

The age of artificial emotional intelligence

Dagmar Schuller, Björn Schuller

Angaben zur Veröffentlichung / Publication details:

Schuller, Dagmar, and Björn Schuller. 2018. "The age of artificial emotional intelligence." *Computer* 51 (9): 38–46. <https://doi.org/10.1109/mc.2018.3620963>.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



The Age of Artificial Emotional Intelligence

Dagmar Schuller, audEERING

Björn W. Schuller, Imperial College London, University of Augsburg, and audEERING

Science fiction often portrays future AI technology as having sophisticated emotional intelligence skills to the degree where technology can develop compassion. But where are we today? The authors provide insight into artificial emotional intelligence (AEI) and present three major areas of emotion—recognition, generation, and augmentation—needed to reach a new emotionally intelligent epoch of AI.

Humans can be quite emotional about technology and computing systems, especially in their moments of failure. A key scene in the 1999 film *Office Space* shows three professional programmers violently destroying a copy machine that annoyed them for too long. Certainly, most of us have experienced an emotional moment when interacting with computing devices, such as when facing a frozen or blue screen. But will it soon make a difference whether we praise or blame our devices? Will emotion in human-computer interaction (HCI) shift from unidirectional to bidirectional?

When it comes to artificial intelligence (AI), emotion and emotional intelligence are not usually the first things that come to mind. Rather, some of the most far-spread applications of AI we presently encounter in everyday life are found in computer perception and natural language processing. Next, one tends to think of abilities such as knowledge representation (learning, planning, reasoning, and problem solving) or motion and manipulation. AI is usually further thought of as aiming for human-like or even beyond human intelligence. But creativity is still largely lacking in AI, and emotional and social intelligence will also need to be added at scale.

However, artificial emotional intelligence (AEI) found its place long ago in TV and film. Think of the HAL 9000 computer in *2001: A Space Odyssey* (1968), which showed the first signs of socioemotional intelligence, and the highly emotional personal AI device Samantha in 2013's *Her*. Another early example is the AI-packed car KITT, who detects in the series pilot of *Knight Rider* (1982) that his driver is in “a slightly irritable mood caused by fatigue.” AI also tends to fall in love in the movies—in *Electric Dreams* (1984), a personal computer and its user fight for the love of the same woman, and in *Wall-E* (2008), the eponymous robot falls in love with another robot. The list of examples is endless: *Short Circuit* (1986); *Star Trek: Generations* (1994); *The Iron Giant* (1999); *I, Robot* (2004); *The Hitchhiker's Guide to the Galaxy* (2005); *Iron Man* (2008); *Moon* (2009); *Ex-Machina* (2014); and *Big Hero 6* (2014), to name a few.

Clearly, these envisioned human-computer dialogs are very different from interactions with today's actual AI such as Alexa, Cortana, Siri, or even non-consumer grade AI such as IBM's Watson, which largely lacks emotional intelligence. This fascination with AEI has led to the emergence of fields such as affective computing,¹ social and behavioral computing,² and emotion-augmented machine learning.³ In this article, we provide insight into the principles of mammalian and artificial emotion and the modeling and history of emotional intelligence, as well as the current state of play in recognizing, generating, and using emotion principles through “emotion augmentation” of current AI systems. We also look at the questions of whether AI will actually have emotions, what we can expect for the future of AI

enhanced by emotions, and the immediate steps to get there.

EMOTIONAL AND SOCIAL INTELLIGENCE

Emotional intelligence—a term first used in a 1964 paper by Michael Belbin—is largely defined as intelligence that is marked by the abilities to recognize emotions (both of others and oneself), generate and adapt emotions, and apply emotional information in goal accomplishment and problem solving. A precondition of these abilities is the capability to differentiate between different emotions.

But why does emotional intelligence play a key role in mammalian lives, and can it be assumed to do so in the next generations of AI? From today's computing point of view, one could be tempted to consider human emotion as “noise” in optimal communication or planning and decision making. However, emotion forms a crucial part of intelligent behavior and plays a key role in a range of cognitive, perceptive, and bodily processes, according to psychological and neuroscientific findings.^{4,5}

Throughout the evolution of mankind, emotion has helped humans survive—one example is the “fight or flight” response. Emotion is our main drive of motivation;⁵ is directly connected to our memory and learning systems; and influences our associations, abstractions, intuition, and reasoning as it helps us move from exploration to exploitation. Being broadly considered as adaptive responses in modern appraisal theory,⁶ emotions also serve to evaluate our environment and monitor our wellbeing, informing ourselves of our current state. As such, emotions fulfill the role of reward and punishment in reinforced learning

and guide our attention, helping with decision making.^{4,5} Finally, emotions play a key role in communication, which is still broadly neglected in spoken dialog systems or general human-computer or human-robot interaction.

