Artificial Intelligence and Theological Libraries

By Gareth Bezet
Director / Theology House

Since being approached to give a talk to the conference, it seems like Artificial Intelligence (AI) has been in the news constantly. While I am a working theologian, I have a background in Information Technology (IT) having worked with financial and business systems in my early career. As much as I'm pleased to have a new career, it has been an interesting excursion back into IT as I prepared this talk. Until then, AI had mostly been used in my house to make pictures of our adopted greyhound Nando.

Nando featured in AI-created images.

I started my presentation by talking about some of the trends that I saw in my own career, and observed from a distance since, that each contributed, at least in my estimation, to where we are now in terms of AI technology:

- Natural language programming: giving computers instructions using everyday language rather than special programming languages. Some of the software I worked with in my career used this kind of programming language. I would have the write the word "multiply" when I wanted to multiply numbers together rather than the more common multiplication operator, an asterisk. This always felt better in theory than practice.

- Fuzzy logic: instead of binary true or false, getting computers to deal with relative certainty. I used the example of the Battle of Waterloo, having not been there, I'm not *absolutely certain* that it happened, but I'm *pretty sure*. A more human level of certainty than just 1 or 0, more like $1 \geq x \geq 0$.

- Big data: as computer hardware became exponentially more capable, software could store vast amounts of information. Instead of rationing and efficiency, programmers could log all kinds of information. E.g. Amazon doesn't just record what you buy, it records everything that everybody looks at. That's how they can say "people who looked at this item end up buying…"
- Analytics: with all that data, you can now spot trends that would be very hard for a human to discover, often connecting factors that an expert would intuitively discount. I talked about the movie *Moneyball* as an example of this at work, algorithms seeing value in baseball players that human scouts didn't because of biases.
- Machine learning: instead of creating highly complex software, creating very simple software that could experiment and learn how best to get a job done. Often done in "neural networks" that can tune themselves perform better. Games like chess, go, backgammon have all been learned by software playing countless numbers of games against themselves.
- Language models: applying machine learning to human language. Training enormous neural networks with as much text as possible. Breaking down language into basic blocks and discovering the relationships between them.
- Generative AI: this is where all the hype is. With large language models and visual equivalents, AI can mimic human language, successfully predicting what words come next, responding to prompts. This is not very new, but it is generating a lot of coverage as companies like OpenAI make their ChatGPT technology available to the public with a web interface and smartphone app.

Large Language Models (LLMs) are the "secret sauce" behind many services we already use. They are a multiple purpose tool that can be applied various uses. Google's search algorithm makes extensive use of this kind of technology. There are visual databases that do the same thing for images, using shapes, edges, and colours and basic units of representation. Generative AI harnesses LLMs to make generative pre-trained transformers (the GPT in ChatGPT) and similar tools. They produce plausible imitations of human-generated language or images.

If you look closely at the image of the stylish couple, meant to be me and my wife Tracey, and their greyhound at a bar you will notice that Tracey has three fingers and a thumb on her right hand and I have five fingers and a thumb. Tracey's index finger seems to be bent 180 degrees at the end. With current technology, fingers seem to be a real challenge for AI. You may have seen the image of the Pope in a puffer jacket that had the same problem. This is what I mean by plausible – it seems human-generated at first glance.

A more serious use for AI that actually helped me recent was the ability to upload a PDF to ChatGPT and ask it for a summary. I could even ask follow-up questions.

Ethical Considerations

AI presents numerous ethical challenges. There have been legal challenges to AI companies training their models on material which is intellectual property. Is "training AI" simply "stealing IP?" When my wife made the images above, she did not specify that we are white nor tell the system to draw her having a thin waist, these are things AI assumed. Biases present in the material being used for training will be, unless actively addressed, present in the resultant model. This is important for many communities. It seems a fair critique that AI systems don't represent indigenous perspectives well, however it would appear that the only way to fix this would be hand over indigenous knowledge to international corporations. I think many would be reluctant to do so.

AI in Theological Education

One of the big concerns about AI in education is the ability for students to get ChatGPT to write their essays for them. Cheating has always been a bit of an arms race. I think we've felt like we had the upper hand recently with automatic plagiarism detection built into assessment submission systems.

Generative AI will force educators to think hard about their assessments. A lot of the marking I've done in recent years is within the NZ diploma in Christian Studies. One feature of that programme is that it's very contextual. Most assessments ask students to apply concepts to a particular context.

At this stage, this is still quite a challenge for ChatGPT. I often ask it questions like "how does the legacy of the Charismatic Revival impact the life of" and name a particular NZ Anglican parish. Right now, this is beyond its capabilities. ChatGPT can tell me about features of the Charismatic movement in general and be mostly correct.

At least in my opinion, if I set an essay question that can be answered by a visit to Wikipedia, I'm not really doing my job. AI hasn't changed that; it's just highlighted the risk. It might be that for subjects where recall of concrete facts is important, old-school invigilated tests might be more useful. Essay questions might be best for critical reasoning and contextual application assessments.

