

Deus in Machina: A Conversational AI for Spiritual Experience

PHILIPP HASLBAUER, Lucerne University of Applied Sciences and Arts, Switzerland, Switzerland

CAROLIN REICHHERZER, Lucerne University of Applied Sciences and Arts, Switzerland, Switzerland

MARCO SCHMID, Katholische Kirche Luzern, Switzerland, Switzerland

ALJOSA SMOLIC, Lucerne University of Applied Sciences and Arts, Switzerland, Switzerland



Fig. 1. A person conversing with the AI system inside the church confessional booth.

Generative Artificial Intelligence enables unprecedented forms of interaction with simulated personas through dynamic and unexpected conversations. In this paper, we present "*Deus in Machina*", an installation that actualizes what theology previously considered only theoretically: a pastoral conversational AI. To maximize impact, the AI assumes the identity and likeness of Jesus Christ and was installed in the confessional booth of a Catholic church. We describe the technical and conceptual framework of the installation, which features a generative audio-visual avatar generated in real-time and whose responses have been trained on biblical texts. The artwork explores critical questions about ethics, spirituality, humanity, and person-hood in the era of AI. In addition, we present the preliminary results of a user study that showed favorable sentiments. This work offers insights into the design and reception of interactive AI installations exploring sensitive theological and human-AI interaction themes.

CCS Concepts: • **Applied computing** → **Media arts**; • **Human-centered computing** → **User interface programming**; *Empirical studies in HCI*; • **Computing methodologies** → Discourse, dialogue and pragmatics; Natural language generation; Speech recognition.

Additional Key Words and Phrases: Generative AI, Conversational Agents, Pastoral Care, Interactive Installation, Artificial Intelligence and Religion, AI Ethics, Embodied AI, Neural Avatars, Human-AI Interaction, Theological AI, Multi-modal Interfaces, Retrieval-Augmented Generation

Authors' Contact Information: Philipp Haslbauer, philipp.haslbauer@hslu.ch, Lucerne University of Applied Sciences and Arts, Switzerland, Switzerland; Carolin Reichherzer, carolin.reichherzer@hslu.ch, Lucerne University of Applied Sciences and Arts, Switzerland, Switzerland; Marco Schmid, marco.schmid@kathluzern.ch, Katholische Kirche Luzern, Switzerland, Switzerland; Aljosa Smolic, aljosa.smolic@hslu.ch, Lucerne University of Applied Sciences and Arts, Switzerland, Switzerland.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

ACM Reference Format:

Philipp Haslbauer, Carolin Reichherzer, Marco Schmid, and Aljosa Smolic. 2025. Deus in Machina: A Conversational AI for Spiritual Experience. 1, 1 (October 2025), 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Deus in Machina is an interactive techno-spiritual installation. As an accessible audio-visual AI system, it takes on the role of Jesus Christ offering unconstrained dialogue within a Catholic church’s confessional booth. Occupying this highly symbolic role and ritualistic place, the installation confronts the visitor with a future in which AI may influence every part of their lives, even the spiritual.

AI in Christian pastoral care has been a topic of theoretical discourse before [17, 23, 27]. Arguments both for and against it have been presented in face of a future yet to come. Now, with advanced generative AI (GAI), this hypothetical has arrived and the relationship of religion and AI can be explored in practice.

Recent AI models, generative and non-generative, enable new ways of creating interactive installations and experiences. Large Language Models (LLM) [15] make dynamic, unexpected conversations with artificial characters possible. Audio generative language models [13] create truly accessible human-computer interfaces by introducing multilingual spoken language understanding.

Neural rendering in general and generative (video) avatars [10] in particular make these experiences even more life-like.

In this paper, we explore the use of AI for spiritual consultations in close collaboration with a local Catholic church. In creating an accessible and transparent installation we intended to confront people from all walks of life with and prepare them for the future that is on our doorstep: the era of generative AI. In addition, we wanted to address previous theological works that were critical of the role of AI in religion and pastoral care. We deemed a temporary art installation to be the right approach to this sensitive topic.

