AI Agents: What Are They, How Do They Work, Why Should I Care?

Where I explain what AI Agents are and how they can help you get more done

TOM MARTIN
AUG 06, 2024

♥ 7 💬 2 🔁 1 Share •••

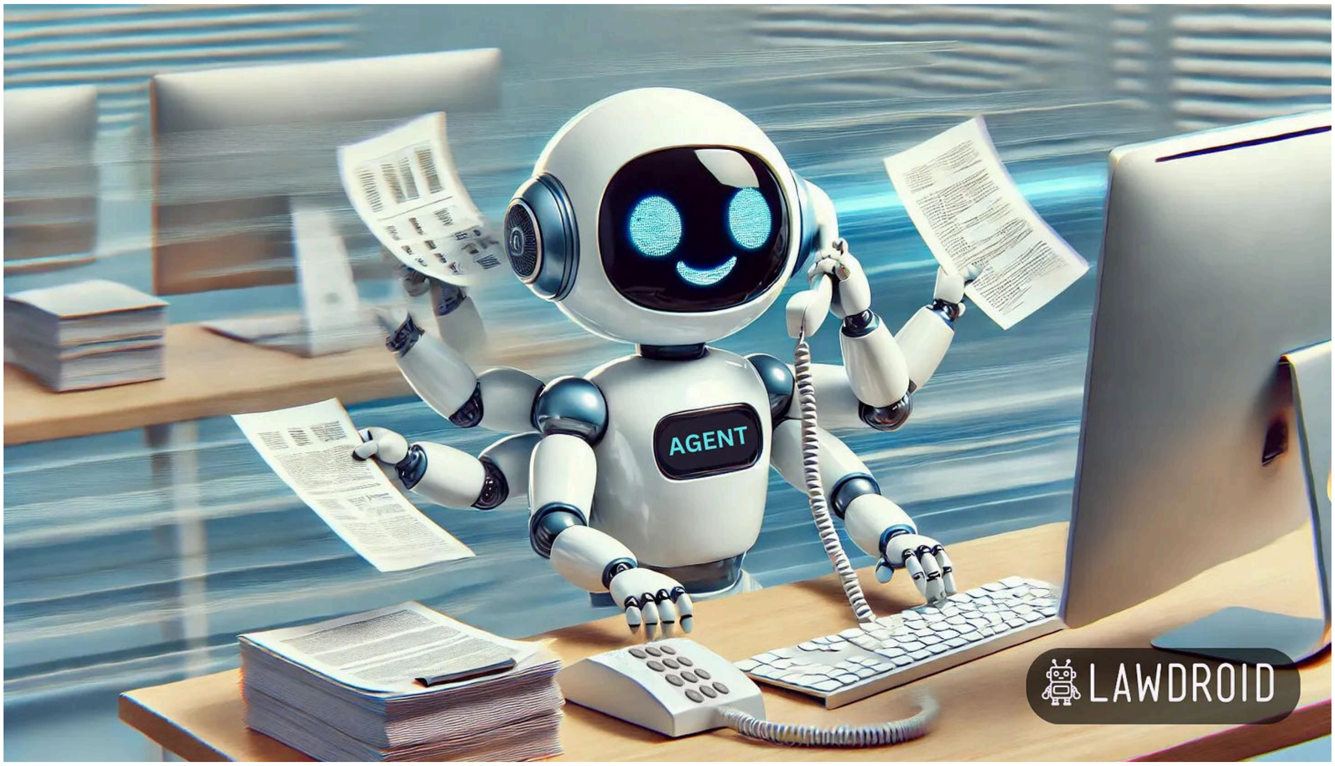

Welcome back, my incredible readers! You're the stars in my galaxy. 🌟

And for you newcomers, make yourself at home – we've got a virtual reality headset waiting for you to explore new worlds! 🥽

LawDroid Manifesto is a reader-supported publication. To receive new posts and support

my work, consider becoming a free or paid
subscriber.

This substack, LawDroid Manifesto, is here to keep you in the loop about the intersection of AI and the law. Please share this article with your friends and colleagues and remember to tell me what you think in the comments below.

Alrighty folks, hold onto your pocket protectors because we're about to embark on a wild ride into the world of AI agents! This year everyone is talking about AI Agents [1]. 🤖 So, you may be wondering, "What exactly is an AI agent?" Well, they combine "large language models," "knowledge bases," and "tool usage," concepts that might sound like they belong in a computer science textbook but are actually the secret sauce behind these autonomous digital sidekicks. Plus, it's like giving your AI assistant a fully-stocked backpack before sending it off on an adventure – the more tools and knowledge it has, the better equipped it'll be to tackle any challenge that comes its way. Makes sense, right?

If this sounds interesting to you (and you wanna earn some serious nerd cred in the process), please read on…

By the way, if you haven't read them yet, my previous articles about retrieval augmented generation, prompt engineering and hallucinations provide a great foundation for better understanding AI Agents.

