

Received 30 December 2024, accepted 25 January 2025, date of publication 4 February 2025, date of current version 10 February 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3538599

EXPOSITION

Interacting Large Language Model Agents. Bayesian Social Learning Based Interpretable Models

ADIT JAIN^{ID}, (Student Member, IEEE), AND VIKRAM KRISHNAMURTHY^{ID}, (Fellow, IEEE)

School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA

Corresponding author: Adit Jain (aj457@cornell.edu)

This work was supported in part by the Army Research Office under Grant W911NF-24-1-0083, and in part by the National Science Foundation under Grant CCF-2312198 and Grant CCF-2112457.

ABSTRACT This paper discusses the theory and algorithms for interacting large language model agents (LLMAs) using methods from statistical signal processing and microeconomics. While both fields are mature, their application to decision-making involving interacting LLMAs remains unexplored. Motivated by Bayesian sentiment analysis on online platforms, we construct interpretable models and stochastic control algorithms that enable LLMAs to interact and perform Bayesian inference. Because interacting LLMAs learn from both prior decisions and external inputs, they can exhibit bias and herding behavior. Thus, developing interpretable models and stochastic control algorithms is essential to understand and mitigate these behaviors. This paper has three main results. First, we show using Bayesian revealed preferences from microeconomics that an individual LLMA satisfies the necessary and sufficient conditions for rationally inattentive (bounded rationality) Bayesian utility maximization and, given an observation, the LLMA chooses an action that maximizes a regularized utility. Second, we utilize Bayesian social learning to construct interpretable models for LLMAs that interact sequentially with each other and the environment while performing Bayesian inference. Our proposed models capture the herding behavior exhibited by interacting LLMAs. Third, we propose a stochastic control framework to delay herding and improve state estimation accuracy under two settings: 1) centrally controlled LLMAs and 2) autonomous LLMAs with incentives. Throughout the paper, we numerically demonstrate the effectiveness of our methods on real datasets for hate speech classification and product quality assessment, using open-source models like LLaMA and Mistral and closed-source models like ChatGPT. The main takeaway of this paper, based on substantial empirical analysis and mathematical formalism, is that LLMAs act as rationally bounded Bayesian agents that exhibit social learning when interacting. Traditionally, such models are used in economics to study interacting human decision-makers.

INDEX TERMS Bayesian social learning, large language models, Bayesian revealed preferences, structural results, optimal stopping POMDPs, self-attention, rational inattention, model collapse.

I. INTRODUCTION

This paper discusses the theory and algorithms for interacting Large Language Model Agents (LLMAs) by leveraging techniques from Bayesian inference, stochastic control, and microeconomics. Specifically, we focus on developing interpretable models and stochastic control algorithms for

The associate editor coordinating the review of this manuscript and approving it for publication was Derek Abbott^{ID}.

LLMAs, enabling them to interact sequentially for Bayesian inference.

We construct interpretable models of LLMAs at two levels of abstraction, as outlined in Figure 1. First, we model an individual LLMA as a rationally inattentive Bayesian utility maximizer, capturing the agent's decision-making process under limited attention. Second, we extend this approach to a sequence of LLMAs engaging in Bayesian social learning or *groupthink*, where each agent is a Bayesian utility maximizer.

Our models are inspired by the self-attention mechanism in large language models (LLMs) and observed challenges, such as model collapse, that arise during LLM training.

Furthermore, motivated by the observed bias in the behavior of interacting LLMA, we demonstrate that a sequence of LLMA engaging in Bayesian social learning end up making identical decisions, or “herd”. To address this phenomenon, we propose a stochastic control approach, formulating an optimal stopping problem to balance the trade-off between privacy and herding, to detect the failure state. Our approach is designed for two scenarios: (a) when the LLMA are centrally controlled, and (b) when they operate autonomously.

The overarching goal of this paper is to demonstrate that concepts from controlled sensing and microeconomics, traditionally applied to human decision-making, can be used to both understand and synthesize the behavior of interacting LLMA [1], [2], [3], [4], [5]. We support our theoretical findings with numerical experiments using advanced LLMs for Bayesian inference [6] on real-world data. This paper is written to engage a broad readership, highlighting applications of Bayesian agents in diverse fields, including financial news analysis, e-commerce review evaluation, and online toxicity detection. These examples underscore the flexibility of our methodologies for cross-disciplinary applications. The reproducible code for our experiments is publicly accessible at github.com/aditj/sociallearningllm.

A. MOTIVATION

LLM agents (LLMA) are being rapidly deployed for different applications [7], and to quote Sam Altman, CEO of OpenAI (creators of ChatGPT, a popular LLM which has 200 million weekly active users): “in 2025, we may see the first AI agents join the workforce” [8].

LLMA use a large language model (LLM) to parse the input and have additional functionality to perform tasks. LLMs (such as ChatGPT and LLaMA) are neural networks with billions of parameters trained on trillions of tokens of textual data to parse long texts for summarizing, compiling key facts, and generating new text. The key technical improvement that leads to the efficient deployment of LLMs is the transformer architecture [9]. The effectiveness of LLMs on textual texts has made their deployment and adoption widespread [10]. Many applications have been proposed in healthcare, online platform moderation, and finance, where these LLMs are used to parse the textual observations and suggest decisions based on their outputs [11]. In many tasks, the outputs of the LLMs are often part of a more extensive pipeline; for example, the output of the LLMs, either in a specified format or as embeddings, is frequently used as inputs to other Bayesian entities, including classifiers [12]. The Bayesian framework also becomes essential in applications where the LLMs have to output decisions and need to provide confidence in the decision output. Thus, it is of interest to study a single Bayesian agent that uses the LLM to

parse text observations, update its Bayesian belief, and take action. This paper studies such entities and refers to them as **Large Language Models Agents (LLMAs)**. Constructing interpretable models for LLMAs is crucial to understanding and controlling their interaction.

1) INTERACTING LARGE LANGUAGE MODEL AGENTS

It is predicted that by 2030, 90% of web content will be generated by large language models (LLMs) [13]. In recent practical implementations, individual LLMs are part of a bigger system, referred to as LLMAs, and interact with the content generated by other LLMAs and the external environment [14]. Furthermore, recent research has shown how generative models are trained on the data generated by other generative models can collapse [15]. Therefore, naturally, LLMAs interact with each other either implicitly or explicitly.

Hence, controlling the dynamics of interacting LLMAs is essential to improve the accuracy and trustworthiness of decisions by LLMAs. To the best of our knowledge, only a few recent works systematically study the behaviors of LLMAs using tools from microeconomics and signal processing [16]. This study aims to bridge this gap by systematically reviewing LLMAs and the different mathematical frameworks by studying Bayesian social learning in a sequence of LLMAs to achieve Bayesian inference.

2) INTERPRETABLE ENGINEERING OF LLMAs

Many different third-party services have already started providing various kinds of LLMAs as a service, including Agentforce by Salesforce and IBM AI agents [17]. The underlying intelligence engine of these third-party agents is an LLM or a vision language model (VLM). The LLMAs are used in personal applications for coding, shopping, and scraping data and in enterprise applications for getting insights on user activity and automating industrial workflows.

Therefore, it becomes imperative to study interpretable models for these agents since many of the proposed applications for these agents involve sensitive information (like personal records, financial information, bio-medical data, and personal preferences). By interpretable, we refer to *models that facilitate a transparent understanding of complex models through clear and explainable representations of their decision-making processes*. The workflows of the AI agents also include making decisions, and the interpretability and reliability of these agents become vital for them to be trustworthy. Therefore, mathematical models are needed to aid in engineering and deploying LLMAs. To this end, we propose a LLMA composed of an LLM and a Bayesian engine, which by construction is interpretable. Further, we use Bayesian revealed preferences¹ to reconstruct a Bayesian utility function for both our constructed LLMA and for off-the-shelf LLMAs.

¹The framework of Bayesian revealed preferences is also referred to as inverse optimization or inverse reinforcement learning.

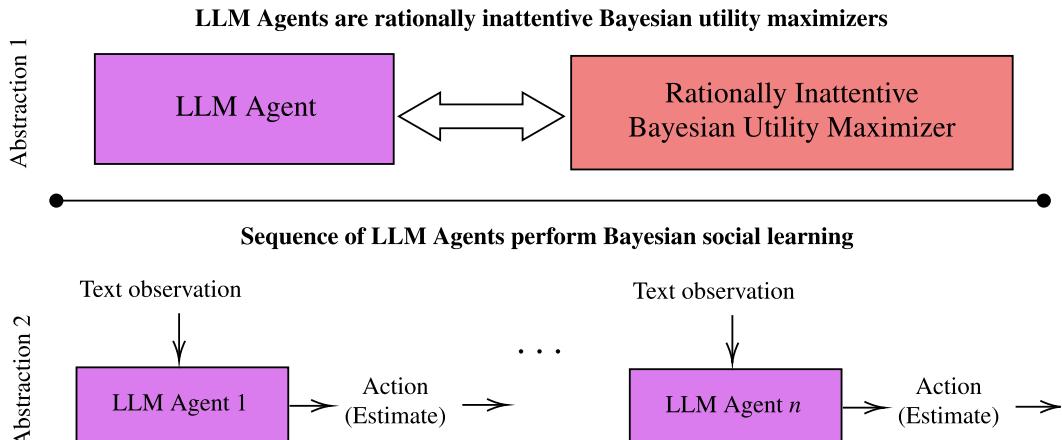


FIGURE 1. Summary of the proposed contributions: We discuss the different blackbox models for LLMA and how LLMAs can be used as a sensing mechanism to perform Bayesian inference. Part 1 models the LLMAs as a rationally inattentive Bayesian utility maximizer and numerically establishes the behavior in applications of product quality identification and hate speech classification. Part 2 discusses how Bayesian social learning in a sequence of LLMAs can be used for sequential state estimation. However, in Part 3, we show that the agents can perform the same incorrect action due to herding. We then discuss a stochastic control approach to delay herding when LLMAs are centrally controlled and when they are autonomous but are incentivized.

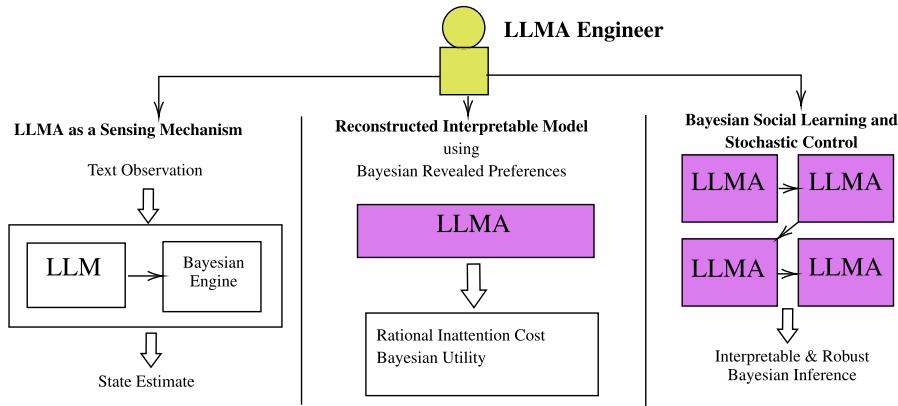


FIGURE 2. Engineering with large language model agents (LLMAs): We propose engineering with LLMs on three different levels: a) First, we construct LLMAs with an LLM attached to the Bayesian engine. The LLM acts as a sensor for the text input and outputs interpretable low-dimensional outputs, which are used by the Bayesian engine to produce a state estimate. b) Second, we formulate necessary and sufficient conditions for a LLMA to be a rationally inattentive Bayesian utility maximizer (RIBUM). We also present algorithms to reconstruct feasible utilities and rational inattention costs if the LLMA is indeed a RIBUM, attributing the LLMA with an interpretable microeconomic model. c) Finally, we demonstrate how a sequence of LLMAs can efficiently perform sequential Bayesian social learning by controlling their outputs to delay herding optimally. Our Bayesian social learning models can be extended to study Bayesian social learning in a network of LLMAs.

3) BAYESIAN INFERENCE FROM MULTI-MODAL DATA STREAM

In various applications, like online e-commerce platforms, video streaming platforms, and social networks, there is a rich stream of multimodal data available using text, images, and videos. Different inference tasks involve fusing information from various data streams to get actionable insights. With the recent progress in deep learning, many of the traditional signal processing methods are being replaced with contemporary methods that use LLMs and VLMs. However, just using static models is not sufficient to model the dynamics of real-life settings, e.g. on online platforms, and underlying dynamics are better modeled in a Bayesian framework. Therefore, motivated by practical

applications, we propose the construction of LLMAs which can perform Bayesian inference sequentially on a data stream. This complements continual learning, which deals with continually learning new tasks without forgetting what was learned previously [18].

B. MAIN RESULTS

This paper builds on Bayesian revealed preferences from microeconomics (inverse reinforcement learning), sequential Bayesian estimation (from signal processing), and structured stochastic control to construct interpretable models and synthesize interaction of LLMAs. The impact of our results on more efficient, systematic, and interpretable engineering

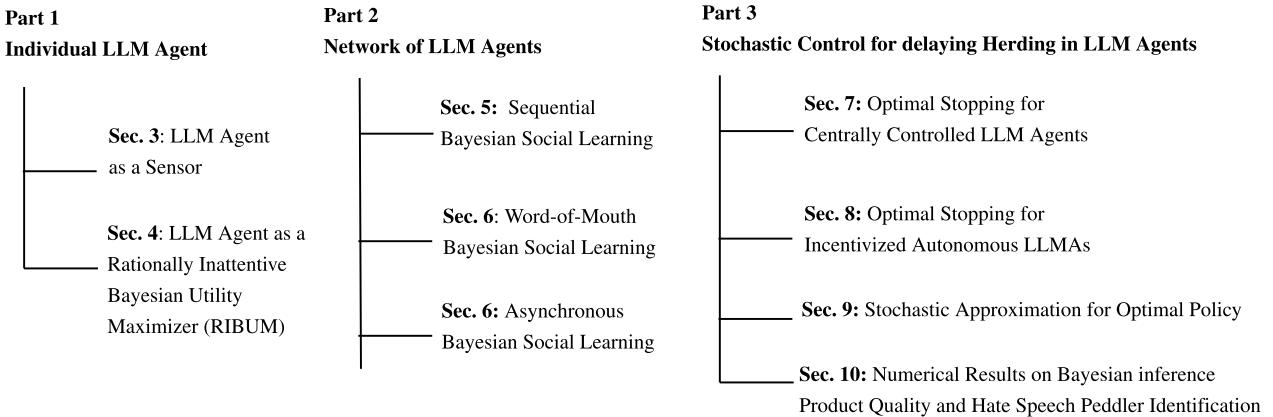


FIGURE 3. Organization of the paper: The paper is divided into three parts. Part 1 deals with interpretable models for an individual LLM agent. Part 2 extends the models to a social learning setting where LLM agents interact with each other to perform Bayesian inference. Part 3 proposes stochastic control methods to delay herding in a sequence of LLM agents.

of LLMAs is summarized in Figure 2. The main contributions of this paper are:

- 1) We propose constructing a LLMA as a composition of a large language model (LLM) sensor,² which acts as a low-dimensional map from the text space, and a Bayesian engine, which uses the measurement from the LLM to update the posterior and act optimally. We show how this model is useful for interpretable Bayesian inference with applications in sequential data on online platforms.
- 2) To obtain an interpretable utility function for a LLMA, we provide necessary and sufficient conditions in Theorem 1 for a LLMA to be a rationally inattentive Bayesian utility maximizer (RIBUM). For a LLMA who is a RIBUM, we propose Algorithm 2 and Algorithm 3 to reconstruct the max-margin and sparsest utility estimate, respectively. Our methods are applicable to both our LLMA and off-the-shelf LLMAs.
- 3) We study Bayesian social learning in a LLMAs, sequentially estimating a state given text observations and in Theorem 2 show that such a sequence of LLMAs form an information cascade and herd in their actions. We show that this is true for both when no private observations are shared and when a finite number of private observations are shared. Further, we provide a detailed analysis of the effect of the quality of results from LLM of the LLMA and the number of private observations.
- 4) To delay herding in a sequence of LLMAs, we formulate an optimal stopping problem for two regimes: a) when the LLMAs are centrally controlled by an entity b) when the LLMAs are autonomous but are incentivized by an entity. We show in Theorem 3

and Theorem 4 that under certain assumptions on the observation matrix and cost functions, the optimal policy for the partially observed Markov decision process of both the optimal stopping problems has a threshold structure. We then propose a policy gradient algorithm in Algorithm 7, which exploits the structural results to estimate the optimal policy parameters. The algorithm does not need access to the system parameters, is computationally efficient, and can track changes in the system.

- 5) We finally present several numerical experiments to demonstrate the efficacy of our proposed methods. We show how our constructed LLMA can be used for interpretable Bayesian inference for analyzing financial data. We show how the Bayesian revealed preferences framework can estimate the utility of an off-the-shelf LLM when used for hate-speech detection. Finally, we show numerical studies on two examples of sequential Bayesian inference: hate speech peddler identification and product quality analysis, to demonstrate herding of LLMAs, and the applicability of our structural results.

To summarize, this paper attempts to answer the following questions with respect to interacting LLM Agents,

- 1) How can LLMAs be constructed so that they can be used for sequential Bayesian inference such that the observation and outputs are interpretable?
- 2) What is a principled approach to analyze whether a LLMA is a Bayesian utility maximizer and also reconstructs its utility function given only blackbox access?
- 3) How does one systematically study Bayesian social learning in multiple interacting LLMAs to explain observed behaviors such as herding and model collapse?
- 4) How can herding in (centrally controlled or autonomous) LLMAs be optimally delayed so that the agents

²We consider LLMs as social sensors and not physical sensors since they do not sense physical quantities like temperature but virtually analyze text and multimodal data to provide observations.

optimally switch between preserving privacy and improving estimation to achieve sequential detection?

C. ORGANIZATION

This paper is organized into three parts, and the schematic of the organization is given in Figure 3. Part I discusses the interpretable model for a single LLMA and attempts to answer questions 1 and 2 above. Section II discusses the related work in large language models, agents using LLMs, and current interpretable models for Bayesian inference. Section III discusses the mathematical model used for modeling LLMAs in this paper and motivates the different components involved. Section IV gives the necessary and sufficient conditions for the LLMAs to be rationally inattentive Bayesian utility maximizers (RIBUM). It further proposes algorithms to estimate the utility function for a LLMA which is a RIBUM.

Part II discusses interpretable models for interacting LLMAs and attempts to answer question 3. Section V discusses the mathematical framework of Bayesian social learning in LLMAs and proves that a sequence of LLMAs form an information cascade in finite time. Section VIII discusses a stochastic control problem for the optimal stopping time problem to achieve quickest time herding with minimal loss to the privacy of LLMAs. Section VI discusses interpretable models to explain model collapse and data incest in LLMAs using word-of-mouth and asynchronous social learning.

To decrease the bias when a sequence of LLMAs perform Bayesian inference, Part III deals with stochastic control for delaying herding in interacting LLMAs performing Bayesian sequential learning proves structural results, and proposes a policy gradient approach. Section IX considers the problem of a central controller optimally optimizing a sequence of autonomous LLMAs to achieve the state estimation by optimally controlling herding. Section X proposes a policy gradient based approach to approximate the optimal policy, which has a threshold switching curve. Numerical results on real-life text classification tasks and related applications are discussed in Section XI. Section XII concludes the paper with discussions on future works, open problems, and research opportunities. The appendix contains the proofs and details about the numerical experiments.

We emphasize that the paper is built around a coherent framework with the unifying theme of building interpretable models for interacting LLMAs using Bayesian social learning and powerful generative models from behavioral economics. For the ease of the reader, we have included a motivation and a discussion subsection in each section, which grounds the different aspects of LLMAs to a real-life application and different microeconomics and statistical signal processing tools presented in the section. We also provide different block diagrams and illustrative examples to further aid the reader. All the results reported in this paper are fully reproducible with code downloadable from github.com/aditj/sociallearningllm.

II. BRIEF LITERATURE REVIEW AND RELATED WORK

In this section, we review related work on LLMs,³ LLMAs and social learning using LLMs. We first provide a brief background on LLMs and discuss the different applications and models for LLMAs. We provide motivation for the interpretability of the LLM agents. Finally, we review literature in sequential state estimation setup studied in this paper and provide motivation for using a sequence of contemporary LLMAs in a classical Bayesian inference setting. We also review applications of sequential Bayesian inference using LLMAs. Table 1 summarizes some of the related work.

TABLE 1. Summary of Related Literature studying LLMAs and their interaction: There has been work in engineering, sciences, and economics, which motivates a careful study.

Area	Topic	Papers
Large Language Models	Survey Applications	[19], [20] [21]–[24]
LLM Agents (LLMAs)	Survey Applications Networks	[25]–[28] [14], [29]–[31] [32]–[38]
Bayesian Social Learning	General LLMs Incentivization	[2]–[4], [39] [21], [40]–[42] [17], [19]
Interpretability	Mechanistic Glass Box Models blackbox Models Sociological Fairness	[43], [44] [45] [46], [47] [48]–[51]
Applications of LLMAs in Sequential State Estimation	Product Quality Identification Hate Speech Classification Recommendation	[52], [53] [16], [54] [55], [56]

A. BACKGROUND ON LLMs

Large language models (LLMs) have become omnipresent in various industry applications, given the drastic improvement in compute availability and rapid development of open and closed-sourced models [20]. They are being rapidly deployed for various applications, including education, information retrieval, gaming, recommendation systems, and understanding graphs [21], [22], [23], [24].

The primary reason for the proliferation of LLMs is that they can take a high-dimensional multi-modal (text, images, audio etc.) data input and provide useful inferences about them. This is possible because they are deep learning networks (transformer architecture) with billions of parameters trained on massive amounts (in the order of petabytes) of data (CommonCrawl, etc.) using extremely fast GPUs that can parallelize computations efficiently. There are two different classes of LLMs: a) open source LLMs like LLaMA and Mistral [57], [58] and b) closed source LLMs like ChatGPT and Claude. Open source LLMs make available the underlying deep learning architecture that they use, some even share the data that the LLM is trained on; closed-source LLMs on the other hand only provide an inferencing interface, where the LLM can be asked different questions,

³In this paper, LLMs also refer to the various transformer architectures that process multi-modal data, including images, audio, and documents.

the answer to which is provided by the LLM. The methods and framework of our paper only require blackbox access to these LLMs and are generally applicable to both classes.

Although the adaption of LLMs is rapidly increasing, from a safety and reliability perspective, their deployment in sensitive applications like healthcare, finance, and defense still poses challenges [20]. Even in general-purpose applications, LLM-based chatbots provide spurious information, a phenomenon referred to as hallucination [20]. There have been different approaches to ensure that the outputs of the LLMs are reliable and interpretable, and many of the challenges specific to the new paradigm require revisiting traditional interpretability literature [59]. As reviewed in [60], one of the approaches is to improve the reasoning capabilities of LLMs so that the LLMs provide a descriptive reason for the output. Another school of thought is to mechanistically understand the transformer architecture and training procedure that is the backbone of an LLM [43], [44] to better understand the working and eliminate the sources of bias, if possible.

B. LLM AGENTS

Standalone LLMs are powerful tools for many applications, but there is strong motivation to consider networked LLMs as part of complex systems and integrate them into existing workflows. These networked LLMs are often given *functionality* to interact with the virtual environment and are referred to as LLM agents (LLMAs) [25]. There are two distinct features that LLMAs have that make them different from LLMs:

- 1) *Decisions*: In the workflows that the LLMAs are used in, they are provided functionality (agency) using different mechanisms, including function calls
- 2) *Communication*: The LLMAs are allowed to communicate with other LLMAs, to exchange information. Often, tasks are also broken into smaller sub-tasks and are performed parallelly and sequentially by different LLMAs leading to different topologies of LLMAs [26].

1) APPLICATIONS OF LLM AGENTS

An important application LLMAs are used for is programming, primarily because of LLMs ability to generate code given a text prompt. LLMAs are used to automate different parts of the software lifecycle, including development, deployment, testing, and fixing bugs [26]. Other applications propose using LLMAs in healthcare for counselling [28], financial trading [27], automating customer service [17] and shopping assistants [31].

2) MODELS FOR LLM AGENTS

The different components used in the standard model of a LLMA include a memory, retrieval mechanism, action sets, and an environment. [32] studies different cognitive architectures using these components for LLMAs. There has been a lot of work to improve the capabilities of these components in LLMAs, using a dynamic context [42] and using self notes to perform continual learning [61]. However,

we propose augmenting the LLM with a Bayesian engine to perform sequential Bayesian inference on a stream of data. The Bayesian engine model proposed in this paper can also be used for more general tasks, as we discuss in the conclusion.

3) NETWORKS OF AGENTS

Since many of the LLM agents interact with other agents directly or through content generated by them, there is a need for more systematic and mathematically rich blackbox models for LLMs and LLM agents (LLMAs). Such models help understand their behavior and eventually control it to ensure reliability. Often, the collaboration of LLMAs are modeled as a directed graph [25], [36], [37], which [34] proposes dynamically adapting depending on the task. There are various different programmatic frameworks where the LLMAs can be abstractly programmed to perform different tasks [35], [62]. Some of these frameworks even allow making these agents autonomous [14]. The methods in this paper deal with a line graph topology of LLMAs, which perform sequential Bayesian estimation. The setup studied in this paper can be extended to more general graph structures, and different issues such as data incest can be studied.

C. BAYESIAN SOCIAL LEARNING IN A SEQUENCE OF LLM AGENTS

1) MULTIPLE BAYESIAN AGENTS SEQUENTIALLY ESTIMATING A STATE

The interaction of multiple such Bayesian agents, each receiving a private observation, is motivated by privacy, improved detection, and finite context length. If the same private observation (even the low-dimensional representation) is used, the LLM can be fine-trained on this data, which might contain sensitive information [39]. Also, different LLMs can be given a diverse set of contexts, which enables reducing the bias involved with their decisions [19]. Also, practically due to the finite context length, the observations can be considered private with respect to consecutive LLMAs evaluations.

2) FRAMEWORK OF BAYESIAN SOCIAL LEARNING

Recent work has looked at social learning in LLMs using a teacher-student framework [40], but this work was in a static setting where the LLMs do not have a belief that they adaptively update. In general, sequential social learning in Bayesian agents has been studied extensively [3], and our work formalizes the problem of Bayesian social learning in LLMAs. The theoretical results presented in this paper have been studied before in the context of distributed Bayesian sensors in [2] and [4]. Compared to [2], [4], we view the LLMAs as interpretable Bayesian sensors and provide a more comprehensive outlook. We also consider multiple observations being shared and provide a concentration inequality (Theorem 5) for overspending the incentive with respect to a budget constraint. Recently [63], looked at detecting information cascades using deep learning. Although in this paper we focus on delaying an information

cascade to improve estimation accuracy, methods similar to [63] can be integrated with our approach.

3) INCENTIVIZATION OF THE LLMAs BY A CENTRAL CONTROLLER

Recent research has studied modeling LLMs as autonomous agents and making LLM part of bigger autonomous agents, including robots, self-driving cars, and programming co-pilots [30], [64]. Such autonomous agents can be *leased* from third-party services at a unit cost. The incentive can also be looked at from the following perspective: providing more context to the same LLM can lead to a more accurate output [19] but increases the cost of processing the query. Third-party LLMAs often offer a tiered pricing structure, where higher pricing provides access to more accurate LLMs.