HUMAN EMOTION

How do humans experience, process, and apply emotion? Neuroimaging studies have shed significant light on the structural processing of human and general mammalian emotion. Emotion is a multifaceted and complex phenomenon, but one structure consistently associated with emotion is the amygdala. Studies have shown that the amygdala is highly connected within the brain, sharing inward and outward projection with components of the limbic system and others directly and indirectly.³ Thus, the amygdala (and emotion) is involved in multiple neuronal processes, such as the processing of raw (thalamus), object level (cortex), and contextual (hippocampus) data. The amygdala is considered a unitary system in this article, although it is of complex structure and can also influence other behaviors such as motor behavior and automatic responses via other systems.

The amygdala is thought to be able to evaluate environmental stimuli, shifting focus toward emotionally associated features. The formation of these emotional associations is called emotional learning. The amygdala is likely responsible for storing stimuli-emotional response patterns, and has also been shown to participate in the formation of declarative memory (memories that can be consciously remembered), providing “emotional coloration.” In addition, emotional arousal provokes secretion of the neuromodulator norepinephrine, which enhances

memory and learning. The evidence of the amygdala's (and thus, emotion's) role in attention and memory supports the idea that emotion is critical to evaluating situations and retrieving the most important stimuli to survive in an environment (such as the need to eat). Overall, research shows ample evidence of the amygdala (and with it, emotion) as a crucial centerpiece in the neuroeconomic decision-making process.

MODELING EMOTION

The optimal representation of emotion is discussed in the psychological literature. In technical applications, however, few models have found broad usage. In early days, categorical approaches prevailed such as Paul Ekman's "Big 6" discrete emotion categories (anger, disgust, fear, happiness, sadness, and surprise).⁷ Yet, such categories often oversimplify the subtle nature of real-world emotion and tend to be too limited in coverage.⁸ Continuous modeling by dimensions is gaining momentum. The most commonly applied dimensions are arousal, valence, and dominance.⁷ Emotion classes can be mapped to areas in the space spanned by these dimensions. An example is the emotion of fear being marked by high arousal, negative valence, and low dominance. The use of these two approaches—categories or continuous dimensions—is particularly popular in emotion recognition and generation. In reinforcement learning, appraisal-based approaches⁶ are far spread. Additional and mixed models exist, such as turning categories into dimensions, or tagging-based models, where more than one category might be present at the same time.

Another important aspect in modeling is temporal evolution. This can

be challenging to represent, as different modalities might operate on different time levels. For example, raw audio has a different sampling rate (usually above 8 kHz) than most physiological parameters, such as an electrocardiogram signal (usually below .5 kHz), or video or depth-information frames (usually below 100 Hz). Similarly, feature sampling rates largely differ (audio is often around 100 Hz, whereas linguistic analysis based on feature vectors are typically sampled at less than 1 Hz). In cross-modal modeling, one might therefore encounter compromises, and a typical emotion-update frequency in time-continuous modeling—which is gaining momentum⁷—is around .1 Hz to 1 Hz.

Finally, the type of emotion needs to be modeled. For human emotion recognition by computers, emotion as perceived by others is usually assessed in addition to self-assessed "felt" inner emotion. Further aspects can be modeled, such as the degree of acting, regulating, or suppressing an emotion; the degree of intentionality of acting, regulating, or suppressing an emotion; the degree of prototypicality of the emotion; or the degree of discrepancy. This could also affect AI's "inner" and "outer" emotion—what the AI "feels" versus how it is perceived to be "feeling." AI could then also, for example, suppress its emotion.

HISTORY OF COMPUTERS AND EMOTION

The term "affective computing" was coined in 1995 by MIT Media Lab's Rosalind W. Picard in her seminal book of the same name.¹ However, the concept dates back slightly earlier—for example, the first patent for automatic speech emotion recognition was filed in 1978 by John D. Williamson.⁹

In 1989, Janet E. Cahn—also from the MIT Media Lab—wrote about "the generation of affect in synthesized speech,"¹⁰ and in 1992, Hiroshi Kobayashi and Fumio Hara released their work on neural networks automatically recognizing the six basic facial expressions.¹¹

As for work on rendering emotion in the face, the European SEMAINE project (which ended in 2010) provided the first real-time system able to recognize user emotion and generate adapted agent output in a 2D audiovisual input-output chain for emotionally intelligent dialogs. Similar projects include University of Southern California's SimSensei (which began in 2011) and ARIA-VALUSPA (Artificial Retrieval of Information Assistants—Virtual Agents with Linguistic Understanding, Social Skills, and Personalized Aspects), which began in 2015.