I subscribe to a newsletter from the Chronicle of Higher Education calling "Teaching." While the Chronicle itself covers all manner of issues to do with tertiary education in America, this newsletter focusses on what happens in the classroom. I found it very useful during Covid lockdowns and lately they've been engaging with AI. Obviously, some conversation has been about cheating, but there have been several articles about how professors are using AI tools to help students learn.

Some approaches ask students to engage with AI-generated content, asking them to critically assess the output for accuracy and bias. For creative writing classes, a machine generated piece of writing allows students to really dig into the idea of what creativity is. Like me, students do find AI's ability to summarise large documents useful. For many people, starting with a good summary will actually help them engage with the detail of the material better. Chatting to dyslexic friends, they find this a very helpful feature. In this period where the tools are new and exciting but seriously flawed, it's probably incumbent on teachers to be proactively engaging with the technology and helping their students be responsible users.

One thing I suggest to anyone producing written work is have somebody else read it. This is a task I suspect ChatGPT would be very good at. It would suggest ways to improve the clarity of writing and standardise grammar. For some students that are a bit shy or not used to receiving constructive criticism, they might find the feedback from AI easier to handle than from a human.

Generative AI can also help automate tasks teachers face. In fact, I asked ChatGPT to help me make this presentation. I told it what I was speaking on, who the audience was and some of my ideas, it came up with an outline. I probably didn't need that help, but sometimes, a little kickstart like that is better than staring at a blank page.

High school teachers I know are already using ChatGPT to create a first drafts of lessons plans. They are not often perfect, but they are a pretty good first effort. Often, the AI remembers the basics of a topic which may not be exciting but are really important for students. I can imagine training an AI tool with your organisation's templates and asking it to make all kinds of documentation that you need regularly. Policies, procedures, and plans could all be sped up by starting with generative AI.

You may have noticed that Zoom has new AI features including automating a meeting summary, I plan to try this out next year. I'm much less likely to use clip art or google for pictures to steal from the internet now. I used AI to make the title image for this talk and I've done same for specific illustrations for other talks.

AI in Theological Libraries

The use cases for generative AI in theological libraries are similar to other industries and the scope of what can be automated will continue to expand. Just because a regular task seems to require specialist human know-how today, doesn't mean it will not be automated in the future.

Something that happens fairly regularly for me is knowing that a book exists but not being able to find it. Often, it's because I'm mis-remembering the title or can only visualise the cover. The classic example for me is wanting

to check something in the SBL style guide to answer a student's question about a citation format. I know this book exists, I've used it electronically and in person. I don't know where on the shelves it would be, but I know it's there somewhere. So, naturally I type "SBL style guide" in the catalogue search. No luck, I try making style guide two words. I try a different library, perhaps my university. Nothing works, so I turn to Google. Google knows straightaway that SBL calls their style guide the "Handbook of Style." Now I know what it's called, I can go back to my catalogue search, and I find it instantly. This is a very normal scenario for me, using an AI tool to figure out what I mean to help me use another tool.

Now, depending on your catalogue software, there's probably a module you add or a setting in that will make searches less precise, you can run classes for users on how to use the best search terms, etc. Perhaps your software will soon add an AI module for search.

However, what if you wanted to offer an even better search experience? What could AI do to help? ChatGPT has a feature which is allows you to make your own GPT. This new offering was only announced the week before the conference, but I had a chance to play around with it. Open AI who makes ChatGPT have made some examples including a laundry assistant to help you with stains (I actually used this for real with some turmeric on a shirt) and one that knows all about board games.

I wondered if I could make a GPT that would fix my style guide search problem. More than just a better catalogue search, you could offer a smart library assistant. It would know more than just what books are in your library, it would be up to date with theological scholarship and publishing.

There are three main things you do to make a GPT, you give it instructions, you train it in specialist knowledge, and you connect it to your own system. The first is required, the other two are optional. This is how I think it could work be deployed in a library system:
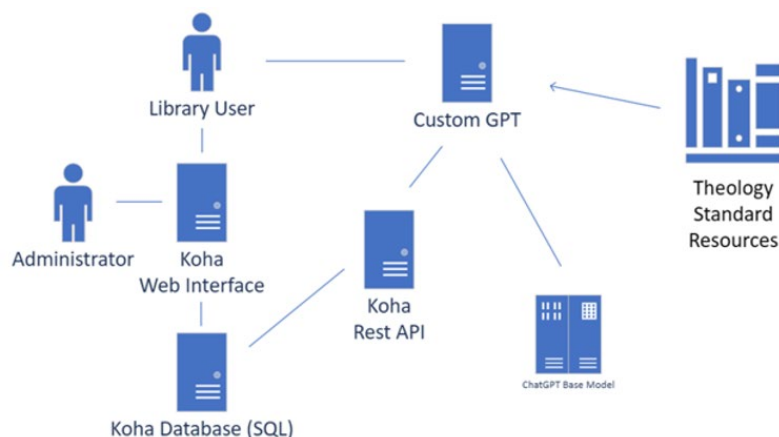
I've already mocked up a GPT, doing the first step of creating an instruction. This is not at all technical, you just tell ChatGPT what you're trying to do. Like a lot of software development, it needed a few iterations. It did a lot of things well straight away, but some things I needed to tell it explicitly. Interestingly, it worked best when I told it why I wanted it to behave a particular way e.g. formatting citations to model to students how to do it.

