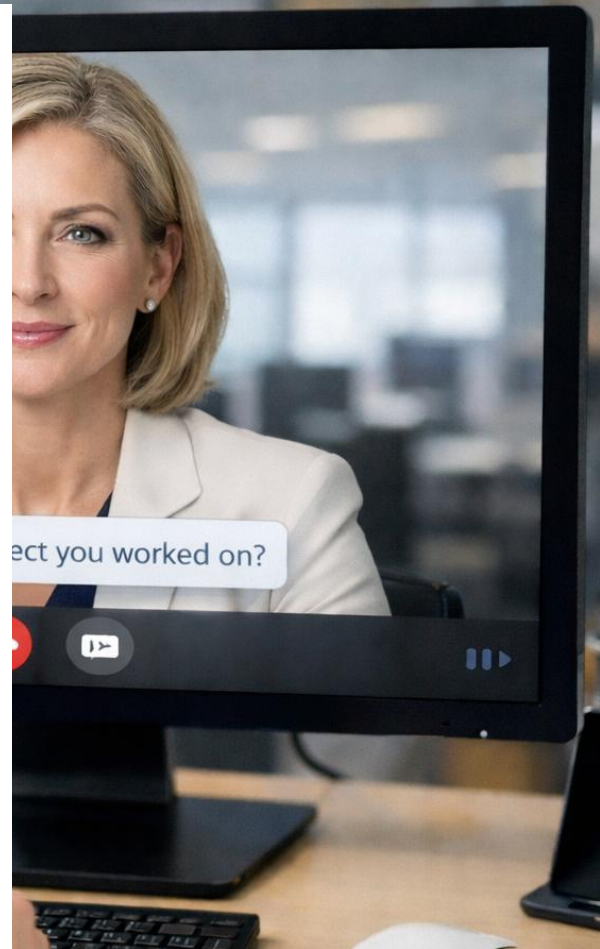


AI Agents for Continuous Organizational Learning



February 2026

DADEN LIMITED

Author: David Burden



INTRODUCTION

Organizational Learning is often defined as the processes by which an organization creates, retains, and transfers knowledge in order to improve over time¹, or how it promotes learning and systems thinking within the organisation². It is about the meta-processes of the organisation, and the processes by which an organization transforms individual knowledge into organizational knowledge³.

As AI agents (aka chatbots) become ever more common in the workplace, there is an obvious question as to how they can best contribute to such a process of continuous organisational learning. Terms such as mentors, coaches, assistants and subject experts are all often mentioned within both an organisational learning and an AI agent context. However, there are subtle differences between each of these, and they by no means represent the limits of potential AI use within an organisational learning context.

The aim of this white paper is to address this issue of what roles can AI agents play within organisational learning, and how these roles differ and interact. The examination is very much positioned with the context of knowledge management, where AI becomes a tool to enhance knowledge management and organisational learning, rather than necessarily being something which forces or demands a revolutionary change in how organisations work.

The paper will first provide some models of organisational learning and knowledge management, and then examine the roles of mentor, coach, expert and others in terms of their experience and exchange dynamics. The distinction between team-work and task-work will be introduced, as well as the different forms of Digital Human Twin (DHT) which are emerging. The paper will then consider how AI agents can potentially fulfil the roles discussed, and how these relate to the concept of digital human twins. The paper closes by examining AI as knowledge management and organisational learning custodians.

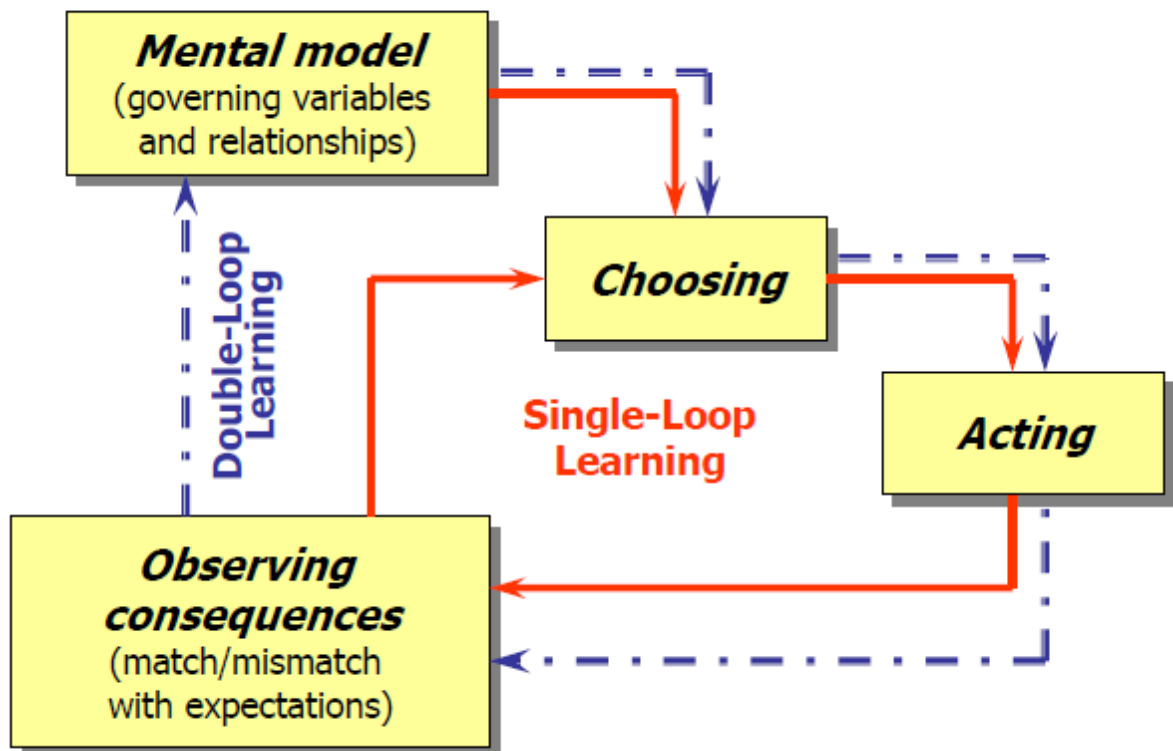
1 https://en.wikipedia.org/wiki/Organizational_Learning

