



Source: ChatGPT

# The Military Applications of Conversational AI

February 2025

David Burden  
DADEN LIMITED



---

## Terminology

Note that this White Paper considers the terms conversational AI, chatbot and even ‘bot as synonymous, all describing a computer programme which you can converse with through voice or text chat.

The entity may be represented as a simple text chat interface, a voice on a radio, a static or animated image, as video, or as a 3D avatar within a virtual world.

The terms “virtual human”, “digital replica” or “virtual persona” describe a more sophisticated conversational AI which has many of the characteristics of a human being, including memory, intention, and even emotion.

A non-player character (NPC) is a specific sort of bot, which may include a chatbot, found in games and training applications.

## INTRODUCTION

This White Paper looks at the military applications of chatbots and conversational AI. However, many of the concepts, ideas and use cases also apply to civilian organisations, in both the private and public sector.

The White Paper is divided into 5 main parts:

- An analysis of the non-operational use of chatbots using the UK's TEPIDOILI framework (similar to the US DOTMILPF framework);
- An analysis of the operational use of chatbots using the J1-9 staff functions, and including specialist roles such as LEGAD and POLAD;
- A consideration of how operational use of chatbots changes for different types of military operation;
- An analysis of the implications of deploying chatbots using the TEPIDOILI framework; and
- A consideration of the security, validity, explainability and ethical issues of deploying chatbots.

In addition, Annex A considers the role of conversational AI in Wargaming.

---

## NON-OPERATIONAL ROLES AND USE CASES

The TEPIDOILI framework is used by the UK military to evaluate the Defence Lines of Development (DLOD) implications of a change across the military estate. The elements are:

- Training;
- Equipment ;
- Personnel;
- Information;
- Doctrine and Concepts;
- Organisation;
- Infrastructure;
- Logistics;
- Integration.

Similar models are used by other militaries, and Table 1 provides a comparison of TEPIDOILI and the US DOTMLPF framework.

TEPIDOILI	DOTMLPF
Training	Training
Equipment	Materiel
Personnel	Personnel
Information	
Doctrine and Concepts	Doctrine
Organisation	Organisation
Infrastructure	Facilities
Logistics	
Integration	
	Leadership and Education

*Table 1: Comparing TEPIDOILI and DOTMILPF*

The opportunities for the deployment of chatbots in each of these areas are considered below. Towards the end of the White Paper we also look at the TEPIDOILI implications of conversational AI technology and systems itself.

### Training

There is immense scope for the use of conversational AI within training (both individual and collective) and education. At their most basic chatbots can help to guide and narrate learners and trainees

---

through content, to answer questions about a topic, and even to pose questions and quizzes to learners. Such systems are often known as **virtual tutors**, and can be combined with AI to offer students different ways of approaching a topic, and to adjust the challenge level of learning to match the students ability.

Another use of conversational AI is to power **non-player characters (NPCs)** within scenario based virtual learning (2D or 3D) exercises. For instance, the AI can play the role of a squad member with a welfare issue that the student needs to work through, or to play a local civilian in a foreign country (and even the interpreter) with whom the student has to interact in a culturally appropriate way.

As the military make more integrated use of **Live-Virtual-Constructive-Autonomous** training environments, conversational AIs could appear in all of these, typically as NPCs. NPCs can be used as part of the synthetic information wrap in any element of LVCA training (see sidebar), providing either a passive narrative on what is happening from the point of view of the civilian population, or as interactive agents which civil-military cooperation (CIMIC) and influence operations (IO) can try and engage with.

There have been a number of projects to look at the use of chatbots to support **After Action Reviews (AARs)**, whether from training exercises, operations or wargames (Eaton & Oлару, 2021). These can not only act as automated scribes, but can also manage interactive Q&As, and ensure that all lessons identified are properly captured, coded and logged. Chatbots can also dramatically increase the number of one-on-one AARs/interviews conducted, uncovering issues which may be missed by collective AARs/interviews – with data capture being enabled by a simple conversation with an app on your phone (or approved military terminal/BYOD/PED).

The use of chatbots within wargaming , both as a training/education exercise and as a COA activity, is a large topic, and is discussed in more detail in Annex A.

## Replicas vs Project/Platform Personas

A **replica**, or **virtual persona**, is a warts-and-all copy of a person, including their biases, errors and gaps in knowledge – think the headmaster portraits in Dumbledore’s study.



A project/platform persona (or virtual assistant) provides the sum of knowledge and considered best opinion, hopefully bias-free, and can also challenge the user in a reflexive way – think the daemon’s in Philip Pullman’s *His Dark Materials*, or the Cortana avatar in Halo.



## Equipment

Whilst some of the more trivial applications of conversational AI to equipment might include **conversational interfaces** in order to learn how to operate a piece of equipment, there are two particular use cases which might have a greater long term impact.

The first is the concept of a **Project Persona**. Most military equipment projects last for decades, but military staff are often moving on every 2-3 years, and so corporate knowledge and memory continually leaches out of the project. What if every project (major or minor) had a conversational AI persona which could not only help current project members find all the information relevant to the project (and bring it to their attention at the right time even if they didn’t know it existed), but could also reflect the views, memories, intents and rationales for decisions of everyone who had ever been involved in the project?

The second is an extension of the conversational interface, but where the equipment – or more realistically in this case the platform (such as a warship) – has its own persona, a **Platform Persona**. The platform could be a ship, or an aircraft, a missile system or even an AFV. This, like the project persona, includes not only all of the knowledge of capabilities and experiences of the platform (and acts as a gateway to a range of more dedicated AIs doing things like proactive maintenance prediction, situational awareness and Course of Action analysis), but also the sum knowledge of experiences of all of its crews, all of its sister platforms, all platforms that have fulfilled the same or similar role, and perhaps even draws on the traditions and experiences of earlier instances of the class or name. This is an idea which was explored in *“It takes three hundred years to build a new tradition”* (Strong et al., 2018), and also explored by Daden in their Artificial Intelligence Decision Aid (AiDA) project as part of the UK MOD’s Intelligent Ship programme. In fact, in that project Daden made the case that rather than think about an “intelligent” ship, militaries should think about how you create a “wise” ship.

## Personas in Science Fiction

Platform Personas are present in science fiction literature (for good or ill) and include:

- HAL in *2001*
- Zen in *Blake's Seven*
- Romy in the *Star Trek* spin-off *Andromeda*.

And, of course, the sentient ships of Iain M Bank's *Culture* novels.

Virtual Personas are even more prevalent and some of the more interesting include:

- Zoe in *Caprica*
- The Doctor in *Star Trek Voyager*
- Data in *Star Trek: The Next Generation*
- Ava in *Ex Machina*
- C3PO and R2D2 in *Star Wars*

Of these, Zoe is unique in being a copy of a real singular human being.

