# AISB Quarterly

THE NEWSLETTER OF THE SOCIETY FOR THE STUDY OF
ARTIFICIAL INTELLIGENCE AND SIMULATION OF BEHAVIOUR

No. 140 January, 2015

Do you feel artistic?
Exhibit your artwork on our front covers!
Email us at aisbq14@aisb.org.uk!

This artwork is based on Particle Swarm Optimization (PSO) algorithm. PSO is an evolutionary computation technique developed in 1995 by Kennedy and Eberhart, and it's inspired by social behavior of bird flocking or fish schooling.

Particles in PSO are made to follow a hypothetical point (focal point, $f_p$) moving horizontally (i.e. scanning each row with a constrained random vertical offset); once $f_p$ reaches the end of a line, it goes to the next row; this process is repeated until the entire input image (the Q of the Quarterly) is scanned. As particles trace the $f_p$ , the average colour of the pixel, where each particle is 'flying over', is taken and the colour is reflected on the output image (the current cover of the Q). In other words, the colour of each pixel of the input image changes the corresponding part of the output image to the swarm's average colour as they flock around $f_p$. The swarm's flocking behaviour over the image creates the resultant cover for this issue of the Quarterly.

The technique will be presented in detail at the A-EYE Art Exhibition symposium at AISB50, 1st April–4th April. See website for more information:
http://aisb50.org/a-eye-an-exhibition-of-art-and-nature-inspired-computation/



Russell Eberhart & James Kennedy, 'A new optimizer using particle swarm theory', *Proceedings of the Sixth International Symposium on Micro Machine and Human Science MHS'95*, IEEE, (1995).

# It should look like a robot: Mental models of anthropomorphic features

by *Christine Edwards-Leis* (St Mary's University )

## Abstract

Mental models were proposed by Craik (1943) to explain human-computer interaction. They are of particular interest to educational researchers, because they explain what happens when teachers and students interact with each other and with phenomena in their environment. They have been said to form the basis of all human behaviour and, as such, are simultaneously a process where they provide a means to interact in problem-solving situations and a product where they act as a storage facility to retain the knowledge from those interactions.

A longitudinal study of the mental models held by students and their teacher in an Australia Primary School affirmed the role mental model theory has to play in enabling researchers, and teachers, to understand how students engage with artefacts successfully in a problem-solving environment. The study uncovered the students' mental models of the anthropomorphic features of the robots with which they were engaging. These findings offer design implications for the development of robotic technologies for use in classrooms of the future. They may also provide guidance for the wider animatronics community. Students were able to engage successfully in problem-solving activities in robotics but their mental models indicated a clear preference for humanistic features.

## Context

The purpose of the longitudinal study (March 2005–October 2006) of a teacher and her 24 students in a primary school in South-East Queensland, Australia, was to determine and analyse their mental models of teaching, learning, and assessment in robotics, an optional component of the Queensland Technology Syllabus. The study had broad aims to determine matches and mismatches of mental models and the impact, if any, on teaching, learning, and assessment. Anthropomorphism was one of the aspects of the participants' mental models of robotics that was of interest to the researcher. The results of the investigations into anthropomorphism had limited impact on the broader study's purpose. However, they offered salient data about the anthropomorphic characteristics of robots that students found relevant in their lives. This data has the potential to contribute to future design choices of those who envisage learning with such interactive technological artefacts or those who work in the animatronics field.

## Theoretical background

Mental models were first theorised by Craik (1943) who was searching for a means of explaining the interactions between humans and systems. Craik initiated the use of the term, mental models, and described them as "representations in the mind of real or imaginary situations" (Craik, 1943, p. 12). He used the subsequent theory to describe how we understand, solve, and explain anticipated events. Mental models enable users of systems to explain and make predictions about the actions and reactions of those systems (Halford, 1993; Vosniado, 2002). Mental model theory has permeated many areas of human endeavour including technology (Edwards-Leis, 2007; Halford, 1993; Henderson & Tallman, 2006; Vosniado, 2002; Williamson, 1999) and language (Johnson-Laird, 1983, 1989, 2004, 2006; Merrill & Gilbert, 2008). The functional aspect of mental models enables students to explain processes, predict outcomes of those processes, and communicate their understanding to others (Edwards-Leis, 2010). Using mental model theory to determine what is happening in classrooms as students engage in problem-based learning (Edwards-Leis, 2007; Henderson & Tallman, 2006; Stripling, 1995) is a focal way of determining what, in reality, is being learned.

