# Building Pedagogical Conversational Agents, *Affectively* Correct

Michalis Feidakis[1], Panagiotis Kasnesis[1], Eva Giatraki[2], Christos Giannousis[1],
Charalampos Patrikakis[1] and Panagiotis Monachelis[1]
*[1]Department of Electrical & Electronics Engineering, University of West Attica, Thivon 250 & Petrou Ralli street, Egaleo GR 12244, Greece*
*[2]6th Primary School of Egaleo (K. Kavafis), 6 Papanikoli St, Egaleo GR 12242, Greece*

Keywords: Conversational Agents, Chatbots, Pedagogical Agents, Affective Feedback, Artificial Intelligence, Deep Learning, Affective Computing, Reinforcement Learning, Sentiment Analysis, Core Cognitive Skills.

Abstract: Despite the visionary tenders that emerging technologies bring to education, modern learning environments such as MOOCs or Webinars still suffer from adequate affective awareness and effective feedback mechanisms, often leading to low engagement or abandonment. Artificial Conversational Agents hold the premises to ease the modern learner's isolation, due to the recent achievements of Machine Learning. Yet, a pedagogical approach that reflects both cognitive and affective skills still remains undelivered. The current paper moves towards this direction, suggesting a framework to build pedagogical driven conversational agents based on Reinforcement Learning combined with Sentiment Analysis, also inspired by the pedagogical learning theory of Core Cognitive Skills.

## 1 INTRODUCTION

During the last decade, ubiquitous learning (MOOCs, Webinars, Mobile Learning, RFIDs & QRcodes, etc.) has made significant steps towards the sought democratization of education (Caballé and Conesa, 2019), also drifted by the latest ICT advancements (i.e. IoT, Cloud Computing, 5G) (El Kadiri et al., 2016). However, the auspicious new educational paradigms fail to encounter distance learning (d-learning) or e-learning common "headaches" such as low students' engagement and poor immersion rates (Afzal et al., 2017). The rising technology of Affective Computing (AC) promises to contribute significantly towards the motivation and engagement of the *isolated* remote learners (Graesser, 2016).

AC has emerged as a new research field in Artificial Intelligence (AI) designated to "sense and respond" to user's affective state (Picard, 1997) in various domains, and especially in education, in which addresses three key issues (Feidakis, 2016):

1. Detect and recognize the learner's emotion/ affective state with high accuracy.
2. Display emotion information through effective visualisations for self, peers and others' emotion awareness.
3. Provide cognitive and affective feedback to improve task & cognitive performance, as well as learning outcomes.

The latter (affective feedback) remains unexplored, still constituting a big challenge: How to contend with the learner's isolation encountered in Virtual Learning Environments (VLEs), often leading to dropouts, while system's feedback suffers from:

- **Empathy –** The system is usually unaware of the learner's affective state and affect transitions, i.e. the student's confusion steadily increases, and possibly transit to frustration (Afzal et al., 2017).
- **Brevity –** Response is outdated because there is no need to address the problem anymore, or worse, the student *left the building* (Feidakis, 2016).
- **Sociality –** New learning tools (MOOCs) often isolate the learner who ultimately ends up in cognitive deadlocks (Caballé & Conesa, 2018).

The recent findings of Deep Learning –especially when applied to the fields of Natural Language Processing (NLP) (Devlin et al., 2018) and Reinforcement Learning (RL) (Silver et al., 2018) – premises the capacity to develop smart, artificial agents able to manipulate the above-mentioned issues. Nevertheless, the following questions arise:



Feidakis, M., Kasnesis, P., Giatraki, E., Giannousis, C., Patrikakis, C. and Monachelis, P.
Building Pedagogical Conversational Agents, Affectively Correct.
DOI: 10.5220/0007771001000107
In *Proceedings of the 11th International Conference on Computer Supported Education (CSEDU 2019)*, pages 100-107
ISBN: 978-989-758-367-4

1. How we built agents that are triggered **consistently**? –in terms of transparency in affect detection and brevity in response.
2. How we train agents to reply **adequately**? – regarding the impact of response.

In the current paper, we present a framework model to build agents based on Deep RL models, empowered by Question-Answering (QA) pedagogical learning strategies, also showing *some respect* to the respondent's feelings through Sentiment Analysis of short dialogues. We first review current State-of-the-Art of Artificial/ Conversational Agents, Chatbots and Affective/ Empathetic Agents, together with Sentiment Analysis of short texts and posts. Then, we present our proposal towards the enquiries set. This new framework will guide the implementation of a new model in our future steps, also prompting for more contributions and collaborations.