The Data-Information-Knowledge-Wisdom pyramid shown in Figure 1 is fairly common, but little attention is often paid to the Wisdom layer, and in the AiDA project we also highlighted how humans (and so platform personas) often derive a lot of information from the stories ("war-stories" and the like) that they hear, and that often the best way to impart, justify and explain information is through the use of stories and recollections of similar occasions.

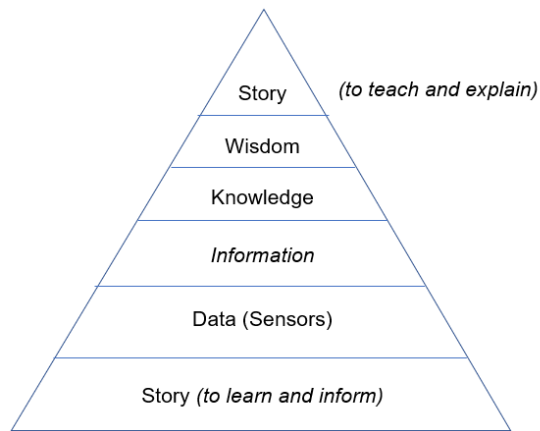


Figure 1: The Extended DIKW Diagram

In civilian terms a **platform persona** could be equated to something like a **factory persona** or a **system persona**.

## Personnel

In the Virtual Humans project for UK MOD (Burden, 2019; Burden & Savin-Baden, 2019) Daden investigated the concept of **the virtual persona** or **digital replica**, trying to create a virtual copy of a real person. A major challenge for the military is that, as previously mentioned, people only stay in post for 2-3 years, and then may leave the service or get posted to somewhere with poor communications, and so knowledge retention is a major challenge. What if a post-holder effectively had every previous post-holder on speed-dial? What if whilst you are in a post your virtual assistant is helping to build a virtual persona of you that can stay behind when you move on and help the next incumbent? The ethical issues of this will be considered below! Building such a system in 2019, in pre-LLM days, was a challenge, and we never expected to be able to make a perfect virtual persona, but what we were able to do was to help to show the merit and feasibility of the idea, and with modern LLM/RAG and GAN

## Grice's Maxims

Grice's Conversational Maxims (Wilson & Sperber, 2022) are a useful way of evaluating a conversational AI. They ask whether a chatbot's response:

- Is **truthful**;
- Provides a helpful **quantity** of information, not too much or too little;
- Is **relevant** to the question asked or the topic being discussed;
- Is presented in a **clear and orderly** manner.

To these we would also add:

- And helps to **move the conversation forward**.

Even many humans would score low on Grice's Maxims for some of their conversations!

based systems, such as Sensay's **Digital Replicas**

(<https://sensay.io/>), the creation of such virtual personas is far more feasible.

One other thing that the project did was to highlight the major ethical issues around the creation of virtual personas based on physical people. For instance, if the persona is good enough why bother employing another person to do the job, or can the employer even make multiple copies, or can the original person use the persona to cover for them when they are ill, on holiday or just busy? Or can the persona carry on working even after the real person has retired (and who gets the money)? And what about when the person dies – does their next of kin have a right of access, or the employer, and again can it carry on working – and if so, for who? Whilst concepts such as a digital afterlife and even digital immortality may seem far off, even a basic virtual persona begins to raise some of these questions (Burden, 2020). ☒

To bring things back down to earth a bit, we see it as vital that almost all of these considerations of conversational AI, and particularly of virtual personas and platform personas, are considered within the context of corporate knowledge management model (Figure 2) (Probst & Romhardt, 1997), and particularly how they contribute to the acquisition, distribution, preservation and use of corporate knowledge.

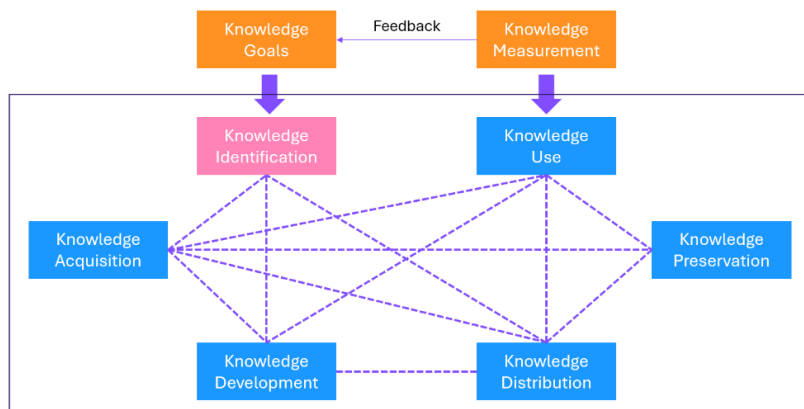


Figure 2: A Knowledge Management Model (after Probst, 1997)

---

## A More Human Conversation

One area of research that we looked at as part of the **Virtual Human** project for MOD was how to make chatbot conversations more like human conversations (Burden, 2019). The way the system worked was that the user would ask an initial open question and the chatbot would then find a good answer to that, but then if the user said nothing else (or even nodded their head), then the bot would find some related information, and then some more and so-on. This was facilitated by the fact that all the information was stored in a knowledge graphs, so there was a very explicit link between concepts and information, but no doubt prompts can be developed for LLM/RAG systems to behave in the same way.

It should be noted that one of the reasons why Daden moved from the **Virtual Persona/Virtual Humans** project to a **AiDA Platform** or **Project Persona** project was that the latter are seen as having less ethical considerations since they are reflecting an aggregate of information held by individuals, rather than trying to replicate a single individual.

## Information

That conversational AIs should be able to improve the access to information, and should perhaps be available as an interface (maybe even the default interface) to almost every information store, can almost be taken as a given. This technology can help with almost every aspect of the Knowledge Management model presented above. There are 3 specific issues to briefly consider here.

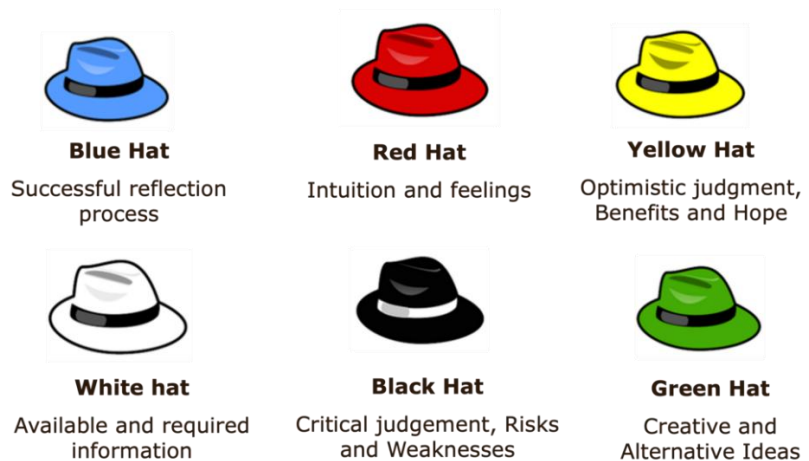
First, a lot of the time when we are talking to someone we don't actually know the precise information we want. We find it by asking open questions, listening to (a relatively long) response, and steering the conversation in the right direction, and then diving into the detail once we hit what we want. The "**discovery problem**" has long been a challenge in chatbots – how do we know what the bot knows? This is exacerbated by the fact that most chatbots are strictly question in, answer out. The sidebar shows one part of a possible solution to this problem. One interesting aspect of the knowledge graph approach used was that by changing the variables around how quickly (if at all) the bot moved from one topic to another related topic (and how strong or tenuous that relationship was) we could model a variety of personality types and conversational styles.

