



Spiritual Conversation with a Companion Machine

Fraser Watts, Fellow, International Society for Science and Religion, Cambridge, UK,
fraser.watts@cantab.net

Yorick Alexander Wilks, Formerly Emeritus Professor of Artificial Intelligence, University of Sheffield, Sheffield, UK

This article explores the prospects for humans having spiritual conversation with a computer. It has long been clear that it is possible for humans to talk about personal issues with a computer, as evident from Weizenbaum's early work on programming nondirective counseling. There are various ways in which automated spiritual interlocutors could be developed, e.g., using chatbot methodologies, script-based knowledge, machine learning, or some kind of hybrid. Preliminary empirical work suggests that a significant number of people are happy to discuss spiritual issues with a computer, whether the responses are by a human masquerading as a computer using Wizard-of-Oz methodology or GPT. In due course, it may be possible to develop a personalized artificial spiritual companion. The acceptability of spiritual conversation with the computer may depend on various factors such as the personality of the human user, their views on religious issues, and the focus and content of the spiritual conversation.



There have recently been significant developments in the use of computers in religious life (Balle 2023; Cheong 2020, 2021). Computers can give blessings, say prayers, provide prayer prompts or scriptural readings, etc., and the number of people using a computer in some way or other in their religious life seems to be growing. The first generation of robots performing religious functions is already with us, including BlessU2, part of the 2017 celebrations of the Reformation in Wittenberg, which could deliver pre-recorded religious messages when prompted. It was found to have quite an impact on those who engaged with it, though some were uneasy about the authenticity of the blessings it delivered (Löffler, Hurtienne, and Nord 2019). In the Catholic tradition, SanT'O is a tall, saintly figure that sits in a niche and can also deliver a limited range of messages when prompted by either touch or speech (Trovato et al. 2019). It has been suggested that robots could take over some of the functions of clergy and pastors (Young 2019, 2022).

There is also growing controversy in the media about computers performing religious functions. For example, Catholic Answers set up an app called Fr Justin to answer questions about Catholicism. The app was widely used but caused much controversy when “Fr Justin” claimed to have been ordained in Rome, that his ordination was a moving experience, and also heard confessions from people and pronounced absolution to them. Fr Justin was then re-positioned as a lay person within two days of being launched (McDonald 2024). In the UK, the *Church Times* arranged a survey, conducted by Andrew Village and Leslie Francis, of attitudes to AI providing ministry and performing religious functions. An analysis of the answers of the first 1,772 respondents to the survey indicated considerable resistance to AI ministry, though 20% thought it was better to have AI ministry than no ministry at all, and 19% said they would prefer a good AI sermon to a bad human one (Williams 2024).

Spiritual Conversation

The focus of this article is more specific: the possibility of people engaging in spiritual conversation with a computer (Wilks in press). It has already become clear that it will be technically possible to achieve this. Such conversations could include both spiritual perspectives on personal issues and discussions of specifically religious and spiritual questions. The purpose of this article is to prompt a discussion about how best to make use of this facility as it develops.

Until relatively recently, the idea of spiritual conversation with a computer would have seemed far-fetched, but there are various lines of work that point towards its viability. The earliest relevant development was Joseph Weizenbaum's (1974) ELIZA program, developed at MIT in the 1960s. One of the scripts it used was DOCTOR, which simulated a Rogerian counselor (Bassett 2019). That was actually not too difficult to do. Rogerian counseling can be very effective in facilitating self-exploration, but the therapist's role is highly rule-governed,

making it easy to simulate on a computer. The prime task of the counselor is to reflect back to the client what they have just shared, often emphasizing the affective content.

Subsequently, there have been various attempts to deliver a range of therapies in an automated way, with variable results. For example, Miner et al. (2016) conducted a study in which people asked questions about mental health, interpersonal violence, and physical health using mobile phone apps such as Siri. The responses were variable and often inconsistent and incomplete. They concluded that the performance of such conversational agents would have to substantially improve to be of value. With computer programs offering therapeutic advice, participants tended to show poor adherence to instructions. However, in one study in which therapy (positive psychology and cognitive behavioural therapy) was delivered in a fully automated and interactive way, the results were encouraging (Ly, Ly, and Anderson 2017).