# 2 LITERATURE REVIEW

## 2.1 Artificial Conversational Agents and Chatbots

There have been almost three decades since intelligent agents have been introduced, and AI researchers started to look at the "whole agent" problem (Russell et al., 2010). Nowadays, "AI systems have become so common in Web-based applications that the -bot suffix has entered everyday language" (Russell et al., 2010, p. 26). The consistent evolution of AI together with the prompt establishment of IoT have coined many new areas to invest such as smart homes, smart cities and recently smart advisors or virtual assistants (Walker, 2019). Artificial Conversational Agents are software agents trying to answer a particular question through information retrieval techniques and by engaging the user into understanding the nature of the problem behind the question.

Chatbots are one category of Conversational Agents (Radziwill and Benton, 2017), hence, applications that use AI techniques to communicate and process data provided by the user. Chatbots look for keywords into user's questions, trying to understand what the user really wants. With the help of AI, Chatbot applications keep evolving based on historical data to cope with new input information. Chatbots such as Amazon Alexa ("Amazon Alexa," 2019), Google Assistant ("Google Assistant," 2019), Apple's Siri ("Apple Siri," 2019), and Facebook Jarvis (Zuckerberg, 2016) are steadily increase their market share in people's everyday applications.

There are several tools based on information retrieval techniques to create a Chatbot agent that come with respective developer's tools. Dialogflow ("Dialogflow," 2019) consists of 2 related components: (i) *intents*, to build arrays of various questions (e.g. Where is the Gallery?), and (ii) *entities*, sets of words (tokens) that help the agent to analyse spoken or written text (e.g. product names, street names, movie categories). It also provides an API to integrate the application in certain media or services such as Facebook Messenger, WhatsApp, Skype, SMS etc. IBM's Watson ("IBM Watson," 2019) works similarly to Dialogflow, allowing the user to connect to many other applications through Webhook or APIs, however it requires subscription (the free version is limited up to 1000 intents per month).

In education, Chatbots have been recently applied to various disciplines such as physics, mathematics, languages and chemistry, usually inspired by gamification pedagogical strategies. For instance, through machine learning-oriented Q&As, a Chatbot is able to evaluate the learner's level each time and fine-tunes the game's level of difficulty (Benotti, Martínez, & Schapachnik, 2014).

## 2.2 Affective Agents

The provision of affective feedback to users, in response to their *implicit* (automatically and transparently) or *explicit* recognition of affective state, remains prominent topic. Learners need to perceive a reaction from the system, in agreement with their emotion sharing, immediately or after a short period (Feidakis et al., 2014), triggering an affective loop of interactions (Castellano et al., 2013).

Since 2005, "James the butler" (Hone et al., 2019), an Affective Agent developed in Microsoft Agents and Visual Basic, managed to reduce negative emotions' intensity (despair, sadness, boredom). Similarly, Burleson (2006)) developed an Affective Agent able to mimic facial expressions. The agent managed to help 11-13 years old students to solve cognitive problems (i.e. Tower of Hanoi) by providing affective scaffolds in case i.e. of despair. AutoTutor comes with a long list of published results (D'Mello et al., 2011) successfully demonstrating both cognitive and affective skills (empathy). Moridis and Economides (Moridis and Economides, 2012) report on the impact of Embodied Conversational Agents (ECAs) on the respondent's affective state (sustain or modify) though corresponding empathy, either parallel (express harmonized emotions) or reactive (stimulate different or even contradictory emotions). EMOTE is another tool that can be integrated in existing agents, enriching their emotionality towards their improved learning performance and emotion wellbeing (Castellano et al., 2013).

In (Feidakis et al., 2014), a virtual Affective Agent provided affective feedback enriched with task-oriented scaffolds, according to fuzzy rules. The agent managed to improve student cognitive performance and emotion regulation. A visionary review of artificial agents that simulate empathy in their interactions with humans is provided by Paiva et al. (2017), delivering sufficient evidence about the significant role of affective or empathetic responses when (or where) humans, agents, or robots are collaborating and executing tasks together.

### 2.3 Sentiment Analysis

Sentiment Analysis plays a significant role on products marketing, companies' strategies, political issues, social networking, as well as on education. Main goal is to mine for opinions in textual data e.g. a web source, and extract information about author's sentiments. The derived data are oriented mainly to polarity depiction of the sentiments using data mining and NLP techniques (Cambria et al., 2013).

