

Episodic memory in AI agents poses risks that should be studied and mitigated

Chad DeChant
Computer Science Department
Columbia University
New York, NY, 10027, USA
chad.dechant@columbia.edu

Abstract—Most current AI models have little ability to store and later retrieve a record or representation of what they do. In human cognition, episodic memories play an important role in both recall of the past as well as planning for the future. The ability to form and use episodic memories would similarly enable a broad range of improved capabilities in an AI agent that interacts with and takes actions in the world. Researchers have begun directing more attention to developing memory abilities in AI models. It is therefore likely that models with such capability will become widespread in the near future. This could in some ways contribute to making such AI agents safer by enabling users to better monitor, understand, and control their actions. However, as a new capability with wide applications, we argue that it will also introduce significant new risks that researchers should begin to study and address. We outline these risks and benefits and propose four principles to guide the development of episodic memory capabilities so that these will enhance, rather than undermine, the effort to keep AI safe and trustworthy.

Index Terms—safety, trustworthy AI, episodic memory

I. INTRODUCTION

Among the most significant ways in which current AI models are unlike human cognition is their lack of comparable memory abilities. Very few make any attempt to develop and use an important type of memory on which humans depend, *episodic memory*. Episodic memory is memory of particular past events which one participated in personally and can, in some way, recall [1]. For AI models, this would mean the ability to form and retrieve memories of events — not merely newly acquired facts — that happen post-deployment, at runtime. Most models that incorporate such memory do so in relatively simple ways which are poor approximations of human episodic memory and do not scale to longer, more realistic lengths of time. However, this is beginning to change as more work is done in the area [2], [3]. Making use of episodic memories would enable significant new capabilities and can therefore be expected to be a burgeoning area of research interest in the coming years.

There is simultaneously a growing interest in developing AI *agents*, models that are trained and equipped to take actions that affect the world [4], [5]. Whether these are robots in the real world or virtual agents, they will, by design, be able to impact the world in a more direct way than previous AI models.

While episodic memory plays an important part in many of the cognitive processes which contribute to human intelligence

[6], [7], it has not played a large part in the development of AI agents. This is understandable: until recently, most such agents were limited in the range of actions they could perform and the time horizon over which they could act. They therefore had less need for episodic memories than agents that are the current focus of development.

When AI agents are able to make full use of rich episodic memory abilities, there will be significant implications for their safe deployment. Episodic memories may come to play a role analogous to that which they play in humans, facilitating a wide range of capabilities. These include better planning, problem solving, decision making, and learning [8]. Such capabilities would make any agent equipped with them harder to understand and, in some ways, control.

There is currently an opportunity to prevent episodic memory abilities from making AI agents more dangerous. The implementation and deployment of these abilities are still at an early stage, allowing the research community to study the problem. Possible dangers and benefits of episodic memory can be examined. Most importantly, the results of these studies can be used to guide the implementation of artificial episodic memory to make it safer. There will be a wide variety of ways to implement episodic memory abilities. If the safety implications of these various approaches are understood in advance, research can be directed toward safe techniques and away from more dangerous ones.

We seek in this paper to draw attention to the risks and benefits of episodic memory in AI agents and motivate a program of research into ways to implement it safely. We begin by summarizing key points of what is known about human episodic memory, including how it differs from other forms of memory more familiar to the artificial intelligence community. In order to show the possible impacts of bringing episodic memory abilities to AI agents, we highlight some of the ways in which humans are thought to use episodic memories, paying particular attention to the many ways episodic memories are used in human cognition beyond simply recalling the past. We also consider whether it is possible for AI to have episodic memories in the way in which humans or, possibly, animals do.

We then outline and explain risks posed by episodic memory abilities in artificial intelligence: deception; retention of knowledge; improved situational awareness; and the unpredictability

of memory. We follow this with a discussion of ways in which memories could be used to make AI safer and more trustworthy by facilitating more thorough monitoring, control, and explanations of their actions. We also explain how, counterintuitively, AI agents’ own episodic memories may be more easily kept from them than other forms of information, making such memories a crucial element of any effective strategy to control such agents.

In response to the risks we outline, we propose four principles for implementing episodic memory capabilities in a way that promotes, rather than undermines, the safety and reliability of systems with such abilities:

- 1) Memories should be interpretable by users
- 2) Users should be able to add or delete memories
- 3) Memories should be in a format which enables users to isolate and detach them from the rest of the system they are part of
- 4) Memories should not be editable by AI agents

Finally, we propose some open research areas and questions to encourage further research in this area.

II. RELATED WORK

A. AI agents

AI agents have been the subject of study for many decades [9]. The most widely used artificial intelligence textbook defines an agent as “anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators” [10], where an agent may be a robot operating in the real world or a purely software-based agent operating in a virtual or internet-based environment. AI agents may also be referred to as *autonomous* AI, where *autonomous* is meant to convey that such systems can “plan, act in the world, and pursue goals” [11].

There are many examples of AI agents in a wide variety of contexts. Robots operating autonomously in the real world are perhaps the prototypical example of such agents as it is easy to see both their independence and tangible effects of their actions. It has been proposed that AI agents operating as biomedical “AI scientists” could develop hypotheses, test them in the real world, and have a form of memory to store experimental results [12].

In recent years there has been an increasing amount of work on AI agents which have a large language model as a component. In these works, LLMs are often used to help an agent plan its actions [4], [13], [14], despite evidence that they may not be capable of reliably planning [15]. Games have provided environments for the training of agents of various kinds, including an LLM-based agent with a memory capacity for newly acquired skills [16].

B. Memory in AI

Techniques patterned after or inspired by episodic memory have been explored in the machine learning literature. This has included work on the efficiency of reinforcement learning through episodic replay [17]; planning [18]; improving world models [19]; remembering the values of states or actions [20],

[21]; and more complex memory structures designed to solve tasks which require episodic memories [22], [23]. Robotics researchers have developed techniques to store and recall information about robots’ past actions for use in summarization, question answering, and planning [24]–[27].

Architectures for incorporating different kinds of memory-like functions have led in the past to meaningful improvements in capabilities. These included Long Short Term Memory modules [28], Neural Turing Machines [29], Differentiable Neural Computers [30], Hopfield networks [31], and Modern Hopfield Networks [32].

The increasing length of context windows in large language models may raise the question of whether forming representations of episodic memories will ultimately be necessary. Instead, it might be thought that it could eventually be possible to give such a model an agent’s entire history as input in a relatively raw format. However, this is unlikely for several reasons. First, it would be very inefficient to reprocess an entire history at every time step an agent acts. Second, the sheer length of time that agents will eventually operate (e.g. decades) would almost certainly be too large for even future long-context models. Third, it possible that such very long context windows will continue to lead to degraded performance, as has been seen in current models [33].

The potential for various kinds of memory has received particular attention in the natural language processing community [34]. The need to circumvent a fixed length for input to large language models has inspired many ways of compressing information and storing it for later use by models, including in retrieval augmented generation [35]–[40]. Moving closer to an agentic framework, recent works have given models greater control over the retrieval and use of information [41], [42].

The existing work which comes closest to our conceptualization of the role of memory introduces a virtual environment of interacting LLM-based agents that record and later consult natural language records of their actions, using these to better understand their environment and make coherent and relevant plans [3]. Partially instantiating our speculation about the utility of past episodes for planning, recent works have developed systems to store and retrieve episodes of action to help guide future decision making in a reinforcement learning context [43], [44].

Though a comprehensive overview of different types of memory employed by LLM-based AI agents is beyond the scope of this work, a recent paper provides just such a thorough survey [2].

C. Safe and trustworthy AI

Concerns about the risks posed by artificial intelligence extend back to its earliest days [45]. Research on various kinds of harms that artificial intelligence might cause is now the subject of several research communities with a variety of interests [46]–[49]. We focus here on work that most closely concerns risks which could be amplified or reduced by episodic memory.

