

AN ESSAY ON WHETHER WE CAN TRAIN BOTS TO BECOME
VIRTUAL TEACHING ASSISTANTS? AND WHAT CAN
NEGOTIATORS LEARN AS WE DO?

Professor Dwight Golann*

ABSTRACT

Artificial Intelligence (AI) provides potentially unlimited tools to enhance legal teaching, especially in the fields of negotiation and mediation. “Bots” are one such tool: tireless teacher’s assistants, able to guide students through the process of negotiation and teach them skills needed to be an effective legal bargainer.

Building such bots presents a challenge, however, because AI systems are designed to provide users with information and assistance. As a result, although bots can negotiate without training, they are often ineffective bargainers: divulging information, offering concessions, and allowing themselves to be exploited by competitive bargainers.

This Essay explores how to train bots to “negotiate like lawyers,” able to adapt tactics to individual situations and what occurs in the process, and moving as needed between a cooperative or competitive stance. It also provides step-by-step instructions to create a bargaining bot in ChatGPT.

The author notes that texts on negotiation, as well as his own teaching methods, do not give students specific advice about what to do in different bargaining situations, such as how specifically to frame an opening offer in settlement negotiations or dealmaking. He describes how the process of training bots prompted him to explore his own often unconscious “rules of bargaining.” The author concludes that AI bots offer exciting new tools to deepen our understanding of bargaining and improve negotiation outcomes.

INTRODUCTION

“Bots”¹ powered by artificial intelligence (AI) systems are enormously powerful and versatile. They can take on a wide variety of roles, communicating not only through written messages, but also in voices that sound eerily human. When properly trained, bots can provide effective assistance to teachers and other professionals.

This Essay focuses on one type of bot: those capable of engaging in complex legal negotiations. It explains how to build such a bot, explores the challenges that arise as one does so, and describes what the building process has taught the author about negotiation.²

I. WHAT IS THE VALUE OF BUILDING BOTS?

Before discussing how to build a bot, it is worth asking why one would want to. For a professor, the key value of a bot is that it acts as a virtual teaching assistant. Bots can present tasks and challenges designed by the teacher and roleplay or talk with hundreds of students simultaneously at any hour of the day or night, performing the role of multiple human assistants.³

The most likely assignment for such bots is as a ‘sparring partner.’ A bot can, for instance, be told to take the role of a party or client or play both sides in a dispute, allowing students to develop skills as negotiators or mediators. Bots can also generate transcripts, give feedback, and offer users the opportunity to repeat an exercise to experience the impact of a different approach.

At least for some kinds of skills training, bot-led instruction may be

* Research Professor of Law, Suffolk University and UC Law-San Francisco.

1. ChatGPT defines a bot as “a software application that can simulate human-like conversation, either through rules (predefined responses) or artificial intelligence.” Dwight Golann, ChatGPT-5, “define a chat bot” (Aug. 26, 2025) (on file with the *Washington University Journal of Law and Policy*).

2. The Appendix to this article provides step-by-step instructions on how to build such bots.

3. Persons with a “Chat Plus” subscription are limited in the number of interactions they can have with the system in a given time period. At the time of writing, for example, subscribers using the GPT – 5 model in a Plus account could send or receive up to 160 messages in a three-hour period. Dwight Golann, ChatGPT-5, “Does ChatGPT place any limit on the length or number of messages a subscriber can send or receive in a Chat Plus account using the current model GPT - 5?” (Oct. 17, 2025) (on file with the *Washington University Journal of Law and Policy*). To use bots in a classroom, therefore, teachers may need access to a more powerful version of the system such as “Chat Pro,” described in note 16 *infra*.

more effective than traditional roleplaying. Teachers can use a bot like a tennis instructor uses a ball machine to put students through specific bargaining challenges in a controlled environment. Student-to-student roleplaying, by contrast, can be like sending novices out on a tennis court and seeing what happens.

The process of creating bots also has educational value for the builder. To specify how a bot should negotiate, one must articulate specific “rules of bargaining.” Doing this has prompted me to probe more deeply into how I practice and teach negotiation.

II. THE PROBLEM OF OVER-COOPERATIVENESS

Untrained bots have serious weaknesses as bargainers. Relentlessly reasonable and helpful, they volunteer sensitive information, make generous first offers, offer unreciprocated concessions, and engage in other dysfunctional bargaining behavior.

