YXM830-23J
Advance your independent learning

Candidate: Matthew Moran
Emotional AI in education: applications and implications
Tutor: Christine Abbott

Emotional AI in education: applications and implications

Introduction

Hello everyone.

My name is Matthew Moran, I am a student at The Open University studying the MSc in Technology Management. My presentation draws on an investigation undertaken during my studies, and is entitled:

'Emotional AI in education: applications and implications.'

Agenda

In this presentation, I will:

- Outline what we mean by 'emotional AI' and how it is defined and understood by researchers
- Explain briefly how emotional AI systems work computationally
- Show why this is significant for education, giving some examples of applications in education
- We will then give an overview of the current state of the art of emotional AI, and possible near-future trajectories
- We will outline some barriers to adoption of emotional AI in education
- And, finally, I will describe some implications for the theory and practice of innovation management.

The investigation was motivated by my awareness, as a practising educational technologist working in higher education, of a wave of research publications on emotional AI in education during the late 2010s. I was curious to understand the current state of research, and why, despite various benefits claimed by researchers, we have not seen wide diffusion of emotional AI in education. Emotional AI is still little known and understood in this community of edtech researchers and practitioners, so I hope you find the results of my investigation to be informative.

Definition

What is emotional AI?

'Emotional AI' (also 'emotion AI') is a term used by computer scientists to refer to systems with, in McStay's (2018) words, 'the capacity to see, read, listen, feel, classify and learn about [human] emotional life' computationally, using machine-learning and deep-learning techniques, so improving the quality of human–computer interaction.

Emotional AI is best understood as a development of affective computing. Affective computing was first defined by Picard (1997) as 'computing that relates to, arises from, or deliberately influences emotion or other affective phenomena'. Affective computing systems detect, recognise and simulate human affects, by analysing data from multiple modes of human expression, such as the face, voice, text, gesture and physiology. Affective computing is now a significant field of interdisciplinary research and development.

When we say 'emotional AI' we are really referring to applied, state-of-the-art affective computing systems that use advanced machine-learning and deep-learning models to detect, analyse, respond to, learn about, and predict human affects from multiple data sources.

How does it work?

At a high level, emotional AI systems consist of hardware and software designed to:

- collect emotion data from users (through audio, video and text inputs, and physiological and neuro-physiological signals)
- prepare, process and fuse the data
- categorise emotions through comparison with large datasets
- generate emotion-aware responses, and
- express these responses through system outputs.

The earliest systems of the 1990s and 2000s were unimodal; in other words, they processed data from a single mode (text, audio or video) using machine-learning techniques. Contemporary systems are multimodal; that is, they process, sequence and fuse data from multiple modes simultaneously, using machine-learning and deep-learning techniques, including recurrent neural networks such as long short-term memory.

While early systems were engineered to detect your emotional state from a still picture of your face at a given moment in time, current systems combine and triangulate data from face, voice, gestures, skin conductance, and a range of other signals continuously over longer periods of time, and therefore discern your emotional state precisely, and respond naturally and appropriately.

So what?

Why is this significant?

The influence of emotions on learning is widely researched. Our emotions affect our motivation to learn, our attention, our ability to self-regulate, and memory processes. There is a broad consensus among researchers in the fields of cognitive neuroscience and education that emotions influence cognitive and social processes in learning. Positive emotions have been demonstrated to promote states beneficial to study, and the student's ability to plan, monitor and evaluate their learning. Negative emotions have been shown to impair the student's motivation, attention, comprehension, information recall and self-efficacy.

Imagine for a moment that I am not a real person, but a Virtual Human generated by a state-of-the-art emotional AI. At every moment of this presentation, I am detecting and analysing how you express your affective state in your face, voice, text, gestures, physiologically and even neuro-physiologically. And based on my continuous analysis of you, I am adapting my presentation to enhance your attention, engagement, and your memory of what I am saying. In doing so, I am enabling you to become a more effective and successful learner. And I can do this for every individual in the audience simultaneously.

We know that emotions are significant precursors of learning performance and academic achievement. And so it is surprising, at least to me, that there is not greater awareness of emotional AI among educational technologists; and surprising, too, that we have not seen wider diffusion and adoption of affective computing within digital learning environments.

Emotional AI in education

That said, there is now a body of examples of mostly experimental applications of affective computing in learning environments and systems. I will now give some representative examples.

These take several forms. We find examples of affective computing applications in:

- social robots
- serious games
- intelligent tutoring systems
- immersive virtual environments, and
- Virtual Humans.

Let's look at these briefly.