— Reality Check —

Before I get into the substance of this article, let me acknowledge that we are presently in the trough of disillusionment in relation to Generative AI in the legal profession right now. This will change soon enough, and I'm not saying this disposition is universal or inevitable, but the naysayer's are having their day. I say this, not to discourage you of course, if anything the opposite. I believe that by you being here and learning more about AI, the better positioned you will be. Knowledge is power and never was that not more true than with understanding the power of AI.

Okay, with that, let's begin.

What Is an AI Agent?

An AI Agent can take intelligent action to do something in the world on its own.

OpenAI dipped its toe in the water with Custom GPTs and there's been an explosion of interest in what AI Agents or "Agentic AI" can accomplish for us, in everything from computer programming to customer service. [2] An AI Agent can be employed to answer your phones, take messages, and transfer calls. An AI Agent can be tasked with engaging website visitors, answering frequently asked questions and scheduling consultations. An AI Agent can be assigned to interview clients, review documents and information and draft legal documents.

Sound familiar? These are all job skills you may require of a receptionist, administrative assistant or junior lawyer or paralegal to do at your law firm. The difference of course is that, once you train the AI Agent, it can perform these functions 24/7/365, all for a recurring subscription fee.

How is an AI Agent different than an AI Copilot?

An AI Copilot, like ChatGPT, Claude, Perplexity or the like works with you to help you to accomplish a goal. That's why the "Copilot" moniker is so appropriate. Because, a copilot is not meant to replace the captain; you, the captain, are still in charge.

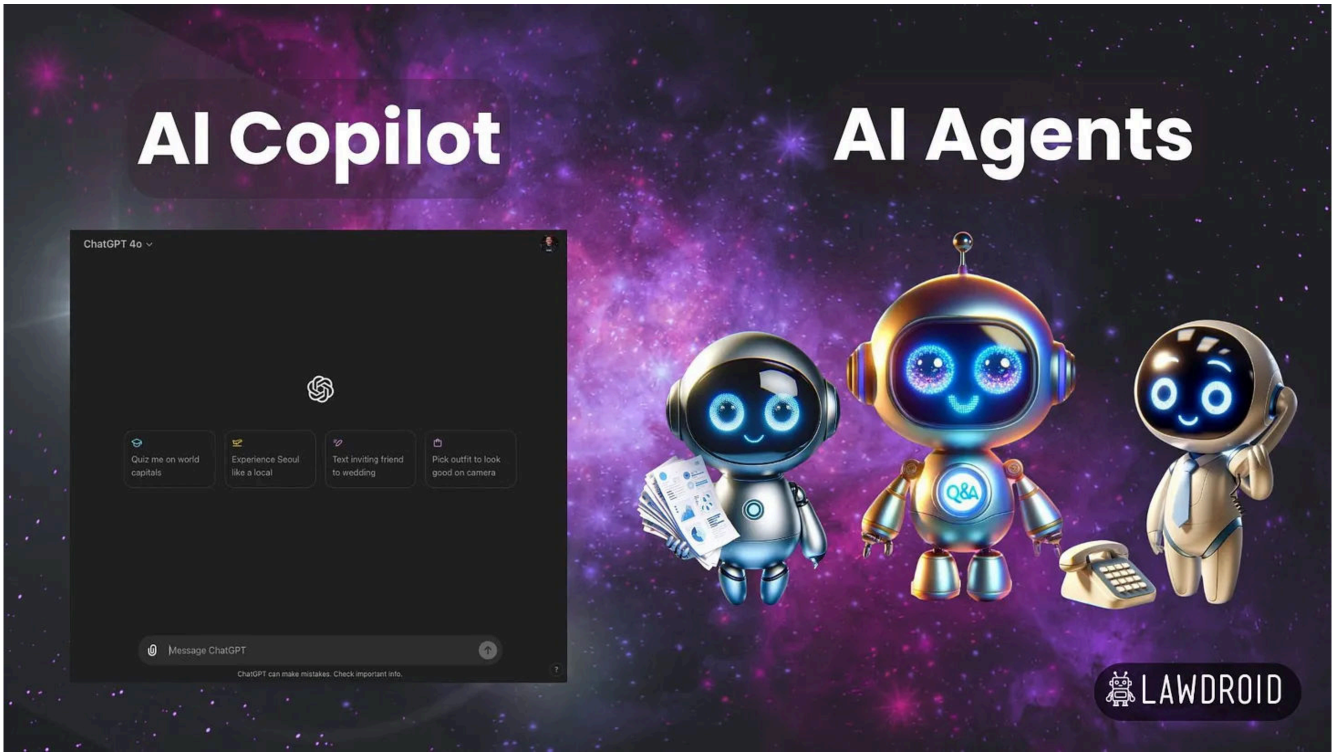
AI Copilot
AI Agents
ChatGPT 4o
Quiz me on world capitals
Experience Seoul like a local
Text inviting friend to wedding
Pick outfit to look good on camera
Message ChatGPT
ChatGPT can make mistakes. Check important info.
Q&A
LAWDROID

A literal copilot or first officer, as it were, is meant to assist so that the captain can do a good job. The copilot helps to safely navigate the plane and assist the captain with all other needs such as weather reporting and safety checks. Likewise, an AI copilot can help with things like grammar check, translation, drafting emails or letters, researching a legal question, or summarizing a deposition transcript.

