

Emotional Computation in Artificial Intelligence Education

Sara Owsley Sood

Department of Computer Science
Pomona College
185 East Sixth Street
Claremont, CA 91711
sara@cs.pomona.edu

Abstract

After decades of work towards creating *artificial intelligence*, some researchers are now attempting to create machines that are *emotionally intelligent*. The standard definition of artificial intelligence is the ability for a machine to “think or act humanly or rationally” (Russell, 1995). The push towards *emotionally intelligent* machines is an exciting addition to this field, in the direction of machines that truly “act humanly.” In this article, I will provide a survey of topics and resources for teaching emotional computation in artificial intelligence courses. I will show that this topic area serves as an excellent application of modern artificial intelligence techniques, and is an important aspect of modern research in AI. The interdisciplinary nature of work in this space is not only compelling in the classroom, it may also lead to wider interest in the field of Computer Science.

Introduction

Many researchers in the fields of Artificial Intelligence and Human Computer Interaction have projected that the machines or intelligent agents of the future must connect on an emotional level with their users (Norman, 2004). This is based on the notion that an intelligent and successful human is not only strong in mathematical, verbal and logical reasoning, but is able to connect with other people. Much recent work in this area has focused on empowering agents with the ability to both detect emotion via verbal, non-verbal, and textual cues, and also express emotion through speech and gesture. In this article, I will present evidence of this movement towards systems with emotional intelligence while also showing why and how this topic should be included in introductory artificial intelligence courses.

Emotional Intelligence

The concept of *Emotional Intelligence* became prominent in the late 1980’s; however, Thorndike discussed a similar concept called *social intelligence* much earlier, in 1920

(Thorndike, 1920). While one’s *social intelligence* is typically defined by their “ability to understand and manage other people, and to engage in adaptive social interactions” (Kihlstrom, 2000); *emotional intelligence* deals specifically with one’s ability to perceive, understand, manage, and express emotion within oneself and in dealing with others (Salovey, 1990). Salovey and Mayer define five domains critical to *emotional intelligence*: knowing one’s emotions, managing emotions, motivating oneself, recognizing emotions in others, and handling relationships. A common measure of Emotional Intelligence is EQ (emotional intelligence quotient), as gauged by a myriad of widely published EQ tests.

In the late 1990s, many AI and HCI researchers began to take the notion of emotion and emotional intelligence quite seriously. The Affective Computing Lab within the MIT Media Lab was founded by Rosalind Picard following the publishing of her 1997 book titled “Affective Computing,” in which she laid out the framework for building machines with *emotional intelligence* (Picard, 1997). Picard, along with many other researchers in this space, has built machines that can both detect, handle, understand and express emotions.

Before discussing the theories and applications of machines that are emotionally intelligent, it is important to first understand that emotion is an important aspect of intelligence. The evidence that necessitates this work comes from a few different fields. I will provide references to some of this evidence prior to moving into the work that has been completed in building machines that are emotionally intelligent.

Is “emotional intelligence” a contradiction in terms? One tradition in Western thought has viewed emotions as disorganized interruptions of mental activity, so potentially disruptive that they must be controlled.

Salovey and Mayer, 1990

As Salovey and Mayer express in the quote above, the common notion is that emotion is a hindrance to intelligent thought. Much work in the field of Affective Neuroscience has provided empirical evidence that this is not the case,

and indeed the opposite is true. Affective Neuroscience is the study of the processing of emotions within the human brain. Researchers in this field have shown that emotion plays a crucial role in problem solving and other cognitive tasks within the brain (Damasio, 1994). In addition to the empirical evidence that arises from Affective Neuroscience, many publications from popular psychology have backed this notion as well, arguing that emotional intelligence is critical to a person's success in many aspects of life (Gardner, 1993; Goleman, 1997).

Models of Emotion

Prior to understanding efforts in this space, one must have an understanding of the various models of emotion that are incorporated into systems. The choice of model is completely dependent on the task at hand; namely what dimensions of emotion can be gleaned from the available input signal, what model lends itself best to internal reasoning within a system, and what type of emotional expression the system aims to accomplish.

The simplest model is one of *valence* (positive or negative) and *intensity*, where sentiment is represented as a single score between -1 and +1, where -1 denotes the most intense negativity and +1 corresponds to the most intense positive score. A slightly more complex model adds the dimension of dominance (a scale from submissive to dominant). In this model, the intensity dimension is called "arousal" (a scale from calm to excited). This more complex model is commonly known as the VAD model, which stands for valence, arousal, and dominance (or PAD where valence is replaced by the synonym "pleasure") (Bradley, 1999; Mehrabian, 1996). This model is commonly used in measuring emotional reactions in humans as these dimensions lend themselves well to this task.

