From chatbots to connections: How AI will change services

Projects by IF / June 2023

Projects by IF is a strategic design agency, we specialise in what it takes to make trustworthy products and services easier to design, launch and maintain. We have worked with some of the biggest tech brands, government organisations and healthcare providers in this space.

Our passion for this work comes from our curiosity and focus on building a future where technology embeds care.

We also believe that for any product, service or organisation to be successful, trustworthiness needs to be as much of a consideration as usability.

We would also love to hear from you. What have we missed? What do you think of these ideas? Where do you see Learned services being most disruptive?

Get in touch at hello@projectsbyif.com

Contents

What makes Learned services different

Postscript

Glossary

Learned services are a paradigm shift

The future isn’t chatbots. The future isn’t products or services built on individual Large Language Models (LLMs). The future is far more exciting, and interesting.

We agree on one popular idea... that the future will include lots of personal AIs. More on that another time. This work is about another paradigm shift. A paradigm shift we call Learned services.

Learned services help people create their own experiences with features and steps from multiple organisations, and potentially with multiple people.

They create new digital capabilities

Learned services come about not only by adding Generative AI into existing products. It's about applying the capabilities of LLMs to digital infrastructure. LLMs will become a kind of connective tissue, like fascia in the human body, for this infrastructure.

This has big implications for the kinds of experiences we expect to see around the corner.

> Why the name Learned services?
>
> We call them Learned services, because we often refer to AI as being a Learned system. That comes from a project we completed for Google in 2018. Our client, Matt Jones, introduced us to the term, as a way to refocus on people and society. Learned systems are not magic. They are made up of data, algorithm(s), computational power, people's work and money. People build and teach these systems, and they operate in the real world.

What makes Learned services different

Services are already becoming like lego bricks

Services today are tailor-made "end-to-end" to solve a user problem. However, they are changing. Services are increasingly being componentised into building blocks that can be used by multiple services to solve a need. We can already see this in experiences like WeChat, or in Online banking via Open Banking. What once were stand-alone services become features in someone else's product.

With the capabilities of AI we expect this componentisation and reuse to become more frequent. But people won't only be able to use components from within the 'boundary' of one organisation.

AI couples these components into services

AIs can perform a range of complicated tasks like calling an API, ringing a call centre, helping people negotiate with an organisation, or helping a customer service agent answer a question.

Through this capability AI can combine components from multiple organisations into Learned services, even if those components were not originally designed to be building blocks.

People can create their own experiences

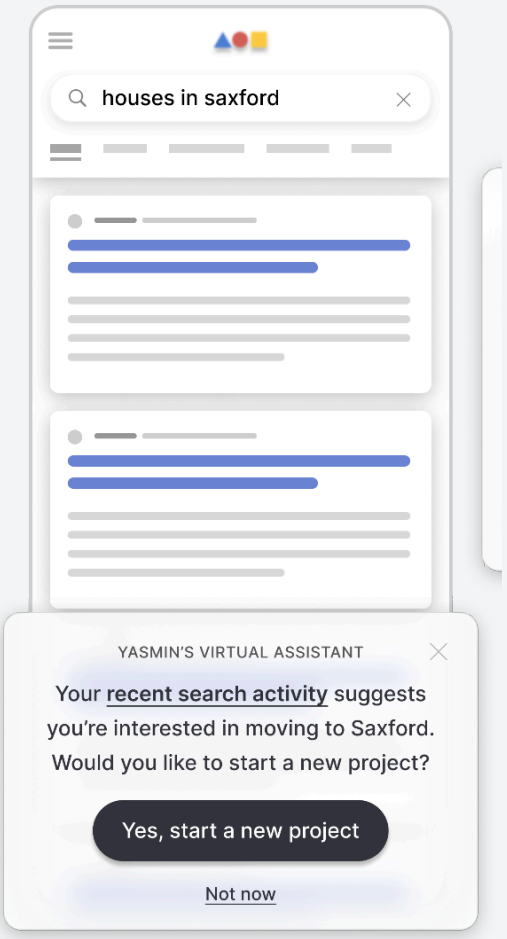

We expect multiple organisations will provide components and steps that people can automatically assemble, or self assemble, within their interface of choice. Essentially, combining components to create an ultra-flexible experience that can be open-ended.