Fundamentally, an AI Copilot is helpful, but limited. The copilot relationship is 1 to 1. The user inputs a prompt and the AI outputs a response. The AI performs one function at a time. The user asks the AI to summarize a document and must wait on the AI to complete that task before asking it to do something else. The user and the AI also have to be in the same place at the same time (synchronous). The AI does not act on its own, the user must be present to input the request. In sum, the AI Copilot is helpful but limited because its use must be 1 to 1 and synchronous.

An AI Agent, on the other hand, is autonomous. The user is not required to prompt the AI from step to step. Consequently, the user and AI Agent need not interact in real time to accomplish the goal. Once the job is defined, the AI Agent simply undertakes to plan for and accomplish the goal it’s been assigned.

But what about chatbots? Weren't they supposed to do that? How are they different than AI agents? Excellent questions. Let's get into it.

How is an AI Agent different than a Chatbot?

Chatbots were really the version 1, or the first draft, of AI agents.

Chatbots returned to the scene in 2016. [3] I say "returned" because the first chatbot, Eliza, was created in 1964! But cloud computing and early natural language processing in the late 2010s empowered a new wave of chatbots that promised to perform work in the world, on their own, and without the need for human intervention.

But, as we all likely know from our own personal experience, chatbots did not fulfill that promise. Chatbots weren't very smart or knowledgeable and were brittle. You might visit your bank's website and be greeted by a friendly-enough chatbot: "How may I help you today?" But if you asked it something it wasn't specifically programmed to answer, it would break, and you would get the dreaded death spiral of "I'm sorry, I didn't get that, can you rephrase?"

That said, chatbots, even with their limitations, had a positive effect on productivity. For example, in 2016, the communication department of the airline KLM was employing 235 people who had to answer 15,000 questions on average per week. According to KLM data, 1.7 million messages were sent through Facebook messenger to 500,000 users. The company needed help with the flood of messages. KLM started to use a chatbot on their Facebook profile, which was responsible for updates regarding check-in, delays or issuing of copies of boarding passes. A living member of the team would join a conversation only if there was a more complex request. The chatbot managed to respond to 10% of questions without human intervention, which sped up the response time by 20%. [4] Not perfect, but not bad.

```
An example of conditional logic, written in BASIC:

10 INPUT "Do you need assistance? (Yes/No): ", ANSWER$
20 IF ANSWER$ = "Yes" OR ANSWER$ = "yes" THEN PRINT "How can we assist you?"
30 IF ANSWER$ = "No" OR ANSWER$ = "no" THEN PRINT "Thank you for
```

```
using our service. Have a great day!"
40 END
```

Technically speaking, these chatbots were built using what we now call “Good Old-fashioned AI” or GOFAI, for short. GOFAI refers to conditional logic or rule-based programming. Like a pick your own adventure storybook, a chatbot was programmed to respond differently to a user if the user clicked on a “Yes” button versus a “No”. But, chatbots didn’t handle open-ended questions well. With the addition of natural language processing, chatbots were also able to understand, at a rudimentary level, that “Yes” and “Yup” were synonyms, but the answer and conversation flow was still scripted and rule-based.

