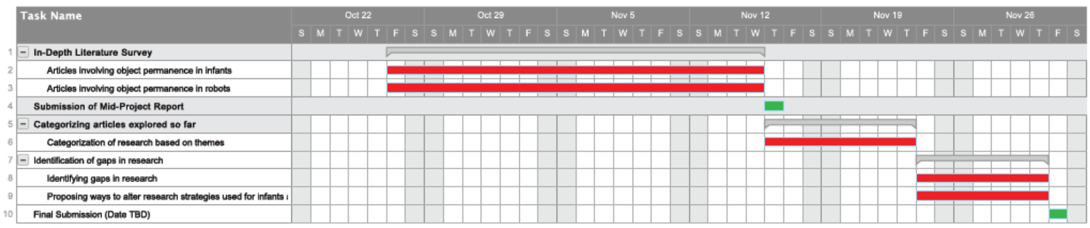

Figure 1: Project Timeline

References

[1] Andréa Aguiar and Renée Baillargeon. Developments in young infants' reasoning about occluded objects. *Cognitive psychology*, 45(2):267–336, 2002.

[2] Renee Baillargeon. Representing the existence and the location of hidden objects: Object permanence in 6-and 8-month-old infants. *Cognition*, 23(1):21–41, 1986.

[3] Sarah Bechtle, Guido Schillaci, and Verena V Hafner. First steps towards the development of the sense of object permanence in robots. In *Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 2015 Joint IEEE International Conference on*, pages 283–284. IEEE, 2015.

[4] Richard S Bogartz, Jeanne L Shinskey, and Thomas H Schilling. Object permanence in five-and-a-half-month-old infants? *Infancy*, 1(4):403–428, 2000.

[5] Sandra T De Blois, Melinda A Novak, and Melanie Bond. Object permanence in orangutans (pongo pygmaeus) and squirrel monkeys (saimiri sciureus). *Journal of Comparative Psychology*, 112(2):137, 1998.

[6] Andrew Lovett and Brian Scassellati. Using a robot to reexamine looking time experiments. In *Proceedings of the Third International Conference on Development and Learning*, pages 284–291, 2004.

[7] David Meger, Per-Erik Forssén, Kevin Lai, Scott Helmer, Sancho McCann, Tristram Southey, Matthew Baumann, James J Little, and David G Lowe. Curious george: An attentive semantic robot. *Robotics and Autonomous Systems*, 56(6):503–511, 2008.

[8] Yuko Munakata, James L McClelland, Mark H Johnson, and Robert S Siegler. Rethinking infant knowledge: toward an adaptive process account of successes and failures in object permanence tasks. *Psychological review*, 104(4):686, 1997.

[9] Richard J O'Connor and James Russell. Understanding the effects of one's actions upon hidden objects and the development of search behaviour in 7-month-old infants. *Developmental science*, 18(5):824–831, 2015.

[10] Irene M Pepperberg and Florence A Kozak. Object permanence in the african grey parrot (psittacus erithacus). *Learning & behavior*, 14(3):322–330, 1986.

[11] Jean Piaget and Margaret Cook. *The origins of intelligence in children*, volume 8. International Universities Press New York, 1952.

[12] Deb Roy, Kai-Yuh Hsiao, and Nikolaos Mavridis. Mental imagery for a conversational robot. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34(3):1374–1383, 2004.

[13] Michael Tomasello and Michael Jeffrey Farrar. Cognitive bases of lexical development: Object permanence and relational words. *Journal of child language*, 11(3):477–493, 1984.

[14] Juyang Weng, Yilu Zhang, and Yi Chen. Developing early senses about the world:" object permanence" and visuoauditory real-time learning. In *Neural Networks, 2003. Proceedings of the International Joint Conference on*, volume 4, pages 2710–2715. IEEE, 2003.