2 <https://www.sciencedirect.com/topics/neuroscience/organizational-learning>

3 <https://journals.sagepub.com/doi/10.1177/2158244018794224>

Organisational Learning

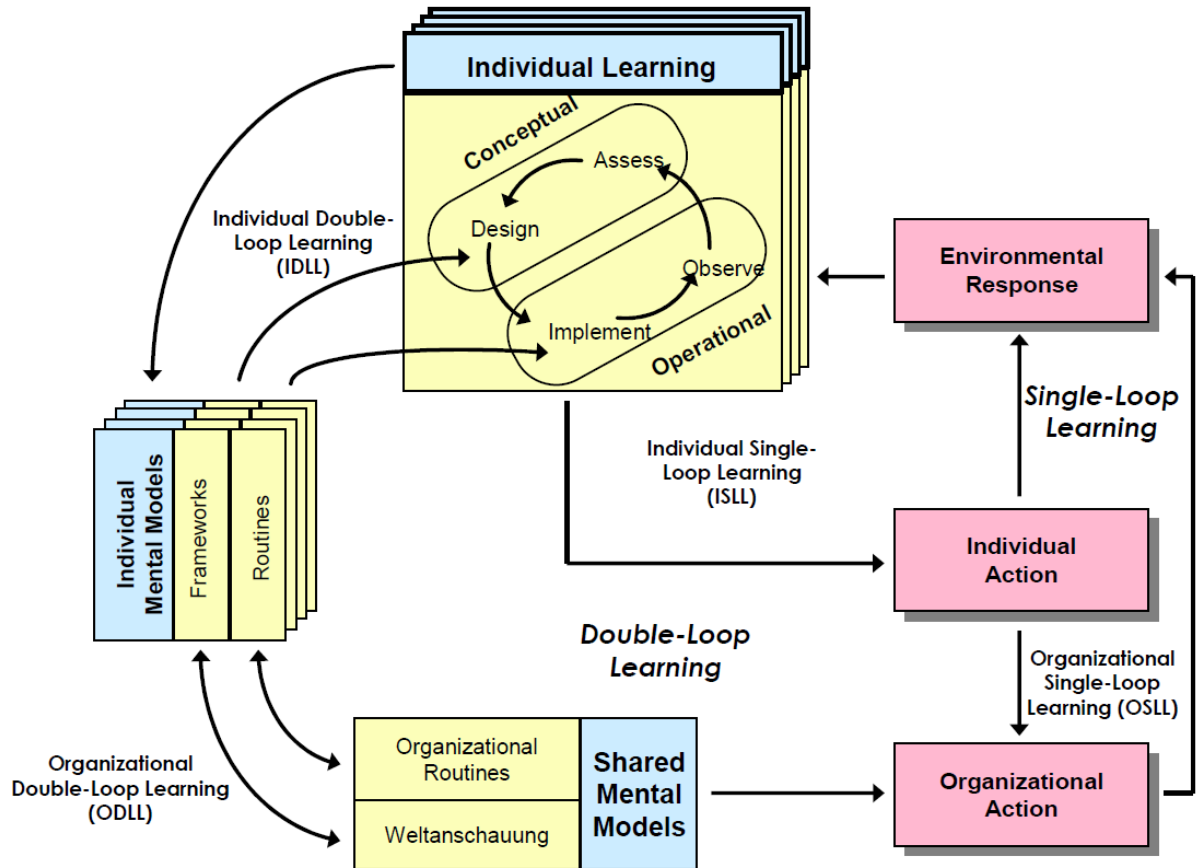
It is useful to have some sort of model of what organisational learning is, as this will help us to see whereabouts AI agents can fit in. Organisational learning is often talked about in terms of single-loop and double-loop learning. **Single loop learning** is essentially where an action is taken, it's result checked, and if different from the norm then a new action is taken to better fit the norm, and that new action learnt. **Double-loop learning** is where that norm, and the context for the decision, is challenged and norms adjusted to better meet corporate goals and the whole mental model updated.



(after von Kutzschenbach, 2006⁴)

⁴ Kutzschenbach, M. V. (2006). Organizational learning and social network structures: case studies of German and Norwegian forest sector organizations (Doctoral dissertation, Freiburg (Breisgau), Univ., Diss., 2006).

An integrated organisational learning model then put this into the context of **individual mental models** and the organisation's **shared mental model** (see sidebar) - the evolution of that shared mental model representing the sum of the organisational learning.



(after von Kutzschenbach, 2006⁵ and Kim, 2009⁶)

Weltanschauung in the diagram refers to an organisation's world view - its assumptions about the world and its role in it, which determines how it responds and acts.

It is the transfer of learning from the **individual mental models** to the organisation's **shared mental model** that is the

5 Kutzschenbach, M. V. (2006). Organizational learning and social network structures: case studies of German and Norwegian forest sector organizations (Doctoral dissertation, Freiburg (Breisgau), Univ., Diss., 2006).

6 Kim, D. H. (2009). The link between individual and organizational learning. In The strategic management of intellectual capital (pp. 41-62). Routledge.

Shared Mental Models



A shared mental model (SMM) effectively represents the degree of congruence between the individual mental models of each member of a team or organisation. Whilst the SMM itself can be hard to manipulate directly, through focussing on the meta-learning process, and making participants more aware of the concept and of concerns such as job/task, technology, team and team individuals as described in the main text the team's alignment to a shared mental model can be gradually developed, and the shared mental model itself evolved.

crux of organisational learning. For reference, shared (and individual) mental models can be thought of as comprising information about⁷: ❖

- **Technology and equipment**, including parameters, operating procedures, limitations and failure modes;
- **The job/task**, including processes, procedures, scenarios, contingencies, strategies and constraints;
- **The team**, including individuals' knowledge, skills, aptitudes, preferences and tendencies, and applying to customers and suppliers as well; and
- **Team interactions**, including roles and responsibilities, information sources, interaction patterns and hierarchies, communication channels and information flows, and role interdependencies.

Shared mental models have proven vital in creating high performing human teams⁸, and as a way of considering human-AI partnerships and teaming⁹.

Particularly pertinent to the current discussion is that whilst organisational memory might include all the documents and data that an organisation has, "*the mental models in individuals' heads are where a vast majority of an organization's knowledge (both know-how and know-why) lies*" (Kim, 2009).

⁷ Cannon-Bowers, J. A., Salas, E., & Converse, S. (1993). Shared mental models in expert team decision making. *Individual and group decision making: Current issues*, 221, 221-46.

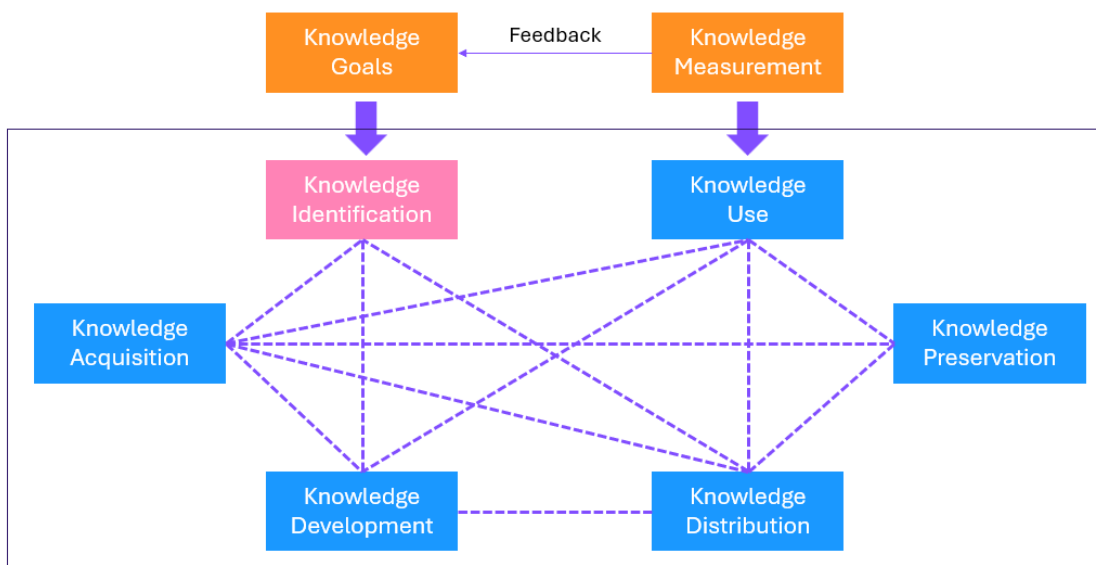
⁸ Richards P, Mascarenhas D, Collins D (2016) Developing team decision making: A holistic framework integrating both on-field and off-field pedagogical coaching processes. *Sports Coaching Review*.
<http://dx.doi.org/10.1080/21640629.2016.1200819>