A more commonly known model is Ekman's six emotions model – happiness, sadness, anger, fear, surprise and disgust (Ekman 2003). This six dimensional model is intended to characterize emotional facial expressions and is typically used in systems that intend to express emotion in interaction with users. A mapping between the VAD and Ekman models exists in order to facilitate building systems that both detect and express emotion. For example, a low valence, high arousal, and low dominance VAD score maps to fear in the Ekman model, whereas low valence, high arousal and high dominance maps to anger (Liu, 2003).

Machines With Emotional Intelligence

With an established need for work in this area, I can now present the ways in which researchers are making efforts towards building emotionally intelligent machines. There

are various models/definitions of emotional intelligence, but they all boil down to the ability to connect with others by detecting, expressing, managing and understanding the emotions of oneself and others. Efforts in building machines that are emotionally intelligent center around a few key efforts: empowering the machine to detect emotion, enabling the machine to express emotion, and finally, embodying the machine in a virtual or physical way. Projects that incorporate all of these aspects also require the additional ability to handle and maintain an emotional interaction with a user, a large added complexity. In the sections that follow, I will provide examples of systems that approach these tasks – examples that are intended to expose students to work in the space of emotional intelligence.

Detecting Emotion

Work in the space of automated approaches to detecting emotion has focused on many different inputs including verbal cues, non-verbal cues including gestures and facial expressions, bodily signals such as skin conductivity as well as textual information. The end goal in building systems that are able to detect an emotional response from a user, is to handle/understand that response and act accordingly – a problem that is larger, and less understood than the problem of simply detecting the emotional responses/expression in the first place.

There are many modern research systems that can be used to exemplify this concept in a classroom setting, including systems that detect emotion in speech (Polzin, 1998; Yu 2001), in facial expressions and gestures (Gunes 2005), in bodily cues (Strauss, 2005), and in text (Pang, 2002; Turney 2003). While there is a wealth of examples of projects in this space, in my AI course, I typically introduce the notion of detecting emotion by presenting my own work in this space, detecting emotion in text (Owsley, 2006; Sood, 2007). This work fits well into the framework of an AI course as it uses many techniques that are traditionally taught in AI courses.

The system, call *RTS (Reasoning Through Search)*, uses a machine learning approach for the task of classifying a given piece of text within the valence/intensity model of emotion. Other similar systems have been built based on the other models of emotion discussed in the previous section – however, this system was limited to the seemingly simple task of ranking a selection of text on the scale from -1 to +1 (extremely negative to extremely positive). The system is trained on 106,000 movie and product reviews, where star ratings (from 1 to 5) served as truth data.

Given the task at hand, a purely statistical approach performs poorly because of the disparities between the emotional connotations of some words between domains/contexts. For example, when describing a politician, the term "cold" typically has a negative

connotation, whereas it is positive for a beverage to be “cold.” In order to build a “general” (not domain specific) emotional classifier, a case based approach is more appropriate to avoid the effects of averaging. The *RTS* system combines the benefits of the Naïve Bayes Model as well as Case Based Reasoning and various techniques from Information Retrieval. The end system classifies an unseen review on the valence/intensity model with 78% accuracy.

I have found the *RTS* system to be an engaging example to use in teaching machine learning via Naïve Bayes and Case Based Reasoning. The task of emotional classification is simple enough that students easily understand the challenges, yet complex enough to allow students to explore the tradeoffs of various machine learning approaches. By providing students with a training set of movie and product reviews, they are able build a Naïve Bayes classifier and learn the shortcomings of the model first hand within a medium sized programming assignment.

Expressing Emotion

Emotional expression within computer systems is typically focused on applications involving speech and/or facial expressions/gestures. Again, a plethora of work exists in this space, all of which would engage students in a classroom setting, including systems that attempt to automate gestures and expressions for an avatar (Breazeal, 1998; Breazeal, 2000), and those that enhance emotional expression through computer generated speech (Cahn, 1990). To introduce students to this concept, my work in the latter serves as a good example of machines that express emotion. This work is presented in a digital theater installation called *Buzz* (Sood, 2007), a system that has now moved online (www.buzz.com).