There are 3 main approaches to implement Sentiment Analysis:

- **Linguistic Approach:** A sentiment is exported from text analysis according to a lexicon. In this case, three levels of text (document-level, sentence-level, aspect-level) are analysed and algorithms based on lexicons produce the sentiment score (Feldman, 2013).
- **Machine Learning Approach:** Employs both supervised and unsupervised learning towards the classification of texts. It has become quite popular since it is used for many applications, such as movie review classifier (Cambria, 2016). Different algorithms have been applied over the years, such as Support Vector Machines, Deep Neural Networks, Naïve Bayes, Bayesian Network and Maximum Entropy MaxEnt, with Deep Learning approaches appearing to give better results (Howard and Ruder, 2018).
- **Hybrid Approach:** A reconciliation of the two above methods (Cambria, 2016).

Nowadays, State-of-the-Art Sentiment Analysis techniques are based on supervised machine learning approaches and in particular Deep Learning algorithms. Most of these approaches use word embeddings, such as word2vec (Mikolov et al., 2013), which are vectors for representing the words, and are trained in an unsupervised way. Afterwards, they are fed to a deep Recurrent Neural Network (RNN) (Rumelhart et al., 1987) – a family of neural networks used for processing sequence values– empowered with a mechanism named LSTM (Long Short-Term Memory) (Hochreiter and Schmidhuber, 1997) to enhance the memory of the network. In addition, recent advances (Devlin et al., 2018) exploiting the attention mechanism (Vaswani et al., 2017) –focusing more on some particular words in a sentence– seem to achieve higher accuracy in the Sentiment Analysis task.

In education context, students' textual data could be analysed and provide valuable feedback to an agent, who can act interactively, either by rewarding a successful task accomplishment or encouraging the students after a failed attempt. The application of the aforementioned techniques to short posts extracted from educational platforms (i.e. LMS, MOOC), could provide useful information regarding the correlation between students' sentiments and performance (Tucker et al., 2014).

## 3 ARCHITECTURAL APPROACH

According to (Russell et al., 2010), the definition of an AI agent involves the concept of rationality: "*a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has*" (p. 37). A learning agent comprises two main components: (i) the **learning element**, for making improvements, and (ii) the **performance element**, for selecting external actions. In other words, the performance element involves the agent decisions, while the learning element deals with the evaluation of those actions. Designing an agent requires the selection of an appropriate and effective learning strategy to train the agent. Learning strategies and models that have been validated for years, in real education settings, could extend the design horizons of artificial pedagogical agents. In next paragraphs we unveil this perspective.

A popular strategy to guide a learner in alternative learning paths is to use Question-Answer (QA) models, which are quite easy to implement, especially when short default answers are deployed (Yes/No, Multiple choices, Likert-scales, etc.). In learning settings, QA models constitute a common way to deploy formative assessment i.e. according to a rubric. Also, they can be nicely shaped as valuable scaffolds to unlock preexisting knowledge or semi-completed conceptual schemas (Vygotsky, 1987). In all cases, the overload shifts to formulate the right questions.

In their Academically Productive Talk (APT) model, Tegos and Demetriadis (2017) emphasize the orchestration of teacher-students talks and highlight a set of useful discussion practices that can lead to reasoned participation by all students, thus increasing the probability of productive peer interactions to occur. Moreover, ColMOOC project (Caballé and Conesa, 2019) constitutes a recent exertion towards a new pedagogical paradigm that integrates MOOCs with Conversational Agents and Learning Analytics, according to the Conversation Theory (Pask, 1976).

Next paragraphs provide our framework to build agents, (i) to give the right answers, (ii) based on pedagogical models, (iii) also considering affective factors. Our design is grounded on Deep Learning, thus, generating responses through the imitation of training datasets.

## 3.1 Goal-Oriented Question-Answering (QA) Model

In supervised learning the algorithm given as input a labeled dataset *X* tries to predict the output *Y*. For example, in Sentiment Analysis the algorithm is given as input a short text (e.g. review) and its goal is to predict if it is negative or positive. Since, during the training process the label is given (i.e. 0 for negative and 1 for positive) the algorithm learns from its mistakes and updates its weights in every iteration, until it converges.