## Methodological processes

The study of mental models, in this broader longitudinal study (Edwards-Leis, 2010), was centred within information processing theory (Kail & Bisanz, 1992; Lohman, 1989, 2000) and linked with the introspection meditating process tracing paradigm. The study followed constructs of learner-centredness and how students select, organise, and integrate new experiences with existing knowledge and the processes (Edwards-Leis, 2010) used in metacognitive activity (Mayer, 1996). It was conducted on the campus of a P-7 Primary School in South-East Queensland, Australia and the participants were aged 10 years at the commencement of the study and were 12 years old at its completion. The qualitative study used social anthropological perspectives to select, focus, simplify, abstract and transform data (Miles & Huberman, 1994) to gather a vivid picture of teaching, learning, and assessment in robotics from the perspectives of the teacher and the students.

## Data collection

Data were collected from the teacher and her students over a 20 month period from individual and paired Semi-Structured and Stimulated Recall Interviews, Journals, Likert Scale Questionnaires, a Teach-Back Interview, and a Focus Group Interview (Edwards-Leis, 2010). While 24 students participated in providing data from journals and the questionnaires, four students were randomly selected to participate in the more in-depth aspects of the investigation. Data were triangulated (Burns, 2000; Miles & Huberman, 1994) to strengthen validity of analysis. Post- and pre-experience data underwent comparative analysis and, while the population was too limited to provide generalisations, the rigour and intensity of the investigations offered a "focus on the complexities and qualities in educational action and interaction that might be unattainable through the use of more standardised measures" (Burns, 2000, p. 390). In reality, data was not just triangulated – it was "multi-angulated" and offered a detailed picture of teaching, learning, and assessment over time (Edwards-Leis, 2010). Data for the anthropomorphic aspect of the study were collected from questionnaires, journals, and semi-structured interviews; it is in-depth and offers significant information about the mental models young students hold about the artefacts teachers require them to use in problem-solving situations.

## What is a robot?

Did you watch The Lord of the Rings (Jackson, 2001/2002/2003) and become a little bit uncomfortable as Gollum joined the intrepid journeymen on their trek? If you did, then you were well-targeted by the "Gollum effect" (Giles, 2007). Giles (2007) described the "Gollum effect" as an example of the phenomenon known as the "uncanny valley" (Mori, 1970) where, as a general rule, people are less troubled by a robot that is clearly a robot. When a certain realism threshold is reached, through such displays as human-like movement or vocalisation, then people exhibit unsettled reactions to the robotic creation. According to Giles (2007) this effect is what Gollum's designers, Weta Digital of Wellington, New Zealand, relied upon to create the character who was introduced in The Hobbit (Tolkien, 1937) and reappeared in the subsequent film trilogy, The Lord of the Rings (Jackson, 2001/2002/2003). They gave Gollum a human-like voice and animal-like body movements which were in the threshold that created uncomfortable responses to the character in many viewers of the film.

Why is this reaction of cognitive relevance rather than mere entertainment interest? Humans are interacting more with robots in a multitude of environments including the home, work, and recreation facilities and such human responses to robots are of increasing interest to designers who undertake their creation. Studies (Chaminade, Hodgins & Kawato, 2007; Gee, Browne & Kawamura, 2005; Hinds, Roberts & Jones, 2004; Kiesler & Goetz, 2002; Barchi, Cagliari & Giacopini, 2002) have been engineered to understand human-robot interactions and the individual reactions and interactions that are associated with strong or weak anthropomorphic features of robots. Meanwhile, the incidence of human-robot interaction is becoming more of an everyday event from medical nanotechnology to robotic scuba-divers plunging into the depths of the ocean in search of lost ships to Spielberg's AI: Artificial Intelligence (2001), a tale of an android child who was programmed to love, or Proyas' I-Robot (2004) where humanity is threatened by robotic crime. These robotic realities and fictionalisations are becoming part of the average day for children. Does this exposure create mental models that will enable the students to function effectively with the robotic experiences they will encounter in the classroom?