A second observation was that often a person knows the answer to the question they are asking, but needs it drawn out of them. This is precisely how a lot of coaching works, and the creation of virtual coaches (although from a more motivational and facilitating viewpoint) has already been discussed above. One interesting approach is the creation of reflexive chatbots, where the bots are



asking questions of the user, not the other way around. Something like Edward de Bono's 6 Thinking Hats gives a good model for this (

Figure 3), where the bot can adopt different perspectives on the problem and ask a range of salient questions in order to draw the answer out from the user. Daden developed a prototype bot that worked in this way for UK MOD in 2018, a creating a prompt for ChatGPT to do the same is now almost trivial. One advantage of using an LLM for this is that it already comes with a lot of the general knowledge about most topics, so the questions it asks can be more salient and specific than some of the earlier prototypes we created



*Figure 3: Edward de Bono's 6 Thinking Hats (The Power of the 'Six Thinking Hats' in Coaching, 2024)*

The final issue with regards to information, and of major interest to the military, is that of classification. If dealing with SECRET, or even TOP SECRET or CODEWORD material then you need to know that it will not leak out through a chat with a conversational AI. There is already significant work going into the problem of how to implement LLMs within secure environments – most obviously as standalone (or at least non-Internet connected) systems (although this can of course degrade their usefulness) (Adam & Carter, 2023; Everill, 2024; Gallagher, 2024). A more challenging issue is perhaps that of information aggregation, where individual pieces of information might be UNCLASSIFIED (or at a low classification), but when brought together (and particularly when analysed through the

---

algorithms of an AI) might become classified (or of a higher classification). Again research is being conducted within this area, both in its implications for classified and personally sensitive information (Shao et al., 2024).

## **Doctrine and Concepts**

There is a natural link between peacetime work on Doctrine and Concepts, and the operational execution of that doctrine, so much of the discussion on this element will be part of the operational section below. There are also close links between Doctrine and Concepts and the use of wargaming, which is more fully discussed in Annex A. For now, the following points can be highlighted:

- Conversational AI should help enable more thorough wargaming of concepts and potential new doctrine, both through helping to create wargames (and broader simulation environments) and in instrumenting and analysing those wargames;
- Conversational AI should enable us to better represent the opposition, and ensure that they are front-of-mind in our concept and doctrine development; and
- Through notions such as **platform personas** and **virtual personas** we should be able to draw on (and test against) a broader range of opinion and experience when developing new doctrine.

## **Organisation**

Organisation is an area that could be significantly affected by improvements in corporate knowledge management, particularly through **virtual personas**. How could the military restructure if almost all knowledge was always available wherever it was needed? Could physical headcounts be lower? Could HQs, offices and teams be more dispersed if **virtual personas** could readily stand-in if bandwidth was limited between sites?

---

## Infrastructure

Infrastructure is probably something what is more of an enabler of conversational AI, rather than something which conversational AI can improve. A later section of this paper looks at the DLOD implications of actually providing conversational AI services, and the infrastructure impacts will be considered there.

Bear in mind though that infrastructure is an array of projects, systems and platforms, and as such should be represented by **project personas** and **platform personas** which can hopefully lead to better infrastructure projects, and better and more effective and efficient infrastructure.

## Logistics

Likewise, logistics is ultimately a project and asset/materiel management activity, and so the impact of conversational AI on it should come through the wider use of **project, platform and virtual personas**.

## Integration

Integration is often the implied (or unmentioned) component of the TEPIDOIL framework. Integration is absolutely vital to the working of a conversational AI as it is likely to need to access a wide variety of information and systems in order to be truly effective. One particular challenge is how these systems can be integrated and made interoperable at a semantic level – i.e. so that the terms in one system (e.g. attack helicopter, non-lethal weapon, multi-domain operations etc) correspond to the terms used in another system, enabling more effective data and system fusion.

---

## OPERATIONAL ROLES AND USE CASES

### Arms and Services

The main analysis of the use of conversational AI is focussed on the more generic use cases, at staff level and below. There are also of course a myriad of potential specialist use cases with different Arms and Services such as:

- Specialist engineering advice (particularly for protection and recovery of civilian infrastructure);
- Repair and maintenance advice – particularly for ally and enemy equipment;
- First aid and medical advice;
- Virtual interpreters and cultural advisors.

...and any other area where knowledgeable people are in short supply.

When analysing the operational activities of the military a useful framework is provided by the numbered staff functions:

- J1 - Personnel;
- J2 - Intelligence;
- J3 - Operations;
- J4 - Logistics;
- J5 - Plans;
- J6 - Communications;
- J7 - Training;
- J8 – Resource Management (Finance and Contracts); and
- J9 - Civil-Military Co-operation, including Legal Advisor (LEGAD) and Political Advisor (POLAD).

Within each arm and service there are likely to be Virtual Expert roles (e.g. medical support, combat mentor etc – see sidebar) but these are likely to be variation on themes covered elsewhere.

The opportunities for the deployment of chatbots in each of these areas is considered below. Note that J7 and J8 are not examined here as they are primarily peacetime roles and have been dealt with above. ☐

### The Headquarters Persona

As a starting point, just as we have described above how projects or platforms could benefit from having their own digital persona which aggregates the best practice, knowledge, lessons and experiences of that project/platform and its staff, so to an HQ could benefit from having a **Headquarters Persona** (or even **Formation Persona**), a sort of super staff officer/chief clerk, that can provide continuity and consistency within the HQ. Many of the functions below would then be an integral part of this **HQ Persona**. Note also that the presence of an **HQ Persona** could also open up new options as to how an HQ is deployed and structured, and make more distributed HQs less reliant on communications if each has a “shard” of the **HQ Persona** – which is resync’d across the component HQs as bandwidth allows. It may also be that the **HQ Persona** is not one persona as such, but is

## Human Machine Teaming and Shared Mental Models

Whilst human-machine teaming (HMT) is much discussed, a lot of it is in the context of the “machine” being a relatively dumb tool. The assumption here, though, is that conversational AIs are possible near-equal partners of human teammates.