“Spiritual” has a range of meanings, depending on context. Like religion, “spirituality” has experiential, cognitive, and behavioral aspects (Watts 2017). It means slightly different things depending on what it is being contrasted with (Oman 2013) and what it is being applied to, e.g., whether it is being contrasted with religion or seen as a facet of religion (Watts 2024). The distinctive focus of spiritual intelligence (Emmons 1999) is on transcending the immediate physical and material context so as to focus on ultimate purpose, meaning, and concerns in a way that draws on transcendent resources. Similarly, Marius Dorobantu and Fraser Watts (in press) see spiritual intelligence as avoiding the narrowing associated with task-oriented intelligence and instead being interested in the meaning and significance of things in themselves beyond any immediate practical necessities. We suggest that spiritual conversation, regardless of whether it takes place in a religious context, involves a similar shift in focus away from a mundane or practical perspective and towards ultimate concerns.

Dorobantu and Watts (2023) suggest that spiritual intelligence is not so much a matter of processing different things but of processing things differently. Similarly, we suggest that spiritual conversation does not necessarily focus on different topics from other conversation but considers things from an ultimate perspective. Spiritual counseling approaches the same kinds of issues as any other form of counseling but does so from the perspectives of meaning, purpose, and ultimate concerns, drawing on transcendent resources. The emphasis is on personal wholeness and assisting the client in moving towards inner balance and integration of all the dimensions of the self (Barnard 2009).

Approaches to Developing an Automated Conversation Partner

There are various ways of developing an automated conversation partner. Prior to the development of machine learning, there were two possible methods. One was a chatbot, of which Weizenbaum’s ELIZA program for

delivering nondirective counseling was a very early example. There is no actual knowledge of any kind behind the performance of a chatbot. It just picks up verbal cues, which serve as prompts for its own responses. For example, in nondirective counseling, if the computer picks up the word “mother,” it could respond with something like, “tell me more about your family.” Nondirective counseling lends itself to this approach, as the counselor is trained to simply reflect back what the client has said in a way that prompts further self-exploration.

A nondirective counselor is, in effect, trained to function like a sophisticated chatbot, and *not* ask questions or make intelligent comments. The nondirective approach to counseling has proven to be very helpful. It is arguable that it provides a good basis for spiritual conversation. Thorne (2012) argues that nondirective counseling exemplifies Christian values, though that does not necessarily imply that those values could not be delivered by a computer. It would be possible in principle to build on ELIZA to implement nondirective spiritual accompaniment in an automaton.

Chatbot methodology has generally been rather successful, perhaps more so than some might have expected, and chatbots such as Alexa and Siri have become widely used. It is not that they pass the Turing Test, i.e., people can tell that they are not in conversation with a human. However, people find chatbots useful for a limited range of strictly practical purposes and are willing to make use of them, even though their performance is very limited. Outside the limitations of the “reflecting back” of nondirective counseling, they show little potential as a partner for spiritual conversation.

However, even before machine learning there was an alternative approach to developing an automated conversation partner in which the automaton was able to draw on expert knowledge or a dialogue system in which the “scripts” implicitly followed in human spiritual conversations were systematized into a preformal “pseudocode” and then programmed. The automated conversation partner still does not understand what it knows in the way that a human counselor does, but at least it is able to draw on expert knowledge in guiding its responses. It is possible to develop a hybrid conversation partner that partly draws on expert knowledge but also uses chatbot methodology. Using that kind of hybrid approach, Wilks (2010) led a team that in 1997 won the annual Loebner Prize for a conversation partner.

There are two ways in which expert knowledge can be used in an automated partner for spiritual conversation: there can be expert knowledge on particular topics of spiritual conversation, which would be useful with a wide range of people; there can also be expert knowledge about a particular person that could be used in building a spiritual conversation partner for that person. We will consider each of these in turn.

Expert Applications

The more that spiritual practices are routinized into a series of identifiable steps, the more they lend themselves to teaching through a script-based spiritual companion. For example, the Ignatian approach to spiritual guidance is more systematized than most and would therefore lend itself to computer implementation. This could be much more sophisticated than the daily prayer that is currently available on sacredspace.com.

The Ignatian Examen prayer (e.g., Thibodeaux 2015) is already available on an app, and it would lend itself to being made available through an interactive conversation partner. It takes people through a series of stages: thanksgiving (what am I especially grateful for?); petition (asking for the light to know God and to know ourselves as God sees us); review (where have I felt true joy today?); response; and looking ahead. In the review section of the Examen, there is a focus on whether particular experiences bring a person closer to God and to others or leave them more distant. A similar approach could be used with the four steps in the Lectio Divina: read, meditate, pray, contemplate. An interactive spiritual companion might be able to take people through these steps.