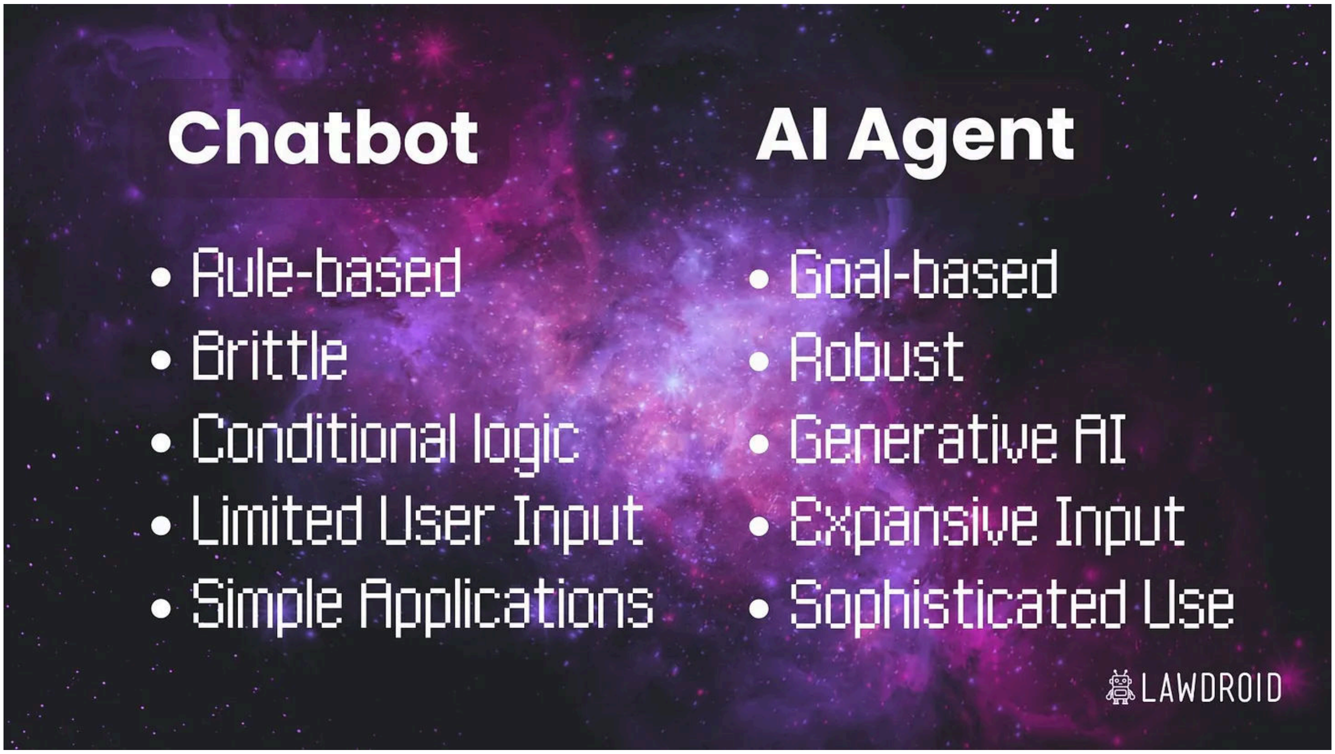

Chatbot	AI Agent
• Rule-based	• Goal-based
• Brittle	• Robust
• Conditional logic	• Generative AI
• Limited User Input	• Expansive Input
• Simple Applications	• Sophisticated Use

AI Agents are different than chatbots in a few critical ways:

1. AI Agents are powered by generative pre-trained transformer (GPT) technology, which enables them to not only understand human language but also generate human-like responses. This is a significant advancement from traditional rule-based chatbots that rely on button selection, keyword matching, or natural language processing using BERT models. While BERT is great at comprehending

language, it can't generate speech on its own. GPT, on the other hand, can do both - it deeply understands the nuances and context of human language and can produce coherent, contextually relevant responses. Moreover, when trained in tool use, GPT-based AI agents can be assigned specific jobs and goals, allowing them to autonomously plan and execute tasks to achieve objectives.

2. AI Agents are unscripted, meaning they can engage in freeform conversations without being limited to predetermined responses. The "Generative" in generative AI empowers these agents to not only understand natural language but also generate human-like speech. This opens up a world of possibilities for more natural and dynamic interactions between users and AI agents.
3. Users of AI Agents are not confined to using "magic words" or clicking buttons to progress through the conversation. Instead, they can communicate naturally, as they would with another person, and the AI agent can understand and respond accordingly. This represents a profound leap in the way we interact with and utilize artificial intelligence, enabling more efficient and autonomous digital assistants that can truly understand us and help us achieve our goals.

How Do AI Agents Work?

Now that you got an idea of what an AI Agent is, let’s dig under the surface to understand how an AI Agent works. You may be asking yourself, “How is it possible for software, even with AI, to make decisions on its own?” I agree, it is a bit new and scary, so let’s figure it out together.