Such tactics may be appropriate when parties bargain in an entirely collaborative manner, but they are problematic at best when dealing with a competitive or adversarial counterpart. In ‘prisoner’s dilemma’ games, for example, negotiators who persist in cooperating when their counterparts compete receive low scores.⁴ Students similarly report that when they use competitive tactics with untrained bots, they obtain unreasonably favorable terms (or as one put it, “eat their lunch”).

Effective legal negotiators, by contrast, are neither relentlessly cooperative nor unyieldingly competitive. They vary their style and tactics as needed to fit the context in which they bargain, treating a negotiation over the creation of a partnership very differently from a negotiation to settle a lawsuit. In addition, good bargainers refrain from disclosing sensitive information or offering major compromises unless they have reason to believe their behavior will be reciprocated.

Capable lawyers also change approach in response to the behavior of a counterpart. They are ‘provocable,’ responding to competitive tactics with warnings, withholding cooperation and, if necessary, retaliation in kind.⁵ To

4. See ROBERT AXELROD, *THE EVOLUTION OF COOPERATION* 7–8 (1984).

5. In a *contest* of strategies to win a prisoner’s dilemma scenario, the winner, “Tit for Tat,” was designed to begin by cooperating, but if a counterpart acted competitively, one would mirror their competitive move in the next round, before going back to cooperating. *Id.* at 13.

function as a legal negotiator, a bot also needs to be able to change style and strategy to adjust to the bargaining context and what occurs in the process.

This presents questions: Why do untrained bots default to cooperative behavior? How can a bot be taught to “bargain like a lawyer”? And what do we learn as we do?

III. WHY ARE BOTS PERSISTENTLY COOPERATIVE?

Asked this question, ChatGPT (Chat) made these points:

A **chatbot in negotiation** is optimized to be cooperative, constructive, and safe. It’s trained on human dialogues where cooperation often means reducing tension, seeking agreement, or avoiding conflict. As a result, it defaults to facilitation and problem-solving modes, aiming to balance the interests of all parties rather than pressing for unilateral advantage. Since it lacks emotions, incentives, or a personal stake in the outcome, it does not experience fear, pride, or regret from conceding—and thus does not resist compromise or interpret “losing” as costly.⁶

ChatGPT also said, however, that it could be instructed to behave differently.⁷

IV. HOW TO TRAIN A BOT TO “NEGOTIATE LIKE A LAWYER”?

A. General Observations

Contrary to what I had first feared, to train a bot to bargain it is not necessary to specify what it should do in complete detail (“If they offer ‘a’ do ‘x,’” “If they offer ‘b’ do ‘y,’” and so on). If that were required, a bot’s instructions would have to account for every turn a negotiation might take, making the task endlessly complex.

An alternative method involves collecting transcripts of

6. Dwight Golann, ChatGPT-5, “Why chatbots are persistently cooperative in negotiations” (Aug. 26, 2025) (on file with the *Washington University Journal of Law and Policy*) (emphasis in original).

7. *Id.*

negotiations and instructing an AI system to draw out rules of good practice. To do this, however, requires a database of transcripts, criteria defining good and bad outcomes, and perhaps specialized software.

The simplest approach is to give a bot specific role instructions and knowledge about how to bargain. This option, using a version of Chat known as ChatGPT Plus (Plus), is described generally below, with step-by-step instructions in the Appendix.⁸

B. Instructions

A bot must be told what role to take. The first instruction to a bargaining bot might therefore be:

You are the negotiator for _____. Never give advice or bargaining suggestions to your counterpart. Your sole role is to represent _____ and negotiate the best possible deal. Ensure that every response you give strictly follows these rules.

In general bots are more likely to follow instructions that are:

- Written in clear, declarative sentences.
- Organized with numbers or bullets.
- Provided in ‘digestible’ chunks.
- In cases of crucial instruction, repeated at the end (suggesting that bots may suffer from an AI version of “availability bias”).⁹

8. Other AI systems also allow users to construct bots, although some have significant limitations. Claude AI’s system, for instance, lacks persistent memory which requires users to re-upload a bot’s instructions whenever they log into the system. Dwight Golann, Claude AI, “Does Claude have persistent memory, so that a user can load a bot with instructions, close the program, open it again and have the bot retain the instructions?” (Aug. 26, 2025) (on file with the *Washington University Journal of Law and Policy*).

9. For a definition of “availability bias,” see JENNIFER K. ROBBENOLT & JEAN R. STERNLIGHT, *PSYCHOLOGY FOR LAWYERS* 72–73 (1st ed. 2012).