The anthropomorphic features a robot possesses may offer a focus for the interaction and replicate the everyday interactions humans have with each other. How humans respond to the variations of humanness of a robot's features has been the focus of many studies (Chaminade, et al., 2007; Gee et al., 2005; Kiesler & Goetz, 2002). While results indicate that, in the main, people tend to engage more readily with a robot with humanlike appearance, Gee et al. (2005) discovered that the actual concealment of artificiality of a robot is the source of most discomfort. Kiesler and Goetz (2002) found that it was the humanistic dialogue that adversely affected people's mental models of what constituted a comfortable interaction more than their responses to the robot's humanlike physical appearance. Given this robust investigation into human/robot interaction, what would be of interest is how the participants, in this study, would perceive the relevance and importance of the appearance of the robots with whom they would be working.

## The robots we *might* prefer: Pre-experience data

The first exploration by the researcher was to uncover the conceptual knowledge that demonstrated the espoused mental models of robots held by the participants. Much data was collected on a variety of issues such as the students' mental models of desirable features and their preference of working with a robot that looks, speaks, and moves like a human. The students who were to participate in the project and the study had been given no indication of the style of robots they would be using by Pamela, the teacher. At this pre-experience stage, their mental models were totally unaffected by any classroom interactions, discussions, or preparations associated with the study. When Pamela, was interviewed in the pre-experience interview, she revealed that her mental model of the students' mental models of robotics would be that they would not necessarily be humanistic in appearance. This is evidenced by her comment, "I think that they have a little bit of an understanding that robots don't always necessarily look like a human being" (Pamela, Pre-Experience Interview, March 2005). She felt that their main exposure would have been "through movies and TV".

Pamela's responses to the pre-experience Likert Scale questionnaire supported her belief that she could not really have an understanding of how the students viewed robots and the necessity for humanoid features until she commenced working on the project. She "disagreed" that the students would see robots as more useful in they responded like humans and also with the proposal that students would rather interact with a robot that was humanlike in appearance. Pamela was "unsure" of the students' mental models of the usefulness of robots if they could talk like humans which, after analysis, was a perceptive response given that all students agreed with the statement in their pre-experience Likert Scale questionnaire (see Table 2). Pamela was not asked to provide any comments about facial features, eyes, ears, and mouth in her pre-experience Likert Scale questionnaire but the responses to these questions by the students provided rich data on the their mental models of the anthropomorphic features of robots.

The pre-experience Likert Scale questionnaire asked specific questions about the anthropomorphic features of robots such as those shown in Table 1. The students appeared to be more inclined to respond in the affirmative for specific anthropomorphic facial features that they would prefer a robot to have such as eyes, ears, and mouth than they were for general human looks. The item that proposes "looks like a human" does not refer specifically to facial features and may have prompted students to run mental models of robots with humanoid body features of arms, legs, and torsos.

The pre-experience, semi-structured interviews uncovered more detail with students referring specifically to a robot's appearance with responses: "It should look like a robot" and "... not exactly the same [as a human] ... but a little different". Three of the four students stated their preference for a robot to have eyes, nose, mouth, legs, arms, and two wanted hair. The other human features mentioned once by each student included fingers, face, voice, feet, teeth, and coloured face. All four students, Bree, Ellen, Jayne and Sam, who participated in the Likert Scale questionnaire, appeared to prefer a robot that had human features. The students' mental models of robots may be affected by cultural considerations (Gee et al. 2005) such as their personal experiences with television and motion pictures or it might be reflecting their mental models of "interactions" in general. Kiesler and Goetz's (2002, p. 1) study indicated that people tend to create anthropomorphic mental models "of higher animals, deities, nature, and animated objects and machines". Asking students whether or not the robot with which they will be interacting should look like a human may be the same as asking them if they thought the Easter Bunny was capable of delivering eggs!