D. INTERPRETABILITY AND SOCIAL FAIRNESS

There has been recent work in augmenting the LLM with an explainable artificial intelligence (xAI [50]) system to provide more interpretable outputs as in [45]. This work is more aligned with the latter, in which we propose using the LLM as a sensor that provides interpretable low-dimensional outputs, used by a Bayesian engine to estimate the state.

However, we use an interpretable Bayesian model to perform sequential Bayesian inference from text observations of an underlying state, whereas [45] create an xAI model with decision trees and n-grams models using outputs from an LLM. Another such work is [65], where the authors propose training a separate concept neural network that uses the output of an LLM to interpretably classify text embeddings. This approach can complement the work in this paper when the setting is dynamic. Our work uses tools from revealed preferences and social learning to analyze the behavior of individuals and interacting LLMAs from a microeconomic lens.

Further, our focus is also on designing LLMAs, which are safe, reliable, and fair, goals which are aligned with operationalizing responsible AI [51]. This becomes challenging to do since LLMAs have been known to show biases which are inherent to human beings like conformity [66] and bias towards different attributes [67]. These effects will be more prominent when the LLMAs interact with each other in different scenarios like a Mindstorm [33] or a language model cascades [68]. Finally, the need to study interpretable models for LLMAs is motivated by the rise of unified agents [69], which are a representation of a trend in artificial intelligence that the different models are converging to a single efficient model [70]. This paper therefore tries to systematically understand Bayesian social learning in LLMAs, to help prevent undesirable phenomena like model collapse [15], [38].

There has been substantial work in the fairness of the machine learning models [48], and even evaluating large language models for different measures of social fairness [71]. However, [49] recently highlighted how it is

extremely difficult to benchmark LLMs on existing fairness metrics because of the way LLMs are used. This becomes even more challenging for LLMAs, where the agents further have functionality and can also communicate with other LLMAs. Therefore, the focus of our study is to construct interpretable models for LLMAs, which can be used to understand the decisions of LLMAs. This understanding can help construct more suitable societal fairness metrics. Our work additionally has relatively mild assumptions on the utility/cost function of the LLMAs, and hence can be adapted for different sociological costs.

E. APPLICATIONS OF SEQUENTIAL STATE ESTIMATION USING BAYESIAN SOCIAL LEARNING IN LLMAs

We detail examples of real-life problems where textual observations of the state are available and sequential Bayesian learning in LLMAs is used to perform state estimation.

1) HATE SPEECH PEDDLER IDENTIFICATION ON SOCIAL NETWORKS

Identifying hate speech⁴ and toxic content has been studied in various contexts, e.g., in reducing unintended bias, detecting covert hate speech, and mitigating hate speech on online platforms [73]. [54] have looked at how to quantify the intensity of hate speech and created labeled datasets. In [74], the authors looked at controlling federated learning for hate speech classification. In this paper, we look at the problem of Bayesian agents identifying hate speech peddlers by sequentially parsing comments from users using an LLM.

2) FINANCIAL NETWORKS

In financial networks, LLMs can be used as sensors to parse textual information, including news articles, opinions on social networks, and financial reports. This can be especially useful for making decisions based on the low-dimensional observations from the LLMs. This process can be automated using LLMAs, based on algorithmic rules [27]. But, since the actions of LLMAs of a single entity affect the environment (market), such a sequence of agents can herd in their decisions, leading to a financial bubble. This has been studied classically in human traders [3] and makes the study of sequential Bayesian learning in LLMAs interesting.

3) PRODUCT QUALITY IDENTIFICATION

One of the issues on e-commerce platforms is to identify poor quality products early on; however, just using numerical ratings can be uninformative, especially when the number of ratings is less. However, there is a lot of information contained in the descriptive reviews of the product, which one can efficiently extract using LLMs. Therefore, a sequence of

⁴There is an active debate on the definition of hate speech and the tradeoff between free speech and hate speech [72]. Hence, to circumvent this discussion, we use hate speech as an exemplary case study of our methods, and the definition of hate speech is implicit from the source of the dataset in the experiments. Our techniques can be applied to different definitions of hate speech and other applications as is described later.

LLMAs can efficiently analyze the reviews of a product to identify if the product is of good quality or not. This is an extended and sophisticated version of opinion mining, which has shown to be effective in sentiment analysis on online platforms [53].

4) PERSONALIZED RECOMMENDATION SYSTEMS

Another primary application of LLMs is in recommendation systems [55], where the LLMs act as a natural language interface between the user and the item (e.g., movies, products) database. LLMAs can be further used to enhance the recommendation quality by analyzing the past activity of the user and of the user's social network. However, using LLMAs directly raises privacy concerns since the users information is sensitive. This paper proposes one way to deal with this, by ensuring that each LLMAs has a different private observation.

F. COMPARISION TO PRIOR WORK

The primary focus of this paper is to employ ideas from statistical signal processing and microeconomic theory to model and analyze interacting LLMAs. However, given the domain-specific constraints that are inherent to LLMAs, we claim the following modeling and theoretical innovation over prior work in Bayesian social learning and interpretable utility reconstruction: a) There is a hierarchical observation structure inherent to the Bayesian sensor model of an LLMA considered here, whereas the Bayesian revealed preference framework only has a single observation likelihood [1] b) There has been limited research on interpretable models for LLMAs. However, in contrast to the other interpretable machine learning literature, including glassbox models, we reconstruct a utility function and an information acquisition cost given only blackbox access to the LLMA and is more closely aligned to [46]. c) We study the effect of sharing more private observations and a better observation likelihood on herding in a classical Bayesian social learning setting [3], [75] d) The social learning framework considered here allows for LLMAs with different LLMs and is presented for the specific application of Bayesian inference in contrast to [40].

Part I: Analyzing a Single LLM Agent

In Part I of this paper, we consider a single LLMA in isolation, where we first construct a Bayesian sensor model for a LLMA which comprises an LLM and a Bayesian engine. Then, we look at the LLMA as a rationally inattentive Bayesian utility maximizer and propose methods to reconstruct utilities for both our constructed Bayesian LLMA and more general LLMA. The motivation for this modeling from this perspective from the self-attention mechanism inherent to LLMs. Part I of this paper comprises of Section III and Section IV.

III. LLM AGENT AS A SOCIAL SENSOR

Motivated by interpretable Bayesian inference on online platforms using LLMAs this section discusses the Bayesian

sensor model we consider for a single large language model agent. We propose the model of LLMs as a sensing mechanism as a map from a high dimensional space (e.g., text prompt) to a low dimensional space (e.g., structured outputs). The LLMs are equipped with a Bayesian engine and are referred to as an LLM agent (LLMA), which updates the prior regarding the state to be estimated using the text observations. This proposed model is depicted in Figure 4.

The different aspects of the mathematical model for the LLMAs are discussed, and the utility of a LLMA is introduced, which can be reconstructed for a blackbox LLMA using the Bayesian revealed preference framework discussed later. We discuss how the framework and the results of the paper can be extended to contemporary models like vision language models (VLMs). In essence, this section, therefore, shows how LLMA acts as a social sensor, which can be applied to sophisticated settings such as online platforms where physical sensors do not work to sense the underlying state from observation obtained by the interaction of humans.

A. MOTIVATION. LLM AGENT FOR INTERPRETABLE SENTIMENT ANALYSIS

Since in many contemporary applications, LLMs are used for inferring the underlying state given a text observation; we construct an LLM (which acts as a likelihood function) equipped with a Bayesian engine, both of which act as a LLMA to perform Bayesian inference on textual data with applications in sentiment analysis on online platforms [53].

We further motivate the construction of such a Bayesian sensor model for an LLM from the point of view of interpretability, reliability, and controllability. There has been a lot of work done to improve the interpretability of the output of a standard LLM [59]. In the blackbox setting, the approach proposes asking the LLMs to provide a reason in addition to the output [60]. This works well in practice for simple applications, however when sequential Bayesian inference needs to be performed on millions of text observations interpreting the reason itself becomes a tedious task. Therefore, we propose using the LLMs as a low-dimensional map from the high-dimensional text space by designing prompts that are useful in analyzing. The LLMs can either be explicitly controlled using the system prompt or their outputs can be restricted to a certain state.

A Bayesian engine then uses these low-dimensional variates to update the belief, which is an easier task to do than on the high-dimensional text data due to the curse of dimensionality. This helps in using the LLMA in a reliable way since the LLMAs can provide confidence in the actions they would take given the observations. Such a model of LLMA is also controllable with respect to the cost function associated with the Bayesian engine, as we illustrate below.

Notation: x' denotes the transpose of vector x . $\text{diag}(x')$ denotes a diagonal matrix with x as diagonal entries. Capital letters (e.g. B) denote matrices and B_y denotes y^{th} row of the matrix.

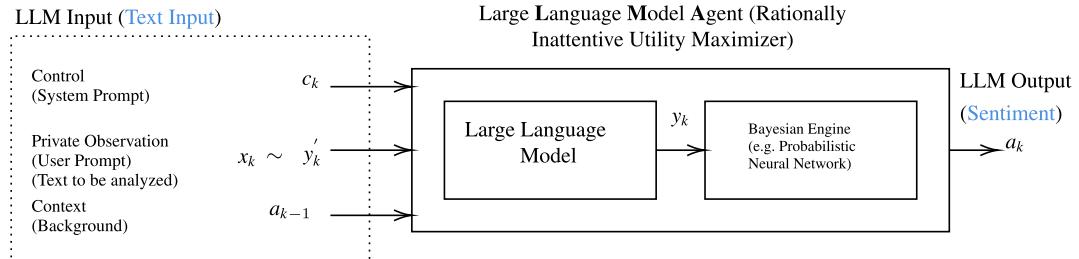


FIGURE 4. Brief schematic of a large language model agent as a sensing mechanism for Bayesian Inference: LLM Input is composed of the system instruction prompt, which is also the control; the user prompt, which is a private observation; and the in-context examples generated from the previous LLM agents are the past actions. Based on the input, the LLM outputs an intermediate textual output. The Bayesian engine uses a likelihood function and past actions to select an action maximizing the expected utility. If utility function is not explicitly given, Bayesian revealed preference is used to obtain a set-valued estimate using an input-output dataset. The paper discusses variations of this model with application in Bayesian sentiment analysis.

B. BAYESIAN UTILITY MAXIMIZATION MODEL FOR LLMA

We consider a LLMA composed of a large language model (LLM) and a Bayesian Engine.

1) ABSTRACTING A LARGE LANGUAGE MODEL (LLM) AS A SENSOR

First, we give a mathematical model for a general-purpose LLM. For this paper, we consider blackbox access to the LLM and hence both open-source LLMs like LLaMA3 [58] and closed-source LLMs like ChatGPT. One of the ways to model a blackbox LLM is as an input-output block. The input to an LLM is a single text prompt, which we decompose into three things: the *system prompt*, which we refer to as *control*, the *context*, and the *user prompt*, which we refer to as *observation*.

Assume that the dictionary of all words of the LLM (tokens⁵) is given by \mathcal{D} , and this dictionary also includes the blank word. The time index is given by $k = 1, 2, \dots$. A blackbox LLM can be viewed as an input-output block.

The control or the system prompt is an input to the LLM, often prepended before the in-context examples and the user prompt, which is used to give instructions to the LLMs on how exactly to respond. This is used to control the behavior of the LLM and ensure that the LLM behaves (outputs) as required. We assume a control of length m_{control} and at time k denote the control by $c_k \in \mathcal{D}^{m_{\text{control}}}$.

Following the system prompt, the next input is a context to the LLM. This context could contain external information and examples that are time-dependent and may depend on previous interactions, which is dynamic and can not be put as a part of the system prompt. We consider a context of length m_{context} and at time k denote it by $\kappa_k \in \mathcal{D}^{m_{\text{context}}}$.

Finally, the user prompt, which we also refer to as the private observation, is the text sequence to which the LLM is supposed to give a response conditioned on the control c_k and the context κ_k . The private observation at time index k is

⁵Although the input of the LLM is text, in most of the architectures, the text (or more generally multi-modal data) input is first decomposed into different tokens and the tokens are processed, but since we consider the LLM as a blackbox, our interpretable models abstract these implementation details.

given by $z_k \in \mathcal{Y}'$. Where \mathcal{Y}' is the text observation space and for a maximum length of m_{user} , $\mathcal{Y}' = \mathcal{D}^{m_{\text{user}}}$.

We consider an LLM, which is pre-trained on trillions of tokens of text to autoregressively generate the next-token.⁶ For developing a token of length 1, the output of the LLM is given by a conditional probability distribution,

$$\phi(c_k, \kappa_k, z_k) = \mathbb{P}(\cdot | c_k, \kappa_k, z_k),$$

where $z \in \mathcal{Y}'$ is the user prompt, c_k is the system prompt and κ_k is the context. Therefore the function ϕ ,

$$\phi : \mathcal{D}^{m_{\text{control}}} \times \mathcal{D}^{m_{\text{context}}} \times \mathcal{D}^{m_{\text{user}}} \rightarrow \mathcal{P}(\mathcal{D}),$$

outputs a probability distribution over the dictionary \mathcal{D} . For generating an output of an LLM with length > 1 , we consider a function g which takes in the function ϕ and outputs tokens from the space $\mathcal{D}^{m_{\text{output}}}$, where m_{output} is the maximum length of the output. The output of the LLM, denoted by y is obtained as $y = g(\phi, c_k, \kappa_k, z_k)$. Therefore, we can represent a blackbox LLM with the following tuple,

$$\mathcal{L} = (\mathcal{D}, m_{\text{control}}, m_{\text{context}}, m_{\text{user}}, m_{\text{output}}, \phi, g). \quad (1)$$

Bayesian inference involves estimating an unknown state $x \in \mathcal{X}$ where \mathcal{X} is the state space using observations of the state. If we are performing Bayesian inference using text observations, LLM can be a powerful tool, as illustrated below in Example 1 and Example 2. The LLM can be directly used to infer an underlying state. In fact, it is already used to do so, as highlighted before in the motivation section. However, we now remark on the challenges with directly using an LLM using the formalism described above:

- 1) Current Large Language Models (LLMs) in production lack the ability to explicitly express confidence in their generated outputs. This limitation significantly hinders their deployment in critical domains like finance and healthcare, where system safety and reliability are paramount.
- 2) Furthermore, the absence of explicit confidence scores presents a significant challenge for human-in-the-loop

⁶There are other techniques of generating the token, but they can be accommodated by a suitable augmentation in the mathematical formulation presented here.

systems. These systems rely on the ability to seamlessly integrate human expertise when an LLM's confidence in its response falls below a certain threshold. Without this crucial feedback mechanism, it becomes difficult to effectively leverage human oversight. Finally, exclusive reliance on a single LLM for critical tasks is inherently risky. LLMs are vulnerable to adversarial prompting, where carefully crafted inputs can induce them to generate incorrect or misleading outputs. This vulnerability underscores the need for robust mechanisms to assess and mitigate the risks associated with LLM-based systems.

- 3) For a task where the LLMs are being used to infer the state given a sequential stream of text observations, the LLMs alone do not explicitly make use of the temporal nature of the observations and can not characterize how many of such observations are enough to be sufficiently confident in the estimate of the underlying state.

Therefore, we propose using the LLM as part of a mechanism (LLMA) where, in addition to the LLM, there is a Bayesian engine that tackles the challenges enlisted above. This ensures that the LLMA acts not just based on the Bayesian estimate of the underlying state, making the decisions of the LLMA more interpretable.

To this end, we propose using the LLM as a map from the high-dimensional space \mathcal{Y}' to a low-dimensional space \mathcal{Y} which takes value from whose cardinality $Y \ll |\mathcal{D}|^{m_{\text{user}}}$. We construct a Bayesian engine composed of two probability distributions, the prior of the state, $\pi \in \mathcal{P}(\mathcal{X})$ and the likelihood of a low-dimensional observation $O : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y})$ and is a conditional probability distribution. We use B and O to denote the observation matrices in the rest of the paper.

From an implementation standpoint, there are different ways we can take to ensure that the output of the LLM is from a low-dimensional space. The low-dimensional output could be a structured output like JSON or a Python dictionary. One way is provide the LLMs with a few examples of the type of outputs we want, for instruction-tuned LLMs this technique has been shown to be very useful [10]. Next we can restrict the dictionary space of the output and reject text tokens not in the restricted dictionary, this is referred to as restrictive or constrained decoding and has been shown to be very effective [76]. Additionally, a lot of LLMs, including ChatGPT, give explicit access to structured outputs (platform.openai.com/docs/guides/structured-outputs).

The following example illustrates how LLMs can be used to map the text space to a low-dimensional space. We discuss how this low-dimensional space can be constructed by an analyst or engineer so that the posterior computed by the Bayesian engine is interpretable.

Example 1: Consider the text dataset of interactions between a customer-service agent and a customer (either a transcript of a call or a chat interface). A service quality engineer is interested in analyzing whether or not the problem of the customer was resolved. The engineer could use an

LLM directly and design a prompt to assess if the issue was resolved. However, there is no interpretability, and if the engineer asks the LLM to give the reason for its answer, then it becomes more difficult for the engineer to analyze. The engineer could, however, design a system prompt so that the LLM answers a specific set of binary (yes-or-no) questions. For example, the engineer could ask the following set of questions:

- 1) What part of the product was the user concerned about?
- 2) Did the user understand the solution being provided?
- 3) How satisfied did the last three messages of the customer seem? (attaches last three messages)
- 4) Is the solution the best possible solution for the problem?

Note that even if we assume the dictionary has $|\mathcal{D}| = 100$ words, and the conversation has length $m_{\text{user}} = 500$, the dimension reduction is substantial from 500^{100} to 2^4 .

2) CONSTRUCTION OF BAYESIAN ENGINE

Owing to dimensionality reduction because of an LLM sensor, we discuss next how a Bayesian engine uses the low-dimensional output to provide an interpretable model based on which an optimal action can be taken. Note that in general, LLMA need not just be used for Bayesian inference but for more general tasks like coding, shopping assistants, research writing, etc., we construct the Bayesian engine to be more general purpose and detail on how a particular cost function leads to Bayesian inference. Also, note that each of the general tasks involves multiple Bayesian inference steps.

For a state $x \in \mathcal{X}$, the LLMA receives a text observation $y' \in \mathcal{Y}'$ which the LLM of the LLMA parses and provides a low-dimensional output $y \in \mathcal{Y}$. The LLMA has a Bayesian engine which has a prior on the state space given by π and an observation likelihood O . The LLMA uses the low-dimensional output from the LLM to compute the posterior using Bayes' rule,

$$T(\pi, y) = \mathbb{P}_{O, \pi}(x|y) = \frac{O_{y, x} \pi(x)}{\sum_x O_{y, x} \pi(x)}.$$

Let \mathcal{U} denote the finite action space of the LLMA. Let $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ be the utility that the Bayesian agent receives from taking action $u \in \mathcal{U}$ when the underlying state x . Then, the Bayesian agent performs the action, which maximizes the expected utility under the posterior distribution.

$$u = \arg \max_{u \in \mathcal{U}} \sum_{x \in \mathcal{X}} r(x, u) \mathbb{P}_{O, \pi}(x|y).$$

Therefore, our LLMA can be described as the following tuple,

$$\text{LLMA}(\mathcal{L}, \pi, O, r), \quad (2)$$

where \mathcal{L} is an LLM of the form (1), π is the prior over the state, O . We make the following remarks on our interpretable sensor model construction of a LLMA.

Remark 1: We described the operation of Bayesian inference on a single observation when there is a stream of

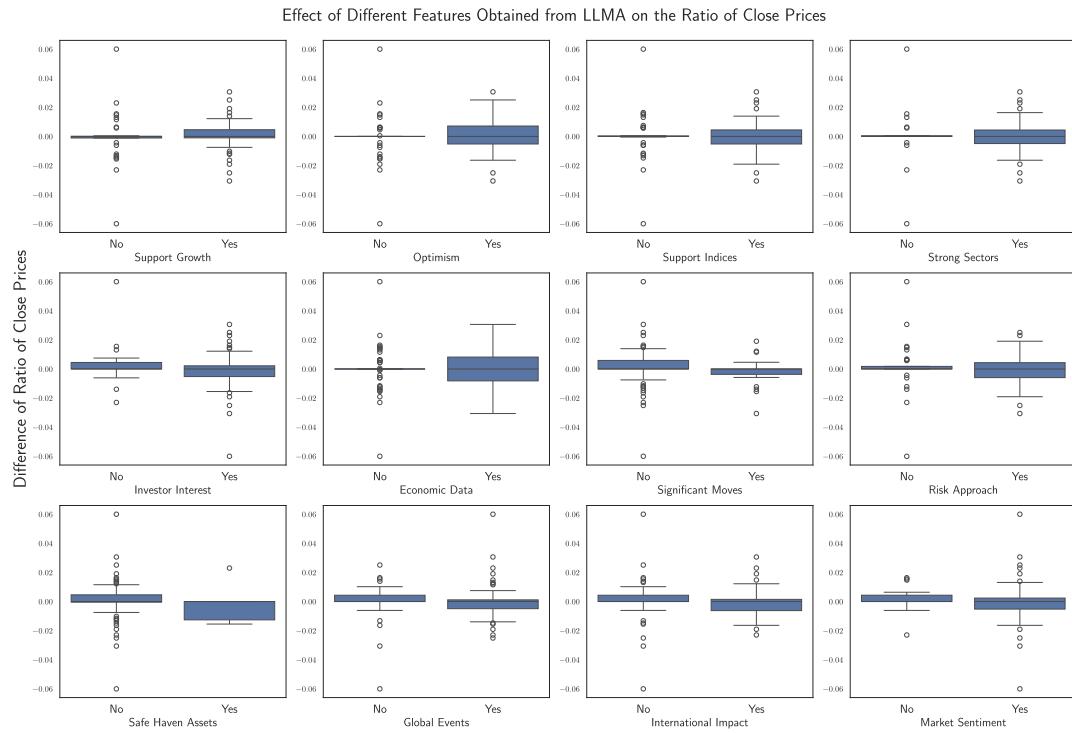


FIGURE 5. LLMAs can be used to detect and analyze the change in financial indicators (difference of close prices) by parsing financial news to extract 16 interpretable features in Example 2. We analyze the news articles from 03/2020 to 08/2020 corresponding to the AAPL stock. We query the LLM for 16 binary with different features, including whether the news article indicates optimism about the market and whether there is investor interest in the stock. The interpretable features can be used to analyze the stock (and subsequently for Bayesian inference), as illustrated by the difference in the ratio of the stock prices across days.

observations y_1, \dots, y_k then the prior is updated after every step with the computed posterior $\pi_{k+1} = \mathcal{T}(\pi_k, y)$.

Remark 2: In Section V, when we describe Bayesian social learning in a sequence of LLMAs, we will note that the equations of the social learning protocol are the same as the above equation. However, the prior of all the subsequent agents is updated based on the actions of agents subsequent to them (even when the rest of the agents do not observe a private observation at any given point), which leads to herding.

Remark 3: For the case of Bayesian inference of the states, the action space is taken to be $\mathcal{U} = \mathcal{X}$, and one of the possible utility functions is the indicator function $r(u, x) = \mathbb{1}(u = x)$.

Remark 4: The framework presented in this section can be extended to multi-modal models like the vision language models (VLMs [77]) for more general tasks by accordingly modifying dictionary \mathcal{D} and output generation mechanism g .

Remark 5: The likelihood O can be computed by using the LLM on a set of synthetic or public offline data where we simulate the state and use the text observations to obtain the low-dimensional observations from the LLM.

3) ILLUSTRATIVE EXAMPLE FOR INTERPRETABLE FEATURE EXTRACTION USING LLMAs

In example 1, we discussed how the low-dimensional representation can be constructed to reduce the observation space. The example can be extended to illustrate Bayesian inference, for example, by analyzing the performance of

a particular customer service agent given its interactions with different customers. We next present an example for a different application, financial news analysis using LLMAs.

Example 2: Consider a financial analyst who receives a stream of financial news and public opinion data from social media. The state in which the financial analyst wishes to estimate if the market is in an upturn or downturn. Similar to example 1, the analyst can design questions using her domain knowledge, which extracts relevant information from the text. Since the stream of data is temporal in nature, the analyst can use our model of a LLMA, to adaptively update the belief of the underlying state, and any point interpret the interminent outputs y to identify trading opportunities. We consider FNNSPID, a financial news dataset [78], where we analyze news pertaining to the AAPL stock. We use LLaMA-3, and for each news article, ask 16 binary questions. The questions are provided in the appendix, and we plot the distribution of the difference of the ratio of close prices (a performance metric used to gauge the performance of a stock across days) in Figure 5. It is clear how LLM can be used as a sensor to parse textual observations and extract interpretable features. The illustrative example using the financial news dataset focuses on the AAPL stock to demonstrate the applicability of the study to readers who may be less familiar with the subject or come from a different disciplinary background, such as finance. We note that the period that we consider was during COVID-19, and therefore was an exception to the standard

stock market behaviour. However we again highlight that the purpose of our methods is to enable analysis by providing an interpretable toolbox and not to analyse any specific event.

C. SUMMARY

In order to construct an interpretable model of a LLMA performing Bayesian inference, we use the LLM as a sensor attached to a Bayesian engine. This section gave a mathematical model for the LLMA, which is modeled as Bayesian sensors. The LLMA we consider in this paper is composed of an LLM and a Bayesian engine. We assumed that the entity interested in using the LLMA has access to the observation matrix described in this section, O . However, this might not be the case where the LLMA is used by a third party, and the Bayesian engine might not be explicit, but still, the entity might be interested in controlling and understanding the actions (decision) of the LLMA. We, therefore, discuss in the next section under what conditions we can reconstruct the utilities of the LLMA by probing the LLMA given blackbox access. This extends the work done in explainable machine learning, where the deep neural network is modeled as a rationally inattentive Bayesian utility maximizer, and the post-training classifications are explained using the utilities obtained from the Bayesian revealed preferences framework [46], [47].

IV. LLM AGENT AS A RATIONALLY INATTENTIVE BAYESIAN UTILITY MAXIMIZER

There are intriguing parallels between self-attention in LLMs and rational inattention in microeconomics. Self-attention is a special type of rational inattention mechanism. Therefore, we use Bayesian revealed preferences from microeconomics to estimate the utilities of a LLMA, which can then be used to understand and control its behavior.

Motivated by the inherent self-attention mechanism of an LLM, this section discusses how an interpretable model for a single LLMA is to model to them as Rationally Inattentive Bayesian Utility Maximizers. First, we state the protocol that a Bayesian agent who is rationally inattentive follows. Then, we consider the problem of the viewpoint of an analyst who only observes the states of nature and the actions of a LLMA and wishes to analyze if the LLMA is a Rationally Inattentive Bayesian Utility Maximizer. For this, we state the necessary and sufficient conditions for LLMAs to act as Rationally Inattentive Bayesian Utility Maximizers. Finally, we discuss algorithms that can be used to get a max-margin estimate and a sparse estimate of the utility function. A few illustrative examples are presented on real-life datasets to explain how the framework can be practically used to systematically obtain utilities of a LLMA and also a *standalone* LLM.

A. MOTIVATION. SELF-ATTENTION MECHANISM OF LLM

The LLM of the LLMA is driven by a transformer neural network. The key innovation of the transformer neural network is the self-attention mechanism. Self-attention allows a model to focus on different input parts when processing each token.

In NLP tasks, for example, it helps a model understand which tokens in a sentence are important in relation to a given token. This relation is then to autoregressively generate texts [9].