The Human-Machine Interaction Network on Emotion (HUMAINE) was created in 2004 and ran until 2007, when it moved into a non-funded network. It finally became the UK-based Association for the Advancement of Affective Computing (AAAC) in 2014 (when co-author Björn Schuller was president of AAAC). In 2005, the first International Conference on Affective Computing and Intelligent Interaction (ACII) was held, and it has run biannually since. In 2010, *IEEE Transactions on Affective Computing* was launched, and it remains the field's main journal. The first open competition event in the field was the Interspeech 2009 Emotion Recognition Challenge, based on voice acoustics and linguistic content analysis and initiated and co-organized by Björn Schuller, who enriched this concept in 2011 with the first ever Audio/Visual Emotion Challenge (AVEC). Several

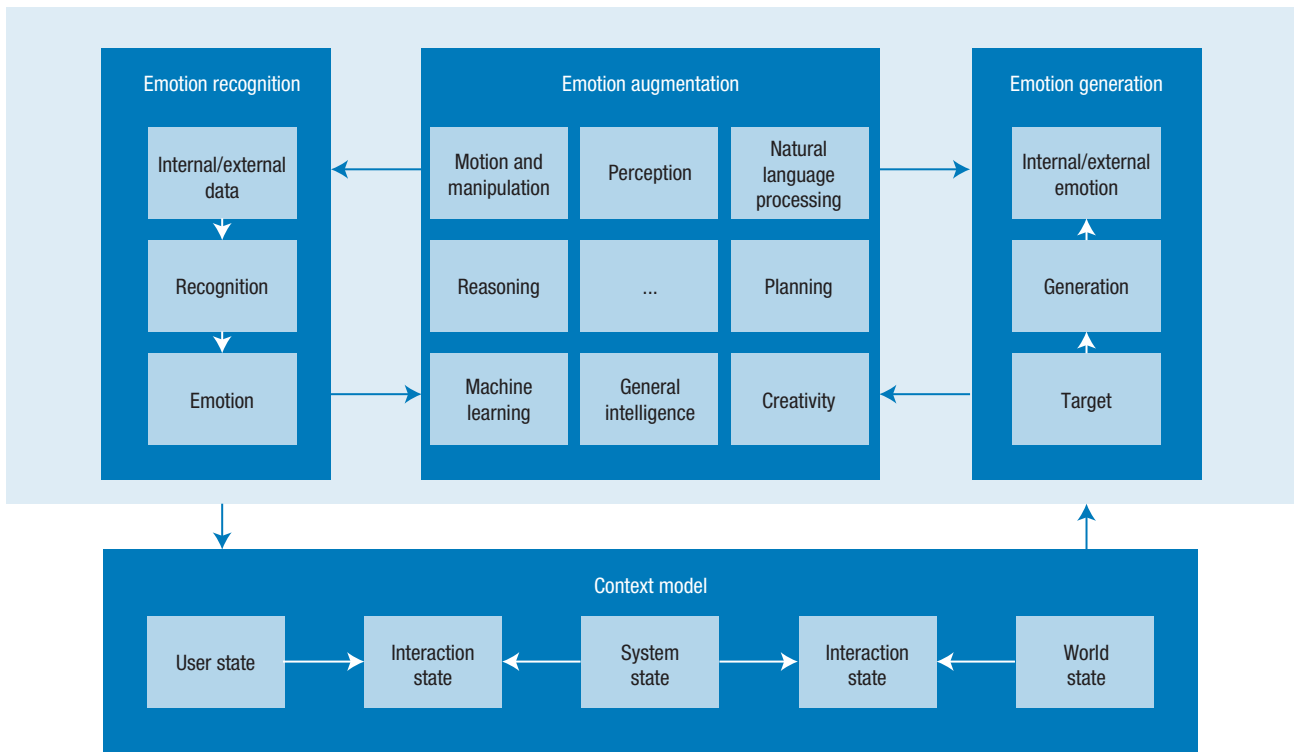


FIGURE 1. Main areas for the integration of emotional intelligence into artificial intelligence (AI).

similar challenges have appeared, including EmotiW, FERA, MEC, and the OMG Emotion Challenge, in addition to events for text analysis related to sentiment and emotion. The first physiology-based challenge was held in 2015 by AVEC.