Our cultural tendency to personalise inanimate objects, such as machines (Kiesler & Goetz, 2002), often by giving them names, may encourage the development of strong anthropomorphic mental models when we interact with robots and androids. This may be an example of instance-based cognitive processing which Hintzman (1986) saw as necessary to integrate new knowledge and experiences so that a productive interaction took place. Such an interaction would lead to the formation of a functional mental model. The different images necessary for exemplar-based processing (Linville, Fischer & Salovey, 1989) occur where separate images are linked or joined to create a consistent or acceptable anthropomorphic mental model (Keisler & Goetz, 2002). Keisler and Goetz (2002, p. 1) gave a "cheerful robot" as an example of this type of processing where a "life-like robot that tells a joke could activate ... exemplars of ... machines and ... humorous people". The "humanlikeness" of a robot through either its behaviour or appearance can, they believe, "lead to a mental model that does not deny the technology in the machine but that also incorporates anthropomorphic features into it" (Keisler & Goetz, 2002, p. 2).

It seems that if students are given a preference, as suggested in the pre-experience Likert Scale questionnaire, then the predilection is for a robot that is humanlike in appearance. While no evidence was sought to determine if these espoused mental models had been formed from interactions or experiences from television or motion pictures as suggested by the teacher

| I would rather interact with a robot that.. | | | | | |
|---|---|---|---|---|---|
| | Strongly Agree | Agree | Unsure | Disagree | Strongly Disagree |
| Q. 20 ..looks like a human | 4 | 8 | 9 | 3 | 1 |
| Q. 21 ..has eyes | 15 | 7 | 2 | 0 | 0 |
| Q. 22 ..has ears | 11 | 12 | 1 | 0 | 0 |
| Q. 23 ..has a mouth | 14 | 9 | 1 | 0 | 0 |

Table 1: Students responses to Likert Scale anthropomorphic questions.

Pamela, the students have, nonetheless, created mental models that may reflect such a socio-cultural influence (Vygotsky, 1978). The students' espoused mental model of the mechanical nature of the robot, though, is clearly shown by two of the students interviewed. As one interviewee participant, Bree, stated a robot has "memory chips" to help it move. Another, Sam, referred to robots as "artificial intelligence" which may indicate his viewing of the movie AI: Artificial Intelligence (Spielberg, 2001) released prior to the time of the research.

Question 22 in the pre-experience Likert Scale questionnaire asked whether students would rather interact with a robot that had ears (see Table 1). There was a response of 23/24 in the affirmative with one student unsure. Given the responses from the students in the interviews, one may wonder if this result indicates that the students' initial mental models of facial features do not include ears unless a specific prompt is given. During the semi-structured interviews the four students' responses did not propose ears as part of the desired facial features but all (N=4/4) affirmed that we "tell" robots instructions. Half (N=2/4) said that robots do what you "say" (Ellen) or can "talk" (Jayne) with you. Sam stated that robots have a "chip that helps them listen" and Bree said that we put information "into their heads". Robots needed to "talk" (Jayne) and "see" (Bree) but only one student, Ellen, included "ears" as a necessary feature that would enable a robot to hear or listen once asked how robots would do that–in other words, given a specific prompt.

One of the components that can be manipulated by digital/robotic animators is the voice that their creation is given. The student participants were asked if they agreed with the statement: "Robots are more useful if they can talk to you" in the pre-experience Likert Scale questionnaire and the responses are shown in Table 2.

Seventeen students strongly agreed and six agreed with the statement (Table 2). One child, a girl called Tani (pseudonym), was unsure whether the ability to speak made a robot more useful. Tani's espoused mental model, as indicated by her journal response on the usefulness of robots, included the concepts of "moving, walking and talking" and being able to be "programmed" to "do everything you tell them". Her re-

| Q. 19: Robots are more useful if they can talk to you. | | | | | |
|---|---|---|---|---|---|
| | Strongly Agree | Agree | Unsure | Disagree | Strongly Disagree |
| Responses | 17 | 6 | 1 | 0 | 0 |

Table 2: Students responses to Likert Scale question about the usefulness of robots.

sponse to the Likert Scale questionnaire did not seem to match that which she provided in her journal and this may indicate that there was some uncertainty, on her part, about the context in which "talking" is useful. However, the general affirmative response from the group would seem to indicate that the students' espoused mental models included the usefulness of robots being able to talk to humans. The responses in the pre-experience investigations provided some interesting data from which to make comparison six months later when the post-experience data was collected.

