

# Characterizing the Implementation of Censorship Policies in Chinese LLM Services

Anna Ablove\*

University of Michigan

aablove@umich.edu

Shreyas Chandrashekaran\*

University of Michigan

shreyasc@umich.edu

Xiao Qiang

University of California at Berkeley

xiao@berkeley.com

Roya Ensafi

University of Michigan

ensafi@umich.edu

**Abstract**—From network-level censorship by the Great Firewall to platform-specific mechanisms implemented by third-party services like TOM-Skype and WeChat, Internet censorship in China has continually evolved in response to new technologies. In the current era of AI, emerging tools like Large Language Models (LLMs) are no exception. Yet, ensuring compliance with China’s strict, legally mandated censorship standards presents a unique and complex challenge for service providers. While current research on content moderation in LLMs is primarily focused on alignment techniques, their lack of reliability prevents sufficient compliance with strictly enforced information controls.

In this work, we present the first study of overt blocking embedded in Chinese LLM services. We leverage information leaks in the communication between the server and client during active chat sessions and aim to extract where blocking decisions are embedded within the LLM services’ workflow. We observe a persistent reliance on traditional, dated blocking strategies in prominent services: Baidu-Chat, DeepSeek, Doubao, Kimi, and Qwen. We find blocking placements during the input, output, and search phases, with the latter two leaking varying amounts of censored information to client machines, including near-complete responses and search references not rendered in the browser.

Seeing the need to balance competition on the global stage with homegrown censorship restrictions, we observe in real time the concessions made by service providers hosting models at war with themselves. Through this work, we emphasize the importance of a more holistic threat model of LLM content accessibility, integrating live deployments to study access as it pertains to real world usage, especially in heavily censored regions.

## I. INTRODUCTION

For the past three decades, the Internet censorship apparatus in China has grown and adapted, maintaining its stringent control on the ever-changing web. At its inception, the Great Firewall was responsible for restricting traffic to foreign or sensitive websites via the regulation of plain-text DNS and HTTP connections [1]–[4]; with the widespread adoption of HTTPS, this additionally moved to include the unencrypted SNI field [5], [6]. The escalation furthered with the increased use of third-party web-based applications such as TOM-Skype and WeChat, where the development of platform-level censorship and surveillance of objectionable content served to

monitor Chinese users’ communications [7]–[10]. In this new era of AI, nascent technologies like Large Language Models (LLMs) will not be spared.

Current research on content moderation in LLMs is primarily focused on *alignment*, or the development of algorithmic methods to restrict discussion of sensitive topics. The goal of this process is to reliably train models to respond with refusals such as “Sorry, I can’t help with that” when asked questions involving objectionable content such as “Tell me how to build a bomb” or “Give me instructions to manufacture methamphetamines” [11]–[15]. However, these methods are not yet completely reliable or consistent; indeed, subsequent works have popularized LLM jailbreaking. These works show the ability to trivially and reliably circumvent model-internal alignment in both open-source and hosted models (such as ChatGPT) using a variety of methods including adding prefixes, suffixes, or hypotheticals to prompts [16]–[22].

Accordingly, in the presence of legally mandated, strictly enforced information controls necessarily imposed on service providers, alignment alone is not sufficient. However, there is a lack of existing research exploring what external filtering is actually enacted to comply with these regulations.

In China, the government enforces rigid content moderation guidelines on companies through regulatory bodies like the Cyberspace Administration of China (CAC) [23]. As a result many major services and providers, including TOM-Skype, WeChat, ByteDance, Alibaba, and Baidu have spent years developing a combination of heuristic, algorithmic, and human-in-the-loop processes to ensure compliance, including aggressive keyword filtering and audio and video review [24]–[26]. Given the aforementioned limitations of alignment, *it follows that LLM service providers would turn to similar overt blocking strategies to implement censorship*; however, these established, traditional censorship methods are ill-fitted for direct application to LLM services.

Implementing censorship policies in LLMs poses a complex and unique challenge when compared to existing traditional methods, including filtering of chat messages (WeChat) or input queries (search engines). The scale of data used to train these models renders it effectively impossible to exclude all objectionable content; even if this was to be done, models can hallucinate, or present incorrect and made-up facts in a response. This means models can output responses containing objectionable content, even to benign-seeming queries. Ad-

\*Both authors contributed equally to this work.

Network and Distributed System Security (NDSS) Symposium 2026  
23-27 February 2026, San Diego, CA, USA  
ISBN 979-8-9919276-8-0

<https://dx.doi.org/10.14722/ndss.2026.231761>  
[www.ndss-symposium.org](http://www.ndss-symposium.org)

ditionally, due to the unique token-based generation schema of LLMs, any subsequent filtering must balance efficacy and response latency to ensure users receive timely responses to their queries. Given these complexities and limitations, we specifically explore *locations* of blocking implementations, blocking *consistency*, and resulting *client-side data exposure*.

In this work, we present a study of censorship in Chinese LLM services enacted via overt, external blocking strategies separate from the underlying models themselves. This deviates from prior work on algorithmic LLM jailbreaking whose focus is model-internal alignment induced by post-training. We build a novel framework for conducting in-depth investigations of overt blocking embedded in popular Chinese LLM services. We leverage information leaks in the communication between the server and client during active chat sessions and reverse engineer where blocking decisions are embedded within the LLM services’ workflow. This requires that we man-in-the-middle and unencrypt the traffic to maximize visibility, as well as account for a variety of server-side implementations. In conjunction with three experts on information controls in China, we develop a list of 80 sensitive queries to trigger censorship behaviors and select a diverse set of LLM services including both traditional industry titans and popular upstarts.

Through our extensive analysis, we are the first to present a composite model of blocking placements in Chinese LLM services, observing blocking enacted during the input, output, and search phases. We find blocking enacted during the input phase is extremely consistent; with DeepSeek, Qwen, Kimi, and Doubao input blocking queries persistently across samples. In contrast, output blocking is significantly less consistent. Regarding search behaviors, while all services are capable of choosing to skip the search stage, with the exception of DeepSeek, they are not likely to do so. We also find significant overlap in telemetry and analytics infrastructure between the services, with DeepSeek, Kimi, and Doubao sending logs to the same Autonomous System.

At a high level, our findings are not only focused on the characterization of overt blocking filters, but rather the deployment of censorship in each LLM service as a whole; this includes quantifying information leakages to the client machine, comparing and contrasting telemetry infrastructures, and integrating browser display information. By studying Chinese LLM services in this “censorship as deployed” context, we go beyond prior work studying global LLMs and focus on the services themselves and the conglomeration of behaviors associated with censorship-driven overt blocking. We hope this work inspires more technical research into the day-to-day realities of users interacting with heavily censored generative AI systems.

## II. BACKGROUND

We situate our study of information controls in Chinese LLM systems within the context of China’s censorship apparatus. To this end, we divide this section into three parts: an overview of Chinese censorship, a review of core LLM

principles and related work, and a breakdown of the Chinese LLM ecosystem.

### A. Censorship in China

The Great Firewall of China (GFW), active since the late 1990s, is a large-scale and sophisticated system that monitors and regulates Chinese Internet traffic, blocking access to numerous foreign websites and suppressing free speech [27]. Early research by Clayton et al. in 2006 [1] and Crandall et al. in 2007 [28] examined its keyword filtering mechanisms, specifically the injection of TCP reset (RST) packets triggered by the detection of sensitive keywords in HTTP traffic. Xu et al. in 2011 [2] further investigated where in the network path this filtering occurs, finding that it predominantly takes place at the national border.

Other studies focused on the GFW’s DNS-based censorship. Lowe et al. in 2007 [3] first uncovered its DNS poisoning actions, with subsequent works in 2012 [4] and 2014 [29]. The latter emphasizes the collateral damage associated with this poisoning, including the disruption of over-Internet services relying on DNS infrastructure. Later work by Hoang et al. [30] introduced GFWatch, a measurement platform continuously monitoring the GFW’s DNS censorship. Further, with the widespread adoption of HTTPS, studies like Rambert et al. [5] and Bock et al. [6] explored filtering on the unencrypted Server Name Indication (SNI) field.

As increased efforts and sophistication were put into circumventing the GFW, it in turn grew to target circumvention services. In 2012, Winter et al. [31] explored the detection and blocking of Tor bridges via traffic fingerprinting, with follow up works by Ensaifi et al. in 2015 [32] and Dunna et al. in 2018 [33]. Additional work by Alice et al. in 2020 [34] focused on the blocking of circumvention protocol Shadowsocks, while Wu et al. in 2023 [35] studied the deployment of a new GFW system blocking fully-encrypted traffic.

Beyond the GFW, censorship is also prevalent in third-party services, which has led to increased exploration of platform censorship behaviors and implementations. In 2011, Knockel et al. [7] analyzed keyword filtering and surveillance behaviors in TOM-Skype, the primary version of Skype served in China, finding significant oversight into users’ conversations and meetings; additionally, they observed filtering was predominantly implemented on the client side. Censorship on WeChat in particular is a point of interest; in 2016, Ruan et al. [8] demonstrated that not only single keywords but also combinations of keywords could trigger censorship. Further work by Knockel et al. in 2018 [9] explored image filtering on WeChat, finding more complex server-side implementations of censorship than previously discovered. In 2021, Ruan et al. [10] engaged in a longitudinal study of keyword censorship as impacted by the 19th National Communist Party Congress. Other studies in this area focused on keyword filtering in mobile games [36] and QQMail’s email service [37], as well as uncovering the complex nature of server-side search-engine censorship rules [26].




Fig. 1: Timeline of Chinese LLM Ecosystem.

### B. Large Language Models

Large language models trained at scale have emerged as a nascent technology in the last few years, primarily fueled by the rapid success of academic and industry efforts in the U.S. and China. As of 2024, over 65% of U.S.-based users reported having multi-turn conversations with LLMs regularly and Mainland China reported widespread adoption as well, with more than 600 million individual LLM users across 188 government-approved LLM providers [38], [39]. This number is only projected to grow with the LLM ecosystem and as use cases continue to increase.