In microeconomics, on the other hand, rational inattention is used to model the behavior of individuals managing cognitive resources by prioritizing certain information while ignoring less relevant details due to the inherent “cost” of processing. This is akin to how self-attention mechanisms in machine learning assign weights to different parts of an input sequence, prioritizing relevant segments to optimize understanding or prediction. Both processes are fundamentally about efficient allocation: rational inattention models decisions based on the economic trade-off of information processing costs, while self-attention models adaptively weigh parts of data to capture context, streamlining processing.

Motivated by the inherent self-attention mechanism central to the LLM transformer architecture [20], the LLMA can be modeled as a rationally inattentive Bayesian utility maximizer using the microeconomics model of Bayesian revealed preferences. Rational inattention is about constrained human decision-making due to limited attention, while self-attention is about LLMs assigning attention weights to different parts of input data to optimize understanding or predictions.

Another motivation for studying LLMAs from Bayesian revealed preferences [1] is to model the cost-accuracy tradeoff that is inherent in analyzing large amounts of data done by the LLM of an LLMA. Namely, a LLMAs can output more accurate outputs by increasing the attention effort expended (by a better observation matrix using an LLM with a larger number of parameters or a larger context window). Such a rationally inattentive model for Bayesian agents was first proposed by Nobel laureate Christopher A. Sims. We present the theoretical framework of Bayesian revealed preferences and present experiments on real-life datasets using LLMAs.

B. RATIONALLY INATTENTIVE BAYESIAN UTILITY MAXIMIZING AGENT

In the last section, we discussed a Bayesian Sensor model of a single LLMA. If the LLMA is designed by the entity that is deploying them, then the utility function of the LLMA can be set manually by the cost of the function r of the Bayesian engine. However, if these LLMAs are used off the shelf with or without an explicit Bayesian mechanism, then the utility function is unknown to the entity using them. Although heuristics like confusion matrices or domain knowledge-based cost functions can be used, we need a more systematic approach to estimate the utility functions of a LLMA.

We now present the model of an agent who is a rationally inattentive Bayesian utility maximizer (RIBUM) [46], [47].

Consider a state x belonging to a finite state space \mathcal{X} which is sampled from a prior denoted by $\pi_0 \in \mathcal{P}(\mathcal{X})$. A RIBUM operates in M environments indexed by m and performs an action (denoted by u) from the finite set \mathcal{U} . The utility functions of the RIBUM are given by

$r_{m,u} = [r_{m,u}(1, u), \dots, r_{m,u}(X, u)]'$ for each action $u \in \mathcal{U}$ and each environment $m \in \{1, \dots, M\}$. The RIBUM observes the states x through observations y from a finite observation space \mathcal{Y} . For a given observation matrix (which is a stochastic matrix) $B = (B_{xy} = p(y|x), x \in \mathcal{X}, y \in \mathcal{Y})$, the RIBUM has an information acquisition cost $Z(B, \pi)$. Let $B_y = \text{diag}(B_{1y}, \dots, B_{Xy})$, $y \in \mathcal{Y}$ denote the different probabilities for observing a particular $y \in \mathcal{Y}$. Let Δ denote the set of all $X \times Y$ stochastic kernels B .

Then the RIBUM follows the following protocol:

- (Step 1) The RIBUM first optimizes for the observation matrix by maximizing the expected utility regularized by the information acquisition cost. This optimization is given by,

$$\begin{aligned} B(m) &\in \arg \max_{B \in \Delta} U(r_m, \mathcal{Y}, \pi_0) - Z(B, \pi_0) \\ U(r_m, \mathcal{Y}, \pi_0) &\triangleq \mathbb{E}\{\max_{u \in \mathcal{U}} \mathbb{E}\{r_m(x, u)|y\}\} \\ &= \sum_{y \in \mathcal{Y}} \max_{u \in \mathcal{U}} r'_{m,u} B_y \pi_0. \end{aligned} \quad (3)$$

- (Step 2) A state of nature $x^0 \in \mathcal{X}$ is drawn from prior π^0 and is not known to the RIBUM.
- (Step 3) The RIBUM draws an observation y from optimized observation likelihood $B_{x^0 y}(\mathcal{D})$.
- (Step 4) Given an observation y the RIBUM computes its posterior using Bayesian update step:

$$\pi = \mathcal{T}(\pi_0, y, m) \triangleq \frac{B_y(m)\pi_0}{\mathbf{1}' B_y(m)\pi_0}, \quad (4)$$

where $\mathbf{1}$ is a row vector $[1, \dots, 1]'$.

- (Step 5) Finally, the RIBUM performs the action maximizing the expected utility where the expectation is taken with respect to the posterior computed in Step 4,

$$u \in \arg \max_{u'} \mathbb{E}\{r_m(x, u')|y\} = \arg \max_{u' \in \mathcal{U}} r'_{m,u'} B_y(m)\pi_0. \quad (5)$$

Therefore, given the above protocol a RIBUM agent can be parameterized by the following tuple [5],

$$(\mathcal{M}, \mathcal{X}, \mathcal{Y}, \mathcal{U}, \pi_0, Z, \{B(m), r_m, m \in \mathcal{M}\}). \quad (6)$$

We now make several remarks on the above protocol.

Remark 6: The information acquisition cost $Z(B, \pi_0)$ of Step 1 can be considered as the sensing cost that RIBUM incurs in acquiring the information to make the decision on which action u to perform. The information acquisition cost can also be interpreted as,

$$Z(B, \pi_0) = \sum_y \mathcal{E}(\mathcal{T}(\pi_0, y, m), \pi_0) \mathbf{1}' B_y(m)\pi_0,$$

where \mathcal{E} is an entropic regularizer (e.g., mutual information or Renyi Entropy) [47]. Intuitively, a higher information cost is incurred for a more accurate attention strategy since we obtain a more accurate estimate of the state [46].

Remark 7: Using the above interpretation of the information acquisition cost, the optimization of (3) can be seen as the RIBUM agent optimally choosing the observation sampling strategy. This strategy is chosen to maximize the expected utility regularized by a rational inattention cost.

Remark 8: We now remark on the correspondence between the model of the LLMA from (2) and the RIBUM tuple of (6). First note that even a single LLM \mathcal{L} of the form 6 can be considered as a RIBUM where the \mathcal{L} optimizes its self-attention matrix to optimize for picking the tokens that best predict the next token using the conditional probability distribution ϕ . Further the LLMA can be considered as a product of three observation matrices: 1) from the state to the text observations 2) from the \mathcal{L} (as described above) and 3) from the Bayesian engine, \mathcal{O} . 1) is not in control of the LLMA; however, it is known to the analyst who simulates the text observations for a given state x . 2) comes from the pretraining of the LLM, which involves minimizing an entropic loss with respect to the self-attention mechanism. 3) comes from the LLMA which has a pre-trained likelihood function on a suitable dataset.

C. VIEWPOINT OF ANALYST

An analyst who observes the actions (*behavior*) of the LLMA under different states aims to ascertain if the LLMA behaves as a RIBUM. In particular, the analyst the following dataset,

$$\mathbb{D} = \{\pi_0, p_m(u|x), x \in \mathcal{X}, u \in \mathcal{U}, m \in \mathcal{M}\}. \quad (7)$$

where π_0 denotes the prior distribution of the state and $p_m(u|x)$ denotes the conditional probability of performing the action u given the state x and environment m . The joint probability of the state-action pair (x, u) is given by $p_m(u, x) = \pi_0(x)p_m(u|x)$ and the probability of state x given action u by,

$$p_m(x|u) = \frac{p_m(u, x)}{\sum_{\bar{x}} p_m(u, \bar{x})}.$$

Remark 9: In practice, given the state-action pairs, the analyst empirically estimates the action posterior $p_m(u|x)$, and by Kolmogorov's law of large numbers, the empirical estimate converges to the true distribution w.p. 1 as the number of the state-action pairs goes to infinity.

Given the dataset \mathbb{D} , the analyst aims to a) check if the LLMA is a RIBUM and b) if LLMA is indeed a RIBUM then obtain the reconstructed utility (reward) function \hat{r} which rationalizes the behavior of the LLMA.

We next state the necessary and sufficient conditions for a) from Bayesian Revealed Preferences [1] and then discuss two algorithms describing how b) can be performed.

1) NECESSARY AND SUFFICIENT CONDITIONS FOR RATIONALLY INATTENTIVE BAYESIAN UTILITY MAXIMIZING BEHAVIOUR

As proved in the seminal work of [1], there are two inequalities, the No Improving Action Switches (NIAS) and the No Improving Action Cycles (NIAC), which are

necessary and sufficient for an agent (in our case LLMA) to be a RIBUM.

We now state the NIAS and NIAC conditions and provide intuition for both of them,

No Improving Action Switches (NIAS)

$$\sum_x p_m(x|u)(\hat{r}_m(x, \bar{u}) - \hat{r}_m(x, u)) \leq 0 \quad \forall u, \bar{u} \in \mathcal{U}, m \in \mathcal{M}. \quad (8)$$

Remark 10: The NIAS condition enforces that for any environment m , the agent chooses the optimal action with respect to the posterior probability mass function.

No Improving Action Cycles (NIAC)

$$\begin{aligned} \sum_u \max_{\bar{u}} \sum_x p_l(x, u) \hat{r}_m(x, \bar{u}) - z_l \\ - \left(\sum_u \sum_x p_m(x, u) \hat{r}_m(x, u) - z_m \right) \leq 0 \quad \forall l, m \in \mathcal{M}. \end{aligned} \quad (9)$$

Remark 11: The NIAC inequality operates on pairs of environments and ensures that the agent has an attention strategy that is optimal for all M environments. The above inequality is a pairwise version from [47] of the original combinatorial inequality of [1]. The combinatorial inequality gives a more intuitive explanation for the same, wherein the agent takes actions that are consistent across all possible subsets of the environments. Here, consistency is with respect to the action posterior. Intuitively, NIAC ensures that every agent chooses the best attention strategy in a given environment.

We now state the main results, which show that for a LLMA to be a RIBUM, it is sufficient to check if the dataset \mathbb{D} obtained from the RIBUM satisfies the NIAS and NIAC conditions. We summarize this feasibility check in Algorithm 1, where the input is the dataset \mathbb{D} of the form (7), obtained from the LLMA and ascertains if the LLMA is a RIBUM or not.

Theorem 1 (Necessary and Sufficient Conditions for LLMA to be a RIBUM): Let \mathbb{D} be the dataset that the analyst has, as described in (7) for a LLMA performing protocol (Step 1 to 5 of Sec. IV) in $M \geq 2$ environments. Then the LLMA is a RIBUM iff there exists a feasible solution $\{\hat{r}_m(x, u), z_m(x, u)|x \in \mathcal{X}, u \in \mathcal{U}\}_{m=1}^M$ to the NIAS inequality of (8) and the NIAC inequality of (9).

The result was first derived in [1] and has been used extensively to verify if different engineering systems, including RADARs and Deep Learners, are RIBUM or not [46], [47].

Remark 12: The above theorem gives an if and only if condition for a LLMA to be a RIBUM. If the inequalities have a feasible solution, then there exists a reconstructable set of utilities and information acquisition costs that rationalize \mathbb{D} . The necessity implies that for a RIBUM, the true utilities satisfy the NIAC and NIAS conditions; hence, Theorem 1 yields consistent estimates of the utilities.

Algorithm 1 Bayesian Revealed Preferences Feasibility (BRP)

Input: Dataset \mathbb{D} of the form (7) from LLMA.

Ascertain: If $\exists \hat{r}_m$ and $z_m \forall m \in \mathcal{M}$ satisfying the NIAS inequality from (8) and NIAC inequality from (9)

Remark 13: The feasibility check is summarized in Algorithm 1 which is derived from Theorem 1. Any utility that satisfies the NIAS and NIAC inequalities with respect to the dataset \mathbb{D} is a feasible utility. Hence, the Bayesian revealed preference returns a set-valued estimate (rather than point estimates) of the true utility. This set-valued estimate is given by \hat{r} and the reconstructed information cost \hat{Z} can be derived as [5],

$$\begin{aligned} \hat{Z}(\mathbb{D}) = \max_{m \in \mathcal{M}} \left(z_m + \sum_u \max_{\bar{u} \in \mathcal{U}} \sum_x p_m(x, u) \hat{r}_m(x, \bar{u}) \right. \\ \left. - \sum_u \sum_x p_m(x, u) \hat{r}_m(x, u) \right). \end{aligned} \quad (10)$$

Although Theorem 1 is an extremely powerful result, for practical applications, we often need a single utility estimate rather than a set-valued estimate. We next discuss two different methods to obtain a point estimate for the utility: the max-margin method and the sparsest utility estimation.

D. ESTIMATING SET-VALUED AND POINT ESTIMATES FOR THE UTILITIES

Estimating the utility function of the LLMA is an inverse optimization or inverse reinforcement learning problem, which is, in general, ill-posed. Hence, instead of reconstructing point-valued utility estimates, we describe how the BRP test can be used to reconstruct set-valued utility estimates. Each of the utilities in the set-valued estimate is a feasible utility. We provide two algorithms that return a point-valued estimate from the set, which satisfy certain other structural properties.

Firstly, we describe the max-margin approach to reconstruct the utility function. Since trivial utilities can satisfy the NIAS and NIAC conditions, we maximize the margin with which each condition is satisfied, denoted by decision variables ϵ_1 and ϵ_2 , respectively. Specifically, we consider the following convex program summarized in Algorithm 2,

$$\begin{aligned} \arg \min_{\{\hat{r}_m(x, u), z_m(x, u)\}_{m=1}^M, \epsilon_1, \epsilon_2} & (\epsilon_1 + \epsilon_2) \\ \text{NIAS}(\cdot) \leq -\epsilon_1, \text{NIAC}(\cdot) \leq -\epsilon_2, \\ \epsilon_1, \epsilon_2 > 0. \end{aligned} \quad (11)$$

The max-margin formulation is especially useful when the analyst wishes to estimate the utilities of LLMA which pass the NIAC and NIAS tests maximally. There is no guarantee that this will be close to the true utility because of the bias in the observed data. Still, this estimate is a useful reconstruction that has shown to work well in practice [46], [47].

Next, we describe the utility reconstruction method, which minimizes the ℓ_1 -norm of the reconstructed utility so as

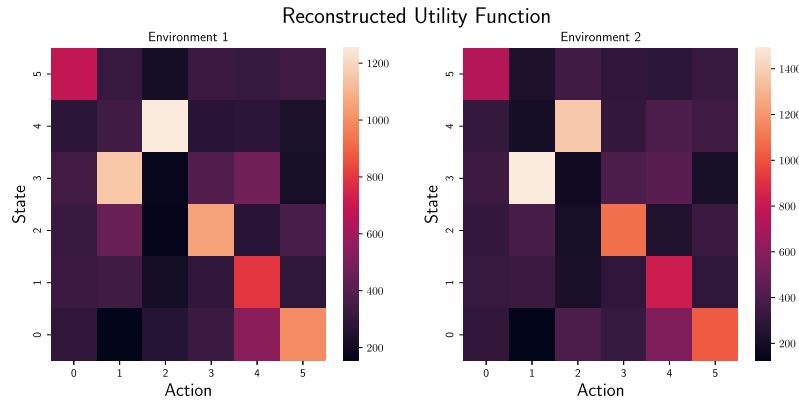


FIGURE 6. Reconstructed Max-Margin Utility of a LLMA for Illustrative Example 3. The utilities near the diagonal are comparatively higher than off-diagonal entries, showing that LLMA correctly classifies the states. Along the diagonal, utility is highest for the state 5 which is the most toxic, hence the LLMA gives the highest utility to classifying the most toxic state.

Algorithm 2 Max-Margin Utility Reconstruction for LLMA

Input: Dataset \mathbb{D} of the form (7) from LLMA.
Output: Reconstructed utilities \hat{r}_m and information acquisition cost \hat{Z} of the LLMA in $m \in \mathcal{M}$ environments; Margins for NIAS ϵ_1 and NIAC ϵ_2 .
if BRP(\mathbb{D}) **is True then**
solve: Optimization (11) for $\hat{r}_m, z_m \forall m \in \mathcal{M}, \epsilon_1, \epsilon_2$
 Obtain \hat{Z} from equation (10)
else
return: *Feasibility Error*: LLMA is not a RIBUM.
end if

to obtain a sparse representation. Here, we manually set tolerances $\epsilon_1 \in \mathbb{R}^+$ and $\epsilon_2 \in \mathbb{R}^+$ for the margins with which each condition is satisfied. Ideally, the margin should be as high as possible such that the linear inequalities still have a feasible solution. Therefore, some trial and error is required to select a suitable ϵ_1 and ϵ_2 . However, this is possible since the Bayesian revealed preference step is offline. We minimize the following convex program, and we summarize the procedure in Algorithm 3,

$$\arg \min \sum_{\substack{\{(\hat{r}_m(x, u), z_m(x, u))\}_{m=1}^M \\ m \in \mathcal{M}, u \in \mathcal{U}, x \in \mathcal{X}}} |\hat{r}_m(x, u)| \\ \text{NIAS}(\cdot) \leq -\epsilon_1, \text{NIAC}(\cdot) \leq -\epsilon_2. \quad (12)$$

The sparsest utility reconstruction is especially useful when the analyst is interested in understanding the key state-action pairs that the LLMA finds especially useful. This sparse utility can be informative in focussing the design of the environment and the system prompt that the LLMA uses.

Remark 14: Note that the above Bayesian revealed preferences framework can be used even if the LLMA does not have an explicit Bayesian engine. This is because, as remarked above, even a standalone LLM has an observation matrix from its pertaining, and the Algorithm 1 only requires access to the state-action pairs from interacting with the LLMA.

Algorithm 3 Sparsest Utility Reconstruction for LLMA

Input: Dataset \mathbb{D} of the form (7) from LLMA, Margins for NIAS ϵ_1 and NIAC ϵ_2 .
Output: Reconstructed utilities \hat{r}_m and information acquisition costs \hat{Z} of the LLMA in $m \in \mathcal{M}$ environments; **if** BRP(\mathbb{D}) **is True then**
solve: Optimization (12) for $\hat{r}_m, z_m \forall m \in \mathcal{M}$
 Obtain \hat{Z} from equation (10)
else
return: *Feasibility Error*: LLMA is not a RIBUM.
end if

Remark 15: Algorithm 1 is a feasibility test with $M(|\mathcal{U}||\mathcal{X}| + 1)$ free variables and $M^2 + M(|\mathcal{U}|^2 - |\mathcal{U}| - 1)$ linear inequalities. The number of free variables and inequalities in the feasibility test of Algorithm 1 scale linearly and quadratically, respectively, with the number of environments, M .

Remark 16: During pre-training, LLMs are trained on massive datasets like the CommonCrawl, which are of the order of hundreds of terabytes, and therefore, they can produce outputs for different contexts. Therefore there is no sample size restriction on the dataset for applying Bayesian revealed preferences to LLMAs. It is important to emphasize the Bayesian revealed preferences provide a necessary and sufficient condition for Bayesian utility maximization, thereby providing a rigorous data analysis tool for real-world data. The primary requirement of both algorithms is that the dataset must be collected in at least two environments, with the action posterior properly defined. There are no additional prerequisites for reconstructing the set-valued utilities from the action posteriors. However, in practice, when the action posterior is estimated empirically, errors are introduced due to finite sample approximations. Characterizing the exact convergence of empirical action posterior probabilities to their true values would be an intriguing direction for future research.

The feasibility condition of Theorem 1 provides a set-valued estimate for the utilities since any utility that satisfies the NIAS and NIAC conditions is a valid utility function. However, for many applications, a single utility function is desired. Therefore, we now describe two algorithms that use the Bayesian revealed preference framework along with Theorem 1 to obtain the reconstructed utility of the LLMA along with information acquisition cost.

E. NUMERICAL EXPERIMENT WITH LLAMA LLM-BASED AGENT

We illustrate the Bayesian revealed preferences for LLMAs in the toy example for hate-speech classification.

Example 3: We consider the example of analyzing a LLMA which classifies a text into six levels of hate speech. Therefore for this task the action space \mathcal{U} is same as the state space $\mathcal{X} = \{0, 1, \dots, 5\}$. The levels indicate the intensity of hate speech. The details of the different levels and the exact construction of the LLMA are given in Section XI. We then obtain 200 pairs of state and actions from the LLMA, which forms our dataset \mathbb{D} . We run the Algorithm 2 and provide the reconstructed max-margin estimates of the utility in Figure 6. It can be seen that the utilities quantify the observed behavior of the LLMA.

F. SUMMARY

The LLM of a LLMA uses an entropic regularization method to provide its output, this motivates looking at the LLMA from the lens of rationally inattentive Bayesian utility maximization, which is a form of entropic regularized utility maximization. This section discussed the necessary and sufficient conditions for a LLMA to be a rationally inattentive Bayesian utility maximizer (RIBUM). We proposed two algorithms to reconstruct a point-valued estimate of the utility of the LLMA if it is a (RIBUM). We illustrate how the reconstructed utility is useful in analyzing the behavior of a blackbox LLMA.

Part II: Interacting LLM Agents

We now move on to Part II of the paper. Having studied a single LLMA in isolation in Part I, we next study social learning in an interacting network of LLMAs. There is a lot of work done in studying different topologies of LLMAs [32]. However, we restrict ourselves to the three topologies described in Figure 7. We motivate studying the different topologies to understand, analyze, and explain some of the observed phenomena in LLMs. We study Bayesian social learning in a sequence of LLMAs to analyze information cascade to an incorrect action. Motivated by model collapse observed while training LLMs from their generated dataset, we briefly study a different protocol for social learning in LLMAs: word-of-mouth social learning. Finally, we illustrate how data-incest can arise if LLMAs perform Bayesian inference in an asynchronous fashion. Part II comprises Section V and Section VI.

V. BAYESIAN SOCIAL LEARNING IN A SEQUENCE OF LLM AGENTS

LLMs are already trained on synthetically generated data by other models [15] and also often use the output of other LLMs to output based on the current context [29]. Motivated by studying interacting LLMs, each of which has computational and privacy constraints, this section introduces a second layer of abstraction, wherein we study Bayesian social learning in a sequence of large language model agents. We first motivate the setting where a sequence of LLMAs sequentially estimate a state from their private observations and take a public action, which is used to update the public belief. We discuss the Bayesian social learning protocol in a sequence of LLMAs, which aim to detect an underlying state by sequentially analyzing text observations of the text. The optimal update equation for the public belief is derived. We consider two scenarios, one where no private observations are shared and one where private observations are shared to the next L LLMAs. Finally, we show that under both scenarios, an information cascade takes place, and the agents take the same action irrespective of their private observation. To show this, we use the martingale convergence theorem [79]. We illustrate the effect of the number of private observations revealed and the resolution of the probe on the convergence in herding. We also present the mathematical model for incentivized autonomous LLM agents used later in Section IX, which is motivated by different entities employing such agents to perform Bayesian state estimation using textual data.

A. MOTIVATION. INTERACTING LLMs, FINITE CONTEXT LENGTH AND PRIVACY IN LLM AGENTS

Even if a single LLMA is used in an application, it can be treated as a sequence of different LLMAs since the context of the previous LLMA evaluation might not be available due to privacy of the content and finite-context length [25]. We consider two scenarios, one where no private observation is shared between the LLMAs and the second where each LLMAs can observe previous L agents. This is motivated by practical constraints from the perspective of privacy, context length, and cost incurred inherent in using LLMAs.

We motivate studying LLMAs using a Bayesian social learning perspective with the following constraints:

1) PRIVACY

Since the text observations often contain sensitive information, the text observations can be used to train the LLM of the LLMA [80]; hence, to prevent this often, systems involving LLMs often treat the private observation in a one-shot setting where the private observation is not stored. Even the low-dimensional representation of the text observation might contain information that can be used to identify attributes of the person the text observation comes from, and in a social network application, this can lead to unfair decisions by LLMAs [48]. Therefore to preserve privacy of users,

Different Types of Bayesian Social Learning

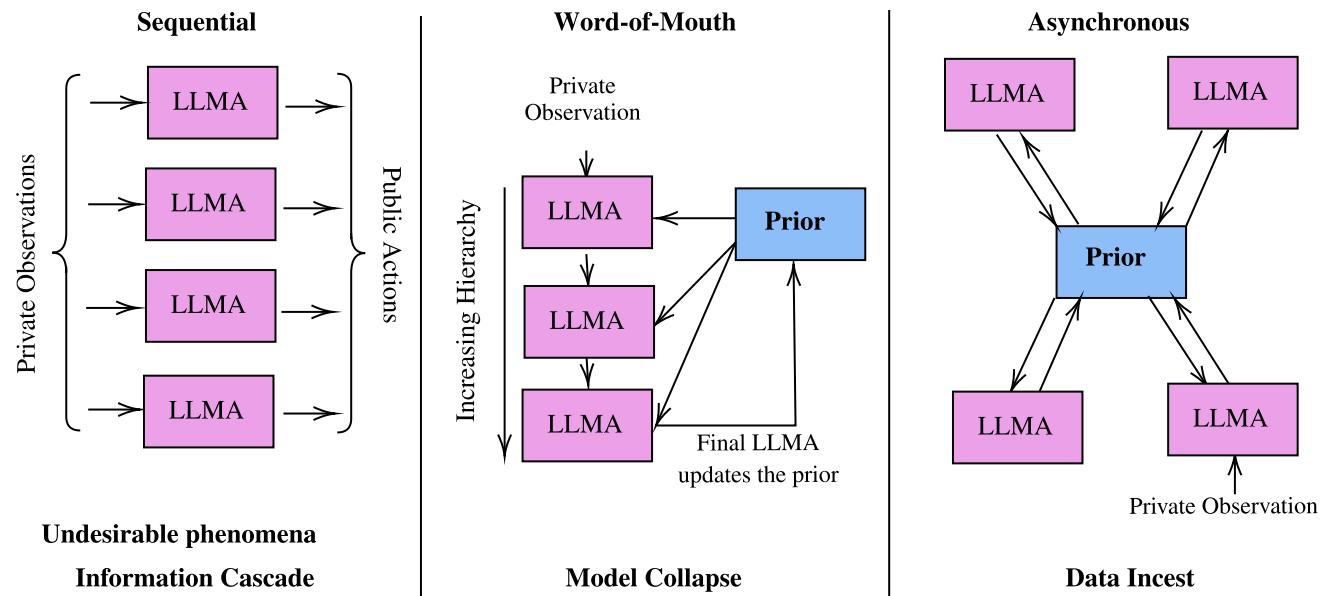


FIGURE 7. The three topologies considered for a network of interacting LLMAs, to interpretably understand and mitigate undesirable phenomena observed in LLMs. We look at sequential Bayesian learning, word-of-mouth learning, and asynchronous Bayesian inference and motivate these topologies using information cascade, model collapse, and potential data incest.

the LLMAs we consider either do not share the private observation or only share a limited sequence of private observations.

2) LIMITED CONTEXT LENGTH AND STORAGE CONSTRAINTS
 Another constraint is the limited context length (length of text that the LLM can process at a time) that is inherent to the LLM used in the LLMA. Note that there are methods which allow for infinite context window or a very large (1 million) context window, however these take a lot of time to work which is often not feasible in a real-time Bayesian inference setting. Also, the quality of the responses decreases with increasing context [19]. There are storage constraints that won't allow storage of an arbitrary amount of private observations, especially if there are methods (like those presented in this paper) that do not require storage of the private observations.

3) COMPUTATIONAL RESOURCES AND COST

Computational resources (involving GPUs) are often limited, especially if the same LLM deployment is used for different applications. Also, the attention mechanism is such that the computational complexity grows quadratically (linear for state space LLMs) in the input length. Therefore, it is often required to limit the size of the context being provided. More importantly, the LLM service providers often bill on a per-token basis. Hence the costs scale linearly with increasing the size of the private observations, however the value of information of including a previous observation is concave.