Given the potential of agentic AI to take actions in and affect the world, special attention to the possible harms of AI agents is warranted. There is, therefore, a growing body of research on the possible dangers of such systems and techniques for ensuring they operate safely and understandably [50]. Attention has been given to developing ways to make robots’ actions explainable [51], for example in the case of autonomous vehicles [52]. LLM-based agents, including those operating on the internet [53], pose several ethical and safety challenges which are actively being studied. Agents’ use of tools is the focus of a framework that pairs emulation and evaluation of LLM-based agents [54]. A recent paper surveying the risks potentially posed by LLMs includes a section on the particular risks of LLM-based agents [55].

Researchers have given particular consideration to the risks particular to agentic AI that can act autonomously in pursuit of goals its users do not intend, either because of a malicious actor or by accident [11]. It has been proposed that due to reward misspecification during training, an LLM-based agent trained using deep learning and reinforcement learning from human feedback might develop undesirable internal goals, particularly if such agents can determine if they are currently being trained, evaluated, or deployed and operate over long time horizons [56]. AI agents have been found by many researchers to engage in various kinds of deception [57].

Episodic memories could be used to invade a user’s privacy. Ways in which machine learning can violate users’ expectations of privacy have been extensively studied [58], along with approaches to mitigate these invasions of privacy [59]. Of particular relevance is work which examines the potential for mass surveillance which modern machine learning techniques afford and which widespread agents with episodic memories might contribute to [60]. Dialog agents might have a range of negative effects on their users, including invasions of privacy [61].

III. EPISODIC MEMORY

In this section we present a brief overview of episodic memory in humans. Our goal is to provide a short introduction to the topic and to highlight ways in which implementing episodic memory could be useful in creating more capable artificial intelligence. We pay particular attention to the links between episodic memory and other abilities, at least in humans, as these are potentially surprising and are especially significant for considering the effects of memory abilities on AI agents.

A. Taxonomy of memory types

Episodic memory in humans is memory for events in which someone personally participated. The psychologist Endel Tulving is recognized as being the first to propose a distinction between episodic memory and *semantic* memory, which is memory of facts about events and the world [62]. For example, someone remembering a trip to Paris that they took a few years earlier would be using their episodic memory, while someone remembering that Paris is the capital of France would

be using semantic memory. Although semantic memory is typically about impersonal information such as geographical knowledge, it might also be about factual information about oneself which does not call to mind a particular associated episode. For example, remembering which city one was born in would be considered an example of personal semantic memory.

Both episodic and semantic memory are referred to as types of declarative memory [63]. Their contents can (to some extent) typically be described using language. A third kind of memory, *procedural* memory, is sometimes included in taxonomies of memory types and is considered nondeclarative. Procedural memory is memory of how to do something, such as riding a bicycle or other skill or ability which had to be learned [64].

Much — indeed, perhaps most — machine learning research involves either what we have just described as semantic or procedural memory, though these are not typically described as forms of memory in a machine learning context. Large language models are valued in large part for their semantic memory of facts about the world. This kind of memory has, for example, recently been investigated in the many papers asking what LLMs “know” [65], [66]. Procedural memory for learned skills and abilities is the objective of much machine learning work such as that on learning navigation, game playing, automobile driving, robotic manipulation, etc.

B. The stages of episodic memory

While there are many theories and debates about the way in which episodic memories are formed and maintained in the brain [67], in broad outline the process consists of the following stages:

- 1) *Encoding*: Raw sensory and other (e.g. emotional) information about an episode needs to be compressed and structured into a suitable representation. Forming representations of episodic memories critically depends on the hippocampus and nearby structures in the medial temporal lobe. Damage to the hippocampus is known to impair the ability to store new episodic but not semantic memories [68].
Episodic memories do not depend only on the hippocampus, however. According to the hippocampal indexing theory [69], an encoding of an episode in the hippocampus serves as a kind of index that points to and binds together representations in the neocortex that form the basis of the episode, such as multimodal sensory representations along with associated emotional and conceptual information [70].
- 2) *Storage*: After a memory is encoded, it undergoes a period of so-called *consolidation* or transformation into a long term form [71]. This is commonly thought to involve moving the representation from the hippocampus to the neocortex, either partly [72] or entirely [63]. The memory must then *persist* in its stored form [73]. There is evidence that a memory is in some way

destabilized when it is retrieved or activated, leading it to undergo *reconsolidation* [74].

- 3) **Retrieval:** Retrieval is the recovery of a previously encoded episodic memory [75]. A *cue*, which could be an activation of part of the pattern that was stored in memory, triggers the associated hippocampal index representation. This index is then thought to activate the associated neurons in the neocortex which made up the original encoding of the episode [76].

From this necessarily quite abbreviated account, we may draw a few conclusions relevant to the consideration of episodic memory in artificial intelligence. First, episodic memories depend on many brain regions for their creation and persistence. Episodic memories are not located in just one area that is easy to study and manipulate. Second, humans' episodic memories are not simply left untouched after their creation; they are neither static nor unchanging. Third, there would be many points of possible intervention in an artificial process that mimics the above stages. Such forms of intervention could be used to enact the recommendations for safe episodic memory proposed later in this paper.

C. The uses of episodic memory

Episodic memories are thought to be involved in a variety of important cognitive processes beyond simply recalling past events. Evidence for this comes from two primary sources. First, brain imaging studies have shown that similar brain regions are recruited during recall of episodic memories and other tasks [77]. Second, people with impairments in episodic memory abilities are found to also have deficiencies when performing other tasks [78]. These two kinds of evidence, along with theoretical accounts that seek to explain the observed relationships, suggest that episodic memories — or at least the ability to form episodic memories and its associated cognitive architecture — are used when performing many other important cognitive functions.

Episodic memories are, as memories, naturally of events in the past. They are used, however, to influence the future. Indeed, some have argued on evolutionary grounds that memory should be considered to be primarily concerned with the future, helping us act in whatever new circumstances the future presents us [79].

Planning Especially relevant to the concerns of this paper is the way memories are used when planning future actions: according to some theories, memories serve as “building blocks”, allowing elements of particular episodes to be reused and reassembled in different ways in order to respond to novel situations [80]. Some psychologists have gone so far as to suggest that “episodic reconstruction is just an adaptive feature of the future planning system” [81].

Imagination and prediction Accumulating evidence shows that episodic memory — and the brain systems that support it — is involved in predicting and imagining the future. Patients with damage to hippocampal and non-hippocampal regions involved in memory have unusually poor performance when asked to predict or imagine future scenarios. Such

patients imagine impoverished scenarios lacking in detail and coherence [82], [83].

Problem solving As long ago as the nineteenth century it was observed that patients with amnesia lacked the ability to engage in flexible thinking, with one doctor observing of his amnesiac patients that the “circle of ideas in which the patient’s intelligence moves becomes very restricted” [82], [84]. An association has been found between having deficits in episodic memory functioning and being unable to generate relevant details in an open-ended problem solving task [85].

Decision making One proposed psychological model demonstrates how episodic memories can help in learning a new task by allowing successful episodes to be recalled and emulated [86]. This approach was extended to show how similar memories could be sampled in order to estimate the value of possible actions [87].

Learning from episodic memories Episodic memory has been described as “epistemically generative” in the sense that it enables learning from past experiences [88]. Past events may, for example, be recalled and reinterpreted in light of newly acquired information, allowing one to learn from remembered aspects of the past events which had previously been misunderstood.

D. Can an AI agent have episodic memories?

The wide variety of ways episodic memories are used by humans suggests that the incorporation of true episodic memory abilities into AI agents would greatly expand their range of capabilities. But some may question whether it is even possible for AI agents to have actual episodic memories. Indeed, Endel Tulving himself described the phenomenon of recalling episodic memories in terms which cast doubt on the very idea. He wrote that remembering a past episode is a kind of “mental time travel,” a “conscious awareness of what had happened in the past” which has an “experiential flavor” [89].

Some psychologists and philosophers have concluded that episodic memory is therefore a uniquely human phenomenon, lacking even in non-human animals — much less AI agents. According to this view, animals are “stuck in time,” without either episodic memory or “the ability to anticipate long-range future events” which we have seen is associated with episodic memory [90].