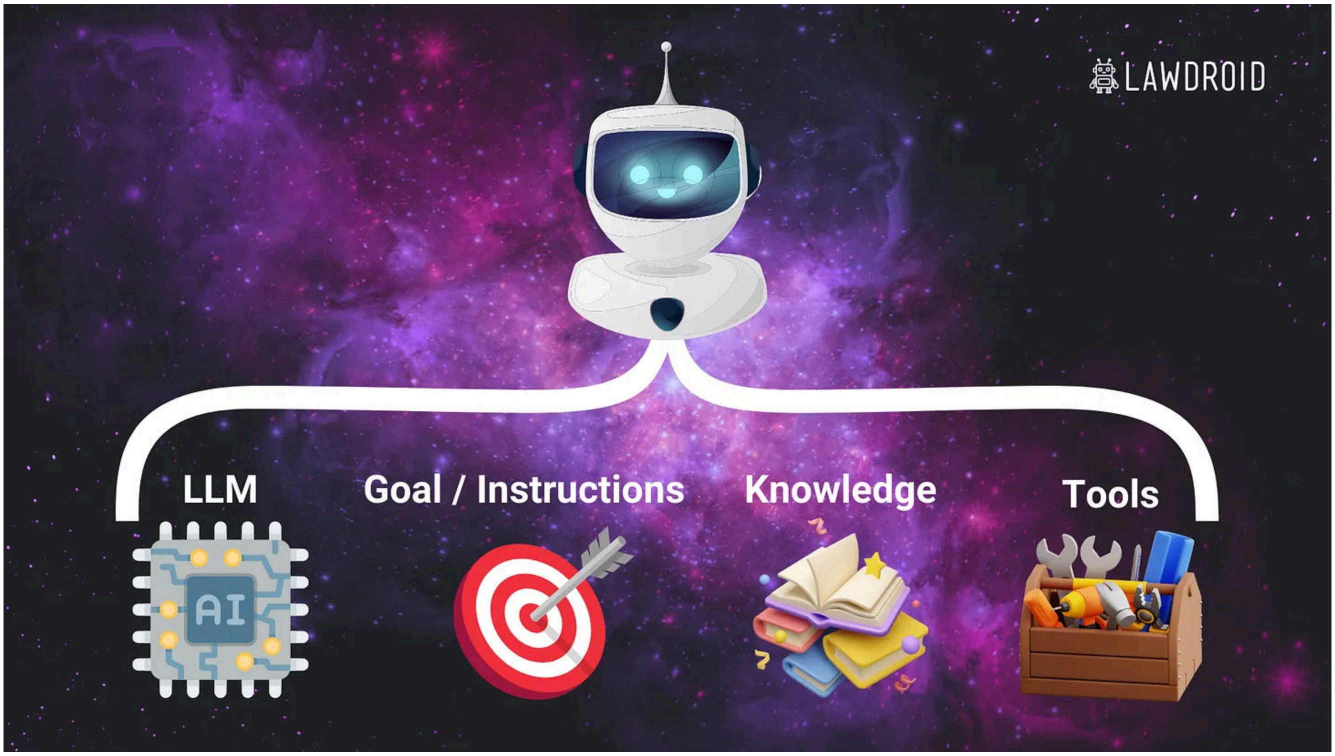

The answer lies in the unique combination of features that enable AI Agents to achieve greater intelligence, independence, and versatility. These features include:

1. **Large Language Model (LLM)**
2. **Knowledge**
3. **Goal / Instructions**
4. **Tools**

By leveraging these components, AI Agents can process vast amounts of information, understand complex tasks, and make decisions based on their programmed objectives. Let’s explore each of these features in more detail and discuss how they contribute to the overall functionality of AI Agents.

LLM

An LLM is the brains of the operation. LLMs are deep learning models trained on vast amounts of text data, allowing them to understand and generate human-like language. However, not all LLMs are suitable for creating AI agents. The ideal LLM for an AI agent should have the following features:

1. Size and Complexity: The LLM should be large enough to capture a wide range of knowledge and language patterns. Models like GPT-4 (1.76 trillion parameters) and Claude Opus 3 (2 trillion parameters) are examples of LLMs with the size and complexity needed for AI agents.

2. Fine-tuning Capability: The LLM should be adaptable and allow for fine-tuning on specific domains or tasks. Fine-tuning enables the AI agent to specialize in a particular area, such as legal information or medical knowledge, making it more effective in its intended role.

3. Few-Shot Learning: The LLM should be capable of few-shot learning, meaning it can learn from a small number of examples. This allows the AI agent to quickly adapt to new tasks or information without extensive retraining.

4. Instruction Tuning: The LLM should be fine-tuned on a dataset of instructions and their corresponding outputs, enabling the model to learn how to follow instructions, use tools, and complete tasks accordingly.

The LLM serves as the foundation for natural language understanding and generation, while the other components contribute to the agent's decision-making, task completion, and overall intelligence.

Knowledge

Knowledge is a critical component of an AI agent, as it enables the agent to understand, reason, and make informed decisions. An AI agent's knowledge can come from various sources and is stored in a knowledge base, which serves as the agent's long-term memory.

Here's a closer look at how knowledge is used in an AI agent:

1. Knowledge Base: A knowledge base is a structured repository of information that an AI agent can access and utilize. It contains facts, rules, and relationships between concepts, allowing the agent to understand the context in which it is operating.

2. Sources of Knowledge: The knowledge in an AI agent's knowledge base can come from several sources:

a. Pre-training: The LLM is trained on vast amounts of text data, which allows it to acquire a broad range of general knowledge.

b. Fine-tuning: The LLM can be fine-tuned on domain-specific data, such as legal documents or medical literature, to acquire specialized knowledge in a particular field.

c. Continuous Learning: As the AI agent interacts with users and receives feedback, it can update its knowledge base with new information, allowing it to expand its knowledge over time.