In building a team of digital actors for *Buzz*, I wanted to empower them to convey emotion in their computer-generated voices. Standard text-to-speech engines are rather flat, emotionless; they wouldn't make for a compelling performance. I chose to augment an off-the-shelf text-to-speech engine with a layer of emotional expression (Sood, 2007). This layer was informed by Cahn's work outlining the acoustic parameters of speech that change with different emotional states (Cahn 1990). The problem of *when* to express emotion in speech was solved by a simple mood classifier (trained on a set of mood labeled blog posts from livejournal.com). The actual transformations of the speech were done by an audio post-processing tool that altered the audio file of speech based on the evidence from Cahn's work. The end result is a system that actual conveys emotion (consistent with the content of what they are saying) in its voice.

Embodiment

Finally, the embodiment of a system facilitates a more personal connection between machine and user. People often attribute other human characteristics to a system when it perceives it as somewhat human/animal looking. This not only results in emotional connections, but it makes users more forgiving when the system makes a mistake. Many online systems in e-commerce, tutoring and training applications have recently begun embodiment as a way to engage/connect with users.

While there are many example systems in this space ranging from robotic seals to game based avatars, I find the most compelling example to be Kismet (Breazeal, 1998; Breazeal 2000), a robot created in the humanoid robotics group at MIT. The reason Kismet is such a great example is that it is an embodied system that detects, manages and expresses emotion in a social interaction with a human. While this is not the only system in this space, I think it is a compelling example to introduce students to the state of the art in all aspects of emotional intelligence.

Why teach emotional intelligence?

Artificial Intelligence is a field that is young and constantly changing. In order for students receive an education that enables them to make future contributions in AI, it is my opinion that Introductory Artificial Intelligence courses have to balance historical work with the current state of problems and approaches in the field of AI. Creating machines with emotional intelligence is emerging as a prominent area in Artificial Intelligence, and I feel that students could benefit from an introduction to this topic within the confines of an introductory artificial intelligence course. Additionally, the tasks involved in creating emotionally intelligent machines lend themselves well to AI techniques typically included in introductory courses.

In addition to being an engaging topic, building machines with emotional intelligence is a task that is inherently multidisciplinary. Many have hypothesized that multidisciplinary work can have wider appeal, bringing women and other underrepresented groups into the field of computer science (Margolis, 2002). “In particular, some studies indicate that students from underrepresented groups in computer science, such as women and minority students, who show early interest in computer science are more likely to pursue career paths in computer science when they are exposed to computational applications which demonstrate social relevance” (Tjaden, 2007). The concept of emotion and machines with emotion is a great example of an application that demonstrates “social relevance.”

Current artificial intelligence courses

In this paper, I have provided a survey of topics and resources on the emerging topic of building machines with emotional intelligence. Additionally, I have provided reasoning as to why this topic is important and should be included in introductory AI courses, but I have yet to address the issue of how to incorporate this topic in the framework of existing artificial intelligence courses. Through my experiences in two different courses, I have explored two methods of introducing this topic.

The first method involves exploring this topic towards the end of an introductory AI course. Russell and Norvig's text "Artificial Intelligence: A Modern Approach" (Russell, 1995) is widely used in introductory courses. Chapter 26 of this text addresses the philosophical backings of AI. When addressing the question "Can machines think?", we teach Turing and Searle. We now have the questions "Can machines feel?" and "Can machines understand emotion?" These questions are equally important in the philosophical foundations of the field and can be extended to discuss efforts in building machines with emotional intelligence.

The second method, which I have found more compelling, is to introduce the topic of emotional intelligence on the first day of class, when "defining AI." Throughout the semester, I tie in examples within the confines of AI techniques I'm teaching. For example, when teaching Machine Learning, we discuss ways to classify the emotional content of text; when teaching speech generation, I present Janet Cahn's work on emotional speech and show these findings were used in Kismet to make a more emotional voice (Cahn, 1990; Breazeal, 1998); when wrapping up the chapters on agent based work, we discuss the embodiment of agents and the effects/reactions from users.

Conclusion

The creation of machines that are empowered with emotional intelligence is a research area that is growing within the field of artificial intelligence. I have provided a survey of topics and references of research toward this end. My experiences introducing this topic in artificial intelligence courses have been positive and I have found it to be an engaging and important area of study.

References

M. M. Bradley and P. J. Lang. *Affective norms for English words (ANEW): Stimuli, instruction manual, and affective ratings*. Technical Report C-1, Center for Research in Psychophysiology, University of Florida, Gainesville, Florida, 1999.