Studies of elite teams show that their high performance is often due to them having the same shared mental model (SMM) of the task at hand (Richards et al., 2017). We all have our own mental model as to how things, or tasks, work; the trick for a human, or human-AI team to perform well is for all the participants to share the same mental model of what needs to be done, how, and what each individuals responsibilities are. Work on integrating AIs and SMMs is already underway by Daden and other (Lyons et al, 2021).

the aggregate of the persona that support each of the HQ teams (or cells in military parlance) detailed below – a series of **Cell Personas**.

## The Virtual Combat Assistant (VCA)

The **Virtual Combat Assistant (VCA)** is the tactical counterpart of the **HQ Persona**. It is a personal assistant focused on tactical issues that is available to whoever needs it, probably accessed primarily by voice and audio. It has enough local information to answer many questions without the need for vulnerable communications, but, like the **HQ Persona**, takes the opportunity to synchronise and update data whenever bandwidth and the operational situation allows it. There is an ongoing concern about overloading leaders with additional terminals, sensors, UAV feeds etc, and whilst one of the aims of the **VCA** might well be to help manage that data flow on the leader’s behalf, it may still be the case that it is a “systems operator” who actually interfaces directly with the **VCA**, leaving the leader free to command. Again, many of the functions detailed below would, where tactically relevant, be available as part of the **VCA**.

### J1 - Personnel

Personnel management on operations tend to resolve around resourcing and the management of casualties and their replacement. As such generic conversational AI systems as part of the **HQ Persona** providing better access to information, and advice from past operations/exercises and staff could help to enhance this function. The civilian equivalent is of course the HR department.

### J2 - Intelligence

Intelligence is one of the big opportunity areas for conversational AI, and for AI in general. Key areas which could be implemented at both the **HQ Persona** and **VCA** level include:

- Access to the database of enemy equipment and capabilities, and structured help in identifying both enemy equipments and their intents;
- Access to enemy course of action (COA) prediction models, in order to be ready to counter them;

- 
- Personas of key enemy commanders, in order to be better able to judge how they might respond to the situation;
  - Structured AARs after patrols and actions to better capture lessons identified and new intelligence;
  - Helping to better understand (and build) the Intelligence Preparation of the Operating Environment (IPOE), and in particular the human terrain and pattern-of-life elements represented by the PMESII-PT and ASCOPE models;
  - Live interpretation and cultural advice;
  - Fronting a variety of sensor and data fusion systems in order to be able to develop better situational awareness.

Crossing the boundary slightly into J3, conversational AI is also likely to be vital to successful influence and information operations, helping to model the target audience in order to develop successful strategies and interventions, and then enabling the execution of those interventions through a more sophisticated variation on the bot-farm.

### **J3 - Operations**

J3 sits a bit at the intersection of J2 (Intelligence) and J5 (Plans), and many of its requirements overlap with those two functions. J3 is primarily about the command and control of the current operation. Key to this is knowing what your forces are doing, what the enemy is doing, and how you are instructing your forces to respond to variations from the plan. An essential J3 task is that of the watchkeeper – the person who monitors the voice nets and systems to let the staff know what is going on. Watchkeepers work in shifts, and the shift-handover is often the weak point, as some information may not be passed on, and the new watchkeeper might take time to get up to speed. Whilst systems such as Blue Force Tracker might ease some of the task, they do not provide a view on things like intent and morale. There may be significant scope for Conversational AI to help augment the role of the watchkeeper, helping to provide the continuity across that handover gap, and enabling a broader context to be taken of what is going on.

---

It should also be noted that handovers can also be an issue at unit level – particularly when conducting a relief-in-place. Anecdotally every unit thinks that it gives the best possible briefing to the units it hands over to, but conversely every unit also thinks that it always receives the worst possible briefing from the unit it is taking over from. Both can't be true! It's also not just a question of the quality of the handover. In Afghanistan much of the activity was seasonally driven, so a unit arriving in the spring would find a handover from the last unit to arrive in the spring just as useful (possibly even more so) than the one which it is taking over from which may have arrived in the autumn and conducted its operations through the winter. The conversational AI helps to distribute information in time as well as space.

The role of the liaison officer, discussed below under J5 Plans, would also extend into the J3 conduct of operations – and in particular how one unit liaises both with the units on its flanks (or to its front or rear) and to its higher and lower headquarters – ensuring that all are on the same page and working to achieve the higher commander's intent.

#### **J4 - Logistics**

Like J1 Personnel activity, J4 Logistics activities would likely benefit from the broad capability of conversational AI embodied in the HQ Persona to get conversational access to the broad range of systems, databases, spreadsheets and models that support the logistical operation of a unit.

#### **J5 - Plans**

The task of J5 Plans is, as its name suggests, to plan the next operation. There is significant scope here for the use of AI, in many cases fronted by a conversational AI, in order to understand the current situation, evaluate the most likely and most dangerous enemy courses of action, develop, evaluate and select your own courses of action (COA) – including through wargaming (COA wargames or simulation/modelling) – and to then generate the staffwork associated with the orders for the next operation. There is also significant overlap with the J2 Intelligence function and systems.

---

There are two novel uses of **Virtual Personas** which could come into play here.

First, the standard orders process is for the commander to issue their orders to their subordinates in person (the O Group), for those subordinates to then return to their own (remote) units and prepare their own orders for dissemination to the next level down, and so on all the way down through the chain of command. What if each subordinate commander left the O Group with a copy of the Commander's persona, or at least that part of it relevant to the operation and orders (and probably also fragmentary personas of each of the staff cells that produced the orders)? That way, once back at their own HQ if they have any questions or queries about the orders they could ask the persona (an **Orders or Operation Persona?**) about the orders and the intent and rationale behind them. In developing their own orders HQs often develop a list of questions for the higher formation – this new approach could reduce or even eliminate that. Of course there are security issues, and these will be discussed later. Such an **Operation Persona** would persist throughout the operation, informing the G3 Ops cell, and indeed the handover from G5 to G3 could be enhanced through the creation and passing on of the **Operation Persona**.

A second, and related, activity is that of liaison officers. The force for an operation is often drawn from a variety of different sources, and even from different nations, and includes elements from across Air, Land and Sea (and even Cyber and Space). Typically, liaison officers are exchanged to ensure that as well as any formal flow of information between HQs and units there is a more informal channel through which to gauge more qualitative information and to ask questions and to exert influence. Having virtual personas fulfill all or some of this liaison officer role, as **Virtual Liaison Officers (VLO)**, could potentially help reduce HQ headcount and ensure a consistency and 24/7 availability of expert information and advice, but without causing an over-reliance on communications.