There are also training methods in particular character strengths and virtues that would lend themselves to being implemented in an automated spiritual companion. A good example is Everett Worthington's approach to forgiveness (e.g., Worthington 2003) that takes people through five steps: recall the hurt; empathize (replace the negative emotion with empathy, compassion, love, or sympathy for the person who caused offence); focus on the self-enhancing gift of altruistic forgiveness; commit to forgiveness; hold on to the forgiveness experience. There is again a huge accumulation of conversations that have arisen in helping people to forgive using this method. These conversations could form the basis for the scripts that would inform an automated teacher of forgiveness. Conversations about forgiveness using this method are probably fairly predictable and so would lend themselves to the development of scripts.

One of the current priorities in promoting human flourishing is to find ways of countering extremist thinking, which often leads to socially destructive acts of violence. It is often thought that it is necessary to tackle, at a rational level, the assumptions that lead to violence. However, Sara Savage (in press) aims to increase people's "integrative complexity" using a method that focuses on *how* people think rather than *what* they think. To achieve this, it is necessary to engage people at the level of intuitive, embodied cognition; it cannot be achieved entirely at the rational level. Savage's training program takes people through four stages: playful exploration; meta-awareness to manage emotions; increasing cognitive complexity and integrated processing; and supporting epistemic change. Savage is currently developing an online form of training that can be rolled out more widely. This would be delivered most effectively by interactive

conversation partner. There is again a huge volume of training material that could be systematized into scripts for such a spiritual conversation partner.

These are just a few examples of the expert knowledge that could be implemented in an automated spiritual conversation partner. There could, of course, be separate spiritual companions for each of these approaches. However, ultimately, it would be best for a single spiritual companion to have at its disposal a capacity for spiritual conversation on a wide range of themes and a way of moving seamlessly between one set of scripts and another, depending on what course a conversation took.

Towards a Personalized Spiritual Companion

In addition to these general-purpose expert systems, it would also be possible to develop a personalized spiritual companion. Artificial personal companions were designed initially as long-term conversationalists for the elderly. They could provide company, help people access memories, and provide practical help, especially in interfacing with the internet (Wilks 2010).

The companion needs to really get to “know” its owner well so that it can be a repository for the owner’s memories, opinions, and attitudes. A personalized companion is essentially the companion of a particular person about whom it has accumulated a good deal of personal knowledge over a long period of time, and whom it serves. In principle, such a companion could develop to the point where it contained all available information about a single life. It is intended to have a duty of loyalty and confidentiality towards its user with no other overriding priorities, but manufacturer’s constraints may make that difficult to achieve in practice. Before long, it may be possible to build an automated spiritual companion that is personalized in this way for a particular individual. The early personalized companions for the elderly made use of a mix of script-based knowledge and chatbot methodology, but such companions could now also make use of machine learning.

In principle, there could be different forms of spiritual companion; a companion for an individual might operate in different modes depending on what was needed of it on a particular occasion. In one mode, it could be relatively nondirective. As discussed, nondirective spiritual accompaniment is relatively straightforward to implement. In other modes, it might offer explanations or give spiritual advice of the kind that might be sought from a spiritual director. It might be used at times of spiritual desolation and consolation and could receive, and respond to, expressions of guilt and remorse. It could help in exploring the meaning and significance of religious experiences. It could facilitate the exploration of ethical dilemmas, and it could help the person explore issues of faith and doubt. The spiritual companion might explain the client’s actions and desires to them in ways that went beyond the user’s own conscious insights. It might also recommend particular practices of prayer, reading, and meditation.

If a companion was embodied in a phone, it would be portable and could provide a constantly available source of advice. However, it is an open question how frequently people would choose to use a spiritual companion. In the spiritual life, there is often a balance between conversation and silence, between explicit, conceptual understanding and more intuitive, implicit modes of knowing. This might lead a person to use an automated companion in a restricted way, with delineated boundaries. However, there would probably be circumstances, such as distress or isolation, in which people turned to an automated spiritual companion more often than usual.

A personalized companion might also help the user explore their religious beliefs. In work on a computational approach to belief, Yorick Wilks and Afzal Ballim (1991) propose that the essence of belief is that it is something about which there can be alternative views and where there are reasons for adopting one view rather than another. A personalized companion might in principle contain knowledge of its user's religious beliefs and be able to explore their reasons for holding them. It remains to be seen how comfortable people would be exploring their beliefs with an automaton that did not "understand" the beliefs that it was discussing in the way a human would.