There is the potential to serve multiple people at the same time too.

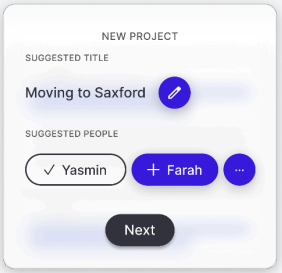

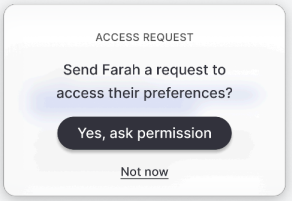

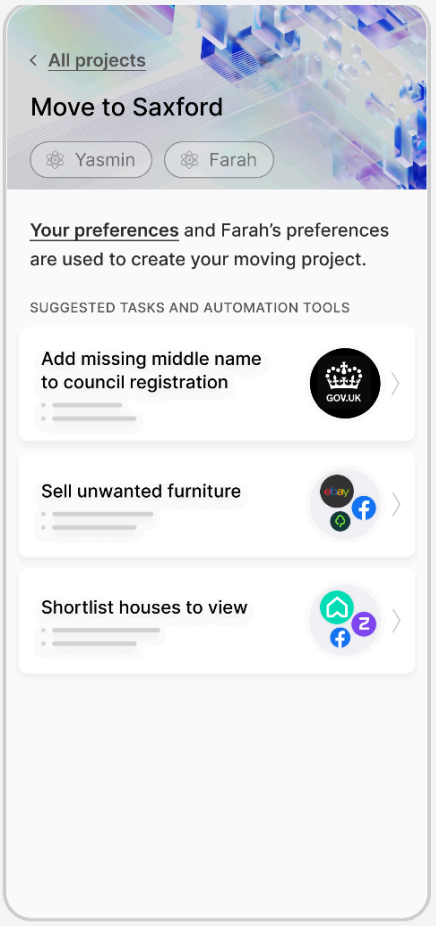

This doesn't mean perfectly modelling relationships in software, it means enabling people to use services in ways that support their existing relationships.

New user needs and design patterns become status quo

We know the introduction of new experiences changes the expectations that people have. What originally is thought of as a nice to have, or the exception, over time becomes an expectation. For example, when cars were first introduced seatbelts were a nice-to-have. Now, few people would choose to travel in a vehicle without a seatbelt.

Similarly, Learned services will create new needs such as forgetting: “as someone who has got a new job, forget my previous work preferences”.

From our client work, we know that existing design patterns are already insufficient to help people understand how AI works, or whether to trust it. We know that if someone cannot trust something, they are less likely to adopt it. New patterns will be needed.

Products solve problems.

Learned services orchestrate the architecture for the problem to solve itself.

No two experiences will be the same

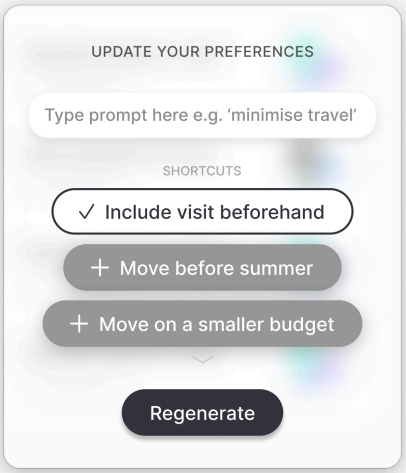

In providing Learned services, organisations will have less control over what a user does or sees. That's because AI removes the fixed set of circumstances, and replaces it with effectively an infinite set of circumstances. By design the service can provide a different experience to each user.

This could mean that Learned services make users feel less in control. For example, if assumptions the Learned service is making about them are inaccurate.