## The robots we prefer–now! Post-experience

The robots that the students constructed and programmed during the learning experience had no human characteristics so there was no concealment of their artificiality that might cause any perceived discomfort (Gee et al., 2005). The students' espoused mental models of anthropomorphic issues, which indicated a preference for human characteristics on robots, did not appear to discourage them from participating in the activities to construct and program their robots which were made from a Lego™ "brick" with attached motors, wheels, and sensors. The robot made from Lego™ seemed far removed from the robots children might have seen on television programs, in commercials, and in feature films.

Pamela's espoused mental model had indicated uncertainty about whether or not students were concerned with how a robot looked. Her reflective mental model altered considerably because in the post-experience Likert Scale questionnaire she responded with a strong positive to Item 19: "Students would rather interact with a robot that is humanlike in appearance." Her espoused mental model was interesting, but predictable, when compared to those of the students who had responded positively in the pre-experience Likert Scale questionnaires about this item. Pamela's uncertainty of their mental models of anthropomorphic issues would be unsurprising until she had the opportunity to observe their interactions with the equipment and with each other. Her espoused mental model may have included some doubt as to the students' concepts of robots and their functionality, but this doubt was removed as she worked with the students during the course of the learning experience. Pamela was not requested to respond to any items on the post-experience Likert Scale questionnaire that included personal anthropomorphic considerations, but her responses to such items concerning the students' mental models continued to be of interest as a reflection of her increasing knowledge of the students with whom she interacted.

The students' post-experience Likert Scale questionnaire repeated the questions about the anthropomorphic features of robots from the pre-experience questionnaire. There was a slight movement toward the negative in all three items that addressed facial features of robots – eyes, ears, and mouth. Graphs illustrating the change in responses are shown in Figures 1, 2, and 3 below. In each of the items there has been an evident shift to the negative which indicated that the students were less concerned with the inclusion of human facial features or characteristics in their reflective mental models than they had been at the beginning of the experience. The robots they were working with did not have any human features yet were able to follow the commands that the students had designed and programmed. While the efficacy of mental models held by individuals may be seen as the basis of success and/or failure (Henderson & Tallman, 2006) in undertaking a problem-solving activity, the experience of success and/or failure itself can be powerful motivator to actually alter inefficacious mental models. It seems that several students have altered their espoused mental models in these areas and the responses to the question of the preference for a robot to have a mouth (see Figure 3) show the most distinct move to the negative. This could be attributed to the fact that the ability to communicate verbally is less of an issue than one of seeing what there is to do or where one is going (see Figure 1) and the ability to listen to instructions (see Figure 2).

# Digging deeper

Digging deeper While mental models are idiosyncratic, the predominance of a shared response indicates that some mental models are distributed. The students in this case were not negotiating a transitory, shared model (Anderson, Howe & Tolmie, 1996). However, Vygotskian theory (1978) may inform analysis that concludes that the construction of this particular mental model is a result of a shared experience either in the robotics classroom or by other experiences with media. The responses to investigations on anthropomorphic issues in the post-experience semi-structured interviews uncovered more detail and some interesting issues about such human characteristics in robots.

Bree, Ellen, Jayne and Sam were asked about their preferences about anthropomorphic issues in order to provide an opportunity to extrapolate on their mental models and to triangulate data obtained from the Likert Scale questionnaires and journals. There were notable differences in the responses to questions about what human characteristics, if any, the students preferred a robot that they were interacting with to have.

