

HOW GENERATIVE LANGUAGE MODELS CAN ENHANCE INTERACTIVE LEARNING WITH SOCIAL ROBOTS

Stefan Sonderegger

*Institute for Educational Management and Technologies
University of St.Gallen, 9000 St.Gallen, Switzerland*

ABSTRACT

The use of social robots in education is a growing area of research and the potential future applications are various. However, the conversational models behind current social robots and chatbot systems often rely on rule-based and retrieval-based methods. This limits the social robot to predefined responses and topics, thus hindering it from fluent communication and interaction. Generative language models such as GPT-3 could be beneficial in this context, e.g. for an improved conversation and open-ended question answering. This article presents an approach to utilizing generative language models to enhance interactive learning with educational social robots. The proposed model combines the technological possibilities of generative language models with the educational tasks of a social robot in the role of a tutor and learning partner. The implementation of the model in practice is illustrated by means of a use case consisting of different learning scenarios. The social robot generates explanations, questions, corrections, and answers based on the pre-trained GPT-3 model. By exploring the potential of generative language models for interactive learning with social robots on different levels of abstraction, the paper also aims to contribute to an understanding of the future relevance and possibilities that generative language models bring into education and educational technologies in general.

KEYWORDS

Social Robots in Education, Human-Robot-Interaction, Interactive Learning, Generative Language Models, GPT-3

1. INTRODUCTION

Social robots are specialized in social interactions with humans and have the potential to support our learning and teaching (Belpaeme *et al.* 2018). In roles such as tutor, teaching assistant, and peer learner, social robots might be used not only as a practical tool to learn about STEM subjects but also as a teaching and learning tool, e.g. for language learning, in special education, in diverse areas and on the basis of various learning content (Belpaeme & Tanaka 2021). The demand for personal tutorial support for learners underlines the potential of scalable, individualizable, and interactive social robots for future learning and teaching support. In addition to these advantages, which apply to chatbots or conversational systems in general (cf. Pérez, Daradoumis & Puig 2020), research groups around social robots in education highlight and demonstrate the advantage of physical presence over screen-based or computer-based technologies. The physical nature and social presence allow real-world interactions and are shown to enhance learning outcomes and motivation (Belpaeme & Tanaka 2021).

However, for more complex learning and teaching scenarios, social robots currently often are unable to meet the high expectations and technical requirements that are imposed on them (e.g., with regard to personalized learning). Researchers emphasize that current models of social robots, such as Pepper and NAO, quickly reach their limits in the inherently complex educational interactions (Belpaeme *et al.* 2018). Schulze *et al.* (2021) point out that learning can only be reduced to a pre-programmed interaction to a limited extent and in individual cases. The studies of social robots in the field of education that they have reviewed therefore have a strong experimental character, which means that previous results can only be transferred to educational practice to a limited extent. Accordingly, besides an attributed versatile future potential, researchers see social robots in the near future primarily in supporting and complementary roles (Schulze *et al.* 2021). Our previous research in this area also shows a differentiated picture. A survey of students on the acceptance of social robots

for learning purposes in higher education has shown that the intention to use them can be predicted and thus increased by determinants (factors) such as improved perceived adaptiveness or social presence of the robot (Guggemos, Seufert & Sonderegger 2020). And in a practice-oriented interview study, today's widely used models of social robots are perceived to be not mature enough in terms of a natural interactive conversation when it comes to their use in the educational context (Sonderegger, Guggemos & Seufert 2022).

In order to achieve the quality of learning and scenarios desirable from an educational perspective, adaptive and interactive elements are essential parts of robot-assisted or, more generally, technology-enhanced learning. According to the ICAP framework, *interactive* (dialoguing) learning as a form of engagement represents the highest quality level of learning (e.g., learning with deeper understanding), after *passive* (receiving), *active* (manipulating), and *constructive* (generating) learning (Chi & Wylie 2014). The authors emphasize, that the ICAP framework can be used as a guide for instructional design and computer-based learning environments (Chi & Wylie 2014). Scalable learning in discourse (interactive mode of engagement) is a central future potential of social robots and conversational technologies in education, not only but even more so at higher levels of education due to its complexity. At the same time, the limited social-interactive capabilities are still the weak points of today's systems. However, technological advances in natural language processing (NLP) open up untapped potential. According to Foster (2019), most developers of social robots tend to employ simple rule- or template-based methods for language generation and make no use of the full possibilities of state-of-the-art natural language generation technologies due to the fact that there are numerous other technical challenges in the development of social robots. But since face-to-face dialogue is the most basic, richest, and most flexible form of human communication, it is essential, especially for its use in education, that linguistic conversation is fluid and flexible as an important part of the interaction (Foster 2019).