However, in RL the algorithm (i.e. AI agent) has as input a set of data *X*, but the desired output is defined as a reward *r*. In particular, the input is considered to be the state of the agent *s*. To maximize the reward function, the agent selects an action *a* from the action space *A*. It should be noted, that there can be two types of rewards, long-term and short-term (instant). Let's take for example, the *Pac-Man* game:

- The state of an agent (i.e. Pac-Man) is the locations of the ghosts, its own location and the existing dots in the maze.
- The action space is consisted of the moves the agent can make (i.e. up, down, left, right).
- The instant rewards can be 1 for eating a dot, -100 for been eaten by a ghost and 0 for doing none of them. However, there can be a long-term reward equal to 100 for completing the level.
- Finally, the strategy of the agent (i.e. which move it should select given its current state) is called policy, and is denoted with $\pi$.

In the current paper we advocate that RL combined with Sentiment Analysis, can be used to develop a pedagogical driven conversational agent. In this case the whole task of the agent can be formulated as follows:

- *Goal*: Help the student to complete the test successfully without providing the direct answer to him/her.
- *State:* The text input given by the student and his/her current affective state.
- *Action*: The hint that is provided to the student in text format.
- *Reward*: When the student will successfully complete the task (e.g. if the hint will be helpful).
- *Policy:* What strategy the agent should select, given the affective state of the student (e.g. frustration) and the answer(s) he/she provided.

Therefore, the learning task is expressed in a goal-oriented way, meaning that the agent tries to achieve the student's comprehension through dialog. Of course, this is not the first goal-oriented approach in a QA task (see Rajendran et al., 2018). The proposed method includes two phases: (i) the agent tries to learn how to perform dialog with the student and detect his/her affective state given an annotated dataset (supervised learning), and (ii) the agent learns through trial and error –based on the incoming rewards– to assist the student (reinforcement learning) (Figure 1).

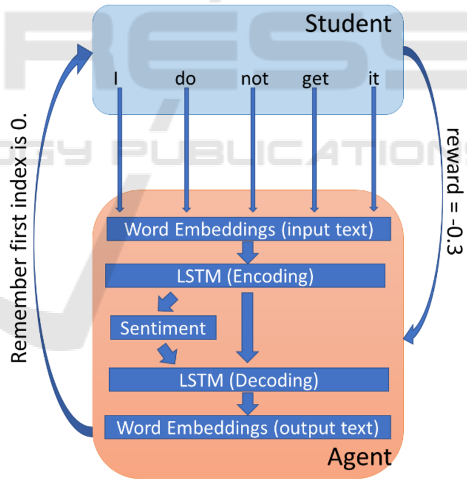


Figure 1: Goal-Oriented Question-Answering Agent.

In (Figure 1), the student sends a text answer in the agent, which converts it into embeddings (i.e., numerical vectors). Next, the embeddings are processed by an LSTM layer, which is both used to encode the incoming data and recognizes the students affective state (i.e., confusion). The outputs are fused together using again a LSTM layer and produce the output word embeddings; these are mapped into actual words so that a helpful hint is sent to the user. During the whole process the agent receives a reward based on the student performance.

## 3.2 Pedagogical Supervision

Human reasoning constitutes a synergistic association of ideas, elaborating a "non-stop" mining of reasonable correlations between existing and input data, in a continuous restructuring, or *reforming* of cognitive schemas (assimilation-accommodation-adaptation, known from the Piaget's Theory of Cognitive Development (Piaget et al., 1969). Reforming comprises two main operations: the convergent *critical* and the divergent *creative* thinking.

Elaborate *creative* thinking on machines, seems intangible, nevertheless, it is already in the AI agenda. In ("What's next for AI," 2019) it is mentioned as the "ultimate moonshot for artificial intelligence" highlighting the fact that AI researchers should not confine to machines that only think or learn, but also create (Russell et al., 2010). The integration of emotional and affective appraisals in computer intelligence, constitutes a big step, since creativity depend heavily on emotional weighted decisions related to motivation, perversion, insight, intuition, etc.

Critical thinking sounds more tangible. For instance, the IBM Watson ("IBM Watson," 2019) cognitive platform managed to analyse visuals, sound and composition of hundreds of existing horror film trailers and accomplished an AI-created movie trailer for 20th Century Fox's horror flick, *Morgan* ("What's next for AI," 2019).

The vision and challenge as well, involves an agent holding the capacity, not only to beat a chess grandmaster (Silver et al., 2018), but also to guide the learner in a collaborative mining of creative solutions, *out* and *beyond* the respective problem. Q&As methods – i.e. the Socratic method – are moving towards that direction by breaking down a subject into a series of questions, the answers to which gradually distill the answer a person would seek.