Commercial start-ups focused on human emotion recognition include Affectiva, which was founded in 2009 and is physiology- and video-focused; audEERING, which was founded in 2012 and is audio-focused; and RealEyes, which was founded in 2007 and is video-focused. These have turned into million-dollar organizations and find a growing number of peers in a rapidly growing market. Consumer products, however, are still sparse and largely unknown by the public.

THREE BUILDING BLOCKS OF AEI

To date, AEI research has largely focused on automatic (human) emotion recognition and emotion generation for conversational agents and robots. Emotion augmentation of learning algorithms and dialog management has been attempted, albeit on a much smaller scale. Figure 1 shows emotion recognition, emotion generation, and emotion augmentation, which can be seen as the major building blocks for AEI.

Emotion recognition

The recognition of emotion by computing systems has more than two decades of history primarily focusing on the recognition of human emotions.

Emotion has grown to a mainstream topic in music, sound, images, video, and text. The prevailing modalities for emotion analysis are acoustic speech and spoken (or written) linguistic content; facial expression; body posture and movement such as gait; and physiological measurement such as heart rate, skin conductance, or even brain activity. Emotion in haptic interaction has also been researched.

Earlier emotion-recognition systems were marked by a variety of expertly crafted features such as those extracted by openSMILE (audio) or OpenCV (video) at varying sampling frequencies.⁷ As early as 2008, there was already a wide variety of machine-learning algorithms, including support vector machines and

kernel machines, hidden Markov models and more general graphical models, and neural networks such as (deep) recurrent neural networks with long short-term memory.

Today's approaches are increasingly focusing on deep end-to-end learning, such as the End2You toolkit (which allows for recognition from raw data) or shallow representations such as spectral transformations across these modalities at state-of-the-art performance as measured. A reappearing observation is the complementarity of the modalities: voice acoustics are known to carry information on arousal and dominance—facial expression

synthesis of emotional speech and facial expression, which dates back almost three decades. Further examples of emotion-dependent generation include text and haptic feedback. However, as opposed to analysis, synthesis approaches are traditionally more rule-based and less data-trained.⁷ These primarily focus on emotional speech synthesis such as MARY Text-to-Speech (MaryTTS); visual agent rendering including emotion-driven facial expression, body posture, and movement such as the Greta engine; or emotional text production. However, current tendencies steer toward increasing the use of deep learning,

by the principles of emotion. However, similar to emotion generation, such dialog management is currently mostly rule-based. In addition, one finds a number of mostly unidirectional emotion input or output only examples. For example, user emotion is partially already employed in some forms of system-state adaptation, such as in video games.

Most additional examples of emotion augmentation exist in the context of emotion-augmented machine learning (EML),² which aims at exploiting emotion for efficient learning as a bio-inspired principle. Likewise, including the principles of the amygdala in AI¹² would follow in the footsteps of artificial neural networks that have been inspired by mammalian neural networks, or (more specifically) convolutional neural networks inspired by the mammalian visual cortex. In fact, traditional machine-learning algorithms are usually designed to minimize an error function, and emotion as a guideline is rarely considered.

Even though affective computing has yet to provide deeper insight into how to model and generate artificial emotion, it seems obvious that there is huge potential in the integration of such principles in machine learning to increase the efficiency of learning—for example, through faster convergence or lower computational cost by emotional intuition for pruning or similar emotion-driven learning guidance. At the same time, better overall solutions might be found, and confidence measurement in the result of a learning algorithm can equally be informed from its emotional state.

One can roughly structure the approaches for embedding emotion principles into machine learning or

**THERE IS HUGE POTENTIAL IN THE
INTEGRATION OF AFFECTIVE COMPUTING
PRINCIPLES IN MACHINE LEARNING TO
INCREASE THE EFFICIENCY OF LEARNING.**

and spoken words are strongly related to valence. Human parity levels are partially met or about to be met—even under “in the wild” conditions⁸—according to the benchmarks and competitions in the fields mentioned above.^{7,8} Note that the emotion to be recognized is usually external from an AI's perspective (such as the emotions displayed by humans and animals), as AI's “internal emotion” is not yet very complex.

Emotion generation

Similar to the analysis side in terms of emotion recognition, the synthesis side of emotion generation has a longer tradition—particularly for the

such as exploiting WaveNet in emotional speech synthesis. This largely targets “external” emotion, which can be observed from the outside, rather than “inner” AI or system emotion.