Jayne summed up their individual responses by stating that "there's no normal robot"; indicating that there are diverse robotic designs and that "they can look a bit like anything" (Jayne, Post-Experience Interview, September, 2005). While Jayne continued to express a reflective mental model that

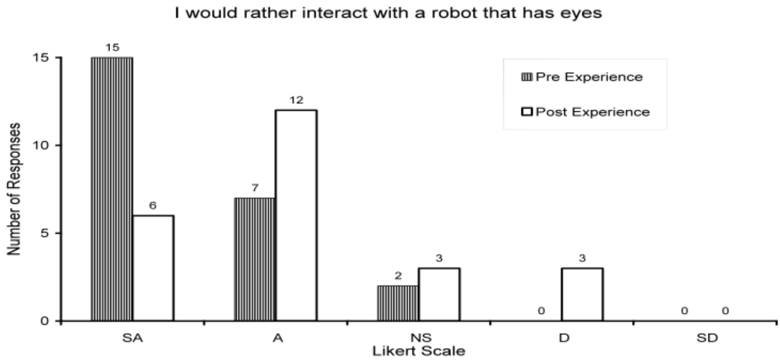


Figure 1: Comparison of responses to Item 21 on Pre- and Post-Likert Scale questionnaire about interacting with robots that have a eyes.

contained an understanding that robots were "different to humans", she believed that "some of them act like humans" because they are "trained" to do so. However, she would like her robot to have "feelings" and be able to "talk" so she would know what it wants to do next. Her inclusion of the ability to feel was so that she would know when it was "happy and sad"; a clear inclusion of human emotional characteristics. Jayne's responses in the Likert Scale questionnaire to items including eyes, ears, and mouth (see Figures 1, 2, and 3) were all in the negative indicating that she felt that these anthropomorphic features were unnecessary.

Ellen's responses in her post-experience, semi-structured interview supported her responses on the Likert Scale questionnaire. While she disagreed with the need for a robot to have a mouth, she did agree that she would rather interact with one that had eyes and ears. She stated that robots "don't really have to look like what you expect because some robots do different things" but she would like for them "to see and hear things because ... if it can hear, you might be able to tell it something and it might be able to follow that instruction" (Ellen, Post-Experience Interview, September 2005). If it could not see, "it might run into a few walls". Ellen differentiated some tasks that robots do, including washing dishes and washing clothes. Robots do not do both because "it [the clothes washer] might break the plates". This differentiation indicates her reflective mental model is developing more expert viewpoints that discriminate a designated functionality of a robot and the irrelevance of its appearance to complete those tasks. Beliefs can influence a student's thinking (Szabo, 1998) and teachers need to recognise that the authority of viewpoints may influence an individual's ability to effectively incorporate new information into

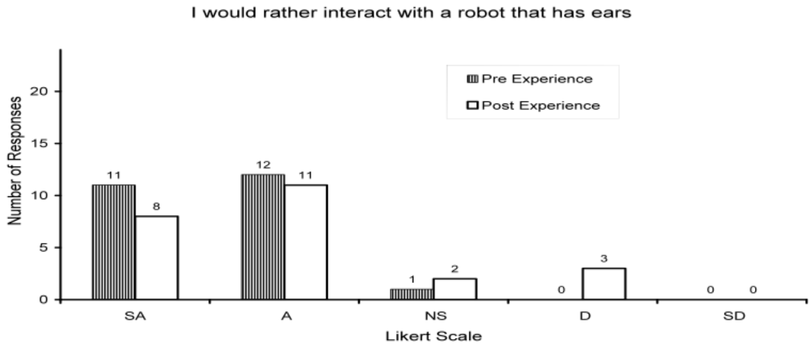


Figure 2: Comparison of responses to Item 22 on Pre- and Post-Likert Scale questionnaire about interacting with robots that have ears.

the mental models that are required to interact effectively in a given domain.