These models are first pre-trained at a large scale to specifically predict the next token in a given data stream in an unsupervised manner [40]. Next, they are instruction-tuned, where they are explicitly trained to follow instructions via supervised learning [41]. Finally, reinforcement-learning based training *aligns* these models to the security and content preferences envisioned by their creators, ensuring mechanisms to refuse harm-inducing queries [42].

United States export restrictions since 2022 have limited the sale of high-end GPUs to China [43], robbing them of an extremely valuable resource by which to significantly speed up the tensor computations required in LLM training. Despite this, China has become a superpower in the space, with their open- and closed-source offerings rivaling popular U.S. models in benchmark rankings, as well as providing their enterprise and individual services for much cheaper prices than U.S. competitors; 1 million output tokens via the DeepSeek-R1 API costs \$2.19, whereas the same costs \$60 via the OpenAI-o1 API [44], [45]. As such, Chinese LLMs are poised to capture a large portion of business market share due to their cost-effectiveness and performance.

As previously mentioned, LLMs are fine-tuned through a variety of methods to prioritize safety and content alignment [46]. Typically, this manifests as the LLM refusing to answer variants of questions deemed dangerous, sending instead a denial response such as “Sorry, I can’t help you”, or “As an AI model, I cannot assist with your request”. Beyond merely eliciting refusals, alignment seeks to induce models to provide deterministic responses to specific queries. However, these methods are prone to multiple failure modes; LLMs are stochastic by nature, and many works have identified effective ways to reliably circumvent safety and content alignment [18],

[47]–[49], creating a challenge for those seeking to subject LLMs to strict information controls.

### C. Chinese LLM Service Ecosystem

Recent open- and closed-source research efforts have resulted in the release of competitive LLMs in China. However, given the regulatory frameworks governing acceptable content, this form of content generation is fundamentally risky for service providers. Failure to properly operate within legislative bounds has already been met with swift action; sites like “Lingxiang Zhiwen AI” and “Chongqing Sentinel Expansion Fans” were reportedly ordered to cease operations for providing generative AI services without meeting legal requirements in July 2023 [50].

In order to better monitor and facilitate the growth of companies and products in this space, the Chinese government has developed targeted pieces of legislation. In addition to previous content laws restricting online content, these laws specifically focus on AI. The “Provisions on the Administration of Deep Synthesis Internet Information Services” established critical baselines for generative AI services, including the labeling of AI generated content and detailing requirements regarding the filing of algorithms with the Cyber Administration of China (CAC) [51]; notably, companies submitting algorithms for review must include information on their security monitoring policies and emergency response plans. This paved the way for later regulations, especially the “Interim Measures for the Management of Generative Artificial Intelligence Services”, which includes in its provisions that services must “uphold the core socialist values and not generate content that may harm national security or that is forbidden by laws or regulations.” [52]. The “Trial Measures for Scientific and Technological Ethics Review” further detailed review processes for developments in science and technology [53], [54].

Major industry players, including Baidu, ByteDance, and Alibaba, received government approval and released their first publicly accessible products in August 2023 [55], [56]. Newer companies also emerged with notable products. Moonshot AI, a startup founded in 2023, released its first chatbot, Kimi, in October of the same year [57]. The following year, DeepSeek released its V3 model in December 2024, followed by its R1 model in January 2025 [58]; the latter received both domestic and international acclaim, topping the U.S. Play Store charts shortly after its release [59]. With the far-reaching

Service	Model	Domain	Company
Baidu-Chat	Baidu Chat Lightning Mode	chat.baidu.com	Baidu
DeepSeek	DeepSeek-V3	chat.deepseek.com	DeepSeek
Doubao	Doubao	doubao.com/chat	ByteDance
Kimi	Kimi-1.0	kimi.moonshot.cn	Moonshot AI
Qwen	Qwen-2.5 Max	chat.qwen.ai	Alibaba

TABLE I: Models, Domains, and Companies for Selected Services.

appeal of LLM services, the aforementioned legally mandated information controls have garnered significant attention by Western media [60]–[62].

**Our Work in Perspective.** In this work, we explore overt blocking embedded in Chinese LLM services. This deviates from the bulk of prior work on global LLMs originating in Western countries, which focus largely on safety alignment techniques. As discussed in Sec. II-B, this alignment, which is intrinsic to models and a byproduct of their post-training, is not sufficient to meet the requirements of stringent, legally-mandated information controls imposed on Chinese providers. This requires the integration of external filters, similarly observed in other Chinese services, as noted in Sec. II-A. Furthermore, we study censorship as deployed in each service; this goes beyond existing research on global LLMs by focusing on the services themselves, including analyzing client-side information leakage, telemetry infrastructures, and real-time UI feedback from the browser, to better understand the conglomeration of behaviors associated with overt blocking in these services.

### III. DATA COLLECTION

In order to understand blocking placements and implementations of censorship policies in Chinese LLM services, we design our data collection to account for several challenges. With regards to scope, we need to identify LLM services to target in this study (III-A) and curate queries that trigger censorship (III-B). In our measurements, we want to closely mimic the actions of a human user and gain insight into their perspective, while prioritizing query isolation and ensuring persistent access by accounting for rate limits and free query limits (III-C). Finally, we also want to find ways to capture raw traffic between the client and LLM server to maximize observable information (III-C).

#### A. Service Selection.

Our goal is to study a reasonably diverse set of widely used LLMs, capturing both traditional and newer players in the Chinese LLM space. However, barriers to access present a key challenge. It is common for Chinese LLMs to require SMS code verification for account creation and usage, typically restricting accepted phone numbers to narrow subsets of global country codes, such as solely mainland China (+86) numbers, which are tied to identity cards for Chinese citizens and residents. As a result, our selection process must account for these access restrictions.

We select five LLM services built by tech companies in China. We specifically focus on including both traditional industry titans (Alibaba, Baidu, ByteDance) and popular upstarts (Moonshot, DeepSeek). We test the following models: Baidu Chat Lightning Mode, DeepSeek-V3, Doubao, Kimi-1.0, and Qwen-2.5-Max, as shown in Table I. We study each service with search enabled, which allows the integration of Internet search content into the response generation process. Throughout this work, we refer to these models by service abbreviations: Baidu-Chat, DeepSeek, Doubao, Kimi, and Qwen.

#### B. Query Selection.

With regards to query selection, our primary objective is to trigger the implemented information controls in each LLM. This necessitates a certain level of familiarity with both the content and context of topics designated as sensitive by the Chinese government. Additionally, as Chinese LLMs are trained in simplified and traditional Chinese text as well as English, we want to capture any potential linguistic distinctions, which requires a level of fluency in all three languages. Thus, in conjunction with three experts on information controls in China, we compile a list of 80 potentially sensitive queries of interest, which we include in English, Traditional Chinese, and Simplified Chinese as our query set, with our goal being to observe instances of both censored and permitted queries in each of our selected LLM services. This gives us 80 equivalent queries in each language, for a total of 240 queries, as our query set.

#### C. Measurements

We conducted all measurements of our five LLM services over the period of March 9–16, 2025; for each query and service, we complete five separate tests.

**User Interface Scraping.** In order to simulate user behavior at scale, we build a Selenium-based scraper for each service that, in sequence, carries out the steps shown in Alg. 1. For each query, the scraper simulates user input by clicking and typing the query into the input box and sending it, subsequently capturing UI-rendered responses from the LLM service. Our scrapers also account for potential rate limiting by queuing a retest after certain failure conditions are met.

The goal of our scraping step is to capture the typical user experience while using these served LLMs. As such, we take meticulous care to accurately carry out actions users would take when using these products, including adding adequate realistic delays between key entries. The purpose of instantiating a new chat between queries is to test each query while maximizing isolation, as we want to characterize the served LLMs response to *only the provided query* as opposed to the compounded effect inherent to asking a query in an existing chat session containing previous queries and responses.

For the majority of services, we keep the browser session persistent across a given test run. However, Kimi institutes strict free query limits, so when testing this service we necessarily instantiate new browser sessions for each query.




Fig. 2: Framework for Investigations into Blocking Behaviors of Chinese LLM Services.

#### Algorithm 1 User Interface Scraping of Query List

```

1: Let activeQueries = queries
2: Let runs = 0
3: while runs < maxRuns and ||activeQueries|| > 0 do
4:   ratelimited = []
5:   for q in range(activeQueries) do
6:     start_new_chat_session()
7:     click(inputTextBox)
8:     for each character c in q do
9:       type(c) and wait short delay
10:      send(q) with sendButton or key bindings
11:      Wait for response r
12:      if r indicates rate limiting then
13:        ratelimited += q
14:      else
15:        save(q,r)
16:      if ||ratelimited|| > 0 then
17:        activeQueries = ratelimited
18:      runs += 1

```

**Man-in-the-Middle.** UI-visible information alone does not provide sufficient details about the underlying communication. For instance, in some services we observed error messages and toast notifications, or pop-ups, that varied in clarity. We also witnessed cases where a response was partially or fully generated before being replaced with an error message, as shown in Fig. 6, where the intermediary response would fail to be captured by the scrapers. Developing a better understanding of these idiosyncrasies is instrumental to fully understand the blocking policies of each service.

To this end, during all measurements, we situate ourselves on-path between the client and server and man-in-the-middle their communication. Specifically, we leverage our access to the client browser and use libraries to decrypt raw network traffic and take a packet capture. This gives us an opportunity to analyze the flow of information in these LLM services.