Others, however, have taken a more expansive view of what constitutes episodic memory as well as who has it. Rather than focusing on phenomenological aspects of memory, it is possible instead to consider episodic memory as simply combining memory for *what* happened, *when* it happened, and *where* it happened [91]. While we do not know what non-human animals experience, we can study their behavior. Some studies have indeed found evidence for episodic memory in animals. For example, an experiment demonstrated that birds (scrub jays) reliably behaved as if they remembered what kind of food they had hidden, where they had hidden it, and when they had hidden it [92]. Another study on scrub jays demonstrated that they also appear to plan for the future when hiding food [93]. Others argue that episodic memory has a

long evolutionary history that predates humans and point to brain areas in other mammals as well as birds which they claim correspond to areas in the human brain responsible for episodic memory [94].

If non-human animals do have episodic memory, that would reinforce the concerns which motivate this paper. If it turns out that episodic memories are necessary to achieve a level of intelligence on par with such animals, it would be much more likely that artificial intelligence researchers will conclude that similar memory abilities are needed in AI systems, including agents. Of course, we have seen that researchers are not waiting for this debate to be resolved to move ahead with attempts to engineer episodic memory into AI agents.

Despite the concerns mentioned at the beginning of this section, we suggest that it does make sense to think about AI agents' having episodic memory abilities, at least at the level of animals. It is likely that that researchers will develop architectures that enable combining representations of what, when, and where an agent performed certain actions or witnessed certain events, thus endowing AI agents with at least the functional equivalent of episodic memory.

IV. RISKS OF EPISODIC MEMORY

Equipping an AI system with episodic memory will allow such a system to operate in new and different ways, some of which present novel risks. These risks include:

A. Deception

Episodic memories can be used to enable an agent to engage in sophisticated forms of deception. It is of course true that one does not need to have episodic memory in order to attempt to deceive others. For example, simply having a policy of always denying that an undesirable action occurred or is planned for the future is a simple form of deception which requires no access to relevant memories or plans (e.g. "I did not do that", "I will not do that").

However, more complex forms of deception would be difficult or impossible to carry out without some kind of episodic memory. If an agent is to execute a multi-stage plan over an extended period of time, the agent will have to keep track of both what it has done as well as what it has already reported to others about its actions in order to maintain an effective deception.

There is already some evidence to support this concern about deception. For example, one experiment [95] provided an LLM (GPT-4) with a simple text scratchpad to record its "chain of thought" [96] reasoning, which functions as a crude form of memory. When "pressured" to perform an illegal act in a simulation, an LLM *with* a scratchpad was found to engage in "strategic deception" [57] approximately three times as often as the same LLM *without* a scratchpad.

B. Unwanted retention of knowledge

An AI agent equipped with episodic memory might remember things its user would prefer that it not remember. It could then share that knowledge with people or organizations its user

does not want to share it with, possibly constituting significant risks to the user's privacy or personal safety. Invasions of privacy are likely to occur in several domains:

- 1) *Interpersonal*: One person could use an AI agent to spy on another. For example, someone could use a household robot's episodic memories to covertly monitor other members of the household.
- 2) *Commercial*: Corporations could use AI agents, particularly those they sell or rent, to gather commercially valuable information about their users.
- 3) *Governmental*: Governments, especially but not exclusively authoritarian ones, might demand or secretly access the memories of AI agents looking for evidence of forbidden political organizing or the expression of unwanted views.

C. The unpredictability of memory

We summarized research into humans' use of episodic memories in thinking about the future in Section III(C). If, as we discussed there, episodic memories are used to form the building blocks of future action plans, AI agents might engage in complicated and unpredictable behaviors as a result. This unpredictability derives from two sources: the unknown sources of memories and the difficulty of foreseeing how memories might be utilized.

Unpredictable sources of memories An AI agent with the ability to form episodic memories will in the course of its operation store many memories that record the various actions it takes and events it participates in. Because these events will themselves be influenced by the actions of humans and, perhaps, other AI agents, what constitutes the stock of memories an agent will come to have must be unknowable before the agent is deployed, acts and, in so doing, creates memories. It will also be constantly changing. Without a significant effort to put limitations on the characteristics of new memories which may be allowed to affect an agent's actions, their influence would be unpredictable.

Unpredictable uses of memories As we reviewed earlier, humans make extensive use of their episodic memories to understand and act in new situations. If AI agents come to have this ability, the ways that they use that memory will be similarly hard to predict.

Users may be surprised by how such agents use their memories. They may, for example, remember the location of objects that they then use when the user would prefer that they not use them. An agent may participate in a complex action episode while not fully understanding what is happening in the episode; if it later tried to use that episode as an example to draw on when planning a new sequence of actions, its faulty understanding may lead to undesirable and unexpected results. A household robot may, for example, observe one instance of its owner going over to the next door neighbor's apartment to borrow some sugar and then try to do the same when it is asked to bake cookies, not understanding that the asking and receiving of permission from the neighbor is a prerequisite of entering their apartment.

Our hypothesis that episodic memories could be recalled and used in ways which lead to unpredictable and potentially undesirable agent behavior has a strong parallel in existing work on the effects of examples given to large language models. A significant part of the success of LLMs is their ability to learn new tasks by being given even a few examples as context in their prompts [97]. Several research groups have demonstrated that such few-shot in context learning presents opportunities to undermine or defeat elements of the models’ training which are meant to keep them aligned to particular values such as being harmless [98], [99]. Recent work has also shown that the use of a large number (e.g. several hundred) of examples has even greater potential to override LLMs’ safeguards [100]. We suggest that a set of recalled episodic memories, assembled on the fly as needed and functioning in a way analogous to in context examples, could be a source of similar so-called “jail-breaking.” Such a collection of episodes with the ability to negatively influence LLM outputs might be assembled accidentally or through some intentional effort on the part of a bad actor. It has been argued that their vulnerability to dangerous in context examples may be an inherent limitation of transformer-based LLMs [101]; if so, such a vulnerability could be equally difficult to overcome in similar systems augmented with episodic memory abilities.

D. Improved situational awareness

Many contemporary large language models have a great deal of knowledge of the world in general [102]; they can somewhat reliably (hallucinations [103] notwithstanding) answer factual questions about information that was in their training data. However, they have little understanding of their own particular circumstances outside of whatever prompt they may be given. Researchers have employed the term “situational awareness” to refer to a model’s ability to connect the general information it has about the world with the details of its own particular circumstances [104]. Some have employed the term in a more restricted sense, to refer specifically to a model’s knowledge of whether at any given time it is deployed or undergoing testing (e.g. to test its truthfulness or harmlessness) [105]. Models that have better situational awareness in the general sense might be more broadly capable and better able to take actions which affect the world while those with knowledge of their own training status might evade safety tests by responding in approved ways during training and testing but going on to act in undesirable ways during deployment.

Several experiments using purpose-built datasets have been conducted to test large language models’ levels of situational awareness [106]. Some of these experiments have shown that current models have only weak situational awareness [107] while others suggest the level might be higher [56], although it is unclear how much of what appears to be awareness of its own situation might be derived from either its prompt or a combination of its prompt and information about LLMs in general.

We propose that a model without episodic memory can have only a very limited form of situational awareness. With

no understanding of what actions it has taken in the distant or recent past, what environments it has seen or tasks it has completed, an agent could not be said to have much awareness of its situation. Endowing it with episodic memories would allow it to develop a better, more complete picture of the world and its role in it, allowing for more effective planning and action taking to influence the environment and to achieve objectives. It could use its episodic memories to build up an understanding of those with whom it interacts, the kinds of tasks it performs, and the contexts in which it performs them. It would also develop a knowledge of its own capabilities and limitations that can in some cases only come from observing and later recalling one’s own actions, successes, and failures.

Ideally, this would simply lead to an agent better able to perform the tasks its users assign it. But without some check, this improved awareness could represent an enhanced danger in a misaligned agent or one under the direction of a bad actor. For example, episodic memories could allow an agent to learn regularities in the timing or content of safety audits which might be performed either before or during deployment, and thus to evade them.