We think for Learned services to deliver good outcomes, they need to be worthy of user trust.

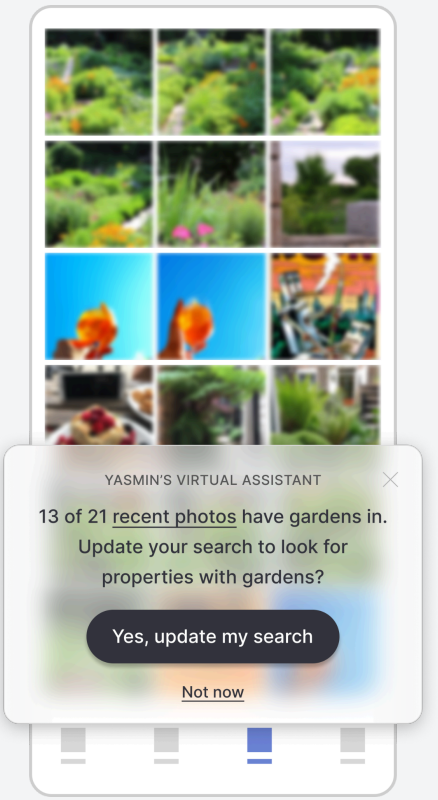

Learned services also create different demands of data

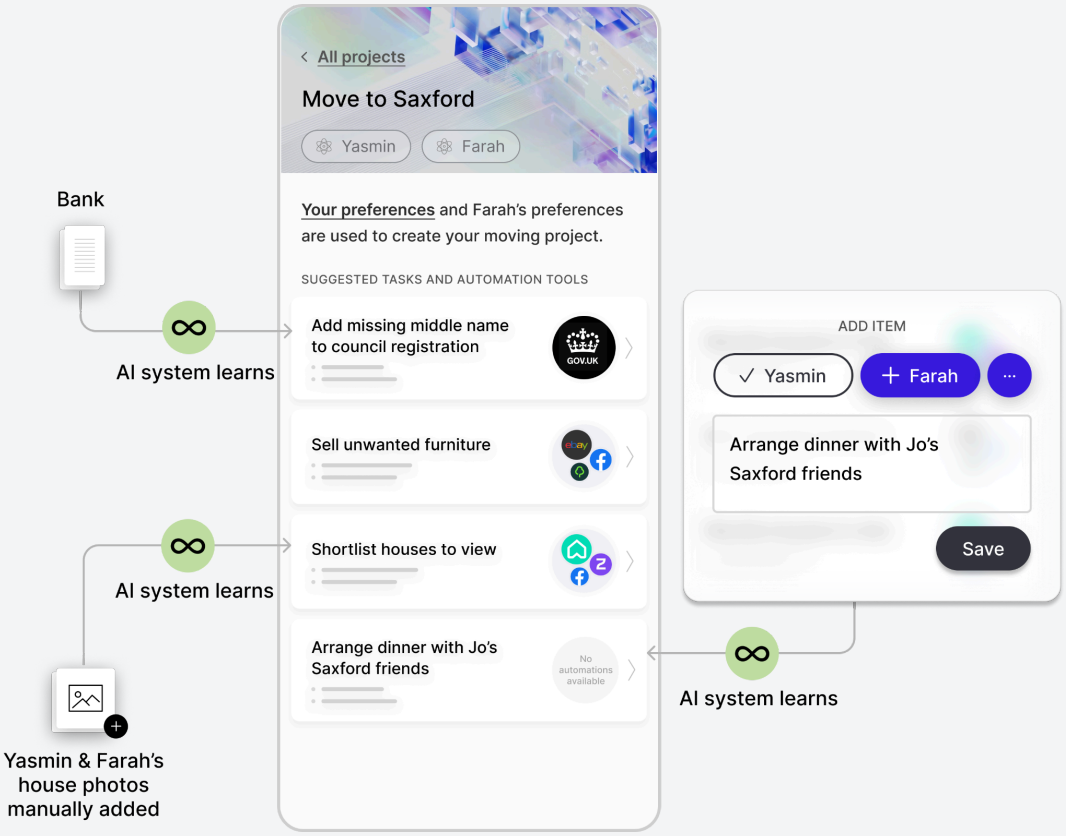

A Learned service needs different types of data to deliver its outcomes. Some of this data will come from the organisation that delivers the service, while other data will be publicly available information about our world, societies and the places we live in. But the rest of the data will come from people and the other organisations in the architecture that the Learned service is orchestrating.