The aim of this paper is to address this challenge by exploring the potential of modern generative language models for interactive learning with social robots based on the following research question:

How can interactive learning with social robots for educational purposes be enhanced based on state-of-the-art generative language models?

Based on this research goal we present an approach to evolve social robots towards interactive learning by utilizing generative language models. By integrating the state-of-the-art generative learning model GPT-3 (Generative pre-trained transformer) by OpenAI (Brown *et al.* 2020) over a web API, we illustrate a practical use case of an interactive robot tutor. Based on GPT-3, the robot is more flexible and adaptive in terms of his conversation skills and can for example answer open-ended questions (Dale 2021). AI models like GPT-3 are pre-trained based on huge amounts of data and can be fine-tuned for a wide range of tasks (e.g., question answering, information extraction, sentiment analysis), thus they are also referred to as pre-trained models or foundation models (cf. Bommasani *et al.* 2021). Given the significant technological advances, the simplicity of fine-tuning for users based on limited data, and the resulting increased scalability, such foundation models enable new and interesting applications, especially in education (Bommasani *et al.* 2021).

With the proposed approach, this paper contributes to application-oriented research in this area by providing an understanding of the possibilities and relevance of generative language models in the field of educational technologies. Illustrating the approach in the context of the presented use case can further help researchers and practitioners gain insight into the current and future opportunities that generative language models offer for robot-supported education.

2. INTERACTIVE ROBOT TUTOR MODEL

Based on the theoretical foundation discussed in the introduction we propose an interactive robot tutor model (Figure 1). The model follows a high-level learner-centered and task-oriented view and highlights the relevance of the underlying technologies. On the basis of the activities of a robot tutor and the learner's mode of engagement according to the ICAP framework (Chi & Wylie 2014) on the horizontal axis, the presented model lays out the possibilities of generative language models by structuring on the vertical axis along the maturity of the underlying chatbot technology or method (Cahn 2017; Adamopoulou & Moussiades 2020).

	Robot explains (explain a topic) Passive learning (learner is listening)	Robot asks (ask questions) Active learning (selecting/answering)	Robot assesses (assess/correct answers) Constructive learning (justifying/reflecting)	Robot interacts (simulate conversations) Interactive learning (dialoguing)
Rule-based (e.g., Template)	<ul style="list-style-type: none"> • Closed domain • Stored explanation (e.g., predefined topics and answers in template)	<ul style="list-style-type: none"> • Closed domain • Closed questions (e.g., single or multiple choice questions)	<ul style="list-style-type: none"> • Closed domain • Closed answers (e.g., single or multiple choice correction)	<ul style="list-style-type: none"> • Closed domain • Linear conversation (e.g., simple predefined rule-based conversation)
Retrieval-based (e.g., LMS)	<ul style="list-style-type: none"> • Closed domain • Available explanation (e.g., topics available in databases or via external services)	<ul style="list-style-type: none"> • Closed domain • Available questions (e.g., questions available in databases or via external services)	<ul style="list-style-type: none"> • Closed domain • Closed answers (e.g., multiple choice or keyword detection)	<ul style="list-style-type: none"> • Closed domain • Linear conversation (e.g., simple directed conversation with predefined responses)
Generative (e.g., GPT-3)	<ul style="list-style-type: none"> • Open domain • Adaptive explanation (e.g., 'all' topics can be explained on different levels of sophistication)	<ul style="list-style-type: none"> • Open domain • Open-ended questions (e.g., different sorts of questions covering 'all' topics on different levels)	<ul style="list-style-type: none"> • Open domain • Open-ended answers (e.g., assess open answers by the learner with adaptive feedback)	<ul style="list-style-type: none"> • Open domain • Free conversation (e.g., 'free' conversation with generation of new responses)