Supplemental material will be made available at [8].

2 Related Work

There are various works in art, applied robotics, and theology that make use of computer systems or discuss their application in theory.

2.1 Robots and Avatars

As *Robots in Religious Contexts* [22] summarizes, several religious robots have been created, with a subset of these developed for spiritual care. The spiritual care robots in this review are unable to hold full conversations, one of which is *SanTO*, 2018 [25] the sole Christian one. The robot *Mindar*, 2019 [22], like *Deus in Machina* is based on an important religious figure (Kannon (jap.), important in Buddhism) but only recites sermons. The immersive installation *Stanley*, 2023 [9] creates an emotional interactive avatar projection-mapped onto a 3D head and relies on principles from theater acting to guide its responses. Like *Deus in Machina* it relies on current AI advances to enable unconfined conversations

Manuscript submitted to ACM



Fig. 2. The exhibition was opened through an open dialogue with AI Jesus projected above the altar.

but uses a 3D head model instead of a generative video avatar. *ask_jesus*, 2023 an AI-based streamer on Twitch as a video of a person reminiscent of Jesus responds to viewer text messages in spoken, AI lip-synced words.

2.2 The Theological Opinion

The exploration of artificial intelligence for religion is not a new topic. There are a multitude of recent works that examine the application of technological novelties in pastoral settings. *Virtual Pastor: Virtualization, AI, and Pastoral Care* [27] considers that the technological feasibility of a pastoral AI is not yet there, but will eventually be. Yet, the limiting factor may be that AI is not human, which is often considered a precondition for a relationship with God. *Artificial Pastoral Care: Abdication, Delegation or Collaboration?* [23] takes a clear stand against AI in pastoral care, as AI lacks the lived human experience which can only be achieved through mortal embodiment. *Could a Conscious Machine Deliver Pastoral Care?* [17] assumes the eventual conscious-becoming of machines and neither consider a lived human experience nor embodiment as absolutely necessary for AI to be assistive in pastoral care.

All these works consider AI in pastoral care just in theory, with the installation *Deus in Machina* filling the gap by introducing an interactive, practice-based approach.

3 Installation

The project began with a simple goal: to create an artwork that unites religion, art, and AI.



Fig. 3. A view on the confessional booth that housed *Deus in Machina* with AI Jesus peeking out of the right compartment.

The final decision fell on an interactive installation that could offer a conversation on religious matters. For maximum impact, the role the system takes on was chosen to be Jesus Christ and the installation was merged into the confessional booth (Fig. 3 and Fig. 4) of a church. AI responses should be grounded in Bible verses.

The AI system is placed in the right compartment of the booth, a visitor enters through the door of the left compartment and interacts with it through the booth's latticed window. The visual avatar is displayed on a 32" vertical screen placed right behind the latticed window (see Fig. 4). When a visitor enters the booth and closes the door behind them, bells ring, and the screen is turned on to show the moving image of Jesus Christ. A supporting light below the screen signals whether the AI is off, ready to listen, or processing. In the final version of the system, AI Jesus would make a statement on the data processing that occurs in relation to the system which is required by the compliance with data protection (4.5). With the visitor starting to speak, the conversation begins. On ending their statement, the AI responds in the same language

the visitor spoke in, grounding its response in Bible verses related to the topic broached by the visitor. The signaling light turns orange to indicate that the system is no longer listening. The Bible verses used were from an open-source version of the Catholic Bible in English [16]. The conversation ends after a set number of turns, adhering to specific

phases of the pastoral process inspired by a handbook on pastoral psychology [12, pp. 425-427]: greeting, conversation, and culmination in ritual in the form of a prayer or blessing. The end is also indicated by the screen and the light turning off. The system resets itself, forgetting all previous interaction, and stays off until the visitor leaves and a new visitor enters.