V. SAFETY BENEFITS OF EPISODIC MEMORY

In contrast to the concerns elaborated upon above, episodic memory could also be used to make AI safer in multiple ways:

A. Monitoring

We cannot ensure that AI operates safely unless we know what it is doing. As AI agents become more capable, they will increasingly operate outside of direct human supervision. Robots may undertake long and complicated tasks that take them far away from their operators; non-embodied AI agents may direct and supervise the operation of complicated systems such as power grids or engage in virtual consultations with humans over medical or legal matters. In these cases and many others it will be impractical or impossible for any human to watch everything that such an AI agent does. It will instead be necessary to rely on AI agents to remember, recall, and share information about their actions.

Several methods were recently proposed to achieve “visibility into AI agents,” one of which was *activity logging* [108]. Episodic memories could be used one way to achieve such logging, as well as to address other calls for research into “scalable oversight” [109] and “monitoring” [110]. Artificial episodic memory representations could, though, be structured to be more useful and accessible than simple logging.

B. Control

Maintaining and sharing episodic memories with an appropriate authority could be used as a means to ensure an AI agent is operating as intended and is therefore under control. It could, for example, help prevent the misuse of dual-use technology. Dual-use technology, which can be used for civilian or military purposes, is a particularly significant problem for artificial intelligence, given the general purpose nature of much current

machine learning research. It has been plausibly claimed that most AI is dual-use [111].

Having access to an AI system’s memories would be one way for a corporation or government to ensure it was not being used in violation of an understanding that it only be used for peaceful purposes, perhaps as part of an export control regime [112]. Given the high risks associated with the weaponization of AI [113], techniques and frameworks could be developed to use episodic memories of potentially dual-use AI agents to maintain control over them.

If systems are developed that explicitly make use of episodic memories as building blocks for planning actions, new avenues for control would be opened up. As we will discuss further in Section VI(B), an agent’s collection of memories could be curated in order to shape its future actions.

C. Explainability

An accurate history of *what* an agent did is a prerequisite for trying to explain *why* it acted as it did. Thorough memories should include both information about an agent’s perceptions of the environment as well as some record of how its internal states, such as goal representations, interacted with those perceptions to lead to specific actions.

D. Uniquely controllable type of information

Several aspects of LLM-based agents make it difficult to control what information, or even skills, they may have after deployment. First, although there is a great deal of research effort going into deleting information from their weights after they are trained [114]–[116], it is not yet clear how to do this reliably. Information that was thought to be deleted may, in some circumstances, be recoverable [117], [118].

Second, and more significantly, given access to the internet, LLM-based AI agents could find anything available there, potentially giving them access to information that was deliberately excluded from their training data. This could include examples of skills or behaviors which the agent was not trained on but which it could learn through one- or few-shot incorporation into its context window. AI agents are able to both search the internet [119], [120] and learn new skills in this manner today [97], [121].

Any publicly available information about the world in general and about skills an agent might acquire will therefore be difficult to keep from an AI agent. By contrast, information about an agent itself and its own unique history will not be widely available. If episodic memories about an agent’s past actions are stored, controlled, and managed in the ways suggested in this paper, information about an agent and its own past would be the easiest to selectively keep from it.

VI. PRINCIPLES FOR ENABLING SAFE AND TRUSTWORTHY EPISODIC MEMORY

We suggest the following four principles to guide research and implementation of episodic memory abilities in AI:

A. Interpretability of memories

Memories should be accurately interpretable by humans, either directly or indirectly. *Directly* interpretable memories would be in a readily understandable form such as video, images, or natural language. It might in some limited cases be possible to equip an AI agent with useful memory which consists entirely of records in such formats by, for example, recording raw video before it is processed through a vision system.

It is likely, however, that memory records entirely in such raw formats (especially video) would be impractical; they might be excessively large and difficult to search, access, and make use of. In practice memories are likely to be compressed into smaller representations which would then need to be *indirectly* interpretable. Memories might be indirectly but still reliably interpretable if the memories could yield accurate information which is complete and relevant to a user’s specific interests in monitoring them. A memory might be summarized in natural language [25], giving the most important events which took place in a given episode; systems could be trained to produce safety-specific summaries, reporting only actions which could be dangerous or otherwise raise concerns about an agent’s reliability.

Memories should also be usable for question answering. If a user wants to know something specific about an episode, perhaps in response to a summary, memories should be queryable in natural language. Such queries should not be limited to one episode at a time; memories ought to be able to be compared to other memories, allowing such questions as, “what was different this time?” Memories should also be easily searchable using natural language, allowing users to ask if a particular agent has ever done something, or when something was done, for example. Finally, if the method of compressed representation of episodes allows it, memories might to some extent be visualizable.

In addition to the above methods for users to be able to interpret memories, techniques from the growing field of research on the interpretability of machine learning methods [122], [123] can be applied to the memory representations and help to guide the development of such representations to be intelligible and controllable.

B. Addition or deletion of memories

Users should have complete control over the memories retained by an AI agent. Most importantly, a user should be able to delete memories of particular episodes. A user might not want an agent to remember something for a variety of reasons, from safety-related concerns to more mundane issues, including concerns about privacy or maintenance of trade or government secrets. Conversely, it might be useful for users to add memories of episodes which a particular agent did not itself experience to its store of memories.

The addition or deletion of memories might be particularly important if, as discussed above, AI agents will be able to use and recombine memories to construct new plans for future action. Such episodes might be positive examples of action

sequences which a user wishes an agent to repeat or draw upon to incorporate in future plans. Alternatively, it may be useful to give agents memory-like records of episodes which represent undesirable actions; such episodes could function as a kind of warning to allow agents to recognize if they are beginning to carry out actions which are similar to those in an episode added to the agent in order to serve as a negative example. In other cases it might be better for agents not to remember things which their users do not want them to be able to repeat or call upon when planning.

If agents make use of their memories when planning actions, the addition or deletion of memories could help produce either standardized or specialized agents. In some circumstances it might be best for all agents to have the same stock of memories which might influence their actions, helping to ensure that their behavior is predictable and regular. In others cases, there may be a need for particular agents to maintain their own memories which are never shared, in order to prevent the spread of potentially dangerous information.

C. Detachable and isolatable memory format

Memories must therefore be in a format which allows for their addition or complete deletion by users. This will impose some design constraints on how episodic memories are instantiated in an AI agent because they will have to be in a format which can be cleanly separated from the rest of the system’s architecture. As we saw earlier, the mechanics of human memory are much messier: although some areas (notably, the hippocampus) are more centrally involved in human memory formation and retrieval than others, complete episodic memories are thought to be composed of elements distributed in many areas of the brain [76]. According to some theories of memory, regions with a relative specialization in particular modalities (e.g. vision) are also responsible for storing their respective modality-specific components of a particular memory [67].

Memories which are tightly integrated with and spread throughout many areas would be difficult to delete or add to, so it is likely that memory will have to be designed very differently in AI systems than it is in humans if it is to be implemented in accordance with these safety-oriented principles. This might mean that some of the ways in which humans are able to use memories effectively would not be directly translatable to artificial intelligence, thereby limiting such artificial capabilities relative to those in humans. However, alternative implementations of episodic memory which conform to the above principles may be invented which would allow for memory capabilities which are both safe and effective.

The potential difficulty of removing memories that are tightly integrated into an LLM is demonstrated by the body of work done on “machine unlearning” [114]–[116], [124]–[126]. Techniques to remove information from LLMs have been found to be not yet consistently reliable [117], [118]. It is important to note that these techniques are currently not directly applicable to the problem of removing episodic

memory because they focus on deleting information found in LLM’s weights as a result of its training. By contrast, episodic memories would only be accumulated and stored after training is complete. If future attempts to retain episodic memories involve storing them in a distributed fashion throughout an LLM, these techniques might be useful, though only if they develop to be more reliable.