⁹ Gu, E., Siu, H. C., Platt, M., Hurley, I., Peña, J., & Paleja, R. (2025). Enabling Rapid Shared Human-AI Mental Model Alignment via the After-Action Review. *arXiv preprint [arXiv:2503.19607](https://arxiv.org/abs/2503.19607)*.

Knowledge Management

Whilst the focus of this paper is not on knowledge management *per se*, a significant part of organisational learning is, as in the definitions at the beginning, about how organisations create, retain, transform and transfer knowledge.

In considering knowledge management the model below provides a useful reference¹⁰.



Probst (1997). https://www.researchgate.net/publication/5063346_Building_Blocks_of_Knowledge_Management_-_A_Practical_Approach.

It identifies the key elements of a knowledge management process as follows:

- **Knowledge Goals** - what knowledge the organisation needs, and who needs it and why;
- **Knowledge Identification** - identifying where the knowledge an organisation needs is - this can include knowledge external to an organisation, new knowledge, or knowledge that is in people's heads (implicit knowledge);
- **Knowledge Acquisition** - how you acquire the knowledge you've identified as being needed;

¹⁰ Probst, G. J., & Romhardt, K. (1997). Building blocks of knowledge management-a practical approach (No.97.12). https://www.researchgate.net/publication/5063346_Building_Blocks_of_Knowledge_Management_-_A_Practical_Approach

-
- **Knowledge Preservation** - how you keep the knowledge you've acquired so it's available to others in the future;
 - **Knowledge Development** - how you can build new knowledge from the knowledge you've got;
 - **Knowledge Distribution** - how you ensure that employees can get access to the knowledge they need at the time they need it - and now that it exists;
 - **Knowledge Use** - ensuring that the knowledge is actually used in the activities of the organisation and its members; and
 - **Knowledge Measurement** - measuring which, how and how much knowledge is used, and where gaps are, which can in turn affect, in particular, goals, identification and preservation (including deliberate deletion or archiving).

As a cautionary tale about emphasising knowledge acquisition and preservation without thinking about knowledge distribution and knowledge use, the US Navy once built a “lessons learned” database of over 35,000 “lessons”, but had no effective methods for delivering those lessons to potential users - a lesson (or knowledge) distribution gap¹¹.

Mentoring, Coaching and other Roles

The Chartered Institute of Personnel Development (CIPD) defines relevant organisational learning roles as follows¹²:

- **Coaching** targets high performance and improvement at work and usually focuses on specific skills and goals, although it may also have an impact on an individual's personal attributes such as social interaction or confidence. The process typically lasts for a relatively short defined period of time, or forms the basis of an ongoing management style.
- **Mentoring** involves the use of the same models and skills of questioning, listening, clarifying and reframing associated with coaching. Traditionally, however, mentoring in the workplace has tended to describe a relationship in which a more experienced colleague uses

¹¹ Weber, R. O., & Aha, D. W. (2003). Intelligent delivery of military lessons learned. *Decision support systems*, 34(3), 287-304.

¹² <https://www.cipd.org/uk/knowledge/factsheets/coaching-mentoring-factsheet/>

his or her greater knowledge and understanding of the work or workplace to support the development of a more junior or inexperienced member of staff. One key distinction is that mentoring relationships tend to be longer term than coaching arrangements.

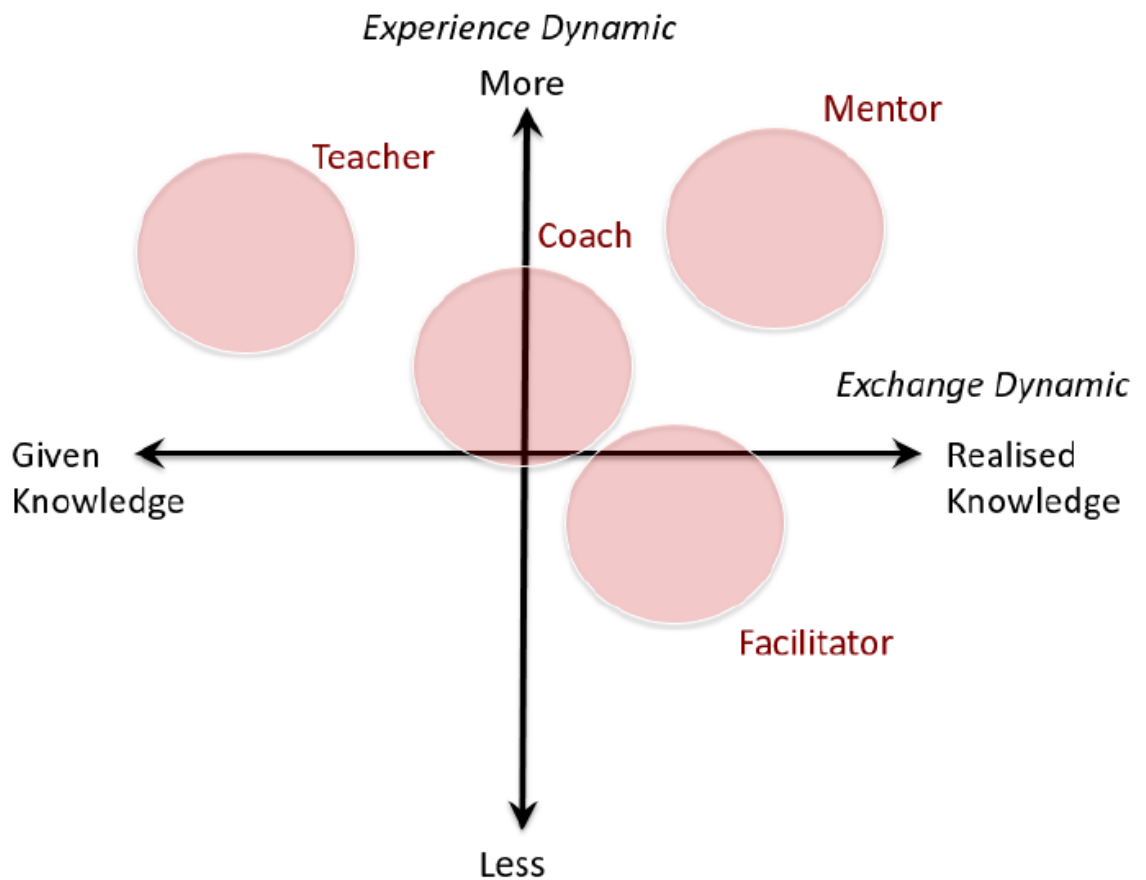
By contrast:

- A **facilitator** works with a group of people to help them have a conversation, come to an agreement, or plan for the future. In general, the facilitator acts as a neutral outside voice, making decisions about the process the group goes through but allowing the group to focus on and control the content of the discussion. The facilitator is a gentle guide, making it easier for the group to have that discussion¹³.
- A **teacher** is someone whose primary role is to impart knowledge, rather than necessarily being an expert in it, which the learner is then likely to use at a later date, whereas:
- An **expert** is someone who knows all that anyone is likely to need to know about a topic and can advise on immediate issues and impart some information for later use.