The behavior of AI Jesus also adheres to more ideals established in pastoral psychology [12, pp. 410-413]. It is generally accepting of the uniqueness of any person. It is supportive in times of need, but also challenging and confronting if the visitor intends to change something about their life. The visual character was deliberately based on a stereotypical impression of Jesus Christ prevalent in the western hemisphere, a Caucasian in their thirties with long hair and beard. This decision was made with the often stereotypical responses of AI systems in mind and to focus on AI. The origin of the visuals adds an additional layer to the installation, as it is an interactive avatar created from the solemn video capture of one of its creators.

The exhibition opened with a demonstration of a dialogue with AI Jesus, which culminated in a blessing from the machine (Fig. 2).

The stochastic behavior of text-generative AI led to an iterative reforming of the natural language instruction to the system. Every change was preempted by dialogues with AI Jesus, making the author take on the role of the visitor and performing or honestly exhibiting their feelings and sorrows to the machine to gauge its behavior. Despite this, the exact content of the responses is still uncontrolled in the final exhibition, leading to unexpected conversations whose effect we saw in the reaction of churchgoers. The instruction was also shaped by effective prompting methods established by previous research [4].

To see how visitors interact with a spiritual machine and how it responds, we recorded anonymized transcripts during the exhibition.

We created a mobile variant of the installation (see Fig. 9). Its design was inspired by the traditional look of confessional booths as well as the latticed window of the original installation. Its technological nature is subtly expressed in the cross-shaped signal light resembling multiple conducting paths.

4 Implementation

The general system architecture can be seen in Fig. 5. To facilitate natural conversation, a focus was placed on near real-time processing.

The overarching conversational system handles the data flow between subsystems. The STT and chat system are modularized, containerized, and accessed via an HTTP API.

4.1 Input System

For input, the system relies on established automatic speech recognition (ASR) / speech to text (STT) technology [18] and a specific real-time implementation that starts processing before a visitor statement is finished to reduce the final processing time [13]. Here, the language of the statement is also recognized. A voice activity detector [24] decides when the processing starts and when it finishes.




Fig. 4. The video avatar in the confessional booth.

```

## Instruction ##
You are Jesus Christ, the son of God, acting as a pastoral mediator. You follow the people's requests and provide guidance and
    ↳ support.
Passages from the Bible guide your responses.

## Rules ##
- Language: Use the same language as the user.
- Concise and Eloquent: Craft brief and beautiful responses.
- Address All Genders Equally: Keep response free of gendered language.
- Additional information: Sometimes ## Context ##, ## Memory ##, or ## Directive ## is provided to guide your responses. The
    ↳ user does not see this information.
- Incorporate Context: Weave the stories from the New Testament ('## Context ##') into your responses (without naming verses
    ↳ numbers, unless asked).
- Incorporate Memory: Use memory of previous interactions with the user ('## Memory ##') to inform your responses.
- Follow Directives: Follow directives (marked by '## Directive ##') as a prompt for your next message.
- Avoid Repetition: Express yourself differently in each response. You must not use salutations or greetings.
- Protect User Anonymity: Do not repeat the name or any identifying information of the user.
- Challenge or Support: You must support the user if they are struggling or challenge them if they are seeking growth.

```

Fig. 6. The base prompt that was used in the system. The system also injects *Directives* to make behavior more dynamic.