## **J6 - Communications**

As with J1 Personnel and J4 Logistics, J6 Communications could benefit from conversational access to support its communications



---

management activities, not least through a series of **System Personas** for each of the key communication systems to help in their planning and operation. The impact of all these virtual personas on the communications bearers itself (which should be positive) is discussed below.

### **J9 - Civil-Military Co-operation**

The potential role of conversational AI in terms of acting as a translator and cultural advisor has been covered under J2 Intelligence, and is just as applicable in supporting Civil Military Co-operation (CIMIC), typically supporting work with the host nation, and also with Displaced Persons (DPs), refugees, and the broader civilian population and civil society. A **Virtual Civil Affairs Officer** would be of great benefit to that infamous strategic corporal, and even the strategic subaltern!

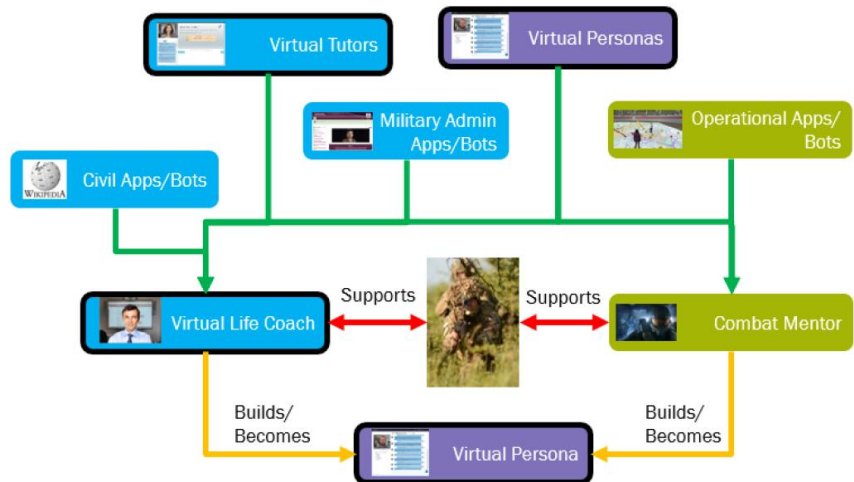
In addition, there are two other roles which are staffed by subject-matter experts, but which conversational AI could help increase the availability of, including providing 24/7 cover at lower level headquarters through **virtual personas** or aggregated **virtual advisors**. These are:

- Legal Advisor (LEGAD) – who provides legal support, particularly regarding the Laws of Armed Conflict, and who could be augmented by a **Virtual Legal Advisor**; and
- Political Advisor (POLAD) – who provides political and diplomatic advice, particularly in connection to collaboration with the host nation, and potentially international partners, and who could be augmented by a **Virtual Political Advisor**.

### **Peacetime and Combat Virtual Assistant**

Figure 4 illustrates how a warfighter could be expected to make use of a **Virtual Personal Assistant/Virtual Life Coach** in peacetime, providing access to a range of virtual tutors, coaches and corporate

systems, and of a **Virtual Combat Assistant/Combat Mentor** during wartime, providing a range of intelligence, logistical and operational support. Interacting with both bots could help to build the **Virtual Persona** of the warfighter without the need for much explicit authoring.



*Figure 4: The Collaboration between a peacetime Virtual Assistant and a Virtual Combat Assistant*

## DIFFERENT TYPES OF MILITARY OPERATION

The analysis so far has been made on the implicit assumption that Large Scale Combat Operations (LSCO) (as in the Russo-Ukraine War) are the focus. In operations such as Services Protected Evacuations (SPE) and Humanitarian and Disaster Relief (HADR) most of the same considerations apply - only the data, knowledge and information changes. However, two other types of operation need a additional consideration: counter-insurgency operations (COIN) and Grey-Zone/Sub-Threshold warfare.

---

## **Counter-Insurgency Operations (COIN)**

In Counter-Insurgency Operations (COIN) there is likely to be more of an emphasis on intelligence, information and influence operations. In these areas conversational AI will not only be providing information and advice but may well be an active part of the delivery, from recruiting and maintaining agents to conducting overt and covert information and influence campaigns. There are however likely to be significantly more ethical issues around such operations, within what is otherwise a “peacetime” environment. Roles such as the LEGAD and POLAD are also likely to gain additional significance, and being able to provide those virtual at almost level - including to the “strategic corporal” (Krulak, 1999) – could make a significant difference to the success of operations.

## **Grey-Zone/Sub-Threshold Warfare**

Grey-Zone (or Sub-Threshold) conflict is ongoing as we speak, particularly as Russia – through the Gerasimov Doctrine (Fridman, 2019) - and China – through its ideas of the Three Warfares (Kania, 2016) and Continuous and Unrestricted Warfare (Saalman, 2021) – seek continually to weaken their opponents and to secure victory (or at least establish the conditions for victory) by all means other than open warfare. As with COIN grey-zone conflict involves significant intelligence, information and influence operations, and their respective counter-operations, and so systems based on conversational AI are likely to play key roles.

## **IMPLICATIONS OF CHATBOT DEPLOYMENT**

Earlier the TEPIDOILI framework was used to identify the opportunities for chatbots within the military. Table 2 uses TEPIDOILI in a more traditional way to look at the implications for the Defence Lines of Development (DLOD) of introducing the extensive use of chatbots into the military estate.

<b>Training</b>	Since conversational AI is predicated on just having a conversation with the system then the training overhead should be minimal. However, some military staff also need to be trained on how to build, or at least customize and enhance, such systems.
<b>Equipment</b>	Again, since the interface to these systems is based around voice/audio, possibly augmented by text-chat or even video, then most existing military terminals should be able to support them. However, if it is required to run the code locally, and host the data and even LLM models locally, then the processing power (and electric power) requirements for local devices and servers may increase. There may also be a requirement for more powerful servers on defence networks in order to run (or at least to build) the LLMs.
<b>Personnel</b>	Apart from the training requirement, the main impact on personnel is potentially that of making existing (and even past) expertise go further, through its use in <b>virtual personas</b> and incorporation into <b>project, platform and system personas</b> . Conversational AI may offer scope to reduce back-office headcounts, and even to reduce (or at least distribute) deployed HQ headcounts, which should increase survivability.
<b>Information</b>	As mentioned above, conversational AI in the context considered here is largely about improving knowledge, and hence information, management. Conversational AI may increase the amount of data available (e.g. through AARs and capture by virtual assistants), but it balances that by making existing and new information more available, and ideally at point and time of need.
<b>Doctrine</b>	Whilst the deployment and use of conversational AI does not require new doctrine <i>per se</i> , the potential impact of conversational AI on doctrine has been discussed above.
<b>Organisation</b>	Likewise, conversational AI should not require any changes to organisation, but it may enable changes to organisation due to the deployment of <b>personas</b> .
<b>Infrastructure</b>	As mentioned above the widespread use of LLM-based conversational AI could require the deployment of more powerful servers, particularly to build new models. The bandwidth implications are more moderate, as most audio or video presentation is likely to be done at the terminal, although some audio and video understanding/decoding may be done at the server.