In the exemplary work of Marzano (1988), 24 Core Cognitive Skills (CCS) have been classified and grouped into four basic categories of data processing, namely *collection*, *organisation*, *analysis* and *transcendence* (Table 1). Based on this work, Kordaki et al. (2007) proposed an architecture for a Cognitive Skill-based Question Wizard.

It is in our intentions to build, evaluate and evolve a framework that will employ both supervised and RL approaches to simulate tutor-learner conversations. The intelligent tutor will guide the learner to solve learning tasks, evolving Cognitive Skills based on the 24 CCS. In Table 2 we provide examples of questions that could be used to develop basic Cognitive Skills in agents.

Table 1: Marzano 24 CCS.

| | |
|---|---|
| Data Collection | 1. Observation<br>2. Recognition<br>3. Recall |
| Data Organization | 4. Comparison<br>5. Classification<br>6. Ordering<br>7. Hierarchy |
| Data Analysis | 8. Analysis<br>9. Recognition of Relationships<br>10. Pattern Recognition<br>11. Separation of Facts from Opinions<br>12. Clarification |
| Data Transcendence | 13. Explanation<br>14. Prediction<br>15. Forming Hypotheses<br>16. Conclusion<br>17. Validation<br>18. Error detection<br>19. Implementation-Improvement<br>20. Knowledge organization<br>21. Summary<br>22. Empathy<br>23. Assessment /Evaluation<br>24. Reflection |

Table 2: Examples of question-models (adjusted from Kordaki et al., 2007).

| **List of Basic Cognitive Skills** | | **Examples of question-models** |
|---|---|---|
| **Data collection** | ***Observation*** | *A* is a list of integers<br>*A* = [2, 1, 7, 0]<br>0 1 2 3 ←index<br>List index starts from 0<br>List increments by 1<br>What is the value of *A*[2]? |
| | ***Recall*** | *Remember A index starts from 0.*<br>Print (*A*[4])<br>Is this correct? |
| **Data organization** | ***Comparison*** | *A*[0] > *A*[1]<br>Is this correct? |
| | ***Classification*** | *a* is integer, *b* is real<br>*c* = *a*+*b*<br>*c* is real<br>Is this correct? |
| **Data analysis** | ***Recognition of Relationships*** | *a* >*b*, *b*>*c*<br>*a*>*c* Is this correct? |
| | ***Pattern recognition*** | *A* = [2, 1, 7, 0], B = [3, 5, 4]<br>*C* = [*A*, *B*]<br>*C*[1][0]=3 Is this correct? |
| **Data transcendence** | ***Prediction*** | $a=2x+y$<br>if $x=2$, $y=\begin{cases} 1, \text{ if } p<0.1 \\ 0, \text{ if } p>=0.1 \end{cases}$<br>What is the value of *a*? |
| | ***Forming Hypotheses*** | If (it rains) and (I won't take umbrella)<br>then I get wet |

### 3.3 Affective Factor

As already mentioned (subsection 2.2), an Affective Agent seeks for affective cues to appraise the respondent's sentiments. Next step involves the agent's decision to provide an answer *affectively correct*. An affective response should be able to change learning paths when i.e. boredom or frustration is recognized (Feidakis, 2016).

Zhou et al. (2017) provide a dataset of 23,000 sentences collected from the Chinese blogging service Weibo and manually annotated using 5 labels: anger, disgust, happiness, like, and sadness. Such datasets provide a starting point to enrich agent's short dialogs with *emotion hues*.

In previous work (Feidakis et al., 2014), a Sentiment Analysis mechanism was implemented classifying short posts in Web forums and Wikis, according to 6 states (Figure 2), based on Machine Learning and NLP.

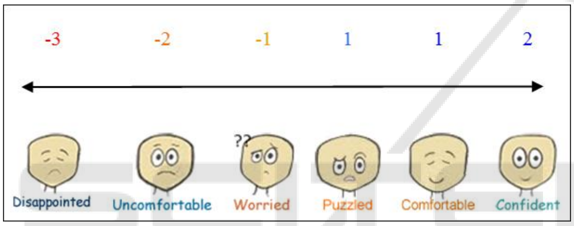


Figure 2. Sentiments' classifiers (Feidakis et al., 2014).