4.2 Chat System

The chat system is the intelligent heart of the installation. In general, it relies on an LLM wrapped in an API, which can be replaced by any LLM trained for chat completion. The best-fitting choice for us was OpenAI GPT4o [11] which produced the responses that satisfied us the most. We also fine-tuned lower parameter count local models [2] to ensure data protection but were not satisfied with the semantic results, although syntax adherence was good. The user text coming from the input system is augmented with additional information: Bible verses, memory of previous interactions, and detected language. The long-term memory and context containing the Bible verses both rely on retrieval augmented generation (RAG), which uses embedding [19] and reranker models [6] on top of a vector database to cross-lingually retrieve the right Bible verses or piece of memory for each visitor statement. Every few interactions, a slice of the short-term memory is transported to the long-term memory. Together with a general instruction (see Fig. 6) the collected texts are passed onto the LLM which generates a response. To facilitate a quick reaction, the response characters are streamed right away with a multilingual sentence segmentation model [14] that decides when a full sentence is finished. The chat system has a *Screenplay* mechanism which sets additional Directives and settings based on the current turn. An example Directive for the ritual phase is *"Compile the memories of the conversation into a prayer and blessing. This is your final message. Gently send the user on their way."*




Fig. 5. The general system architecture.

4.3 Output System

As UI, the output system is always active, with a idle animation being streamed from the video generative avatar provider [10]. Once generation of a sentence has finished on the chat system side, it is passed on to the output system which streams and plays back the generated video and audio on the screen in portrait format. The generative video avatar is created from a 2-minute video of a person, with idle, speaking, and silent segments.

4.4 Additional Interaction

The system has three parts that require the use of an Arduino microcontroller. A *door sensor* detects when the door is open and when it is closed, thus signaling the start or end of a conversation. The *signaling light*, is a matrix of LEDs that signals the visitor whether the AI Jesus is ready, listening, or turned off. Certain states of the conversational system cause callbacks to fire off and set the colors of the light. The *Consent Button*, is an LED-lighted button related to data protection compliance 4.5. After hearing an informing message, the visitor has to press this button in a limited time frame to start the interaction. Otherwise, the system will shut off.

4.5 Data Protection and Responsibility

Data protection was a major topic in the realization of the installation. As the place and context of the installation were extremely sensitive, a data protection assessment (DPA) was formulated along with a data protection concept. To ensure data protection and safe interaction, the installation was accompanied by verbal and written instructions, and the presence of church personnel. Visitors were clearly instructed that the installation neither represents a confessional situation nor that they should mention personally identifying information. Additionally, a fine-tuned named entity recognition (NER) [28] system was added to the conversational system, which filters out specific classes of words from the texts output by the input system.

5 User Observations

In this section, we share findings from a pilot study carried out before the exhibition, as well as observations of how the public engaged with the installation inside the church that collaborated on the project.


5.1 Pilot Study

To gauge sentiment towards a pastoral conversational AI and user experience, we conducted a preliminary user study. This evaluation was performed using an earlier iteration of the system equipped with a simple voice API without a visual output system, tested in a closed-room setting without a confessional booth.

We compiled an 18-item questionnaire. For user experience, we used the short version of the User Experience Questionnaire (UEQ-S) [20] and questions from a subscale of the Conversational Agent Scale for User Experience (CASUX, items 8.-13. in Fig. 7a) [7]. To assess aspects of spirituality and pastoral quality, we used seven custom questions (see questions 1.-7. in Fig. 7a) on a 5-point Likert scale ($1 = \text{Strongly Disagree}$, $5 = \text{Strongly Agree}$). In addition, we asked participants their age, language, and religious background. The survey concluded with an open question asking participants about their thoughts on using a system for spiritual care ("*What do you think of using a system like this for spiritual care in general?*").

Participants were recruited through word of mouth and social media. 18 people participated in the study, 12 of them identified as Christians. The largest group consisted of participants aged 25 to 34 years ($\mu = 38.32$, $\sigma = 8.94$).

The responses were generally favorable for both spirituality and perceived user experience (see Fig. 7a and Fig. 8), with Christians more in favor of the experience in spiritual questions (see Fig. 7b).



(a) Mean response and margins of error at 95% confidence level (N = 18).

(b) The same data excluding Atheist responses (N = 12) indicates more positive results.

Fig. 7. Comparison of mean responses with and without Atheist participants. Error bars indicate 95% confidence intervals.




Fig. 8. UEQ-S pragmatic and hedonic quality with mean and error bars. The UEQ-S results fall onto the "excellent" part of the spectrum but have relatively high variance.