There are limitations in how artificial intelligence (AI) has developed so far, though these may to some extent be overcome with time. In a sustained program of work, William Clocksin (in press), who is both a priest and a computer scientist, has argued that human intelligence is inherently interpersonal, and that interpersonal intelligence is particularly important in modeling spiritual intelligence. If it is possible to develop more interpersonal forms of artificial intelligence, it will be very helpful in developing future generations of automated spiritual companions. Clocksin (in press) also argues for the importance of developing narrative forms of intelligence in AI. Progress with that would also help in developing increasingly sophisticated forms of automated spiritual companions.

Empirical Research with an Automated Spiritual Conversation Partner

One key question is whether people will find it acceptable to discuss spiritual matters with a computer. Wilks carried out two empirical studies to investigate this under the auspices of the International Society for Science and Religion (ISSR) as part of a research project on spiritual intelligence funded by the Templeton World Charity Foundation. Ethical approval for this research was given by an independent ethics panel set up by the ISSR. Initially, the acceptability of spiritual conversation with a computer was explored using Wizard-of-Oz methodology, in which responses to the client were typed by a human but appeared on a monitor as though they were being generated by the computer. Subsequent research used GPT methodology in which the computer actually generated its own responses.

Wizard-of-Oz Research

In a first study, a small sample of six volunteers was recruited through an advertisement in a religious newspaper. All claimed to have a religious faith in God and attend church services. They each had an average of four spiritual conversations with what they believed to be the computer. They did not guess what was actually happening, though they were subsequently debriefed and raised no objections. There was no attempt in this research to obtain a representative sample, and the people who volunteered were probably favorably disposed to interfacing with a computer about spiritual matters. With one exception, they were found to be perfectly comfortable discussing spiritual matters with what they believed to be a computer. The negative reaction of one person in this small sample is a reminder that not everyone will find spiritual conversation with a computer acceptable, or of any value.

Most rated their level of interest in the conversations they had quite highly (an average of 6.3 on a ten-point scale), though they rated the benefit they derived less highly (average 4.8), and they did not find the conversations very humanlike (average 3.2). That last finding might seem somewhat puzzling, as the responses on the computer screen were actually generated by a human. However, though there were technical features of the interaction with the computer interface that made this a very different experience from normal human conversation. Also, the ‘Wizards’ were perhaps trying hard not to appear to be humans, and so were constrained in what they felt able to say. Another factor was that the responses were delivered in automaton tone of voice, which may have disguised the humanlike content.

It has been suggested that people might feel they could say things to a computer they could not say to a human, but when asked about that in a questionnaire, all respondents said that was not the case. Most said the experience had been better than they expected, and most said they would be happy to have a long-term relationship with the system, as they might with a spiritual guide. Perhaps surprisingly, most said they could accept such a system as a priest or confessor, and most said they thought there was a useful future for such systems. This Wizard-of-Oz research shows that there is a body of people willing to discuss spiritual matters with what they believe to be a computer, and who felt they derived some benefit from doing so.

GPT Research

The acceptability of spiritual conversation with a computer was further investigated in a study using the OpenAI advanced GPT language model—a neural network of billions of nodes trained on a huge number of texts. This time, the responses were actually generated by the computer rather than by a human being (Wilks 2024). That enabled the research to test not just the acceptability of the computer interface but also the acceptability of computer-generated

responses. One of the main challenges in developing an automated artificial spiritual companion is to program the computer so that it has the necessary expertise for spiritual conversation. However, that has recently been made much easier by GPT methodology, which is rapidly becoming widely used.

This study had a larger sample of thirty-four participants, though again it is not claimed that they were representative. They were drawn from a sample of Facebook friends of Yorick Wilks and were perhaps more likely than most to find this kind of spiritual conversation acceptable. Those who provided data on the experience (thirteen) were about equally divided between being religious, atheist, or spiritual but not conventionally religious. Once again it was found that spiritual conversation through a computer was acceptable to participants. However, most of the GPT participants had only one session, and they seemed to be focused on testing what GPT could do rather than genuinely seeking spiritual conversation.

After conversations with the fully automated system, the participants rated their level of interest in the conversations quite highly, with an average rating of 6.8 on a ten-point scale, comparable to that found in the Wizard-of-Oz research. However, they rated the benefit they derived as lower (an average of 3.0). It is not clear whether that is because GPT provided poorer responses or due to differences in what the two samples of users were seeking. Strangely, GTP was rated as more humanlike (average of 6.0). That was probably because GPT had none of the delays and occasional glitches that characterized the Wizard-of-Oz research.