In this work, we need to extend our tasks to integrate emotion or affective oriented *lemmas* (inferences), following a machine learning approach. In Table 3 we provide an example of agent's responses, while detecting first frustration and then interest in user's response:

Our contribution lies also on a new conceptual model reflecting learners' affective states in e-learning context deploying both (i) dimensional, and (ii) label representation models. In the former, candidate dimensions already evaluated in learning settings (Feidakis, 2016) are *valence* (positive-negative), *arousal* (high-low) and *duration* (short-long). The latter involves affective states that have been classified in educational context such as *inspiration(high-five), excitement, frustration, anger, stress* (more emotive) or, *engagement(interest), boredom, fatigue, confusion* (more affective). Both approaches will be prototyped in a hybrid conceptual model addressing emotive/affective states together with their tentative transitions in time (Feidakis, 2016).

Table 3: An example of affective responses.

| Agent | User |
|---|---|
| *A* = [2, 1, 7, 0]<br>*B* = [3, 5, 4]<br>*C* = [*A*, *B*]<br>*C*[1][0]=3<br>Is this correct? | No |
| Remember first index is 0. | I am not sure I got it |
| OK. Do you want to try something else and come back later? | No |
| I like that! Here is a hint for you: *Lists can take also other lists as values.*<br>What is the value of *C*[1]? | *B* |
| *A* is a list | [3, 5, 4] |
| | |
| Nice! And the [0] of *A* | 3 |
| Great! That's the correct answer! *C*[1][0]=3 | |

## 4 CONCLUSIONS AND FUTURE WORK

The proliferation of digital assistants, which at the present is targeting the domains of productivity and domotic automation (i.e Google home, Amazon's Alexa), has paved the ground for the introduction of digital assistants in several other domains (e.g. personal training). The education domain, presents particularities and needs that arise from the need to integrate the assessment, monitoring and exploitation of cognitive and affective skills of the trainees. At the same time, RL inherently supports the level of interactivity required, so that the process of integrating machine learning techniques in the process of supporting the educational process is performed in line with Sentiment Analysis and Cognitive Skills assessment.

The current State-of-the-Art of Deep Learning can provide simple models and exploit large amounts of real data to train agents to act as valuable assistants, customizing them to the needs of each case through specialized datasets like bAbI (Weston et al., 2015) or Emotional STC (ESTC) (Zhou et al., 2017). In this paper, the above approach is proposed, towards the provision of a framework for supporting pedagogical driven conversational agents. The model of course needs a lot of data towards its realistic use, since the agent need to perform natural language understanding and Sentiment Analysis, before responding to the user based on the optimal policy (both state and action spaces i.e., input and output text, are quite big, if we consider how many word combinations exist).

Though the potential that the proposed approach is quite high, leading to the introduction of highly interactive tools of (pedagogically) added value, the introduction of such a framework will introduce new challenges, especially when it comes to the protection of personal information. For this, both the legal/regulatory and the technical frameworks which will ensure the protection of personal and sensitive information should be taken under serious consideration, and implementations adopting privacy by design approaches should be adopted. As for the former, the recently applied GDPR in EU (679/2016) provides the basis, over which compliance to ethical and legal standards of any implementation can be evaluated. As for the latter, the ability of increased capabilities of terminal devices (Edge Computing – Shi et al., 2016) is a promising candidate for limiting the range over which collected data for applying the machine learning techniques are used.

## REFERENCES


Afzal, Shazia, Bikram Sengupta, Munira Syed, Nitesh Chawla, G. Alex Ambrose, and Malolan Chetlur. "The ABC of MOOCs: Affect and Its Inter-Play with Behavior and Cognition." In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, 279–84. San Antonio, TX: IEEE, 2017. https://doi.org/10.1109/ACII.2017.8273613.

Amazon Alexa [WWW Document], URL https://developer.amazon.com/alexa (accessed 3.21.19).

Apple Siri [WWW Document], URL https://www.apple.com/siri/ (accessed 3.21.19).

Benotti, Luciana, María Cecilia Martínez, and Fernando Schapachnik. "Engaging High School Students Using Chatbots." In *Proceedings of the 2014 Conference on Innovation & Technology in Computer Science Education - ITiCSE '14*, 63–68. Uppsala, Sweden: ACM Press, 2014. https://doi.org/10.1145/2591708.2591728.

Burleson, Winslow. "Affective Learning Companions: Strategies for Empathetic Agents with Real-Time Multimodal Affective Sensing to Foster Meta-Cognitive and Meta-Affective Approaches to Learning, Motivation, and Perseverance." PhD Thesis, Massachusetts Institute of Technology, 2006.

Caballé, Santi, and Jordi Conesa. "Conversational Agents in Support for Collaborative Learning in MOOCs: An Analytical Review." In *Advances in Intelligent Networking and Collaborative Systems*, edited by Fatos Xhafa, Leonard Barolli, and Michal Greguš, 384–94. Springer International Publishing, 2019.