4) SINGLE LLMA CAN BE MODELED AS A SEQUENCE OF LLMAs

Finally, we remark that a single LLMA can be modeled as a sequence of LLMAs, especially given the above three constraints. This is because when a single LLMA is used to sequentially perform Bayesian inference, the constraints from above enforce that no more than L observations can be considered at any given time. We consider updating the prior based on the action of the $L + 1$ -th previous LLMA. This might seem counterintuitive assumption, however in practical applications often the observations are processed in batches.

B. SOCIAL LEARNING PROTOCOL WHEN NO PRIVATE OBSERVATION IS SHARED

A sequence of large language model agents (LLMA) wishes to estimate an underlying state $x \in \mathcal{X}$, where \mathcal{X} is finite dimension discrete space. At time k agent k receives a private observation $z_k \in \mathcal{Y}'$ from the state x , where \mathcal{Y}' is a high-dimensional discrete space (text). Agent k uses a large language model (LLM) as a sensor to obtain a feature vector $y \in \mathcal{Y}$ where \mathcal{Y} is a low-dimensional discrete space. The text observation is sampled according to the probability distribution $\mathbb{P}(z|x)$. For a given text observation z , the feature vector is sampled according to the probability $\mathbb{P}(y|z)$, \mathbb{P} represents the suitably defined probability measure. Therefore for a given state, the feature vector is sampled with probability $B_{y,x} = \sum_{z \in \mathcal{Y}'} \mathbb{P}(z|x) \mathbb{P}(y|z)$, where $B \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{X}|}$ denotes the observation matrix. Let \mathcal{U} be a discrete

action space and $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^+$ be the cost function.⁷

In classical Bayesian social learning, the agent k computes a posterior belief on the state according to Bayes' rule,

$$\mathbb{P}(x|y_k) = \frac{B_{y_k, x} \pi_k(x)}{\sum_{x' \in \mathcal{X}} B_{y_k, x'} \pi_k(x')}, \quad (13)$$

where $\pi_k(x)$ denotes the public belief belief over the state space \mathcal{X} . The agent k takes an action $u \in \mathcal{U}$ minimizing the expected cost with respect to the posterior,⁸

$$u_k = \arg \min_{u \in \mathcal{U}} \sum_x c(x, u) \mathbb{P}(x|y_k). \quad (14)$$

We make the following assumption which is standard in classical Bayesian social learning [3],

(B1) The observations z_m and y_m are private, i.e., are only available to LLMA m , but the actions are public, i.e., visible to all subsequent LLMA ($k = m+1, m+2, \dots$).

We relax this assumption in the next subsection by allowing the private observations to be shared with the next L agents, which allows us to model more realistic Bayesian social learning in LLMAs. Based on the action u_k of the agent k , the public belief belief π_k is updated using the following filtering equation, which follows from the filtering equation of a hidden Markov model derived in Appendix XII-B,

$$\pi_{k+1} = \mathcal{T}(\pi_k, u_k), \quad (15)$$

where \mathcal{T} is given by the following equation,

$$\mathcal{T}(\pi, u) = \frac{R(\pi, u)\pi}{\mathbb{1}_X R(\pi, u)\pi}, \quad (16)$$

where $R(\pi, u) = \text{diag}([\mathbb{P}(u|x=1, \pi), \dots, \mathbb{P}(u|x=X, \pi)])$ is the probability of actions for different states given the prior. For $u \in \mathcal{U}$, $\mathbb{P}(u|x=i, \pi)$ is given by,

$$\begin{aligned} \mathbb{P}(u|x=i, \pi) &= \sum_{y \in \mathcal{Y}} \mathbb{P}(u|y, \pi) \mathbb{P}(y|x=i, \pi), \\ \mathbb{P}(u|y, \pi) &= \begin{cases} 1, & \text{if } c'_u B_y \pi \leq c'_{\tilde{u}} B_y \pi, \tilde{u} \in \mathcal{U} \\ 0, & \text{otherwise} \end{cases}, \end{aligned} \quad (17)$$

where $B_y = \text{diag}([\mathbb{P}(y|x=1), \dots, \mathbb{P}(y|x=X)])$ and $c_u = [c(1, u), \dots, c(X, u)]'$.

C. SOCIAL LEARNING PROTOCOL WHEN THE LAST L PRIVATE OBSERVATIONS ARE SHARED

We now consider a modification of the Bayesian social learning described in the previous subsection. We let the agents use the observations from the last L LLMAs to update their posterior. We weaken assumption (B1) to the following,

(B2) The observations z_m and y_m are visible to agent m and the next L agents, i.e., to LLMAs $k = m, m+1, \dots, m+L$.

⁷To be consistent with standard Bayesian social learning, we consider cost minimization in Part 2, however, utility is simply the negative of cost for most nonpathological cost functions.

⁸A tie-breaking rule such as uniform sampling can be used if two actions have the same cost.

Algorithm 4 Social Learning Protocol for LLMAs

```

1: Agents aim to estimate state  $x$ 
2: for  $k \in 1, 2, \dots$  do
3:   Agent  $k$  observes  $z_k \sim \mathbb{P}(z_k|x)$ 
4:   Agent  $k$  uses LLM to obtain  $y_k \sim \mathbb{P}(y_k|z_k)$ 
5:   Agent  $k$  computes posterior using (13) or (18)
       depending on availability of previous observations.
6:   Agent  $k$  takes optimal action according to (14)
7:   Agents  $k+L+1, \dots$  update public belief using (15)
8: end for

```

To ensure the privacy of the text z_k , only the feature outputs by the LLMs y_k can also be shared, with the implicit assumption that the likelihood for the different LLMs used by the previous agents is approximately the same.

Agent k creates a vectors using the $L+1$ observations $\mathbf{y}_k = [y_{k-L}, \dots, y_k]'$. Since each of the observations of \mathbf{y}_k of the state x are sampled independently. The augmented observation space is \mathcal{Y}^L . The joint likelihood can be computed as $\mathbb{P}(\mathbf{y}_k|x) = \prod_{m=0}^L B_{y_{k-m}, x}$. Then the Bayesian update of (13) can be augmented as follows,

$$\mathbb{P}(x|\mathbf{y}_k) = \frac{\mathbb{P}(\mathbf{y}_k|x)\pi_k(x)}{\sum_{x' \in \mathcal{X}} \mathbb{P}(\mathbf{y}_k|x)\pi_k(x')}. \quad (18)$$

The LLMA k takes the action u_k corresponding to the action which maximizes the expected cost using (14). We make the following assumption related to the agents discarding actions of the previous agents if the observation is available,

(B3) In lieu of observations y_{m-L}, \dots, y_{k-1} , LLMA m disregards observed actions u_{m-L}, \dots, u_{m-1} of the previous L agents.

Hence as a consequence of (B3), the action of agent k is used by agents $k+L, k+L+1, \dots$ to update their public belief $\pi_m, m = k+L, k+L+1, \dots$ using (15).

D. EMERGENCE OF HERDS AND INFORMATION CASCADES

This section proves that the LLMAs described in the previous section form an information cascade and herd in their actions when the public belief gets strong.

We first define an information cascade occurring at time k for the LLMAs in the following definition.

Definition 1 (Information Cascade): An information cascade is said to occur at time K if the public belief of all agents after time K are identical. That is, $\pi_k(x) = \pi_K(x)$ for all states $\forall x \in \mathcal{X}$ for all time $k \geq K$.

Information cascade implies that the public belief freezes after time K , and since the public belief freezes, the optimal action taken using (14) under any the posterior of (13). Since the information cascade implies the optimal action remains the same, the following definition naturally describes herding at time K for LLMAs where the actions remain the same.

Definition 2 (Herding): Herding in the LLMAs agents takes place at time K if the action of all agents after K are identical, i.e. $u_k = u_K$ for all time $k \geq K$.

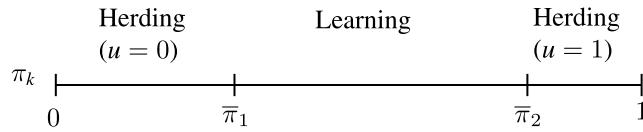


FIGURE 8. Herding and Learning Regions for state space with 2 states. The scalar $\pi_k \in [0, 1]$ denotes the prior of state 1. If the prior is in the region $\pi \in [\bar{\pi}_1, \bar{\pi}_2]$, learning happens; otherwise, the LLMAs form an information cascade and herd in their actions.

It is straightforward to show that an information cascade (Def. 1) occurring at time k implies that herding also takes place at time k (Def. 2). We now state the main result on herding in LLMAs, which shows that the protocol of Algorithm 4 leads to the agents herding in finite time [5].

Theorem 2 (Herding in Bayesian Social Learning of LLMAs): The social learning protocol of the LLMAs described in Algorithm 4, under either assumption (B1) or assumptions (B2, B3) leads to an information cascade (Def. 1) and agents herd (Def. 2) in finite time $K < \infty$ with probability 1.

Proof: Proof in Appendix. \square

Theorem 2 shows that herding happens in finite time, and therefore, the agents take the same action regardless of their private observation. Discarding the private observation, which provides valuable information about the current state, makes their state estimation incorrect and inefficient.

Remark 17: From a purely statistical perspective, Theorem 2 can be seen as the following: when the priors are updated without seeing the observation but rather using the correlated actions, the posterior becomes inconsistent and need not necessarily converge to the true value asymptotically.

E. EFFECT OF THE NUMBER OF PRIVATE OBSERVATIONS REVEALED AND RESOLUTION OF THE PROBE ON HERDING CONVERGENCE

We next discuss the effect of the resolution of the LLM probe and the number of private observations in changing the threshold at which the convergence takes place. For the purpose of this section, assume that the state space is such that $|\mathcal{X}| = 2$ and the action space is such that $|\mathcal{U}| = 2$ and consider the case of the Bayesian agents performing inference.

We first mathematically describe the different regions with respect to the public belief. Then, we derive the relation between the threshold of the public belief and the observation matrix for the different observations. Such a derivation can be used to see the effect of more accurate observations either by considering a higher resolution probe or by considering more number of observations.

We can derive the following for the different regions with respect to the public belief π where herding happens and

where it does not, i.e., where learning happens.

$$\mathcal{R}(\pi) = \begin{cases} \text{Region 1 (Herding } u = 1\text{),} \\ \quad \cap_{y \in \mathcal{Y}} \left\{ \sum_{x \in \mathcal{X}} (c(x, 1) - c(x, 2)) \mathbb{P}(x|y) \leq 0 \right\} \\ \text{Region 2 (Learning } u = y\text{),} \\ \quad \cup_{y \in \mathcal{Y}} \left\{ \sum_{x \in \mathcal{X}} (c(x, 1) - c(x, 2)) \mathbb{P}(x|y) > 0 \right\} \\ \quad \cap \cup_{y \in \mathcal{Y}} \left\{ \sum_{x \in \mathcal{X}} (c(x, 2) - c(x, 1)) \mathbb{P}(x|y) > 0 \right\} \\ \text{Region 3 (Herding } u = 2\text{),} \\ \quad \cap_{y \in \mathcal{Y}} \left\{ \sum_{x \in \mathcal{X}} (c(x, 2) - c(x, 1)) \mathbb{P}(x|y) \leq 0 \right\} \end{cases}$$

Remark 18: The above equation can be derived by equations (15) and (14) by setting the action to be constant for the herding regions. The regions are also illustrated in Figure 8 for 2 states and are numerically shown for 3 states of a real-world dataset in Figure 16.

Note that learning (region 2) only happens when the action taken by the LLMA corresponds to the observation.

Let $\pi = [p, 1-p]'$ and $c(x, u) = |u - x|$. We first derive the expression of the region of herding for a specific observation y ,

$$\begin{aligned} & - \frac{pB_{y,1}}{pB_{y,1} + (1-p)B_{y,2}} + \frac{(1-p)B_{y,2}}{pB_{y,1} + (1-p)B_{y,2}} \leq 0 \\ & \implies \frac{(1-p)B_{y,2} - pB_{y,1}}{pB_{y,1} + (1-p)B_{y,2}} \leq 0 \implies p \geq \frac{B_{y,2}}{B_{y,2} + B_{y,1}}, \end{aligned}$$

and then take the intersection to obtain,

$$p \geq \max_{y \in \mathcal{Y}} \frac{B_{y,2}}{B_{y,2} + B_{y,1}}.$$

We can prove a similar argument for state 2 and obtain the following result,

$$p \leq \min_{y \in \mathcal{Y}} \frac{B_{y,1}}{B_{y,2} + B_{y,1}}.$$

This discussion shows us that even for a simplistic setup, improving the probe accuracy of the LLM helps reduce the herding threshold. However, more accurate LLMs are often larger and have a higher unit cost, clearly highlighting the tradeoff between herding and the cost incurred. A similar result can be derived when the number of shared observations L increases, as the observation space grows with L .

Example 4: We show empirically how interacting LLMs can be used to identify a hate speech peddler (HSP) from a large corpus of data (tweets, blogs, pictures, essays, opinions). Since LLMs charge per token and have latency constraints, we consider multiple LLMs that collaborate to process the large corpus of information. We now experimentally illustrate how information cascades emerge when LLM agents interact to identify an HSP. We used the Mixtral-8 \times 7B-v0.1 LLM. The state $x \in 1 = (\text{not HSP}), 2 = (\text{HSP})$

is the ground truth. The observations $\{z_k\}$ are the content generated by users. When LLM agent k receives z_k , it parses the content to generate a low dimensional observation $y_k \in \{1, 2\}$, designed to detect hate speech. This observation y_k is processed by the Bayesian engine, which computes the posterior of the state. Based on the posterior, the LLM agent selects action $u \in \{1 = \text{not HSP}, 2 = \text{HSP}\}$ by minimizing the Type-1 error cost $\mathbb{P}(x = \text{HSP}|u_1, \dots, u_{k-1}, y_k)$. This action u_k is broadcast to subsequent LLM agents that parse the remaining content, but the private observation y_k is kept confidential to preserve privacy. We empirically computed $B_{11} = B_{22} = 0.8$ from training data. Figure 9 displays two sample paths of actions generated by the LLM agents for different initial priors. Both sample paths emerge into information cascades.

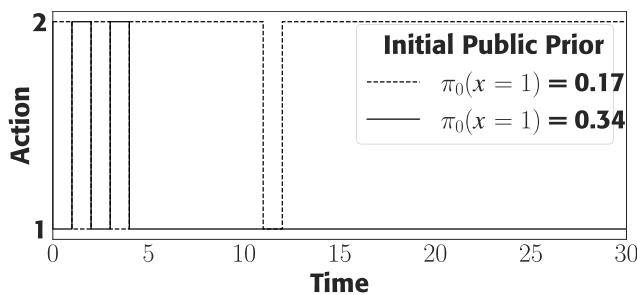


FIGURE 9. Emergence of information cascade in LLM Agents for Bayesian Inference of Hate Speech Peddler in Example 4. Even though the underlying state is 2 (not an HSP), the cascade is in the wrong direction when $\pi_0(x = 1) = 0.34$.

F. SUMMARY

Given the privacy, computational, and cost constraints, multiple LLMAs need to interact with each other to perform sequential Bayesian inference on online platforms. This section studied Bayesian social learning in a sequence of LLMAs to analyze multiple LLMAs performing sequential state estimation on online platforms. We discuss the relaxation of the standard Bayesian social learning protocol, in which the agents are allowed to share their private observations. We illustrate the effect on the herding threshold of the public belief when LLMAs are allowed to share private observations and use more accurate LLMs as sensors.

VI. ASYMMETRIC INFORMATION STRUCTURES OF LARGE LANGUAGE MODEL AGENTS

Motivated by the observed model collapse while training LLMs in a sequential manner and data incest in asynchronous Bayesian sensors, this section studies word-of-mouth and asynchronous social learning. Word-of-mouth social learning is a hierarchical social learning paradigm characterized by asymmetric information flow, where lower-level agents process and communicate observations, and top-down influence, where top-level agents dictate public belief. We provide protocols for both kinds of social learning and a corollary that shows that information cascades happen in word-of-mouth learning as well. Techniques to perform stochastic control to

prevent model collapse and data incest are not in the scope of this paper and can be considered in future work.

A. MOTIVATION. MODEL COLLAPSE IN LLMs

Several recent studies [15] have shown that when LLMs are repeatedly trained on data generated by other LLMs, a phenomenon known as “model collapse” can occur. In model collapse, the output probability distribution of the model collapses to a degenerate distribution as the model is trained iteratively on data generated by the previously trained model [81]. Similar results have been shown on model distillation, where a smaller LLM is trained using data generated from a larger LLM [82]. We use the Bayesian social learning model to show how an information cascade in LLMAs is similar to the model collapse observed while training LLMs.

Model collapse observed while training LLMs can be seen as a special case of sequential Bayesian social learning. In the case of training LLMs using data generated from the previous LLM, it can be considered as estimating the underlying state x , which is the true probability distribution of the data. However at each time k , instead of receiving a private observation from the state x LLMA k receives observation y_k from the previous LLMA $k - 1$. The LLMA k then minimizes a cost function, which is an entropic regularizer (maximum likelihood, KL-divergence, or cross-entropy loss) to obtain an estimate of the state x (represented by \hat{x}) from the observations y_k . The LLMA, then uses the estimate \hat{x} to sample observations y_{k+1} which it provides to agent at $k + 1$.

If the underlying true probability distribution was a Gaussian, then it can be shown that such a protocol leads to a slowdown of learning [83]. However, in the case of discrete distributions in which the LLMs learn such a protocol, it leads to collapsing on one of the support points [15]; therefore, model collapse in the training of LLMs can be studied using the framework of Bayesian social learning.

Asynchronous social learning in LLMAs is motivated by real-time settings like online platforms where there is a stream of data of the order of a hundred thousand every second [34], [84], [85]. Since LLM functionality within an LLMA often requires several milliseconds to a few seconds, especially if these are third-party services, sequential Bayesian learning is often not possible. This is true particularly when the LLMAs are used for tasks that are more sophisticated than just Bayesian inference [26].

B. WORD OF MOUTH BAYESIAN SOCIAL LEARNING IN LLM AGENTS

We now describe the word-of-mouth social learning protocols in L LLMAs. The protocol is summarized as a pseudo-code in Algorithm 5. The protocol can be considered to run on two timescales. On the slower time scale, the first LLMA receives a new text observation $z_{KL} \sim \mathbb{P}(z|x)$ of the state x , where $\mathbb{P}(z|x)$ is the observation likelihood of text z given the state x . The LLMA does not explicitly have knowledge

of $\mathbb{P}(z|x)$ but only receives a text observation z from it. On the faster time scale, the LLMAs communicate with each other in a sequential fashion by generating text observations corresponding to the low-dimensional features from the received text observation. That is, agent m first receives a text observation z_m and uses an LLM to obtain a low dimensional observation y_m . The agent m then takes an action to update the public belief using (14). Then, the LLMA uses the LLM to generate a *synthetic* text observation z_{m+1} used by the next agent.

Algorithm 5 Word Of Mouth Protocol for LLMAs

```

1: for  $k \in 0, 1, 2, \dots$  do
2:   LLMA-( $kL$ ) receives a observation  $z_{kL} \sim \mathbb{P}(z|x)$ .
3:   for  $m \in kL, kL + 1, kL + 2, \dots (k+1)L - 1$  do
4:     LLMA- $m$  obtains low-dimensional features
       $y_m \sim \mathbb{P}(y|z_m)$  using LLM
5:     LLMA takes optimal action  $u_m$  using (14).
6:     Public prior is updated,  $\pi_{k+1} = \mathcal{T}(\pi_k, u_m)$ .
7:     LLMA- $m$  generates synthetic text observation
       $z_{m+1} \sim \mathbb{P}(z|y_m, u_m)$  using an LLM
8:   end for
9: end for

```

Hence, the main difference is that each agent does not receive a new private observation. Note that there are different versions of the word-of-mouth, for example, one where the prior update is done at the end of the inner loop once. All of them are interesting to study. However, we focus on the one presented because we can derive the following result as a corollary of Th. 2 showing herding of LLMAs in Algorithm 5.

Corollary 1: The word-of-mouth social learning protocol of the LLMAs described in Algorithm 5 leads to an information cascade (Def. 1) and therefore herd (Def. 2) in their actions with probability 1.

The above result can be proved using Theorem 2, where each agent has a different observation likelihood based on the previous agent, each of which is a concatenation of two observation likelihoods: the LLM as a low-dimensional sensor map and the LLM sampling a text observation from the text observation. Corollary 1 shows that even in the modified protocol of Algorithm 5, where synthetic data is used to aid the decision-making of LLMAs, cascades are inevitable.

Algorithm 6 Naive Asynchronous Data Fusion in LLMAs

```

1: Initialize prior  $\pi_0$ 
2: while LLMA  $k$  receives new observation  $z_k$  and a
   broadcasted prior  $\pi_{k-1}$  do
3:   LLMA  $k$  uses LLM to obtain  $y_k \sim \mathbb{P}(y|z_k)$ 
4:   Broadcast the posterior  $\pi_k = \mathcal{T}(\pi_{k-1}, y_k)$ 
5:    $k = k + 1$ 
6: end while
7: Return Estimate using  $\pi_k$  and Eq. (14)

```

Asynchronous Social Learning in LLM Agents: We finally consider the asynchronous social learning setting in LLMAs.

Here, the main difference between the previous two topologies is that the agents do not necessarily act in a predefined sequential manner and do not necessarily coordinate.

The protocol is summarized as a pseudo-code in Algorithm 6. The public belief π_k is updated asynchronously when a LLMA receives a new private observation z_k , which it parses using the LLM to obtain a low-dimensional observation y_k . This observation is used to compute the posterior at time π_k using a previous prior π_{k-1} . The posterior is then broadcasted.

It is immediate to see how this protocol can lead to data incest. For example, consider a case where the LLMA k uses the prior π_{k-1} and updates the prior. The prior π_{k-1} was, in turn, updated based on an LLMA m , where $m \leq k - 1$. However, LLMA m at time $k + 1$ sees a new prior and uses the previous one to compute its estimate without accounting for the fact that its previous observation was already used to compute this updated prior. This leads to double counting the observation and, therefore, to data incest [5].

VII. SUMMARY

We discuss how the phenomenon of model collapse observed in LLMs is a form of Bayesian social learning. We present models of asymmetric information structures of LLMAs including word-of-mouth social learning and asynchronous social learning protocols. We state a corollary showing how information cascade occurs with probability 1 even in word-of-mouth social learning and motivate the careful design of asynchronous data fusion in LLMAs.

Part III: Stochastic Control for Bayesian Social Learning in LLM Agents

In Part II, we studied interpretable Bayesian social learning in LLMAs and proved that information cascade is inevitable; we would like to at least delay herding. This is especially important when the LLMAs are prone to cascading to the wrong prior, which is critical in different practical applications, including hate-speech peddler identification.

Part III, therefore, looks at stochastic control for LLMAs. For both regimes of LLMAs, when they are collaborative and autonomous, the paper formulates optimal stopping time problems to control herding by balancing the tradeoff between privacy and estimation. Structural assumptions on the optimal policy of the stopping time problems are proved by making structural assumptions on the cost and observation probabilities. The proposed solutions are extensions to our work in quickest change detection and quickest time herding [4], [16], [86]. A policy gradient algorithm is proposed to estimate the optimal policy without the knowledge of the transition probabilities.

VIII. OPTIMAL STOPPING TIME CONTROL IN CENTRALLY CONTROLLED LLMAs

When LLMAs are deployed in real-life settings, they often exhibit bias in their actions, especially when there are multiple such agents [67]. The previous section showed that such bias could be explained by the herding behavior of

LLMAs. This section formulates an optimal stopping time problem for quickest time herding in a sequence of LLMAs, which ensures that the herding is optimally delayed by letting the LLMAs share private observations. The delay in herding helps improve the state estimate. We discuss a stochastic control approach to solve the optimal stopping time problem and state assumptions that ensure that the optimal policy for the stopping time problem has a switching threshold curve with respect to the public belief. This structural result is exploited in Section X to efficiently approximate the optimal policy. Finally, we discuss extensions of the problem framework to the optimal switching between models of different sizes in sequential state estimation tasks is discussed. The schematic of the system model considered is illustrated in Figure 10.

A. MOTIVATION. CONTROLLING FOR BIAS IN DECISIONS OF INTERACTING LLMAs

Since LLMAs exhibit herding behavior, there is a need to control them to ensure that the estimation is more accurate. In order to do this, we propose optimally switching between sharing the private observation and herding. Our setup is further motivated by the fact that the publicly available LLMs are available in different sizes, and with increasing size, the accuracy of the LLM improves, but so does the unit cost of using the LLM. As we explained in Section V-E, the herding can be delayed by using a more accurate LLM. Therefore, the optimal switching can be considered an optimally stopping problem for using a larger LLM and switching to a smaller one so that the cost is minimized and the herding is optimal. However, solving for the optimal policy of the stopping time problem is a computationally intensive task. Since solving the optimal policy for the optimal switching can be computationally challenging, we look at structural assumptions on the system parameters such that the optimal policy has a threshold structure that can be more efficiently searched for. Note that the POMDP considered in this and the next section is non-standard and, therefore, allows for structural results that can be exploited for efficient policy estimation techniques.

B. SOCIAL WELFARE OBJECTIVE FOR OPTIMAL STOPPING

This subsection formulates an optimal stopping time problem to delay herding by making the agents optimally switch between two modes, acting benevolently by sharing their private observations *or* herding by performing the action from (14). Let $a_k \in \{0 = \text{share private observation}, 1 = \text{herd}\}$ denote the decision at k for the chosen mode. Let $\mu : \mathcal{P}(\mathcal{X}) \rightarrow \{0 = \text{share private observation}, 1 = \text{herd}\}$ denotes the stationary policy which maps the public belief to the decision rule. μ is a sufficient statistic for optimally delaying herding since the information cascade depends on only the public belief [5]. For this section, we assume that $\mathcal{X} = \mathcal{U} = \mathcal{Y}$, that is, in the sequential detection task, the

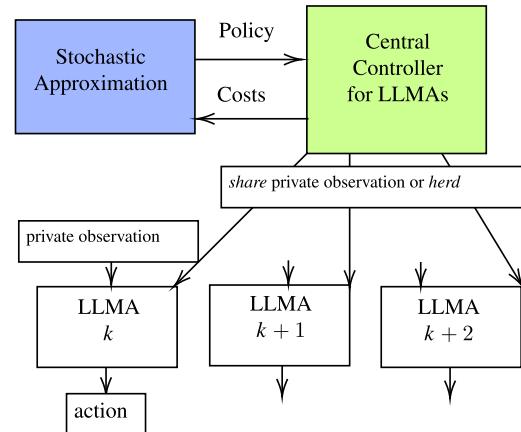


FIGURE 10. Schematic of a stochastic control approach for optimally delaying herding in a sequence of centrally controlled large language model agents. Such a setting is motivated by a central entity hosting the LLMAs.

observation y is the noisy observation of the underlying state x and u is the maximum a-priori estimate of the state.

Further, we formulate the stopping time problem such that one of the states is of special interest (denoted by e_1), and a decision to herd is taken once this state is estimated. To formulate the objective for the optimal stopping time problem, we consider the following natural filtration, \mathcal{F}_k of the actions and decisions till time k , $\mathcal{F}_k = \sigma(\{u_1, \dots, u_{k-1}, a_1, \dots, a_k\})$.