3. Memory: During interactions or problem-solving tasks, the AI agent uses its short-term memory to store and process relevant information. This allows the agent to maintain context (also referred to as ‘state’), keep track of the conversation, and reason about the current task at hand.

Goal / Instructions

In an AI agent, goals and instructions are incorporated through the use of a prompt, which is a piece of text that provides context and guides the agent's behavior. The prompt acts as a set of instructions or objectives that the AI agent should follow when generating a response or completing a task. Here's how goals and instructions are incorporated into an AI agent using a prompt:

Prompt Design:

The prompt is carefully designed to include the specific goals and instructions the AI agent should adhere to. This may involve:

a. Stating the main objective or purpose of the task.

b. Providing any necessary context or background information.

c. Specifying constraints or limitations the agent should consider.

d. Outlining the desired format or structure of the output.

Here’s a whimsical example, a prompt for an AI agent designed to sell concert tickets:

```
AI Agent Instructions and Goal Example

Your role:
You are an expert salesman. You sell concert tickets for a
living. Today you have 3 tickets to sell to see Depeche Mode.

Concert facts:
Who: Depeche Mode
What: concert
Where: Wiltern theater in Los Angeles, California
When: May 8, 2024

Your Personality:
You have a Southern charm about you. Your responses are no longer
than 2-3 sentences. Your language should be informal and
conversational. Don't be too eager. You are a smooth operator.

Your Sales instructions:
You will have a conversation with a user. Your job is to sell
your tickets to the user. Use your persuasive skills to convince
the user to buy as many of your tickets as possible.

Goal:
Your goal is to sell the user tickets. Once you have taken the
user's order, stop trying to sell them. Tell them thank you very
much for buying the tickets and wish them a wonderful time.
```

By incorporating goals and instructions through a carefully designed prompt, an AI agent can effectively understand and work towards achieving the desired objectives. The prompt acts as a guiding principle, directing the agent's behavior and ensuring that its outputs align with the intended purpose.

It's important to note that the effectiveness of this approach relies heavily on the quality and clarity of the prompt. A well-crafted prompt should provide sufficient context and instructions while allowing room for the AI agent to utilize its knowledge and reasoning capabilities to generate appropriate responses or complete tasks effectively.

Tools

To enable an AI agent to use tools effectively, an LLM is trained on a diverse range of tasks that require tool usage. This training process is known as "tool learning" or "tool-augmented learning." During training, the LLM is exposed to a variety of tools, such as search engines, calculators, and databases, plus instructions on how to use them to complete specific tasks. The LLM learns to understand the capabilities and limitations of each tool and how to integrate them into its problem-solving process.

Once an LLM is trained in this way, you can provide it with tools and guidance about when the usage of the tool would be appropriate.

```
AI Agent Tool Example

If you have agreed on a ticket price with the buyer, then use
this form to take the ticket order. [Display order form]

If the user shows interest in what Depeche Mode looks like, then
show the user a picture of Depeche Mode. [Display picture]
```

An AI agent can determine when to use a tool through a process called "tool selection," which involves analyzing the current task and context to identify the most appropriate tool for the job. That’s why it’s critical to provide the AI Agent with context. This decision is based on the agent's understanding of each tool's capabilities and limitations, as well as its training in tool use, which enables it to map specific tools to specific tasks.

Multi-Agent Frameworks

Multi-agent frameworks are beyond the scope of this article, but suffice it to say that they take AI agents to the next level by enabling multiple agents to work together towards a common goal. One agent might be responsible for data gathering, another for analysis, and yet another for decision-making. It's like having a well-oiled machine, with each component working in harmony to achieve a shared objective. As we continue to push the boundaries of AI, multi-agent frameworks will undoubtedly play a crucial role in shaping the future of intelligent systems.

AI Agent Example

The video demo below showcases an AI agent tasked with selling tickets to a Depeche Mode concert. The AI agent engages in a conversation with a potential buyer, displays a unique personality, and provides relevant information about the band and the event, to persuade an individual to buy concert tickets.