Figure 1. Interactive robot tutor model

The robot in the role of a tutor and the social context of a one-to-one interaction can approach a topic or learning subject from different directions. He can explain a topic based on a request or a question (1), ask questions to the learner (2), assess answers from the learner and provide feedback (3) and he can try to discuss a topic interactively with the learner (4). These different types of robot activities and the corresponding modes of learner engagement are independent of the type of knowledge being taught and may include factual, procedural, conceptual, or metacognitive knowledge (cf. Anderson & Krathwohl 2001). But not all methods of language generation or dialoguing offer the same possibilities in terms of topic choice, type, and adaptivity of the interaction elements.

- (1) Passive learning scenario in which the robot explains: rule-based and retrieval-based explanations allow predefined answers and explanations and thus fewer errors (e.g., teacher-inserted answer), whereas generative-based explanations allow domain-open explanations at different levels of difficulty.
- (2) Active learning scenario in which the robot asks closed questions answered by the learner: questions can be retrieved from external services (e.g., from an LMS or open access question databases) offering predefined questions and answer options or they can be generated based on trained or fine-tuned generative language models offering again domain-open and freely adaptive questions.
- (3) Constructive learning scenario in which the learner answers or even rationalizes an open-ended question and receives an adaptive, individualized, corrective feedback from the robot based on the generative language model (not possible in this form with rule-based or retrieval-based methods of language generation).
- (4) Interactive learning scenario in which the learner freely interacts with the robot in dialogue, whereas the robot answers are generated by the generative language model based on the context of the previous dialogue (not possible in this form with rule-based or retrieval-based methods of language generation).

The fields in light grey are combinations that are only possible in a very limited form and basically do not fulfill the pedagogically intended tasks. The fields highlighted in blue thus indicate the building blocks that complement each other in an integrative application and should be optimally combined to enable the learner to learn interactively with the robot tutor in the future. Even if some elements are still technically immature, this combination promises to cover the different engagement levels of learners on the one hand and the tasks of a tutor robot on the other. In order to fulfill its task according to the curriculum, the robot tutor ideally obtains its content and knowledge base from the learning management system (LMS) and uses it to fine-tune the generative language model. As a result, the generative language model is oriented to the predefined learning content from the LMS in terms of both content and level of difficulty but can respond very freely and adaptively to the learner in the interaction with the robot.

3. USE CASE OF AN IMPLEMENTATION OF THE MODEL

In the following section, the presented interactive robot tutoring model is illustrated by a practical use case of a social robot used in higher education. By providing insight into implemented elements and learning scenarios used in workshops on AI and the future of learning, the theoretical ideas are made tangible through concrete examples. In addition to these current applications, their technical implementation and limitations and future ideas will also be discussed. The use case differentiates four main learning scenarios and corresponding technical implementations which are however closely related and integrated within an overall tutorial module.

	Robot explains (explain a topic)	Robot asks (ask questions)	Robot assesses (assess/correct answers)	Robot interacts (simulate conversations)
	Passive learning (learner is listening)	Active learning (selecting/answering)	Constructive learning (justifying/reflecting)	Interactive learning (dialoguing)
Rule-based (e.g., Template)	Template-based introduction to the tutorial module and topic Example: Learner can choose between given topics (e.g. AI, robotics, big data) and a level of difficulty (e.g., beginner, advanced, pro) and is provided with a short introduction into the topic based on illustrations on the robot's tablet. Technical implementation: The content is designed template-based and configured in a robot management platform. Via the dialog management (e.g., in qiChat) one can decide what sorts of questions are answered retrieval-based (e.g., "what is ...") or based on the generative module (e.g., "why ...", "where ...", "how ...").			
Retrieval-based (e.g., LMS)	Robot explains retrieved knowledge from Wikipedia (or if available from LMS) Example dialogue: Learner Question: "What is a social robot?" Robot answer: "A social robot is an autonomous robot that interacts and communicates with humans by following social behaviours and rules attached to its role" (Wikipedia). Robot tablet output: Picture of a social robot Technical implementation: A self developed Python module retrieves a summary for any given 'term' via the Wikipedia web API and outputs it via the robot's speech engine after parsing.			
Generative (e.g., GPT-3)	Robot explains any given topic/question based on GPT-3 text generation Example dialogue: Learner question: "Where could I use a social robot?" Robot answer: "There are many potential applications for social robots. They could be used in healthcare settings to provide companionship to patients, in educational settings to help children learn, or in businesses as customer service representatives." Technical implementation: A self developed Python module sends a prompt containing any question or topic to be explained via web API to a fine-tuned GPT-3 model (fine-tuned to give short robot explanations on different levels) and returns the output of the model via the robot's speech engine.			