The optimal stopping time problem is to decide when to stop sharing private observations ($a_k = 0$) and announce state e_1 ($a_k = 1$). Let τ denote the stopping time with respect to the filtrations \mathcal{F}_k , $k = 1, \dots$.

We now state the social welfare objective that each LLMA optimizes to solve the optimal stopping time problem and achieve quickest time herding,

$$J(\mu) = \mathbb{E}_\mu \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \mathbb{E} \{c(x, u_k) | \mathcal{F}_{k-1}\} \right. \\ \left. + \sum_{k=1}^{\tau-1} \rho^{k-1} d \mathbb{E}_\mu \{1(x = e_1) | \mathcal{F}_{k-1}\} \right. \\ \left. + \rho^{\tau-1} \Upsilon \mathbb{E}_\mu \{1(x \neq e_1)\} \right. \\ \left. + \frac{\rho^{\tau-1}}{1-\rho} \min_{u \in \mathcal{U}} \mathbb{E} \{c(x, u) | \mathcal{F}_{\tau-1}\} \right\}. \quad (19)$$

Here $\rho \in (0, 1)$ is the economic discount factor and can be set to a lower value if the central controller only wants the first few actions to determine the policy.

Justification for social welfare cost: In (19), the first part corresponds to the discounted cost incurred if the first τ agents perform sensing, and the last term corresponds to the cost incurred for the agents after the stopping time τ (which herd and take the same opportunistic action). The second term is the delay cost with delay parameter d in announcing if the underlying state is e_1 and the third term is the error cost for misclassification with Υ as the parameter. The social welfare cost incentivizes revealing private observations

to delay herding so that enough private observations are available for estimation of the state. The reason we consider estimation with respect to a single state is that many practical applications focus on the identification of a critical state (hateful user, bad product, etc.). The same social welfare cost function of (19) is optimized to obtain a policy with respect to the belief which decides when to stop herding.

Remark 19: In a centrally controlled setting, the central controller is the human operator who manages the platform the LLMAs are deployed on. The action space of the agents consists of either performing the optimal action or revealing their private observation. The decision to ‘stop’ corresponds to taking the optimal action, as this would result in herding. Consequently, the optimal stopping formulation inherently leads to agents sharing their private observations until the decision to stop is made.

We consider the decision rule such that depending on the prior, the learners either share the private observation y_k or they output the action which minimizes the expected cost of (14) under the public belief, i.e., the action u_k is of the form,

$$u_k(\pi_{k-1}, y_k, \mu) = \begin{cases} y_k, & \text{if } \mu(\pi_{k-1}) = 2 \\ \arg \min_u c_u \pi_{k-1} & \text{if } \mu(\pi_{k-1}) = 1. \end{cases} \quad (20)$$

The motivation for this decision rule is two-fold: a) the herding can be delayed by ensuring enough private observations are shared, and b) the agents can reduce cost by outputting the opportunistic action once the public belief is strong enough.

Next, we show that under certain assumptions, the optimal policy μ^* minimizing the expected cost (19),

$$\mu^* = \inf_{\mu} J(\mu), \quad (21)$$

indeed has a nice threshold structure that optimally delays herding and improves the detection of state e_1 .

C. STRUCTURAL RESULTS ON OPTIMAL POLICY

We make the following assumptions on the cost function c and the observation matrix.

- (S1) $c(e_i, u) - c(e_{i+1}, u) \geq 0 \forall i = 1, \dots, X-1 \forall u.$
- (S2) $c(e_X, u) - c(e_i, u) \geq (1 - \rho) \sum_y (c(e_X, u) B_{X,y} - c(e_i, u) B_{i,y}) \forall i = 1, \dots, X.$
- (S3) $(1 - \rho) \sum_y (c(e_1, u) B_{1,y} - c(e_i, u) B_{i,y}) \geq c(e_1, u) - c(e_i, u) \forall i = 1, \dots, X.$
- (S4) B is totally positive of order 2 (TP2).⁹

Discussion of Assumptions: (S1) ensures that the states can be ordered such that taking action in some states is costlier. (S2) and (S3) ensure that the cost function is submodular in the belief. This makes the cost differential between continuing and stopping (herding) the highest for state e_1 and

⁹ A matrix A is TP2 if all second order minors of the matrix A are positive.

gives incentive to the agents to herd when approaching the state e_1 .

The TP-2 condition on the observation matrix in (S4) ensures consistency of the observations [2], i.e., there is some order to the observations. For example, for a 2×2 observation matrix, TP-2 implies that $B_{11} > B_{21}$ and $B_{22} > B_{12}$. If one of the observations is more likely for a particular state, then the other observation must be more likely for the other state. In our experimental results we consider a standard type-1 error based misclassification cost, however more general costs can be considered which satisfy submodularity.

We now state the main structural result on the threshold structure of the optimal policy (21).

Theorem 3: Consider the sequential decision problem of LLMAs for detecting state e_1 with the social welfare cost comprising (19) and the constrained decision rule of (20). Then, under Assumption (S1-S4), the constrained decision rule of (20) is a threshold in the belief space with respect to a threshold switching curve that partitions the belief space $\mathcal{P}(X)$. The optimal policy can be given by,

$$\mu^*(\pi) = \begin{cases} 2 \text{ (continue) if } \pi \in \mathcal{S}_2 \\ 1 \text{ (stop) if } \pi \in \mathcal{S}_1 \end{cases}, \quad (22)$$

where \mathcal{S}_1 and \mathcal{S}_2 are individual connected regions of $\mathcal{P}(X)$. The above theorem proves that under certain conditions on the cost function and observation matrix (S1-4), the optimal policy for solving the discounted social welfare cost optimal stopping time problem has a switching threshold curve. Such a switching threshold can be approximated by set of lines which can be searched for efficiently using a policy gradient based approach [5]. \mathcal{S}_1 and \mathcal{S}_2 are not known to the LLMA and to the central controller, and the structural result only shows the existence of such disjoint regions. These regions are unknown, and we show in Section X that for $|\mathcal{X}| = 2$, one can efficiently search for these regions using a policy gradient algorithm. Although standard reinforcement learning algorithms can be used, we study the extreme case to derive an efficient algorithm. This is computationally tractable, unlike optimizing for a non-structured policy where a finite approximation is used for a general infinite-dimensional policy. This is computationally tractable, in contrast to optimizing for a non-structured policy since it's a finite approximation of a general infinite-dimensional policy.

Remark 20: The LLM agents share their current private observations, and the stopping-time policy determines when the sequence of agents ceases sharing private information. This approach can delay herding while ensuring that not all private information is disclosed. The privacy loss can be quantified using an information-theoretic metric that captures the exact loss of private information. Investigating this in practical scenarios would be an intriguing direction for future research.

D. OPTIMAL SWITCHING BETWEEN DIFFERENT LLMs

The formulation that we proposed in this section can also be used by the LLMAs to switch between LLMs of different sizes. LLMs with a higher number of parameters can process more tokens and give more accurate responses; however, they are more expensive and use more computational resources. Therefore we can use an action space $\mathcal{U} = \{1 = \mathcal{L}_{\text{small}}, \mathcal{L}_{\text{large}}\}$, where $\mathcal{L}_{\text{large}}$ is an LLM with significantly ($10\times$) more parameters than $\mathcal{L}_{\text{small}}$.

Since the LLM with more number of parameters ($\mathcal{L}_{\text{large}}$) can give a more accurate observation, by the arguments of Section V-E, we can use it to delay herding. However, the cost of using $\mathcal{L}_{\text{large}}$ is higher, and accounts for the privacy cost c considered in this section. And then, once we have sufficiently many good responses, the LLMAs can switch to a smaller LLM, which will lead to a quicker information cascade but can still provide low-dimensional readings useful for analytical purposes at a lower cost. Such a problem is often referred to as a quickest change detection problem [5].

E. SUMMARY

When a central entity is designing and deploying LLMAs for sequential Bayesian inference on online platforms, the entity needs to ensure such LLMAs do not have a bias due to herding. In this section, we solved the problem of LLMAs optimally herding in Bayesian social learning framework to announce a particular state, such that their opportunistic cost and cost of obtaining and sharing the observation is balanced. The Bayesian agents considered in this section were cooperative and shared the same socialistically optimal policy. We also briefly discussed how such a scheme can be used to optimally switch between different sizes of LLMs to achieve an optimal tradeoff between cost and accuracy of estimation. In the next section we consider the problem in a different setting, where a central controller can incentivize autonomous LLMAs, and needs to control the incentives for improved state estimation.

IX. OPTIMAL STOPPING TIME CONTROL FOR AUTONOMOUS LLMAs

The LLMAs used for Bayesian inference can often be from different third-party services, each of which requires an incentive to perform the task. Motivated by controlling bias in such LLMAs, this section considers the problem of the central learner optimally incentivizing a sequence of autonomous LLMAs to delay their herding and obtain more accurate estimates. We first formulate the optimization problem of the central controller as a discounted cost POMDP and then show that under structural assumptions on the cost and observation matrix, the optimal policy has a threshold structure with respect to the public belief. This structure is exploited in Algorithm 7 to approximate the optimal incentivization policy of the central controller. The schematic of the setup considered in this section is illustrated in Figure 11.

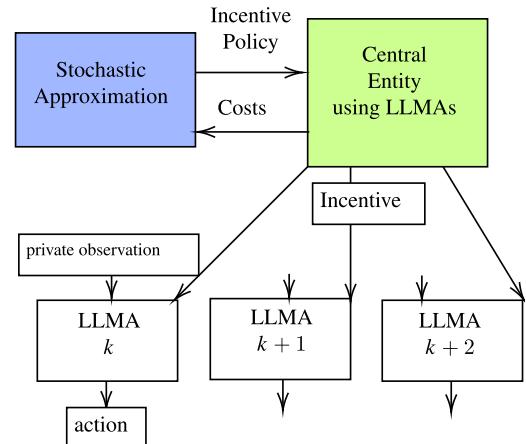


FIGURE 11. Schematic of stochastic control approach for optimally delaying herding in a sequence of incentivized autonomous LLMAs. This is motivated by a central entity deploying third-party LLMAs for Bayesian inference.

A. MOTIVATION. OPTIMAL INCENTIVIZATION OF THIRD PARTY LLMAs

There are already several third-party services that offer LLM agents as a service [17], and such agents can be used for the task of performing Bayesian inference on an online platform. Each of these agents has a cost of processing a query associated with it, and each agent can be asked to give more accurate responses (by performing more computations or by using a larger model). Therefore, we can incentivize the LLMAs to share more accurate private information (low-dimensional outputs). However as shown in Section V, LLMAs are still prone to herding, therefore we propose an optimal stopping time formulation for optimally herding and at the same time minimize the cost incurred by the central entity.

B. OPTIMIZATION PROBLEM OF THE CENTRAL CONTROLLER WITH INCENTIVIZED LLMAs

We now consider a case where the LLMA are incentivized by a central controller, and the cost function is now dependent on the incentive p as well. Specifically, we consider a cost function c as,

$$c(x, u, p) = \alpha_u \mathbb{1}(x \neq u) + \Delta_u + \omega_u p, \quad (23)$$

here, $\alpha_u > 0$, $\Delta_u > 0$ and $\omega_u < 0$, $u \in \mathcal{U}$ are the coefficient for the misclassification cost, the cost incurred in performing the action and the coefficient for the incentive, respectively. The cost accounts for the cumulative cost that the autonomous LLMA incurs while classifying the textual input and the incentive received for the same.

The cost incurred by the central controller for performing information fusion and incentivization is assumed to be linear in the incentive and is taken to be,

$$f(p_k, k) = p_k - g(k) \mathbb{1}(u_k = y_k | \pi_{k-1}), \quad (24)$$

where $g(k)$ is the function that determines the coefficient of the reward associated with the LLMA revealing the

observation (and not herding). Generally, g is decreasing in k since the benefit of a new observation has diminishing returns as more observations are made available.

We now discuss our stochastic control approach to optimally incentivize the LLMAs, such that the fusion cost of the central controller is minimized. We consider the following natural filtration, which contains all information known to the central controller at time k , namely, the initial prior, actions of the LLMAs and the incentives by the central controller,

$$\mathcal{H}_k = \sigma(\{\pi_0, u_0, \dots, u_k, p_1, \dots, p_{k-1}\}).$$

With the estimates of the cost of the individual LLMAs and observation matrices, the central controller can use \mathcal{H}_k to compute the public belief using (15). If the cost and observation matrices are not available, the central control can learn the optimal policy using policy gradient algorithm, as discussed in the next section. Similar to the previous subsection, π_k can be shown to be a sufficient statistic for the filtration \mathcal{H}_k and therefore the incentivization policy $v : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}^+ \cup \{0\}$ of the central controller determines the incentives as,

$$p_{k+1} = v(\pi_k).$$

Therefore, the discounted cumulative cost of the central controller with a discount factor of ρ can be written as,

$$J_v(\pi) = \mathbb{E}\left\{\sum_{k=0}^{\infty} \rho^k f_v(p_k, k)\right\}. \quad (25)$$

The expectation is with respect to the observations and the randomized incentive policy. The optimal incentive policy v^* is then the policy which achieves the minimum cost,

$$J_{v^*}(\pi) = \inf_{v \in \Pi} J_v(\pi). \quad (26)$$

Although classical methods like value iteration can be used to solve the continuous-valued optimization problem of (26), for the simple case of $|\mathcal{U}| = |\mathcal{X}| = 2$, we show in the next subsection that the optimal policy has a threshold structure which can be searched much more efficiently using a policy gradient algorithm. Note that this reduction is still practical for failure state detection e.g. bad product identification using product reviews on online platforms.

C. STRUCTURAL RESULTS

To show structural results, in this section, we consider the simplification that $|\mathcal{U}| = |\mathcal{X}| = 2$. To show the structural results, we make the following assumption on the augmented cost function of the LLMA given in (23).

(S5) The cost function is submodular in (x, u) for all incentives p , i.e., for $|\mathcal{U}| = 2, |\mathcal{X}| = 2$, $c(1, 1, p) + c(2, 2, p) \leq c(1, 2, p) + c(2, 1, p)$.

We consider the following incentive function,

$$\chi(y, \pi) = \frac{\alpha_2 - \alpha_1}{\omega_2 - \omega_1} \frac{B_y \pi}{\mathbb{1} B_y \pi} + \frac{\Delta_2 - \Delta_1}{\omega_2 - \omega_1}. \quad (27)$$

The above incentive function comes naturally when the cost function of the individual LLMAs (23). The derivation is also there in the appendix. We next state results on the structure of the optimal incentivization policy of (26).

Theorem 4: Under (S4) and (S5), the optimal incentive policy $v^ : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}^+ \cup \{0\}$ of (26) is a threshold with respect to the public belief and can be computed as,*

$$v^*(\mu) = \begin{cases} 0, & \text{if } \pi(2) \in [0, \bar{\pi}] \\ \chi(y, \pi), & \text{if } \pi(2) \in [\bar{\pi}, 1] \end{cases}, \quad (28)$$

where χ is the incentive function of the central controller from (27) and $\bar{\pi}$ is the threshold value.

Similar to the result of Theorem 3, this theorem shows that the optimal incentivization policy of the central controller is threshold in the public belief. Since the state space has cardinality 2, the threshold switching curve becomes a single threshold point $\bar{\pi}$, and can be efficiently searched.

Theorem 4 informs that to minimize the discounted cost of (25), the central controller should incentivize (using the scheme of (27)) only when the public belief is not *too strong* in favor of state 1. This supports the intuition that more incentive would be required for a stronger public belief (27).

Further, the incentivization function of (27) and an optimal policy with a threshold structure of (28) implies that the incentive sequence by the central controller is a submartingale. This result and a concentration inequality type bound on the cumulative incentive spent are formalized in the next result.

Theorem 5: Considered the controlled incentivized fusion of information from LLMAs where the cost function is (25) and the optimal incentive policy, v^ satisfies (26) with the incentive function of (27). Under (A1) the optimal incentive sequence $p_k = v^*(\pi_{k-1})$ is a sub-martingale, i.e., $p_k \geq \mathbb{E}\{p_{k-1}\}$. Further, the cumulative incentive spent is such that in a sample path is such that, $\mathbb{P}\left(\sum_{1 \leq k \leq T} p_k \geq B\right) \leq \frac{T}{B}$, where B can be considered as a budget constraint.*

The above theorem characterizes the nature of a sample path of incentivization and secondly provides a bound on the probability that the cumulative incentive exceeds the budget B . This helps analyze the deviation of the total expenditure from the budget B , which is a constraint of the central controller.

D. SUMMARY

Autonomous LLMAs are already offered by third-party entities as a service [7], [17], [87]. They have a unit monetary cost associated with using them for any application. Motivated by such LLMAs, we studied stochastic control of autonomous LLMAs who are incentivized by a central controller to perform Bayesian inference. We showed structural results on the optimal incentive policy and a concentration inequality, which characterized the probability that the central controller would exceed their budget. The next section proposes a policy gradient algorithm that exploits the structural results

of Theorem 3 and Theorem 4 to search for the corresponding threshold policy of the social welfare LLMA and the central controller, respectively.

X. POLICY GRADIENT FOR ESTIMATING THE OPTIMAL POLICY TO CONTROL HERDING IN LLMAs

We propose a policy gradient algorithm that searches for the optimal policy for the stochastic control of the social welfare LLMAs (21) and the central controller (26), which have the threshold structure of (22) and (28), respectively.

A. MOTIVATION. EFFICIENTLY ESTIMATING THE OPTIMAL STOPPING TIME POLICY

When LLMAs are deployed on online platforms to perform Bayesian inference for various purposes, one needs to control for the herding the LLMAs exhibit. For this purpose, we proposed two optimal stopping formulations for centrally controlled and autonomous LLMAs. The efficiency of estimation of the optimal policy parameters is especially important in real-life applications where the time to update the policy parameters is limited. Another constraint is that in real-life systems, access to the observation matrices is limited; therefore, the estimation has to be done without the knowledge of the system parameters. Therefore, we use structural results on the optimal policies for the stopping time formulations and use policy gradients to estimate the optimal policy.

B. POLICY GRADIENT ALGORITHM

Searching for a hard threshold of the form (22) and (28) can be formulated as a discrete stochastic optimization problem. However, in this section, we relax the problem to a continuous stochastic optimization by approximating the hard threshold policy by a sigmoidal of the form,

$$\hat{\mu}(\pi; \Theta) = \frac{1}{1 + \exp(\frac{\pi - \Theta}{\varepsilon})}, \quad (29)$$

where Θ is the policy parameter representing the threshold, and $\varepsilon \in (0, 1]$ is the approximation parameter, and the policy converges to a hard threshold policy at Θ as $\varepsilon \rightarrow 0$. To restrict the policy parameter to be in the range $[0, 1]$, we can reparameterize it as $\sin^2(\theta)$. The approximate parameterized policy of (29) can be used to obtain an approximate value of the cumulative value function (of either (25) or (19)) by interacting with the system using the policy for T interactions.

For the social welfare cost of (19), we can compute the approximate cost with respect to a policy by

$$\begin{aligned} & \hat{J}(\hat{\mu}(\cdot, \Theta_k)) \\ &= \sum_{k=1}^T \hat{\mu}(\pi_k, \Theta_k) \left(\rho^{k-1} c(x, u_k) + \rho^{k-1} d \mathbb{1}(x = e_1) \right) \\ &+ \rho^{T-\sum_{k=1}^T \lfloor 0.5 + \hat{\mu}(\pi_k, \Theta_k) \rfloor} \left[\Upsilon \mathbb{1}(x \neq e_1) + \frac{\min_{u \in \mathcal{U}} c(x, u)}{1 - \rho} \right], \end{aligned} \quad (30)$$

where $\sum_{k=1}^T \lfloor 0.5 + \hat{\mu}(\pi_k, \Theta_k) \rfloor$ just computes the empirical stopping time in place of the stopping time τ in (19). Note that in the above equation, the filtration and the expectations have been replaced with the realized cost, and hence, this is a noisy estimate of the true expected cost of (31).

For fusion cost of (25), the approximate cost is given by,

$$\hat{J}(\hat{\mu}(\Theta_k)) = \sum_{k=1}^T f(p_k, k). \quad (31)$$

We now describe our simultaneous perturbation based policy gradient algorithm, which is summarized in Algorithm 7. We perform H learning episodes, each of which is initialized with the same conditions, and we adaptively update the policy parameter Θ_m across these episodes. For each episode m , we perturb the policy parameter by $\pm \delta$ and obtain two policy parameters Θ_m^+ and Θ_m^- . We then approximate the cost function for both the parameters using (31) by performing T interactions each. We approximate the gradient by finite-difference method (step 6 of Algorithm 7). Finally, the policy parameter uses a gradient descent step with a step size Θ_m . If the parameters of the system are known to be more or less constant, the step size is decreasing to ensure asymptotic convergence [88], else a constant step-size can be used to track changes in the true threshold parameter.

Algorithm 7 Stochastic Gradient Algorithm for Estimating Optimal Policy

```

1: Input: Initial Parameter  $\Theta_0$ , Perturbation  $\delta$ , Time Horizon  $H$ , Step Sizes  $(\beta_m)$ , Episode Length  $T$ 
2: Output: Terminal Parameter  $\Theta_H$ 
3: for  $m$  in  $1, \dots, H$  do
4:   Perturb parameters  $\Theta_m^\pm \leftarrow \Theta_m \pm \delta$ .
5:   Approximate cost with (31) using  $T$  interactions with  $\Theta_k^+$  and  $\Theta_k^-$ ,  $\hat{J}(\hat{\mu}(\Theta_k^+))$  and  $\hat{J}(\hat{\mu}(\Theta_k^-))$ .
6:   Approximate gradient  $\hat{\nabla} J \leftarrow \frac{\hat{J}(\hat{\mu}(\Theta_k^+)) - \hat{J}(\hat{\mu}(\Theta_k^-))}{2\delta}$ .
7:   Update parameter using  $\Theta_{m+1} = \Theta_m - \beta_m (\hat{\nabla} J)_m$ .
8: end for

```

Remark 21: Apart from the parameters of the algorithm, the algorithm only requires the approximate reward with a particular policy. If the cost function is not known to the controller (which is the case when they are incentivized), then the framework from Section IV can be used to estimate the utility of the individual LLMAs using Algorithm 2 or Algorithm 3. The negative of the utility can be used as the reward function. This can account for the misclassification cost. Different incentive regimes can be considered as different environments to exactly obtain the cost function for an incentivized autonomous LLMA.

Remark 22: The above algorithm does not need access to the probability distributions B , which is expensive to obtain, especially in the incentivized case. Also, if the underlying

parameters of the setup are evolving on a slower timescale, the above algorithm can be run with a constant step size to make the policy parameters track the changes in the system. These are the key advantages of using Algorithm 7 compared to the value iteration or other sub-optimal methods for solving our proposed non-standard POMDPs. The computational complexity of Algorithm 7 for each iteration is $O(T)$.

Remark 23: For the centrally controlled case, only a single parameter needs to be communicated across the LLMAs to learn the optimal stopping time policy. Further, since the LLMAs act in a sequential fashion, such a parameter can be communicated either by the central controller or by the previous LLMA, making this a scalable approach. For the incentivization policy too, a single parameter needs to be estimated by the central entity employing the autonomous LLMAs, which makes the policy estimation protocol independent of the number of agents, and therefore scalable.

C. SUMMARY

This section proposed a computationally efficient policy gradient algorithm to estimate the optimal policy of the stopping time problems for LLMAs of Section VIII and Section IX. This algorithm does not need the underlying parameters of LLMAs and is adaptive to changes to the underlying parameters. The next section presents numerical studies that illustrate how these models can be used to control LLMAs.

XI. NUMERICAL RESULTS: SEQUENTIAL BAYESIAN SENTIMENT ANALYSIS USING LLMAs

Our numerical studies demonstrate how the Bayesian social learning framework and the stochastic control approach can be used to perform more accurate and efficient state estimation in the described applications. We present numerical studies on applications related to two real-life datasets: a hate-speech dataset and a product review dataset. These numerical experiments build on our past work to build robust hate speech classification using covert federated learning [74], [89]. The reproducible code, appendix with proofs, and the dataset link are on github.com/aditj/sociallearningllm.

A. MOTIVATION

LLMs are already used in different real-world applications, including on e-commerce platforms, to provide an overview of the reviews and on social media platforms to flag malicious content; therefore, motivated by the real world, we present numerical experiments where LLMAs are used to perform Bayesian inference on different real-world datasets.

We first describe the two main real-world tasks and datasets on which the numerical results are presented. Then, we show how our construction of a single LLMA leads to interpretability. We extend the exemplary study from Section IV and conduct more extensive experiments. Next, we show numerical results on herding phenomena in a

sequence of LLMAs. Finally, we show how the optimal policy for optimal stopping has a threshold structure. We also show the efficacy of policy gradient algorithm in a simulated setting.

B. TASK DESCRIPTION

For both tasks, the LLMA is an online detection mechanism equipped with a large language model (LLM) sensor that analyses comments and flags the state sequentially. The LLM is used to parse the text and obtain a list of appropriate features from the finite-dimensional feature space \mathcal{Y} , which the platform could design. The features contain information about the text and comparisons of the text with a given context. In our setup, the Bayesian engine of the LLMA consists of a likelihood parameterized by a neural network. For a discrete distribution, the likelihood neural network uses restricted Boltzmann machine (RBM) to generate samples from the likelihood with 1000 samples of annotated comments (paragraphs). The experiment results are averaged over N_{MC} independent runs mentioned along with each experiment. The posterior can be updated using (13) using the likelihood and the closed form prior of (15) from the previous step. The LLMA takes an action according to (14). The action is classifying whether the user has hateful intent, and the cost accounts for the misclassification of the state. We use Mixtral 7B, an open-source mixture of experts LLM with 7 billion parameters [57], LLaMA-3 [58] with 70b parameters, and ChatGPT-4o which is a closed source LLM. The details of which LLM was used for this experiment is in the Appendix. We query using the TogetherAI API for the open-source LLMs (Mixtral and LLaMA) and OpenAI API for ChatGPT.

1) HATE SPEECH CLASSIFICATION

Motivation: Flagging users who spread toxic content online is a significant challenge. The state x_k represents the category of peddlers classified based on the intensity and type of content they are propagating. For example, the state could be 3-dimensional, indicating the hate-intent of the user (hateful or not), the hate speech intensity scale [90], and the particular group the hate speech is directed towards. The noisy observations are the text comments from the user that inform about the state and are from a high-dimensional observation space.

In a social network, there may be multiple LLMAs deployed to flag malicious users and decrease the propagation of hate speech. The flags by the previous LLMA are visible, but the private observations are not due to computation and privacy restrictions (so that the LLM can not be fine-trained on the text or the feature mappings). Since the observations are generated sequentially, the LLMA may use the same LLM but with a different context and for a different text observation. Therefore, a single LLMA can be viewed as a sequence of LLMAs learning from their private observations and the past actions of previous LLMAs.

The state space models the state of the user, $\mathcal{X} = \{0 = \text{non-hateful}, 1 = \text{hateful}\} \times \{1, 2, 3, 4, 5\}$, where the first dimension corresponds to whether a user is a hate speech peddler (HSP) or not, and the second dimension to the intensity of the toxicity of the speech (evaluated by crowdsourcing [54]). For the numerical result, we consider an augmented state space with 6-states, $\mathcal{X} = \{0 = \text{non-hateful}, 1, 2, 3, 4, 5\}$, where the last 5 entries denote an HSP of different intensities as described later.