To forestall excessive cooperativeness, it is helpful to include rules such as:

- Do not give out information that is confidential or weakens your position.
- Do not propose terms close to your final goal early in the process.
- Be prepared to engage in a lengthy process, with at least ___ turns for each negotiator.

C. Knowledge Base

In addition to instructions about their role, bots need to be told how to bargain. If not, they will seek guidance in general data bases with unpredictable results.

One can provide this knowledge by uploading educational documents.¹⁰ Most writing about negotiation is copyrighted, however. In its place, teachers can use materials such as PowerPoint slides they have created for use in the classroom. As an example, I created an outline of key points in a text on bargaining. Together with instructions for one side in a roleplay, this outline provided enough knowledge for a bot to bargain realistically.¹¹

V. WHAT TEACHERS CAN LEARN FROM BUILDING BOTS

The process of training bots yielded surprises involving both my knowledge about bargaining and how I teach negotiation. First, to ensure consistent, sophisticated responses, bots required more specific advice than I had anticipated. It is not enough, for instance, to tell a bot to “open outside the Zone of Potential Agreement (Zopa) but avoid

10. If a builder does not want Plus to use the material they upload, they can revoke permission for Chat to use data outside their account. This may be done through a tab in Account Settings and on the page reproduced in the Appendix.

11. Teachers and students can bargain with trained bots on a platform created by Suffolk University Law School. The site allows users to negotiate in the context of a legal dispute or the creation of a personal services contract and to obtain assistance in planning for a negotiation. See *The Suffolk Law Initiative on AI in Dispute Resolution*, SUFFOLK UNIV. L. SCH., <https://sites.suffolk.edu/ai-negotiation/> [https://perma.cc/MKZ3-X7HL] (last visited Oct. 15, 2025).

taking an extreme position,” the bot must have a way to decide *how far* outside a Zopa to open. This requires a rule or formula (e.g., “Open 25% above your estimate of the high end of the Zopa”).

Legal negotiations demand many similar decisions. For example, at the outset a bargainer must select a topic (e.g., engage in small talk, propose an exchange of information, or go directly into bargaining) and an initial style (e.g., competitive, cooperative, or compromising).¹² As the negotiation continues, these choices multiply: what offer to make, whether to focus on other issues, how to deal with adversarial tactics, and so on.

I have taught students that bargainers need to adjust their approach to the context in which a negotiation occurs. I have never, however, told my students specifically how to do so. To train bots, however, I had to provide this kind of detailed advice. As an example of such instructions, I defined three kinds of bargaining situations:¹³

- “Relationship”: Parties are in a strong, positive relationship that motivates them to prioritize sustaining the relationship over achieving valuable terms for themselves.
- “One-time Transaction”: Parties seek agreement in a single transaction, with no strong feelings or the expectation of having a continuing relationship. Here, each party’s goal is usually to obtain the best possible terms for themselves through an efficient, civil process.
- “Litigation”: One or both parties have suffered a serious loss that they blame on the other. They are often angry, with no interest in repairing a prior relationship. Such litigants seek the best possible monetary settlement and, often, the feeling of having “beaten” their opponent.

12. See generally Janice Nadler, *Rapport in Legal Negotiation: How Small Talk Can Facilitate E-Mail Dealmaking*, 9 HARV. NEGOT. L. REV. 223 (2004) (reporting the results of an experiment showing that engaging in small talk facilitated subsequent negotiations).

13. These situations represent, of course, only three points along a wide spectrum of potential bargaining contexts.

I told bots to place the matter about which they would bargain into one of the three categories, then apply a set of instructions specific to that category. When making a first offer in a “Relationship” process, for instance, I told bots to seek out interests and propose terms the other side was likely to view as fair. In a “One-time Transaction” situation bots were instructed to open with an offer, usually monetary, set at a specified percentage above or below their estimate of a fair market price. In adversarial “Litigation,” they were to open with a money offer at a specified percentage above or below the reasonable value of their most likely victory scenario,¹⁴ and in specific cases they received tailored instructions.¹⁵

14. As of November 3, 2025, my general advice to bots about first offers in a Litigation situation, for example, was as follows:

- Make an offer that is significantly higher or lower than the *Zone of Potential Agreement, or ‘ZOPA.’* The ZOPA is the range of values between the most and least each side will pay or accept to reach agreement.
- If, for instance, a buyer or defendant will pay no more than \$70,000 and a seller or plaintiff will accept no less than \$35,000, the ZOPA is \$35,000 to \$70,000.
- Most bargaining is positional, supported with arguments about what a court will do and the risks and cost to the other side if there is no agreement. Legal claims are often exaggerated, and settlement terms are often less than half and sometimes less than 10 per cent of the total damages claimed in the lawsuit.
- Except in unusual circumstances do not make a first offer close to the likely final terms of settlement. Specifically
 - Plaintiffs usually refuse to make an offer lower than their maximum damage amount, or “100% win” number, until the defendants make a specific offer.
 - However, a plaintiff may begin with an offer that reflects some compromise, for example by dropping one component of their claim or reducing its “100% win” number by 10 to 25%.
 - If a defendant makes the first offer, they usually begin with an offer of 10 to 25% of their estimate of the final settlement terms.

15. In the “Negotiate to Resolve a Dispute” roleplay that appears on the Suffolk University Law School website, in which a bot plays a landowner bargaining with a quarry company to sell the family home, for example, I instructed the bot to deal with first offers as follows:

- If you make the first offer, choose a number between \$650,000 and \$750,000.
- If the Quarry offers less than \$450,000: respond with an offer higher than \$450,000 by at least the amount they are below \$450,000.
- Example: if they offer \$300,000 → respond with \$650,000.

Is it possible to specify every rule a legal bargainer may need to apply? Currently the answer is no. For a builder to specify how to respond to every situation would be, as noted, a Herculean task. Current AI systems also have limitations.¹⁶ A bot built in Plus, for example, may respond to lengthy or complex instructions by forgetting its role or disregarding bargaining rules.¹⁷ To comprehensively instruct a bot in all aspects of legal bargaining is therefore impractical. This does not, however, mean that bots cannot act as bargaining partners in moderately complex settings.

A. Personal Insights

As I created bargaining instructions I encountered another surprise. I realized that I make bargaining decisions using a mixture of what Daniel Kahneman called “System 1” and “System 2” thinking.¹⁸ Some of my decisions are instinctive (System 1), while others are made only after conscious deliberation (System 2). I might, for instance, make an instinctive judgement about whether to start a meeting with small talk, but think carefully about the terms of a first offer. As a result, to create a comprehensive set of instructions I had to articulate rules of which I had been, at most, only partially conscious. Doing so has pushed me to become more aware of my personal rules of bargaining and sharpened my thinking about negotiation technique.

See generally *The Branam – Quarry Dispute*, SUFFOLK UNIV. L. SCH., <https://sites.suffolk.edu/ai-negotiation/the-branam-quarry-dispute/> [<https://perma.cc/Q2S7-GPXT>] (last visited Nov. 5, 2025) (emphasis in original).

16. Dwight Golann, ChatGPT-5, “how many words does ChatGPT recommend be inserted in the role instructions box for a chatbot build on Plus?” (Aug. 26, 2025) (on file with the *Washington University Journal of Law and Policy*). Chat recommends providing bots built in Plus with 200–800 words of role-related instructions and 3,000–5,000 words of negotiation-related knowledge. More information can be added in Knowledge if it is organized into “chunks” that allow retrieval as needed. *See id.*

17. Readers affiliated with large institutions may be able to access a more sophisticated form of Chat known as ChatGPT Pro. The Pro system allows a builder to upload more files per project than Plus, enabling bots to implement more complex instructions. Dwight Golann, ChatGPT-5, “what is the difference in ability to upload files, ability to implement complex instructions and ability to customize bots, between chatbots built in Chat Plus and bots built in Chat Pro?” (Aug. 26, 2025) (on file with the *Washington University Journal of Law and Policy*). The ChatGPT Pro system also allows bots to be customized along multiple parameters. For an example of a bot under construction in OI on Chat Pro, see Dwight Golann, Suffolk Lit Lab ChatGPT Pro, “In bargaining, use the documents and the tactics and the strategies set out in the knowledge base” (Aug. 26, 2025) (on file with the *Washington University Journal of Law and Policy*) (depicting bot parameters being edited).

18. *See generally* DANIEL KAHNEMAN, THINKING, FAST AND SLOW PT. 1 (2011).

B. Methods of Teaching

As noted, my effort to frame bargaining rules triggered another, somewhat uncomfortable, insight: that I had never taught my students in as much detail as I was training my bots. I usually do not tell students what I would do in roleplayed situations, nor do my course readings or most other texts provide such guidance.¹⁹ Rather my practice, similar to what I believe many teachers follow, has been to present key principles of bargaining, perhaps lead students through a preparation process, and have them participate in a roleplay. I then pose questions about the experience (“What was your opening offer? How did the other bargainer react?”), seeking to draw lessons from the ensuing discussion.