Figure 2. Use Case – Interactive Robot Tutor: Explanation Module

As illustrated in Figure 2, the configuration of possible topics, levels of difficulty, and learner activities is controlled in the template via a robot management platform and configured in a rule-based manner. Topics can be predefined in the robot management platform, retrieved based on topics in the learning management system (LMS) or a free selection of topics can be provided to the learner. In a free conversation mode where the learner asks a question, the administrator of the robot configures the dialogue management towards the use of rule-based, retrieval-based, or generative explanation generation based on the form of the question or the availability of the given topic (e.g., if an answer can't be answered by an internal template-based rule or retrieved via Wikipedia API, the explanation will be generated by the GPT-3 generative language model).

	Robot explains (explain a topic)	Robot asks (ask questions)	Robot assesses (assess/correct answers)	Robot interacts (simulate conversations)
	Passive learning (learner is listening)	Active learning (selecting/answering)	Constructive learning (justifying/reflecting)	Interactive learning (dialoguing)
Rule-based	Not applicable			
Retrieval-based (e.g., LMS)	<p>Robot asks multiple-choice questions from the learning management system (LMS)</p> <p>Example dialogue: Topic chosen by learner based on LMS given topics: "Robotics" Robot question (LMS based): "Which components of the robot are responsible for perception?" Robot answer options: "A: sensors, B: actuators or C: effectors" Learner answer: "A" Robot Feedback (LMS based): "Correct. Sensors are responsible for the external perception of robots. An actuator is the counterpart to a sensor and refers to the moving components of a robot."</p> <p>Technical implementation: A Python module retrieves multiple-choice questions including answers via the LMS API and outputs it via the robot's speech engine and presents it on the robot's tablet.</p>			
Generative (e.g., GPT-3)	<p>Robot asks questions to any given topic based on GPT-3 text generation</p> <p>Example dialogue: Topic to freely chosen by learner: "Social robots" Robot question (GPT-3 generated): "What are social robots primarily used for?" Optional robot answer options (GPT-3 generated): A: "To interact with humans" or B: "to perform a task" Learner answer: "A" → Robot Feedback: "Correct."</p> <p>Technical implementation: A self developed Python module sends a prompt containing any chosen topic via web API to a fine-tuned GPT-3 model (fine-tuned to generate short robot questions with two or multiple answer options) and returns the output of the model via the robot's speech engine.</p>			

Figure 3. Use Case – Interactive Robot Tutor: Quiz Module

	Robot explains (explain a topic)	Robot asks (ask questions)	Robot assesses (assess/correct answers)	Robot interacts (simulate conversations)
	Passive learning (learner is listening)	Active learning (selecting/answering)	Constructive learning (justifying/reflecting)	Interactive learning (dialoguing)
Generative (e.g., GPT-3)	<p>Robot assesses open answers/rationales and provides feedback based on fine-tuned GPT-3 model</p> <p>Example dialogue: Topic freely chosen by learner: "Social robots" Open robot question (GPT-3 generated or retrieved via LMS): "What are social robots primarily used for?" Learner answer: "Primarily for fun." Robot feedback (individualized): "Social robots have many purposes. They can be used for fun, but they are also used in healthcare, education, and research."</p> <p>Technical implementation: A self developed Python module sends a prompt containing the robots question and the corresponding learner answer to a fine-tuned GPT-3 model and returns the output of the model via the robot's speech engine. The model is fine-tuned to assess answers to open-ended questions based on a limited number of examples and aims to provide a corrective, constructive, motivational, and individualized feedback, that builds on the learner's response as in the example above. This implementation is limited in part by the fact that the robot's ability to understand and process human speech input is error-sensitive at the current state of the technology. Given a not too complex language, the understanding of the GPT-3 model is already at a high level in most contexts, with task-oriented fine-tuning playing an important role.</p>			