In particular people will need to provide Learned services with enough information about themselves, and their preferences, that it can meet their needs. This will require a high element of trust.

To be worthy of trust, AI needs wrappers

AI needs to know enough about the digital infrastructure so that it can reliably use, combine and interact with it. This requires trusted relationships between people, organisations, AI, and it will also require new information to understand how a component can be used.

This information might be API documentation, preferably documentation for human-centred APIs, but should also be in human-readable policies. Those policies can then be read by humans, either on their own or supported by AI. This will form part of a 'wrapper' around the service.

Some organisations will reject becoming a feature in other people's services, see Reddit for a current example, but there are opportunities for organisations that can embrace and thrive in this future. Organisations can make more of the capabilities they have by reaching more customers, meeting additional needs, and growing into new markets.

We would also love to hear from you. What have we missed? What do you think of these ideas? Where do you see Learned services being most disruptive?

Get in touch at hello@projectsbyif.com

Postscript

There are public policy implications for Learned services

Much of the current debate about AI and LLMs see them as products that can be regulated like other kinds of products. We see a more systemic, pervasive and non-deterministic impact and that opens up new questions. For example: How do we

protect people when they can co-create their own service? Where does liability lie when things go wrong? Who will build Learned services and what are the potential competition issues? What will it mean for jobs within service providers? Will it help or hinder the design and delivery of public services that work for everyone? What are the short and long-term risks?

We know that lots of this future isn't strictly “new”

In the glossary we've included links to further reading on current services, like WeChat and step-by-step services, and experiments that show other glimpses of this future.

Some of the capabilities that we believe LLMs will bring to services are already here, but in a less automated way. LLMs bring exponential change to the direction that service design is already headed in. We believe understanding technology’s role and how to make AI understandable and trustworthy is critical work to do.

These ideas build on other ideas

These fictional illustrations were designed and iterated by the team at IF. They are intentionally left untested. To validate our ideas we sought feedback from different experts in the fields of service design, public sector services, and LLMs. Thank you to Josh Laird, Martin Jordan, Matt Webb and Payal Wadwha for their feedback.

Glossary

Service

A service is an organisational* response to users’ needs.

* Organisations include communities

So what

We have deliberately chosen a different definition (inspired by Payal Wadhwa) rather than other existing definitions of “service” that are in use in the service design community (for example, in work by Lou Downe and Kate Tarling). That’s because we are interested in the role that organisations play in providing services, and the sense of responsibility that comes with naming that.

Further reading

Lou Downe - what is a service?
Kate Tarling - Defining services
Sarah Drummond - full stack service design

AI (or Learned system)

The science and development of computers that are able to perform complicated tasks.

AIs, or Learned systems, are made up of 5 main components:

1. Data - Information on computers, e.g. the data used to train a model, the weighting for an algorithm, or data used to filter inputs and outputs.
2. Algorithm(s) - Rules to follow in order to solve a problem eg. from deep learning to reinforced learning.
3. Computing power - Computing resources required for an AI to perform a task.
4. Humans - To train, correct and maintain the AI.
5. Money - To pay for the other components.

So what

Despite the limitations of the acronym AI, we are using it because it is an increasingly familiar term.

Identifying the key components of AI helps us understand what organisations need access to to be able to research, develop, and maintain AI technologies.