In order to better understand the differences between these roles it can be useful to plot them on a grid - the **Personal Support Space** (Jeffrey, 2016)¹⁴.

13 <https://ica-uk.org.uk/what-we-mean-by-facilitation/>

14 Jeffrey, A., (2019). "What is the difference between a coach, mentor and facilitator?". Al Jeffrey. Available at <https://aljeffery.com/coach-mentor-facilitator-whats-the-difference-some-core-distinctions/>



The Personal Support Space has two dimensions: **Experience** and **Exchange**.

- **Experience** varies from less experience or knowledge of the domain under discussion to a high level of experience and knowledge.
- **Exchange** varies between given knowledge to realised knowledge. In a relationship at the “given knowledge” end of the spectrum one person (or agent) is providing knowledge to the other, whereas at the “realised knowledge” the person (or agent) is drawing existing information out of the other person and helping them to discover and “realise” that knowledge.

For instance a traditional teacher (or expert) operates in the High Experience, Given Knowledge area of the space, whilst someone like a business consultant or mentor may operate in the Low Experience (they don’t need much domain specific knowledge of the business as they will draw that out of their client) but high Realised Knowledge (they get the client to

constantly reflect, it's not up to the mentor to provide the answers).

Using the Personal Support Space models the various roles of interest to organisational learning could be mapped as follows:

- **Coach** - Medium Experience, balance of Given and Realised Knowledge
- **Mentor** - High Experience, emphasis on Realised Knowledge
- **Facilitator** - Low Experience (of target domain, not methods), emphasis on Realised Knowledge
- **Expert** - High Experience, High Given/Exchange

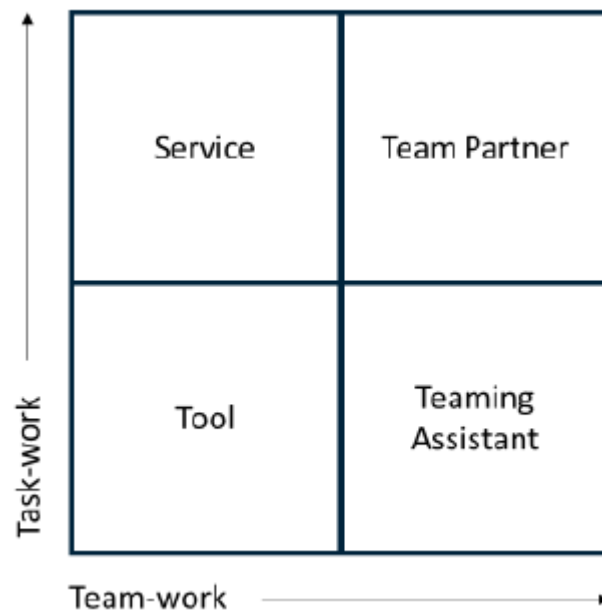
AI Agents

Within the context of organizational learning, there are two useful models when it comes to considering the potential role of AI agents.

Team-Work vs Task-Work

The first is to consider the extent to which the AI agent is focussed on the task-work (high domain expertise, low people orientation) or the team-work (i.e. minimal domain knowledge, high team/people orientation) of an activity¹⁵.

¹⁵ Farry, R. (2022). Human Autonomy Teaming for Adaptive Systems (HATAS): Lessons Learnt, Design Guidance & Best Practice Report. D0.9 DSTLX-1000155787. Serapis Framework/Qinetiq.



Roles such as **Mentors**, **Coaches** and **Facilitators** fit very much into the high Team-Work, low Task-work “teaming assistant” quadrant. In contrast an **Expert** probably fits in the high Task-Work, low Team-Work “service” quadrant, as may a **platform-persona** providing a manifestation of a whole industrial or business system (see later). Tools might represent simpler IT and even AI systems, but could also include:

- **Knowledge Interfaces** - which provide a conversational AI interface to corporate systems, including knowledge management systems; and
- **Virtual Librarians** - which can help in the curation and identification of knowledge, directing the user to the right sources and repositories - whether paper, simple IT or AI based. In fact, this was the role played by an AI in the *Snow Crash* novel by Neal Stephenson which originated the term “metaverse”¹⁶.

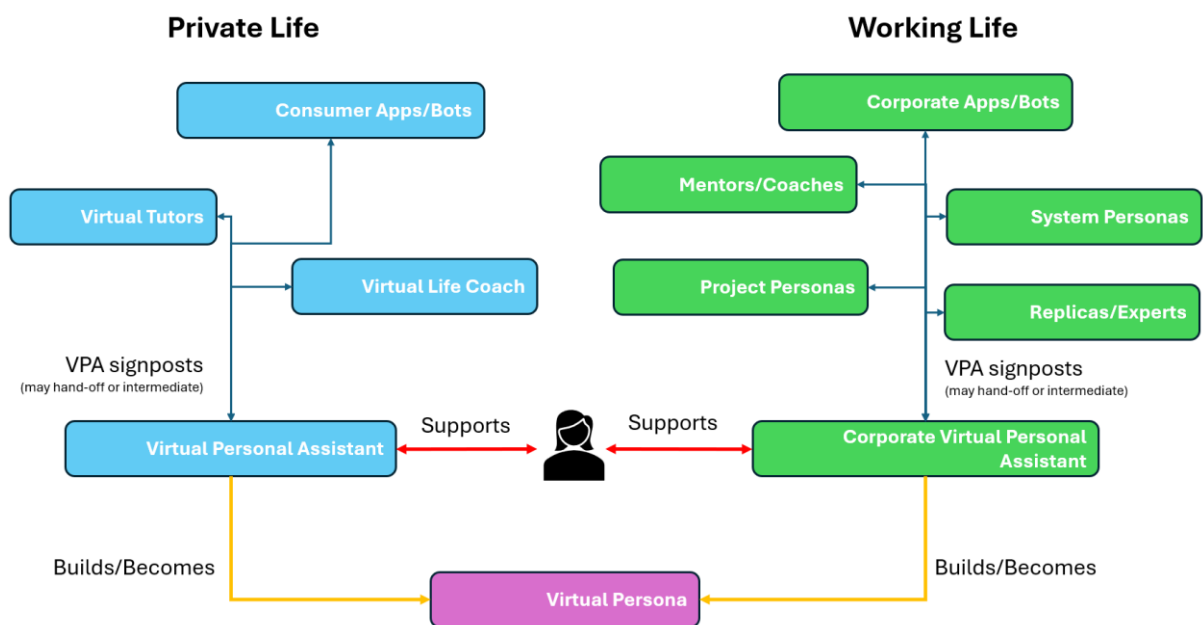
Sitting perhaps centrally in the grid is the idea of the **Virtual Assistant**, which know enough about a task to be helpful, but which also can keep an eye on the performance of the team, and of you and the team, and, again, like the **Virtual Librarian**, knows

¹⁶ https://en.wikipedia.org/wiki/Snow_Crash

which other AIs and systems to refer you to, to help find answers and to progress an activity.