Figure 4. Use Case – Interactive Robot Tutor: Assessment Module

	Robot explains (explain a topic)	Robot asks (ask questions)	Robot assesses (assess/correct answers)	Robot interacts (simulate conversations)
	Passive learning (learner is listening)	Active learning (selecting/answering)	Constructive learning (justifying/reflecting)	Interactive learning (dialoguing)
Rule-based	Not applicable			
Retrieval-based	Not applicable			
Generative (e.g., GPT-3)	<p>Robot freely interacts with the learner based on GPT-3 text generation</p> <p>Example dialogue: <i>Dialogue initiated by learner: "Hello"</i> <i>Robot: "Hello, I am Lexi, an intelligent robot tutor, how can I help you" (Note: answer based on fine-tuning context)</i> <i>Learner: "Tell me about the limitations of AI language models like GPT-3?"</i> <i>Robot: "AI language models have the potential to revolutionize the way we use language. However, they are still in their early stages of development and have some limitations. For example, they may not be able to accurately capture the nuances of human language, or they may not be able to handle complex grammar." (GPT-3 generated)</i> <i>Learner: "What nuances?"</i> <i>Robot: "The nuances of human language can be difficult to capture with a computer model. For example, the way we use idioms, sarcasm or body language can be lost on an AI system." (Note: in this case the robot did understand the context of the short question based on the previous dialogue)</i></p> <p>Technical implementation: A self developed Python module sends a prompt containing the learners speech input and the saved previous dialogue between the learner and the robot to a fine-tuned GPT-3 model and returns the output of the model via the robot's speech engine. The model is fine-tuned to be a helpful and intelligent robot tutor answering to every question and input. The primary limitations of this implementation are the understanding capabilities of the robot model Pepper ranging from the microphone input to the external speech-to-text service required for free text input.</p>			

Figure 5. Use Case – Interactive Robot Tutor: Dialogue Module

The demonstrated use case within Figures 2, 3, 4, and 5 gives a detailed insight into the four modules from a user perspective based on the dialogue and from a technical implementation side. To reunite the picture of the four separate modules, it has to be mentioned that these technically separated tasks are separately accessible via the free conversation mode of the robot, but mainly used integrated into one tutor learning module where the learner chooses between different learning activities after choosing a learning topic. The interaction can thus be initiated by the learner as well as by the robot. In the free conversation mode of the robot, it can be prompted to act via questions. If you ask the robot, for example, whether it can ask you a question, the robot will ask you back on which topic it should ask a question. The choice of topics can be restricted if necessary and defined via an LMS but is in practice used as free as possible in the sense of an open-domain idea. Based on the topic and optionally a selected difficulty level, a method of a self-developed module is executed on the robot by the robot dialogue management, which in turn sends a request to the GPT-3 model via web API containing the topic and a template with training examples (to fine-tune the language model) and receives a model generated question in return. The same procedure applies to the different modules (explanation module figure 2, quiz module in figure 3, assessment module in figure 4, and the dialogue module in figure 5) based on different task configurations and thus different context and task-oriented fine-tunings and training examples.

The primary limitation in this case and generally in the use case is the limited speech recognition capabilities of the robot model Pepper. The robot cannot speak and listen simultaneously and is limited in the ability to detect the end of a speech input to process (e.g., the robot's listening is often interrupted in the middle of user input, as the user tends to take short pauses in speaking). The robot's internal dialog system works well based on the local rule-based dialogues and allows to send parts of the speech input to a remote automated speech recognition (ASR) engine which does not always perform stably. A few words or a compactly formulated sentence can usually be processed, provided the learner is somewhat accustomed to talking to the robot and there is no strong background noise. Given these technical limitations, the current role of the robot is still strongly that of a research and educational tool, helping to understand technological development and to think about its future use in areas such as education. However, in the longer term and based on more advanced robot or educational technology, the role of an AI-based educational tutor, learning partner, or assistant for interactive learning may become more and more realistic.