Emotion augmentation

Compared to emotion recognition and generation, emotion augmentation of AI (such as applying emotion in planning, reasoning, or more general goal achievement) has rarely been attempted in the literature. The bidirectional emotion input/output platforms SEMAINE and ARIA-VALUSPA (and others for emotion-augmented human-computer dialog) are examples of dialog management enriched

even more general AI into four larger classes. These are discussed and exemplified below to provide deeper insight into this crucial realization of today's emotion augmentation of AI. Note that many more examples could have been chosen for each case.

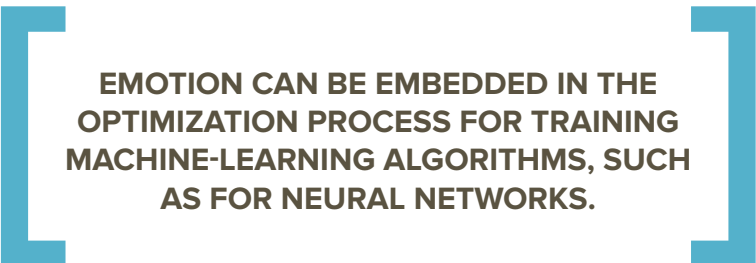
Optimization. It has repeatedly been shown that emotion can be embedded in the optimization process for training machine-learning algorithms, such as for neural networks. For example, Adnan Khashman showed that emotion can be beneficially integrated into a feed-forward neural network.¹³ To this end, he formalized anxiety and confidence and introduced an emotional backpropagation-learning algorithm that is enriched by an emotional bias found in hidden and output layers. He chose to raise the anxiety when novel patterns are presented, relating the network's output to the expected recognition error. Heightened anxiety leads to an enhanced inclusion of the latest error, whereas higher confidence emphasizes previous updates. A face-recognition example showcases a higher recognition rate and shorter execution time based on such emotion-augmented training compared with conventional backpropagation. One can relate the principles to today's attention ideas in deep learning, making it clear that different views on algorithmic implementations can be found. In later work, Khashman allowed for further factors in the input signal or "stimulus" to influence anxiety. Based on the visual cortex that has separate pathways for dorsal or cognitive and ventral or affective information streams, he further improved matters with the introduction of the "DuoNeurons" buildup of a cognitive neuron processing local features and

an emotional neuron (only) for global patterns.

Related efforts were made by Yimin Yang and colleagues,¹⁴ who circumvent the risk of getting stuck in local optima during learning by artificial "feelings" and emotions that control changes from exploration to exploitation. A hormone system further feeds back into the activations of the feelings. The emotions considered are reminiscent of Khashman's, as confidence is regulated (anxiety itself depends on diminishing reward). In addition, fear follows increased anxiety levels to control the strategy of exploration versus exploitation and

conventional multi-layer perceptrons by overcoming plateaus in learning.

Reinforcement learning. The concept of emotion augmentation has been successfully applied in reinforcement learning.^{3,16} An example is agents learning optimal state-action mapping for maximal reward based on principles of the appraisal theory. In this case, appraisals relate to the current state to alter the reward. The crucial factor is considering the common extrinsic motivation in reinforcement learning in addition to intrinsic motivation (in other words, internal or emotion-based reward). It could



**EMOTION CAN BE EMBEDDED IN THE
OPTIMIZATION PROCESS FOR TRAINING
MACHINE-LEARNING ALGORITHMS, SUCH
AS FOR NEURAL NETWORKS.**

"warmth" to control the learning termination based on the iteration number and level of fear. The model further includes the principle of a dominant emotion. Similarly, the learning was observed to converge faster with emotional enhancement.

Other similar approaches include EMANN (Emotional Artificial Neural Network), which is based on a multi-layer perceptron that includes nodes that receive and emit "hormones."¹⁵ The authors' idea is to have the hormone level control the activation threshold, the summation process of the weights, and ultimately the output of a neuron and the output of a given node. EMANN was able to outperform

be shown that appraisal-based emotion incorporation could lead to faster learning compared to conventional agent motivation, as reward can be provided constantly and not only after complete runs. The principle of longer-term emotional states ("moods") can further benefit this effect if there are potential actions without emotional response.³

Other forms of integrating emotional concepts into reinforcement learning are based on the principles of homeostasis and drives by equilibrium as an attraction point for learning agents. Homeostasis is thereby an equilibrium state of an agent's homeostatic variables, including primary

and potentially secondary emotions. Drives experienced by the agent will lead to actions that influence these variables that can lead to a notion of reward.