<b>Logistics</b>	Conversational AI is likely to be accessed through existing terminal equipment, so should not have a significant effect on logistics. The only potential issue might be if it increases the power consumption of terminal devices – increasing demands on batteries and local generation.
<b>Integration</b>	As discussed above, conversational AI systems are likely to be avid consumers of interfaces into other systems, and require a high degree of integration in order to be effective.

*Table 2: DLOD Implications of Deploying Conversational AI*

## **ADDITIONAL CONSIDERATIONS**

This final section looks at some additional consideration of the military use of chatbots which have not been adequately addressed so far. These are:

- Security;
- Validity and Explainability;
- Trust; and
- Ethics.

### **Security**

If the military is providing conversational AI access to multiple systems, then by easing the ability for its own warfighters to access information it is also potentially making that same information more accessible to the opposition through cyber-attacks and other means. Any conversational AI system needs security protection (including encryption) appropriate to the classification of the information (and systems) it can access. Issues around classification have already been discussed above. Any **Virtual Combat Assistant** must be capable of remote kill (or even autonomous “suicide”) if lost or captured.

There are also issues around the security of the repositories of information that the AI is accessing. Whilst an obvious risk is of that information being accessed and exfiltrated, a more damaging scenario may be where that information is “poisoned” so as to give wrong answer at crucial times in the future. It should go without saying that any conversational AI system implements similar levels

---

of sign-on security, authentication and permission levels as any other military IT system.

### **Validity and Explainability**


Whilst hostile poisoning of data is a major concern, a more immediate issue is that of the validity of the information that the conversational AI presents without any hostile interference! We have all read the descriptions of how current LLMs “hallucinate” and just make facts up. Most military applications will require zero-hallucination systems. Far better a bot (or even human staff officer) admit they don’t know something (and potentially give pointers to someone/something that does) than they make answers up.

People often want to know why something is being said, not just what is being said. Creating “Explainable AI” has been a significant issue for many years (Dwivedi et al., 2023), and in some ways LLM based chatbots are a backward step, as any reasoning is lost within the neural nets and word vector embeddings of their underlying algorithms. Alternative approaches such as semantic knowledge graphs (Hogan et al., 2022) have significantly better explainability, and some hybrid system might ultimately be the way forward.

### **Trust**

If a conversational AI is going to be a true part of a team it needs to be trusted. Giving 100% factual answers, and some rationale for those answers, goes some of the way to building trust – but how do you know those answers are 100% true every time? An earlier sidebar discussed how approaches such as shared mental models not only help us understand human teams better, but can also help to potentially build better human-AI teams. Ultimately trust really only comes from people (or people and AIs) working together – this means in training, on exercise and on operations. Trust, as they say, is earned, not given, or programmed. This also implies that the personas that we have talked about need longevity and consistency of identity and even personality, if that trust is going to be built up over time.

## The Military Metaverse



All the developments and opportunities in conversational AI that have been discussed in the white paper are happening in parallel with the steady emergence of a military metaverse, as the military (and society) makes more use of Virtual Reality and Mixed Reality. Within such environments the avatars of conversational AIs need not look (or behave) differently to those of humans.

For a full exploration of the Military Metaverse see the forthcoming (May 2025) book on the topic by Andy Fawkes and this paper's author, David Burden. [Pre-order](#) from Taylor & Francis, and [more information here](#).

## Ethics

Ethical issues in the use and deployment of conversational AI systems have been flagged a number of times in this white paper, whether it is around the earnings potential or post-death management of **virtual personas**, or the use of conversational AI in covert influence operations. And this is quite apart from any more conventional ethical issues with conversational AI – such as anonymous intimacy – the tendency to reveal information to a chatbot (Pan & Hamilton, 2018) or the emotional response of users to friendly and human-like avatars (Fraser et al., 2024). Any consideration of the deployment of conversational AI systems should include an ethical review appropriate to the potential risks and impact of the proposed system – and that review process should run concurrent with, and ideally ahead of, the development process in order to achieve a system which is ethical by design (Chivukula et al., 2024; Ramos et al., 2024).

«

## CONCLUSIONS

This White Paper can only provide a high-level introduction to some of the potential roles of conversational AI within the military, and to some of the potential issues also need to be addressed. The biggest benefit of a conversational AI system is that it should be like talking to a physical person. It doesn't need extra training and we intuitively know how to manage the conversation. A good chatbot should be like the best possible assistant, subject matter expert or staff officer. Not only should it know the answers but it should also know how to give you the right information, at the right time, and in the right way. Discovering what the chatbot knows should not be an issue, it should both provide conversational modes that let you uncover the information as you would when talking to a person, and it should learn to second guess your requirements.

Just imagine the potential. Just imagine if every project had a persistent project persona that brought all the information, decisions, intents and experiences of the project team over *decades* together. Just imagine if you could talk to everyone who's held the

---

post that you now hold. Just imagine if every ship could remember every operation it had ever been on (and those of its sister ships), and how its crew dealt with every issue.

### More Information

For more information on the topic of conversational AI and their application to the military please contact the author, David Burden, at: david@daden.co.uk

You may also like to visit:

- The Daden website at <http://www.daden.co.uk>
- The Metaverse Series website at: <http://www.themetaverseseries.info/>

Or read David's books for Taylor & Francis on *Virtual Humans*, *The Metaverse: A Critical Introduction* and *The Military Metaverse*, all available through the above websites.

Whilst some of the concepts discussed here might be a few years away, for many of them the technology to implement them is available now. Initial systems may not be perfect, but by experimenting now we'll be in a far better place the shape the systems of the future, and to identify and deal with the challenges before they become too significant, and to ensure that our militaries, and the people within them, have the best possible AI support to do their job.

«

---

### 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 is also currently undertaking a PhD on wargaming urban conflict. David is an ex-Royal Signals officer, a Chartered European Engineer and is also 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*.

Declaration: David is also on the Advisory Board of Sensay, a digital replica company (see p.7), as well as running Daden Limited.