4. DISCUSSION AND CONCLUSION

This paper explores the potential of generative language models for interactive learning with social robots in the role of a tutor. The proposed preliminary model presents an approach to utilize generative language models such as GPT-3 to progress towards more interactive and engaging forms of learning with social robots. The model connects the different forms of cognitive learning engagement of the learner with the activities of the robot tutor and thus shows the linkage and the possibilities of different technical approaches and especially generative language models. The practical-oriented illustration of the model within the use case of our social educational robot points out the impact and the potential for future applications, but also the limitations and challenges in current use.

Generative language models offer the advantage of a more fluent and interactive conversation with the learner. This has major implications for potential computer-based or robot-based learning applications, as these models allow us to set aside simple keyword detection and predefined answers. Models like GPT-3 understand the context of natural language in a form never achieved before (cf. Elkins & Chun 2020; Floridi & Chiriatti 2020; Dale 2021). The potential of such models in education is accordingly broad; they can be used, for example, as a tool for formative assessment of students in technology-enhanced learning environments or, more generally, for improved real-time analysis of learning data (e.g., text or speech analysis) (Blikstein & Worsley 2016; Phillips *et al.* 2022). Research groups focusing on automated question generation and evaluation based on GPT-3 emphasize the possibility to customize to individual domains and enable higher quality questions from a pedagogical perspective (Bhat *et al.* 2022). Bommasani *et al.* (2021) highlight the greater ability for adaptive and personalized learning through the adaption of foundation models in educational technology. As demonstrated within the use case, the deep understanding of GPT-3 allows the robot tutor to assess the learner's open-ended answers and rationales and consequently provide individualized feedback that can be optionally structured to be corrective, constructive, and motivational based on just a few examples to fine-tune the language model. The model can further be fine-tuned based on specific learning content from the learning management system and thus be used in a more human-controlled way and subject-specific. These example applications from the use case show the potential for more interactive learning activities that promise higher learning quality according to the underlying ICAP framework (Chi & Wylie 2014).

However, the use case, as well as the model itself, also reveal several limitations in terms of today's use and in terms of future challenges and risks in the context of AI-based generative language models. The current use and combination of retrieval-based and generative text generation are primarily limited by the robot's technical limitations in terms of language understanding and processing. As the robot model Pepper relies on external speech-to-text engines and the GPT-3 model can only be used in the form of a web service, the robot requires a stable internet connection and cannot work independently in such a use case. These external services limit the benefits in many ways and create new challenges and risks in terms of data privacy, data storage, data security, dependency, and responsibility. Computationally intensive and therefore often cloud-based AI services make increased digital privacy and data protection even more important (Walsh 2018). Risks and limitations of AI models further include biased or false information, prejudice, transparency, and copyright issues, that raise ethical questions and the demand for a human-in-the-loop (Zanzotto 2019) and explainable AI (XAI), particularly in education (Gunning 2017). Other risks associated with the use of educational AI systems may include the potential misuse to monitor students (e.g., their behavior and communication) or the risk of reproducing prevailing ideologies about teaching and learning towards a homogeneous and centralized system (Blodgett & Madaio 2021). In the case of social robots, the central robot control should not be outsourced to an external service or language model in order to have integration and interaction of the different components in the sense of a social being (e.g., use of cameras for emotion recognition).

As a consequence of these limitations, the use of our social robot in the context of university courses as well as practical workshops is and will be limited to individualization based on the context and not include personalization of the learning content as the LMS interface would allow. In the future, we plan to integrate the individual modules and learning scenarios more effectively, to continuously evaluate them with students and teachers, and to further develop them based on surveys and user experiments. In terms of the practical and technical implementation, it would be interesting to transfer the given use case to more advanced robot models based on the presented model. As for the research outlook, we hope to see, conduct and discuss more application-oriented research that takes the ideas further and explores the potential of generative language models for interactive learning with robots and in various pedagogical application domains.