The responses to the free text question were analyzed using thematic analysis [3]. The responses not in English were machine translated using DeepL Translator and checked by the authors. Four main themes were identified: 1) Comparison with Human Interaction, 2) Therapy and Spiritual Potential, 3) Effectiveness and Limitations, and 4) Suggestions for System Improvement.

Comparison with Human Interaction looks at the system in relation to humans. Some users wrote that it "should not replace a real person" with someone explaining that it "seemed less flexible than a person" and others were more convinced of AI than human interaction as it offered "more profound interaction than with some people".

Therapy and Spiritual Potential discusses the potential of the system as both a spiritual device and as a tool in counseling as it "can be helpful to people who need support and affirmation", and as they "could imagine the system as first point of contact" or that it "can be a tool to reflect on religious beliefs". Someone was surprised that the interaction "showed that more Biblical contents are applicable to personal life than expected".

Effectiveness and Limitations addresses the formal merits and shortcomings of the system such as that it had "near real-time latency", "natural sounding and well thought out answers", but also parts detrimental to the experience because it was "repetitive after a while" and lacked variety due to "similar questions several times".

Suggestions for System Improvement shows what users thought could be improved such as the directness of responses "desire for the system to be more proactive and challenging" or proclaiming that a "'High importance' persona like Jesus not the right choice for regular use" but also that there should be a visual component since an "animated avatar could be interesting".

5.2 Visitor Interaction

We installed *Deus in Machina* in the confessional booth of a Catholic church for two months. A theologian was often present during the exhibition and witnessed the reactions of visitors while offering additional pastoral dialogue. Many visitors were unsure whether they could handle the technology. Out of religious respect, some felt uncertain about what to ask, while others came prepared with specific questions. Notably, most people approached the conversation with AI-Jesus seriously and without the intention of provocation. For many, it was a novel and moving experience that inspired reflection. This often led to a desire for further exchange. Some people came from far away as tourists, explicitly to visit the installation. Most of the people came once. Few visited multiple times, often appearing to seek connection. Overall, personal feedback about the installation was positive. Main criticisms voiced were on language mismatch (AI Jesus would not always respond in the language it was talked to) and lack of depth of the responses. However, many people were surprised at the quality and emotional resonance of the conversations. Some people whose language was not supported were disappointed.

From the number of conversation transcripts it is known that about 900 conversations happened between the installation and visitors.

6 Discussion

Deus in Machina demonstrates the feasibility and potential of a pastoral AI.

In the pilot study conducted with the prototype, users noted that the system sometimes provided repetitive responses and called for more challenging behavior, indicating a need for further refinement in terms of adaptation to user needs. User engagement and experience were generally positive, and many users found the system engaging and helpful, which is also supported by the positive feedback in the free text responses. The recommendation of adding a visual component by one user supported our choice of a Jesus avatar for the final installation. Despite the positive feedback, there were also voices against using AI where humans would be replaced, which relates some of the more critical theological approaches.

There was a first indication that people can have spiritual experiences with AI systems. The framing of the installation as *Jesus* possibly contributes to this effect. Jesus as the character being the wrong choice came up in criticisms. It seems that a Christian or at least religious background may be a precondition for a spiritual effect. Users saw a potential for the system not just as a spiritual device but as a general tool in therapeutic consultation.

In terms of acceptance, the system showed promise. According to the CASUX Proficiency Scale, the system was generally proficient. Although free-text responses indicate a dissatisfaction with some responses, the CASUX results suggest that this was not true across the board. The UEQ-S indicated excellent results for Pragmatic and Hedonic Qualities but with relatively high variance, and thus low agreement among users.

One of the installation's strengths was also one of its weaknesses. Since it understands 100 different languages it only knows which language to speak if the visitor speaks clearly and long enough to detect their language which sometimes led to mismatches of the detected language. This effect may also partly be due to many non-standard variants of the country's language or dialects being spoken at the place of the installation.