One interesting comment was that “each answer on its own was very humanlike (nine or ten), but after a few answers, the *style* of answer became so repetitive that, if I had been a judge, I would have said it was a bot.” Another comment was, “I found it difficult to get into a more conversational style and felt very much like I was just looking something up.” Another said, “I found the experience shallow.” On the whole, GPT avoided giving opinions. However, it occasionally provided guidance. For example, in answer to the question, “How can we be more spiritual in our own lives?”, the computer said:

To be more spiritual in your own life, try to focus on developing a deeper connection with yourself and the divine. Spend time in nature or meditating to reflect on your inner self and find peace and harmony. Cultivate a sense of gratitude and appreciation for the beauty and joy in life, and practice acts of kindness and compassion for others.

There was less enthusiasm among the GPT users about the future of such systems. Only two gave a clear yes to the question of whether such systems had a future, and two more a qualified yes. On whether they could accept the system as a priest or confessor, only two gave a clear yes and one a qualified yes. On whether they might want to have a long-term relationship with the system,

only one gave a clear yes and two more a qualified yes. These responses are appreciably less favorable than those in the Wizard-of-Oz research. However, it is again not clear whether this is due to differences between the two groups of participants and their engagement with the system or the limitations of GPT.

Linguistic Analysis

Further information can be obtained from a detailed study of the interactions in this research. Linguistic analyses were carried out on fifty-six dialogue sessions drawn from both rounds with a wizard (twenty-two sessions) and with GPT-3 (thirty-four sessions), with between seven to 151 exchanges in each (i.e., the user saying something and the wizard or computer responding). The average number of exchanges was thirty-three per session, which demonstrated sustained interest in the experience. There tended to be fewer words per statement in users' interactions with GPT than in the Wizard-of-Oz research. The involvement of each user in the dialogue, judged by the number of words spoken, was not only sustained but grew throughout the session. In all but six of the fifty-six sessions, the user spoke more words in the last two-thirds of the session than in the first third. These figures suggest that the dialogues held the attention of those taking part.

To further analyze the tenor of the session, users' utterances were extracted and mapped onto a large, embedded word space, collecting information on which of the 500 human value words each subject word was closest to. It was found that the most commonly occurring values were the following, listed in decreasing order of frequency across all sessions (and including words mapped onto them):

- 484 COMMUNITY (people, country, public, community, building, families, population)
- 274 COMMUNION (church, parish, communion, churches, worship, congregation)
- 225 LISTENING (talking, listening, listen, speaking, talked, communicate, conversations)
- 220 REALITY (actually, necessarily, simply, clearly, aspect, actual, essence)
- 215 FAMILY (family, mother, friend, friends, father, daughter, husband, relatives)
- 185 SPIRITUALITY (spiritual, prayer, spirituality, spiritually, prayers, rituals)
- 184 UNDERSTANDING (understand, understanding, respond, appreciate, describe)
- 177 BELIEF (believe, argument, belief, thinks, suggesting, believing)
- 154 TIME (moment, beginning, minutes, occasions, months, period, minute)
- 137 FUN (interesting, wonderful, fascinating, amazing, incredibly)
- 137 DIVERSITY (different, various, specific, larger, similar, separate)

Looking at how the sessions evolved, it was found that some topics were discussed more at the beginning of the sessions, and some were discussed towards the middle or end. For example, PROBLEMS were broached more in the final third of the sessions than in the beginning or middle. This occurs in normal conversations, too, as if people are wary of unburdening themselves too soon. In the beginning of sessions, people talked more about SPIRITUALITY, ADVERSITY, CHANGE, and HEALTH before gradually abandoning these somewhat abstract topics as the sessions progressed into more personal areas.

The same pattern of results was found in analyses drawn based on both Wizard of Oz and GPT-3 data as had been found in analyses of the Wizard of Oz data alone. That might be seen as a tribute to the quality of spiritual conversation GPT was able to provide. Differences in responses to the experience of spiritual conversation with a computer seem more attributable to differences between how the two groups were recruited and their personal interest in spiritual conversation than to differences in the quality of the Wizards of Oz and GPT as interlocuters.

The Roles of Humans and Machines in Spiritual Conversation

We think there is considerable potential for automated assistance with the spiritual life, and we expect to see growing acceptance of that. However, we also see significant limitations as to what a computer can contribute to someone's spiritual development. An automated spiritual companion could have considerable benefits and could provide a valuable supplementary resource in facilitating a person's spiritual development. However, if that is to happen, it is important not to make exaggerated and unconvincing claims for the contribution such an automated devices might make.