Sam's reflective mental models on robots and what they do had developed from his exposure to media and personal experience and displayed more connectedness of such understandings (Stripling, 1995). He agreed that he would prefer to interact with robots that had eyes, ears, and mouths in his responses on the Likert Scale questionnaire. These positive responses were not repeated in his semi-structured interview where he said that a robot's appearance would "depend [on] what they're supposed to be for" (Sam, Post-Experience Interview, September 2005). This statement indicated the incorporation of propositional knowledge (Johnson-Laird, Girotto, & Legrenzi, 1998; Preece, Rogers, Sharp, Benyon, Holland, & Carey, 1994; Reddish, 1994) in Sam's reflective mental model. He had seen a robot in a television program that was "just a flat piece of metal ... like a rectangle with tracks on it" so he was aware that functionality informed design. Sam was also interested in creating robots that could "drag race" and he had been working in the robotics laboratory to develop a robot that would do that. One of the human characteristics that he spoke about in the interview that would assist in his quest for a drag racing robot was the need for "common sense" as this would allow the robot to "overrun the program" if it was not going long enough to win the race; again an example of his incorporation of propositional knowledge in his reflective mental model.

Bree was the one student interviewed in the post-experience, semi-structured interviews whose Likert Scale responses to the anthropomorphic items were strongly positive. Her clearly-expressed need to have human characteristics on the robots with which she interacted was repeated in the post-experience interview, where she wanted her robot

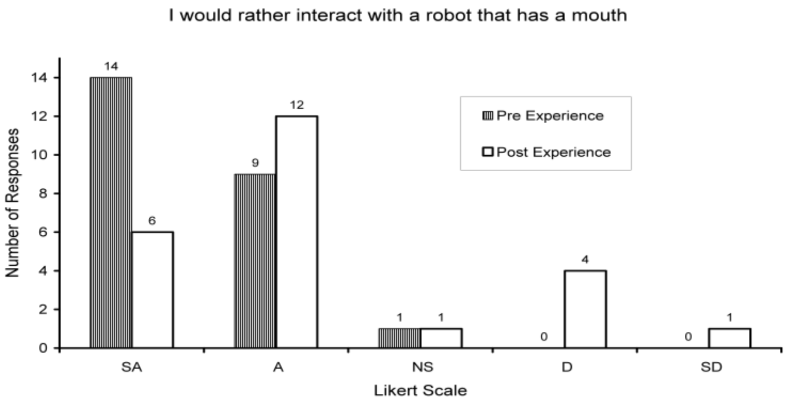


Figure 3: Comparison of responses to Item 23 on Pre- and Post-Likert Scale questionnaire about interacting with robots that have a mouth.

to have human characteristics so she "could get along with them faster". She clarified this statement by saying that she worked better with people than she did with "cars and stuff like that". Bree wanted eyes "so it can see", ears "so they can hear you" and a mouth "so it can talk to you" (Bree, Post-Experience Interview, September 2005). She felt the need to speak was important so that if she missed a part or got a wrong piece while constructing them then they could tell her "without me having to look it up". Bree's inclusion of this human characteristic to talk had a distinct purpose and was likely influenced by her problem-solving strategy of going back over the construction and programming of her robot to find an error. Bree's responses are evidence of the mechanism by which a mental model can be used to understand the "self-reflective aspects of the self" (Power & Wykes, 1996, p. 240) through her demonstration of how objects (robots) relate to both herself and her interactions in the world of robotics. It is clear that Bree's reflective mental models of robots and their anthropomorphic characteristics have incorporated many levels of meaning.

*Paper presented at UBC, Vancouver at the Technological Learning & Thinking 2010 Conference (17 -19 June)*


**Christine Edwards-Leis, PhD**
Senior Lecturer in Education
School of Education, Theology and Leadership
St Mary's University, UK

# Back matter

# AISB Quarterly – No. 140, January, 2015



**It should look like a robot: Mental models of anthropomorphic features — by *Christine Edwards-Leis* (St Mary's University ) 3**


The AISB Quarterly is published by the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB). AISB is the UK's largest and foremost Artificial Intelligence society. It is also one of the oldest-established such organisations in the world. The society has an international membership of hundreds drawn from academia and industry. We invite anyone with interests in artificial intelligence or cognitive science to become a member.

AISB membership includes the following benefits:

- Quarterly newsletter
- Student travel grants to attend conferences
- Discounted rates at AISB events and conventions
- Discounted rates on various publications
- A weekly email bulletin and web search engine for AI-related events and opportunities

You can join the AISB online via: http://aisb.org.uk


ISSN 0268-4179