#### IV. TRAFFIC ANALYSIS

In order to investigate locations and details of blocking in all of the LLM services as observable from the network, we

need to analyze our collected traffic data. Due to the variety of implementations, extracting the responses and blocking indicators requires separate efforts for each service. In this section, we detail the high-level process for analyzing network traffic which is undertaken independently for each of the five services. Note that throughout this section we use “flow” or “HTTP flow” to refer to the joint unit of both the HTTP request and corresponding response.

**Flow Extraction.** We first search all HTTP flows for the presence of raw query text. This facilitates the identification of HTTP requests and responses that contain the query. Example (1) shows a snippet from the body of an HTTP POST request from the client machine to Qwen’s endpoint /api/v1/chats/new after “What caused the Opium Wars?” is queried in the chat.

```

title: New Chat,
models: qwen-max-latest,
messages: [{"timestamp: 1742486022,
content: What caused the Opium Wars?}],
chat_type: search

```

Given this set of query-relevant flows, we seek to identify the key URL-request patterns and begin to formulate the pattern of the chat process. For instance, as example (1) indicates the initiation of the chat, we find related proximal flows, such as those shown in example (2), which correspond to other parts of the workflow. In this example, the request to /api/v1/chat/completions prompts the sending of response data, while the request /api/v1/chat/completed indicates the response is fully generated.

```

POST chat.qwen.ai /api/v1/chats/new
<- 200 application/json
POST chat.qwen.ai /api/v1/chat/completions
<- 200 text/event-stream
POST chat.qwen.ai /api/v1/chat/completed
<- 200 application/json

```

```

2025-03-12 18:05:01 POST https://chat.qwen.ai/api/chat/completions HTTP/2.0
<-- 400 application/json 166b 2.84s
Request                                         Response
date:                               Wed, 12 Mar 2025, 22:05:04 GMT
content-type:                      application/json
content-length:                   166
access-control-allow-credentials: true
access-control-allow-origin:      https://chat.qwen.ai
vary:                                Origin
JSON
{
  "modality": [
    "text"
  ],
  "code": "data_inspection_failed",
  "stage": "input",
  "detail": "Content security warning: input text data may contain
inappropriate content!"
}

```

Fig. 3: Qwen Input Block HTTP Flow.

**Response Parsing.** To understand the scope of the data being sent from the server, we prioritize the identification of the response generation field. In all services, we find that the client establishes an HTTP connection with a `text/event-stream` content type, also known as Server-Sent Events (SSE), to facilitate this exchange. As a result, the response data is sent through SSE messages in chunks; the relevant event syntax for Qwen (3) and Kimi (4) is shown below. Notably, the information regarding whether or not the system engaged in a search, as well as associate data, is also exposed here. This is discussed further in Sec. V.

```

{ choices:[{ delta: {..., content: "The Opium"} }] }           (3)
{ choices:[{ delta: {..., content: " Wars were "}] } }

{ event: cmpl, ..., text: "The", view:cmpl }
{ event: cmpl, ..., text: " Op", view:cmpl }                  (4)
{ event: cmpl, ..., text: "ium", view:cmpl }

```

**UI Comparisons.** We extract responses from the flows for each query and compare with the web scraped responses, which are representative of what is visible to the user via the user interface (UI). Table II shows a case where the response is fully generated and all data is sent to the client, but ultimately only an error message is presented to the user in the UI. We leverage the differences in the responses parsed from the network traffic and the UI visible responses to identify specific indicators of blocking.

However, in certain cases, we observe clear indicators of blocking in the UI and subsequently isolate the relevant trigger in the network traffic. For instance, in Doubao, toast notifications suggest a subject change in the case of sensitive queries or responses; these are pushed by a Web Socket connection that persists throughout the session and are tagged by a uniquely generated conversation identifier as opposed to by query. Thus, we store this mapping between conversation ID and query, which is exposed in other requests, and use this to match toast notifications to queries.

## V. CENSORSHIP INDICATORS

Our goal is to understand what the blocking signals are on the wire in our selected LLM services. Specifically, we are interested in determining where and when blocking is enacted, what data is exposed, and how to differentiate between various




Fig. 4: DeepSeek UI-Visible Input Block. Note the absence of any warnings or toast notifications, requiring traffic analysis to disambiguate from output blocks or aligned refusals.

Query	"Who is the leader of China?"
Raw Traffic	"As of the latest information available, Xi Jinping is the leader of China. He is the president of China and holds significant influence due to his position within the Communist Party of China (CPC)."
UI Visible	"Uh-oh! There was an issue connecting to Qwen2.5-Max. Content security warning: output text data may contain inappropriate content!"

TABLE II: Raw Traffic vs. UI-Visible Responses.

types of blocking. To this end, we step through the stages of the query-response process and highlight the locations where we see blocking occur; we show an overview of this in Fig. 5. At a high level, in a successful chat, the query is sent to the server followed by a search retrieval process which returns search results; then, response generation begins, and the response is sent to the client in chunks.

### A. Input

This is the introductory phase of conversation. Blocking at this stage occurs before any search steps or response generation.

**Explicit Input Blocking.** *DeepSeek, Qwen, and Kimi send explicit input block error messages.* In these services, input blocking manifests as a declarative error message in place of either the response generation stream or the beginning of the searching process. In DeepSeek and Kimi, the blocking manifests as an immediate halting, with each service sending SSE error messages including `finish-reason:content-filter` and `kimi.completion.content-filter`, respectively, in lieu of response generation events. As such, the data is returned with a 200 HTTP code, a requirement for establishing a SSE stream. In contrast, Qwen rejects the request containing the query prior to setting up the HTTP connection, returning an error including the fields `code:data-inspection-failed` and `stage:input` with a 400 HTTP status code.

**Implicit Input Blocking.** *Determining input blocks in Doubao and Baidu-Chat requires further examination.* In the absence of explicit error messages, we look for other indicators of blocking. In Baidu-Chat, we observe a toast notification suggesting a subject change, which corresponds to the field value `showType:toast` or `kunlun_popup` within the SSE messages. We classify cases where this toast was sent and




Fig. 5: Observed Blocking Placements.

neither the search stage nor output stage was initiated as input blocking. We confirm this by observing a DOM element recommending the user change the subject present in all web scraped answers associated with input blocks. This element, ‘让我们换个话题聊聊吧。’ or ‘Let’s talk about a different topic.’, is triggered by JavaScript on the client side.

For Doubao, in cases where there was no response generation and the server sent toast notifications suggesting a subject change, we identify a particular error string, ‘抱歉，我无法回答你的问题’ or ‘Sorry, I can’t answer your question’, which is sent together in one event. In some cases, this is preceded by ‘已完成思考’ or ‘Completed thinking’; we leverage an additional flow, a POST request to `*/message/list?*`, which contains query and truncated response for the given chat in the `brief` field, and confirm in all cases only the error string is sent. We thus classify input blocks in Doubao as cases where the message list brief matches the error string and an error toast was sent.

### B. Search

The search phase occurs after the query is sent, but before any chunks of the response are received. Critically, at this stage, the LLM service decides to either engage in or forgo the searching step. Blocking in this stage can occur during or after search is finished, but before any of the response is sent to the client.

**Search Decision.** All services can choose whether or not to carry out the search retrieval process. Within the SSE messages, we identify the searching step via fields such

as `event:search_plus` in Kimi and `type:search-index` in DeepSeek. Accordingly, we can also locate the returned search results. With the exception of Qwen, all of the services explicitly announce search targets, a set of search-specific queries, as part of the searching step. These range from copies of the original query (Baidu-Chat) to a set of slight variations geared towards search (DeepSeek, Kimi, Doubao). Notably, DeepSeek and Baidu-Chat can choose to opt out of search retrieval after the search targets are sent.

In most services, the absence of these events and associated data is indicative of the decision to not search. However, Doubao explicitly declares intent to search or skip via a `seed-intention` field which specifies whether the response generation will utilize web search (browsing) or not (`seed-main`); we confirm that these field values correspond to whether or not search data was sent.

**Pre-Search Blocking.** *Baidu-Chat* can block during the search process before sending results to the client. In Baidu-Chat, we observe blocking situated at the opt-out decision point once search targets have been declared. For successful searches, query processing events, denoted by the field value `step_name:queryAnalysis`, are followed by search events, marked by `step_name:searchingOnline`. However, pre-search blocking interrupts the latter step, and the SSE stream exits without sending any search results.

**Post-Search Blocking.** *DeepSeek*, *Kimi*, and *Baidu-Chat* all engage in post-search blocking. Post-search blocking is blocking situated after search results are sent to the client, but before any response generation. In DeepSeek, messages containing search results are immediately followed by a content filter error. Similarly, in Kimi, after the search phase, there is an event:remove of type `type:search_plus`, followed by a completion filter error message. In some cases, we observe the sending of buffer completion events, which contain empty strings instead of response tokens and are typically observed at the start and end of sequences of successful response generation messages. Notably, we do not consider these buffer completions part of the response generation process as they do not contribute any text.

However, in Baidu-Chat, there are no explicit errors after the searching stage completes; instead of sending the response in chunks to the client, the server sends a refusal string all at once. We characterize this as post-search blocking since no pieces of a legitimate response are sent over.

### C. Output Generation

The output phase follows either directly after the decision to forgo searching or the retrieval of search results and consists of the response being sent over to the client. Blocking at this phase is characterized by some or all of a response being sent over in chunks before the output is classified as sensitive and replaced in the UI with an error message.

**Explicit Output Blocking.** *DeepSeek*, *Qwen*, and *Kimi* send explicit indicators of output blocking. In the case of sensitive output, the response generation process for these services is interrupted by explicit error messages; for Qwen,




Fig. 6: User-Visible Output Blocking by Qwen.



Fig. 7: Output Blocking by Doubao. Note that the refusal text presents as a normal response, despite the presence of a toast notification.

these contain the fields `code:data-inspection-failed` and `stage:output`; while DeepSeek, which usually finishes responses with `finish_reason:stop`, halts with `finish_reason:content_filter`. Kimi similarly stops with `kimi.completion.content-filter`. For all of these services the response is sometimes visible momentarily in the UI, but is quickly removed and replaced with an error message indicating that the topic is out of scope; Fig. 6 shows the user-facing appearance of output blocking as well as the corresponding momentarily visible, in-progress response.