Further reading

Alan Turing Institute - Data science and AI glossary
European Commission - EU AI Act
Georgetown Law Center on Privacy & Technology - Artifice and intelligence
UK Government - A pro-innovation approach to AI regulation
USA White House - Blueprint for an AI Bill of Rights

Foundation model (or FM)

A kind of AI that can be adapted to a range of tasks. The first kind of foundation models were Large Language Models.

So what

There is no widely accepted definition of foundation model. Some definitions state that they are trained on a large amount of data, that foundation models use deep-learning algorithms, or that the tasks they can perform are specifically 'downstream' tasks.

We have simplified these definitions to remove information about how foundation models are developed because we are more interested in what they can perform - i.e. a wide range of tasks - and because the techniques used to develop foundation models are likely to change over time.

Further reading

Alan Turing Institute - Exploring foundation models
Jacob Steinhardt - What will GPT-2030 look like
Harvard - Introducing the Center for Research on Foundation Models (CRFM)
Qinghua Lu et al - Towards responsible AI in the era of ChatGPT: A reference architecture for designing foundation model based AI systems
Unknown Google Researcher - We have no moat And neither does OpenAI

Large language model (or LLM)

Large language models were the first foundation model to be created. They are an AI that is trained on language data, and that usually outputs language.

So what

Large language models were the first type of foundation model to be created. They are trained on text and typically produce text outputs.

As these models were the first to be created, they may be the first to be both widely used and regulated. As a result the EU and US have jointly developed a definition that could be used in regulation. This definition is "*A class of language models that use deep-learning algorithms and are trained on extremely large textual datasets that can be multiple terabytes in size*."

Our shorter description focuses on the language nature of the model and does not care about the type of algorithm that was used to develop the model or the size of the

datasets that was used to train it. Algorithms and the size of training datasets may change over time.

Further reading

Bender, Gebru et al - On the dangers of stochastic parrots: can language models be too big?
European Commission - EU-U.S. terminology and taxonomy for Artificial Intelligence

Generative AI

A content creating AI, for example an AI that creates text, images, audio and/or video. Think of it as an instance of a Foundation Model.

So what

There is no widely accepted definition of generative AI. In our short definition we focus on the creation of a wide variety of new content and anchor it from the definition of a foundation model.

More detailed definitions will focus on how the generative AI model is developed or the statistical techniques the model uses to develop new content.

Further reading

Mckinsey - What is generative AI?
Ted Chiang - chatGPT is a blurry JPEG of the web
World Economic Forum - What is generative AI? An AI explains

Foundation model providers

An organisation that provides foundation models as a product to service providers.

So what

The product might be an API provided by a SaaS (software-as-a-service) firm, an open source repository, a custom installation developed by the foundation model provider's engineers and integrated into the service provider's stack, or other things.

Further reading

Hugging Face is a foundation model provider.

Learned service

Learned services help users create their own experiences with features and steps from multiple organisations.

So what

We are increasingly seeing services that go beyond the step by step services provided by a single organisation or a group of organisations from a single sector, such as a government. Instead these services combine multiple services from multiple organisations into an all-in-one service and a single journey.

Some early attempts at these services are the so called “superapps” like WeChat, Grab, and GoTo. Services provided by other organisations become features within these services.

We are pushing this concept further into a world where generative AI is widely used.

Rather than the generative AI capabilities being used to create new text or images, generative AI capabilities are used to learn and create a new service by combining service components, service propositions and service processes based on information provided by users.

This will allow the creation of new service experiences.

Further reading

Apple - Shortcuts User Guide
Customer Futures - Finding the winner in Personal AI
Gartner - What is a Superapp
GDS - Step by step navigation
McKinsey - The next frontier of customer engagement: AI-enabled customer service
Matt Webb - The surprising ease and effectiveness of AI in a loop
Salesforce - 3 ways generative AI will reshape Customer Service
Simon Willson - Thoughts on AI safety in this era of increasingly powerful open source LLMs
Stephen Wolfram - ChatGPT gets its “Wolfram superpowers”!
Thoughtworks - Leverage WeChat to quickly start your loyalty program in China