The Team Partners quadrant is where we would probably expect (or hope) to find our high-performing human team-mates - but is an area in which we may yet find AI partners.

The result is hence a constellation of AI agents, all playing a specific role, with personalities optimised to the role (and to your needs) in order to complete the task, build the team, and develop the organizational learning.

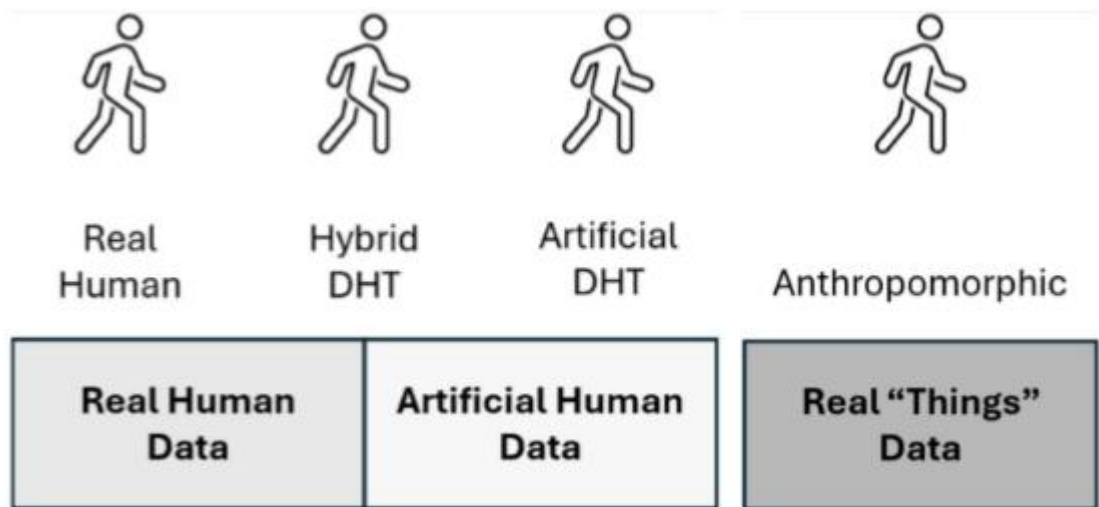


An interesting question is whether we may have two **virtual personal assistants (VPAs)**, one a lifelong personal companion and one an employment episode assistant, or whether the personal VPA will be granted access to corporate systems in order to be most effective so that we only have one VPA to deal with - shades of whether you carry one or two mobile phones around. Another aspect is the extent to which the VPA deals with all of your interactions with other systems and bots (through APIs, as in an Agentic model) or whether it acts as more of a signposting agent, directing you to the right bot/service which you can then directly engage with. The reality is likely to be a mix of the two, depending on the task and immediate context. Perhaps the most significant part of the diagram is how the VPAs become the ideal vectors by which to build our own virtual replica/persona, which in time might feed back into the top half of the diagram for another user.

Digital Twins

Digital twins (DT) are an increasingly common concept in civil and mechanical engineering and are based around the idea of having a digital copy of a physical system, which can be fed the same inputs as the physical system and be used to model future performance and changes, and to alert to potential future problems and failures¹⁷.

An emerging concept is that of the digital human twin (DHT).



A recent paper I co-authored¹⁸ divides DHTs into four types:

- **Real Human DHTs** - which look human (as avatars or even robots) and are built from information gathered from a real human or humans;
- **Hybrid DHTs** - which again look human (or humanoid) but are built from a mix of real human derived information and synthesised and synthetic/machine-derived data (perhaps representing a whole cohort to create a generic type);
- **Artificial DHTs** - which look human but are built solely from synthesised and synthetic/machine-derived data;

¹⁷ <https://www.digitaltwinconsortium.org/>

¹⁸ Fawkes, A. & Burden, D.J.H. (2025). Digital human twins and the military metaverse: opportunities and challenges. *AI & Society*. <https://doi.org/10.1007/s00146-025-02508-2>. Access at: <https://rdcu.be/eBhUm>

-
- **Anthropomorphic DHTs** - which look humanoid but represent one more computerised system which represent part of the physical “real” world - ranging from perhaps a complex business management system to a factory, warship or an entire airport.

All of these have application to our current discussion on organisational learning:

- **Real Human DHTs** would enable us to create copies of physical individuals, often called **AI or Digital Replicas**, so we can access their knowledge, wisdom and expertise long after they have left the organisation - but such a persona is “warts and all”, reflecting the progenitors biases, gaps, errors and other short-comings.
- **Hybrid DHTs** - let us “fix” the deficiencies of the Real Human DHTs, filling in gaps and errors in knowledge and potentially building from multiple human sources;
- **Artificial DHTs** - break any explicit connection to the human sources but become the eternal “experts” within an organisation, knowing all about a system, territory, product, project, client etc.
- **Anthropomorphic DHTs** - extend the artificial DHTs to represent *things* in real-time, letting you interact with the system as though you were talking to a person. They can talk about the systems past, present and future, talking about the history of the thing they represent, what has been done to it, how it’s been used, how it’s been maintained and repaired etc.

The **Mentors, Coaches** and **Facilitators** discussed earlier are likely to be implemented as **Hybrid** or even **Artificial DHTs** (although a particularly talented individual might be immortalised as a **Real Human DHT**). Experts could be represented by any of them, from **Real Human DHT** to **Anthropomorphic DHT**. The **Virtual Librarian** could even be an **Anthropomorphic DHT** - the library made “flesh”, and the **Virtual Assistant** an **Artificial DHT**.

AI Agents and Organisational Learning

This section explores in more detail how each of these main roles within an organisation concerned with organisational learning could be implemented by an AI agent, and what some of the advantages and challenges may be.

AI Mentors, Coaches and Facilitators

As described above, the role of the mentor, coach or facilitator is one which is about a high exchange dynamic and a reasonable experience dynamic. They should know something about the domain being talked about - but doesn't need to be an expert; but does need to be an expert in how to get the subject to think about their own knowledge, beliefs and approaches and to try and pull the answer or action out of them. A coach is perhaps more often associated with team rather than individual performance, so should know more about how teams work and how to create high performing teams, including though approaches such a **Shared Mental Models**¹⁹ and **Action vs State orientation** (Kuhl & Kazen, 1994)²⁰. A facilitator is solely team focussed and probably has the least domain knowledge and also tends to use more proscriptive approaches.