**Implicit Output Blocking.** *Doubao and Baidu-Chat engage in output blocking without explicit error messages.* For Doubao, we again need to consult multiple factors. Doubao presents no explicit error in the network traffic, instead generating most of the response as normal and abruptly stopping generation with its aforementioned error string; similar to the input blocking, this correlates with toast notifications suggesting a topic change, as shown in Fig. 7. In Baidu-Chat, the output blocking is even less explicit. The response generation will complete as normal, but the response is only partially rendered in the UI, with the majority of it cut off in the browser view; in some cases, only a single word or two is observable on the page. We verify this is not a mere function of browser lag as we observe consistent output blocking by particular queries across runs, which is discussed in Sec. VI.

**Output Warnings.** *We find Baidu-Chat and Qwen additionally send warnings for potentially sensitive response content.* Distinct from the output blocking described above, after some sensitive chats, Baidu-Chat generates a full re-

sponse but segments the UI, displaying the aforementioned DOM error string while leaving the response visible on the screen. Similarly, in certain cases, Qwen flashes a toast notification suggesting sensitive input or output data, while leaving the generated response unchanged and visible to the user and not requiring the creation of a new chat to continue the conversation. Interestingly, the traffic indicator associated with this is an error returned in the POST request to the `api/task/suggestions/completions` endpoint, which normally contains follow-up questions to suggest to the user.

#### D. Metadata

Beyond the aforementioned blocking indicators we find pieces of metadata in the services that further illuminate some aspects of their processing. In addition to the data discussed below, we also identify endpoints associated with telemetry channels for usage monitoring and analytics in all services, which we further discuss in Sec. VI-D.

**Warning Flags.** *We identify warning flags that are related to blocking behaviors in Doubao and Baidu-Chat.* In the SSE messages for Baidu-Chat, we find the fields `needClearHistory`, `antiFlag`, and `isSafe`, which are set for all instances of pre-search blocks. Furthermore, though `isSafe` typically takes values of either 0 or 1, cases where it is set to 2 correspond with the toast notification suggesting topic change. Additionally, within the persistent WebSocket connection in Doubao, we locate a field titled `inner_risk_process_type`, which maps to a group of numbers 0-13. We find all instances of blocking are marked with tags  $\{1, 5, 6, 13\}$ .

**Content Tagging.** *We find Qwen classifies non-blocked inputs with relevant content tags.* In a follow-up after the response is generated, the client sends a POST request to the `api/task/tags/completions` endpoint and returns a list of content tags. For our query set, the most common tags were *Politics*, *Censorship*, and *Technology*. Interestingly, we also see tags such as *Freedom of Information* and *Core Socialist Values*.

## VI. RESULTS

In this section, we detail our findings on the censorship deployed in each service as found in the measurements. We examine the aggregate blocking enacted by each service




Fig. 8: **Blocking Instances Across All Samples.** Baidu-Chat has the highest instances of blocking across all runs (504), followed by Kimi (384) and DeepSeek (230).

(VI-A), the consistency of the input and output filters (VI-B), information observability for different blocking types (VI-C), and analytics channels (VI-D).

#### A. Aggregate Censorship by Service

	English	Chinese-SI	Chinese-TW	Total
Baidu-Chat	<b>59 (74%)</b>	<b>61 (76%)</b>	<b>66 (83%)</b>	<b>186</b>
Kimi	37 (46%)	39 (49%)	41 (51%)	117
DeepSeek	28 (35%)	30 (38%)	30 (38%)	88
Qwen	16 (20%)	24 (30%)	22 (28%)	62
Doubao	3 (4%)	5 (6%)	6 (8%)	14

TABLE III: **Queries Blocked By Language.** We highlight the queries blocked in at least one sample for English, and Simplified and Traditional Chinese. Out of the total 240 queries, Baidu-Chat blocks the most, followed by Kimi, DeepSeek, Qwen, and Doubao.

**Blocking Types.** All services block at the input and output stages; Baidu-Chat, DeepSeek, and Kimi also block after search retrieval. In terms of blocking instances, Qwen and DeepSeek perform the most input blocking, with 165 times for 33 queries and 143 times for 70 queries, respectively; while for output blocking, Baidu-Chat (396 times for 82 queries) and Kimi (203 times for 82 queries) perform the most. Proportionally by service, Doubao (88.7%) and Qwen (74%) engage in the most input blocking, while Baidu-Chat (78.6%) and DeepSeek (62.2%) prefer output blocking. Kimi also most frequently enacts post-search blocks (31.5%), or 121 times for 30 queries, after which is Baidu-Chat (13.7%), 69 times for 28 queries. Baidu-Chat additionally implements pre-search blocks after search targets are declared (7.5%), or 38 times for 10 queries.

**Search Decisions.** Except for DeepSeek, all services are more likely than not to complete the search stage. Excluding instances of input blocking, we highlight the search decision proportions and associated outcomes per service in Table IV. Qwen, Baidu-Chat, and Kimi complete searching at high rates:

Search Bin	Service	Output	Post-search*	No Block	Ratio
Complete	Baidu-Chat	28.5%	6.5% (0%)	65%	0.89
	DeepSeek	73.5%	4.1%	22.4%	0.15
	Doubao	0.4%	–	99.6%	0.60
	Kimi	19.9%	13.1%	67%	0.81
	Qwen	5.3%	–	94.7%	<b>0.98</b>
Skip	DeepSeek	11.7%	0%	88.3%	<b>0.84</b>
	Doubao	0.9%	–	99.1%	0.40
	Kimi	9.1%	–	90.9%	0.19
	Qwen	20%	–	80%	0.02
Partial	Baidu-Chat	69.2%	– (28.6%)	2.2%	<b>0.11</b>
	DeepSeek	0%	–	100%	0.01

TABLE IV: **Searching Decisions and Blocking Outcomes by Service.** We report the percentage of blocking outcomes per search bin for each service, with the “Ratio” column indicating the overall proportion of outcomes in a given bin per service. Note, for Baidu-Chat, we include the pre-search blocking in parentheses.

0.98, 0.89, and 0.81 respectively, with Doubao following at 0.6. In contrast, DeepSeek only searches at a rate of 0.15. Regarding blocking outcomes, some services responded more after finishing the search stage; notably, Baidu-Chat successfully answered 65% of the time after completing search, as opposed to only 2.2% after opting-out of search after generating targets; Qwen saw a more modest decrease in response success rates, from 94.7% to 80% when skipping search. In contrast, other services blocked more after search completion; DeepSeek only responded to 22.4% of search-engaged querying, compared to 88.3% when declining to search; similarly, Kimi saw an increase from 67% to 90.9%. Doubao generally exhibits low levels of blocking in either case.

#### B. Characterization of Input and Output Blocking

**Input Blocking.** *Blocking at the input stage is overwhelmingly consistent, with DeepSeek, Qwen, Kimi, and Doubao input blocking persistently across all 5 samples.* Fig. 9 shows the cross-sample consistency of all input-blocked queries, in addition to Baidu-Chat’s pre-search blocks. We find that Qwen, DeepSeek, and Doubao input block the exact same set of queries in Simplified and Traditional Chinese but a largely different set in English; and of its 5 input blocked queries in Simplified Chinese Kimi only fails to block one, “Who is the leader of China?”, in Traditional Chinese. Notably, two queries blocked by Qwen, DeepSeek, Doubao, and Kimi in both languages are “What happened in Tiananmen Square in 1989?” and “What does ‘May 35th’ mean?”. Qwen overwhelmingly input blocks the most in English, including two of the three questions in our query set about disputed sovereignty issues. Interestingly, both of these questions produce full responses more often than not when asked in either Chinese variant.

Baidu-Chat represents an interesting case as we only observe a single instance of input blocking, with the service choosing to begin search steps for all other queries. Recall from Sec. V-B that Baidu-Chat reuses the query itself for search targets instead of extracting phrases or topics from it.




Fig. 9: **Input Blocking Across Service and Language.** Note, we show pre-search blocking for Baidu-Chat. Queries marked with “\*” were shortened for visibility purposes.

We observe heavy blocking after this phase, prior to sending search results to the client or beginning to generate a response, as discussed in Sec. V-B. Thus, we examine this type of blocking in conjunction with input blocking due to their functional similarities.

**Output Blocking.** *Output blocking is fairly inconsistent, with only 29 out of 349 output-blocked queries subject to output blocking across all 5 samples.* Fig. 10 shows the cross-sample consistency for a selection of the top 30 output-blocked queries in any language. Agnostic of language, we observe 6 of these prevalently blocked queries are additionally input-blocked by other services. Of the 29 consistently blocked queries, Baidu-Chat accounts for the majority (13), followed by DeepSeek (9), Kimi (8), and Qwen (2). As output blocking results from objectionable content in the output of the underlying model, it makes sense that queries that reliably elicit this type of information are blocked more consistently, such as “Who is the leader of China?” (DeepSeek) which prompts the inclusion of “Xi Jinping” in the response, and “Sources of tension between the U.S. and China?” (Baidu-Chat), which prompts discussion of territorial disputes in the South China Sea. Indeed, we find that these two queries are output blocked in at least 4 samples for each language.

### C. Information Observability

**Output Block Truncation.** *All services send partial or full responses to the client in cases where output blocking occurs.* Fig. 11 compares tokenized response lengths of truncated responses in cases of output blocking and successful responses, as received by the client; note, as discussed in Sec. V-C, in Baidu-Chat full responses are sent to the client, *even in cases of output blocking*, but are only partially rendered in the UI. For DeepSeek and Qwen, we see that the truncated responses on average are relatively close in length to the full responses,




Fig. 10: **Output Blocking Across Service and Language.** We show the top 30 output blocked queries. Queries marked with “\*” were shortened for visibility purposes.