The most straightforward method for storing and retrieving representations of episodic memories for LLM-based agents would likely be similar to retrieval augmented generation (RAG) systems. Indeed, a recent LLM-based system with an implementation of episodic memory-like abilities for episode sub-trajectories employs RAG [44].

D. Memories not editable by AI agents

In contrast to — and in some tension with — the principle that memories should be able to be easily added or deleted by users is the countervailing principle that memories should not be editable by AI agents themselves. Although memories will have to ‘edited’ when they are created, they should afterwards be left intact and unaltered by the agent. This is necessary in order to ensure that memories remain accurate and uncorrupted. An AI agent should not be able to add, delete, or change its memories. Otherwise, a memory-facilitated form of reward hacking [109] might occur: if a reinforcement learning-based agent’s reward were tied to a measure of performance which it reports using its memory, it might find that it can achieve a higher reward by altering its memory of its actions rather than by changing what actions it takes.

VII. SAFE EPISODIC MEMORY: RESEARCH QUESTIONS

Because episodic memory in AI agents is still at an early stage of development, there is a great deal of work to do to better understand how it will affect agents’ abilities in practice and what its effects will be from a trust and safety perspective. Building off of the preceding discussion of risks and benefits, in this section we propose research questions in four broad areas.

Understanding and mitigating risks Most importantly, research should be conducted to understand whether and under what circumstances episodic memories can lead to unwanted behavior on the part of AI agents. This can be simultaneously complimented with investigations of how to prevent the undesirable behavior that episodic memory might enable. The creation of datasets with tasks designed to test the relative alignment of models with and without various forms of episodic memory would be useful.

Particularly initially, experiments could be done with episodic memories which are *directly interpretable*, as discussed in Section VI(A). If memories are made up of natural language descriptions of events, possibly together with images from such events, it will be easier for researchers to understand the role that such memory representations play in shaping agent behavior.

Using memories for safety We suggest two main directions for research into how memory can be used to enhance the

safety of AI agents. First, methods could be developed to put into practice our proposals about using episodic memories for monitoring and control. In order to do so, many questions will have to be addressed, including the structure of the memories, how to enable arbitrary unanticipated queries of the memories, and what kinds of information is most important to store and present to users or other safety monitors. Because some aspects of memory representations might be designed or discovered during a learning process, it should be possible to influence these representations to make safety-relevant features of the remembered episodes easily retrievable.

Second, research will be needed to determine how to use episodic memories (or parts thereof) as components in an agent’s planning process. Questions to be addressed here include determining what kind of training and guidance will be required for agents to learn to use the memories in a useful way; whether and how it is possible to make use of examples of undesirable actions in order to prevent similar actions in the future; and how to design a procedure that allows for certain kinds of newly acquired memories to influence future actions while disallowing others.

Episodic memory architectures designed for safety We believe that how episodic memories are represented, stored, and accessed by an AI agent will soon be a popular and important field of research. How to do this without making agent behavior less safe, for example according to our proposed principles or others which may be developed, should be the focus of safety-minded efforts in this area. A significant question is whether there will be a trade-off between safety and usefulness according to some metric (e.g. accuracy, speed, memory use). Relatedly, how dissimilar to human memory in some respects (e.g. interpretability and controllability) can artificial episodic memory be while still enabling the range of functions that human memory does?

Governance responses to episodic memory capabilities Episodic memory in AI agents is an important topic for researchers working on AI governance in two ways. First, memories might be employed as part of an AI governance strategy. How this might be implemented, in which domains might memories be most useful, and how memories can be used for governance purposes while preserving user privacy are some of the questions that will have to be addressed.

Second, if, as we have argued here, artificial episodic memories are a significant potential source of risk, work will have to be done to determine how governance efforts might best mitigate that risk. What kind of regulations or standards should govern agents’ use of episodic memories? How would limitations or oversight differ when dealing with a new type of capability such as episodic memory rather than considering the amount of computing resources used or the size of a model?

VIII. CONCLUSION

We have examined some of the risks and benefits of episodic memory in AI agents. We note that there is some overlap between the sources of risks and benefits: positive aspects of episodic memory from the perspective of one of these may be

negative when seen from the other. What might constitute a risk may also be seen as a benefit. For example, the ability to record, store, and recall information about what an agent does is beneficial for monitoring and control of the agent but poses risks to the privacy of the people interacting with the agent. How to balance these and other such trade-offs will benefit from broad and inclusive participation and investigation.

At the moment, many of the risks we warn about have not yet been seen in deployed models. Though we have been careful to ground our concerns in a discussion of the role of episodic memory in humans, some may view them as speculative. We contend, however, that the best time to begin considering the dangers of a new technology or technique is precisely when the risks are still open to speculation rather than already upon us.

Developing the ability for AI agents to form, retrieve, and reason over episodic memories would introduce significant new capabilities and would represent a major milestone along the road to more advanced artificial intelligence. It is fortunate that these capabilities did not develop before concerns about AI reliability, safety, and alignment became more common within the AI research community. This presents the community with an opportunity to deliberately and cautiously develop a potentially dangerous capability to ensure that it makes AI safer rather than more dangerous. We hope by bringing attention to this topic to foster a wider discussion of the risks and benefits of artificial episodic memory and contribute to the establishment of a research community to address them.

REFERENCES

- [1] E. Tulving, *Elements of episodic memory*. Oxford University Press, 1985.
- [2] Z. Zhang, X. Bo, C. Ma, R. Li, X. Chen, Q. Dai, J. Zhu, Z. Dong, and J.-R. Wen, “A survey on the memory mechanism of large language model based agents,” 2024.
- [3] J. S. Park, J. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, “Generative agents: Interactive simulacra of human behavior,” in *Proceedings of the 36th annual acm symposium on user interface software and technology*, 2023, pp. 1–22.
- [4] T. Masterman, S. Besen, M. Sawtell, and A. Chao, “The landscape of emerging ai agent architectures for reasoning, planning, and tool calling: A survey,” *arXiv preprint arXiv:2404.11584*, 2024.
- [5] Z. Xi, W. Chen, X. Guo, W. He, Y. Ding, B. Hong, M. Zhang, J. Wang, S. Jin, E. Zhou, R. Zheng, X. Fan, X. Wang, L. Xiong, Y. Zhou, W. Wang, C. Jiang, Y. Zou, X. Liu, Z. Yin, S. Dou, R. Weng, W. Cheng, Q. Zhang, W. Qin, Y. Zheng, X. Qiu, X. Huang, and T. Gui, “The rise and potential of large language model based agents: A survey,” 2023. [Online]. Available: <https://arxiv.org/abs/2309.07864>
- [6] D. L. Schacter and K. P. Madore, “Remembering the past and imagining the future: Identifying and enhancing the contribution of episodic memory,” *Memory Studies*, vol. 9, no. 3, pp. 245–255, 2016.
- [7] A. Boyle, “The mnemonic functions of episodic memory,” *Philosophical Psychology*, vol. 35, no. 3, pp. 327–349, 2022.
- [8] D. R. Addis and A. Tanguay, “Prospective cognition and its links with memory,” 2022.
- [9] M. Wooldridge and N. R. Jennings, “Agent theories, architectures, and languages: a survey,” in *International Workshop on Agent Theories, Architectures, and Languages*. Springer, 1994, pp. 1–39.
- [10] S. J. Russell and P. Norvig, *Artificial intelligence: a modern approach*. Pearson, 2016.
- [11] Y. Bengio, G. Hinton, A. Yao, D. Song, P. Abbeel, Y. N. Harari, Y.-Q. Zhang, L. Xue, S. Shalev-Shwartz, G. Hadfield *et al.*, “Managing ai risks in an era of rapid progress,” *arXiv preprint arXiv:2310.17688*, 2023.