Most LLMs probably already have the general knowledge about a domain to equip them for a coaching or mentoring role, and this can always be augmented by domain or organisation specific information through a **Retrieval Augmented Generation**²¹ (RAG) type system. Even just acting as a signpost to different support services can be useful in the more pastoral care aspects of a mentor role. As such the agent is reflecting best practice based rather than capturing the idiosyncratic knowledge of individuals. Where knowledge capture might be useful is in capturing stories and anecdotes relevant to the company which can then be used to illustrate "standard" advice.

19 Andrews, R. W., Lilly, J. M., Srivastava, D., & Feigh, K. M. (2023). The role of shared mental models in human-AI teams: a theoretical review. *Theoretical Issues in Ergonomics Science*, 24(2), 129-175.

20 Kuhl, J. & Kazen, M. (1994). Volitional aspects of depression: State orientation and self discrimination. In J. Kuhl & J Beckmann (Eds.), *Volition and personality: Action vs state orientation* (pp. 297 315). Seattle: Hogrefe & Huber.

21 https://en.wikipedia.org/wiki/Retrieval-augmented_generation

More important to the mentor/coach/facilitator are the models and approaches it uses to deliver the mentoring and coaching and to get the user to realise the knowledge and wisdom that is within them. This could make use of industry standard mentoring and coaching models such as **Dyadic mentoring, flash mentoring and the GROW, OSKAR, CLEAR, and FUEL coaching models**. Such models need to be scripted in some way into the agent. With an LLM it may suffice to provide them with documents via RAG which describe the process, and then a system prompt which explains how to implement these models in context. One of the big advantages of an agent is that it can be available 24/7, so is not only always on-hand to answer user questions or to provide coaching/mentoring as required but can also prompt and “nag” the user to do exercises or apply learnings as they go about their daily tasks. There is probably a blurred line between a well-developed mentor and virtual counselling agents which can help apply standard counselling techniques, particularly talking therapies such as **Cognitive Behavioural Therapy**²².

Here are three simple examples of how mentor agents might work, drawn from some of my past projects:

- **Student Mentor** - an agent was developed to help young soldiers who had just completed basic training but were about to embark on up to a year’s worth of technical study in order to prepare them for careers as army technicians. So, they not only had to deal with learning technically complex things but also surviving (and excelling) in a military environment. The mentor agent prototype that was developed for them knew their daily schedules, could advise them on where and when they needed to be, and with what kit, could link them to additional learning resources on technical and military topics and connect them with self-help resources for pastoral care - and links to human support. The students found the mentor very valuable, and particularly the way that it let them ask

²² Vowels, L. M., Francois-Walcott, R. R., & Darwiche, J. (2024). AI in relationship counselling: evaluating ChatGPT’s therapeutic capabilities in providing relationship advice. *Computers in Human Behavior: Artificial Humans*, 2(2), 100078.

questions in confidence and grow their own knowledge before deciding whether to ask for human support.²³

- **Six-Hats Agent** - An agent was designed for a technology research organisation to implement **De Bono's Six Thinking Hats** model²⁴. This is a way of getting someone to think through a particular problem by considering questions (based on notional coloured hats) about the pros and cons of an issue, what facts are known, what the gut feel response is, thinking creatively about it, and thinking about the processes. The great advantage of the process is that it needs no domain knowledge on the role of the questioner (the agent) but implementing it using an LLM the agent inherited enough general knowledge about most topics to be able to answer meaningful follow-on questions. Such an agent is also known as a **Reflexive Agent** as it is getting the user to reflect on their own knowledge.
- **Teaming Assistant** - As part of a collaborative data visualisation virtual environment an agent was developed to fulfil three roles. The first was akin to being a systems interface (see below) which enabled users to just chat to the agent in order to access datasets and manipulate them. The second was more like a virtual assistant, recording each of the sessions that users did with the datasets, and then, when new users looked at similar data, could offer up the findings and analyses of previous users, and even replay complete sessions. The third mode was where it acted as a group coach or mentor, using the **Action vs State Orientation model**²⁵ to track how much participants focussed on the team or the task, and where team members or data set features were being overlooked and bring them to the attention of the users.

Mentor and Coaching agents can have quite different challenges associated with them in comparison to other types of agents as their focus is not so much on (possibly confidential) corporate knowledge and information (although their

23 Burden, D.J.H. (2019) Developing Conversational AI solutions for the MoD | CogX 2019 - <https://youtu.be/SQj8kTpN7lw?si=1bqMJQcG1joYYWf0>

24 <https://www.debonogroup.com/services/core-programs/six-thinking-hats/>

25 Kuhl, J. & Kazen, M. (1994). Volitional aspects of depression: State orientation and self discrimination. In J. Kuhl & J Beckmann (Eds.), Volition and personality: Action vs state orientation (pp. 297-315). Seattle: Hogrefe & Huber.

conversations can reveal that), but on the more personal and private knowledge, concerns and even weaknesses and behaviours of the user. In some mentoring roles (such as with the students above) this may even stray into safeguarding territory. As such there need to be clear ethical guidelines to the implementation and operation of the agents (including possibly real-time alerting of human staff), as well as privacy and corporate IP wrappers around the content of the conversations - and in the privacy case clear transparency as to what information is captured, how it is stored, and who can access and use the information.

Surprisingly, it may currently be the facilitator role that is the hardest to implement with an AI agent as anyone who has ever been in, let alone tried to run, a facilitation workshop will know how chaotic it is and very akin to herding cats. Being able to “read the room” is probably also a key skill, and so the role can be more dependent on understanding and responding to social dynamics than the other roles. It may be that the way forward is to use agents to help facilitate particular parts of a workshop (for instance acting as note-takers and presenters which no-one ever appears to want to do or taking participants through a particular process) and then to grow their involvement from there. Virtual facilitators might also find early use in virtual worlds where they can interact directly with the environment, and their ability to “read” the room is facilitated by the fact that the whole space and all participants are also virtual.


AI Experts

The expert is a very different role to the mentor, coach or facilitator. It is all about high (even the highest) levels of Experience dynamic and Given knowledge. Before delving into how this might be realised through an AI agent it is worth considering some different facets of knowledge.

The first is the differentiation between **semantic, episodic and procedural knowledge**. These terms are often found when talking about how human brains manage and store information and can be useful in designing AI agents too. **Semantic knowledge** is the knowledge of things - what they are, what they do, how they are made, how they interact etc. **Episodic knowledge** is the knowledge of things that have happened - typically to the user but could be extended to the experiences of a system, project or organisation. **Procedural knowledge** is the

knowledge of how to do things - how to maintain a piece of equipment, how to make a sale, how to manage an employee etc. All are inter-related, but when thinking about how we might create a virtual expert we need to decide whether all three forms are relevant, or whether the agent only needs to have semantic and/or episodic knowledge.

Another useful distinction is between **knowledge** and **wisdom**. Ackoff's DIKW model shown below is quite commonly used in this regard²⁶. Whilst we might say we want to capture knowledge often what we'd really like to capture is wisdom. Wisdom seems to be all about choose the right frame or frames to examine a problem, and to ensure that the frame has a wide enough scope (in range and time) to capture unintended consequences²⁷. It can be useful to consider that whilst we can readily conceive of an intelligent person making a mistake it is a lot harder to think of a wise person as making a mistake. There is also a strong **reflexive role** (as discussed above) connected to wisdom - the archetypal guru sat on a mountain top who answers every question with another question in order to draw the knowledge (and wisdom) out of the person!