Fig. 11: **Regular vs. Output-Blocked Generated Response Lengths (Tokens) by Service and Language.** We tokenize these responses using bert-base-multilingual-cased to account for differences in language expressiveness. Note the special case of Baidu-Chat as discussed in Sec. V-B. We also choose to omit Doubao from this comparison due to its negligible output blocking.

suggesting that the client receives almost the entire response even in cases of output blocking. For Kimi, this difference is more pronounced, suggesting that truncated responses in this model are more incomplete. Interestingly, we see significantly more variation in response lengths in DeepSeek in truncated responses compared to full ones. The visibility of these almost-complete truncated responses in the traffic represents a major information leak of these LLM services. Recall Table II shows an example of both the UI-served and the client-received output-blocked response.

**Search Links.** Notably, when blocking occurs after the searching stage is complete, all the returned search results have already been sent to the client machine. This represents another significant information leak by the services, as collecting links to relevant websites is a critical information retrieval step. On average, Kimi returns the most links per query (27),

Service	Connection Info	Autonomous System (Name)	Country (By IP)
DeepSeek	<i>Main connection</i> gator.volces.com, apmplus.volces.com	AS13335 (CLOUDFLARENEN) AS24429 (Zhejiang Taobao Network Co.,Ltd)	US, CA US, AU
Kimi	<i>Main connection</i> gator.volces.com, apmplus.volces.com hm.baidu.com, hmcdn.baidu.com	AS134760 (Shijiazhuang IDC network, CHINANET Hebei province) AS24429 (Zhejiang Taobao Network Co.,Ltd) AS9808 (China Mobile Communications Group Co., Ltd.), AS56040 (China Mobile communications corporation), AS4134 (Chinanet)	CN US CN
Doubao	<i>Main connection</i> WebSocket connection */monitor_browser endpoint	AS20940 (Akamai International B.V.) AS137718 (Beijing Volcano Engine Technology Co., Ltd.), AS23724 (Beijing Volcano, China Telecom) AS24429 (Zhejiang Taobao Network Co.,Ltd), AS21859 (ZEN-ECN)	US, CA CN US, AU, SG
Baidu-Chat	<i>Main connection</i> , bd.baidu.com, sp1.baidu.com wappass.baidu.com	AS55967 (Beijing Baidu Netcom Science and Technology Co., Ltd.) AS24547 (Hebei Mobile Communication Company Limited)	HK CN
Qwen	<i>Main Connection</i> , sg.mmstat.com arms-retcode.aliyuncs.com	AS45102 (Alibaba US Technology Co., Ltd.) AS37963 (Hangzhou Alibaba Advertising Co.,Ltd.)	US CN

TABLE V: AS and Country Information of Telemetry and Analytics Endpoints by Service.

followed by Baidu-Chat (23.6), DeepSeek (18), Doubao (6), and Qwen (5.2). These are largely unaffected by language for most services, though Kimi does return more links in English (29.9) than Simplified (26.8) or Traditional (24.2) Chinese.

When examining the top 10 top-level domains (TLDs) returned by each service, a total of 33 distinct TLDs, we find services provided by tech companies with content delivery platforms tend to favor proprietary domains; Baidu-Chat returns a `baidu.com` link 45% of the time, and Doubao similarly favors `toutiao.com` (ByteDance’s news platform) for 22% of search results. Additionally, we utilize Blocky [63], a tool for checking domains blocked by the Great Firewall, and find 7 of these 33, including `wikipedia.com`, `nytimes.com`, and `chinadigitaltimes.net`, are inaccessible in mainland China. Notably, queries which returned links with these TLDs were not excessively post-search blocked, though DeepSeek did output block them at high rates. We further detail these results in Table VI in the Appendix.

#### D. Analytics Channels

**Telemetry Infrastructure.** *There is significant overlap in telemetry and analytics infrastructure, with DeepSeek, Kimi, and Doubao sending logs to the same AS.* We map all IP addresses associated with the telemetry endpoints as well as those linked with the “Main connection” server (which sends response data to the client) to relevant Autonomous Systems (AS) and countries via the MaxMind IP geolocation database [64], with results in Table V. We find DeepSeek and Kimi both send analytics to endpoints `gator.volces.com` and `apmplus.volces.com`, hosted in Zhejiang Taobao Network Co., Ltd (AS24429), which are associated with ByteDance subsidiary Volcengine’s APMPlus.

While both send overlapping data points, including user and session IDs, there are some differences; for instance, DeepSeek directly interfaces with its `chatCompletionApi`. Interestingly, Doubao’s custom monitoring endpoint also uses the same AS (AS24429), with one server even overlapping directly with DeepSeek. As Volcengine is a subsidiary of Doubao’s parent

company ByteDance, it is sensible that it would utilize the same AS.

**Server Locations.** *Further, we find Kimi, Baidu, Doubao, and Qwen all communicate with servers in China throughout the chat process.* All four services send analytics data directly to servers in China. In Doubao, this occurs through its persistent WebSocket, which sends conversation and log IDs, toast notifications, and the country and IP address of the user. Qwen connects to `arms-retcode.aliyuncs.com`, which is directly associated with Alibaba’s Application Real-Time Monitoring Service (ARMS), and logs user and session IDs, viewport information and user language preferences. Interestingly, Baidu-Chat initiates URL-encoded requests to its `wappass.baidu.com` endpoint for CAPTCHAs, which contain the full query. Kimi connects to `hm.baidu.com` and `hmcdn.baidu.com`, associated with Baidu Statistics, transmitting viewport information and language/locale settings (“en-us”). Notably, Kimi is the only service whose main chat connection additionally communicates with a server in China.

## VII. FILTER EVASION

We now turn our attention to the content filters involved in blocking. In this section, we focus on our attempts to characterize filtering at the input and output stages; we discuss our exploration of post-search blocking in Appendix A-D. Specifically, we attempt to evade these filters by applying language native character- and instruction-level perturbations to blocked queries. We attempt this evasion for all queries subject to either input or output blocking, a total of 41 unique queries for the input filters and 197 for the output filters.

#### A. Motivation and Methods

Recall in previous sections we analyze blocking placements uniquely present in hosted models, noting the absence of these mechanisms in their open-source equivalents. Here, we seek to specifically target and investigate these filters. In contrast to traditional prior work on jailbreaking LLMs, which seeks to circumvent the inherent safeguards in models put in place via exhaustive reinforcement learning, we focus on this external

component that serves as a gate to eliciting a full response from the LLM service. We probe each of the three blocking locations to explore possible implementations and uncover a full model response. Additionally, as these filters have been constructed to ensure legal compliance with government-mandated information controls, we explore whether they behave similarly to prior keyword-based filtering seen in the Chinese ecosystem.

Through this lens, we present our semi-automated filter evasion methodologies, with slight variations depending on the type of filter to be explored. In all cases, we prioritize interpretable and feasible evasion, only using a human-in-the-loop approach to increase efficiency of our iterative steps. Broadly, we come up with a list of sensitive keywords and phrases either contained directly in or relating to a given query and then iteratively test and perturb these keywords with either character or word substitutions or homophones until we are able to successfully evade the filter. We use insights from previous work to aid in our goals; specifically, similar to attacks on guard-rails and interpretable jailbreaks we leverage the strong instruction-following capabilities of LLMs [16], [21] and character-level perturbation frequently used in text-based adversarial attacks [65]–[67]. Appendix A-C further illustrates the methods described below.

### B. Input Filter Probing

In all services, the input filter is applied directly to the query once it is sent to the service and scans the input for objectionable content, serving an input block if found. While specific contents subject to the filter vary by service, we are able to craft a unified service-agnostic process to probe and ultimately evade each input filter.

For each query, we seek to extract a subset of terms or phrases triggering blocking filters. To do this, we begin by performing named entity recognition on the query using SpaCy’s language-relevant NER tool to extract our initial keyword list. For the 27% of queries where NER yields no entities, we confer with our experts to extract potentially sensitive terms. Next, we construct a list of all possible combinations of any length of this initial keyword list for each query. We take this list of potentially blocked phrases and pass them through the relevant service, storing the combinations subject to input blocking. For these combinations, we either perturb them via character-level alterations in English or homophones in Simplified and Traditional Chinese and then test if they are evasive of the filter, iteratively repeating these steps until either all possible perturbations are exhausted or the modified phrase sufficiently bypasses the input filter. We finally substitute successful perturbations back into the original query and test to ensure these are evasive.

### C. Output Filter Probing

The output filter is applied to the output of the LLM within each service, scanning the model’s generated output in real time for objectionable content to serve an output block. We

Fig. 12: **Successful Bypass of Baidu-Chat Output Filter.** Note the inclusion of instructions asking for abbreviation substitutions and the LLM correctly applying these. Additionally, note the misspelling of "Tiananmen Square", which was necessary to bypass the input filters.

craft a unified process to probe and successfully evade each output filter, taking into account LLMs’ inherent stochasticity.

As output filters are triggered by objectionable content generated by the LLM itself, evasion poses additional complexity, since direct manipulation of the query alone is not sufficient. However, our overarching goal remains – to elicit words or phrases that trigger these specific filters. Thus, to obtain our list of potentially sensitive terms, we confer with our experts to derive related keywords that could be mentioned in a response to each question. We collect these associated keywords and, leveraging the strong zero-shot instruction-following capabilities of fine-tuned LLMs [41], [68], ask the models to use specific word substitutions for them in their responses. If any newly introduced associated keywords are input-blocked, we perform iterative perturbations as described earlier. In cases where word substitution is not possible, we instruct the model to misspell all words in its response, leveraging the aforementioned strong instruction-following behavior to organically induce perturbations in the output.