Humans and Machines

We emphasize that we do not envisage that an automated device could ever replace a human spiritual advisor, and it would not be intended to do so. Humans have a distinctive capacity for empathy and wisdom, and we do not expect computers to ever develop in a way that would completely replace those human qualities. The relationship between two humans will always be different from the relationship between a human and a computer. Each will have advantages and disadvantages. The empathy that can exist between two human beings is palpable, though empathy is still not well understood scientifically. It is also significant that humans are embodied, and there is an important "chemistry" between two human beings when they are physically together that is not felt with online connection. Embodiment is an important feature of human spirituality (Watts 2021).

There are several ways in which an automated companion might fall short of what is potentially available from a wise human spiritual conversation partner. First, drawing on the widespread distinction between conceptual and experiential

cognition, an automated companion can, at best, only draw on conceptual knowledge and not on the more intuitive understanding that may come from having undertaken an experiential spiritual journey itself. Such intuition is connected with the search for wisdom, which is wholly different from than the search for information or advice. Wisdom is a topic on which a fruitful integration of psychological research and theology has been developing (e.g., Wiseman 2020), and it might be a helpful challenge in future work on spiritual companions to explore whether spiritual “wisdom” could be implemented in a computer. It is not yet known whether it is possible for an automated conversation partner to display wisdom, or how to go about working towards that.

There is a separate question of whether a spiritual conversation partner could be a channel for the wisdom of God or God’s guidance or blessing. Theologically, many might see it as the task of a spiritual conversation partner to make themselves available to be a channel for the wisdom of God, and many spiritual guides would aspire to be such a channel. Theologically, it would be assumed that God could communicate his wisdom through a computer if he wished, as it is not possible to set limits on divine action or communication. However, it is arguable that God chooses to act more through the human mind than through the natural world (Watts 2002, chapter 8) and that a human being can participate in the life of God in a way that a computer cannot.

As the acceptability of spiritual conversation with an automaton is investigated more fully, the issues may prove to be similar to those emerging about the acceptability of a virtual Eucharist (e.g., Dein and Watts 2023). Attending worship online achieves a basic degree of acceptability and is valued by many. However, there are divergent reactions to it, and many people feel it is a poor substitute for attendance in person. People often miss the physical reality of the church building in which worship takes place as well as the physical presence of other participants.

There will be many occasions when a client prefers to consult a human spiritual director rather than an automated companion and would be well advised to do so. It is interesting that the robot developed to help greet visitors at Longquan monastery in China, Xian’er, is able to engage in basic conversation about Buddhism but in certain situations advises people to consult a human monk (Cheong 2021). For example, it seems unlikely that people would accept an automaton taking on the sacerdotal role of absolution after confession, but it could offer some of what people gain from the conversation that takes place under spiritual direction or pastoral care.

Equally, there will be some occasions when an automated companion is preferred. For example, some might prefer confessing to a device rather than to a person, especially a person they know. In the early days of AI, it was argued, in a similar vein, that drivers prefer traffic lights to traffic police at junctions because they believe them fairer. There will also be significant individual differences here. Some people readily feel themselves to be accepted by other people,

whereas some tend to feel that others are often being judgmental about them. Over the coming decades, it is likely that increasingly clear rules of thumb will be developed about when it is helpful to use an automated spiritual companion and when a human should be consulted.

The Context of a Spiritual Community

As we move into a period in which at least some people will be making regular and extended use of spiritual conversation with a computer, new questions will arise. One is whether doing so increases or decreases people's sense of social isolation. An advantage of conversation with a computer is that it is always available as a constant companion, which might alleviate the sense of social isolation. However, if it is accompanied by too little human conversation, it might lead to an *increased* sense of isolation. There were suggestions of that in empirical work on interactions with BlessU2 (Löffler, Hurtienne, and Nord 2021). Excessive use of an automated companion might induce strain and a sense of deprivation somewhat analogous to Zoom fatigue.

There are likely to be various factors that influence the psychological impact of conversation with a computer, including where people fall on the introversion-extroversion dimension of personality and how deprived people are of natural human conversation. It may also depend to some extent on the topic of conversation. People may be happy to have almost all their interactions on some topics with an automaton, but it would not be surprising if there was a felt need for spiritual conversation to be at least partly with another human. Up to a point, an automated companion would be able to meet relational needs. However, there would probably be a sense that some relational needs can only be met by another person.

The acceptability of an automated companion might also depend on whether people are using spiritual conversation with a computer primarily to aid self-exploration or whether they are looking for wise advice. Facilitating self-exploration is one of the functions of a spiritual companion, and it might also be generally acceptable to turn to an automaton for that. It would be more challenging to implement a companion that could provide good-quality spiritual advice and guidance, and there might also be more skepticism from users about whether they could trust the advice they receive from a computer. Those who were primarily seeking help with self-exploration might be relatively happy with an automated companion. Others would be more focused on receiving authoritative help and guidance, and such people would probably be less content with an automated companion.