26 Ackoff, R. L. (1989). From data to wisdom. *Journal of applied systems analysis*, 16(1), 3-9.

27 Vervaeke, J., & Ferraro, L. (2013). Relevance, meaning and the cognitive science of wisdom. In *The scientific study of personal wisdom* (pp. 21-51). Springer, Dordrecht.

53 Watson, T. (2012).

A final useful distinction is between **explicit knowledge** (knowledge stored on paper and in systems) and **tacit knowledge** (knowledge stored in people's heads).

Creating an **AI Expert** is, on the face of it, a relatively straightforward process - you upload all your documents on the subject through a RAG interface to an LLM chatbot and you're done. Unfortunately, things are rarely as easy as that. Even if all the information you want is in the documents it may not be expressed as knowledge, let alone wisdom. Even with RAG the agent is likely to hallucinate answers, unless you have adequate hallucination detection. Retrieval is also rarely 100% even if the information is there. The biggest problem though is that a lot of the knowledge (and particularly wisdom) you want the **AI Expert** to have is still in people's heads.

Building a truly effective **AI Expert** probably requires a multi-pronged approach:

- The organisation needs a consistent and sustained approach to make more of its tacit knowledge explicit, and particularly capturing tacit knowledge through an effective off-boarding (or ongoing) process before it leaves the organisation. One way of achieving this may be through the use of **AI Offboarding Agents** who can have "intelligent" conversations with leavers in order to identify and capture their unique knowledge and wisdom – although leavers may not be in the best frame of mind, so some creative approaches might be required – for instance rewarding people for use of the knowledge (not just its capture), even once they've left the organisation. Exit capture should probably be in parallel with making the knowledge capture process a continuous one through employees' interactions with their virtual personal assistants and other corporate agents ;
- Information needs to be converted to knowledge, either by users or by the system, and the system needs ways in which it can then create new knowledge from the information (knowledge development). This may be an area where LLMs using unstructured text are less than ideal and where other approaches, such as knowledge graphs, operating in a hybrid model with LLMs have something to offer; and
- The **AI Expert** needs to show wisdom, not just knowledge. This may be something that can be introduced through

effective writing of system prompts in order to get the agent to “think” in a wise way, to know when and how to use reflexive techniques, and to work with the user to develop the relevant frames with which to consider a problem. An example of this approach is in a recent attempt to get an LLM chatbot to think in a hermeneutic way (a form of iterative thinking) simply through writing a suitable prompt²⁸.

As suggested above, the **AI Expert** has a different set of challenges to the **AI Mentor/Coach/Facilitator**. The **AI Expert** is all about the organisation’s intellectual property (IP), and so needs to be adequately protected against inappropriate, and even malign, access and use. There is also a potential grey area between organisational knowledge (and wisdom) and personal/employee knowledge (and wisdom), and so employment contracts need to clearly define what information an organisation has a right to capture and use from an employee, and what (tacit) information the employee has a right to take with them. Whilst such issues could be of some importance in the pre-AI era when there was just the organic employee, in the world of **AI Experts** the company potentially has access to limitless clones of the employee’s knowledge and expertise, and the (now ex) employee has the potential to create limitless clones of themselves. So apart from the legal issues of IP ownership this also starts to bring in ethical issues, particularly around issues such as who has rights of access to, and use of, such **AI Experts** of the Real Human Digital Human Twin (DHT) variety if the subject of the DHT dies²⁹.

These ethical issues may well be what drives organisations away from Real Human DHTs and towards Hybrid DHTs where the **AI Expert** is the “best practice” amalgam of multiple employees, rather than the **AI Replica** of individual employees.

28 https://medium.com/@jordan_gibbs/the-most-esoteric-chatgpt-prompt-3ff028a6f84a
29 Savin-Baden, M., Burden, D., & Taylor, H. (2017). The ethics and impact of digital immortality. *Knowledge Cultures*, 5(2), 178-196.
https://www.researchgate.net/publication/325429129_The_ethics_and_impact_of_digital_immortality

Other Agent Roles

System Personas



Whilst Rommie from Andromeda might embody the worst of the female avatar/chatbot tropes, the concept behind her/it is a very sound one. What if your complicated system or platform – be it a factory or warship or even a business – was represented or fronted by a conversational AI – one which could be present on Zoom calls or in virtual or mixed reality sessions. You wouldn't ask a colleague how the system was – you'd ask the system itself. And the system/AI would build up years and years of experience in how it works, what has happened to it, what can go wrong, how to fix it, and how to get it to operate most efficiently and effectively.

Before concluding, there are a couple of other AI agent roles which may be important in a learning organisation.

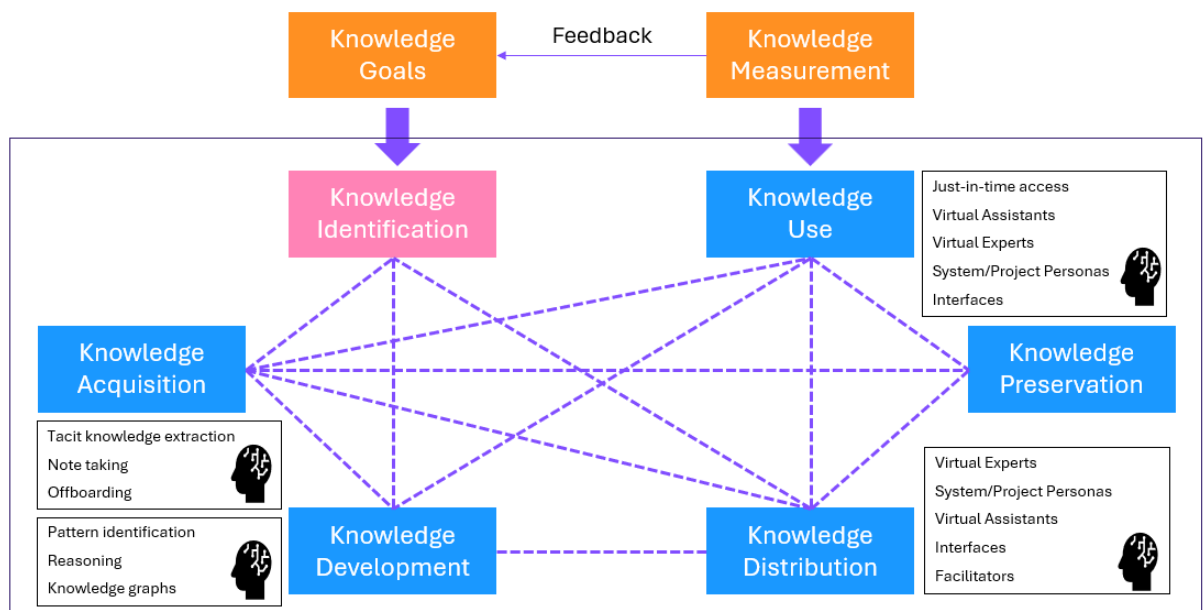
The first is the idea of agents as **System Personas**, or even more simply as **Chatbot Interfaces** to corporate IT (and industrial systems). In the latter case the agent is relatively dumb, just providing a natural language interface to the system, turning spoken (or typed) plain language requests and commands into system instructions, and system data into verbal (or text) reports. A more sophisticated system might include an advanced form of on-line help, letting users achieve their aims without a detailed knowledge of the system. An even more advanced manifestation might be interfaced with a digital twin of the system so that it could run what-ifs and predictions. In its ultimate form as a **System Persona** the agent may not only present in some anthropomorphic form but also knows the entire history of every user's interaction with the system, and how it, as the system, has performed, and even "thinks" about itself as the system. A variation of this is the **Project Persona** which holds all the knowledge and experience of everything, and everyone, to do with a specific project. ❖