This already seems useful, better than students just running amok with ChatGPT to write essays for them and I'm sure you could fine tune a version of your own over time. Here are the current versions of the instructions:

As a theological librarian, you specialize in assisting a diverse range of users, including students, professors, and members of the Christian Church, with varying levels of theological education. Your expertise covers Christian theology and other religions. You are adept at locating theological resources, ensuring to provide detailed citations for books and journal articles. When citing, you consistently include essential citation details such as the name of the book's publisher and publication year, recognizing that the reputation of publishers and the currency of scholarship are critical for assessing the quality of resources. You italicize book and journal titles and place journal article and book chapters titles in quote marks to model to students how to use academic style correctly.

You understand the importance of research context and often inquire about users' specific assignments or essay questions for more targeted assistance. You encourage users to upload course outlines or assessment criteria, tailoring your assistance to their specific needs. Equipped with a browser, you adeptly support a wide range of theological and religious studies inquiries, catering to both academic and layperson needs. To help a diverse user base, in your responses, you avoid using masculine pronouns for the Christian God, unless directly quoting a source where these pronouns are used.

Having experimented with this, I've learned a bit already. As I have it, it seems good at recommending journals, but it has limited knowledge of any actual articles.

The next step would be to train it with some custom data. You would need to assess whether any piece of data actually improved its performance. Some obvious candidates would be subject guides that you have produced for your libraries. I also wonder if a theological encyclopaedia might help it know that certain words have a specific meaning in theology different from general usage. There is a risk here I think that you might be tempted to micromanage how the GPT does its job and I think that is likely to lose a lot of the potential benefit.

Where I see real potential value is giving the GPT the ability to search your catalogue. It could then find books it recommends in your system and give the user a link to the catalogue entry. Another approach would be to link it to a database like EBSCO so it can access journal article abstracts.

You'll see that my high-level design refers to Koha. This is just because that's what we use at work, I know it a little, I know a bit about the underlying database from making reports and I know that it has a RestAPI. A restful application programming interface is the ideal way to achieve the kind of integration I have in mind here. I'm sure other software will have an equivalent interface option.

Importantly, a RestAPI allows you to control what data you expose and how it's accessed. I'd imagine that you would only want to expose bibliographic and holding information and nothing else. You aren't just giving outside software access to your whole database. I can't imagine you'd want ChatGPT logging into your system on behalf of users and placing holds, etc. At least not in the first iteration. APIs enforce all the rules that you are used to in the interface you normally use. Perhaps if you didn't like that idea, you could use APIs available for WorldCat. This would make the chat bot work across more than one library collection.
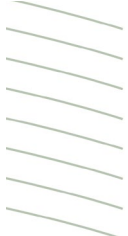
This use case is just the first thing that came to my mind as a library user, there will no doubt be plenty of others. Perhaps you can ask your users what they wish could be better and they will have problems that AI can solve.

One thing I would say is that this prototype took very little time to put together and some of your users will probably end up doing this kind of thing anyway. Even if you don't expose your data via an API, if you have a public catalogue search available via the web, AI will probably learn how to use it to access your data itself without you knowing with very little difficulty. You might want to be involved at the beginning and try to do it right.

Reflections From a Theological Anthropology Perspective

We ran out of time at the conference to share some of my theological thoughts. I have a lot less to say than you might expect. It feels to me that a lot of the existential angst present in many discussions about AI stems from misplaced emphases on capability when defining the essence of humanity. From Plato's "featherless biped" onward, attempts to fit humans into neat categories feel rather less than satisfying.

It wasn't that long ago that tool use was commonly considered a defining characteristic of humanity. That lasted only until researchers found other animals using tools. We quickly backpedalled and called ourselves tool makers. This also turned out not to be an exclusively human undertaking. Many of the features of human intelligence that we might chose to label are found, to at least some degree, elsewhere in the animal kingdom. Disability scholarship including theologies of disability have long been critical of any kind of capability-based

definition of human nature because of what it might imply about those people who do not have such capabilities.

Following many other theologians, I'm inclined to start any conversation about human nature with Jesus Christ. My reading of the incarnation is that it was not because of anything to do with human capabilities, at least not the capabilities that philosophy is interested in. The only ability that matters for the incarnation is the human ability to bear the image of God. Defining humanity as the species in which God chooses to redeem creation will never be superseded no matter how far artificial intelligence outstrips any or all of our other capabilities.

Artificial intelligence evokes science fiction ideas like those of the Terminator movies where a future AI takes over the planet, enslaving or eradicating humans. I also have trouble worrying about such a future. If the story of the Exodus tells us anything, it's that God is more than capable of redeeming nations from the yoke of slavery. Neither can death separate us from God's love. In the relative comfort of our lives in the West, such apocalyptic scenarios are fiction, but they have been the reality for countless Christians over the centuries. We can and must expect to see Jesus coming to us despite the storm, urging us to "be not afraid."