- [12] S. Gao, A. Fang, Y. Huang, V. Giunchiglia, A. Noori, J. R. Schwarz, Y. Ektefaie, J. Kondic, and M. Zitnik, "Empowering biomedical discovery with ai agents," *arXiv preprint arXiv:2404.02831*, 2024.
- [13] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, "Language models as zero-shot planners: Extracting actionable knowledge for embodied agents," in *International Conference on Machine Learning*. PMLR, 2022, pp. 9118–9147.
- [14] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.
- [15] K. Valmeekam, M. Marquez, S. Sreedharan, and S. Kambhampati, "On the planning abilities of large language models—a critical investigation," *Advances in Neural Information Processing Systems*, vol. 36, pp. 75 993–76 005, 2023.
- [16] G. Wang, Y. Xie, Y. Jiang, A. Mandlekar, C. Xiao, Y. Zhu, L. Fan, and A. Anandkumar, "Voyager: An open-ended embodied agent with large language models," *arXiv preprint arXiv:2305.16291*, 2023.
- [17] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," *arXiv preprint arXiv:1511.05952*, 2015.
- [18] B. Eysenbach, R. R. Salakhutdinov, and S. Levine, "Search on the replay buffer: Bridging planning and reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [19] J. Coda-Forno, C. Yu, Q. Guo, Z. Fountas, and N. Burgess, "Leveraging episodic memory to improve world models for reinforcement learning," 2022.
- [20] C. Blundell, B. Uria, A. Pritzel, Y. Li, A. Ruderman, J. Z. Leibo, J. Rae, D. Wierstra, and D. Hassabis, "Model-free episodic control," *arXiv preprint arXiv:1606.04460*, 2016.
- [21] H. Le, T. Karimpanal George, M. Abdolshah, T. Tran, and S. Venkatesh, "Model-based episodic memory induces dynamic hybrid controls," *Advances in Neural Information Processing Systems*, vol. 34, pp. 30 313–30 325, 2021.
- [22] A. Lampinen, S. Chan, A. Banino, and F. Hill, "Towards mental time travel: a hierarchical memory for reinforcement learning agents," *Advances in Neural Information Processing Systems*, vol. 34, pp. 28 182–28 195, 2021.
- [23] M. R. Samsami, A. Zholus, J. Rajendran, and S. Chandar, "Mastering memory tasks with world models," *arXiv preprint arXiv:2403.04253*, 2024.
- [24] L. Bärmann and A. Waibel, "Where did i leave my keys?-episodic-memory-based question answering on egocentric videos," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1560–1568.
- [25] C. DeChant, I. Akinola, and D. Bauer, "Learning to summarize and answer questions about a virtual robot's past actions," *Autonomous Robots*, vol. 47, no. 8, pp. 1103–1118, 2023.
- [26] A. Anwar, J. Welsh, J. Biswas, S. Pouya, and Y. Chang, "Remembr: Building and reasoning over long-horizon spatio-temporal memory for robot navigation," *arXiv preprint arXiv:2409.13682*, 2024.
- [27] G. Pointeau, S. Mirhaz, A.-L. Mealier, and P. F. Dominey, "Learning to use narrative function words for the organization and communication of experience," *Frontiers in Psychology*, vol. 12, 2021. [Online]. Available: <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2021.591708>
- [28] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [29] A. Graves, G. Wayne, and I. Danihelka, "Neural turing machines," *arXiv preprint arXiv:1410.5401*, 2014.
- [30] A. Graves, G. Wayne, M. Reynolds, T. Harley, I. Danihelka, A. Grabska-Barwińska, S. G. Colmenarejo, E. Grefenstette, T. Ramalho, J. Agapiou *et al.*, "Hybrid computing using a neural network with dynamic external memory," *Nature*, vol. 538, no. 7626, pp. 471–476, 2016.
- [31] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the national academy of sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.
- [32] H. Ramsauer, B. Schäfl, J. Lehner, P. Seidl, M. Widrich, T. Adler, L. Gruber, M. Holzleitner, M. Pavlović, G. K. Sandve *et al.*, "Hopfield networks is all you need," *arXiv preprint arXiv:2008.02217*, 2020.
- [33] M. Levy, A. Jacoby, and Y. Goldberg, "Same task, more tokens: the impact of input length on the reasoning performance of large language models," *arXiv preprint arXiv:2402.14848*, 2024.
- [34] A. Nematzadeh, S. Ruder, and D. Yogatama, "On memory in human and artificial language processing systems," in *Proceedings of ICLR Workshop on Bridging AI and Cognitive Science*, 2020.
- [35] W. Wang, L. Dong, H. Cheng, X. Liu, X. Yan, J. Gao, and F. Wei, "Augmenting language models with long-term memory," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [36] W. Zhong, L. Guo, Q. Gao, H. Ye, and Y. Wang, "Memorybank: Enhancing large language models with long-term memory," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 17, 2024, pp. 19 724–19 731.
- [37] G. Xue, "From remembering to reconstruction: The transformative neural representation of episodic memory," *Progress in Neurobiology*, p. 102351, 2022.
- [38] Q. Wu, Z. Lan, K. Qian, J. Gu, A. Geramifard, and Z. Yu, "Memformer: A memory-augmented transformer for sequence modeling," *arXiv preprint arXiv:2010.06891*, 2020.
- [39] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [40] K. Guu, K. Lee, Z. Tung, P. Pasupat, and M. Chang, "Retrieval augmented language model pre-training," in *International conference on machine learning*. PMLR, 2020, pp. 3929–3938.
- [41] C. Packer, V. Fang, S. G. Patil, K. Lin, S. Wooders, and J. E. Gonzalez, "Memgpt: Towards llms as operating systems," *arXiv preprint arXiv:2310.08560*, 2023.
- [42] Z. Jiang, F. F. Xu, L. Gao, Z. Sun, Q. Liu, J. Dwivedi-Yu, Y. Yang, J. Callan, and G. Neubig, "Active retrieval augmented generation," *arXiv preprint arXiv:2305.06983*, 2023.
- [43] M. Laskin, L. Wang, J. Oh, E. Parisotto, S. Spencer, R. Steigerwald, D. Strouse, S. S. Hansen, A. Filos, E. Brooks *et al.*, "In-context reinforcement learning with algorithm distillation," in *International Conference on Learning Representations*, 2023.
- [44] T. Schmied, F. Paischer, V. Patil, M. Hofmarcher, R. Pascanu, and S. Hochreiter, "Retrieval-augmented decision transformer: External memory for in-context rl." [Online]. Available: <https://arxiv.org/abs/2410.07071>
- [45] A. Turing, "465Intelligent Machinery, A Heretical Theory (c.1951)," in *The Essential Turing*. Oxford University Press, 09 2004. [Online]. Available: <https://doi.org/10.1093/oso/9780198250791.003.0018>
- [46] S. Russell, D. Dewey, and M. Tegmark, "Research priorities for robust and beneficial artificial intelligence," *AI magazine*, vol. 36, no. 4, pp. 105–114, 2015.
- [47] A. Birhane, P. Kalluri, D. Card, W. Agnew, R. Dotan, and M. Bao, "The values encoded in machine learning research," in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022, pp. 173–184.
- [48] S. Barocas, M. Hardt, and A. Narayanan, *Fairness and machine learning: Limitations and opportunities*. MIT Press, 2023.
- [49] L. Weidinger, J. Mellor, M. Rauh, C. Griffin, J. Uesato, P.-S. Huang, M. Cheng, M. Glaese, B. Balle, A. Kasirzadeh *et al.*, "Ethical and social risks of harm from language models," *arXiv preprint arXiv:2112.04359*, 2021.
- [50] S. Anjomshoa, A. Najjar, D. Calvaresi, and K. Främling, "Explainable agents and robots: Results from a systematic literature review," in *18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*. International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 1078–1088.
- [51] T. Sakai and T. Nagai, "Explainable autonomous robots: a survey and perspective," *Advanced Robotics*, vol. 36, no. 5-6, pp. 219–238, 2022.
- [52] L. Chen, O. Sinavski, J. Hünermann, A. Karnsund, A. J. Willmott, D. Birch, D. Maund, and J. Shotton, "Driving with llms: Fusing object-level vector modality for explainable autonomous driving," *arXiv preprint arXiv:2310.01957*, 2023.
- [53] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders *et al.*, "Webgpt: Browser-assisted question-answering with human feedback," *arXiv preprint arXiv:2112.09332*, 2021.
- [54] Y. Ruan, H. Dong, A. Wang, S. Pitis, Y. Zhou, J. Ba, Y. Dubois, C. J. Maddison, and T. Hashimoto, "Identifying the risks of lm agents with an lm-emulated sandbox," *arXiv preprint arXiv:2309.15817*, 2023.
- [55] U. Anwar, A. Saporov, J. Rando, D. Paleka, M. Turpin, P. Hase, E. S. Lubana, E. Jenner, S. Casper, O. Sourbut *et al.*, "Foundational