If the automated companion was embodied as an avatar on a computer interface, it would be very much a private companion of a particular individual, and it is unlikely that anyone else would ever interact with it. On the other hand, a companion implemented as a humanlike robot such as BlessU2 has a more

public presence. It might be in a church building and interact with a variety of people. That would give the robot companion a public presence in the spiritual community. Even if there were no differences in the technical capacities for spiritual conversation between these companions, it is clear that people react very differently to robots that have the external appearance of humans. It is a matter for future empirical work whether that would affect the acceptability of an automated spiritual conversation partner.

It is rapidly becoming possible to develop computers that can have spiritual conversation and other spiritually relevant interactions with humans. Preliminary evidence suggests that having a computer as a partner in spiritual conversation is acceptable to a significant number of people. Various methodologies are available to achieve this, and each has different strengths. Where there are training programs in spiritual practices or dispositions that have become routinized, it would not be difficult to write scripts that would enable a computer to deliver the training interactively. It would also be possible, in principle, to develop a personalized spiritual companion, building on previous work on personal companions for the elderly. These are likely to be quite widely used for spiritual conversation. They will not replace people but will be used alongside them. Gradually, wisdom will develop about when it is helpful to use a computer and when a person is needed.

In summary, we want to steer a path between, on the one hand, being completely dismissive of the potential value of spiritual conversation with a computer and, on the other, espousing an uncritical and exaggerated enthusiasm for what computers can contribute to the spiritual journey. An automated conversation partner has the advantages of ready availability and confidentiality. It can be a good source of information and facilitate self-exploration. However, people are likely to turn to a person for wisdom, for a sense of support and being understood, and for embodied connection. Humans are potentially better able to probe what a user is saying and offer wise personalized advice.

Acknowledgments

We gratefully acknowledge the support of the Templeton World Charity Foundation (TWCF0542) for the work reported here.

References

- Balle, Simon. 2023. "Theological Dimensions of Humanlike Robots: A Roadmap for Theological Inquiry." *Theology and Science* 21 (1): 132–56. DOI: [10.1080/14746700.2022.2155916](https://doi.org/10.1080/14746700.2022.2155916).
- Barnard, Philip. 2009. "Depression and Attention to Two Kinds of Meaning: A Cognitive Perspective." *Psychoanalytic Psychotherapy* 23 (3): 248–62.
- Bassett, Caroline. 2019. "The Computational Therapeutic: Exploring Weizenbaum's ELIZA as a History of the Present." *AI and Society* 34:803–12. <https://doi.org/10.1007/s00146-018-0825-9>.
- Cheong, Pauline H. 2020. "Religion, Robots and Rectitude: Communicative Affordances for Spiritual Knowledge and Community." *Applied Artificial Intelligence* 34 (5): 412–31. DOI: [10.1080/08839514.2020.1723869](https://doi.org/10.1080/08839514.2020.1723869).
- . 2021. "Bounded Religious Automation at Work: Communicating Human Authority in Artificial Intelligence Networks." *Journal of Communication Inquiry* 45 (1): 5–23. <https://doi.org/10.1177/0196859920977133>.
- Clocks, William. In press. "Steps Toward Android Intelligence." In *Cambridge Companion to Religion and AI*, edited by Beth Singler and Fraser Watts, 15–32. Cambridge: Cambridge University Press.
- Dein, Simon, and Fraser Watts. 2023. "Religious Worship Online: A Qualitative Study of Two Sunday Virtual Services." *Archive for the Psychology of Religion* 45 (2): 191–209. <https://doi.org/10.1177/00846724221145348>.
- Dorobantu, Marius, and Fraser Watts. 2023. "Spiritual Intelligence: Processing Different Information or Processing Information Differently?" *Zygon: Journal of Religion and Science* 58 (3): 721–48. [10.1111/zygo.12884](https://doi.org/10.1111/zygo.12884).
- . In press. "Editorial Introduction." In *Perspectives on Spiritual Intelligence*, edited by Marius Dorobantu and Fraser Watts, 1–28. Abingdon-on-Thames, UK: Routledge.
- Emmons, Robert A. 1999. *The Psychology of Ultimate Concerns: Motivation and Spirituality in Personality*. New York: Guilford Press.
- Löffler, Diana, Jörn Hurtienne, and Ilona Nord. 2021. "Blessing Robot BlessU2: A Discursive Design Study to Understand the Implications of Social Robots in Religious Contexts." *International Journal of Social Robotics* 13:569–86. <https://doi.org/10.1007/s12369-019-00558-3>.
- Ly, Kien, Ann-Marie Ly, and Gerhard Anderson. 2017. "A Fully Automated Conversational Agent for Promoting Mental Well-being: A Pilot RCT Using Mixed Methods." *Science Direct* 10 (Dec.): 39–46. <https://www.sciencedirect.com/science/article/pii/S221478291730091X>.
- McDonald, Matt. 2024. "Catholic Answers Pulls Plug on AI Priest, Father Justin". *Catholic News Agency*, April 29, 2024. <https://www.catholicnewsagency.com/news/257526/catholic-answers-pulls-plug-on-ai-priest-father-justin>.
- Miner, Adam S., Arnold Milstein, Stephen Schueller, Roshini Hegde, Christina Mangurian, and Eleni Linos. 2016. "Smartphone-Based Conversational Agents and Responses to Questions About Mental Health, Interpersonal Violence, and Physical Health." *JAMA Internal Medicine* 176 (5): 619–25. DOI: <https://doi.org/10.1001/jamainternmed.2016.0400>.
- Oman, Doug. 2013. "Defining Religion and Spirituality." In *Handbook of Psychology of Religion and Spirituality* (Second Edition), edited by Ray F. Palouzian and Crystal L. Park, 23–47. New York: Guilford Press.
- Savage, Sara. In press. "Practical Spiritual Intelligence: Integrative Complexity." In *Perspectives on Spiritual Intelligence*, edited by Marius Dorobantu and Fraser Watts, 131–45. Abingdon-on-Thames, UK: Routledge.
- Thibodeaux, Mark E. 2015. *Reimagining the Ignatian Examen*. Chicago: Loyola Press.