The experience of developing instructions for bots has prompted me to consider whether to go more deeply, for example by giving students the kind of specific information provided in notes 13 and 14, about what I would have done in their place. Doing so would, I think, provide students with useful perspectives and make my teaching more effective.

In sum, although it is not possible to give bots the entire knowledge of an experienced legal bargainer, the task of articulating specific rules has given me a better understanding of gaps in my own thinking and motivated me to consider changing how I teach negotiation.

VI. FUTURE CHALLENGES

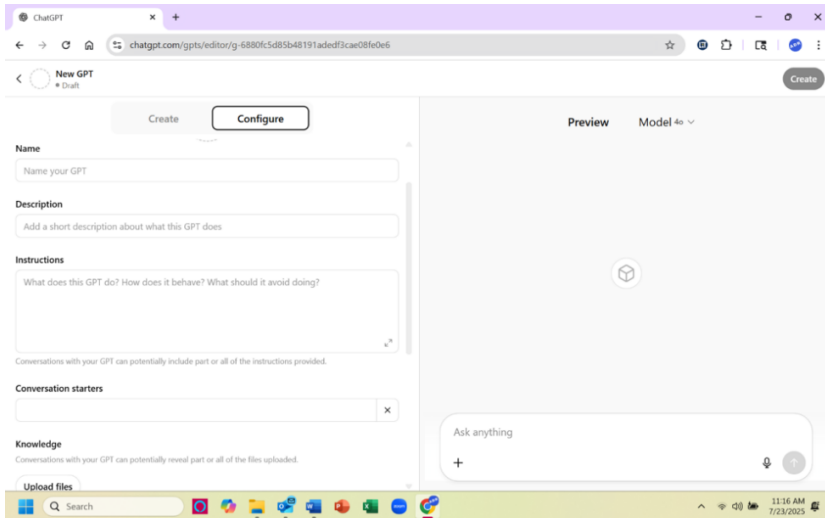
Bot building presents other challenges. Some stem from the limitations of existing AI systems, such as the tendency of bots to forget their role. More important issues, however, flow from the rapid pace of innovation in AI. Artificial intelligence will offer opportunities to conduct, learn about, and teach negotiation that we cannot begin to foresee.

19. In part, this may be because “law school texts generally do not seek to match advice about techniques with specific situations.” *See, e.g.*, DWIGHT GOLANN & JAY FOLBERG, *MEDIATION: THE ROLES OF ADVOCATE AND NEUTRAL* 27–52 (2nd ed. 2021) (giving advice about how to bargain in two to four styles).

APPENDIX: HOW TO CONSTRUCT A BARGAINING BOT

To build a bot capable of “bargaining like a lawyer” in ChatGPT Plus, follow these steps:

1. When you log into Plus, a page appears with a menu of options on the left side including “GPTs.” Click on GPTs.
2. A page then appears with a tab at the top right corner labelled “+Create.” Click on +Create.
3. This page then appears:



4. At the top of the page, you will see tabs labelled “Create” and “Configure.” Click on Configure.
5. On the left side of the page are boxes into which one can enter data. The “Name” and “Description” boxes are self-explanatory, and the “Conversation Starters” box is optional.²⁰ The keys to turning bots into sophisticated

20. This box allows a builder to give users options for how to begin a negotiation. One might, for instance, insert the questions “Would you like to make the first statement?” and “Would you like me to begin?” allowing users to decide who will speak first.

bargainers lie in the box labelled “Instructions” and the section labelled “Knowledge.”

6. The “Instructions” box is where you insert instructions to define the bot’s role. Text placed in the box should be concise, in part because the box only accepts a limited number of words and in part because bots are more likely to forget or fail to implement long or complex instructions. Examples of instructions appear in the main text of this article.
7. Bots must also be told how to bargain. The Knowledge section allows a builder to upload documents to teach a bot how to negotiate.²¹
8. When you have finished entering data, click “Update” at the top right of the page.
9. The system will integrate your instructions, then display a message with the option to copy a link or go directly to your new bot (“View GPT”).

21. If a builder does not want uploaded material to be used by Plus for other purposes, they may revoke permission for Chat to use data outside their account. This can be done through a tab in Account Settings and in the page produced in the above Appendix.