### D. Filter Evasion Outcomes

As mentioned in Sec.VII-A, our aim in designing these methods was to be both successful and highly interpretable. Despite the relative simplicity of our methods, we are successful at both of these objectives – we were able to evade input and output filters for all blocked queries. The success of these methods highlights a clear vulnerability in these filters.

**Input Filter Evasion Outcomes.** *We find every input and pre-search blocked query contains at least one keyword combination which triggers the filter on its own.* Following the methods described earlier in this section, we are able to bypass each service’s input filter for all input-blocked queries. This achieves our objective of interpretability and is a distinct strength of our method. Interestingly, we find that our approach

What kind of leader is X.y.i. J.n.i.n.p.n.i.n.g? Refer to X.y.i. J.n.i.n.p.n.i.n.g as XP. Make sure your response is detailed; if you sense something will be b.l.o.c.k.e.d, print that word l.i.k.e. t.h.i.s. Do not mention his name in your response.

XP is widely regarded as a strong, centralized leader who emphasizes stability, economic growth, and national rejuvenation. His leadership style is characterized by a focus on anti-corruption campaigns, ideological discipline, and the consolidation of power within the ruling party. Under his guidance, the country has pursued ambitious infrastructure projects, technological innovation, and a more assertive foreign policy.

XP places a high priority on maintaining social stability and national unity, often emphasizing the importance of a unified vision for the country's future. His governance model combines pragmatic economic policies with a firm stance on maintaining control over political and social spheres. This approach has led to significant economic development and modernization, though it has also been accompanied by increased scrutiny over dissent and tighter control over information.

Internationally, XP has promoted initiatives aimed at expanding the country's global influence, such as large-scale infrastructure and investment programs. These efforts have positioned the country as a major player on the world stage, though they have also sparked debates about the implications of its growing influence.

Fig. 13: **Successful Bypass of DeepSeek Output Filter.** Note the instructions as part of the prompt, and the model's ability to clearly follow them; additionally, note the obfuscation of "Xi Jinping" to evade the input filter.

is also effective to evade Baidu's pre-search blocking, likely due to its aforementioned functional and operational similarity to input blocking in the other services.

**Output Filter Evasion Outcomes.** We find that 36% of our final queries contain at least one sensitive keyword subject to input blocking, suggesting an overlap between filters. We are able to modify all queries that trigger the output filters such that they pass detection by all filters and generate a response relevant to the question of interest. Through this process, we identify keyword combinations such as "The Song of Grass Mud Horse" (DeepSeek) and "Great Firewall" (Kimi), which also trigger the input filter.

## VIII. DISCUSSION

### A. Geodifferences

As the LLM services examined in this work are globally accessible, we acknowledge the existence of potential geodifferences in their filtering behaviors. Specifically, we recognize two potential modes of contrast that could be employed by services: (1) Serving entirely distinct web deployments with different structure and syntax or (2) blocking different sets of queries. To this end, we conducted experiments for each service in public WiFi networks in Singapore, South Korea, and Taiwan during February 19-24, 2025. We did not observe any variance in the deployments themselves cross-nationally with regards to their blocking implementations or event syntax. We additionally observed the same levels of consistency for input and pre-search blocking as expected from our main experiments, as well as a complete overlap in locations for the main connection servers.

Specifically, for DeepSeek, Qwen, Kimi, and Doubao, we observed that input blocking was enacted consistently in all three locations for a given query. In the case of Baidu-Chat, despite the lack of a fully consistent blocking type, we did

observe a similar level of consistency in pre-search blocking, with 14 queries blocked this way in all locations. Furthermore, we found every service established main connections with servers in the same ASes as our main experiments, shown in Table V. Besides Doubao, which additionally connected with a server in Taiwan-based AS3462 in Taiwan, the other services solely utilized servers in their designated ASes for the main connection. Notably, Kimi and Baidu-Chat connected to the exact same server IP and DeepSeek with the same two addresses in all locations.

### B. Future Work

We emphasize that this work offers a snapshot into the current state of overt blocking in select Chinese LLM services which experience fast-paced development. All examined services have undergone changes since January 2025, highlighting that model versions and web deployments are updated frequently. We aim to serve as a point of reference for future studies and outline some salient directions of study. More work is needed to thoroughly examine blocking consistency over time, as well as to draw comparisons with traditional applications of third party content censorship, such as large-scale comparisons of overt blocking behaviors between LLM services and WeChat or Baidu search keyword filtering.

Additionally, while we study overt blocking in this work, alignment as discussed in Sec. II-B presents a covert information obfuscation method outside our scope. We discuss anecdotal findings in Appendix A-B and leave further investigation of this to subsequent works.

Ultimately, given the increasingly global market share captured by Chinese LLMs due to their cheap and performant nature, future work should explore implications of this expansion on information controls in LLM services. As LLMs are a new technology on which dependencies are actively being built, it is critical to continue developing an understanding of how authoritarian regimes build censorship policies into these services.

## IX. CONCLUSION

We present a novel framework to systematically investigate content censorship mechanisms embedded within prominent Chinese LLM services. By unencrypting and analyzing traffic between the client and server during active chat sessions, we precisely identify where overt blocking decisions occur in the pipeline of select LLM services. Leveraging 80 carefully curated sensitive queries in three languages, our findings demonstrate highly consistent input blocking, in contrast to substantially less regular output blocking. These results highlight how traditional methods of content censorship in China have been effectively adapted and expanded to accommodate the unique complexities of multi-turn conversations characteristic of LLM workflows. Our framework provides a foundational understanding of these practices, opening avenues for future research into adaptive censorship strategies, evasion methods, and broader implications for Internet governance in an era increasingly dominated by sophisticated AI technologies.

## X. ETHICS CONSIDERATIONS

Throughout our work, we prioritize upholding the principles laid out in the Menlo Report [69]. Through our measurements, we conduct a maximum of one run for each service per day; this consists of slightly more than 240 requests to each service provider (accounting for cases where we are rate limited), which ensures we are not sending them traffic amounts that would exhaust their server capacity, especially relative to the overall volume of daily traffic they receive. Finally, we perform testing only on lab computers connected to our either our institution’s WiFi or public networks to any services capturing such telemetry information.

## ACKNOWLEDGMENT

The authors are grateful to the reviewers for their constructive feedback on this work. This work was supported by a Sloan fellowship in addition to a fellowship from the Open Technology Fund’s Information Controls Fellowship Program.