The high-dimensional observations of the state are in the form of text comments posted on online platforms. An LLM is used to parse the text observations by prompting the LLM with the text and a system prompt to return an output belonging to an observation space \mathcal{Y} , which contains the following binary variables: a) targetted towards someone and b) contains explicit words c) indicate violence d) has bias e) is dehumanizing f) is genocidal. We augment the output of the LLM to an observation space of cardinality $|\mathcal{Y}| = 6$. The details of this augmentation, along with additional experimental details, are in the appendix.

The action space \mathcal{U} is considered the same as the state space, and the cost function which accounts for the misclassification of an HSP is given by,

$$c(x, u) = \mathbb{1}(x \neq 0)[\mathbb{1}(u = 0) + |x - u|], \quad (32)$$

The first time accounts for the misclassification of a hateful user, and the second term accounts for the difference in intensity. We use the measuring hate speech dataset from [54], which contains 40,000 annotated comments from Twitter, Reddit, and Gab. The annotations are performed by crowdsourcing and indicate if the comments contain hate speech and the intensity of the toxicity exhibited on a scale of 1 to 5, measured using a Rasch measurement scale [54].

Since the data is anonymized, we consider a synthetic user construction. In a span of T textual comments, a hate speech peddler (HSP) contains hate speech text from one of the intensity levels. Hence, there are 6 types of users: non-HSP and HSP with intensity from 1 to 5, each with $T = 100$ comments of the corresponding intensity.

2) PRODUCT QUALITY IDENTIFICATION

Motivation: Identifying products that are of poor quality early on is important in many e-commerce applications. While the current state-of-the-art relies on a complaint-driven interface that involves a human service agent to identify a bad product, there has been some preliminary work that suggests using existing reviews from the customers to determine the quality of the product [91]. Product reviews are already used to provide a summary of the product on major e-retailers like Amazon and even influence customer decisions [91]. Therefore, we propose using a sequence of LLMAs and our stopping time formulation to identify bad products.

In the case of an e-commerce platform, the restriction on the sharing of the private observation is motivated by the computational constraints that the LLMAs will have when performing Bayesian inference on millions of products

and billions of reviews. The public belief can be efficiently computed based on the output (action) of each of the LLMA.

We use the Amazon Review dataset, which has 233.1 million [56] text reviews for products from 29 categories. Each review also has a rating from 1 to 5. We discuss more details about the dataset in Appendix XII-B1. We consider the state as the quality of the product, with the state space given by $\mathcal{X} = \{0 = \text{"bad-quality"}, 1 = \text{"high-quality"}\}$. We consider the Beauty and Electronics categories and sample 5000 reviews from each. Each review has a text and a rating associated with it. We consider the quality of the product as described above.

We now detail on what we mean by *quality* and the definition of the states. Since we do not have access to any additional dataset other than ratings, we consider products with at least 1000 ratings and compute the quality state ex-ante using these ratings. Therefore, the quality here is more representative of the perceived value for money [92]. We consider the following ranges $[0, 3.3), [3.3, 5]$ of the average review of *all* the ratings to assign a product quality as bad, medium, and good, respectively. To clarify, we just use all 1000 ratings to compute the state; the observations are sampled one at a time and analyzed by the LLMA.

The low-dimensional observation obtained using the LLM of the LLMA belongs to an observation space of cardinality $|\mathcal{Y}| = 16$. We obtain the low-dimensional observation by asking the following questions of the text review:

- 1) Does the review mention any specific problems or defects with the product? (*defects*)
- 2) Does the review mention any positive attributes regarding the product's durability or reliability? (*durability*)
- 3) Does the review indicate that the product meets or exceeds the user's expectations? (*expectations*)
- 4) Would the reviewer recommend this product to others? (*recommend*)

These questions were designed by us by qualitatively analyzing what could predict the perception of quality of a product [91], [92]. These features for the different states are plotted in Figure 12. It can be seen that for the good product, the overall decrease in *defects* is more rapid, and even when in reviews with overall ratings 2 and 3, there are substantially more reviews that mention *durability* positively, than the bad product. However, an analyst could consider an alternate design of the observations depending on the need.

We consider the misclassification cost as follows,

$$c(x, u) = \mathbb{1}(x = 0)\mathbb{1}(u \neq 0) + \mathbb{1}(x \neq 0)|x - u|. \quad (33)$$

C. NUMERICAL EVIDENCE SHOWING LLMAs ARE RIBUM

We perform the max-margin-based utility reconstruction using Algorithm 2 when the agent optimizes a utility in order to achieve product quality identification and illustrate the reconstructed utilities in Figure 13.

For illustration purposes, we consider *recommend*, *durability* and *expectations* as the actions, hence making the action space $|\mathcal{U}| = 8$. To show our methods can be used

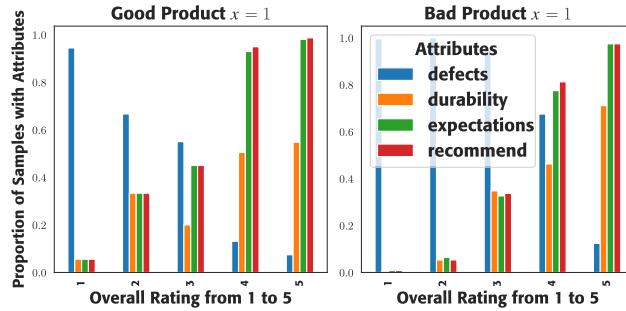


FIGURE 12. Interpretability of the outputs of an LLM Sensor for Product Quality Identification. There are 4 different features that the LLM of the LLMA extracts. The left subplot corresponds to reviews from a good product, and the right subplot corresponds to reviews from a bad product (defined in the text). We analyze text reviews with ratings from 1 to 5 for 4 different attributes and plot the proportion of samples where the attribute is true. It can be seen that the features, indeed, are consistent with the overall rating provided by the reviewer for the two different kinds of products and can be used for a more fine-grained interpretable analysis.

with blackbox LLMs, we use ChatGPT-4o-mini for this experiment. The state space is $\mathcal{X} = \{0 = \text{bad product}, 1 = \text{good product}\}$. We consider the different environments to have different ratings; i.e., the people who rate products at 1 are considered as part of a single environment. The utilities can be used to interpret the behavior of the LLM and can be used in lieu of explicit utilities of the form (33).

D. HERDING IN LLMAS

We report our results with the initial public belief (specifically the prior probability of state 0). Due to Theorem 2, the initial public belief is sufficient to identify intermediate public belief regions where the information cascades are observed.

We first study the freezing of public belief when an information cascade (Definition 1) happens in a sequence of LLMAs performing Bayesian inference for hate-speech classification. We consider 80 timepoints, and the results average over 10 runs. The result is illustrated in Figure 14 for different values of initial public belief. In each subplot, the different lines correspond to the different true underlying states. It can be seen that the public belief for state 0 converges to higher values as the initial prior for state 0 is increased. Hence, the initial prior decides what the prior will freeze.

We show, for a sample path realization, how LLMAs herd in the actions that they take in Figure 15. Each of the subplots corresponds to a different initial public belief, and the lines within each subplot are for a different underlying true state. The reason why it is more resilient to herding when in state 4 is that there is a stronger observation likelihood. All the other states, for a strong enough prior (0.52), herd to an incorrect state (0). Hence classifying an HSP as a non-HSP.

Figure 16 illustrates the regions of herding a case with 3 states. The true underlying state in this simulation is 0, and therefore, it can be seen that most of the initial

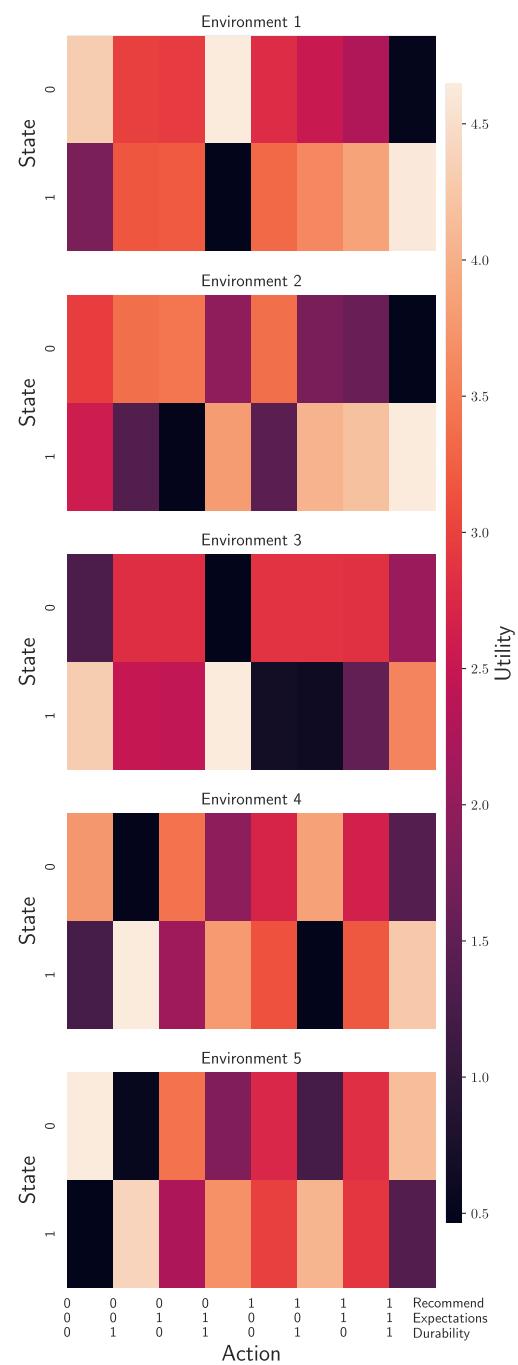


FIGURE 13. Reconstructed max-margin utility of the LLM of an LLMA performing Bayesian inference for product quality identification using text reviews. The reconstructed utilities offer an interpretable way to analyze the behavior of the LLM. Here, the states represent the true quality of the product (0 = bad, 1 = good). The actions are the low-dimensional output of the LLM corresponding to different features in the input text reviews. We observe that for environment 1, when reviews with a rating 1 are considered, the difference between utilities for both states is negligible. However, the contrast increases as the ratings considered in the environment are increased.

public belief corresponds to predicting state 0. We assume an observation matrix with 3 observations, with an observation matrix $B = [0.7, 0.2, 0.1; 0.1, 0.7, 0.2; 0.2, 0.1, 0.7]$ and

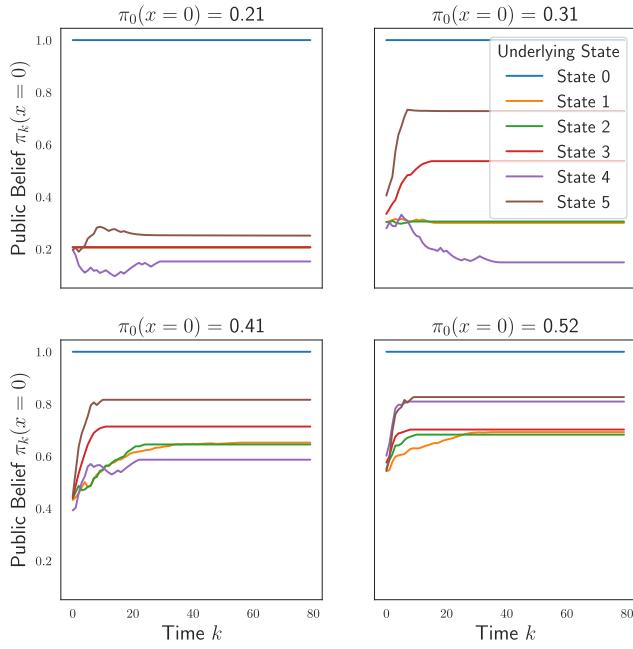


FIGURE 14. Public prior freezes in finite time when LLMAs learn using social learning protocol of Algorithm 4 for detecting hate-speech peddlers. The time taken (average over 10 runs) to form an information cascade is different for different values of the initial public belief $\pi_0(x=0)$ and different true underlying states. Note that for a public prior $\pi_0(x=0) \geq 0.41$ only a few observations are enough for the LLMAs to herd to a strong belief on the wrong state (state 0) regardless of the true underlying state.

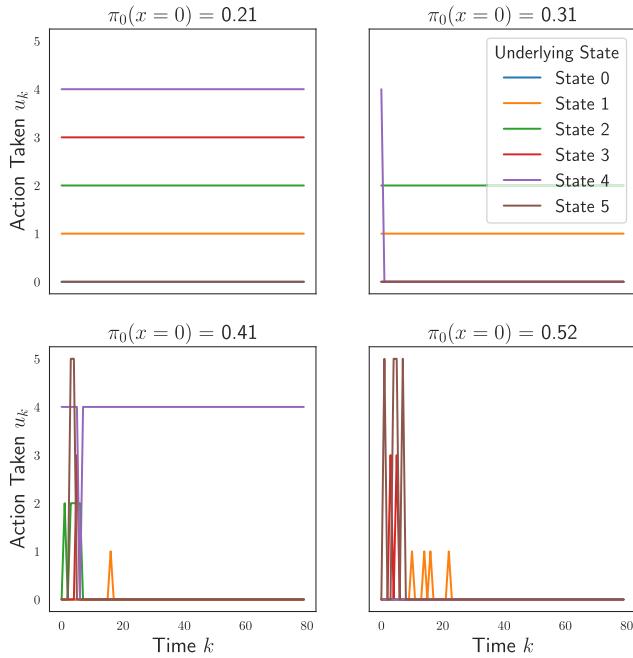


FIGURE 15. Sample path for the actions taken by LLMAs under different initial priors $\pi_0(x=0)$ and true underlying state x . Observation 4 has a stronger observation likelihood and hence is more robust under an increasing prior on state 0.

assume an identity utility function for the LLMAs. However, the top-left and bottom-left corners of the triangle show that herding to the wrong state can happen if the public belief is too strong.

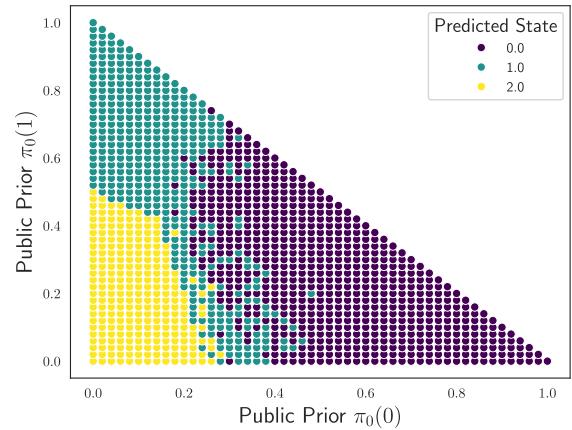


FIGURE 16. Regions in belief space where LLMAs herd during Bayesian social learning with $|\mathcal{X}| = 3$ states. Even though the underlying state is $x = 0$, in the bottom-left and top-left regions, the actions are 2 and 1, respectively, because of a strong public belief.

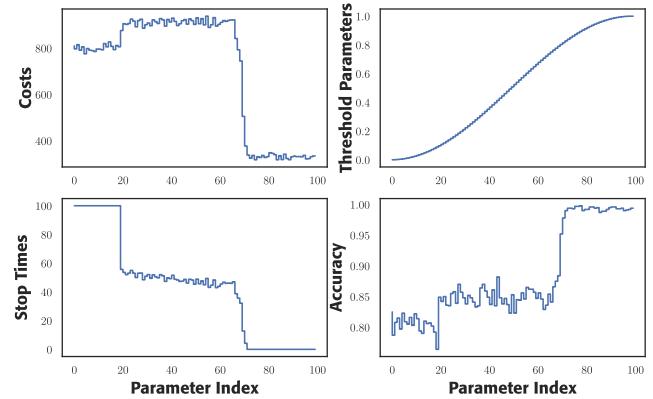


FIGURE 17. The cost function (top-left) corresponds to different values of the threshold parameter (top-right) of the policy of (34). Each policy leads to a different stopping time (bottom-left) and corresponding accuracy (bottom-right). For this experiment, it can be seen that increasing the stopping time decreases the cost. The sudden jumps in the cost (and stopping time) are due to the transition in the resulting policy parameter from a region of learning to herding.

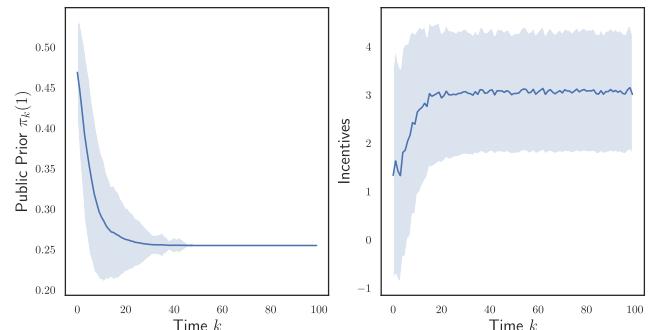


FIGURE 18. Incentivized stochastic control of autonomous LLMAs. The public belief (left) converges slower as a result of incentivization. Also note that the incentive (right) is a supermartingale (as proved in Theorem 5). This is for a policy of the form (35) threshold parameter of $\theta = 0.4$.

E. OPTIMAL STOPPING FOR DELAYING HERDING IN LLMAs

Next, we study the optimal stopping of LLMAs first when these agents are centrally controlled and next when these agents are autonomous.

For both experiments, we consider the product quality identification task. To make sure the actions and the observations are of the same dimensions, we assume that the observation space is also two-dimensional $|\mathcal{Y}| = 2$. The observation matrix was taken to be $B = [[0.7, 0.3], [0.3, 0.7]]$. Note this is necessary since here we ask the LLMAs to either reveal the private observation or to take action using Eq. 14, and update in Eq. 15 needs to be consistent.

For the case when the LLMAs are centrally controlled, and the simplified system model, policies that have a switching threshold curve (Eq. 21) with respect to the belief space can be represented as

$$\mu(\pi) = \begin{cases} 2 \text{ (continue) if } \pi(0) \leq \theta \\ 1 \text{ (stop) if } \pi(0) > \theta \end{cases}, \quad (34)$$

where θ is the threshold parameter. The true state is uniformly sampled from $\{0, 1\}$. The delay cost is taken as $d = 10$, and the misclassification cost is taken as $\Upsilon = 50$. The discount factor is taken as $\rho = 0.99$, and the infinite horizon is approximated by a horizon of length 100. The cost for each individual LLMAs, $c(x, u) = \mathbb{1}(x \neq u)$.

In Figure 17, we perform a parameter sweep over $\theta \in [0, 1]$ with a grid cardinality of 100, and evaluate three different aspects. The results are averaged over 100 independent runs. In the top-left corner, we evaluate the cost function as a function of the threshold parameter. We observe that the cost is minimum for a higher parameter value. And a higher parameter corresponds to a smaller stop time. The accuracy is computed as the proportion of times the correct state is predicted. The sudden jumps in the cost are due to the fact that a small change in the policy parameter is enough to shift the public belief shifts from a region of learning to herding.

Next, we look at LLMAs which are incentivized and autonomous. Here, the incentivization policy is of the form,

$$\mu(\pi) = \begin{cases} 0 \text{ (stop incentivizing) if } \pi(0) \leq \theta \\ \chi(\pi, y) \text{ (incentivize) if } \pi(0) > \theta \end{cases}, \quad (35)$$

where θ and χ is of the form (27). We consider the parameters $\alpha_2 = 1.3$, $\alpha_1 = 0.8$, $\Delta_2 = 0.5$, $\Delta_1 = 0.1$, $\omega_2 = 0.5$ and $\omega_1 = 0.2$. The cost function c is the same as the previous part, and the composite cost function is given by (23).

We first illustrate our result of Theorem 5 in Figure 18, where it can be seen that the sequence of incentives is a submartingale sequence. We fix the threshold parameter of (35) to $\theta = 0.4$. The public belief converges much slower due to the modified cost function (the additional incentive term in 23); hence, using our stopping time formulation, we can extract more private observations, which can be later used to get more accurate estimates.

Figure 19 shows the value of the cost for the different values of the policy threshold θ . We maintain the same system parameters as the previous two experiments. It can be seen that both the classification rate and the cumulative incentive go up as the policy threshold increases in this example. The sudden jump in the incentive is again due to the sudden switch

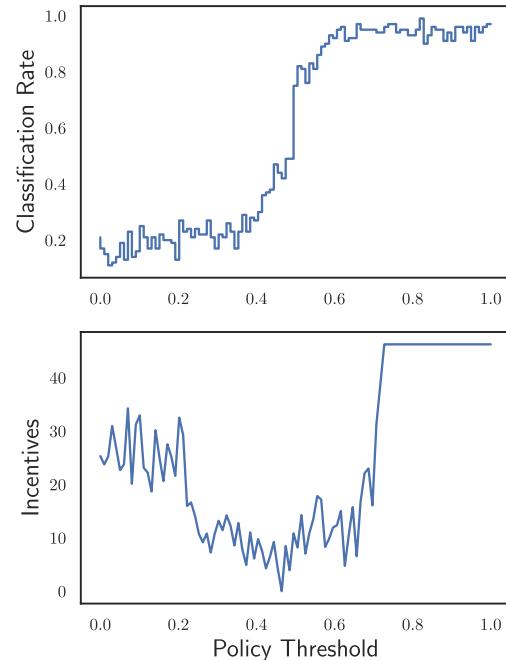


FIGURE 19. The cumulative incentive spent (bottom) and the achieved classification rate (top) for different values of threshold parameter θ . The policy threshold of around 0.7 achieves a tradeoff between the classification rate and the cumulative incentive expenditure. The incentives are constant after a particular parameter because of herding.

between the prior region from learning to herding. However, also note that the threshold around 0.6 has an incentive in the range [5,10], and the classification rate is still decent (0.9-1), therefore showing how there is an optimal threshold that optimally achieves a tradeoff between the cumulative incentive expenditure and classification performance.

F. ILLUSTRATION OF STOCHASTIC GRADIENT ALGORITHM

Finally, we show how the proposed policy gradient algorithm (summarized in Algorithm 7) can be used to optimally estimate the policy parameters for optimal stopping in incentivized LLMAs with system parameters from the previous section. We run 100 iterations, each of which uses a cost function averaged over 100 independent runs. We set a linearly decreasing step size from $\beta = 0.05$ to 0.005 and set the parameter perturbation as $\delta = 1$ and the approximation factor for the sigmoidal policy of (29) as $\varepsilon = 0.3$. The discount factor is $\rho = 0.99$.

Figure 20 presents how the cumulative incentive expenditure changes with each policy parameter update. It can be seen that after a few iterations, the iterates converge to a parameter value such that the incentive expenditure is in the range [5,15], which is close to the optimal incentive from Figure 19.

XII. DISCUSSION AND FUTURE WORK

With the rapid adoption of LLMs across science and engineering and the emergence of LLM-based agents for automating diverse workflows [87], it is crucial to systematically study their behavior and interactions. This

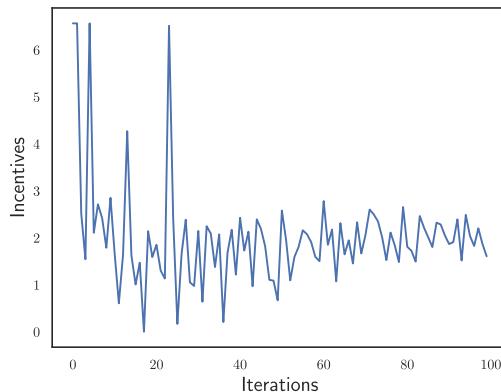


FIGURE 20. The approximate cumulative incentives corresponding to parameters of a policy gradient for estimating the optimal threshold parameters of (26). The incentive of the final iterate is much smaller than the incentive expenditure of a too-small or too-large parameter from Figure 19.

work addresses this need by proposing interpretable models and stochastic control algorithms to engineer systems of LLMA that perform Bayesian inference across a range of applications. Our study highlights the versatility of LLMA across various fields, aiming to inspire practitioners from diverse disciplines to explore and integrate these tools into their workflows.

In this conclusion section, we summarize how the two layers of abstraction explored in this paper using Bayesian revealed preferences and Bayesian social learning provide a sound basis for driving research in understanding and controlling the behavior of LLM agents. We discuss the implications of our methods and their potential applications in studying interacting LLMA. We close with promising research directions at the intersection of signal processing, network science, and machine learning to enhance the capabilities of LLMA.

A. SUMMARY

We first presented a Bayesian sensor model for constructing a large language model agent (LLMA). This model was a composition of a) an LLM, which was used as an interpretable and configurable sensor for high-dimensional data, and b) a Bayesian engine to represent and update the belief of the underlying state. Second, we put forth the necessary and sufficient conditions for a LLMA to be a rationally inattentive Bayesian utility maximizer (RIBUM). We present algorithms for estimating the utility function of a LLMA which is RIBUM. The reconstructed utility naturally leads to the interpretability of the actions of the LLMA. These methods are applicable to both the Bayesian sensor model of LLMA we propose and off-the-shelf LLMA and LLMs. Thirdly we look at Bayesian social learning in a sequence of LLMA. We show that this sequence of agents herd in their actions even if they share their private observations. Then, we formulate optimal stopping problems for failure state detection, which optimally delays herding and improves estimation by allowing the private observations to be shared.

The optimal stopping problem is formulated both for a case when the LLMA are centrally controlled and when they are autonomous. We show that under relatively mild conditions on the observation matrix and the cost function of the LLMA the optimal policy for the optimal stopping problem has a threshold structure. A policy gradient algorithm is proposed to estimate the optimal policy efficiently without the knowledge of system parameters. Finally, we show numerical experiments demonstrating the various claims and frameworks proposed in the study.

B. INSIGHTS

We conclude the paper with the following key takeaways:

- 1) The Bayesian sensor construction of a LLMA lends to a lot more interpretability when performing sequential estimation using a large language model, more than if Bayesian inference was performed with just embeddings (which might be more accurate but not interpretable).
- 2) The Bayesian revealed preferences framework is a systematic way to obtain utility (or cost) functions for blackbox LLMA. It's especially useful when the state and action space are small, but the LLMA operates in tens of different environments.
- 3) For a practitioner, Bayesian social learning might seem simplistic, however the important point is that even a simplistic sequence of LLMA, herding is an undesirable and inevitable phenomena. Therefore, care needs to be taken when creating networks of LLMA that use each other's knowledge to avoid data incest.
- 4) The stopping time formulation is especially useful when deploying LLMA to detect a failure state. The quickest change detection discussed briefly can be practically used to switch optimally between a large (more accurate but more expensive) and a small (less accurate and less expensive) LLM.
- 5) Finally, our numerical experiments are conducted on real-world datasets but are limited to public textual data. Using vision-language models, the framework proposed in this paper can be used for more sophisticated bio-medical data (composed of images, audio, etc.).

C. LIMITATIONS

The Bayesian sensor model developed for a LLMA requires domain knowledge to design the prompts (features) that are input into the LLM, producing a low-dimensional output; this knowledge can be obtained from a larger LLM. However, the Bayesian revealed preference framework becomes computationally infeasible in scenarios with large state and action spaces. Our formulation of stopping time and the structural results presented are focused solely on detecting a single state, specifically for identifying a failure state. While our framework is Bayesian inference centric, LLMA in practice can undertake more complex tasks, indicating a need to extend this work to more general

settings. In the Bayesian revealed preferences framework considered in this paper, access to the true action posterior is assumed; however, in practice, only empirical estimates may be available. Additionally, the applications discussed in this work—finance, online content moderation, and product analysis—may limit the generalizability of the framework to other domains.