Emotion augmentation has also been used in reinforcement learning's action selection.³ Similar to the above example, valence can serve to steer exploration/exploitation behavior. For example, negative valence can lead to consideration of broader action selection and vice versa.³ Frustration can help control the strategy to drive such as by adapting weights of value functions for action selection. An extensive survey on the topic can be found in "Emotion in Reinforcement Learning Agents and Robots: A Survey."¹⁶

Anatomical models of AI. There also exist holistic models inspired by brain systems. Christian Balkenius and Jan Morén¹² suggested the first, which included models of the hippocampus, orbitofrontal cortex (OFC), amygdala, thalamus, and sensory cortex to produce emotional conditioning via the components' interdependencies.³ The (modeled) amygdala serves to learn emotional associations, and the modeled OFC is the contextual inhibitor. This contextual information is injected from the modeled hippocampus by matching stimuli to locations. Further examples include the Brain Emotional Learning-Based Intelligent Controller (BELBIC), a limbic system-inspired algorithm that was successfully applied as a controller in engineering tasks showing good generalization and flexibility, as the algorithm helped adapt to parameter changes and disturbances. Next, brain emotional learning (BEL) implements the mammalian amygdala's short paths that bridge the sensory thalamus and the

long paths to communicate with the frontal cortex of stimuli. The amygdala output is compared with rewards on the input side and can also be used if there is none of the latter. BEL could outperform multilayer perceptrons and fuzzy interference. Similarly, the brain emotional learning-based pattern recognizer (BELPR)¹² outperformed multilayer perceptrons in classification and time-series prediction tasks. In adaptive decayed brain emotional learning (ADBEL), a forgetting process was added to the amygdala—additional extensions such as fuzziness of amygdala and OFC variables in such models can help improve performance. As a final example, the limbic-based artificial emotional neural network (LiAENN) combines several of the above ideas, including anxiety and confidence as emotions, short and long paths, forgetting processes of the amygdala, and emotion suppression via the OFC-amygdala interaction. Again, such approaches were repeatedly outperforming multilayer perceptrons and other machine-learning approaches.³

Cognition and abstraction in learning. A range of further examples demonstrate how emotion augmentation can be used in abstraction and learning, such as adding an "emotional circuit" to agents controlled by a neural network, or emotion-based control for action selection influenced by emotional associations based on the cathexis model that includes an emotion generator, behavior systems, and motor systems as found in the emotionally conditioned robotic agent Yuppy. Further, in the Learning Intelligence Distribution Agent (LIDA), cognitive behavior (and choice of actions) are produced via the emotion-augmented

inclusion of attention, action selection, and motivation. Emotions that are based on event appraisal through emotional association trigger changes in learning and interference. In *Emotion-Augmented Machine Learning: Overview of an Emerging Domain*,³ readers can find further details and pointers to these works. Another example is Sigma,¹⁷ which is based on graphical models augmented by appraisal variables that represent the desire to control attention. Similarly, MAMID¹⁸ includes emotion modeling in an architecture where affective states are based on appraisal processes. These states impact the degree of processing capacities, threat bias, and more. In all of these examples, superiority over non-emotion-augmented peer solutions were observed in the studies introducing these architectures.

THE MILLION DOLLAR QUESTION: WILL AI HAVE EMOTIONS?

In the movie *Her*, the personal AI device Samantha wonders, "Are these feelings even real? Or are they just programming?" While there is a clear distinction between feelings and emotions, one is still tempted to question whether future AI will have real or simulated emotions, which might still fulfill the purpose. While this is largely a philosophical question, Picard spoke of machines driven by emotion, which could include curiosity, ideation, and motivation.¹ Alan Turing stated in 1950 that "if a machine behaves as intelligently as a human being, then it is as intelligent as a human being." One could apply this to human-like AEI. In the 1956 Dartmouth proposal that gave birth to the name AI, John McCarthy stated that "every aspect of learning or any

other feature of intelligence can be so precisely described that a machine can be made to simulate it.” McCarthy seemed convinced that AEI can indeed be reached. In 1976, Allen Newell and Herbert Simon agreed: “A physical symbol system has the necessary and sufficient means of general intelligent action.” Most interestingly, in the formulation of the concept of “strong” AI, John R. Searle postulated in 1980 that “computers given the right programs can be literally said to understand and have other cognitive states.”