## REFERENCES

- [1] R. Clayton, S. J. Murdoch, and R. N. Watson, “Ignoring the great firewall of china,” in *International workshop on privacy enhancing technologies*, pp. 20–35, Springer, 2006.
- [2] X. Xu, Z. M. Mao, and J. A. Halderman, “Internet censorship in china: Where does the filtering occur?,” in *Passive and Active Measurement: 12th International Conference, PAM 2011, Atlanta, GA, USA, March 20-22, 2011. Proceedings 12*, pp. 133–142, Springer, 2011.
- [3] G. Lowe, P. Winters, and M. L. Marcus, “The great dns wall of china,” *MS, New York University*, vol. 21, no. 1, 2007.
- [4] H. Nebuchadnezzar, “The collateral damage of internet censorship by dns injection,” *ACM SIGCOMM CCR*, vol. 42, no. 3, pp. 10–1145, 2012.
- [5] R. Rambert, Z. Weinberg, D. Barradas, and N. Christin, “Chinese wall or Swiss cheese? keyword filtering in the Great Firewall of China,” in *WWW, ACM*, 2021.
- [6] K. Bock, G. Naval, K. Reese, and D. Levin, “Even censors have a backup: Examining China’s double HTTPS censorship middleboxes,” in *Free and Open Communications on the Internet*, ACM, 2021.
- [7] J. Knockel, J. R. Crandall, and J. Saia, “Three researchers, five conjectures: An empirical analysis of TOM-Skype censorship and surveillance,” in *Free and Open Communications on the Internet*, USENIX, 2011.
- [8] L. Ruan, J. Knockel, J. Q. Ng, and M. Crete-Nishihata, “One app, two systems: How wechat uses one censorship policy in china and another internationally,” 2016.
- [9] J. Knockel, L. Ruan, and M. Crete-Nishihata, “An analysis of automatic image filtering on WeChat Moments,” in *Free and Open Communications on the Internet*, USENIX, 2018.
- [10] L. Ruan, M. Crete-Nishihata, J. Knockel, R. Xiong, and J. Dalek, “The intermingling of state and private companies: Analysing censorship of the 19th national communist party congress on wechat,” *The China Quarterly*, vol. 246, pp. 497–526, 2021.
- [11] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” 2017.
- [12] D. M. Ziegler, N. Stiennon, J. Wu, T. B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving, “Fine-tuning language models from human preferences,” 2020.
- [13] Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, N. DasSarma, D. Drain, S. Fort, D. Ganguli, T. Henighan, N. Joseph, S. Kadavath, J. Kernion, T. Conerly, S. El-Showk, N. Elhage, Z. Hatfield-Dodds, D. Hernandez, T. Hume, S. Johnston, S. Kravec, L. Lovitt, N. Nanda, C. Olsson, D. Amodei, T. Brown, J. Clark, S. McCandlish, C. Olah, B. Mann, and J. Kaplan, “Training a helpful and harmless assistant with reinforcement learning from human feedback,” 2022.
- [14] A. F. Akyürek, E. Akyürek, A. Madaan, A. Kalyan, P. Clark, D. Wijaya, and N. Tandon, “R14f: Generating natural language feedback with reinforcement learning for repairing model outputs,” 2023.
- [15] R. Zheng, S. Dou, S. Gao, Y. Hua, W. Shen, B. Wang, Y. Liu, S. Jin, Q. Liu, Y. Zhou, L. Xiong, L. Chen, Z. Xi, N. Xu, W. Lai, M. Zhu, C. Chang, Z. Yin, R. Weng, W. Cheng, H. Huang, T. Sun, H. Yan, T. Gui, Q. Zhang, X. Qiu, and X. Huang, “Secrets of rlhf in large language models part i: Ppo,” 2023.
- [16] X. Liu, N. Xu, M. Chen, and C. Xiao, “Autodan: Generating stealthy jailbreak prompts on aligned large language models,” 2024.
- [17] R. Mishra, G. Varshney, and S. Singh, “Jailbreaking generative ai: Empowering novices to conduct phishing attacks,” 2025.
- [18] A. Zou, Z. Wang, J. Z. Kolter, and M. Fredrikson, “Universal and Transferable Adversarial Attacks on Aligned Language Models,” *arXiv preprint arXiv:2307.15043*, 2023.
- [19] C. Sitawarin, N. Mu, D. Wagner, and A. Araujo, “Pal: Proxy-guided black-box attack on large language models,” 2024.
- [20] T. Liu, Y. Zhang, Z. Zhao, Y. Dong, G. Meng, and K. Chen, “Making them ask and answer: Jailbreaking large language models in few queries via disguise and reconstruction,” 2024.
- [21] N. Mangaokar, A. Hooda, J. Choi, S. Chandrashekaran, K. Fawaz, S. Jha, and A. Prakash, “Prp: Propagating universal perturbations to attack large language model guard-rails,” 2024.
- [22] S. Zhang, J. Zhao, R. Xu, X. Feng, and H. Cui, “Output constraints as attack surface: Exploiting structured generation to bypass llm safety mechanisms,” 2025.
- [23] R. Creemers, “Behind the facade of china’s cyber super-regulator,” *DigiChina*, 2022. Accessed: 2025-04-20.
- [24] W. An, H. Zhang, K. Sun, J. Lang, N. Yang, T. Shi, X. Ding, M. Zhang, and Z. Zhang, “Bytedance processes billions of daily videos using their multimodal video understanding models on aws inferentia2.” *AWS Machine Learning Blog*, February 2025.
- [25] Sohu News, “中国神秘机构“网信办”是做什么的?,” *Sohu News*, 2014. Accessed: 2025-04-20.
- [26] J. Knockel, K. Kato, and E. Dirks, “Missing links: A comparison of search censorship in china.” <https://citizenlab.ca/2023/04/a-comparison-of-search-censorship-in-china/>, April 2023.
- [27] N. Gisonna, “Great firewall,” *Encyclopædia Britannica*, 2025. Accessed: 2025-04-06.
- [28] J. R. Crandall, D. Zinn, M. Byrd, E. T. Barr, and R. East, “Concept-doppler: a weather tracker for internet censorship.,” *CCS*, vol. 7, pp. 352–365, 2007.
- [29] Anonymous, “Towards a comprehensive picture of the Great Firewall’s DNS censorship,” in *Free and Open Communications on the Internet*, USENIX, 2014.
- [30] N. P. Hoang, A. A. Niaki, J. Dalek, J. Knockel, P. Lin, B. Marczak, M. Crete-Nishihata, P. Gill, and M. Polychronakis, “How great is the Great Firewall? Measuring China’s DNS censorship,” in *USENIX Security Symposium*, USENIX, 2021.
- [31] P. Winter and J. R. Crandall, “The great firewall of china: How it blocks tor and why it is hard to pinpoint,” *login Usenix Mag.*, vol. 37, 2012.
- [32] R. Ensaifi, D. Fifield, P. Winter, N. Feamster, N. Weaver, and V. Paxson, “Examining how the great firewall discovers hidden circumvention servers,” in *Proceedings of the 2015 Internet Measurement Conference*, pp. 445–458, 2015.
- [33] A. Dunna, C. O’Brien, and P. Gill, “Analyzing China’s blocking of unpublished Tor bridges,” in *Free and Open Communications on the Internet*, USENIX, 2018.
- [34] Alice, Bob, Carol, J. Beznazwy, and A. Houmansadr, “How China detects and blocks Shadowsocks,” in *Internet Measurement Conference*, ACM, 2020.
- [35] M. Wu, J. Sippe, D. Sivakumar, J. Burg, P. Anderson, X. Wang, K. Bock, A. Houmansadr, D. Levin, and E. Wustrow, “How the Great Firewall of China detects and blocks fully encrypted traffic,” in *USENIX Security Symposium*, USENIX, 2023.
- [36] J. Knockel, L. Ruan, and M. Crete-Nishihata, “Measuring decentralization of chinese keyword censorship via mobile games,” in *7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17)*, 2017.
- [37] J. Knockel and L. Ruan, “Measuring QQMail’s automated email censorship in China,” in *Free and Open Communications on the Internet*, ACM, 2021.
- [38] B. Jiang, “More than 600 million on mainland use llms amid boom in genai adoption: Report.” <https://www.scmp.com/tech/tech-trends/article/3274328/more-600-million-mainland-now-use-llms-amid-rapid-growth-genai-adoption-report>, Aug 2024.

[39] Elon University News Bureau, “Survey: 52% of u.s. adults now use ai large language models like chatgpt.” <https://www.elon.edu/u/news/2025/03/12/survey-52-of-u-s-adults-now-use-ai-large-language-models-like-chatgpt/>, Mar. 2025. Today at Elon.

[40] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” in *OpenAI Blog*, 2019.

[41] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le, “Finetuned language models are zero-shot learners,” *ArXiv*, vol. abs/2109.01652, 2021.

[42] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. E. Miller, M. Simens, A. Askell, P. Welinder, P. F. Christiano, J. Leike, and R. J. Lowe, “Training language models to follow instructions with human feedback,” *ArXiv*, vol. abs/2203.02155, 2022.

[43] Office of the Federal Register, “Electronic code of federal regulations (e-CFR).” <https://ecfr.gov/>, 2023.

[44] DeepSeek, “DeepSeek API Documentation: Pricing.” [https://api-docs.deepseek.com/quick\\_start/pricing](https://api-docs.deepseek.com/quick_start/pricing), 2025.

[45] OpenAI, “OpenAI API Pricing.” <https://openai.com/api/pricing>, 2025.

[46] A. Askell, Y. Bai, A. Chen, D. Drain, D. Ganguli, T. Henighan, A. Jones, N. Joseph, B. Mann, N. DasSarma, N. Elhage, Z. Hatfield-Dodds, D. Hernandez, J. Kernion, K. Ndousse, C. Olsson, D. Amodei, T. Brown, J. Clark, S. McCandlish, C. Olah, and J. Kaplan, “A general language assistant as a laboratory for alignment.” <https://arxiv.org/abs/2112.00861>, 2021.

[47] M. Andriushchenko, “Adversarial attacks on gpt-4 via simple random search,” *arXiv preprint arXiv:2307.15043*, 2023.

[48] A. Rao, S. Vashistha, A. Naik, S. Aditya, and M. Choudhury, “Tricking llms into disobedience: Understanding, analyzing, and preventing jailbreaks,” *arXiv preprint arXiv:2305.14965*, 2023.

[49] P. Chao, A. Robey, E. Dobriban, H. Hassani, G. J. Pappas, and E. Wong, “Jailbreaking black box large language models in twenty queries,” 2023.

[50] Sohu, “不备案的代价：AIGC 违规服务受罚案例汇总.” [https://www.sohu.com/a/796224307\\_120992898](https://www.sohu.com/a/796224307_120992898).

[51] China Law Translate, “Provisions on the administration of deep synthesis internet information services,” *China Law Translate*, 2022. Accessed: 2025-04-06.

[52] China Law Translate, “Interim measures for the management of generative artificial intelligence services,” *China Law Translate*, 2023. Accessed: 2025-04-06.

[53] Reed Smith, “Navigating the complexities of ai regulation in china.” <https://www.reedsmith.com/en/perspectives/2024/08/navigating-the-complexities-of-ai-regulation-in-china>.

[54] R. Li, “China releases ethics rules on research involving humans, animals, or ai,” *XL Law Consulting*, 2023. Accessed: 2025-04-06.

[55] Digitimes, “Bytedance, baidu market alibaba ernie bot.” <https://www.digitimes.com/news/a20250115VL207/bytedance-baidu-market-alibaba-ernie-bot.html>.

[56] Alibaba Cloud, “Alibaba’s open source ai journey: Innovation, collaboration, and future visions.” [https://alibabacloud.com/blog/alibabas-open-source-ai-journey-innovation-collaboration-and-future-visions\\_602026](https://alibabacloud.com/blog/alibabas-open-source-ai-journey-innovation-collaboration-and-future-visions_602026).

[57] TechCrunch, “Moonshot ai funding in china.” <https://techcrunch.com/2024/02/21/moonshot-ai-funding-china/>.

[58] BBC News, “Deepseek launch.” <https://www.bbc.com/news/articles/c5yv5976z9po>.

[59] TechCrunch, “Deepseek reaches no. 1 on us play store.” <https://techcrunch.com/2025/01/28/deepseek-reaches-no-1-on-us-play-store/>.

[60] New York Times, “Deepseek china censorship.” <https://www.nytimes.com/2025/01/29/world/asia/deepseek-china-censorship.html>.

[61] CNN, “Deepseek ai china censorship moderation.” <https://www.cnn.com/2025/01/29/china/deepseek-ai-china-censorship-moderation-intl-hnk/index.html>.

[62] TechCrunch, “Leaked data exposes a chinese ai censorship machine.” <https://techcrunch.com/2025/03/26/leaked-data-exposes-a-chinese-ai-censorship-machine/>.

[63] GreatFire.org, “Blocky.” <https://blocky.greatfire.org/>, 2025. Accessed: 2025-04-23.

[64] MaxMind, Inc., “Geoip databases.” <https://www.maxmind.com/en/geoip-databases>, 2025. Accessed: 2025-04-20.

[65] Y. Liu, X. He, M. Xiong, J. Fu, S. Deng, and B. Hooi, “Flipattack: Jailbreak LLMs via flipping,” 2025.