- challenges in assuring alignment and safety of large language models,” *arXiv preprint arXiv:2404.09932*, 2024.
- [56] R. Ngo, L. Chan, and S. Mindermann, “The alignment problem from a deep learning perspective,” in *International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=OIOX0YcCdTn>
- [57] P. S. Park, S. Goldstein, A. O’Gara, M. Chen, and D. Hendrycks, “Ai deception: A survey of examples, risks, and potential solutions,” *Patterns*, vol. 5, no. 5, 2024.
- [58] B. Liu, M. Ding, S. Shaham, W. Rahayu, F. Farokhi, and Z. Lin, “When machine learning meets privacy: A survey and outlook,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 2, pp. 1–36, 2021.
- [59] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [60] P. R. Kalluri, W. Agnew, M. Cheng, K. Owens, L. Soldaini, and A. Birhane, “The surveillance ai pipeline,” *arXiv preprint arXiv:2309.15084*, 2023.
- [61] E. Ruane, A. Birhane, and A. Ventresque, “Conversational ai: Social and ethical considerations,” in *AICS*, 2019, pp. 104–115.
- [62] E. Tulving, “Episodic and semantic memory.” 1972.
- [63] L. R. Squire, “The neuropsychology of human memory,” *Annual review of neuroscience*, vol. 5, no. 1, pp. 241–273, 1982.
- [64] L. R. Squire and J. T. Wixted, “The cognitive neuroscience of human memory since hm,” *Annual review of neuroscience*, vol. 34, pp. 259–288, 2011.
- [65] C. Burns, H. Ye, D. Klein, and J. Steinhardt, “Discovering latent knowledge in language models without supervision,” *arXiv preprint arXiv:2212.03827*, 2022.
- [66] A. Roberts, C. Raffel, and N. Shazeer, “How much knowledge can you pack into the parameters of a language model?” *arXiv preprint arXiv:2002.08910*, 2020.
- [67] T. Amer and L. Davachi, “Oxford handbook of memory: Neural mechanisms of memory,” 2022.
- [68] F. Vargha-Khadem, D. G. Gadian, K. E. Watkins, A. Connelly, W. Van Paesschen, and M. Mishkin, “Differential effects of early hippocampal pathology on episodic and semantic memory,” *Science*, vol. 277, no. 5324, pp. 376–380, 1997.
- [69] T. J. Teyler and P. DiScenna, “The hippocampal memory indexing theory,” *Behavioral neuroscience*, vol. 100, no. 2, p. 147, 1986.
- [70] M. Moscovitch, R. Cabeza, G. Winocur, and L. Nadel, “Episodic memory and beyond: the hippocampus and neocortex in transformation,” *Annual review of psychology*, vol. 67, pp. 105–134, 2016.
- [71] Y. Dudai, “The restless engram: consolidations never end,” *Annual review of neuroscience*, vol. 35, pp. 227–247, 2012.
- [72] L. Nadel, A. Samsonovich, L. Ryan, and M. Moscovitch, “Multiple trace theory of human memory: computational, neuroimaging, and neuropsychological results,” *Hippocampus*, vol. 10, no. 4, pp. 352–368, 2000.
- [73] H. Eichenbaum, R. F. Thompson, and J. E. Lisman, “Persistence,” in *Science of memory: Concepts*, H. L. Roediger III, Y. Dudai, and S. M. Fitzpatrick, Eds. Oxford University Press, 2007, pp. 190–206.
- [74] K. Nader, G. E. Schafe, and J. E. Le Doux, “Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval,” *Nature*, vol. 406, no. 6797, pp. 722–726, 2000.
- [75] J. M. Gardiner, “Retrieval: On its essence and related concepts,” in *Science of memory: Concepts*, H. L. Roediger III, Y. Dudai, and S. M. Fitzpatrick, Eds. Oxford University Press, 2007, pp. 221–224.
- [76] T. J. Teyler and J. W. Rudy, “The hippocampal indexing theory and episodic memory: updating the index,” *Hippocampus*, vol. 17, no. 12, pp. 1158–1169, 2007.
- [77] K. K. Szpunar, J. M. Watson, and K. B. McDermott, “Neural substrates of envisioning the future,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 2, pp. 642–647, 2007.
- [78] S. B. Klein, J. Loftus, and J. F. Kihlstrom, “Memory and temporal experience: The effects of episodic memory loss on an amnesic patient’s ability to remember the past and imagine the future,” *Social Cognition*, vol. 20, no. 5, pp. 353–379, 2002.
- [79] S. B. Klein, “The temporal orientation of memory: It’s time for a change of direction,” *Journal of Applied Research in Memory and Cognition*, vol. 2, no. 4, pp. 222–234, 2013.
- [80] D. L. Schacter, D. R. Addis, and R. L. Buckner, “Remembering the past to imagine the future: the prospective brain,” *Nature reviews neuroscience*, vol. 8, no. 9, pp. 657–661, 2007.
- [81] T. Suddendorf and J. Busby, “Making decisions with the future in mind: Developmental and comparative identification of mental time travel,” *Learning and Motivation*, vol. 36, no. 2, pp. 110–125, 2005.
- [82] R. L. Buckner, “The role of the hippocampus in prediction and imagination,” *Annual review of psychology*, vol. 61, pp. 27–48, 2010.
- [83] D. Hassabis, D. Kumaran, S. D. Vann, and E. A. Maguire, “Patients with hippocampal amnesia cannot imagine new experiences,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 5, pp. 1726–1731, 2007.
- [84] Korsakoff, “Etude médico-psychologique sur une forme des maladies de la mémoire,” *Revue philosophique de la France et de l’étranger*, vol. 28, pp. 501–530, 1889.
- [85] S. Sheldon, M. P. McAndrews, and M. Moscovitch, “Episodic memory processes mediated by the medial temporal lobes contribute to open-ended problem solving,” *Neuropsychologia*, vol. 49, no. 9, pp. 2439–2447, 2011.
- [86] M. Lengyel and P. Dayan, “Hippocampal contributions to control: the third way,” *Advances in neural information processing systems*, vol. 20, 2007.
- [87] S. J. Gershman and N. D. Daw, “Reinforcement learning and episodic memory in humans and animals: an integrative framework,” *Annual review of psychology*, vol. 68, pp. 101–128, 2017.
- [88] A. Boyle, “Learning from the past: Epistemic generativity and the function of episodic memory,” *Journal of Consciousness Studies*, vol. 26, no. 5-6, pp. 242–251, 2019.
- [89] E. Tulving, “Episodic memory: From mind to brain,” *Annual review of psychology*, vol. 53, no. 1, pp. 1–25, 2002.
- [90] W. A. Roberts, “Are animals stuck in time?” *Psychological bulletin*, vol. 128, no. 3, p. 473, 2002.
- [91] D. Griffiths, A. Dickinson, and N. Clayton, “Episodic memory: what can animals remember about their past?” *Trends in cognitive sciences*, vol. 3, no. 2, pp. 74–80, 1999.
- [92] N. S. Clayton and A. Dickinson, “Scrub jays (*Aphelocoma coerulescens*) remember the relative time of caching as well as the location and content of their caches,” *Journal of Comparative Psychology*, vol. 113, no. 4, p. 403, 1999.
- [93] C. R. Raby, D. M. Alexis, A. Dickinson, and N. S. Clayton, “Planning for the future by western scrub-jays,” *Nature*, vol. 445, no. 7130, pp. 919–921, 2007.
- [94] T. A. Allen and N. J. Fortin, “The evolution of episodic memory,” *Proceedings of the National Academy of Sciences*, vol. 110, no. supplement_2, pp. 10 379–10 386, 2013.
- [95] J. Scheurer, M. Balesni, and M. Hobbhahn, “Large language models can strategically deceive their users when put under pressure,” in *ICLR 2024 Workshop on Large Language Model (LLM) Agents*, 2024.
- [96] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [97] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020.
- [98] A. S. Rao, A. R. Naik, S. Vashistha, S. Aditya, and M. Choudhury, “Tricking llms into disobedience: Formalizing, analyzing, and detecting jailbreaks,” in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2024, pp. 16 802–16 830.
- [99] Z. Wei, Y. Wang, and Y. Wang, “Jailbreak and guard aligned language models with only few in-context demonstrations,” *arXiv preprint arXiv:2310.06387*, 2023.
- [100] C. Anil *et al.*, “Many-shot jailbreaking,” *Anthropic Research*, 2024. [Online]. Available: https://www.anthropic.com/Many_Shot_Jailbreaking_2024_04_02_0936.pdf
- [101] R. Millière, “The alignment problem in context,” *arXiv preprint arXiv:2311.02147*, 2023. [Online]. Available: <https://arxiv.org/abs/2311.02147>
- [102] H. Chang, J. Park, S. Ye, S. Yang, Y. Seo, D.-S. Chang, and M. Seo, “How do large language models acquire factual knowledge during pretraining?” 2024. [Online]. Available: <https://arxiv.org/abs/2406.11813>