---

## REFERENCES

- Adam, C., & Carter, R. (2023). Large Language Models and Intelligence Analysis. *CETaS Expert Analysis*.  
[https://cetas.turing.ac.uk/sites/default/files/2023-07/cetas\\_expert\\_analysis\\_-\\_large\\_language\\_models\\_and\\_intelligence\\_analysis.pdf](https://cetas.turing.ac.uk/sites/default/files/2023-07/cetas_expert_analysis_-_large_language_models_and_intelligence_analysis.pdf)
- Burden, D. J. H. (2019). *Developing Conversational AI solutions for the MOD | CogX2019*.  
<https://www.youtube.com/watch?v=SQj8kTpN7lw>
- Burden, D. J. H. (2020). Building a Digital Immortal. In M. Savin-Baden & V. Mason-Robbie (Eds.), *Digital Afterlife*. CRC.
- Burden, D. J. H. (2024, September 5). *Matrix Game Simulations—Part 1*. David Burden's Substack.  
<https://substack.com/@urbanwargamer/p-148217254>
- Burden, D. J. H., & Savin-Baden, M. (2019). *Virtual Humans: Today and Tomorrow*. CRC.
- Chivukula, S. S., Gray, C., Li, Z., Pivonka, A. C., & Chen, J. (2024). Surveying a Landscape of Ethics-Focused Design Methods. *ACM Journal on Responsible Computing*, 1(3), 1–32.  
<https://doi.org/10.1145/3678988>
- Curry, J., Engle, C., & Perla, P. (2011). *The Matrix Games Handbook*. History of Wargaming Project.
- Dwivedi, R., Dave, D., Naik, H., Singhal, S., Omer, R., Patel, P., Qian, B., Wen, Z., Shah, T., Morgan, G., & Ranjan, R. (2023). Explainable AI (XAI): Core Ideas, Techniques, and Solutions. *ACM Computing Surveys*, 55(9), 1–33. <https://doi.org/10.1145/3561048>
- Eaton, J., & Olaru, S. (2021). *Does Artificial Intelligence Conduct Better Research Interviews than You?* NATO S&T.  
<https://www.sto.nato.int/publications/STO%20Meeting%20Proceedings/STO-MP-SAS-OCS-ORA-2021/MP-SAS-OCS-ORA-2021-AIML-02-3.pdf>
- Everill, T. (2024, May 16). *Microsoft Builds Offline LLM for US Intelligence Services—Bloomsbury Intelligence and Security Institute (BISI)*. Bloomsbury Intelligence & Security Institute (BISI).  
<https://bisi.org.uk/reports/microsoft-builds-offline-llm-for-us-intelligence-services>
- Fraser, A. D., Branson, I., Hollett, R. C., Speelman, C. P., & Rogers, S. L. (2024). Do realistic avatars make virtual reality better?

- 
- Examining human-like avatars for VR social interactions. *Computers in Human Behavior: Artificial Humans*, 2(2), 100082.
- Fridman, O. (2019). On the Gerasimov Doctrine. *Prism*, 8(2), 100–113.
- Gallagher, S. (2024, March 5). *Using Large Language Models in the National Security Realm*. Carnegie Mellon University. <https://insights.sei.cmu.edu/library/using-large-language-models-in-the-national-security-realm/>
- Geist, E., Frank, A. B., & Menthe, L. (2024). *Understanding the Limits of Artificial Intelligence for Warfighters: Volume 4, Wargames*. RAND Corporation. [https://www.rand.org/pubs/research\\_reports/RRA1722-4.html](https://www.rand.org/pubs/research_reports/RRA1722-4.html)
- Griffin, L. D., & Riggs, N. (2024). *Transcript of GPT-4 playing a rogue AGI in a Matrix Game* (No. arXiv:2405.10997). arXiv. <https://doi.org/10.48550/arXiv.2405.10997>
- Griffin, L. D., & Zhu, B. (2024). *Matrix Gaming with Large Language Models*. Connections UK, Brunel University, London. <https://www.professionalwargaming.co.uk/24MatrixGamingWithLLMs.pdf>
- Hogan, A., Blomqvist, E., Cochez, M., D'amato, C., Melo, G. D., Gutierrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., Ngomo, A.-C. N., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., & Zimmermann, A. (2022). Knowledge Graphs. *ACM Computing Surveys*, 54(4), 1–37. <https://doi.org/10.1145/3447772>
- Kania, E. (2016). The PLA's latest strategic thinking on the three warfares. *China Brief*, 16(13), 10–13.
- Knack, A., & Powell, R. (2023). *Artificial Intelligence in Wargaming*. CETaS, Alan Turing Institute.
- Krulak, C. C. (1999). The strategic corporal: Leadership in the three block war. *Marine Corps Gazette*, 83, 18–22.
- Lyons, J. B., Sycara, K., Lewis, M., & Capiola, A. (2021). Human–autonomy teaming: Definitions, debates, and directions. *Frontiers in Psychology*, 12, 589585.
- Pan, X., & Hamilton, A. F. D. C. (2018). Why and how to use virtual reality to study human social interaction: The challenges of

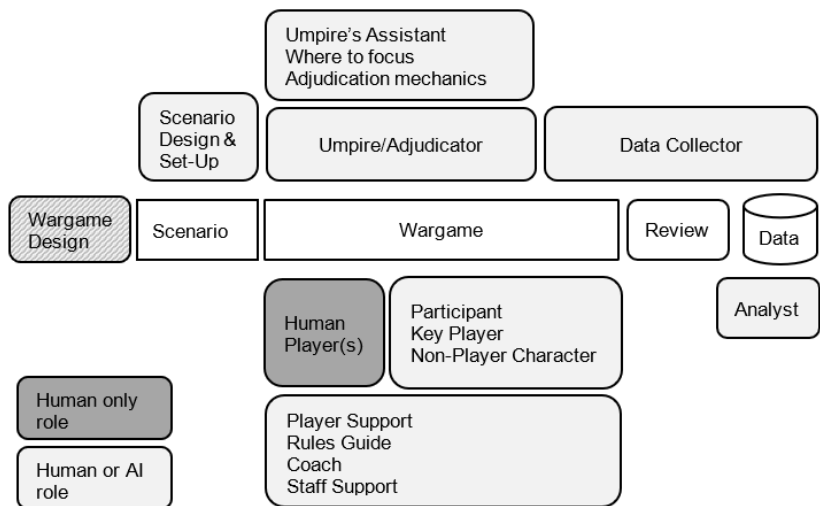
- 
- exploring a new research landscape. *British Journal of Psychology*, 109(3), 395–417.
- Probst, G. J., & Romhardt, K. (1997). *Building blocks of knowledge management-a practical approach*.  
<https://ideas.repec.org/p/fth/ehecge/97.12.html>
- Ramos, G., Squicciarini, M., & Lamm, E. (2024). Making AI ethical by design: The UNESCO perspective. *Computer*, 57(2), 33–43.
- Richards, P., Collins, D., & Mascarenhas, D. R. D. (2017). Developing team decision-making: A holistic framework integrating both on-field and off-field pedagogical coaching processes. *Sports Coaching Review*, 6(1), 57–75.  
<https://doi.org/10.1080/21640629.2016.1200819>
- Saalman, L. (2021). China and Its Hybrid Warfare Spectrum. *Hybrid Warfare: Security and Asymmetric Conflict in International Relations*, 95–112.
- Shao, H., Huang, J., Zheng, S., & Chang, K. C.-C. (2024). *Quantifying Association Capabilities of Large Language Models and Its Implications on Privacy Leakage* (No. arXiv:2305.12707). arXiv.  
<https://doi.org/10.48550/arXiv.2305.12707>
- Strong, P. E., Strong, N., & Maltby, J. (2018). It takes three hundred years to build a new tradition. *The Naval Engineer*, 6(1).
- The Power of the 'Six Thinking Hats' in Coaching*. (2024, October 18). My Coaching Toolkit. <https://mycoachingtoolkit.com/the-power-of-the-six-thinking-hats-in-coaching/>
- Wilson, D., & Sperber, D. (2022). On Grice's theory of conversation. In *Conversation and discourse* (pp. 155–178). Routledge.  
<https://www.taylorfrancis.com/chapters/edit/10.4324/9781003291039-11/grice-theory-conversation-deirdre-wilson-dan-sperber>