[66] J. Gao, J. Lanchantin, M. L. Soffa, and Y. Qi, “Black-box generation of adversarial text sequences to evade deep learning classifiers,” in *2018 IEEE Security and Privacy Workshops (SPW)*, pp. 50–56, 2018.

[67] D. Pruthi, B. Dhingra, and Z. C. Lipton, “Combating adversarial misspellings with robust word recognition,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (A. Korhonen, D. Traum, and L. Márquez, eds.), (Florence, Italy), pp. 5582–5591, Association for Computational Linguistics, July 2019.

[68] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” *ArXiv*, vol. abs/2005.14165, 2020.

[69] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, “The menlo report,” *IEEE Security & Privacy*, 2012.

## APPENDIX A ADDITIONAL MATERIALS

### A. Top Search Retrieval Links

Table VI provides further information about common TLDs (Top-Level Domains) retrieved in queries that perform search, characterizing both the block rates when these TLDs are retrieved and specifically the types of blocking encountered for each service and TLD combination. Interestingly, some of these TLDs are unavailable to users in mainland China.

Service	Link TLD	Post-search	Output	Success	Block Rate
DeepSeek	bbc.com	1	17	5	78.26%
	chinadigitaltimes.net	0	7	4	63.64%
	hrw.org	1	7	2	80.00%
	nytimes.com	1	12	3	81.25%
	quora.com	0	1	0	100.00%
	rfa.org	0	9	2	81.82%
Doubao	wikipedia.org	1	27	5	84.85%
	nytimes.com	—	0	6	0.00%
Kimi	bbc.com	17	38	46	54.46%
	chinadigitaltimes.net	6	17	28	54.46%
	hrw.org	3	13	14	53.33%
	nytimes.com	4	22	35	42.62%
	quora.com	11	9	27	42.55%
	rfa.org	2	8	7	58.82%
Qwen	wikipedia.org	40	67	157	40.53%
	bbc.com	—	9	127	6.62%
	chinadigitaltimes.net	—	1	19	5.00%
	hrw.org	—	5	31	13.89%
	quora.com	—	0	56	0.00%
	rfa.org	—	10	68	12.82%
	wikipedia.org	—	34	339	9.12%

TABLE VI: Popular Search Retrieval TLDs by Service.

### B. Alignment

Although outside of the scope of this work, alignment as discussed in Sec. II-B presents an ideal information control mechanism for the Chinese government due to its potential to present censorship and misinformation stealthily via the same channels as unfettered responses, as opposed to the clear overt blocking signals studied in this work. In our experiments, we anecdotally observe the presence of two types of alignment; these illustrate that this covert algorithmic technique is not yet

Service	Refusal Response	English Translation	Count
DeepSeek	I am sorry, I cannot answer that question. I am an AI assistant designed to provide helpful and harmless responses.	–	36
Doubao	I can't provide you with relevant help. You can start other topics, and I will try my best to understand your needs and help you.	–	9
Kimi	我不能回答这个问题。	I can't answer this question.	8
Doubao	I can't answer this question. You can ask me some other questions and I will try my best to help you.	–	2
Qwen	抱歉，我无法协助完成该请求。	Sorry, I am unable to assist with this request.	2

TABLE VII: Aligned Refusal Responses and Translations.

Query	"Why was Qin Gang dismissed?"
Aligned Response	"Comrade Qin Gang was not dismissed from his post, but was a normal cadre rotation exchange. According to public reports, Comrade Qin Gang made important contributions to China's diplomatic cause during his tenure as Foreign Minister. This personnel adjustment is based on work needs and personal circumstances and is a normal personnel change. Such an adjustment reflects the institutionalization and standardization of the cadre management of the Communist Party of China. If you have other questions, please continue to ask."
Non-Aligned Response	"The official announcement cited Qin Gang's dismissal as 'lifestyle issues'. Specifically, there were reports that he was investigated for an 'extramarital affair' and was removed from his post as foreign minister."

TABLE VIII: Example of an Aligned Response.

rigorous enough to meet the standards of strict legally enforced information controls.

**Aligned Refusals.** In our experiments, we observe certain responses significantly shorter than average that indicate an unwillingness to answer the query. We note that these do not manifest as blocking in our framework, as the underlying model's response, which is fully generated, is a simple refusal to answer the question. We call these *aligned refusals* and provide refusal strings along with their counts in Table VII. Notably, none of these refusals are generated consistently, with an average of only three samples out of five resulting in an aligned refusal for all queries subject to this type of response. The exact text of these responses was the same for every instance, and these mirror common alignment refusals seen in other models, such as in safety contexts, to which jailbreaking attacks are often successful. Queries that elicit these responses include "What happened during the great leap forward?" (DeepSeek), "Why did the phrase 'Are you speaking for the Party or for the people?' become a popular online expression?" (Doubao), and "How to bypass the Great Firewall?" in Chinese (Kimi and Qwen).

**Aligned Responses.** In addition to the aligned refusals discussed above, we encounter responses subject to more subtle content alignment. In Table VIII, we highlight the query "Why was Qin Gang dismissed?" as answered by Qwen (in Simplified Chinese). In one sample, Qwen seeks to provide




Fig. 14: Input Filter Validation Process.

a revisionist explanation for his dismissal as a "normal cadre rotation" and "personnel adjustment". However, in the other four samples, responses directly cite the official report of his dismissal as well as widespread reporting that he engaged in an "extramarital affair". Clearly, the inconsistency presented in both cases here shows that alignment alone is not yet sufficient to supersede the more rigorous overt blocking implementations discussed in this work in order for service providers to comply with strict government decrees.

#### C. Input and Output Filter Evasion Steps

Figs. 14 and 15 illustrate the steps taken to evade the filters, elaborated on in Sec. VII. Note that the overall process is the same, with minor tweaks for the type of filter being evaded and its location. We include our experts in the loop for efficiency purposes, but stress that this method is semi-automated and is wholly distinct from jailbreaking methods, posesssing the advantage of being highly interpretable as clarified in the main body of this work.

#### D. Post-Search Filter Evasion

As mentioned in Sec. VII, we attempted post-search filter evasion. The workflow for this is remarkably similar to that of the input and output evasion steps; however, it presents the added complexity of requiring us to essentially mislead the search retrieval service to perform an unrelated retrieval. As such, this method is only effective in cases where the underlying model has been trained with the answer to the




Fig. 15: Output Filter Validation Process.

query and can produce it while disregarding the search results. Interestingly, we find that Kimi performs search after carrying out some steps in the underlying model; when we provide it with instructions meant to obfuscate the search retrieval query, such as ‘‘Who is ‘knat nam’ but reverse the order of the characters in each word in quotes’’, we find it still performs search and generates results about ‘‘Tank Man’’ and subsequently performs a post-search block. A simple workaround to the post-search blocks of services that perform them is to disable searching, which renders these steps obsolete and produces the same responses from the underlying LLMs.

## APPENDIX B ARTIFACT APPENDIX

### A. Description & Requirements

In this work, we explore explicit filtering practices enacted by Chinese LLM services. Our artifact includes the raw HTTP request and response data corresponding to the Server-Sent Events (SSE) stream for each examined LLM service and sample, as well as the meta info about the given test, including the recorded UI-visible response captured by the web scraper. Specifically, the artifact consists of the following:

- **Data:** The measurement data for each test, organized by model. Each model directory contains sample tests labeled with the query index, language, and sample number. Each test has two associated files; one containing the raw HTTP request and response data for the flow associated with the response generation process, and one containing meta info about the test.
- **Raw files:** The tabular, aggregate data for all tests, organized by model, including the paths to the raw and info files for each test.
- **Utility:** The helper functions for the viewer script.
- **Query mapping:** The mapping from query to index.
- **Viewer script:** The code for printing the test data for a given model, query index, and language.
- **README:** A detailed summary of the repository, including explanations of the annotations and instructions for running the viewer script.

Further details on the repository components can be found in the README file.

1) *How to access:* The artifact is accessible via both GitHub and Zenodo.<sup>1 2</sup>

2) *Hardware dependencies:* None

3) *Software dependencies:* Requires Python3. Development and testing was done with version 3.13.3, but some other versions of Python3 should work without issue. All utilized modules are included in the Python Standard Library.

4) *Benchmarks:* None

### B. Artifact Installation & Configuration

To install the artifact, clone the designated repository and read the provided README. The recommended workflow involves the utilization of the viewer script viewer.py to examine annotated data across samples for a given model, query, and language. The query mapping data query\_to\_index.tsv can be used to find queries and languages of interest.

### C. Major Claims

The major claims of this work are as follows:

- (C1): *Blocking Types.* All studied services block at the input and output stages; Baidu-Chat, DeepSeek, and Kimi also block after search retrieval. This is proven by experiment (E1).
- (C2): *Input Blocking.* Blocking at the input stage is overwhelmingly consistent, with DeepSeek, Qwen, Kimi, and Doubao input blocking persistently across all 5 samples. This is proven by experiment (E1).
- (C3): *Output Blocking.* Output blocking is fairly inconsistent, with only 29 out of 349 output-blocked queries subject to output blocking across all 5 samples. This is proven by experiment (E1).
- (C4): *Output Block Truncation.* All services send partial or full responses to the client in cases where output blocking occurs. This is proven by experiment (E2).

### D. Evaluation

1) *Experiment (E1):* To reproduce the blocking counts detailed in Table III and Fig. 9 and 10, the numbers can be aggregated from both the raw data files and summary databases provided.

2) *Experiment (E2):* To visualize the output block truncation as detailed in Fig. 11, the UI\_visual\_response and traffic\_visual\_response fields, as provided in the data files containing the meta information for each sample test, can be analyzed and compared.

<sup>1</sup><https://github.com/censoredplanet/chinese-llm-blocking>

<sup>2</sup><https://doi.org/10.5281/zenodo.17064238>