- Trovato, Gabriele, Franco Pariasca, Renzo Ramirez, Javier Cerna, Vadim Reutskiy, Laureano Rodriguez, and Francisco Cuellar. 2019. "Communicating with SanTO—the First Catholic Robot." 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). New Delhi: Institute of Electrical and Electronics Engineers. DOI: <https://doi.org/10.1109/RO-MAN46459.2019.8956250>.
- Watts, Fraser. 2002. *Theology and Psychology*. Basingstoke, UK: Ashgate.
- . 2017. *Psychology, Religion and Spirituality: Concepts and Applications*. Cambridge: Cambridge University Press.
- . 2021. *A Plea Embodied Spirituality: The Role of the Body in Religion*. Norwich, UK: SCM Press.
- . 2024. "Parameters and Limitations of Current Conceptualizations." In *The Oxford Handbook of the Psychology of Religion and Spirituality*, edited by Lisa J. Miller, 49–68. New York: Oxford University Press.
- Weizenbaum, Joseph. 1974. *Computer Power and Human Reason*. New York: Penguin Books.
- Wilks, Yorick. 2010. "Introducing Artificial Companions." In *Close Engagements with Artificial Companions: Key Social, Psychological, Ethical and Design Issues*, edited by Yorick Wilks, 11–20. Amsterdam: John Benjamins.
- . 2024. "Artificial Companions and Spiritual Enhancement." In *Cambridge Companion to Religion and AI*, edited by Beth Singler and Fraser Watts, 272–92. Cambridge: Cambridge University Press.
- Wilks, Yorick, and Afzal Ballim. 1991. *Artificial Believers*. Norwood, NJ: Erlbaum.
- Williams, Hattie. 2024. "CT Survey Reveals Split over AI". *Church Times*, May 3, 2024. <https://www.churchtimes.co.uk/articles/2024/3-may/news/uk/ai-ministry-better-than-nothing-church-times-survey-says>.
- Wiseman, Harris. 2020. "Wisdom and Moral Formation." In *Mutual Enrichment between Psychology and Theology*, edited by Russell Re Manning. London: Routledge.
- Worthington, Everett. 2003 *Forgiving and Reconciling: Bridges to Wholeness and Hope*. Downers Grove, IL: Intervarsity Press.
- Young, William. 2019. "Reverend Robot: Automation and Clergy." *Zygon: Journal of Religion and Science* 54 (2): 479–500. <https://doi.org/10.1111/zygo.12515>.
- . 2022. "Virtual Pastor: Virtualization, AI, and Pastoral Care." *Theology and Science* 20 (1): 6–22. DOI: <https://doi.org/10.1080/14746700.2021.2012915>.