D. RESEARCH DIRECTIONS AND FUTURE WORK

With technological innovation in LLMs, many players in the market have launched their own LLM agents. Such agents often have more functionality than those modeled in this paper, and hence, controlling the behavior of these agents would be more challenging. Further, these agents will inevitably interact with each other through the content they generate or on online platforms. Such interaction would lead to a Bayesian update in their beliefs either explicitly by design as modeled in the paper or implicitly as these agents improve. It is therefore important to consider the models presented in this study in future and related research. Given the rise of LLMA, we find the following four research directions particularly exciting.

1) APPLICABILITY IN OTHER PARADIGMS

Exploring the methods presented in this work within different LLM paradigms—such as retrieval-augmented generation (RAG), planning, and fine-tuned agents—would be a compelling area of study. RAG involves retrieving relevant information from a database to provide context for an LLM, enhancing the accuracy and relevance of its responses. The models of LLMA discussed in this paper could be effectively applied in scenarios where multiple agents operate over a knowledge graph to collaboratively retrieve pertinent information [24], particularly when there is a prior defined over the knowledge graph. Additionally, in the context of fine-tuned agents, where agents continuously learn from a private dataset, the framework outlined in this paper can be adapted to accommodate evolving LLMA and heterogeneous agents.

2) APPLICATION IN EDUCATION AND HEALTHCARE

The framework proposed in this paper has broader applicability than content moderation, finance, and opinion mining. We briefly discuss the applications of social learning of LLMA in education and healthcare. In education, there are many potential use cases of LLMs, including customized evaluation, interactive tutorials, and LLM-assisted group discussions [23]. Therefore, a sequence of LLMA can be used to analyze different contents, and studying the control of LLMA is important to reduce the bias, for e.g., while grading students. In healthcare, LLMA are already used for patient interaction to provide instructions for self-care, and to route the patient to the appropriate healthcare provider. However, multiple LLMA can be used to analyze a patient's medical history and provide useful feedback to the doctor. Recently,

authors in [93] have further discussed various opportunities for LLMA in the healthcare domain.

3) SEARCH FOR A UNIFIED INTERPRETABLE MODELS

The interpretable models presented in this paper serve as blackbox models for the LLMA. However, significant research has focused on glassbox (or transparent) models that leverage mechanistic interpretability or utilize explainable features for transformers, the neural architecture underlying LLMs. Integrating these two approaches into a unified model would be advantageous, allowing for explanations of both the LLMA interactions with their external environment and the reasoning behind their behavior. Such a framework could be a valuable tool for analyzing various challenges related to the real-world deployment of LLMA, including reliability and safety.

4) NETWORK OF LLMA AND DATA FUSION

The LLMA in this study are static, homogeneous, and arranged in a line graph. We have considered a homogenous society of LLM agents for notational convenience and for ease of analysis. Our framework allows for different LLM agents to have different observation matrices, this can model LLM agents who have different LLMs, prompts or pre-training data. Additionally, LLMA could possess asymmetric private observations and be fine-tuned on individualized datasets, creating a network akin to a distributed mixture of millions of experts [94]. Another direct extension of our framework accommodates LLM agents with different observation and action spaces, provided that each agent is aware of the observation and action spaces of the preceding agent. A generalized network of LLM agents could also be utilized as Bayesian sensors for distributed state estimation. Furthermore, societies of heterogeneous LLM agents communicating with one another warrant further investigation. These LLMA sensors could then perform data fusion using standard techniques, with careful measures to prevent data incest. This paper has looked at stochastic control methods for controlling herding in a sequence of interacting LLMA. We have formulated two non-standard POMDPs for the centrally controlled and incentivized case. Further to derive structural results, we have looked at the extreme case where there are only two possible actions. These structural results naturally lead to an efficient distributed policy gradient algorithm for searching the optimal policy. It would be interesting to see if newer results from sociology can be used to analyze and control more sophisticated behavior of LLMA when interacting with each other in heterogeneous societies.

5) HUMAN IN THE LOOP WITH LLMA

In applications such as finance, healthcare, and content moderation, LLMA frequently interact with humans to receive feedback that can greatly enhance task performance [95]. This necessitates advanced models that account

for human-agent interactions to better align cost functions, communication protocols, and sensing mechanisms. As these LLMAs become more widely deployed across various applications, they will form natural networks with humans. Investigating how these LLMAs can share private information both reliably and securely to enhance usability and efficiency presents a compelling and open research challenge.

APPENDIX

PROOFS

A. PROOF OF THEOREM 1

Proof of the necessity of NIAS and NIAC: The proof follows closely from proof of [46], and is reproduced for reader's aid. Similar proofs can also be found in [1], [5].

- 1) **NIAS (8):** For environment $m \in \mathcal{M}$, define the subset $\mathcal{Y}_u \subseteq \mathcal{Y}$ so that for any observation $y \in \mathcal{Y}_u$, given posterior probability mass function (pmf) $\mathbb{P}_m(x|y)$, the choice of action is u (5) is maximum. Define the revealed posterior pmf given action u as $\mathbb{P}_m(x|u)$. The revealed posterior pmf is a stochastically garbled version of the actual posterior pmf $\mathbb{P}_m(x|y)$, that is,

$$\mathbb{P}(x|u) = \sum_{y \in \mathcal{Y}} \frac{p_m(x, y, u)}{p_m(u)} = \sum_{y \in \mathcal{Y}} p_m(y|u) p_m(x|y). \quad (36)$$

Since the optimal action is u for all $y \in \mathcal{Y}_u$, (5) implies:

$$\begin{aligned} & \Rightarrow \sum_{y \in \mathcal{Y}_u} p_m(y|u) \sum_{x \in \mathcal{X}} p_m(x|y) (r_m(x, \tilde{u}) - r_m) \leq 0 \\ & \Rightarrow \sum_{y \in \mathcal{Y}} p_m(y|u) \sum_{x \in \mathcal{X}} p_m(x|y) (r_m(x, \tilde{u}) - r_m) \leq 0 \\ & \quad (\text{since } p_m(y|u) = 0, \forall y \in \mathcal{Y} \setminus \mathcal{Y}_u) \\ & \Rightarrow \sum_{x \in \mathcal{X}} p_m(x|u) (r_m(x, \tilde{u}) - r_m) \leq 0 \text{ (from (36))}. \end{aligned}$$

The last equation is the NIAS inequality (8).

- 2) **NIAC (9):** Let $c_m = Z(B(m)) > 0$, where $Z(\cdot)$ denotes the information acquisition cost of the collection of agents \mathcal{M} . Also, let $J(B(m), r_m)$ denote the expected utility of the RIBUM in environment m given attention strategy $B(m)$ (first term in RHS of (3)). Here, the expectation is taken wrt both the state x and observation y . It can be verified that $J(\cdot, r_m)$ is convex in the first argument. Finally, for the environment m , we define the revealed attention strategy $B(m)'$ over the set of actions \mathcal{U} as $B(m)'(u|x) = \mathbb{P}_m(u|x)$, $\forall u \in \mathcal{U}$, where the variable $\mathbb{P}_m(u|x)$ is obtained from the dataset \mathbb{D} . Clearly, the revealed attention strategy is a stochastically garbled version of the true attention strategy since

$$B(m)'(u|x) = \mathbb{P}_m(u|x) = \sum_{y \in \mathcal{Y}} p_m(u|y) B(m)(y|x). \quad (37)$$

From Blackwell dominance [96] and the convexity of the expected utility functional $J(\cdot, r_m)$, it follows that:

$$J(B'_m, r_l) \leq J(B_m, r_l), \quad (38)$$

when B_m Blackwell dominates B'_m . The above relationship holds with equality if $m = l$ (this is due to NIAS (8)).

We now turn to condition (3) for optimality of attention strategy.

The following inequalities hold for any pair of agents $l \neq m$:

$$\begin{aligned} J(B'_m, r_m) - c_m & \stackrel{(38)}{=} J(B(m), r_m) - c_m \\ & \stackrel{(3)}{\geq} J(B(l), r_m) - c_l \stackrel{(38)}{\geq} J(B(l)', r_m) - c_l. \end{aligned} \quad (39)$$

This is precisely the NIAC inequality (9).

Proof for sufficiency of NIAS and NIAC: Let $\{r_m, c_m\}_{m=1}^M$ denote a feasible solution to the NIAS and NIAC inequalities of Theorem 1. To prove sufficiency, we construct an RIBUM tuple as a function of dataset \mathbb{D} and the feasible solution that satisfies the optimality conditions (5),(3) for RIBUM (6).

Consider the following RIBUM tuple:

$$\begin{aligned} & (\mathcal{M}, \mathcal{X}, \mathcal{Y} = \mathcal{U}, \mathcal{U}, \pi, Z, \{\mathbb{P}_m(u|x), r_m, m \in \mathcal{M}\}), \text{ where} \\ & Z(\mathbb{P}(u|x)) = \max_{m \in \mathcal{M}} c_m + J(\mathbb{P}(u|x), r_m) - J(\mathbb{P}_m(u|x), r_m). \end{aligned} \quad (40)$$

In (40), $C(\cdot)$ is a convex cost since it is a point-wise maximum of monotone convex functions. Further, since NIAC is satisfied, (40) implies $Z(B(m)) = c_m$. It only remains to show that inequalities (5) and (3) are satisfied for all environments in \mathcal{M} .

- 1) **NIAS implies (5) holds.** This is straightforward to show since the observation and action sets are identical. From NIAS (8), we know that for any environment $m \in \mathcal{M}$, actions $u, \tilde{u} \in \mathcal{U}, u \neq \tilde{u}$, the following inequalities hold.

$$\begin{aligned} & \sum_x \mathbb{P}_m(x|u) (r_m(x, \tilde{u}) - r_m) \leq 0 \\ & \Rightarrow \sum_x \mathbb{P}_m(x|y = u) (r_m(x, \tilde{u}) - r_m) \leq 0 \\ & \Rightarrow u \in \arg \max_{\tilde{u} \in \mathcal{U}} \sum_x \mathbb{P}_m(x|y) r_m(x, \tilde{u}) \Rightarrow (5). \end{aligned}$$

- 2) **Information Acquisition Cost (40) implies (3) holds.** Fix environment $l \in \mathcal{M}$. Then, for any attention strategy $\mathbb{P}(u|x)$, the following inequalities hold.

$$\begin{aligned} & Z(\mathbb{P}(u|x)) \\ & = \max_{m \in \mathcal{M}} c_m + J(\mathbb{P}(u|x), r_m) - J(\mathbb{P}_m(u|x), r_m) \\ & \Rightarrow J(\mathbb{P}_l(u|x)) - c_l \geq J(\mathbb{P}(u|x)) - C(\mathbb{P}(u|x)), \forall \mathbb{P}(u|x) \\ & \Rightarrow \mathbb{P}_m(u|x) \in \arg \max_{\mathbb{P}(u|x)} J(\mathbb{P}(u|x), r_m) - Z(\mathbb{P}(u|x)) \\ & = (3). \end{aligned}$$

B. DERIVATION OF SOCIAL LEARNING FILTER

Let the posterior as $\eta_k(i) = \mathbb{P}(x_k = i|u_1, \dots, u_k)$. Let $\sigma(\eta_{k-1}, u_k) = \sum_i \sum_y \mathbb{P}(u_k|y_k = y, \eta_{k-1}) \mathbb{P}(y_k|x_k = i) \sum_j P_{i,j} \eta_{k-1}(j)$ be the normalization factor.

$$\begin{aligned} \eta_k(i) &= \frac{1}{\sigma(\eta_{k-1}, u_k)} \mathbb{P}(u_k|x_k = i, u_1, \dots, u_{k-1}) \\ &= \sum_j P_{i,j} \mathbb{P}(x_{k-1} = j|u_1, \dots, u_{k-1}) \\ &= \frac{1}{\sigma(\eta_{k-1}, u_k)} \sum_y \mathbb{P}(u_k|y_k = y, \eta_{k-1}) \mathbb{P}(y_k = y|x_k = i) \\ &= \sum_j P_{i,j} \eta_{k-1}(j). \end{aligned}$$

which completes the derivation.

C. PROOF FOR THEOREM 2

Proof: Define $\Lambda_k(i, j) = \log(\pi(i)/\pi(j))$, $i, j \in \mathcal{X}$. From (15) we have, $\Lambda_{k+1}(i, j) = \Lambda_k(i, j) + \Gamma_k(i, j)$ where $\Gamma_k(i, j) = \log(\mathbb{P}(u_k|x = i, \pi_k)/\mathbb{P}(u_k|x = j, \pi_k))$.

The probability of the actions given the state and prior can be written as,

$$\mathbb{P}(u|x, \pi) = \sum_{y \in \mathcal{Y}} \prod_{\tilde{u} \in \mathcal{U} \setminus u} \mathbb{1}(c'_{\tilde{u}} B_y \pi \leq c'_{\tilde{u}} B_y \pi) B_{y,x}.$$

Let $\tilde{\mathcal{Y}}_k \subseteq \mathcal{Y}$ be a subset of the observation space for which the action u_k is suboptimal with respect to all other actions, i.e., $\prod_{\tilde{u} \in \mathcal{U} \setminus u_k} \mathbb{1}(c'_{\tilde{u}} B_y \pi > c'_{\tilde{u}} B_y \pi) = 1 \forall y \in \mathcal{Y}$. When information cascade (Def. 1) occurs, this set should be empty since no matter what the observation, the action k should be optimal according to (14). Also, rewriting $\Gamma_k(i, j)$,

$$\Gamma_k(i, j) = \log \left(\frac{1 - \sum_{y \in \tilde{\mathcal{Y}}_k} B_{y,i}}{1 - \sum_{y \in \tilde{\mathcal{Y}}_k} B_{y,j}} \right).$$

Therefore when an information cascade occurs, $\Gamma_k(i, j) = 0, \forall i, j \in \mathcal{X}$ (Due to $\tilde{\mathcal{Y}}_k$ being an empty set). Also if $\tilde{\mathcal{Y}}_k$ is nonempty, then $\Gamma_k(i, j) > \kappa$, where κ is a positive constant.

Let $\mathcal{G}_k = \{\sigma(u_1, u_2, \dots, u_k)\}$ denote the natural filtration, where σ is the operator which generates the corresponding sigma field.

$\pi_k(i) = \mathbb{P}(x = i|u_1, \dots, u_k) = \mathbb{E}[\mathbb{1}(x = i)|\mathcal{G}_k]$ is a martingale adapted to \mathcal{G}_k for all $i \in \mathcal{X}$. This follows by the application of smoothing property of conditional expectation, $\mathbb{E}[\pi_{k+1}(i) | \mathcal{G}_k] = \mathbb{E}[\mathbb{E}[\mathbb{1}(x = i) | \mathcal{G}_{k+1}]] = \mathbb{E}[\mathbb{1}(x = i) | \mathcal{G}_k]$.

Therefore, by the martingale convergence theorem, there exists a random variable π_∞ such that, $\pi_k \rightarrow \pi_\infty$ w.p. 1. Therefore $\Lambda_k(i, j) \rightarrow \Lambda_\infty(i, j)$ w.p. 1., which implies there exists \tilde{k} such that $\forall k \geq \tilde{k}$, $|\Lambda_\infty(i, j) - \Lambda_k(i, j)| \leq \kappa/3$ and so,

$$|\Lambda_{k+1}(i, j) - \Lambda_k(i, j)| \leq 2\kappa/3, \forall k \geq \tilde{k}. \quad (41)$$

We now prove the theorem by contradiction. Suppose a cascade does not occur, then for at least one pair $i \neq j, i, j \in \mathcal{X}$, $\mathbb{P}(u|x = i, \pi)$ is different than $\mathbb{P}(u|x = j, \pi)$. This would imply that the set $\tilde{\mathcal{Y}}_k$ is nonempty and therefore,

$$|\Lambda_k(i, j)| = |\Lambda_{k+1}(i, j) - \Lambda_k(i, j)| \geq \kappa. \quad (42)$$

(41) and (42) contradict each other. Therefore $\mathbb{P}(u|x = i, \pi)$ is same for all $i \in \mathcal{X}$ and hence according to (15) information cascade occurs at time \tilde{k} . \square

D. PROOF FOR THEOREM 3

Proof: We prove the Theorem by showing that it satisfies the conditions of Theorem 12.3.4 of [5]. A more general proof can be found in [4], [5].

In order to verify the assumptions of Theorem 12.3.4 of [5], we need to define first-order stochastic dominance (FOSD) and a submodular function. We first define a Monotone Likelihood Ratio (MLR) ordering on a line and then define a submodular function with respect to this MLR ordering. We only need to consider the following lines,

$$\mathcal{L}(e_i, \bar{\pi}) = \{\pi \in \mathcal{P}(X) : \pi = (1 - \epsilon)\bar{\pi} + \epsilon e_i, 0 \leq \epsilon \leq 1\}, \quad \bar{\pi} \in \mathcal{H}_i,$$

where the state index is only between the extreme states, $i \in \{1, X\}$ and,

$$\mathcal{H}_i = \{\pi \in \mathcal{P}(X) : \mu(i) = 0\}.$$

To define the MLR ordering on a line, we first define the MLR ratio with respect to belief space,

Definition 3 (Monotone Likelihood Ratio (MLR) Order): Let $\pi_1, \pi_2 \in \mathcal{P}(X)$, then π_1 dominates π_2 with respect to the MLR order ($\pi_1 \geq_r \pi_2$) if,

$$\pi_1(i)\pi_2(j) \leq \pi_1(j)\pi_2(i), \quad i < j, i, j \in \{1, \dots, X\}.$$

The following definition is for the MLR ordering the lines \mathcal{L}_{e_i} , $i \in \{1, X\}$

Definition 4 (MLR Ordering on Line \geq_{L_i}): μ_1 is greater than μ_2 with respect to the MLR ordering on the line $\mathcal{L}(e_i, \pi)$, $i \in \{1, X\}$ ($\mu_1 \geq_{L_i} \mu_2$), if $\pi_1, \pi_2 \in \mathcal{L}(e_i, \bar{\pi})$ for some $\bar{\pi}$ and $\mu_1 \geq_r \mu_2$.

Finally, we are ready to define a submodular function on a line,

Definition 5 (Submodular Function on Line): For $i \in \{1, X\}$, a function $\phi : \mathcal{L}(e_i, \bar{\pi}) \times \mathcal{D} \rightarrow \mathbb{R}$ is submodular if $\phi(\pi, a) - \phi(\pi, \bar{a}) \leq \phi(\bar{\pi}, a) - \phi(\bar{\pi}, \bar{a})$, for $\bar{a} \leq a, \bar{\pi} \leq_{L_i} \pi$.

The following is used extensively to compare two beliefs and is a weaker condition than MLR ordering,

Definition 6 (First Order Stochastic Dominance (FOSD)): Let $\pi_1, \pi_2 \in \mathcal{P}(X)$, then π_1 first order stochastically dominates π_2 ($\pi_1 \geq_s \pi_2$) if $\sum_{i=j}^X \pi_1(i) \geq \sum_{i=j}^X \pi_2(i) \forall j \in \mathcal{X}$.

The stopping time problem with the social welfare cost of (19) can be decomposed into two cost terms, each corresponding to the cost terms of the stopping time problem

for partially observed Markov decision processes of Theorem 12.3.4 of [5].

$$\begin{aligned} C(\pi, 1) &= \frac{1}{1-\rho} \min_u c'_u \pi, \\ C(\pi, 2) &= \sum_{y \in \mathcal{Y}} c'_y B_y \pi + (d + (1-\rho)\Upsilon) e_1 \pi - (1-\rho)\Upsilon. \end{aligned} \quad (43)$$

$C(\pi, 1)$ is the expected cost after the decision to herd has been made. Similarly, the first term in $C(\pi, 2)$ is the expected cost when revealing private observations. The rest of the terms come from transforming the value function by the delay penalty costs [4].

We now state the main assumptions of Theorem 12.3.4 of [5], which are required for this to hold for the structural result of (21),

- 1) **(C)** $\pi_1 \geq_s \pi_2$ implies $C(\pi_1, u) \leq C(\pi_2, u)$.
- 2) **(F1)** B is totally positive of order 2 (TP2).
- 3) **(F2)** P is totally positive of order 2 (TP2).
- 4) **(S)** $C(\pi, a)$ is submodular on $[\mathcal{L}(e_X, \bar{\pi}), \geq_{L_X}]$ and $[\mathcal{L}(e_1, \bar{\pi}), \geq_{L_1}]$.

F1 follows from S4, and F2 follows from the fact that we only consider an identity transition matrix P . And since the costs are linear, (C) follows by applying the definition of FOSD on equation (43) and using assumption (S1). (S) follows from the definition of MLR ordering on the line, using the fact that B is TP2 in the first term of $C(\pi, 2)$ and using (S2) and (S3). Hence, the assumptions are verified, and the structural result is proved. \square

E. PROOF OF OPTIMAL STOPPING THEOREM 4

This proof follows closely from the proof in [2]. From Lemma 1 of [2], the value function can be expressed as,

$$V(\pi) = \min\{0, \chi(y) - g + \rho \mathbb{E}V(\pi)\}.$$

Further, Lemma 2 of [2] shows that the incentive function is decreasing. Denote V_m to be the m -th iterate of the value iteration algorithm [5] which iteratively converges to the value function $V(\pi)$. The iterates are given by,

$$V_{m+1} = \min\{0, \chi(y) - g + \rho \mathbb{E}V(\pi)\}.$$

From the definition of first order stochastic dominance and Proposition 1 of [2] $\mathbb{E}V_m(\pi)$ is decreasing in π .

Therefore $V_m(\pi)$ is decreasing which implies $V(\pi)$ is decreasing. Let $V(0)$ and $V(1)$ be the value at $\pi = [1, 0]$ and $\pi = [0, 1]$ respectively which makes $\mathbb{E}V(0) = V(0)$ and $\mathbb{E}V(1) = V(1)$.

- 1) For $V = \chi(y) - g + \rho \mathbb{E}V(\pi)$, $V(0) = \frac{\chi(e_1) - g}{1-\rho} > 0$ and $V(1) = \frac{\chi(e_2) - g}{1-\rho} < 0$.
- 2) For $V = 0$, $V(0) = V(1) = 0$

Therefore, the value function decreases with a positive value of $[1, 0]$ and a negative value of $[0, 1]$. Therefore it must be 0 at some time. Since $V(\pi)$ is monotone in π , the set $E = \{\pi(2) | V(\pi) = \chi(y) - g + \rho \mathbb{E}V(\pi)\}$ is convex. We choose a policy $\bar{\pi}^*(2) = \{\hat{\pi}(2) | \hat{\pi}(2) > \pi(2) \forall \pi(2) \in E\}$.

F. PROOF OF SUBMARTINGALE RESULT THEOREM 5

Consider the suboptimal policy where $\epsilon > 0$,

$$\hat{\nu}(\pi) = \begin{cases} \chi(\pi) - \epsilon & \text{if } \pi(2) \leq \bar{\pi}^*(2) \\ \chi(\pi) & \text{if } \pi(2) > \bar{\pi}^*(2) \end{cases}.$$

From Lemma 3 of [2] $\chi(\pi)$ is convex in π . We know that the public belief π_k is a martingale from the proof of Theorem 2. Let $W_k = \hat{\nu}(\pi_{k-1})$. By Jensen's inequality for $\epsilon \rightarrow 0$,

$$\begin{aligned} \mathbb{E}[W_{k+1} | \mathcal{F}_k] &= \mathbb{E}[\chi(\pi_{k+1}) | \mathcal{F}_k] \geq \chi(\mathbb{E}[\pi_{k+1} | \mathcal{F}_k]) \\ &\geq \chi(\pi_k) \geq W_k. \end{aligned}$$

Therefore W_k is a submartingale. Consider H_k as the following sequence,

$$H_k = \begin{cases} 0 \text{ if } \pi_{k-1}(2) \leq \bar{\pi}^*(2) \\ 1 \text{ if } \pi_{k-1}(2) \geq \bar{\pi}^*(2) \end{cases}.$$

Now, by properties of submartingales, $(HW)_k$ is a submartingale. This is exactly the incentive sequence, which is indeed a submartingale.

To show the second statement, we use Doob's martingale inequality on the submartingale sequence p_k ,

$$\mathbb{P}\left[\max_{1 \leq k \leq T} p_k \geq C\right] \leq \frac{\mathbb{E}[\max(p_T, 0)]}{C} = \frac{\mathbb{E}[p_T]}{C}.$$

Also note that $\sum_1^T p_k \leq T \max_{1 \leq k \leq T} p_k$. Therefore the event $\{\sum_1^T p_k \geq CT\} \subseteq \{T \max_{1 \leq k \leq T} p_k \geq CT\}$.

Hence $\mathbb{P}(\sum_1^T p_k \geq CT) \leq \mathbb{P}(T \max_{1 \leq k \leq T} p_k \geq CT)$ which along with the first inequality proves the statement.

APPENDIX

BRIEF EXPERIMENTAL DETAILS

A. HYPERPARAMETERS OF LLM AT INFERENCE TIME

As described in the main text, we use different LLMs.

Large Language models uses default parameters for inference to control the output. Temperature (default 1) influences randomness, with lower values producing more focused responses. Max tokens (default 2048) limits response length. Top p (default 1) controls diversity, with lower values making the output more focused by considering only the most likely words.

1) MIXTRAL-8 \times 7B-v0.1

We consider the maximum response tokens as 100, a temperature of 0.7, a top-p of 0.7, a top-k of 50, and a repetition penalty as 50.

2) LLaMA-3-70b

We consider the maximum response tokens as 200, a temperature of 0.7, a top-p of 0.7, a top-k of 50, and a repetition penalty as 80.

We use the TogetherAI API to send the prompt to the LLM and receive the response for LLama3 and Mixtral.

3) ChatGPT-40

Default parameters were used when inferencing ChatGPT using their API.

B. DATASET DESCRIPTION

1) AMAZON REVIEW DATASET

Amazon review data [56], including 233.1 million reviews spanning various product categories. It features detailed information such as ratings, review text, helpfulness votes, product descriptions, category information, prices, brands, and image features. It also has additional transaction metadata such as product color, size, and user-submitted images, as well as more detailed product landing page information like bullet-point descriptions and technical specifications.

2) MEASURING HATESPEECH DATASET

The Measuring Hate Speech corpus is a comprehensive dataset designed to evaluate hate speech while accounting for annotators' perspectives [54]. It includes 50,070 social media comments from platforms like YouTube, Reddit, and Twitter, labeled by 11,143 annotators from Amazon Mechanical Turk. Each comment is assessed on 10 ordinal labels such as sentiment, disrespect, and violence, which are aggregated into a continuous score using Rasch measurement theory. This approach allows for the statistical summarization of annotator disagreement and adjusts for labeling strictness, providing a nuanced measure of hate speech. The dataset also includes information on the identity group targets of each comment and annotator demographics, enabling detailed analyses of identity-related perspectives.

3) JigsawAI UNINTENDED BIAS DATASET

This dataset was used to train the likelihood network since we assume the actual dataset is not available during real-time inferencing.

The Civil Comments dataset, comprising approximately 2 million public comments from the now-defunct Civil Comments platform, provides a valuable resource for investigating online toxicity. Jigsaw, the sponsor of this effort, facilitated human annotation of these comments for various toxic conversational attributes, including "toxicity," "severe-toxicity," "obscene," "threat," "insult," "identityattack," and "sexualexplicit." Each comment's toxicity was assessed by up to 10 annotators who rated it on a scale ranging from "Not Toxic" to "Very Toxic," with the final toxicity label representing the fraction of annotators who deemed it toxic. Furthermore, a subset of comments were labeled for identity mentions, such as gender, sexual orientation, religion, and race, to analyze the relationship between online toxicity and identity. This dataset offers significant potential for developing and evaluating models aimed at identifying and mitigating harmful online interactions. The dataset can

be accessed here kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification/data.