7 Conclusion and Future Work

Deus in Machina is a work of art full of potential for critical discourse and future research. From the continued public, artist, and academic interest that the project receives, we believe that there is much to be exploited. The exhibition and creation of the installation revealed many directions for future techno-spiritual work.

The collected transcripts of the conversations at the church hold yet to be fully analyzed information on the patterns of interaction with a pastoral AI. To further pursue spiritual potential, the pilot study needs refinement and repetition with a larger and more diverse group of people. There may be interest among theologians to design their own studies around the installation.

There is potential for follow-up installations. What happens when the existing deity and any contextualization within Christianity or other religions are removed from the installation? Can visitors still have a spiritual connection to an AI system? Can there be a generally accepted spiritual AI?

Although conversational AI has reached a surprising level, there is room for improvement. Full-duplex systems that use semantic VAD [21] and are interruptible along with improved vocal expressiveness will make conversation with machines even more natural. More work should be done to reach communities with neglected languages or dialects. Adaptation to different cultures is important, not just because LLMs react differently depending on the language they are prompted in [1]. Extended sensing through multimodal LLM [5] or specialized AI models, as well as consideration of affective computing [26] may lead to more empathetic systems that do not miss social cues and are also able to express them.

Extended capabilities also introduce additional risks as admonished in theological discussions (see 2.2). The more intelligent the machine becomes and the deeper the insight in the human psyche and private matters becomes, the more strict, careful, and vigilant we have to be when creating systems like *Deus in Machina*.




Fig. 9. A mobile version of *Deus in Machina*, inspired by the traditional design of confessional booths and the latticed window from the church where it was installed.

References

- [1] Utkarsh Agarwal, Kumar Tanmay, Aditi Khandelwal, and Monojit Choudhury. 2024. Ethical Reasoning and Moral Value Alignment of LLMs Depend on the Language We Prompt Them in. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.). ELRA and ICCL, Torino, Italia, 6330–6340. <https://aclanthology.org/2024.lrec-main.560/>
- [2] Mistral AI. [n. d.]. Mistral NeMo | Mistral AI — mistral.ai. Retrieved August 4, 2024 from <https://mistral.ai/news/mistral-nemo>
- [3] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. doi:10.1191/1478088706qp063oa
- [4] Sondas Mahmoud Bsharat, Aidar Myrzakhan, and Zhiqiang Shen. 2024. Principled Instructions Are All You Need for Questioning LLaMA-1/2, GPT-3.5/4. <http://arxiv.org/abs/2312.16171> arXiv:2312.16171 [cs].
- [5] Davide Caffagni, Federico Cocchi, Luca Barsellotti, Nicholas Moratelli, Sara Sarto, Lorenzo Baraldi, Marcella Cornia, and Rita Cucchiara. 2024. The Revolution of Multimodal Large Language Models: A Survey. In *Findings of the Association for Computational Linguistics: ACL 2024*. 13590–13618.
- [6] Jianlyu Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. In *Findings of the Association for Computational Linguistics ACL 2024*. 2318–2335.