It seems there is broad expert belief that AEI can be reached in AI systems. A claim for “real” emotion, however, is that it needs a body and a physical connection to the real world. Jürgen Schmidhuber and others allude in this context to the fact that pain sensors already exist in robotics.¹⁹ As partially outlined above, many authors have discussed the relation between AI and emotion and how they are intertwined. These visions include emotions controlling the choice, enabling, intensity, or preventing of AI behavior; attentional and perception mechanisms;²⁰ and a source of reinforcement when establishing functional descriptions of objects²¹ or when learning from humans.²² Many further works address how emotions can influence the reasoning process. For example, Joscha Bach stated just a decade ago that “the project of artificial intelligence is widely regarded as a failure,” but that he sees the principles of AEI as a potential game changer.²³ Like others, Bach argues for emotions impacting learning, memory, perception, and action selection and planning. The abilities of monitoring the actions of an agent are also discussed in John-Jules Ch. Meyer’s “Reasoning about Emotional Agents.”²⁴

Since Newell’s book *Unified Theories of Cognition* was published in 1990,²⁵ the importance of including emotion in blueprints for “soulful” machines has repeatedly been stressed in the literature on affective computing, behavioral computing, cognitive modeling, cognitive science, cognitive systems, adaptive behavior, and many more. AEI has increasingly grown into a mature field that enables computers to recognize and generate emotions at performance levels that can be exploited in real-world applications. Yet, in terms of full bidirectional emotion input/output, audio-visual and textual dialog systems currently prevail. In fact, current AEI is focused mostly on applying emotion augmentation in HCI with the hope of increasing naturalness and building deeper affinity in emotion-augmented retrieval and monitoring. At the same time, emotion augmentation of AI—particularly in machine learning—has been observed to have potential for improved efficiency and even higher accuracy. However, many other fields and applications of AI are yet to benefit from emotion augmentation. Consider, for example, the emotion augmentation box in Figure 1. There, one finds internal emotion recognition and generation, which is largely a “blank spot” in the literature for AI systems, as is the link with diverse contextual information shown in the bottom horizontal box in Figure 1.

To overcome these “blanks” in AEI, it seems crucial to derive methods of finding emotion concepts to include in AI, defining the change of such emotion based on external and internal state evaluation over time, and defining the consequences of such changes in emotion—ideally in reinforced learning, coupling analysis, and

synthesis of emotion so that we do not only recognize but also “generate” emotion, giving us a linked understanding of the underlying concepts.

In this sense, AEI needs to be integrated more seamlessly into future AI, going beyond often isolated and use-case-oriented consideration of affective computing approaches. Accordingly, the main novelty of our perspective is that fully embedded AEI should be a core piece of AI rather than a garnish or an “extra.” At the same time, AEI needs to be holistic in the sense of recognition, generation, and application of emotion and emotion principles, largely uniting the sub-disciplines that are currently considering the embedding of emotion principles in computing systems.

Ultimately, this will lead to a range of ethical, legal, and societal implications that should be addressed from a technical point of view such as with auditable, accountable, explainable, reliable, and responsible AEI. This will help us be better prepared for the advent of fully emotionally intelligent computing systems in the near future.

REFERENCES

1. R.W. Picard, *Affective Computing*, MIT Press, 1995.
2. A. Pentland, “Socially Aware Computation and Communication,” *Computer*, vol. 38, no. 3, 2005, pp. 33-40.
3. H. Strömfelt, Y. Zhang, and B.W. Schuller, “Emotion-Augmented Machine Learning: Overview of an Emerging Domain,” *7th Int’l Conf. Affective Computing and Intelligent Interaction (ACII 17)*, 2017, pp. 305-312.
4. C.E. Izard, *Human Emotions*, Springer, 1977.
5. B.L. Fredrickson, “What Good are Positive Emotions?,” *Rev. General*

ABOUT THE AUTHORS

DAGMAR SCHULLER is a co-founder and CEO of audeERING, an audio intelligence company located in Munich and Berlin that largely deals with artificial emotional intelligence. Her research interests include speech emotion recognition and general audio intelligence and affective computing. Schuller received a Master's degree in economic sciences from Vienna University of Economics and Business. Contact her at ds@audeering.com.

BJÖRN W. SCHULLER is a full professor of artificial intelligence and the head of GLAM—the Group on Language Audio & Music—at Imperial College London, a full professor and ZD.B Chair of Embedded Intelligence for Health Care and Wellbeing at the University of Augsburg, a co-founding CEO and current CSO of audeERING, and a permanent visiting professor at the Harbin Institute of Technology in China. Schuller received a PhD in electrical engineering/information technology from the Technical University of Munich. He is an IEEE Fellow, President Emeritus of the AAAC, and a Senior Member of ACM. He has written or co-authored more than 750 publications and is the editor in chief of IEEE Transactions on Affective Computing. Contact him at schuller@IEEE.org.