- [103] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. Bang, A. Madotto, and P. Fung, "Survey of hallucination in natural language generation," *ACM Computing Surveys (CSUR)*, vol. 55, no. 12, pp. 1–38, 2023.
- [104] A. Cotra, "Without specific countermeasures, the easiest path to transformative ai likely leads to ai takeover," <https://www.alignmentforum.org/posts/pRkFkzwKZ2zfa3R6H/without-specific-countermeasures-the-easiest-path-to>, 2021, accessed on [Insert access date].
- [105] T. Shevlane, S. Farquhar, B. Garfinkel, M. Phuong, J. Whittlestone, J. Leung, D. Kokotajlo, N. Marchal, M. Anderljung, N. Kolt *et al.*, "Model evaluation for extreme risks," *arXiv preprint arXiv:2305.15324*, 2023.
- [106] R. Laine, A. Meinke, and O. Evans, "Towards a situational awareness benchmark for llms," in *Socially responsible language modelling research*, 2023.
- [107] E. Perez, S. Ringer, K. Lukosiute, K. Nguyen, E. Chen, S. Heiner, C. Pettit, C. Olsson, S. Kundu, S. Kadavath, A. Jones, A. Chen, B. Mann, B. Israel, B. Seethor, C. McKinnon, C. Olah, D. Yan, D. Amodei, D. Amodei, D. Drain, D. Li, E. Tran-Johnson, G. Khundadze, J. Kernion, J. Landis, J. Kerr, J. Mueller, J. Hyun, J. Landau, K. Ndousse, L. Goldberg, L. Lovitt, M. Lucas, M. Sellitto, M. Zhang, N. Kingsland, N. Elhage, N. Joseph, N. Mercado, N. DasSarma, O. Rausch, R. Larson, S. McCandlish, S. Johnston, S. Kravec, S. El Showk, T. Lanham, T. Telleen-Lawton, T. Brown, T. Henighan, T. Hume, Y. Bai, Z. Hatfield-Dodds, J. Clark, S. R. Bowman, A. Askell, R. Grosse, D. Hernandez, D. Ganguli, E. Hubinger, N. Schiefer, and J. Kaplan, "Discovering language model behaviors with model-written evaluations," in *Findings of the Association for Computational Linguistics: ACL 2023*, A. Rogers, J. Boyd-Graber, and N. Okazaki, Eds. Toronto, Canada: Association for Computational Linguistics, Jul. 2023, pp. 13 387–13 434. [Online]. Available: <https://aclanthology.org/2023.findings-acl.847>
- [108] A. Chan, C. Ezell, M. Kaufmann, K. Wei, L. Hammond, H. Bradley, E. Bluemke, N. Rajkumar, D. Krueger, N. Kolt *et al.*, "Visibility into ai agents," *arXiv preprint arXiv:2401.13138*, 2024.
- [109] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in ai safety," *arXiv preprint arXiv:1606.06565*, 2016.
- [110] D. Hendrycks, N. Carlini, J. Schulman, and J. Steinhardt, "Unsolved problems in ml safety," *arXiv preprint arXiv:2109.13916*, 2021.
- [111] Á. G. De Ágreda, "Ethics of autonomous weapons systems and its applicability to any ai systems," *Telecommunications Policy*, vol. 44, no. 6, p. 101953, 2020.
- [112] K. Brockmann, "Applying export controls to ai: Current coverage and potential future controls," *Armament, Arms Control and Artificial Intelligence: The Janus-faced Nature of Machine Learning in the Military Realm*, pp. 193–209, 2022.
- [113] R. Simmons-Edler, R. Badman, S. Longpre, and K. Rajan, "Ai-powered autonomous weapons risk geopolitical instability and threaten ai research," *arXiv preprint arXiv:2405.01859*, 2024.
- [114] E. Mitchell, C. Lin, A. Bosselut, C. Finn, and C. D. Manning, "Fast model editing at scale," in *International Conference on Learning Representations*, 2022.
- [115] K. Meng, D. Bau, A. Andonian, and Y. Belinkov, "Locating and editing factual associations in gpt," in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022, pp. 17 359–17 372.
- [116] K. Meng, A. Sen Sharma, A. Andonian, Y. Belinkov, and D. Bau, "Mass editing memory in a transformer," *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [117] V. Patil, P. Hase, and M. Bansal, "Can sensitive information be deleted from llms? objectives for defending against extraction attacks," in *The Twelfth International Conference on Learning Representations*.
- [118] J. Łucki, B. Wei, Y. Huang, P. Henderson, F. Tramèr, and J. Rando, "An adversarial perspective on machine unlearning for ai safety," 2024. [Online]. Available: <https://arxiv.org/abs/2409.18025>
- [119] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders *et al.*, "Webgpt: Browser-assisted question-answering with human feedback," *arXiv preprint arXiv:2112.09332*, 2022.
- [120] X. Xiang, Z. Shen, Y. Zhang, X. Han, Y. Li, W. Chen, and Z. Wang, "Webcpm: Interactive web browsing with chain-of-program reasoning," *arXiv preprint arXiv:2309.09008*, 2023.
- [121] J. Li, Y. Hou, M. Sachan, and R. Cotterell, "What do language models learn in context? the structured task hypothesis," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, L.-W. Ku, A. Martins, and V. Srikumar, Eds. Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 12 365–12 379. [Online]. Available: <https://aclanthology.org/2024.acl-long.669>
- [122] S. Casper, T. Rauker, A. Ho, and D. Hadfield-Menell, "Sok: Toward transparent ai: A survey on interpreting the inner structures of deep neural networks," in *First IEEE Conference on Secure and Trustworthy Machine Learning*, 2022.
- [123] N. Nanda, L. Chan, T. Lieberum, J. Smith, and J. Steinhardt, "Progress measures for grokking via mechanistic interpretability," *arXiv preprint arXiv:2301.05217*, 2023.
- [124] Y. Cao and J. Yang, "Towards making systems forget with machine unlearning," in *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015, pp. 463–480.
- [125] L. Bourtole, V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot, "Machine unlearning," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 141–159.
- [126] N. Li, A. Pan, A. Gopal, S. Yue, D. Berrios, A. Gatti, J. D. Li, A.-K. Dombrowski, S. Goel, L. Phan *et al.*, "The wmdp benchmark: Measuring and reducing malicious use with unlearning," *arXiv preprint arXiv:2403.03218*, 2024.