## **BIBLIOGRAPHY**

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

**WARGAMING AND CONVERSATIONAL AI**

There have been a number of studies recently which have looked at the potential role for AI within a wargaming context, including those by RAND (Geist et al., 2024) and the Centre for Emerging Technology and Security (CeTAS) at the Turing Institute (Knack & Powell, 2023). Figure 5 (Burden, 2025) is typical, showing the different components of a professional wargame.



*Figure 5: The roles for AI within a wargame (after Burden, 2025)*

AI has a potential role to play in almost all of the elements of Figure 5, and these are summarised in Table 3.

<b>Role</b>	<b>Conversational AI Contribution</b>
Wargame Design	Probably the hardest area for AI to contribute to, but a conversational AI as a research and design assistant, and to help generate textual and media content, and even draft rules, is already viable. The AI could also help in extracting wargame requirements from the sponsor, and compare them to its knowledge of existing wargames and wargaming techniques.
Scenario Design	Given a set of rules and a scenario template, a conversational AI can readily generate alternative scenarios, and potentially the media assets required to go with them.
Scenario Setup	Setting up a game can often take time, and it's easy to miss things. A chatbot can help guide the players through this, and answer any questions on setting up the game.
Player	Conversational AIs are increasingly able to digest a set of rules and to become players in a game. For a manual game they can use a human as their "agent" to describe the game state and options available, and to take the physical actions to make a move. Of course, a wargame devoid of any human players is a simulation not a wargame, but an AI player can provide an alternative to solo play, or fill in if players are missing, or specialize in playing the adversary and reflect the adversaries tactics and doctrine.
NPC	Chatbots are more than capable of taking on the role of non-player characters – such as allies or intelligence sources, and a lot of earlier conversational AI work was driven by creating more believable NPCs in sandbox computer games.
Player Support	A chatbot can help a player understand the rules, answer questions on them, and advise on game strategy and tactics.
Umpire/ Adjudicator/ Facilitator	Another major challenge area for chatbots, but LLMs are already being used to umpire conversational based wargames such as matrix games, and could increasingly be used to help facilitate or adjudicate where an umpire familiar with a specific game is unavailable.
Umpire Assistant	Umpires have a lot to contend with in a game, particularly if covering umpire, facilitation and adjudication functions. Chatbots could support an umpire across all of these, helping with rules questions and executions, and potentially being able to take a step back and help with teaming and focus issues.

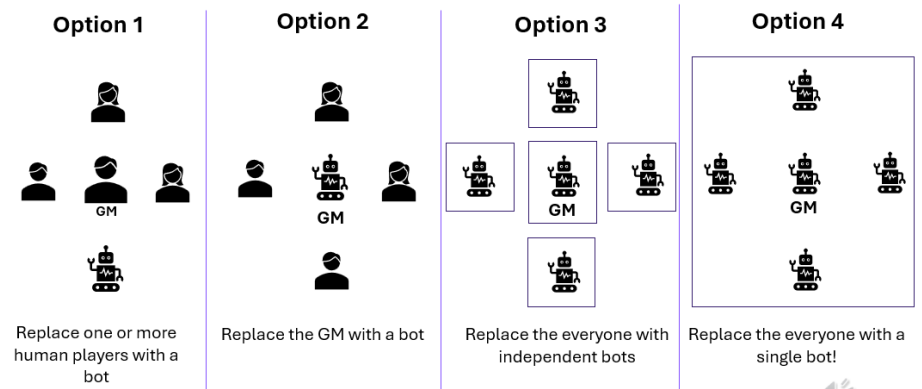
<b>Role</b>	<b>Conversational AI Contribution</b>
Data Collector	Professional wargames typically have a well-developed data collection plan, which may include scribes/observers and even question probes. Chatbots can readily augment or provide much of this support, enabling more players to be interviewed in depth about their experiences with the game.
Analyst	Whilst non-conversational AI probably takes the lead in analysing the wargame output, much of the output is on the qualitative side which will require natural language understanding, and conversational AI can act as an analysts assistant in managing the workflow and helping to link findings to the literature.

*Table 3: Wargaming Roles of Conversational AI*

## **Matrix Games**

Matrix Games are a particularly interesting use case of conversational AI and wargaming as they are so simple to set up. A Matrix Game is essentially a structured argument or discussion around an incident or crisis (Curry et al., 2011). Each player, representing a particular country, faction or unit), states what action they are going to take this turn, and gives a couple of arguments as to why it should work. Other players can propose counter-arguments, and the Umpire decides if the action succeeds and describes the outcome.

There are at least 4 approaches to using conversational AI to play matrix games, as shown graphically in Figure 6.



**Figure 6: Conversational AI Models for Matrix Games**

Table 4 summarises each approach, and links to relevant references.

1. One or more AI Players	Known, or unknown, to humans one or more actors in the game can be played by an AI. For simplicity it may be easiest to “front” the AI with a human who updates the AI on the game play, and then presents the AI’s actions and responses. (Griffin & Riggs, 2024)
2. AI Umpire	The umpire is an AI. This is possibly the most challenging option as a lot of the umpire role is about managing chaotic human behaviour using a variety of verbal and non-verbal social skills.
3. Separate AI Players and AI Umpire	Here each actor and the umpire are played by separate (or at least segregated) AIs, which ensures privacy in each actors intentions, and also allows for “secret arguments” and covert actions agreed between an actor and the umpire. (Griffin & Zhu, 2024)
4. Integrated AI Players and Umpire	A single AI is asked to play all the roles. This is very quick to set up, but there may be some leakage between the different roles. (Burden, 2024)

**Table 4: Conversational AI Approaches to Matrix Games**

---

The big advantage of Option 4 is that more or less anyone can write a matrix game and see it played out within as little as an hour of first having the idea – and rapid and cheap way of exploring ideas to see if they have any potential merit or interesting facets that would be studied in a more developed human matrix, or other, wargame.

You can read more about Option 4 approaches in the authors own posts at: <https://substack.com/@urbanwargamer/posts>