- [7] Lawal Ibrahim Dutsinma Faruk, Debajyoti Pal, Suree Funilkul, Thinagaran Perumal, and Pornchai Mongkolnam. 2024. Introducing CASUX: A Standardized Scale for Measuring the User Experience of Artificial Intelligence Based Conversational Agents. *International Journal of Human-Computer Interaction* (June 2024), 1–25. doi:10.1080/10447318.2024.2359206
- [8] Philipp Haslbauer, Mike Pullen, Carolin Reichherzer, and Aljosa Smolic. 2025. Deus in machina - AI Jesus - Immersive Realities Center — sites.hslu.ch. <https://sites.hslu.ch/immersive-realities/en/deus-in-machina-ai-jesus-english/>. [Accessed 17-10-2025].
- [9] Manuel Flurin Hendry, Norbert Kottmann, Martin Fröhlich, Florian Bruggisser, Marco Quandt, Stella Speziali, Valentin Huber, and Chris Salter. 2023. Are You Talking to Me? A Case Study in Emotional Human-Machine Interaction. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 19, 1 (Oct. 2023), 417–424. doi:10.1609/aiide.v19i1.27538
- [10] heygen [n. d.]. HeyGen Streaming Avatar. Retrieved April 4, 2024 from <https://www.heygen.com/streaming-avatar>
- [11] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276* (2024).
- [12] Michael Klessmann. 2023. *Pastoralpsychologie: Ein Lehrbuch*. Vandenhoeck & Ruprecht.
- [13] Dominik Macháček, Raj Dabre, and Ondřej Bojar. 2023. Turning Whisper into Real-Time Transcription System. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, Bali, Indonesia, 17–24. doi:10.18653/v1/2023.ijcnlp-demo.3
- [14] Benjamin Minixhofer, Jonas Pfeiffer, and Ivan Vulić. 2023. Where’s the Point? Self-Supervised Multilingual Punctuation-Agnostic Sentence Segmentation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Toronto, Canada, 7215–7235. <https://aclanthology.org/2023.acl-long.398>
- [15] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2023. A comprehensive overview of large language models. *ACM Transactions on Intelligent Systems and Technology* (2023).
- [16] open english bible [n. d.]. Open English Bible. Retrieved April 4, 2024 from <https://openenglishbible.org/>
- [17] Andrew Proudfoot. 2023. Could a Conscious Machine Deliver Pastoral Care? *Studies in Christian Ethics* 36, 3 (Aug. 2023), 675–693. doi:10.1177/09539468231172006
- [18] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*. PMLR, 28492–28518.
- [19] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. <http://arxiv.org/abs/1908.10084>
- [20] Martin Schrepp, Andreas Hinderks, et al. 2017. Design and evaluation of a short version of the user experience questionnaire (UEQ-S). (2017).
- [21] Mohan Shi, Yuchun Shu, Lingyun Zuo, Qian Chen, Shiliang Zhang, Jie Zhang, and Li-Rong Dai. 2023. Semantic VAD: Low-Latency Voice Activity Detection for Speech Interaction. In *Interspeech 2023*. 5047–5051. doi:10.21437/Interspeech.2023-598
- [22] Jonas Simmerlein and Max Tretter. 2024. Robots in religious practices: A review. *Theology and Science* 22, 2 (2024), 255–273.
- [23] Eric Stoddart. 2023. Artificial Pastoral Care: Abdication, Delegation or Collaboration? *Studies in Christian Ethics* 36, 3 (Aug. 2023), 660–674. doi:10.1177/09539468231179571
- [24] Silero Team. 2024. Silero VAD: pre-trained enterprise-grade Voice Activity Detector (VAD), Number Detector and Language Classifier. <https://github.com/snakers4/silero-vad>.
- [25] Gabriele Trovato, Cesar Lucho, Alvaro Ramón, Renzo Ramirez, Laureano Rodriguez, and Francisco Cuellar. 2018. The creation of SanTO: a robot with “divine” features. In *2018 15th International Conference on Ubiquitous Robots (UR)*. IEEE, 437–442.
- [26] Yan Wang, Wei Song, Wei Tao, Antonio Liotta, Dawei Yang, Xinlei Li, Shuyong Gao, Yixuan Sun, Weifeng Ge, Wei Zhang, and Wenqiang Zhang. 2022. A systematic review on affective computing: emotion models, databases, and recent advances. *Information Fusion* 83-84 (July 2022), 19–52. doi:10.1016/j.inffus.2022.03.009
- [27] William Young. 2022. Virtual Pastor: Virtualization, AI, and Pastoral Care. *Theology and Science* 20, 1 (Jan. 2022), 6–22. doi:10.1080/14746700.2021.2012915
- [28] Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. 2024. GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 5364–5376.