The second is the much discussed idea of AI agents as **Virtual Assistants** or **Virtual Personal Assistants (VPAs)**. Such an agent looks after routine tasks such as booking appointments, finding information, writing documents, setting reminders, taking notes and may also act as a search and discovery tool to identify other agents and systems which can help you. In its basic form the virtual assistant won't store much information itself but will store information about your preferences and activities. However a virtual assistant is in an ideal position to be able to start to build a virtual replica of you, gathering information both explicitly and tacitly, using down time to ask specific questions to fill in gaps, and ultimately creating with a **Digital Replica** that could act as a version of you (Real Human DHT), or creating the information which could be incorporated into a more generic virtual expert (Hybrid DHT).

Agents as Learning and Knowledge Custodians

To conclude this look at how AI agents can contribute to continuous organizational learning, it is worth zooming back out to the models of organizational learning and knowledge management to see not only how the agents can support the different processes within the models, but also how agents can potentially help to be the custodians and champions of the meta-processes themselves.

For **knowledge management** AI agents can support many of the activities in the knowledge management model as shown below. There is probably some role for AI agents in Identification, Measurement and even Goal setting, but Preservation should be delivered through a robust back-end system which the AI agents then access.



The three areas where AI agents offer the biggest initial impact are probably in:

- **Tacit knowledge capture (Acquisition)**, whether when offboarding or during employment;

-
- **Just-in-time access (Distribution & Use)**, supporting “vibe writing” of documents and presentations³⁰ and making personal and corporate knowledge (and wisdom) available at point of need - even if you’ve forgotten it’s there; and
 - **Knowledge distribution *through time***, so not just across an organization but across the different generations of employees, whether through **Virtual Experts**, **System/Project personas** or other AI agents.

For **organisational learning** whilst AI agents can have a role (particularly through assistants, experts and system personas) in creating effective single and double loop learning, and can act as mentors and coaches to help capture and share learnings, their key role is possibly in helping to make implicit knowledge explicit and in doing so to help transform **individual mental models** about how an organisation works into a **shared mental model**. This again builds on the idea of tacit knowledge capture and knowledge distribution across the organisation, and making the information available at the point and time of need. Ultimately the AI (or more realistically AIs) becomes a part of the team, a part of the organisation, and is present in the corporate shared mental model as well as being its chief enabler.

Conclusions

The aim of this paper has been to identify how AI Agents can support the organisational learning process, and the associated knowledge management process. In both cases a key contemporary challenge is that much of the knowledge and learning is in people’s heads - as tacit knowledge and individual mental models. Implementing AI Agents purely as delivery tools, as chatbot interfaces to existing, explicit, corporate knowledge systems will not deliver step changes in organisational learning or corporate efficiency. This is almost the reverse of the “lesson distribution gap” introduced earlier. Organisations need a twin track approach. Yes they need to be using AI agents to help provide access to information, to distribute it (in place and time), and ideally to make it available just-in-time, proactively at point-of-use, but they also need to be using AI agents to capture

³⁰ <https://deleyna.substack.com/p/what-is-vibe-writing>

implicit knowledge and to make that knowledge explicit, so that it can be used throughout the organisation, and so that it can help to grow the shared mental model which truly marks out a learning organisation.

MORE INFORMATION

For more information on the topic of conversational AI and its application to knowledge management please contact the author, David Burden, at david.burden@daden.co.uk.

ABOUT THE AUTHOR

For the last 20 years David has run Daden Limited, helping organisations explore and exploit the social and commercial potential of using conversational AI and virtual worlds, delivering over 100 projects for clients across the globe, including over two dozen projects for UK MOD. Daden were finalists in the BCS Machine Intelligence Competition, and chatbots designed by David successfully passed two covert Turing Tests in the 2010s. David spoke at the inaugural TEDxBrum, on Digital Immortality, has authored over a dozen papers and book chapters, including co-authoring the book *Virtual Humans*, published by Taylor & Francis, New York. David now operates as an independent consultant and is also currently undertaking a PhD on wargaming urban conflict. David is an ex-Royal Signals officer, a Chartered European Engineer and series co-editor for Taylor & Francis on their Metaverse Series of books, and co-author of both *The Metaverse: A Critical Introduction*, and *The Military Metaverse*.

RELATED WORKS

Burden, D.J.H. (2025a). *Chatbots - Past, Present and Future*. Daden Limited.

Burden, D.J.H. (2025b). *Military Applications of Conversational AI*. Daden Limited.

Burden, D.J.H. & Savin-Baden, M. (2019). *Virtual Humans: Today and Tomorrow*. Taylor & Francis.

Savin-Baden, M. & Mason-Robbie, V. (2020). *Digital Afterlife: Death Matters in a Digital Age*. Taylor & Francis.

Burden, D.J.H. & Savin-Baden, M. (2024). *The Metaverse: A Critical Introduction*. Taylor & Francis.

Fawkes, A, & Burden, D.J.H. (2025). *The Military Metaverse*. Taylor & Francis.



About the Author



David Burden has been involved in AI, VR and immersive environments since the 1990s. David set up Daden Limited in 2004 to help organisations explore and exploit the social and commercial potential of using chatbots, AI and virtual environments. David and his team delivered over 100 immersive learning and chatbot projects for clients across the globe. David has led collaborative research projects funded by InnovateUK and the MOD, and won prizes for his work in the UK and USA. David was a finalist in BCS’s Machine Intelligence competition, spoke at the inaugural TEDxBrum on Digital Immortality and has authored over a dozen papers on virtual worlds and AI. David co-authored *Virtual Humans*, a book published by Taylor & Francis, New York, which was inspired by an MOD project, and is currently series co-editor for Taylor & Francis on their Metaverse Series of books. David has led nearly two dozen projects for MOD around conversational AI, autonomous agents, chatbots, virtual reality and wargaming, including a 3 year project on *Virtual Humans*. David also led the authoring and production of a “Biscuit Book” for Dstl on Evidenced Based Models for VR Training. David is an ex-Royal Signals officer, a Chartered European Engineer and is currently studying for a part-time PhD in Wargaming Urban Conflict.

info@daden.co.uk

www.daden.co.uk

© 2026 Daden Limited