Cambria, Erik. "Affective Computing and Sentiment Analysis." *IEEE Intelligent Systems* 31, no. 2 (March 2016): 102–7. https://doi.org/10.1109/MIS.2016.31.

Cambria, Erik, Bjorn Schuller, Yunqing Xia, and Catherine Havasi. "New Avenues in Opinion Mining and Sentiment Analysis." *IEEE Intelligent Systems* 28, no. 2 (March 2013): 15–21. https://doi.org/10.1109/MIS.2013.30.

Castellano, Ginevra, Ana Paiva, Arvid Kappas, Ruth Aylett, Helen Hastie, Wolmet Barendregt, Fernando Nabais, and Susan Bull. "Towards Empathic Virtual and Robotic Tutors." In *Artificial Intelligence in Education*, edited by H. Chad Lane, Kalina Yacef, Jack Mostow, and Philip Pavlik, 7926:733–36. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013. https://doi.org/10.1007/978-3-642-39112-5_100.

D'Mello, Sidney K., Blair Lehman, and Art Graesser. "A Motivationally Supportive Affect-Sensitive AutoTutor." In *New Perspectives on Affect and Learning Technologies*, edited by Rafael A. Calvo and Sidney K. D'Mello, 113–26. New York, NY: Springer New York, 2011. https://doi.org/10.1007/978-1-4419-9625-1_9.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding." *ArXiv:1810.04805 [Cs]*, October 10, 2018. http://arxiv.org/abs/1810.04805.

Dialogflow [WWW Document], URL http://dialogflow.com (accessed 3.21.19).

El Kadiri, Soumaya, Bernard Grabot, Klaus-Dieter Thoben, Karl Hribernik, Christos Emmanouilidis, Gregor von Cieminski, and Dimitris Kiritsis. "Current Trends on ICT Technologies for Enterprise Information Systems." *Computers in Industry* 79 (June 2016): 14–33. https://doi.org/10.1016/j.compind.2015.06.008.

Feidakis, M. "A Review of Emotion-Aware Systems for e-Learning in Virtual Environments." In *Formative Assessment, Learning Data Analytics and Gamification*, 217–42. Elsevier, 2016. https://doi.org/10.1016/B978-0-12-803637-2.00011-7.

Feidakis, Michalis, Santi Caballé, Thanasis Daradoumis, David Gañán Jiménez, and Jordi Conesa. "Providing Emotion Awareness and Affective Feedback to Virtualised Collaborative Learning Scenarios." *International Journal of Continuing Engineering Education and Life-Long Learning* 24, no. 2 (2014): 141. https://doi.org/10.1504/IJCEELL.2014.060154.

Feldman, Ronen. "Techniques and Applications for Sentiment Analysis." *Communications of the ACM* 56, no. 4 (April 1, 2013): 82. https://doi.org/10.1145/2436256.2436274.

Google Assistant [WWW Document], 2019. URL https://assistant.google.com/ (accessed 3.21.19).

Graesser, Arthur C. "Conversations with AutoTutor Help Students Learn." *International Journal of Artificial Intelligence in Education* 26, no. 1 (March 2016): 124–32. https://doi.org/10.1007/s40593-015-0086-4.

Hochreiter, Sepp, and Jürgen Schmidhuber. "Long Short-Term Memory." *Neural Computation* 9, no. 8

(November 1997): 1735–80. https://doi.org/10.1162/neco.1997.9.8.1735.

Hone, Kate, Lesley Axelrod, and Brijesh Parekh. *Development and Evaluation of an Empathic Tutoring Agent*, 2019.

Howard, Jeremy, and Sebastian Ruder. "Universal Language Model Fine-Tuning for Text Classification." *ArXiv:1801.06146 [Cs, Stat]*, January 18, 2018. http://arxiv.org/abs/1801.06146.

IBM Watson [WWW Document], URL http://www.ibm.com/watson (accessed 3.21.19).

Kordaki, Maria, Spyros Papadakis, and Thanasis Hadzilacos. "Providing Tools for the Development of Cognitive Skills in the Context of Learning Design-Based e-Learning Environments." In *Proceedings of E-Learn: World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education 2007*, edited by Theo Bastiaens and Saul Carliner, 1642–1649. Quebec City, Canada: Association for the Advancement of Computing in Education (AACE), 2007. https://www.learntechlib.org/p/26585.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. "Distributed Representations of Words and Phrases and Their Compositionality." *ArXiv:1310.4546 [Cs, Stat]*, October 16, 2013. http://arxiv.org/abs/1310.4546.