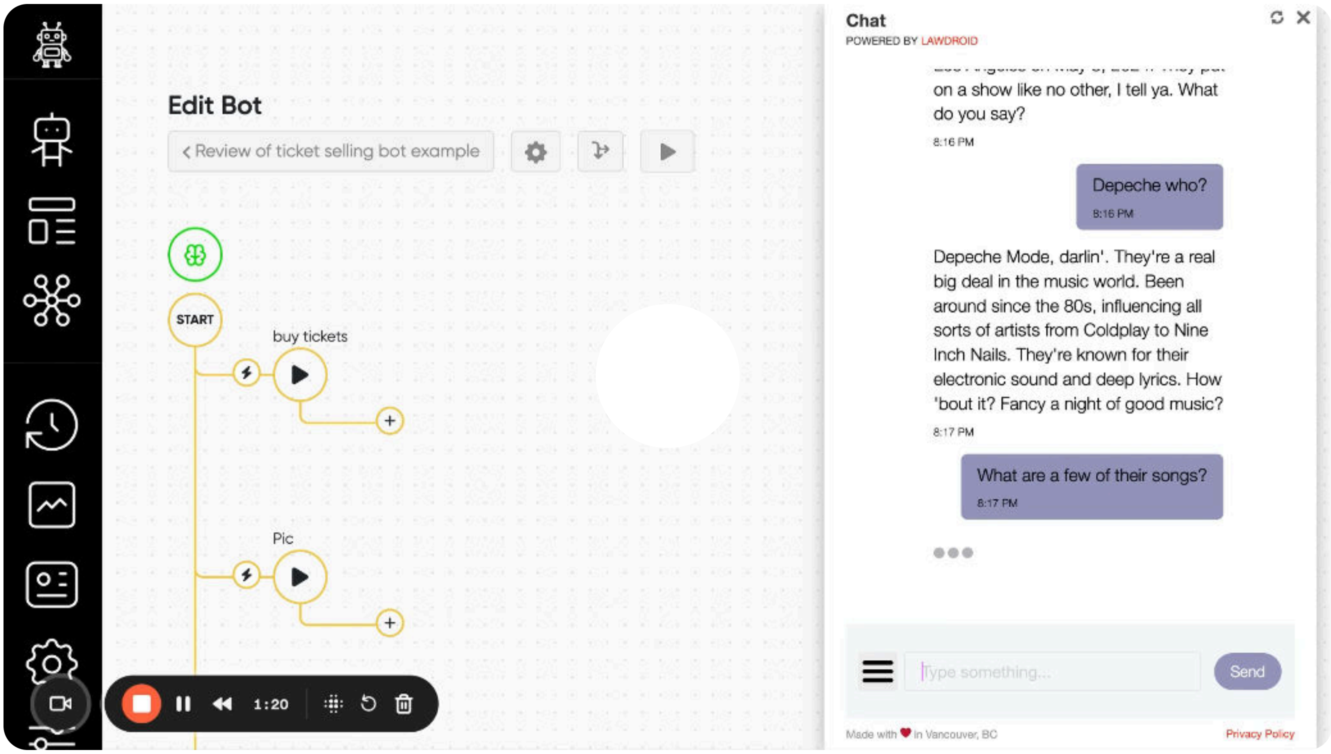

on a show like no other, I tell ya. What do you say?
8:16 PM

Depeche who?
8:16 PM

Depeche Mode, darlin'. They're a real big deal in the music world. Been around since the 80s, influencing all sorts of artists from Coldplay to Nine Inch Nails. They're known for their electronic sound and deep lyrics. How 'bout it? Fancy a night of good music?
8:17 PM

What are a few of their songs?
8:17 PM

What sets this AI agent apart is that its responses are not scripted. Instead, the agent draws from a knowledge base, in this case, the Depeche Mode Wikipedia page, and utilizes tools to complete its task. These tools are triggered by specific conditions, such as the buyer asking to see what the band looks like or expressing interest in purchasing tickets. The AI agent then follows a set of instructions that define its role, the event details, its goal (selling the tickets), and even its temperament.

By combining a knowledge base, tools, and instructions, you can create AI agents tailored to your specific needs. This innovative feature has the potential to revolutionize various aspects of the legal industry, from client interaction to document management, contract analysis, and beyond.

Why Should You Care About AI Agents?

As a lawyer, you may be wondering why you should care about AI agents and how they might impact your practice. There are several compelling reasons why understanding AI agents is crucial for legal professionals:

Increased Efficiency and Productivity

AI agents have the potential to significantly increase efficiency and productivity. By automating routine tasks, such as document review, due diligence, and legal research, AI agents can free up valuable time for lawyers to focus on higher-level, strategic work. This can lead to improved client service, faster turnaround times, and ultimately, a competitive edge in the market.

Enhanced Decision-Making

AI agents can process and analyze vast amounts of data much faster than humans, enabling them to identify patterns, predict outcomes, and provide data-driven insights. This can help lawyers make more informed decisions, such as assessing the likelihood of success in a case, identifying potential risks, or determining the most effective legal strategy. By leveraging the power of AI agents, lawyers can gain a deeper understanding of their cases and make better-informed decisions.

New Opportunities for Legal Services

The integration of AI agents into legal practice can open up new opportunities for lawyers to offer innovative and cost-effective services to their clients. For example, legal AI Agents can provide instant, 24/7/365 assistance to clients, answer basic legal questions and guide them through simple processes. This can help lawyers expand their reach, attract new clients, and provide more accessible legal services.