Social robots

Dragonbot, Tofu and Tega are social assistive robots developed by researchers at MIT for experimental use in preschools. They combine robotics with five degrees of freedom, and Android smartphone 'brains'. The robots communicate and interact with children, detect and analyse facial expressions, and react with emotionally appropriate responses.

Tega was trialled in a Boston preschool in 2017, with children learning Spanish. Tega acts as a learning buddy, giving pre-recorded instructions, hints and encouragement, while the software analyses children's emotional states, and generates appropriate responses, such as happy sounds.

Researchers, led by Breazeal, tested the robots with and without the affective layer. The published results indicate that children were more engaged, learned more Spanish words, and had significantly higher rates of recall when learning with the affective layer than was the case in children who learned without it.

Serious games: Prime Climb

Emotion-aware social robots in learning are quite rare. Much more common are serious games for education and training, which researchers have adapted with affective techniques to enhance the learning experience.

For example, Conati and Maclaren (2009) developed an early affect-aware agent, and integrated it in Prime Climb, an educational computer game, to detect and identify multiple, changing emotions over time. This agent combines two models:

- a theoretical diagnostic model, the OCC model (this recognises 22 emotion types, which, the model holds, are provoked by one's appraisal of a situation and the extent to which we can achieve our goals)
- and a predictive model based on a dynamic Bayesian network.

During gameplay, the agent infers probabilistic relationships between the learner's goals, the learner's actions and their outcomes, and the agent's actions and their consequences for the learner. From this, the agent can detect learners' emotions, and predict why these emotions are provoked, and therefore adjust and respond to learners more effectively, so enhancing the learning experience.

Serious games: Maritime City

Some more recent examples of affective techniques in serious games are Maritime City, and Pandora.

Maritime City is currently used in social work training in the UK. It provides learners with exposure to emotionally charged and challenging scenarios, and enables them to develop skills in interpreting verbal and non-verbal communication, and emotional self-regulation. The National Occupational Standards for Social Work require trainees to develop objective, non-judgmental assessments of clients' emotions. In Maritime City, learners interact with virtual clients. They must evaluate and respond to the information they gain from their interactions, both factual information and the emotional representation. Learners can pause gameplay at any time to analyse their interactions, and gain feedback from peers and instructors, and they can repeat exercises to hone their skills, something that would be hard to achieve with face-to-face placements.

Serious games: Pandora

Pandora is an EU project for training in crisis management. It simulates the very challenging emotional and cognitive experience of crisis managers when dealing with major emergencies, including information overload, emotional intensity, and the pressure of making decisions under pressure.

Pandora places the user in a multimedia simulation of a major emergency control room. We are confronted with video and imagery from the scene, reports from eyewitnesses, and social media postings including nuisance and malicious messaging. Pandora increases stress on the user through the volume of information presented, by demanding decisions be made quickly, and by breaking up or adding noise to incoming information. Pandora aims to build trainees' capability to make effective and appropriate decisions under extreme pressure, and enhance self-efficacy and leadership capabilities.

Intelligent tutoring systems

Another learning environment where we have seen applications of affective computing is intelligent tutoring systems. Probably the most extensively published example here is Java Sensei, an intelligent environment for learning the Java programming language, the development of which is described in detail by Cabada, López and Escalante (2023). Java Sensei uses advanced deep-learning architectures based on convolutional neural networks. Cabada and colleagues' work represents major advances on earlier work with machine-learning techniques such as Bayesian networks.

Java Sensei is built around a modular, client–server architecture, comprising:

- affect recognition module
- fuzzy logic module for pedagogical decision-making
- recommendation module that analyses learner preferences
- and pedagogical agent module that facilitates communication between the system and the learner, directly or via a human educator.

Java Sensei is an example of a state-of-the-art affective multimodal system with capability to fuse data from multiple sources sequentially, and make inferences from one data source to another. Moreover, it is an end-to-end learning system; in other words, it processes data from modes for which it has not been trained. For example, it can learn to process audio signals from labelled datasets of other signals, such as images, using generative adversarial networks (a machine learning framework in which two neural networks compete to make the most accurate predictions).

Immersive virtual environments

Developers of educational and training simulations are now using affective computing technologies in immersive virtual environments. One area where there is currently intense research is in use of affective computing to enhance user sense of agency.

Our sense of agency affects our capacity to regulate and supervise our learning. A learner's ability to regulate their cognitive, affective and behavioural processes while interacting with their learning environment is known to be a precursor to academic success. However, this sense of agency is often lost in immersive virtual environments. Research (including by Aoyagi and colleagues (2021)) demonstrates that increasing the user's sense of agency significantly improves the effectiveness of immersive learning experiences.