Moridis, Christos N., and Anastasios A. Economides. "Affective Learning: Empathetic Agents with Emotional Facial and Tone of Voice Expressions." *IEEE Transactions on Affective Computing* 3, no. 3 (July 2012): 260–72. https://doi.org/10.1109/T-AFFC.2012.6.

Paiva, Ana, Iolanda Leite, Hana Boukricha, and Ipke Wachsmuth. "Empathy in Virtual Agents and Robots: A Survey." *ACM Transactions on Interactive Intelligent Systems* 7, no. 3 (September 19, 2017): 1–40. https://doi.org/10.1145/2912150.

Pask, Gordon. *Conversation Theory: Applications in Education and Epistemology*. Amsterdam; New York: Elsevier, 1976.

Piaget, Jean, Bärbel Inhelder, and Helen Weaver. *The Psychology of the Child*. Nachdr. New York: Basic Books, Inc, 1969.

Picard, Rosalind W. *Affective Computing*. Cambridge, Mass: MIT Press, 1997.

Radziwill, Nicole M., and Morgan C. Benton. "Evaluating Quality of Chatbots and Intelligent Conversational Agents." *ArXiv:1704.04579 [Cs]*, April 15, 2017. http://arxiv.org/abs/1704.04579.

Rajendran, Janarthanan, Jatin Ganhotra, Satinder Singh, and Lazaros Polymenakos. "Learning End-to-End Goal-Oriented Dialog with Multiple Answers." *ArXiv:1808.09996 [Cs]*, August 24, 2018. http://arxiv.org/abs/1808.09996.

Rumelhart, David E, James L McClelland, San Diego University of California, and PDP Research Group. *Parallel Distributed Processing. Explorations in the Microstructure of Cognition*, 1987.

Russell, Stuart J., Peter Norvig, and Ernest Davis. *Artificial Intelligence: A Modern Approach*. 3rd ed. Prentice Hall Series in Artificial Intelligence. Upper Saddle River: Prentice Hall, 2010.

Shi, Weisong, Jie Cao, Quan Zhang, Youhuizi Li, and Lanyu Xu. "Edge Computing: Vision and Challenges." *IEEE Internet of Things Journal* 3, no. 5 (October 2016): 637–46. https://doi.org/10.1109/JIOT.2016.2579198.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, et al. "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go through Self-Play." *Science* 362, no. 6419 (December 7, 2018): 1140–44. https://doi.org/10.1126/science.aar6404.

Stergios Tegos, and Stavros Demetriadis. "Conversational Agents Improve Peer Learning through Building on Prior Knowledge." *Journal of Educational Technology & Society* 20, no. 1 (2017): 99–111.

Tucker, Conrad S., Barton K. Pursel, and Anna Divinsky. "Mining Student-Generated Textual Data in MOOCs and Quantifying Their Effects on Student Performance and Learning Outcomes." *Computers in Education Journal* 5, no. 4 (January 1, 2014): 84–95.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." *ArXiv:1706.03762 [Cs]*, June 12, 2017. http://arxiv.org/abs/1706.03762.

Vygotsky, L. S. *The Collected Works of L. S. Vygotsky, Vol. 1: Problems of General Psychology*. The Collected Works of L. S. Vygotsky, Vol. 1: Problems of General Psychology. New York, NY, US: Plenum Press, 1987.

Walker, M., 2019. Hype Cycle for Emerging Technologies https://www.gartner.com/doc/3885468/hype-cycle-emerging-technologies- (accessed 3.21.19).

Weston, Jason, Antoine Bordes, Sumit Chopra, Alexander M. Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. "Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks." *ArXiv:1502.05698 [Cs, Stat]*, February 19, 2015. http://arxiv.org/abs/1502.05698.

What's next for AI [WWW Document], n.d. URL http://www.ibm.com/watson/advantage-reports/future-of-artificial-intelligence/ai-creativity.html (accessed 3.21.19).

Zhou, Hao, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. "Emotional Chatting Machine: Emotional Conversation Generation with Internal and External Memory." *ArXiv:1704.01074 [Cs]*, April 4, 2017. http://arxiv.org/abs/1704.01074.

Zuckerberg, M., 2016. Building Jarvis [WWW Document]. URL https://www.facebook.com/notes/mark-zuckerberg/building-jarvis/10103347273888091/?pnref=story (accessed 3.21.19).