Breazeal, C. (2000), "Sociable Machines: Expressive Social Exchange Between Humans and Robots". Sc.D. dissertation, Department of Electrical Engineering and Computer Science, MIT.

Breazeal(Ferrell), C. and Velasquez, J. (1998), "Toward Teaching a Robot 'Infant' using Emotive Communication Acts". In Proceedings of 1998 Simulation of Adaptive Behavior, workshop on Socially Situated Intelligence, Zurich Switzerland. 25-40.

Cahn, J. E. *The Generation of Affect in Synthesized Speech*. Journal of the American Voice I/O Society, 1990.

Damasio A.R.. *Descartes' Error: Emotion, Reason, and the Human Brain*. Grosset/Putnam, New York: 1994.

P. Ekman. *Emotions Revealed: Recognizing Faces and Feelings to Improve Communication and Emotional Life*. Henry Holt and Company, New York, NY: 2003.

Gardner, H. *Multiple Intelligences: New Horizons*. Basic Books, New York: 1993.

Goleman, D. *Emotional Intelligence: Why It Can Matter More Than IQ*. Bantam, New York: 1997.

Gunes, H., and Piccardi, M. *Fusing face and body gesture for machine recognition of emotions*. IEEE ROMAN Robot and Human Interactive Communication, 2005.

Kihlstrom, J., and Cantor, N. *Social Intelligence*. in R.J. Sternberg (Ed.), *Handbook of intelligence*, 2nd ed. (pp. 359-379). Cambridge, U.K.: Cambridge University Press, 2000.

LeDoux, J.E.. *Emotion: Clues from the Brain*. Annual Review of Psychology, January 1995, Vol. 46, Pages 209-235.

Liu, H., Lieberman, H., and Selker, T. *A model of textual affect sensing using real-world knowledge*. In Proceedings of the 8th international conference on Intelligent user interfaces, 2003.

Margolis, J., Fisher, A., *Unlocking the Clubhouse: Women in Computing*. MIT Press, 2002.

Mayer, J.D. & Salovey, P. 1993. The intelligence of emotional intelligence. *Intelligence*, 17, 433-442.

Mehrabian, A. *Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in Temperament*. *Current Psychology*, Vol. 14, No. 4. (21 December 1996), pp. 261-292.

Norman, D. *Emotional Design: Why we love (or hate) everyday things*. Basic Books, New York: 2004.

Owsley, Sara, Sood, S., Hammond, K. *Domain Specific Affective Classification of Documents*. AAAI Spring Symposia Computational Approaches to Analyzing Weblogs 2006.

B. Pang, L. Lee, and S. Vaithyanathan. *Thumbs up? sentiment classification using machine learning techniques*. In Proceedings of EMNLP, pages 79–86, 2002.

Picard, R.W.. *Affective Computing*. MIT Press, Cambridge, 1997.

Polzin, T., and Waibel, A.. *Detecting Emotions in Speech*. In Proceedings of the CMC, 1998.

Russell, S.J., Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice Hall, New Jersey: 1995.

Salovey, P. & Mayer, J.D. 1990. Emotional intelligence. *Imagination, Cognition, and Personality*, 9, 185-211.

Sood, S., Owsley, S., Hammond, K., and Birnbaum, L. *Reasoning Through Search: A Novel Approach to Sentiment Classification*. Northwestern University Tech Report Number NWU-EECS-07-05, 2007.

Sood, Sara Owsley. *Compelling Computation: Strategies for Mining the Interesting*. PhD Thesis, 2007.

Strauss, M., Reynolds, C., Hughes, S., Park, K., McDarby, G. and Picard, R.W. "The HandWave Bluetooth Skin Conductance Sensor," The 1st International Conference on Affective Computing and Intelligent Interaction, October 22-24, 2005, Beijing, China.

Thorndike, E.L. 1920. Intelligence and its use. *Harper's Magazine*, 140, 227-235.

Tjaden, Brian. *A multidisciplinary course in computational biology*. *Journal of Computing Sciences in Colleges*, 22(6):80-87, June 2007.

Turney, P.D., and Littman, M.L.. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)*, 21(4):315–346, 2003.

Yu, F., Chang, E., Xu, Y.Q., Shum H.Y.. *Emotion Detection from Speech to Enrich Multimedia Content*. In Proceedings of the Second IEEE Pacific Rim Conference on Multimedia: Advances in Multimedia Information Processing table of contents, p550 – 557, 2001.