Ortiz and Elizondo (2023) have proposed a model for enhancing sense of agency in immersive environments using affective technologies. In this model, the system detects user emotions via multiple modes, including physiological signals from sensors in headsets, and these are mirrored in the user's avatar, so transmitting their emotions and intentions to others with whom they interact virtually. Objects and non-player characters in the virtual environment react and respond appropriately, enhancing the user's sense of bodily and external agency and, therefore, the efficacy of learning in virtual worlds.

Virtual Humans

Finally, we are now seeing Virtual Humans designed using affective computing techniques. Llanes-Jurado and colleagues (2024) have developed voice-based conversational Virtual Humans based on large language models. These are able to have open, emotional conversations in real time that seem fluid and natural. This is the first comprehensive Virtual Human system engineered for emotion elicitation. It comprises 20 virtual humans with five emotional states (neutral, anger, happiness, sadness, and relaxed). Emotions are expressed through the Virtual Humans' behaviour, gestures, facial expressions, and language. The researchers believe Virtual Humans will soon be used in teaching and learning, and indeed, holographic lecturers are already being used by Prof Gary Burnett and colleagues at Loughborough University in the UK.

The state of the art

These examples of emotional AI applications in education are chosen to provide a brief history of the development of affective technologies. In each case, researchers report positive benefits for learners and educators, including improved learner motivation, attention, information recall, self-efficacy, self-concept and mastery orientation. And these examples hint at possible near-future developments.

Current state-of-the-art emotional AI:

- is multimodal – systems can fuse data from multiple sources synchronously, and infer emotions from one data source to another
- is end-to-end – systems can process ’raw’ (unlabelled) data using general unsupervised representation learning (a machine learning technique for identifying patterns in large volumes of data)
- combines machine-learning and deep-learning techniques – while early developments employed machine learning extensively, these tended to require hand-crafting of data features, making reuse difficult; consequently, there is intensive research and development of better-performing deep-learning techniques for feature representation, extraction and sequencing, and cross-domain learning – Wang and colleagues (2022) have published an exhaustive evaluation of the performance of machine- and deep-learning techniques
- affective computing technologies are converging with generative AI and large multimodal language models, a new class of LLM that can accept inputs in audio, video and image formats, effectively giving emotional AIs eyes and ears – Yin and colleagues (2023) have recently published a systematic review of these
- and these technologies are converging with enterprise and consumer hardware such as wearable AI devices, smart glasses and extended-reality headsets.

These most recent developments appear to overcome a long-standing barrier to adoption of emotional AI, namely the invasiveness of input sensors. Emotional AI is now wearable.

Consequently, we are seeing accelerating adoption of emotional AI in domains such as ergonomics and human factors, particularly in the automotive industry, as well as entertainment and marketing.

But not in education.

Why?

Barriers to adoption

Firstly, and as may be obvious, emotional AI has yet to be productised – all the examples I’ve shown are experimental innovations, not marketable products.

Secondly, these systems require advanced hardware and software engineering skills, and expensive compute power, and these resources are simply not available in many institutions.

Throughout my investigation, I've observed among educational technologists a general lack of interest in the role of emotions in learning. As a profession, we prioritise the cognitive and social, and we tend to overlook the emotional.

That said, there is a vibrant community of educational researchers undertaking theory-driven research on the role of emotions on learning in technology-rich environments, and many of these researchers are using digital platforms to gather data. But these researchers are not talking to developers, and we are not talking to them.

Consequently, the technologies that predominate, particularly in higher education, do not lend themselves to integration of affective technologies. The virtual learning environment, and dominant commercial learning platforms, prioritise cognitive and social learning, but their content and experience is non-adaptive, and they cannot respond to the affective state of individual learners as serious games and intelligent tutoring systems do.

And on top of the economic, technical, professional, organisational and theoretical barriers, we now have a significant legal barrier to adoption imposed by the EU AI Act, which bans the use of emotion-detection technology in educational institutions in the bloc.

Is this the end of the road for emotional AI in education?

It's unclear. The convergence of emotional and generative AI with consumer devices like smart glasses may make the EU ban hard to enforce. It's going to be interesting to follow the Act's impact on emotional AI in education.

Implications

I am going to close with some brief reflections on emotional AI as a case study for innovation management.

Emotional AI is problematic when considered in terms of established theories and models of innovation management and technology diffusion.

Firstly, categorisation. What kind of innovation is it?

Emotional AI has been different types of innovation at different times. At its origin in the 1990s, this was breakthrough innovation, a major (but not radical) advance in the technological paradigm.

For two decades, emotional AI research and development proceeded incrementally.

And now it appears to be becoming a radical innovation, which, through convergence with other technologies, presents revolutionary novelty and effects.