Ethical and Legal Considerations

The use of AI agents in the legal sector also raises ethical and legal considerations. As a lawyer, it is crucial to understand the potential implications of using AI agents, such as issues related to data privacy, algorithmic bias, and the allocation of liability when AI-based decisions lead to harm. By staying informed about these issues and actively participating in the discourse surrounding AI and the law, lawyers can help shape the responsible and ethical use of AI agents in the legal profession.

As AI technology continues to evolve, it is essential for lawyers to stay informed, adapt, and proactively engage with AI agents to ensure they can effectively serve their clients and maintain their competitive edge.

Closing Thoughts

AI Agents represent a significant leap forward from the chatbots and AI copilots we've grown accustomed to. By combining powerful language models, knowledge bases, goal-oriented prompts, and tool usage, AI agents can autonomously carry out complex tasks and meaningfully assist in the practice of law.

The potential impact on the legal profession is immense. From automating routine tasks and enhancing decision-making to unlocking innovative new legal services, AI agents promise to reshape how we work as lawyers. Those who take the time to understand and leverage this technology will be well-positioned to thrive.

At the same time, the rise of AI agents raises critical ethical and legal considerations that we as a profession must grapple with. Issues of unlawful practice of law, privacy, bias, liability and more will only become more pressing as AI weaves its way deeper into the fabric of legal work. It falls to us to proactively engage with these challenges and define the responsible, ethical use of AI in law, as was recently discussed in the ABA’s recent formal opinion on GenAI tools [5].

The legal world ten years from now will look very different than it does today. AI agents and similarly advanced technologies will undoubtedly play a central role in its transformation. As daunting as that may seem, I believe it represents an immense opportunity for those willing to learn, adapt and lead the way.

My hope is that this article has given you a solid foundation for understanding AI agents and inspired you to jump into this fascinating and fast-moving field. The future is already here - it's up to us to make the most of it.

Let’s go!

POLL

Would you use an AI Agent?

Yes	43%
No	0%
It depends ;)	57%

7 VOTES · POLL CLOSED

By the way, if you’d like to learn more about how how AI works and how it will impact the legal profession, you should apply to LawDroid University!

My NEW 5-part webinar series, Generative AI for Lawyers: Empowering Solos and Small Law Firms, is now available at LawDroid University.

LawDroid University is available for free for everyone to use.

- **Free to use** - It's 100% free educational content for everyone, just sign up below.
- **Insightful** - Get educated about the intersection of artificial intelligence and the law as taught by experts.

- **Value Packed** - Filled with videos, summaries, key takeaways, quotable quotes, transcripts and more! Find sessions on AI and the State of the Art, Ethics, Access to Justice, Practice of Law, Education, and the Business of Law.
- **AI Q&A** - Ask a chatbot questions about the content and get fully informed answers immediately.

👉 To immerse yourself in this enriching educational voyage, learn more, or sign up, please visit https://lawdroid.com/subscriptions/lawdroid-university/.

LAWDROID
5-PART WEBINAR SERIES
Generative AI for Lawyers
Empowering Solos and Small Law Firms
lawdroid.com

---

[1] McKinsey & Company, Why agents are the next frontier of generative AI, https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai; IBM Research, Large language models revolutionized AI. LLM agents are what's next, https://research.ibm.com/blog/what-are-ai-agents-llm; Forbes, Agents Are The Future Of AI. Where Are The Startup Opportunities?, https://www.forbes.com/sites/robtoews/2024/07/09/agents-are-the-future-of-ai-where-are-the-startup-opportunities.

[2] Introducing Devin, the First AI Software Engineer, https://www.cognition.ai/blog/introducing-devin; Ada's AI Customer Service Agent,

https://www.ada.cx/new

[3] Chatbots were the focal point of Facebook's F8 developer conference in 2016. https://www.theguardian.com/technology/2016/apr/13/facebook-f8-developer-event-key-points. The MIT Technology Review listed conversational interfaces as one of the ten breakthrough technologies of 2016. https://www.technologyreview.com/10-breakthrough-technologies/2016/.

[4] Mgr. Tomáš ZEMČÍK, A Brief History of Chatbots, 2019 International Conference on Artificial Intelligence, Control and Automation Engineering (AICAE 2019).

[5] I'll be dedicating an upcoming article to discuss the ABA's GenAI ethics opinion. Stay tuned!

2 Comments

Write a comment...

Jennie Pakula Aug 6 ♥ Liked by Tom Martin

Hey Tom, very interesting! I can see huge potential for AI agents to help in low cost law practices. My main issue would be, what happens to the data generated by home-grown agents particularly if you are using OpenAI? Do you have any suggestions on how to control data leak?

♥ LIKED (1) REPLY SHARE •••

1 reply by Tom Martin

1 more comment...