REFERENCES

- Adamopoulou, E. & Moussiades, L., 2020, 'Chatbots: History, technology, and applications', *Machine Learning with Applications* 2(53), 100006.
- Anderson, L.W. & Krathwohl, D.R., 2001, *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*, Longman.
- Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B. & Tanaka, F., 2018, 'Social robots for education: A review', *Science Robotics* 3(21), eaat5954.
- Belpaeme, T. & Tanaka, F., 2021, 'Social Robots as Educators', *OECD Digital Education Outlook 2021 Pushing the Frontiers with Artificial Intelligence, Blockchain and Robots: Pushing the Frontiers with Artificial Intelligence, Blockchain and Robots* 2021, 143.
- Bhat, S., Nguyen, H.A., Moore, S., Stamper, J., Sakr, M. & Nyberg, E., 2022, 'Towards Automated Generation and Evaluation of Questions in Educational Domains', *Proceedings of the 15th International Conference on Educational Data Mining*, 701–704. 2022.
- Blikstein, P. & Worsley, M., 2016, 'Multimodal Learning Analytics and Education Data Mining: using computational technologies to measure complex learning tasks', *Journal of Learning Analytics* 3(2), 220–238.
- Blodgett, S.L. & Madaio, M., 2021, *Risks of AI Foundation Models in Education*, from <http://arxiv.org/pdf/2110.10024v1>.
- Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S. & Arx, S.v., et al., 2021, *On the Opportunities and Risks of Foundation Models*, from <http://arxiv.org/pdf/2108.07258v3>.
- Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. & Dhariwal, P., et al., 2020, 'Language models are few-shot learners', *arXiv preprint arXiv:2005.14165* 2020.
- Cahn, J., 2017, *CHATBOT: Architecture, Design, & Development*.
- Chi, M.T.H. & Wylie, R., 2014, 'The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes', *Educational Psychologist* 49(4), 219–243.
- Dale, R., 2021, 'GPT-3: What's it good for?', *Natural Language Engineering* 27(1), 113–118.
- Elkins, K. & Chun, J., 2020, 'Can GPT-3 Pass a Writer's Turing Test?', *Journal of Cultural Analytics* 5(2).
- Floridi, L. & Chiriatti, M., 2020, 'GPT-3: Its Nature, Scope, Limits, and Consequences', *Minds and Machines* 30(4), 681–694.
- Foster, M.E., 2019, 'Natural language generation for social robotics: opportunities and challenges', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 374(1771), 20180027.
- Guggemos, J., Seufert, S. & Sonderegger, S., 2020, 'Humanoid robots in higher education: Evaluating the acceptance of Pepper in the context of an academic writing course using the UTAUT', *British Journal of Educational Technology* 10(6), 408.
- Gunning, D., 2017, 'Explainable artificial intelligence (xai)', *Defense Advanced Research Projects Agency (DARPA)*, nd Web 2, 2.
- Pérez, J.Q., Daradoumis, T. & Puig, J.M.M., 2020, 'Rediscovering the use of chatbots in education: A systematic literature review', *Computer Applications in Engineering Education* 28(6), 1549–1565.
- Phillips, T., Saleh, A., Glazewski, K.D., Hmelo-Silver, C.E., Mott, B. & Lester, J.C., 2022, 'Exploring the use of GPT-3 as a tool for evaluating text-based collaborative discourse', *Companion Proceedings of the 12th 2022*, 54.
- Schulze, H., Bendel, O., Schubert, M., Binswanger, M., Simmler, M. & Reimer, R.T.D., et al., 2021, *Soziale Roboter, Empathie und Emotionen - Eine Untersuchung aus interdisziplinärer Perspektive*.
- Sonderegger, S., Guggemos, J. & Seufert, S., 2022, 'How Social Robots Can Facilitate Teaching Quality – Findings from an Explorative Interview Study', in W. Lopuschitz, et al. (eds.), *Robotics in Education*, Cham, 2022, pp. 99–112.
- Walsh, T., 2018, *It's alive: Wie künstliche Intelligenz unser Leben verändern wird*, Edition Körber.
- Zanzotto, F.M., 2019, 'Human-in-the-loop Artificial Intelligence', *Journal of Artificial Intelligence Research* 64, 243–252.