The implication here is that technologies may move between analytical categories during diffusion, sometimes over extended time periods.

In some models of innovation, such as the well-known SSID model (proposed by Tidd and Besant (2021)), emotional AI is not an innovation at all, since it hasn't found wide-ranging commercial uptake. But do these models of innovation impose a too-rigid distinction between R&D on one hand, and commercial success on the other, a distinction that doesn't correspond to reality?

Emotional AI, like generative AI and spatial computing, appears to occupy a third dimension, somewhere between lab test and marketable product. And increasingly we are seeing speculative products being launched into this space, products like the Humane AI pin, or the Rabbit R1 portable AI assistant.

Furthermore, textbook models of technology diffusion require that a new technology must find market adoption before it moves to new application contexts. Again, emotional AI defies this requirement, passing from serious games to Virtual Humans, and from education to marketing to automotive, before becoming commercially viable. The same models require that fast-evolving 'niche' R&D either breaks through into the socio-technical regime quickly and disruptively, or it fails and dies away. Emotional AI has done neither.

Conventional theories and models of technology innovation place too heavy an emphasis on technology as product. Emotional AI is better understood as technology as possibility. And its possibilities remain mostly unavailable to education.

Conclusions

In this presentation we have:

- learned what emotional AI is, how it works, and why it is significant
- examined some applications in education contexts
- summarised the current state of the art
- reviewed some barriers to adoption in education, and
- considered some implications for theory and practice of innovation management.

I hope you have found this interesting and informative.

Thank you for listening.

References

Aoyagi, K., Wen, W., An, Q., Hamasaki, S., Yamakawa, H., Tamura, Y., Yamashita, A. and Asama, H. (2021) 'Modified sensory feedback enhances the sense of agency during continuous body movements in virtual reality', *Scientific Reports*, 11(1), pp. 2553–2553, article 2553. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/j6vapu/cdi_doaj_primary_oai_doaj_org_article_9570ae86ff274d3c81c46c2c8f6cce89 (Accessed 15 May 2024).

Cabada, R.Z., López, H.M.C. and Escalante, H.J. (2023) *Multimodal Affective Computing: Technologies and Applications in Learning Environments*. Cham: Springer International Publishing. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/la9sg5/alma9953061632502316 (Accessed 15 May 2024).

Conati, C. and Maclaren, H. (2009) 'Empirically building and evaluating a probabilistic model of user affect', *User Modeling and User-Adapted Interaction*, 19(3), pp. 267–303. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/j6vapu/cdi_webofscience_primary_000267665600004 (Accessed 8 January 2024).

Kory-Westlund, J.M. and Breazeal, C. (2019) 'Exploring the effects of a social robot's speech entrainment and backstory on young children's emotion, rapport, relationship, and learning', *Frontiers in Robotics and AI*, 6, p. 54. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/j6vapu/cdi_doaj_primary_oai_doaj_org_article_3775cd63417b48d4be3608ec48131eef (Accessed 15 May 2024).

Llanes-Jurado, J., Gómez-Zaragozá, L., Minissi, M.E., Alcañiz, M. and Marín-Morales, J. (2024) 'Developing conversational Virtual Humans for social emotion elicitation based on large language models', *Expert Systems with Applications*, 246, p.123261. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/j6vapu/cdi_crossref_primary_10_1016_j_eswa_2024_123261 (Accessed 15 May 2024).

McStay, A. (2018) *Emotional AI: The Rise of Empathic Media*. London, New Delhi, Singapore: SAGE Publications Ltd. Available at: https://doi.org/10.4135/9781526451293 (Accessed 30 October 2023).

Ortiz, A. and Elizondo, S. (2023) 'Design of an immersive virtual reality framework to enhance the sense of agency using affective computing technologies', *Applied Sciences*, 13(24), p. 13322. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/j6vapu/cdi_doaj_primary_oai_doaj_org_article_0d4d15a9adc94a008f279d62bdb031cb (Accessed 15 May 2024).

Picard, R. (1997) *Affective Computing*. Cambridge (Mass.): The MIT Press. Available at: https://doi.org/10.7551/mitpress/1140.001.0001 (Accessed 15 May 2024).

Tidd, J. and Bessant, J. (2021) *Managing Innovation: Integrating Technological, Market and Organizational Change*. Hoboken (NJ): Wiley. Available at: https://library-search.open.ac.uk/permalink/44OPN_INST/la9sg5/alma9952607760902316 (Accessed 15 May 2024).

Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T. and Chen, E. (2023) 'A survey on multimodal large language models', *arXiv*, 2306, 13549. Available at: https://doi.org/10.48550/arXiv.2306.13549 (Accessed 15 May 2024).