- Psychology, vol. 2, no. 3, 1998, pp. 300–319.
6. K.R. Scherer, "Appraisal Considered as a Process of Multilevel Sequential Checking," *Appraisal Processes in Emotion: Theory, Methods, Research*, 2001, pp. 92–120.
7. H. Gunes and B. Schuller, "Categorical and Dimensional Affect Analysis in Continuous Input: Current Trends and Future Directions," *Image and Vision Computing*, vol. 31, no. 2, 2013, pp. 120–136.
8. M. Soleymani et al., "A Survey of Multimodal Sentiment Analysis," *Image and Vision Computing*, vol. 35, pp. 3–14, 2017.
9. J.D. Williamson, *Speech Analyzer for Analyzing Pitch or Frequency Perturbations in Individual Speech Pattern to Determine the Emotional State of the Person*, US Patent 4,093,821, 1978.
10. J.E. Cahn, *Generation of Affect in Synthesized Speech*, MIT Media Lab, 1989; www.media.mit.edu/publications/generation-of-affect-in-synthesized-speech-2.
11. H. Kobayashi and F. Hara, "Recognition of Six Basic Facial Expressions and Their Strength by Neural Network," *Proc. IEEE Int'l Workshop Robot and Human Communication*, 1992; doi: 10.1109/ROMAN.1992.253857.
12. C. Balkenius and J. Morén, "A Computational Model of Context Processing," *Proc. 6th Int'l Conf. Simulation of Adaptive Behaviour*, 2000, pp. 256–265.
13. A. Khashman, "A Modified Back-propagation Learning Algorithm with Added Emotional Coefficients," *IEEE Transactions on Neural Networks*, vol. 19, no. 11, 2008, pp. 1896–1909.
14. Y. Yang et al., "Hybrid Chaos Optimization Algorithm with Artificial Emotion," *Applied Mathematics and Computation*, vol. 218, no. 11, 2012, pp. 6585–6611.
15. R. Thenius, P. Zahadat, and T. Schmickl, "EMANN—A Model of Emotions in an Artificial Neural Network," *The Twelfth European Conf. Artificial Life (ECAL 13)*, 2013, pp. 830–837.
16. T.M. Moerland, J. Broekens, and C.M. Jonker, "Emotion in Reinforcement Learning Agents and Robots: A Survey," *Machine Learning*, vol. 107, no. 2, 2018, pp. 443–480.
17. P.S. Rosenbloom, J. Gratch, and V. Ustun, "Towards Emotion in Sigma: From Appraisal to Attention," *Int'l Conf. Artificial General Intelligence (AGI 15)*, 2015, pp. 142–151.
18. E. Hudlicka, "Beyond Cognition: Modeling Emotion in Cognitive Architectures," *Proc. 6th Int'l Conf. Cognitive Modeling*, 2004, pp. 118–123.
19. P. Oltermann, "Jürgen Schmidhuber on the Robot Future: 'They Will Pay as Much Attention to Us as We Do to Ants,'" *The Guardian*, 18 Aug. 2017; www.theguardian.com/technology/2017/apr/18/robot-man-artificial-intelligence-computer-milky-way.
20. D. Cañamero, "Modeling Motivations and Emotions as a Basis for Intelligent Behavior," *Proc. First Int'l Conf. Autonomous Agents (AGENTS 97)*, 1997, pp. 148–155.
21. I. Cos and G.M. Hayes, "Behaviour Control Using a Functional and Emotional Model," *Proc. 7th Int'l Conf. Simulation of Adaptive Behavior (ICSAB 02)*, 2002, pp. 226–227.
22. J. Broekens, "Emotion and Reinforcement: Affective Facial Expressions Facilitate Robot Learning," T.S. Huang et al., eds., *Artificial Intelligence for Human Computing*, LNCS 4451, Springer, Berlin, Heidelberg, 2007.
23. J. Bach, "Seven Principles of Synthetic Intelligence," *Proc. 1st Conf. Artificial General Intelligence (AGI 08)*, 2008, pp. 63–74.
24. J.-J. Ch. Meyer, "Reasoning about Emotional Agents," *Int'l J. Intelligent Systems*, vol. 21, no. 6, 2006, pp. 601–619.
25. A. Newell, "Unified Theories of Cognition," Harvard Univ. Press, 1990.