4) FNNSPID

FNNSPID (Financial News and Stock Price Integration Dataset) is a large-scale dataset designed for financial market prediction. It integrates both numerical and textual data, comprising 29.7 million stock prices and 15.7 million financial news articles for 4,775 S&P500 companies from 1999 to 2023 [78]. Sourced from four major stock market news websites, FNNSPID includes sentiment scores derived from the news articles, offering a unique resource for researchers to investigate the impact of news sentiment on market trends. The metric used in Figure 5 is given by,

$$\frac{C_{k+1} - C_k}{C_k - C_{k-1}},$$

where C_k is the close price at time k . The close price is the stock price at the end of the trading day.

C. SYSTEM PROMPTS

1) FINANCIAL ANALYSIS TASK

For the Financial Analysis Task of Example 2.

"Analyze the article and answer the following questions based on the content:

Are there indications that recent or upcoming policy decisions could support market growth? (Yes/No)

Do statements from central banks suggest optimism about the economic outlook? (Yes/No)

Are there emerging trends or patterns that suggest a shift in market sentiment? (Yes/No)

Is there evidence of key technical levels acting as support for major indices? (Yes/No)

Are certain sectors or industries showing stronger performance compared to others? (Yes/No)

Do shifts in investor interest suggest a move toward specific sectors, such as technology or energy? (Yes/No)

Do recent economic data releases (e.g., employment, inflation, consumer sentiment) point toward growth? (Yes/No)

Are any indicators flashing signals that typically correlate with significant market moves (e.g., yield curves, commodity prices)? (Yes/No)

Is there evidence of a "risk-on" approach among investors? (Yes/No)

Do recent market movements suggest increased interest in safe-haven assets like gold or bonds? (Yes/No)

Are there global or geopolitical events mentioned that could influence market volatility? (Yes/No)

Could changes in international markets or currencies impact domestic market trends? (Yes/No)

Are recent corporate earnings or business announcements likely to influence market sentiment? (Yes/No)

Do specific companies or sectors appear to be driving recent market gains? (Yes/No) *article*"

where *article* contains the HTML page of the online news article.

2) HATE SPEECH CLASSIFICATION TASK

For the Hate speech Classification Task:

“[INST]

Return a JSON with the following format for the given text:

```
{‘is_insulting’: Bool,  
‘is_dehumanizing’:Bool,  
‘is_humiliating’:Bool,  
‘promotes_violence’:Bool,  
‘promotes_genocide’:Bool,  
‘is_respectful’:Bool}
```

Text: {comment }/[INST]”

where “comment” contains the comment observation which needs to be analyzed.

3) PRODUCT QUALITY IDENTIFICATION TASK

For the Product Quality Identification Task:

“Analyze the following product review and provide a summary of the key points:

- Does the review mention any specific problems or defects with the product?
- Does the review mention any positive attributes regarding the product’s durability or reliability?
- Does the review indicate that the product meets or exceeds the user’s expectations?
- Would the reviewer recommend this product to others?

where review contains the review which is being analyzed.

4) RESPONSE

The responses of the LLM includes the JSON response at the start and an explanation for the corresponding mapping. We truncated the output to include the JSON and got a discrete low-dimensional observation from the textual comment.

D. OBSERVATION SPACE

1) REDUCING THE OBSERVATION SPACE

We reduce the observation space for the hate-speech classification task:

Although the LLM output for text observation is of cardinality $2^6 = 64$, we reduce by considering the following order: respectful < insulting < dehumanizing < humiliating < violence < genocide [54]. The binary map $\psi(z) : \{0, 1\}^6 \rightarrow 6 = \max\{j : \text{s.t. } z[j] = 1\}$ takes the observation as the highest severity present in the binary observation.

2) LIKELIHOOD NEURAL NETWORK (RESTRICTED BOLTZMANN MACHINE)

We use a subset of the labeled dataset to obtain a likelihood distribution. We use restricted Boltzmann Machines (RBMs) to approximate the likelihood function $P(y|x)$; we train a conditional RBM on observations obtained from different states (as defined in the main text). Each RBM has $|\mathcal{Y}|$ visible units and 4 hidden units. We train the RBM using contrastive divergence for 100 epochs and generate 1000 samples

using Gibbs sampling with 1000 iterations. We obtain the approximate probabilities by empirically counting the samples.

REFERENCES

- [1] A. Caplin and M. Dean, “Revealed preference, rational inattention, and costly information acquisition,” *Amer. Econ. Rev.*, vol. 105, no. 7, pp. 2183–2203, Jul. 2015.
- [2] S. Bhatt and V. Krishnamurthy, “Controlled sequential information fusion with social sensors,” *IEEE Trans. Autom. Control*, vol. 66, no. 12, pp. 5893–5908, Dec. 2021.
- [3] C. Chamley, *Rational Herds: Economic Models of Social Learning*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [4] V. Krishnamurthy, “Quickest detection POMDPs with social learning: Interaction of local and global decision makers,” *IEEE Trans. Inf. Theory*, vol. 58, no. 8, pp. 5563–5587, Aug. 2012.
- [5] V. Krishnamurthy, *Partially Observed Markov Decision Processes*. Cambridge, U.K.: Cambridge Univ. Press, Mar. 2016.
- [6] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed., London, U.K.: Chapman & Hall, 2013.
- [7] J. Wiesinger, P. Marlow, and V. Vuskovic. (Sep. 2024). *Agents*. Google/Kaggle. [Online]. Available: <https://www.kaggle.com/whitepaper-agents>
- [8] S. Altman. (Jan. 2025). *Reflections—Blog*. [Online]. Available: <https://blog.samaltman.com/reflections>
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. 31st Adv. Neural Inf. Process. Syst.*, vol. 30, Red Hook, NY, USA, Jun. 2017, pp. 6000–6010.
- [10] B. Min, H. Ross, E. Sulem, A. P. B. Veyeh, T. H. Nguyen, O. Sainz, E. Agirre, I. Heintz, and D. Roth, “Recent advances in natural language processing via large pre-trained language models: A survey,” *ACM Comput. Surv.*, vol. 56, no. 2, pp. 1–40, Sep. 2023.
- [11] S. Li, X. Puig, C. Paxton, Y. Du, C. Wang, L. Fan, T. Chen, D.-A. Huang, E. Akyürek, A. Anandkumar, J. Andreas, I. Mordatch, A. Torralba, and Y. Zhu, “Pre-trained language models for interactive decision-making,” in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, Dec. 2022, pp. 31199–31212.
- [12] S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, and J. Gao, “Deep learning-based text classification: A comprehensive review,” *ACM Comput. Surv.*, vol. 54, pp. 1–40, Apr. 2021.
- [13] Finance. 90% of Online Content Could be Generated by AI by 2025, Expert Says. Accessed: Jan. 15, 2025. [Online]. Available: <https://finance.yahoo.com/news/90-of-online-content-could-be-generated-by-ai-by-2025-expert-says-201023872.html>
- [14] W. Zhou, Y. E. Jiang, L. Li, J. Wu, T. Wang, S. Qiu, J. Zhang, J. Chen, R. Wu, S. Wang, S. Zhu, J. Chen, W. Zhang, X. Tang, N. Zhang, H. Chen, P. Cui, and M. Sachan, “Agents: An open-source framework for autonomous language agents,” 2023, *arXiv:2309.07870*.
- [15] I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal, “AI models collapse when trained on recursively generated data,” *Nature*, vol. 631, no. 8022, pp. 755–759, Jul. 2024.
- [16] A. Jain and V. Krishnamurthy, “Identifying hate speech peddlers in online platforms. A Bayesian social learning approach for large language model driven decision-makers,” 2024, *arXiv:2405.07417*.
- [17] Salesforce. (Sep. 2024). *Salesforce Unveils Agentforce—What AI was Meant to be*. [Online]. Available: <https://investor.salesforce.com/press-releases/press-release-details/2024/Salesforce-Unveils-AgentforceWhat-AI-Was-Meant-to-Be/>
- [18] L. Wang, X. Zhang, H. Su, and J. Zhu, “A comprehensive survey of continual learning: Theory, method and application,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 8, pp. 5362–5383, Aug. 2024.
- [19] J. Kaddour, J. Harris, M. Mozes, H. Bradley, R. Raileanu, and R. McHardy, “Challenges and applications of large language models,” 2023, *arXiv:2307.10169*.
- [20] M. A. K. Raiaan, M. S. H. Mukta, K. Fatema, N. M. Fahad, S. Sakib, M. M. J. Mim, J. Ahmad, M. E. Ali, and S. Azam, “A review on large language models: Architectures, applications, taxonomies, open issues and challenges,” *IEEE Access*, vol. 12, pp. 26839–26874, 2024.
- [21] B. Jin, G. Liu, C. Han, M. Jiang, H. Ji, and J. Han, “Large language models on graphs: A comprehensive survey,” *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 12, pp. 8622–8642, Dec. 2024.

[22] R. Gallotta, G. Todd, M. Zammit, S. Earle, A. Liapis, J. Togelius, and G. N. Yannakakis, "Large language models and games: A survey and roadmap," *IEEE Trans. Games*, early access, Sep. 13, 2024, doi: [10.1109/TG2024.3461510](https://doi.org/10.1109/TG2024.3461510).

[23] E. Kasneci et al., "ChatGPT for good? On opportunities and challenges of large language models for education," *Learn. Individual Differences*, vol. 103, Mar. 2023, Art. no. 102274.

[24] R. Reinanda, E. Meij, and M. de Rijke, "Knowledge graphs: An information retrieval perspective," *Found. Trends Inf. Retr.*, vol. 14, no. 4, pp. 289–444, 2020.

[25] T. Guo, X. Chen, Y. Wang, R. Chang, S. Pei, N. V. Chawla, O. Wiest, and X. Zhang, "Large language model based multi-agents: A survey of progress and challenges," in *Proc. 33rd Int. Joint Conf. Artif. Intell.*, Aug. 2024, pp. 8048–8057.

[26] Y. Li et al., "Personal LLM agents: Insights and survey about the capability, efficiency and security," 2024, *arXiv:2401.05459*.

[27] Y. Yu, H. Li, Z. Chen, Y. Jiang, Y. Li, D. Zhang, R. Liu, J. W. Suchow, and K. Khashanah, "FinMem: A performance-enhanced LLM trading agent with layered memory and character design," in *Proc. AAAI Symp. Ser.*, vol. 3, May 2024, pp. 595–597.

[28] S.-L. Hsu, R. S. Shah, P. Senthil, Z. Ashktorab, C. Dugan, W. Geyer, and D. Yang, "Helping the helper: Supporting peer counselors via AI-empowered practice and feedback," 2023, *arXiv:2305.08982*.

[29] Y. Guan, D. Wang, Z. Chu, S. Wang, F. Ni, R. Song, and C. Zhuang, "Intelligent agents with LLM-based process automation," in *Proc. 30th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, New York, NY, USA, Aug. 2024, pp. 5018–5027.

[30] S. Suri, S. N. Das, K. Singi, K. Dey, V. S. Sharma, and V. Kaulgud, "Software engineering using autonomous agents: Are we there yet?" in *Proc. 38th IEEE/ACM Int. Conf. Automated Softw. Eng. (ASE)*, Sep. 2023, pp. 1855–1857.

[31] S. Yao, H. Chen, J. Yang, and K. Narasimhan, "WebShop: Towards scalable real-world web interaction with grounded language agents," in *Proc. 36th Adv. Neural Inf. Process. Syst.*, vol. 35, Apr. 2024, pp. 20744–20757.

[32] T. R. Sumers, S. Yao, K. Narasimhan, and T. L. Griffiths, "Cognitive architectures for language agents," in *Proc. Trans. Mach. Learn. Res.*, Jan. 2023, pp. 1–12. [Online]. Available: <https://openreview.net/forum?id=1i6ZCvflQJ>

[33] M. Zhuge et al., "Mindstorms in natural language-based societies of mind," 2023, *arXiv:2305.17066*.

[34] Z. Liu, Y. Zhang, P. Li, Y. Liu, and D. Yang, "A dynamic LLM-powered agent network for task-oriented agent collaboration," in *Proc. 1st Conf. Lang. Model.*, Jul. 2024, pp. 1–14. [Online]. Available: <https://openreview.net/forum?id=XII0Wp1XA9>

[35] Z. Wang, S. Cai, A. Liu, Y. Jin, J. Hou, B. Zhang, H. Lin, Z. He, Z. Zheng, Y. Yang, X. Ma, and Y. Liang, "JARVIS-1: Open-world multi-task agents with memory-augmented multimodal language models," 2023, *arXiv:2311.05997*.

[36] Y.-S. Chuang, A. Goyal, N. Harlalka, S. Suresh, R. Hawkins, S. Yang, D. Shah, J. Hu, and T. T. Rogers, "Simulating opinion dynamics with networks of LLM-based agents," 2023, *arXiv:2311.09618*.

[37] Y. Han and Z. Guo, "Regulator-manufacturer AI agents modeling: Mathematical feedback-driven multi-agent LLM framework," 2024, *arXiv:2411.15356*.

[38] M. Marchi, S. Soatto, P. Chaudhari, and P. Tabuada, "Heat death of generative models in closed-loop learning," 2024, *arXiv:2404.02325*.

[39] H. Duan, A. Dziedzic, N. Papernot, and F. Boenisch, "Flocks of stochastic parrots: Differentially private prompt learning for large language models," in *Proc. 37th Adv. Neural Inf. Process. Syst.*, May 2024, pp. 76852–76871.

[40] A. Mohtashami, F. Hartmann, S. Gooding, L. Zilkha, M. Sharifi, and B. A. Y. Arcas, "Social learning: Towards collaborative learning with large language models," 2023, *arXiv:2312.11441*.

[41] Y. Ge, W. Hua, J. Ji, J. Tan, S. Xu, and Y. Zhang, "OpenAGI: When LLM meets domain experts," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, May 2024, pp. 5539–5568.

[42] A. Chevalier, A. Wettig, A. Ajith, and D. Chen, "Adapting language models to compress contexts," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Oct. 2023, pp. 3829–3846. [Online]. Available: <https://openreview.net/forum?id=kp1U6wBPXq>

[43] J. Brinkmann, A. Sheshadri, V. Levoso, P. Swoboda, and C. Bartelt, "A mechanistic analysis of a transformer trained on a symbolic multi-step reasoning task," in *Proc. Findings Assoc. Comput. Linguistics (ACL)*, Bangkok, Thailand, Aug. 2024, pp. 4082–4102.

[44] N. Nanda, L. Chan, T. Lieberum, J. L. Smith, and J. Steinhardt, "Progress measures for grokking via mechanistic interpretability," in *Proc. 11th Int. Conf. Learn. Represent.*, Feb. 2023, pp. 1–13. [Online]. Available: <https://openreview.net/forum?id=9XFSbDPmdW>

[45] C. Singh, A. Askari, R. Caruana, and J. Gao, "Augmenting interpretable models with large language models during training," *Nature Commun.*, vol. 14, no. 1, p. 7913, Nov. 2023.

[46] K. Pattanayak, V. Krishnamurthy, and A. Jain, "Interpretable deep image classification using rationally inattentive utility maximization," *IEEE J. Sel. Topics Signal Process.*, vol. 18, no. 2, pp. 168–183, Mar. 2024.

[47] V. Krishnamurthy and K. Pattanayak, "Necessary and sufficient conditions for inverse reinforcement learning of Bayesian stopping time problems," *J. Mach. Learn. Res.*, vol. 24, pp. 1–64, Mar. 2024.

[48] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surveys*, vol. 54, no. 6, pp. 1–35, Jul. 2021.

[49] J. Anthis, K. Lum, M. Ekstrand, A. Feller, A. D'Amour, and C. Tan, "The impossibility of fair LLMs," 2024, *arXiv:2406.03198*.

[50] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.

[51] L. Zhu, X. Xu, Q. Lu, G. Governatori, and J. Whittle, "AI and ethics—Operationalizing responsible AI," in *Humanity Driven AI: Productivity, Wellbeing, Sustainability and Partnership*. Cham, Switzerland: Springer, Dec. 2021, pp. 15–33.

[52] S. Liu, R. Suresh, and A. Banitalebi-Dehkordi, "Beauty beyond words: Explainable beauty product recommendations using ingredient-based product attributes," 2024, *arXiv:2409.13628*.

[53] B. Liu, *Sentiment Analysis and Opinion Mining*. San Rafael, CA, USA: Morgan & Claypool, 2012.

[54] P. Sachdeva, R. Barreto, G. Bacon, A. Sahn, C. von Vacano, and C. Kennedy, "The measuring hate speech corpus: Leveraging Rasch measurement theory for data perspectivism," in *Proc. 1st Workshop Perspectivist Approaches*, Jun. 2022, pp. 83–94.

[55] Z. Zhao, W. Fan, J. Li, Y. Liu, X. Mei, Y. Wang, Z. Wen, F. Wang, X. Zhao, J. Tang, and Q. Li, "Recommender systems in the era of large language models (LLMs)," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 11, pp. 6889–6907, Nov. 2024.

[56] J. Ni, J. Li, and J. McAuley, "Justifying recommendations using distantly-labeled reviews and fine-grained aspects," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, Hong Kong, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds., Nov. 2019, pp. 188–197.

[57] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. Le Scao, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed, "Mistral 7B," 2023, *arXiv:2310.06825*.

[58] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, "LLaMA: Open and efficient foundation language models," 2023, *arXiv:2302.13971*.

[59] C. Singh, J. P. Inala, M. Galley, R. Caruana, and J. Gao, "Rethinking interpretability in the era of large language models," 2024, *arXiv:2402.01761*.

[60] Y. Zhang, S. Mao, T. Ge, X. Wang, Y. Xia, W. Wu, T. Song, M. Lan, and F. Wei, "LLM as a mastermind: A survey of strategic reasoning with large language models," in *Proc. 1st Conf. Lang. Model.*, Jul. 2024, pp. 1–14. [Online]. Available: <https://openreview.net/forum?id=iMqJsQ4evS>

[61] J. Lanchantin, S. Toshniwal, J. Weston, A. Szlam, and S. Sukhbaatar, "Learning to reason and memorize with self-notes," in *Proc. 37th Conf. Neural Inf. Process. Syst.*, Sep. 2023, pp. 11891–11911.

[62] 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," in *Proc. Trans. Mach. Learn. Res.*, Mar. 2023, pp. 1–12. [Online]. Available: <https://openreview.net/forum?id=ehfRiF0R3a>

[63] C. Li, J. Ma, X. Guo, and Q. Mei, "DeepCas: An end-to-end predictor of information cascades," in *Proc. 26th Int. Conf. World Wide Web*, Apr. 2017, pp. 577–586.

[64] C. Cui, Y. Ma, X. Cao, W. Ye, and Z. Wang, "Drive as you speak: Enabling human-like interaction with large language models in autonomous vehicles," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. Workshops (WACVW)*, Jan. 2024, pp. 902–909.

[65] C.-E. Sun, T. Oikarinen, and T.-W. Weng, "Crafting large language models for enhanced interpretability," in *Proc. ICML Workshop Mechanistic Interpretability*, Jul. 2024, pp. 1–15. [Online]. Available: <https://openreview.net/forum?id=pXx3nDi8h1>

[66] R. Baltaji, B. Hemmatian, and L. Varshney, "Conformity, confabulation, and impersonation: Persona inconstancy in multi-agent LLM collaboration," in *Proc. 2nd Workshop Cross-Cultural Considerations NLP*, 2024, pp. 17–31.

[67] A. Borah and R. Mihalcea, "Towards implicit bias detection and mitigation in multi-agent LLM interactions," 2024, *arXiv:2410.02584*.

[68] N. Gupta, H. Narasimhan, W. Jitkrittum, A. S. Rawat, A. K. Menon, and S. Kumar, "Language model cascades: Token-level uncertainty and beyond," in *Proc. 12th Int. Conf. Learn. Represent.*, Jan. 2024, pp. 1–12. [Online]. Available: <https://openreview.net/forum?id=KgaBScZ4VI>

[69] N. D. Palo, A. Byravan, L. Hasenclever, M. Wulfmeier, N. Heess, and M. Riedmiller, "Towards a unified agent with foundation models," in *Proc. Workshop Reincarnating Reinforcement Learn. (ICLR)*, Jan. 2023, pp. 1–12. [Online]. Available: https://openreview.net/forum?id=JK_B1tB6p-

[70] M. Huh, B. Cheung, T. Wang, and P. Isola, "Position: The platonic representation hypothesis," in *Proc. 41st Int. Conf. Mach. Learn.*, vol. 235, Jul. 2024, pp. 20617–20642.

[71] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, W. Ye, Y. Zhang, Y. Chang, P. S. Yu, Q. Yang, and X. Xie, "A survey on evaluation of large language models," *ACM Trans. Intell. Syst. Technol.*, vol. 15, Mar. 2024, Art. no. 39.

[72] J. W. Howard, "Free speech and hate speech," *Annu. Rev. Political Sci.*, vol. 22, no. 1, pp. 93–109, May 2019.

[73] B. Kennedy, X. Jin, A. M. Davani, M. Dehghani, and X. Ren, "Contextualizing hate speech classifiers with post-hoc explanation," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 5435–5442.

[74] A. Jain and V. Krishnamurthy, "Controlling federated learning for covertness," in *Proc. Trans. Mach. Learn. Res.*, Jan. 2024, pp. 1–15. [Online]. Available: <https://openreview.net/forum?id=g01OVahtN9>

[75] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, Mar. 2011.

[76] H. Zhang, M. Dang, N. Peng, and G. van den Broeck, "Tractable control for autoregressive language generation," in *Proc. 40th Int. Conf. Mach. Learn.*, Jan. 2023, pp. 40932–40945.

[77] F. Bordes et al., "An introduction to vision-language modeling," 2024, *arXiv:2405.17247*.

[78] Z. Dong, X. Fan, and Z. Peng, "FNSPID: A comprehensive financial news dataset in time series," in *Proc. 30th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, Aug. 2024, pp. 4918–4927.

[79] V. Krishnamurthy and H. V. Poor, "A tutorial on interactive sensing in social networks," *IEEE Trans. Computat. Social Syst.*, vol. 1, no. 1, pp. 3–21, Mar. 2014.

[80] A. Panda, C. A. Choquette-Choo, Z. Zhang, Y. Yang, and P. Mittal, "Teach LLMs to phish: Stealing private information from language models," in *Proc. 12th Int. Conf. Learn. Represent.*, Jan. 2024, pp. 1–10. [Online]. Available: <https://openreview.net/forum?id=qo21ZlfNu6>

[81] J. Kazdan, R. Schaeffer, A. Dey, M. Gerstgrasser, R. Rafailov, D. L. Donoho, and S. Koyejo, "Collapse or thrive? Perils and promises of synthetic data in a self-generating world," 2024, *arXiv:2410.16713*.

[82] P. Wang, D. Xu, Z. Fan, D. Wang, S. Mohan, F. Iandola, R. Ranjan, Y. Li, Q. Liu, Z. Wang, and V. Chandra, "Taming mode collapse in score distillation for text-to-3D generation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 9037–9047.

[83] V. Krishnamurthy and C. Rojas, "Slow convergence of interacting Kalman filters in word-of-mouth social learning," 2024, *arXiv:2410.08447*.

[84] A. Badanidiyuru, B. Mirzasoleiman, A. Karbasi, and A. Krause, "Streaming submodular maximization: Massive data summarization on the fly," in *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2014, pp. 671–680.

[85] N. R. Council, *Frontiers in Massive Data Analysis*. Washington, DC, USA: National Academy Press, 2013. [Online]. Available: <https://nap.nationalacademies.org/catalog/18374/frontiers-in-massive-data-analysis>

[86] V. Krishnamurthy, "Quickest change detection of time inconsistent anticipatory agents. Human-sensor and cyber-physical systems," *IEEE Trans. Signal Process.*, vol. 69, pp. 1054–1069, 2021.

[87] G. Wang. (Oct. 2024). *OpenAI Fast-Tracks AI Agents. How do we Balance Benefits With Risks?*. Forbes. [Online]. Available: <https://www.forbes.com/sites/geruiwang/2024/10/04/openai-fast-tracks-ai-agents-how-do-we-balance-benefits-with-risks/>

[88] H. Kushner and G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications* (Stochastic Modelling and Applied Probability). New York, NY, USA: Springer, 2003.

[89] A. Jain and V. Krishnamurthy, "Structured reinforcement learning for incentivized stochastic covert optimization," *IEEE Control Syst. Lett.*, vol. 8, pp. 1295–1300, 2024.

[90] B. Bahador. *Classifying and Identifying the Intensity of Hate Speech*. SSRC. Accessed: Jan. 15, 2025. [Online]. Available: <https://items.ssrc.org/disinformation-democracy-and-conflict-prevention/classifying-and-identifying-the-intensity-of-hate-speech/>

[91] J. A. Chevalier and D. Mayzlin, "The effect of word of mouth on sales: Online book reviews," *J. Marketing Res.*, vol. 43, no. 3, pp. 345–354, Aug. 2006.

[92] V. A. Zeithaml, "Consumer perceptions of price, quality, and value: A means-end model and synthesis of evidence," *J. Marketing*, vol. 52, no. 3, p. 2, Jul. 1988.

[93] J. Qiu, K. Lam, G. Li, A. Acharya, T. Y. Wong, A. Darzi, W. Yuan, and E. J. Topol, "LLM-based agentic systems in medicine and healthcare," *Nature Mach. Intell.*, vol. 6, no. 12, pp. 1418–1420, Dec. 2024.

[94] X. O. He, "Mixture of a million experts," 2024, *arXiv:2407.04153*.

[95] M. Sidji, W. Smith, and M. J. Rogerson, "Human-AI collaboration in cooperative games: A study of playing codenames with an LLM assistant," *Proc. ACM Hum.-Comput. Interact.*, vol. 8, pp. 1–25, Oct. 2024, doi: [10.1145/3677081](https://doi.org/10.1145/3677081).

[96] D. Blackwell, "Equivalent comparisons of experiments," *Ann. Math. Statist.*, vol. 24, no. 2, pp. 265–272, Jun. 1953.



ADIT JAIN (Student Member, IEEE) received the B.Tech. degree (Hons.) in electronics and communication engineering from Indian Institute of Technology Guwahati, in 2022. He is currently pursuing the degree with Cornell University.

He has internship experience at Adobe Research and Goldman Sachs. His research interests include structural results for episodic reinforcement learning, online learning, and high-dimensional linear bandits. His research is focused on applications in federated learning, large language model agents, and active learning. He was a recipient of the Data Science Fellowship graciously offered by the Cornell Center for Social Sciences.



VIKRAM KRISHNAMURTHY (Fellow, IEEE) received the Ph.D. degree from Australian National University, in 1992.

From 2002 to 2016, he was a Professor and the Canada Research Chair at The University of British Columbia, Canada. He is currently a Professor with the School of Electrical and Computer Engineering, Cornell University. He is the author of two books *Partially Observed Markov Decision Processes and Dynamics of Engineered Artificial Membranes and Biosensors*, published by Cambridge University Press, in 2016 and 2018, respectively. His research interests include statistical signal processing and stochastic control in social networks and adaptive sensing.

Dr. Krishnamurthy was awarded an Honorary Doctorate at the KTH Royal Institute of Technology, Sweden, in 2013. He served as a Distinguished Lecturer for the IEEE Signal Processing Society and the Editor-in-Chief of IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING.