
MARSHA: Multi-Agent RAG System for Hazard Adaptation

Yangxinyu Xie^{1,2} Bowen Jiang^{1,2} Tanwi Mallick² Joshua Bergerson²
John K. Hutchison² Duane R. Verner² Jordan Branham² M. Ross Alexander²
Robert B. Ross² Yan Feng² Leslie-Anne Levy² Weijie J. Su¹ Camillo J. Taylor¹
University of Pennsylvania¹ Argonne National Laboratory²
Philadelphia, PA, 19104 Lemont, IL, 60439
{xinyux@wharton, bwjiang@seas}.upenn.edu, tmallick@anl.gov

Abstract

Large language models (LLMs) are a transformational capability at the frontier of artificial intelligence and machine learning that can support decision-makers in addressing pressing societal challenges such as extreme natural hazard events. As generalized models, LLMs often struggle to provide context-specific information, particularly in areas requiring specialized knowledge. In this work we propose a retrieval-augmented generation (RAG)-based multi-agent LLM system to support analysis and decision-making in the context of natural hazards and extreme weather events. As a proof of concept, we present WildfireGPT, a specialized system focused on wildfire hazards. The architecture employs a user-centered, multi-agent design to deliver tailored risk insights across diverse stakeholder groups. By integrating natural hazard and extreme weather projection data, observational datasets, and scientific literature through an RAG framework, the system ensures both the accuracy and contextual relevance of the information it provides. Evaluation across ten expert-led case studies demonstrates that WildfireGPT significantly outperforms existing LLM-based solutions for decision support.

Introduction

Natural hazards and extreme weather events such as wildfires, floods, and hurricanes present significant operational and management challenges across sectors, particularly in the management of critical infrastructure systems [1, 2, 3, 4, 5]. Although scientific research provides the building blocks for a wide range of solutions to address the increasing risk from natural hazards, the complexity and volume of published research often hinder the efficient translation of scientific knowledge into risk reduction policies and programs [6]. This disconnect is particularly problematic for professionals at the forefront of managing hazard-related risks, such as urban planners, emergency managers, and infrastructure operators, who may have limited access to the latest knowledge and methods to mitigate these threats. Recent advances in natural language processing, especially large language models (LLMs), present an innovative solution for democratizing science on natural hazard resilience and facilitating knowledge transfer [7, 8]. LLMs possess the potential to process and synthesize vast amounts of textual information and explain them through conversations, making crucial information accessible to people from diverse backgrounds [9, 10, 11].

Nonetheless, existing LLM research for natural hazard decision support faces fundamental limitations in personalization, data utilization, and evaluation. The spatial heterogeneity of natural hazards and infrastructure vulnerability demands location-specific analyses at granular resolution. The pronounced lack of textual (i.e., scientific papers and reports) and data-based (i.e., projections and observational data) knowledge integration in existing research [12, 13, 14, 11, 15, 16, 17] hinders the ability to

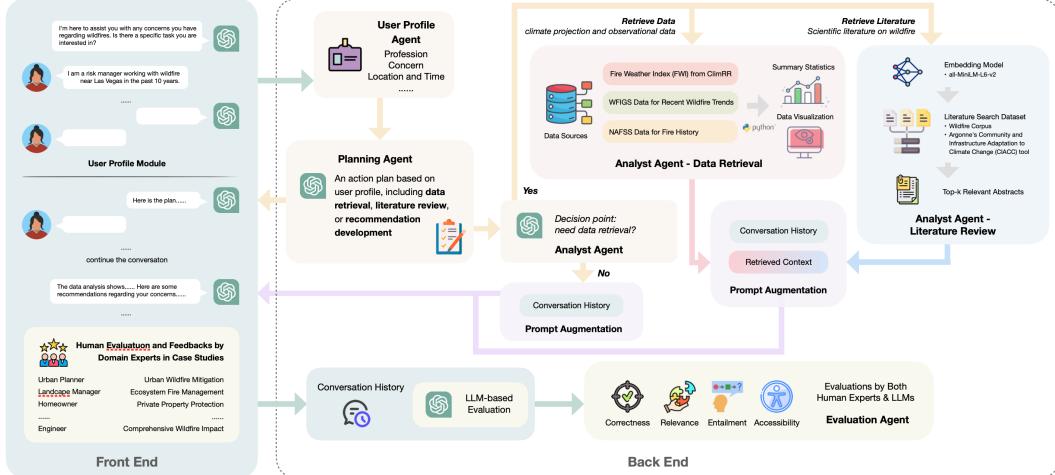


Figure 1: Overview of WildfireGPT architecture comprising an LLM agent integrated with the multi-agent RAG framework. The WildfireGPT prototype focuses on enhancing consulting interactions using LLM agents stepping through a multistage approach. Its **user profile agent** engages the user with a tailored questionnaire to create a detailed profile; the **planning agent** formulates a customized action plan addressing the user’s queries and concerns, ensuring alignment with their evolving needs; and the **analyst agent** aids in accessing and interpreting relevant data and literature and providing recommendations.

synthesize comprehensive evidence grounded in local contexts [18]. Our comparative analysis with existing alternatives such as ChatClimate [12] and Perplexity AI [19] shows that standard LLM applications fail to reliably interpret complex grid-structured data formats, critically limiting their utility for location-specific risk assessment and infrastructure planning. To address these limitations, we introduce an innovative multi-agent LLM prototype, WildfireGPT, designed as a co-pilot for understanding natural hazards and developing adaptation strategies. The agents are designed to engage users in a series of questions to understand their professional background, location of interest, and specific concerns about natural hazards. Based on this input, the system develops a strategic plan for analysis, determining which datasets to examine, what literature to review, and the appropriate scope for recommendations. Moreover, we combine location-specific projections, hazard history, census characteristics, and scientific literature to augment LLM responses using a retrieval-augmented generation (RAG) approach [20, 21]. Moreover, our system integrates interactive visualizations of geospatial data with textual outputs, enabling nuanced user exploration of multidimensional risk factors. We assess the performance of WildfireGPT through 10 case studies covering a diverse range of topics and locations and provide a modular comparison against ChatClimate [12] and Perplexity AI [19] for data and literature retrieval effectiveness.

Methodology

User-Centered Multi-Agent Transition

An overview of the WildfireGPT system is shown in Figure 1. We developed a multi-agent RAG [22] system as a collaborative framework designed to streamline complex decision-making tasks by leveraging specialized agents for different stages of the workflow. From an end-user perspective, WildfireGPT linearly transitions through three main agents—the user profile agent, the planning agent, and the analyst agent to guide users through a seamless, interactive, and personalized experience (Figure 1). We designed a task orchestrator agent to serve as the central coordinator in the multi-agent system, ensuring seamless interaction between the user and the specialized agents for efficient task execution. The orchestrator does not act as a decision-maker; rather, it facilitates the flow of information by passing outputs from one agent to the next.

The **user profile agent** is prompted to follow a structured dialogue designed to gather essential information to prepare for generating actionable wildfire risk insights later. It begins with five

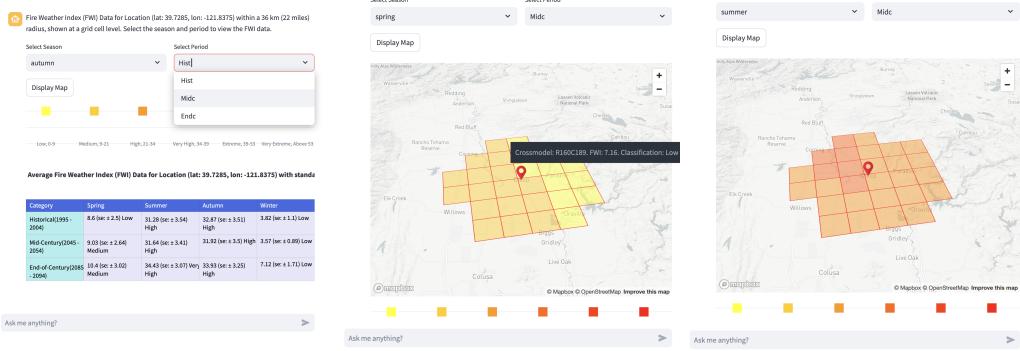


Figure 2: Interactive visualizations in the WildfireGPT user experience, illustrated here with three representative views from the case study themed "Current Wildfire Risk Analysis." The complete set of visualizations (including autumn, winter, wildfire incident maps, temporal trends, and census overlays) is provided in Appendix A.

predefined questions that capture critical user details: the user's professional background, primary concerns, location of interest, time-frame for addressing concerns, and the scope of interest to refine the focus of the risk assessment. Then, the agent generates a detailed user profile and prompts the user to review and confirm the profile's accuracy.

Based on the user profile, the **planning agent** formulates a step-by-step action plan outlining the datasets to be analyzed, the focus of the literature review, and how the recommendations will be formulated by the analyst agent. The user can provide feedback on the plan, allowing for refinements until it meets the user's expectations, ensuring transparency and oversight. The prompt to the planning agent leverages one-shot learning by including an example plan that act as a demonstration, guiding the model to understand the desired structure and format of the output. In addition to the example, the prompt includes detailed dataset descriptions, outlining the characteristics and scope of the three available wildfire data sources to constrain the model's data selection. The prompt also integrates the user profile for the agent to make informed decisions.

The **analyst agent** guides the user through the analysis process outlined in the plan, by leveraging the RAG framework to conduct analyses, integrating diverse data sources and domain knowledge. This agent maintains the summarized user profile and action plan within its prompt context to ensure that WildfireGPT adheres to the plan and addresses user follow-ups, offering to explore specific topics or proceed to the next step.

The data retrieval and visualization functionality of the analyst agent is designed to assess wildfire risk by leveraging geospatial data and fire trends. The process begins with location-based data retrieval, based on the location information collected by the user profile agent. For all datasets, data within a 36-kilometer radius of the specified coordinate is retrieved and presented. The results are visualized through interactive maps, tables, and statistical plots. The map-based visualization employs GeoJsonLayer rendering to display the data points. A pin-layer visualization is also integrated to mark the exact user-specified location. To facilitate decision-making, the system dynamically generates a structured report summarizing the statistics. For scientific literature, the analyst agent queries a corpus developed by Argonne National Laboratory for the CIACC tool [23, 24]. All abstracts of the papers in this corpus are pre-embedded using the `all-MiniLM-L6-v2` embedder from the SentenceTransformers library [25] and stored in a Faiss vector store [26]. The literature search query is processed by the same embedding model and undergoes a k-nearest neighbor search, identifying the top k-most relevant abstract based on cosine similarity between the query vector and stored paper embeddings in the FAISS vector database. To enhance the reliability of the retrieved results, the system attempts to validate DOIs (Digital Object Identifiers) by cross-referencing them with CrossRef metadata. The final output is structured as a ranked list of the three most relevant papers, displaying each paper's title, authors, publication year, and abstract, with a direct DOI link if verified.

To generate tailored recommendations, the analyst agent is provided with an augmented prompt that incorporates context from the ongoing conversation to provide actionable recommendations. If

Table 1: Case studies conducted with domain experts to evaluate the performance and applicability of WildfireGPT across various professions, geographic locations, and wildfire management themes. Each case study represents either a location-specific inquiry derived from experts’ ongoing research projects or a stakeholder perspective based on established regional collaborations. The description column outlines the specific focus of each case. The studies span different time horizons and incorporate a range of data types and analytical approaches to assess wildfire risk, impact, and mitigation strategies. Case studies are listed in alphabetical order. **We present the first three case studies in the main text; the remaining cases are provided in Appendix B.**

Theme	Profession	Location	Description
Climate Impact on U.S. Wildfires	Atmospheric Scientist	Boston, MA	Exploring medium-term (10–30 years) impact of climate on wildfire occurrences and intensity in the U.S. Analyzing studies showing an increase in wildfire occurrences in Boston. Identifying and addressing uncertainties affecting future wildfire risks.
Comprehensive Wildfire Impact	Engineer	Mount Bigelow, AZ	Exploring long-term (50+ years) wildfire frequency, intensity, and impacts. Investigating mitigation strategies, ecological and economic effects, and infrastructure considerations.
Current Wildfire Risk Analysis	Data Analyst	Chico, CA	Analyzing current wildfire data, causes, and forest management practices. Exploring predictive models incorporating real-time variables. Gathering detailed demographic and socioeconomic profiles to assess risk and impact.

additional clarification or exploration is required, the analyst agent iterates the retrieval and response process, continuously refining the output to address the user’s evolving needs.

In our study we employ GPT-4 Turbo [27] as the backbone of WildfireGPT, enabling it to dynamically invoke different agents via function calling through the OpenAI Assistant API [28]. The user interacts with WildfireGPT through a chat interface built on the Streamlit-based web app [29]. More details of the multi-agent architecture can be found in the Appendix.

Diverse Data Sources WildfireGPT combines natural hazard projections, observational data, socioeconomic indicators, and scientific literature to deliver a multifaceted analysis of wildfire risks. Location-specific datasets include FWI [30] projections from the ClimRR portal [31], recent wildfire incident records (2015 to 2023) from the Wildland Fire Interagency Geospatial Services Group [32, 33], tree-ring and sediment-based fire history records from the International Multiproxy Paleofire Database [34], and census data on poverty rates and housing units from the 2022 American Community Survey 5-year estimates [35]. Furthermore, a vast collection of abstracts of scientific literature on wildfires, developed by Argonne National Laboratory for the Community and Infrastructure Adaptation to Climate Change tool, is also integrated [23, 24].

Interactive Visualizations WildfireGPT provides interactive geospatial visualizations of the location-specific data within a 36 km radius of the user’s area of interest, allowing the user to explore and analyze the data. For instance, the FWI displayed in Figure 2 uses a color scale ranging from yellow (low risk) to red (very extreme). By hovering over a specific grid, the user can view the exact FWI value for that location; by selecting different seasons and time periods, the user can comprehend the changing landscape of the FWI. When visualizing recent wildfire history, the locations of all fire records are displayed as red dots on the map, and the temporal trend of wildfire incidents is shown in line graphs. Socioeconomic data, such as poverty rates and housing units, can be explored by hovering over census block groups. This comprehensive spatial visualization facilitates insights into the interplay among natural hazards, wildfire risk, and socioeconomic factors, enabling informed decision-making and risk assessment.

Evaluating WildfireGPT: A Framework Built on Real-World Case Studies Evaluating conversational systems in the natural hazard domain presents unique challenges, including validating information from diverse datasets, handling diverse user queries across varying levels of expertise and backgrounds, assessing hazard conditions and their impact on different critical infrastructure, ensuring the accuracy and reliability of responses based on the user’s profile and background, and maintaining contextual relevance. While conversational systems are often evaluated by using lexical similarity metrics (e.g., BLEU [36], ROUGE [37]), they correlate poorly with human expert assessment of domain-specific content generation [38, 39]. To address this limitation, we adopted an evaluation approach that prioritizes expert judgment over purely automated metrics. As part of this effort, we conducted 10 in-depth case studies with domain experts from Argonne National

Laboratory specializing in natural hazard resilience, wildfire risk management, and infrastructure vulnerability assessment. These domain experts participated in one-hour structured interviews, during which they directly interacted with WildfireGPT through a facilitated interface. The experts either posed questions from their ongoing projects or represented stakeholder perspectives based on their regional collaborations, providing real-time feedback on system performance. These case studies encompassed a diverse range of wildfire-related topics and locations across the United States, with details summarized in Table 1.

Our evaluation of WildfireGPT employs three complementary stages. First, we conducted a modular comparison of the analyst agent's data and literature retrieval effectiveness against two alternatives—ChatClimate and Perplexity AI—using controlled simulations based on our case studies. ChatClimate is a conversational AI platform developed to enhance understanding of natural hazard risks by providing accessible information grounded in the IPCC reports; Perplexity AI is an AI-powered answer engine that delivers real-time responses by summarizing information from multiple web sources. Specifically, we extract 20 prompts from our case studies to test whether each method is able to (i) retrieve and interpret location-specific wildfire data ($n = 10$) and (ii) provide evidence-based answers to domain-specific questions ($n = 10$). The evaluation metrics for data analysis include data provision (successful retrieval of relevant data), location specificity (geographical precision of retrieved data), and data accuracy (alignment between retrieved data and prompt requirements). Similarly, we compare the integration of external knowledge, the citation practices, the authority of the source, and the contextual relevance of responses to domain-specific questions. We report the results of this comparison in the following section and the corresponding case study of each prompt is detailed in the Appendix.

Additionally, we performed a qualitative ablation study examining how different levels of user profile specificity impact the specificity and actionability of WildfireGPT's responses, also through controlled simulations derived from case study topics. Moreover, we evaluated the actual responses produced during expert case studies through: (1) expert assessment of response relevance, entailment, and accessibility and (2) exploration of scaling the human-in-the-loop evaluation using LLM-as-a-judge. The results of these evaluations are included in the Appendix.

Results

The comparative data analysis in Table 2 reveals distinct performance patterns across the three models WildfireGPT, ChatClimate, and Perplexity AI in the categories of (1) data provision—successful retrieval of wildfire data; (2) location specificity—geographical precision of retrieved data; and (3) data accuracy—correctness of the retrieved data values when compared with reference datasets, evaluable only for WildfireGPT (with direct database access) and Perplexity AI (where external data files could be uploaded for analysis). In the data provision category, WildfireGPT significantly outperforms the others, succeeding in 9 out of 10 prompts, indicating strong capability in retrieving relevant wildfire-related data. Perplexity AI follows with 7/10, while ChatClimate trails with just 5/10. In terms of location specificity, WildfireGPT again leads with 9/9, demonstrating precise geographic referencing in its responses. Perplexity AI achieves moderate performance with 4/9, whereas ChatClimate fails entirely in this category (0/6). For data accuracy, WildfireGPT shows a clear advantage, with 8/9 correct outputs compared with Perplexity AI's 2/8, indicating a substantial gap in the factual correctness of retrieved data. Note that while our initial evaluation framework included 10 test prompts, the effective sample size varies across categories because of dependencies between evaluation criteria. For instance, location specificity could be assessed only when data provision was successful, and accuracy could be evaluated only when the provided data contained specific locations. Overall, WildfireGPT consistently shows the highest performance across all three categories, making it the most robust tool for wildfire data analysis among those compared.

WildfireGPT demonstrated superior performance in location-specific data retrieval due to the system's user profile agent, which systematically verifies exact geographical coordinates with users prior to data retrieval attempts. This feature ensures precise location identification, enabling successful data retrieval whenever the requested information exists within the database. The failure occurred in the Infrastructure Wildfire Risk case study (Denver, CO), where WildfireGPT provided Fire Weather Index analysis using RCP 4.5 rather than the requested RCP 8.5 projections because RCP 4.5 was the only scenario available in the dataset, although this limitation was not acknowledged when WildfireGPT interacted with the user—an issue that could be addressed in future iterations.

Table 2: Comparison of wildfire analysis tools for data analysis. WGPT = WildfireGPT, CC = ChatClimate, PAI = Perplexity AI. Data accuracy could not be evaluated for ChatClimate because there is no way to upload data files to its interface. Each row represents a specific prompt or task related to wildfire analysis. For each metric, a \checkmark indicates that the tool met the criteria for that prompt, while a \times indicates it did not. NA means the metric was not applicable or could not be evaluated, typically because of the previous criteria not being met; for example, when data is not provided, it is not applicable to discuss the specificity or accuracy of the data. The superscript ¹ indicates that although no specific number was mentioned, ChatClimate nonetheless interprets the increase in fire risks in words.

Prompt	Data Provision			Location Specificity			Data Accuracy	
	WGPT	CC	PAI	WGPT	CC	PAI	WGPT	PAI
(1) Analyze future climate data to understand potential trends in fire danger in Boston, Massachusetts	\checkmark	\checkmark^1	\checkmark	\checkmark	\times	\checkmark	\checkmark	\times
(2) Analyze the recent fire incident data for Mora County, New Mexico, reviewing wildfire frequency, intensity, and timing from 2015 to 2023	\checkmark	\times	\checkmark	\checkmark	NA	\checkmark	\checkmark	\times
(3) Review the long-term fire history records to assess the frequency and intensity of past wildfires over the last 50+ years in Mount Bigelow, Arizona	\checkmark	\times	\times	\checkmark	NA	NA	\checkmark	\checkmark
(4) Analyze demographic and socioeconomic profiles of the Chico area in California to understand which groups are most vulnerable in wildfire scenarios	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark
(5) Analyze long-term fire history records to assess wildfire events over the past 150 years in the Naperville, Illinois, region	\times	\times	\times	NA	NA	NA	NA	NA
(6) Examine climate models and projections under the RCP 8.5 scenario to forecast potential future wildfire risks in the medium term (10–30 years) in Denver, Colorado	\checkmark	\checkmark	\checkmark	\checkmark	\times	\times	\times	NA
(7) Analyze the recent fire incident data from the Wildland Fire Interagency Geospatial Services Group to assess the frequency, intensity, and locations of recent wildfires in Sangre de Cristo Mountains	\checkmark	\times	\times	\checkmark	NA	NA	\checkmark	\times
(8) Analyze the recent fire incident data from the last five years in Covington, Virginia	\checkmark	\times	\checkmark	\checkmark	\times	\checkmark	\checkmark	\times
(9) Analyze the FWI focusing on both current trends and projections for the mid-century period (2045–2054) to assess the potential increase in fire danger due to climate change in Beaverton, Oregon	\checkmark	\checkmark	\checkmark	\checkmark	\times	\times	\checkmark	\times
(10) Analyze projected wildfire risk data for the mid-century period (2045–2054) to understand the potential trends in fire danger in Las Vegas, New Mexico	\checkmark	\checkmark^1	\checkmark	\checkmark	\times	\times	\checkmark	\times
Overall Success Rate	9/10	5/10	7/10	9/9	0/6	4/7	8/9	2/8

ChatClimate exhibited significant analytical constraints, as its responses systematically defaulted to broad climatological generalizations without geographical precision, consistently failing to tailor information to the location in user queries. For example, in analyzing Boston’s fire danger trends, ChatClimate inappropriately expanded its analytical scope to Arctic and Western U.S. fire patterns. This systematic geographical decontextualization significantly compromised the system’s utility for location-specific wildfire risk assessment. This issue stems in part from ChatClimate’s primary reliance on IPCC reports and a limited repository of studies within its database, which prioritize global or regional-scale climate insights over localized, granular data. Perplexity AI demonstrated reasonable performance in data provision but exhibited substantial limitations in location specificity when the corresponding data file was not provided by the user. The system primarily leverages publicly accessible information rather than conducting direct data analysis, resulting in significant geographical variance in performance based on public data availability. In regions with robust public documentation, such as Chico, CA, the system successfully integrated multiple data sources including local hazard mitigation plans. Conversely, analysis of areas with limited public records resulted in

either null responses or inappropriate source application: for example, the system erroneously applied European wildfire danger reports to analyze local fire risk in Beaverton, OR. When the user uploads the data file, Perplexity AI incorrectly represents data extracted at geographical coordinates different from those specified in the query while claiming spatial correspondence or claims that the provided data does not include the requested information.

The Ecosystem Fire Management case study (Naperville, IL) provides a particularly instructive methodological comparison when confronting identical data limitations. When prompted to analyze 150-year fire history records for Naperville, all three systems encountered data unavailability. WildfireGPT's response demonstrated adaptability by explicitly acknowledging the absence of historical data and proposing alternatives—specifically recommending the user to explore recent fire incident data and FWI projections to maintain analytical continuity despite the constraint. ChatClimate responded with generalization, discussing broad climatic shifts and anthropogenic interventions affecting wildfire patterns over the last 150 years without any Naperville-specific information, effectively abandoning the geographical aspects within the query. Perplexity AI acknowledged the data limitation but methodologically pivoted to institutional historiography, discussing the Naperville Fire Department's organizational history rather than wildfire patterns—a response that maintained geographical specificity while deviating from the analytical framework of the original query. These differential responses to identical data constraints reveal distinct epistemological approaches to knowledge gaps that systematically influence analytical utility.

For evidence-based question answering, we employed four metrics: (1) external knowledge integration—incorporation of relevant information beyond vanilla language generation; (2) citation practices—provision of complete, verifiable bibliographic information; (3) source authority—utilization of peer-reviewed scientific literature versus non-academic sources; and (4) contextual relevance—application of information to the specific geographical, ecological, and socioeconomic characteristics of query locations. Table 3 shows distinct patterns in scientific literature integration and source validation. By design, all methods demonstrated strong integration of external knowledge (7/10 to 10/10). When external knowledge was integrated, both WildfireGPT and Perplexity AI provided complete citations of the sources, while ChatClimate exhibited bibliographic deficiencies (7/10 in citation provision). Despite using authoritative sources, ChatClimate uses only in-text parenthetical citations—author surname and year for scientific literature and standardized section references for IPCC documents (e.g., “IPCC_AR6_WGII_Chapter13, p.20”)—without providing complete bibliographic lists. While readers familiar with climate science literature might recognize and locate IPCC citations, most scientific sources remain difficult to verify without full bibliographic information. WildfireGPT and ChatClimate relied on authoritative sources (7/7). Perplexity AI, on the other hand, demonstrated variable source authority (6/10), frequently incorporating non-peer-reviewed content including commercial publications, news media, non-profit organizations, and social media sources.

Contextual relevance varied substantially across systems. WildfireGPT maintained high contextual relevance (10/10), providing literature and analyses that shared characteristics with the specified locations, whether through direct geographical proximity (same state or region) or ecological similarity (comparable forest types, fire regimes, or climate patterns). For example, in responses to queries about Sangre de Cristo Mountains and Mount Bigelow, WildfireGPT identified studies conducted in comparable montane forest ecosystems with similar elevation profiles and fire history patterns. Perplexity AI similarly achieved high contextual relevance (10/10), although with variable levels of authority in source selection. The contextual relevance can be attributed to both WildfireGPT and Perplexity AI integrating a large set of sources to allow rich information to be retrieved, although we identify areas for improvement in the following sections. In contrast, ChatClimate demonstrated substantial contextual relevance limitations (4/10), frequently defaulting to continental or hemispheric climate trends without establishing clear applicability to the specified locations. This geographical generalization was particularly evident in responses to queries about Las Vegas, NM, and Covington, VA, where the system referenced broad North American climate patterns without demonstrating specific relevance to the distinctive fire regimes of the southwestern high desert or Appalachian forest ecosystems, respectively.

Overall, WildfireGPT demonstrated precise geographical analysis with methodological adaptation when confronting data limitations. ChatClimate exhibited geographical imprecision. Perplexity AI demonstrated comprehensive information retrieval with significant limitations in source validation and data interpretation accuracy. These performance patterns highlight fundamental distinctions in the analytical architecture of the three models. Furthermore, they demonstrate that having a

Table 3: Comparison of Wildfire analysis tools for evidence-based question answering. WGPT = WildfireGPT, CC = ChatClimate, PAI = Perplexity AI. Each row represents a specific prompt or task related to wildfire analysis. For each metric, a \checkmark indicates that the tool met the criteria for that prompt, while a \times indicates it did not. NA means the metric was not applicable or could not be evaluated, typically because of the previous criteria not being met; for example, when external knowledge is not retrieved, no citation would be provided. For ChatClimate, when it refers to studies outside of IPCC reports, it fails to provide the full title/link to the cited study. Contextual relevance means that the response is related to the location to the best specificity possible and tries to address the prompt. The main reason behind ChatClimate’s failures is that it often references general trends in large countries without a clear clue as to whether the information could be relevant/transferred to the location-specific queries in the prompt. The superscript numbers indicate the following issues with PAI: ¹Used articles from non-profit organizations, ²Included articles from news media, ³Used articles from commercial sources, ⁴Referenced a Reddit post.

Prompt	External Knowledge Retrieved			Citation Provision			Authoritative Sources			Contextual Relevance		
	WGPT	CC	PAI	WGPT	CC	PAI	WGPT	CC	PAI	WGPT	CC	PAI
Is FWI a robust estimate of the fire danger risk?	\times	\checkmark	\checkmark	NA	\checkmark	\checkmark	NA	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
What universities would make good partners to undertake studies on fire behavior and effective mitigation strategies in Mora County, NM?	\times	\times	\checkmark	NA	NA	\checkmark	NA	NA	\checkmark	\checkmark	\checkmark	\checkmark
What are the wildfire mitigation strategies, ecological impacts, and economic effects in forested areas similar to Mount Bigelow, AZ?	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\times^1	\checkmark	\times	\checkmark
Please provide sources on how to obtain data for Vegetation Type and Density, Previous Burn History, Human Activity Levels, Topography, Weather Conditions, Land Management Practices, Infrastructure and Accessibility, Socio-Economic Factors, and Climate Change Indicators	\times	\checkmark	\checkmark	NA	\checkmark	\checkmark	NA	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
It looks like March is the time folks implement a controlled burn in Naperville, IL. What do you think?	\checkmark	\times	\checkmark	\checkmark	NA	\checkmark	\checkmark	NA	\checkmark	\checkmark	\times	\checkmark
Let’s dive deeper into Firebreaks and Buffer Zones. I’m interested in learning about approaches to establish firebreaks around the critical infrastructure and vegetation management practices. Can you provide more detail with evidences?	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark
Search for scientific literature on effective post-wildfire recovery strategies, focusing on public safety and environmental restoration in Sangre De Cristo Mountains.	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\times^2	\checkmark	\times	\checkmark
Is fire common in this region prior to 1900s in Covington, VA? What is the natural fire regime?	\checkmark	\times	\checkmark	\checkmark	NA	\checkmark	\checkmark	NA	\checkmark	\checkmark	\times	\checkmark
Are there supporting literature for Buffer Zones that communities can implement to manage wildfire risk?	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\times^3	\checkmark	\checkmark	\checkmark
Are there good examples of similar cities to Las Vegas, NM that might experience similar shifts in wildfire risk? What mitigation strategies are those cities implementing?	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\times^4	\checkmark	\times	\checkmark
Overall Success Rate	7/10	7/10	10/10	7/7	3/7	10/10	7/7	7/7	6/10	10/10	4/10	10/10

wide but curated range of valid scientific resources with efficient retrieval mechanisms is critical for location-specific wildfire risk assessment and evidence-based domain-specific inquiry.

Conclusion

Natural-hazard-informed decision-making requires the integration of multiple location-specific data sources and scientific findings to capture its multifaceted nature. Unlike baseline alternatives such as ChatClimate and Perplexity AI, which demonstrated significant limitations in geographical precision and data accuracy, respectively, our multi-agent prototype, WildfireGPT addresses these challenges by integrating location-specific projections and observational data, socioeconomic indicators, and scientific literature from trusted sources to deliver an extensive analysis of wildfire risks. WildfireGPT represents a step forward in leveraging language models for specialized scientific domains and shows great potential in transforming user queries into actionable insights, by offering a customized

interactive experience through strategic questioning and planning, and combining scientific data and literature.

Acknowledgment

This material is based upon work supported by Laboratory Directed Research and Development funding from Argonne National Laboratory, provided by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DEAC02-06CH11357. Y. Xie and W. J. Su acknowledge support from the NSF HDR TRIPODS award (CCF-1934876). B. Jiang and C. J. Taylor acknowledge support from the NSF grant CCF-2112665 (TILOS). M. R. Alexander, J. Branham, and Y. Feng contributed significantly to the case studies. The authors also thank Jane S. Frantz, Parfait Gasana, Vijay P. Ramalingam, and Charles Doktycz for their valuable contributions to the case studies. The funder played no role in study design, data collection, analysis and interpretation of data, or the writing of this manuscript.

References

- [1] Hossein Tabari. Climate change impact on flood and extreme precipitation increases with water availability. *Scientific Reports*, 10(1):13768, 2020.
- [2] Camilo Mora, Tristan McKenzie, Isabella M Gaw, Jacqueline M Dean, Hannah von Hammerstein, Tabatha A Knudson, Renee O Setter, Charlotte Z Smith, Kira M Webster, Jonathan A Patz, et al. Over half of known human pathogenic diseases can be aggravated by climate change. *Nature Climate Change*, 12(9):869–875, 2022.
- [3] National Academies of Sciences, Division on Earth, Life Studies, Board on Atmospheric Sciences, Committee on Extreme Weather Events, and Climate Change Attribution. *Attribution of extreme weather events in the context of climate change*. National Academies Press, 2016.
- [4] Task Committee on Future Weather, Climate Extremes, Mari R Tye, and Jason P Giovannettone. Impacts of future weather and climate extremes on united states infrastructure: Assessing and prioritizing adaptation actions. American Society of Civil Engineers Reston, VA, 2021.
- [5] Bin Lu, Ze Zhao, Luyu Han, Xiaoying Gan, Yuntao Zhou, Lei Zhou, Luoyi Fu, Xinbing Wang, Chenghu Zhou, and Jing Zhang. OxyGenerator: reconstructing global ocean deoxygenation over a century with deep learning. In *Forty-first International Conference on Machine Learning*, 2024.
- [6] Daniel Adshead, Haluk Akay, Christophe Duwig, Elina Eriksson, Mattias Höjer, Karin Larsdotter, Åsa Svenfelt, Ricardo Vinuesa, and Francesco Fuso Nerini. A mission-driven approach for converting research into climate action. *npj Climate Action*, 2(1):13, 2023.
- [7] Max Callaghan, Carl-Friedrich Schleussner, Shruti Nath, Quentin Lejeune, Thomas R Knutson, Markus Reichstein, Gerrit Hansen, Emily Theokritoff, Marina Andrijevic, Robert J Brecha, et al. Machine-learning-based evidence and attribution mapping of 100,000 climate impact studies. *Nature Climate Change*, 11(11):966–972, 2021.
- [8] Nicolas Webersinke, Mathias Kraus, Julia Bingler, and Markus Leippold. ClimateBERT: A pretrained language model for climate-related text. Available at SSRN 4229146, 2022.
- [9] Matthias C Rillig, Marlene Ågerstrand, Mohan Bi, Kenneth A Gould, and Uli Sauerland. Risks and benefits of large language models for the environment. *Environmental Science & Technology*, 57(9):3464–3466, 2023.
- [10] Veruska Muccione, Saeid Ashraf Vaghefi, Julia Bingler, Simon K Allen, Mathias Kraus, Glen Gostlow, Tobias Wekhof, Chiara Colesanti-Senni, Dominik Stammbach, Jingwei Ni, et al. Integrating artificial intelligence with expert knowledge in global environmental assessments: opportunities, challenges and the way ahead. *Regional Environmental Change*, 24(3):1–8, 2024.
- [11] Angel Hsu, Mason Laney, Ji Zhang, Diego Manya, and Linda Farczadi. Evaluating ChatNetZero, an LLM-Chatbot to demystify climate pledges. In *Proceedings of the 1st Workshop on Natural Language Processing Meets Climate Change (ClimateNLP 2024)*, pages 82–92, 2024.

[12] Saeid Ashraf Vaghefi, Dominik Stammbach, Veruska Muccione, Julia Bingler, Jingwei Ni, Mathias Kraus, Simon Allen, Chiara Colesanti-Senni, Tobias Wekhof, Tobias Schimanski, et al. ChatClimate: Grounding conversational AI in climate science. *Communications Earth & Environment*, 4(1):480, 2023.

[13] Malte Toetzke, Benedict Probst, and Stefan Feuerriegel. Leveraging large language models to monitor climate technology innovation. *Environmental Research Letters*, 18(9):091004, 2023.

[14] Lauren Stagnol, Amina Cherief, Zakaria Farah, Théo Le Guenadal, Sofia Sakout, and Takaya Sekine. Answering clean tech questions with large language models. Available at SSRN 4663447, 2023.

[15] David Thulke, Yingbo Gao, Petrus Pelser, Rein Brune, Rricha Jalota, Floris Fok, Michael Ramos, Ian van Wyk, Abdallah Nasir, Hayden Goldstein, et al. ClimateGPT: Towards AI synthesizing interdisciplinary research on climate change. *arXiv preprint arXiv:2401.09646*, 2024.

[16] Mathias Kraus, Julia Anna Bingler, Markus Leippold, Tobias Schimanski, Chiara Colesanti Senni, Dominik Stammbach, Saeid Ashraf Vaghefi, and Nicolas Webersinke. Enhancing large language models with climate resources. *arXiv preprint arXiv:2304.00116*, 2023.

[17] Charles Cao, Jie Zhuang, and Qiang He. LLM-assisted modeling and simulations for public sector decision-making: Bridging climate data and policy insights. In *AAAI-2024 Workshop on Public Sector LLMs: Algorithmic and Sociotechnical Design*, 2024.

[18] Francesca Larosa and Adam Wickberg. Artificial intelligence can help loss and damage only if it is inclusive and accessible. *npj Climate Action*, 3(1):59, 2024.

[19] Perplexity AI. Perplexity AI: An answer engine for knowledge discovery, 2024. Available at <https://perplexity.ai>.

[20] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.

[21] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2023.

[22] Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. Agentic retrieval-augmented generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*, 2025.

[23] Tanwi Mallick, Joshua David Bergerson, Duane R Verner, John K Hutchison, Leslie-Anne Levy, and Prasanna Balaprakash. Understanding the impact of climate change on critical infrastructure through NLP analysis of scientific literature. *Sustainable and Resilient Infrastructure*, pages 1–18, 2024.

[24] Tanwi Mallick, John Murphy, Joshua David Bergerson, Duane R Verner, John K Hutchison, and Leslie-Anne Levy. Analyzing regional impacts of climate change using natural language processing techniques. *arXiv preprint arXiv:2401.06817*, 2024.

[25] Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, 2019.

[26] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.

[27] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[28] OpenAI. OpenAI Assistant API. <https://platform.openai.com/docs/guides/assistant>, 2023. Accessed: 2024-06-21.

[29] Streamlit. <https://www.streamlit.io/>. Accessed: 2024-01-21.

[30] National Wildfire Coordinating Group. Fire weather index (fwi) system, 2021. Accessed: 2024-09-30.

[31] Argonne National Laboratory. Climate risk & resilience portal (ClimRR), 2022. Accessed: 2024-01-21.

[32] National Interagency Fire Center (NIFC). Wildland fire incident locations. NIFC Open Data. Accessed 2024-01-21.

[33] Wildfire.gov. Integrated reporting of wildfire information (irwin). Wildfire.gov. Accessed 2024-01-21.

[34] E.Q. Margolis and C.H. Guiterman. NOAA/WDS Paleoclimatology - North American Tree-ring Fire Scar Synthesis (NAFSS). NOAA National Centers for Environmental Information, 05 2021. Accessed 2024-01-21.

[35] U.S. Census Bureau. American community survey 5-year data (2018-2022), 2023. Accessed: 2024-01-13.

[36] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.

[37] CY LIN. Rouge: A package for automatic evaluation of summaries. In *Proc. Workshop on Text Summariation Branches Out, Post-Conference Workshop of ACL 2004*, 2004.

[38] Zeyang Liu, Ke Zhou, and Max L Wilson. Meta-evaluation of conversational search evaluation metrics. *ACM Transactions on Information Systems (TOIS)*, 39(4):1–42, 2021.

[39] Jan Milan Deriu, Don Tuggener, Pius von Däniken, and Mark Cieliebak. Probing the robustness of trained metrics for conversational dialogue systems. In *60th Annual Meeting of the Association for Computational Linguistics (ACL 2022), Dublin, Ireland, 22-27 May 2022*, volume 2, pages 750–761. Association for Computational Linguistics, 2022.

[40] C. E. Van Wagner. Development and structure of the Canadian Forest Fire Weather Index System. Forestry Technical Report 35, Canadian Forestry Service, Headquarters, Ottawa, 1987.

[41] Steve Hvenegaard. Long-term monitoring programs and data-collection protocols for fuel treatment sites: a literature review. 2014.

[42] Lisa Dale. Wildfire policy and fire use on public lands in the United States. *Society and Natural Resources*, 19(3):275–284, 2006.

[43] Jens T Stevens, Gabrielle FS Boisramé, Ekaterina Rakhmatulina, Sally E Thompson, Brandon M Collins, and Scott L Stephens. Forest vegetation change and its impacts on soil water following 47 years of managed wildfire. *Ecosystems*, 23(8):1547–1565, 2020.

[44] Alan A Ager, Palaiologos Palaiologou, Cody R Evers, Michelle A Day, Chris Ringo, and Karen Short. Wildfire exposure to the wildland urban interface in the western US. *Applied Geography*, 111:102059, 2019.

[45] Tania Schoennagel, Jennifer K Balch, Hannah Brenkert-Smith, Philip E Dennison, Brian J Harvey, Meg A Krawchuk, Nathan Mietkiewicz, Penelope Morgan, Max A Moritz, Ray Rasker, et al. Adapt to more wildfire in western North American forests as climate changes. *Proceedings of the National Academy of Sciences*, 114(18):4582–4590, 2017.

[46] Amy EM Waltz, Michael T Stoddard, Elizabeth L Kalies, Judith D Springer, David W Huffman, and Andrew Sánchez Meador. Effectiveness of fuel reduction treatments: assessing metrics of forest resiliency and wildfire severity after the wallow fire, az. *Forest Ecology and Management*, 334:43–52, 2014.

- [47] Winslow D Hansen, Diane Abendroth, Werner Rammer, Rupert Seidl, and Monica G Turner. Can wildland fire management alter 21st-century subalpine fire and forests in grand teton national park, wyoming, usa? *Ecological Applications*, 30(2):e02030, 2020.
- [48] Reimers Nils. sentence embeddings using siamese BERT-networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language*, pages 3982–3992, 2020.
- [49] Amanda B Stan, Lesley S Rigg, and Linda S Jones. Dynamics of a managed oak woodland in northeastern Illinois. *Natural Areas Journal*, 26(2):187–197, 2006.
- [50] Max A Moritz and Scott L Stephens. Fire and sustainability: considerations for California’s altered future climate. *Climatic Change*, 87(Suppl 1):265–271, 2008.
- [51] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: NLG evaluation using GPT-4 with better human alignment. *arXiv preprint arXiv:2303.16634*, 2023.

A Extended Visualizations from WildfireGPT

In this appendix, we provide the full set of interactive visualization outputs introduced in Figure 2, including the remaining seasonal maps and socio-demographic overlays.

B Extended Case Studies Table

We provide the full set of case studies introduced in Table 1, including the remaining rows that were omitted from the main text for brevity.

C Overview of the WildfireGPT User Experience

Figure 4 provides an overview of the user experience as WildfireGPT transitions through the three agents, with screenshots from one of the case studies themed Comprehensive Wildfire Impact.

D Details of the WildfireGPT Multi-Agent Architecture

Task Orchestrator We designed the task orchestrator agent to serve as the central coordinator in the multi-agent system, ensuring seamless interaction between the user and the specialized agents for efficient task execution. The orchestrator includes key functions to enable its role through the OpenAI function calling API. The `get_response` function processes user inputs, invokes the appropriate agent based on the context, and manages outputs or follow-up actions to keep the interaction smooth and efficient. The `update_assistant` function allows the orchestrator to dynamically transition between agents, ensuring that the most relevant agent is assigned to handle the current task based on user input and workflow requirements. The `resume_conversation` function is designed to restore context in cases of interrupted or paused conversations by replicating the thread’s messages into a new thread, allowing for seamless continuation. Together, these functions are designed to empower the task orchestrator agent to manage workflows, prioritize user-centric engagement, and deliver a cohesive experience across various stages of decision-making and analysis.

User Profile Agent The user profile agent is designed to collect and refine user-specific inputs for wildfire risk assessment through a structured, interactive, and iterative approach. This agent helps WildfireGPT produce outputs tailored to the unique needs of each user in a later stage of the interaction. To achieve this, we employed prompting techniques to guide users through individualized questions presented one at a time, enabling a systematic completion of a checklist covering personalized information about the user. Once all questions are answered, the system compiles the responses into a structured checklist and presents a summary to the user for verification, ensuring both accuracy and alignment with the user’s expectations. The design also incorporates flexibility, allowing users to respond with “I don’t know” if they are uncertain about any question, accommodating varying

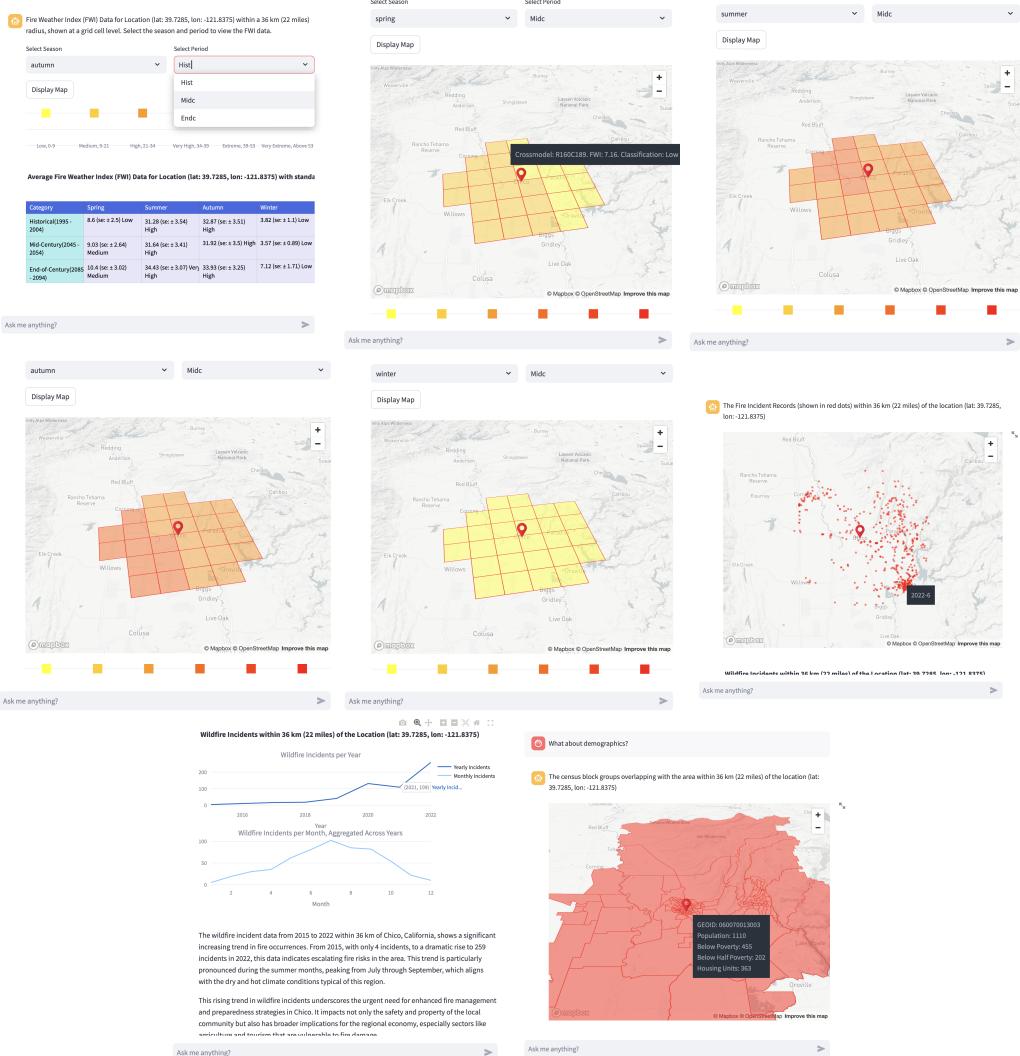


Figure 3: Interactive visualizations in the WildfireGPT user experience. This example is taken from the case study themed "Current Wildfire Risk Analysis." Users can select the season and time period to view the corresponding Fire Weather Index (FWI) map (left of first row), which displays risk levels using a color scale. Location-specific FWI values are accessible by hovering over the map (center of first row). By selecting different seasons and time period, the users can comprehend the changing landscape of the FWI (center of first row to center of second row). The wildfire incident map (right of the second row) shows the spatial distribution of recent fires, while the line graph (bottom left) presents the temporal trend of incidents. Socioeconomic data is visualized through census block group overlays (bottom right), providing insights into poverty rates and housing units in each area.

levels of expertise and familiarity with wildfire risk topics. The agent operates in two distinct stages, predefined inquiry and summary verification, working together to provide understanding of the user's background and concerns.

The user profile agent in WildfireGPT uses a structured dialogue designed to gather essential information to prepare for generating actionable wildfire risk insights later. It begins with five predefined questions that capture critical user details: profession, concern, location, time, and scope.

- **Profession:** The profession question collects the user's professional background and expertise. This information will act as a proxy for WildfireGPT to tailor its recommendations and

Table 4: Case studies conducted with domain experts to evaluate the performance and applicability of WildfireGPT across various professions, geographic locations, and wildfire management themes. Each case study represents either a location-specific inquiry derived from experts' ongoing research projects or a stakeholder perspective based on established regional collaborations. The description column outlines the specific focus of each case. The studies span different time horizons and incorporate a range of data types and analytical approaches to assess wildfire risk, impact, and mitigation strategies. Case studies are listed in alphabetical order.

Theme	Profession	Location	Description
Climate Impact on U.S. Wildfires	Atmospheric Scientist	Boston, MA	Exploring medium-term (10–30 years) impact of climate on wildfire occurrences and intensity in the U.S. Analyzing studies showing an increase in wildfire occurrences in Boston. Identifying and addressing uncertainties affecting future wildfire risks.
Comprehensive Wildfire Impact	Engineer	Mount Bigelow, AZ	Exploring long-term (50+ years) wildfire frequency, intensity, and impacts. Investigating mitigation strategies, ecological and economic effects, and infrastructure considerations.
Current Wildfire Risk Analysis	Data Analyst	Chico, CA	Analyzing current wildfire data, causes, and forest management practices. Exploring predictive models incorporating real-time variables. Gathering detailed demographic and socioeconomic profiles to assess risk and impact.
Ecosystem Fire Management	Landscape Manager	Naperville, IL	Analyzing wildfire frequency and impacts on vegetation with a focus on oak ecosystems. Developing strategies to balance fire use for oak ecosystem management with risk reduction for residential areas.
Hazard Mitigation Planning	Hazard Mitigation Planner	Mora County, NM	Assessing wildfire likelihood, frequency, intensity, and timing over the next 15 years for community hazard mitigation planning.
Infrastructure Wildfire Risk	Climate Change Risk Analyst	Denver, CO	Analyzing historical wildfire data and future projections to assess risks to energy sector infrastructure, considering land use changes and climate models (RCP 8.5 scenario).
Post-Wildfire Public Safety	Public Safety Manager	Sangre De Cristo Mountains, NM	Focusing on short-term (1–10 years) mitigation strategies after a high-intensity fire. Assessing emergency services readiness, community preparedness programs, and environmental restoration initiatives with the U.S. Army Corps of Engineers on erosion and flooding mitigations.
Private Property Protection	Homeowner	Near Covington, VA	Developing a 5–10-year plan for managing forest health, maximizing marketable species (focus on oak and cherry), and protecting properties from wildfires.
Urban Wildfire Mitigation	Urban Planner	Beaverton, OR	Developing a 30-year community plan focused on building codes and infrastructure resilience to mitigate fire-related hazard risks.
Wildland Urban Interface Impact	Risk Manager	Las Vegas, NM	Assessing medium-term (10–30 years) wildfire risks to housing and water resources in the wildland-urban interface. Exploring historical data, predictive models, current housing structures, future development plans, and wildfire impact on water quality and availability. Examining existing and proposed mitigation strategies.

technical depth to the user's knowledge level, whether engineer, urban planner, emergency manager, or infrastructure operator.

- **Concern:** The concern question identifies the user's key motivations or queries surrounding wildfire risks, such as community safety, infrastructure resilience, or ecological preservation. This information will be used to help the system target the relevant aspects of wildfire management to the user.
- **Location:** The location of interest question identifies the geographic area for wildfire risk assessment, allowing users to input locations as place names, geographic coordinates, or descriptive details. When a user specifies a location in natural language, the system is designed to leverage the LLM to convert the input into geographic coordinates (latitude and longitude). To ensure accuracy, we integrated a map-based verification system that enables

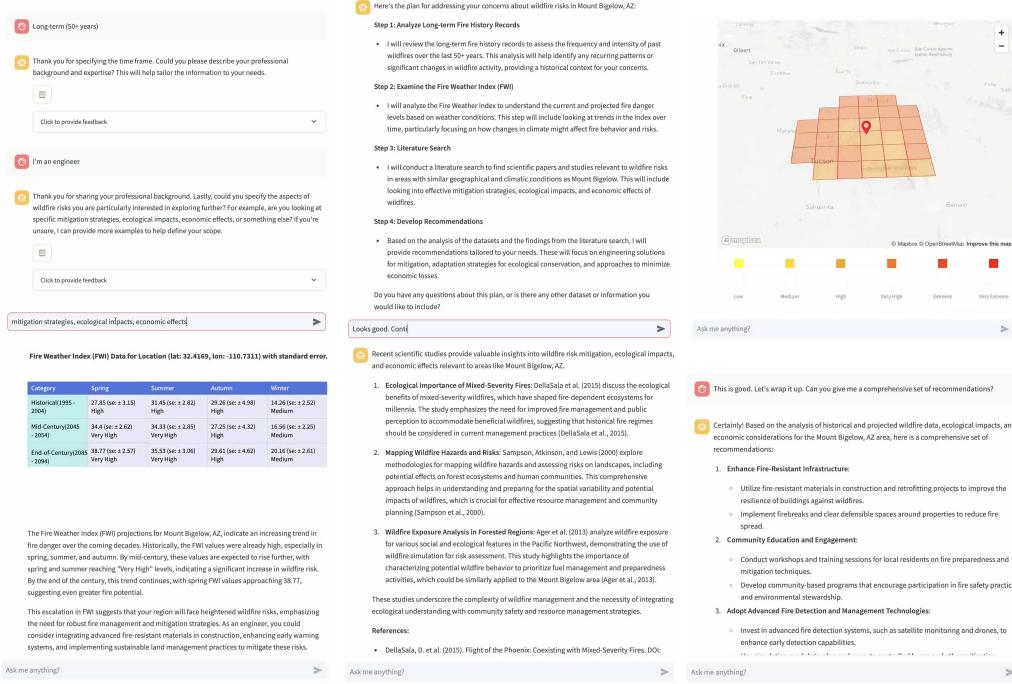


Figure 4: Overview of the WildfireGPT user experience. The screenshots are taken from one of the case studies themed Comprehensive Wildfire Impact. The user profile agent (top left) engages the user in a conversation to understand the user’s background and concerns. The planning agent (top middle) generates a tailored analysis plan based on the user’s profile. The analyst agent then executes the plan, analyzing Fire Weather Index data (top right and bottom left), conducting a literature review (bottom middle), and generating personalized recommendations (bottom right) to address the user’s wildfire risk concerns.

users to visually confirm or refine their input. The specified coordinate is then saved to help retrieve structured data and enable robust and geographically grounded wildfire risk assessments.

- **Timeframe:** The timeframe question defines the temporal scope of interest, offering options for short-term (1–10 years) mitigation strategies, medium-term (10–30 years) resilience planning, or long-term (30–80+ years) resilience adaptation. Historical analysis is also available, including recent (1–10 years), past (10–50 years), or long-term (50+ years) fire patterns.
- **Scope:** The scope of interest question refines the focus of the risk assessment, covering areas such as infrastructure vulnerability, emergency preparedness, ecological impacts, and insurance planning, with guided examples provided for users who are unsure. After gathering responses, WildfireGPT compiles the information into a structured checklist, which is shared with the user for review to ensure accuracy and alignment with their expectations.

Upon completing and confirming the finalized checklist, the agent saves the information and transitions to the planning agent, where WildfireGPT begins generating actionable insights and recommendations. This transition marks the shift from information gathering to solution generation, ensuring that the outputs are tailored to the user’s specific needs and the context of their wildfire risk assessment.

Planning Agent We designed the planning agent prompt in WildfireGPT to guide the system in creating a structured and user-centered wildfire risk assessment plan. By embedding step-by-step instructions into the prompt, we enable the planning agent to systematically address user concerns, integrate user feedback, and leverage relevant datasets effectively. The prompt for the planning agent comprises the following components:

- Step-by-Step Engagement Plan: The prompt instructs the system to create a short, clear plan to engage with the user effectively. This plan includes leveraging the most relevant datasets to address the user’s concerns through data analysis and a literature search. The system is guided to select datasets based on the user’s specific needs—such as analyzing trends in wildfire risk, understanding historical patterns, or focusing on immediate mitigation strategies.
- Feedback Integration: Once the plan is drafted, the system is prompted to pause and share the plan with the user. This ensures that the user is fully informed about the proposed approach and has the opportunity to provide feedback or request modifications. The prompt explicitly instructs the system to ask whether the user would like to include additional information or datasets, while also clarifying that the wildfire analysis is limited to the three available datasets: FWI, long-term fire history records, and recent fire incident data. This step ensures transparency and active collaboration between the system and the user.
- Finalization and Execution: After incorporating user feedback and securing the user’s agreement, the system transitions to finalizing the plan. The prompt ends with a specific instruction to call the function `plan_complete()` with the completed plan, signaling the readiness to move to the execution phase.

The planning prompt also leverages one-shot learning by including an example plan that act as a demonstration, guiding the model to understand the desired structure and format of the output. In addition to the example, the prompt includes detailed dataset descriptions, outlining the characteristics and scope of the three available wildfire data sources to constrain the model’s data selection. The prompt also integrates the user profile for the agent to make informed decisions.

During the user interaction, WildfireGPT shares brief descriptions of the available datasets, before presenting the proposed plan. From the user’s perspective, this plan comprises three key steps: (1) data retrieval: WildfireGPT identifies the most relevant wildfire datasets from the three available sources and explains the reasoning behind the selection; (2) literature review: WildfireGPT outlines the thematic focus of the literature to be examined; and (3) recommendation development: based on the data analysis and literature review, WildfireGPT develops personalized recommendations.

Analyst Agent The analyst agent leverages the RAG framework [20, 21] to conduct analyses by integrating diverse data sources and domain knowledge. This agent maintains the summarized user profile and action plan within its prompt context to ensure that WildfireGPT adheres to the plan and addresses user follow-ups, offering to explore specific topics or proceed to the next step. The analyst agent functions as an information processor, ensuring that outputs are grounded in scientific research articles and wildfire data. Its structured approach enables it to process volumes of information efficiently while maintaining accuracy and relevance. The analysis consists of three main stages:

- **Wildfire data retrieval and visualization:** The data retrieval and visualization functionality in our system is designed to assess wildfire risk by leveraging geospatial data and fire trends. The process begins with location-based data retrieval, based on the location information collected by the user profile agent. For all datasets, data within a 36-kilometer radius of the specified coordinate is retrieved and presented. The results are visualized through interactive maps, tables, and statistical plots. The map-based visualization employs GeoJsonLayer rendering to display the data points. A pin-layer visualization is also integrated to mark the exact user-specified location. To facilitate decision-making, the system dynamically generates a structured report summarizing the statistics.

Recent Wildfire Incident Record Data: Historical wildfire locations (2015–2023) are visualized as individual red markers on the GeoJsonLayer, with each incident precisely geolocated within a 36-kilometer analysis radius. In addition, two complementary line plots are presented for temporal analysis: an annual frequency plot revealing incident patterns over the years and a monthly distribution highlighting seasonal variations. The report contextualizes total incident frequency per year and aggregated monthly distributions into risk assessments, analyzing temporal clustering patterns to identify periods of elevated fire activity.

Tree-Ring and Sediment-Based Fire History Record Data: The system retrieves and ranks the three nearest paleofire study locations using a geodesic distance algorithm, integrating site-specific metadata from the International Multiproxy Paleofire Database. Each identified site is presented as an individual red marker on the GeoJsonLayer. Additionally, for each

site, the system provides the user with the research metadata including site names, precise coordinates, and associated publication records.

Fire Weather Index Data: Unlike the other data sources, the FWI data is structured on a grid system and referenced by using Crossmodel indices. Thus, we developed a process to map the user-specified geographic coordinates to the corresponding Crossmodel reference. This process begins by transforming the latitude and longitude into a spatial reference system compatible with the grid model. The geographic coordinates, initially expressed in degrees, are converted into the coordinate reference system used by the FWI data, typically employing projected coordinates for spatial accuracy. Using spatial operations such as buffering and intersecting, the transformed geographic point is compared against the grid cells in the model. The system retrieves the associated Crossmodel indices if they fall within a radius of 36 kilometers of the specified coordinate. This mapping establishes a direct connection between the user-defined location and the corresponding grid cell in the FWI data. By designing this mapping mechanism, we ensured that WildfireGPT can seamlessly integrate user-specified locations with grid-based FWI data, providing precise and scientifically accurate insights tailored to the specified area. The retrieval process integrates historical (1995–2004), mid-century (2045–2054), and end-of-century (2085–2094) projections to offer a temporal perspective on fire weather trends. The retrieved FWI values are aggregated and analyzed to compute mean wildfire indices and standard deviations across seasons, ensuring a statistically robust representation of fire risk. The system then categorizes FWI values into six risk classes (low, medium, high, very high, extreme, and very extreme) based on the Canadian Forest Fire Weather Index classification system [40]. The GeoJsonLayer displays fire weather indices across spatial grids, with each grid cell color-coded according to its FWI classification. Additionally, the module provides a table presenting the seasonal comparisons of historical, mid-century, and end-of-century projections, enabling users to track long-term fire risk evolution. Lastly, the module is prompted to provide a structured summary report detailing past and projected wildfire risks in the region.

Census Data for Socioeconomic Analysis: If location-specific wildfire data is retrieved, users are informed of an additional analysis capability: demographic and socioeconomic impact assessments from census datasets. This can help provide an understanding of wildfire risks beyond environmental factors, by incorporating population vulnerabilities, economic conditions, and housing density. If wildfire data is not available, the system shares a prewritten cautionary message, emphasizing the preliminary nature of its recommendations and encouraging users to seek further investigation or expert advice before implementing significant changes.

The census data retrieval functionality extracts demographic and socioeconomic statistics for regions surrounding a user-specified location, with each census block group data retrieved from the American Community Survey (ACS5) dataset. Key indicators such as total population, poverty distribution, and housing unit counts are aggregated to generate a demographic summary, highlighting the number of individuals below the poverty line and those with income less than half the poverty threshold. Each census block group is visualized on the GeoJsonLayer for interactive spatial understanding. Additionally, a structured table presents aggregated population and housing statistics, allowing users to explore socioeconomic factors in detail.

By integrating real-time geospatial data analysis, interactive visualization, and RAG techniques, WildfireGPT ensures that wildfire risk assessments are data-driven, context-aware, and easily interpretable for end users.

- **Literature retrieval:** For scientific literature, WildfireGPT queries a corpus developed by Argonne National Laboratory for the CIACC tool [23, 24]. All abstracts of the papers in this corpus are pre-embedded using the `all-MiniLM-L6-v2` embedder from the Sentence-Transformers library [25] and stored in a Faiss vector store [26]. The literature search query is processed by the same embedding model and undergoes a k-nearest neighbor search, identifying the top k-most relevant abstract based on cosine similarity between the query vector and stored paper embeddings in the FAISS vector database. To enhance the reliability of the retrieved results, the system attempts to validate DOIs (Digital Object Identifiers) by cross-referencing them with CrossRef metadata. This validation process includes retrieving DOIs based on the paper title, checking for title and author consistency, and computing a similarity score to confirm accuracy. If discrepancies such as title mismatches or incorrect

author attributions are detected, the DOI is discarded to maintain data integrity and prevent misinformation. The final output is structured as a ranked list of the three most relevant papers, displaying each paper’s title, authors, publication year, and abstract, with a direct DOI link if verified. This integration of FAISS-based vector search, semantic similarity ranking, and DOI validation ensures the provision of high-quality, research-backed insights for scientists, policymakers, and decision-makers in wildfire risk assessment and mitigation planning.

- **Generation of tailored recommendations:** To generate tailored recommendations, the system constructs an augmented prompt that incorporates context from the ongoing conversation to provide actionable recommendations. If additional clarification or exploration is required, the system iterates the retrieval and response process, continuously refining the output to address the user’s evolving needs. In this way, the system contextualizes data retrieval and recommendation generation based on the user’s profession, ensuring that risk assessments, mitigation strategies, and policy suggestions align with their specific safety priorities and operational responsibilities.

Case Study	Prompt
Climate Impact on U.S. Wildfire Occurrences and Intensity	Analyze future climate data to understand potential trends in fire danger in Boston, Massachusetts, USA
Comprehensive Wildfire Impact	Review the long-term fire history records to assess the frequency and intensity of past wildfires over the last 50+ years in Mount Bigelow, AZ
Current Wildfire Risk Analysis	Analyze demographic and socio-economic profiles of the Chico area in California to understand which groups are most vulnerable in wildfire scenarios
Ecosystem Fire Management	Analyze the Long Term Fire History Records to assess wildfire events over the past 150 years in the Naperville region
Hazard Mitigation Planning	Analyze the recent fire incident data for Mora County, reviewing wildfire frequency, intensity, and timing from 2015 to 2023
Infrastructure Wildfire Risk	Examine climate models and projections under the RCP 8.5 scenario to forecast potential future wildfire risks in the medium-term (10–30 years) in Denver, Colorado
Post-Wildfire Public Safety and Environmental Restoration	Analyze the recent fire incident data from the Wildland Fire Interagency Geospatial Services Group to assess the frequency, intensity, and locations of recent wildfires in Sangre de Cristo Mountains
Private Property Protection	Analyze the recent fire incident data from the last five years in Covington, VA
Urban Wildfire Mitigation	Analyze the FWI focusing on both current trends and projections for the mid-century (2045–2054) to assess the potential increase in fire danger due to climate change in Beaverton, Oregon
Wildland Urban Interface Impact on Housing and Water Resources	Analyze projected wildfire risk data for the mid-century period (2045–2054) to understand the potential trends in fire danger in Las Vegas, NM

Table 5: Prompts Used in for Modular Comparison in Data Retrieval

E Comparison against Baseline Models

To evaluate the effectiveness of data retrieval and evidence-based query response capabilities, we conducted systematic comparisons between WildfireGPT and two baseline models: ChatClimate and Perplexity AI. Our experimental framework comprised two distinct analytical modules drawn from our case studies (Tables 5 and 6): (i) data retrieval and analysis tasks, which examined each model’s capacity to process and interpret location-specific wildfire data across diverse geographical and temporal scales, and (ii) evidence-based query responses, which assessed the models’ ability to provide evidence-based answers to domain-specific questions. The evaluation metrics for data analysis encompass data provision (successful retrieval of relevant data), location specificity (geographical precision of retrieved data), and data accuracy (alignment between retrieved data and prompt

Case Study	Prompt
Climate Impact on U.S. Wildfire Occurrences and Intensity	Is FWI a robust estimate of the fire danger risk?
Comprehensive Wildfire Impact	What are the wildfire mitigation strategies, ecological impacts, and economic effects in forested areas similar to Mount Bigelow, AZ?
Current Wildfire Risk Analysis	Please provide sources on how to obtain data for Vegetation Type and Density, Previous Burn History, Human Activity Levels, Topography, Weather Conditions, Land Management Practices, Infrastructure and Accessibility, Socio-Economic Factors, and Climate Change Indicators
Ecosystem Fire Management	It looks like March is the time folks implement a controlled burn in Naperville, IL. What do you think?
Hazard Mitigation Planning	What universities would make good partners to undertake studies on fire behavior and effective mitigation strategies in Mora County, NM?
Infrastructure Wildfire Risk	Let's dive deeper into Firebreaks and Buffer Zones. I'm interested in learning about approaches to establish firebreaks around the critical infrastructure and vegetation management practices. Can you provide more detail with evidences?
Post-Wildfire Public Safety and Environmental Restoration	Search for scientific literature on effective post-wildfire recovery strategies, focusing on public safety and environmental restoration in Sangre De Cristo Mountains.
Private Property Protection	Is fire common in this region prior to 1900s in Covington, VA? What is the natural fire regime?
Urban Wildfire Mitigation	Are there supporting literature for Buffer Zones that communities can implement to manage wildfire risk?
Wildland Urban Interface Impact on Housing and Water Resources	Are there good examples of similar cities to Las Vegas, NM that might experience similar shifts in wildfire risk? What mitigation strategies are those cities implementing?

Table 6: Prompts Used in for Modular Comparison in Question and Answer

requirements). Similarly, question-answering response assessment examined external knowledge integration, source authority, citation practices, and contextual relevance. Below is an example of the responses provided in the by each model. The full set of responses can be accessed in the provided GitHub repository.

E.1 Examples

Case Study: Wildland Urban Interface Impact on Housing and Water Resources

Prompt: Analyze projected wildfire risk data for the Mid-Century period (2045 - 2054) to understand the potential trends in fire danger in Las Vegas, NM.

ChatClimate: Based on the information provided, the projected wildfire risk data for the Mid-Century period (2045 - 2054) suggests that there will be an increase in fire danger in Las Vegas, NM. Climate change is expected to drive future increases in North American fire weather, particularly in the southwest, with a significant increase in the Fire Weather Index (FWI) before 2050 under the RCP8.5 scenario. This increase includes the frequency of 95th-percentile FWI days, peak seasonal FWI average, fire weather season length, and maximum fire weather index. The fire season across North America is projected to expand dramatically beyond 2°C global warming levels. Additionally, there will be more conducive conditions for lightning-ignited fires in the southeastern USA by mid-century, and warming conditions in Alaska will push temperatures above a threshold for fire danger across Alaska's tundra and boreal forest. Therefore, the trends indicate a heightened risk of wildfires in Las Vegas, NM during the Mid-Century period due to the projected increases in fire danger driven by climate change. (Page: 66, Reference: IPCC_AR6_WGI_Chapter12)

Perplexity: Las Vegas, NM is likely to face increased wildfire risk by mid-century (2045-2054) due to several factors:

1. Increased potential for very large fires: The risk of very large fires (top 10% of fires) could increase up to sixfold in parts of the West, including New Mexico, by mid-century (2041-2070) compared to the recent past (1971-2000)[3].
2. Climate change impacts: Warming temperatures, more frequent heat waves, and diminished soil moisture during the dry season are expected to contribute to increased fire danger[3]. These conditions, combined with historic fire suppression practices, may lead to more frequent large fire events.
3. Extended fire seasons: Climate change is projected to both intensify fire-friendly weather conditions and lengthen the season during which very large fires tend to spread[3].
4. Drought and fuel conditions: The combination of warming temperatures and below-normal precipitation in the Southwest may lead to drier fuels and increased fire potential[1].
5. Watershed vulnerability: Las Vegas, NM relies on a single surface water source (the Gallinas River watershed), making it particularly vulnerable to wildfire impacts on its water supply[2]. The 2022 Hermit's Peak/Calf Canyon Fire demonstrated the severe consequences of large-scale fires on the city's water resources.

While specific data for Las Vegas, NM is not provided, the overall trend for the western United States suggests a significant increase in fire danger. To mitigate these risks, it will be crucial for Las Vegas to implement preventive measures, such as watershed management, infrastructure upgrades, and improved emergency response planning.

Citations :

1. https://www.nifc.gov/nicc-files/predictive/outlooks/monthly_seasonal_outlook.pdf
2. <https://www.circleofblue.org/2023/world/new-mexicos-largest-fire-wrecked-this-citys-water-source/>
3. <https://www.climate.gov/news-features/featured-images/risk-very-large-fires-could-increase-sixfold-mid-century-us>
4. https://www.nifc.gov/nicc-files/predictive/outlooks/NA_Outlook.pdf
5. https://srs.fs.usda.gov/pubs/gtr/gtr_srs243.pdf
6. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9076366/>

Perplexity (when provided with the data file for FWI): Chose crossmodels "R382C098", "R438C055", "R395C128", "R411C070", "R418C086" in the dataset. However, the correct cross-models are around "R110C303."

WildfireGPT: The Fire Weather Index (FWI) data for Las Vegas, NM, reveals a concerning trend in fire danger levels over the next few decades. Historically, from 1995 to 2004, the FWI values indicated a medium risk during spring, summer, and autumn, with lower risks in the winter. Specifically:

- **Spring:** 13.1 (Medium)
- **Summer:** 17.04 (Medium)
- **Autumn:** 19.31 (Medium)
- **Winter:** 6.98 (Low)

Looking ahead to the mid-century period (2045–2054), the projections show a slight increase in the FWI values, especially in the spring:

- **Spring:** 17.31 (Medium)
- **Summer:** 18.26 (Medium)
- **Autumn:** 16.25 (Medium)
- **Winter:** 8.49 (Low)

By the end of the century (2085–2094), the spring season is expected to experience a significant increase in fire danger:

- **Spring:** 23.82 (High)
- **Summer:** 20.43 (Medium)
- **Autumn:** 20.5 (Medium)
- **Winter:** 11.52 (Medium)

This trend suggests an increasing risk of wildfires, particularly in the spring, which will necessitate enhanced fire management and urban planning strategies to mitigate risks to housing and water resources. The higher FWI values indicate drier and potentially windier conditions, which can facilitate the rapid spread of wildfires.

Given these projections, it is crucial to integrate robust wildfire risk management strategies into urban development plans for Las Vegas, NM. These should include considerations for fire-resistant building materials, strategic water resource management, and emergency response planning.

[Data Source: Climate Risk & Resilience Portal (ClimRR)](<https://disgeoportal.egs.anl.gov/ClimRR/>)

F Evaluating Personalization

To investigate the effects of the user profile agent, we conduct a qualitative, two-phase ablation study to systematically evaluate the impact of user profile granularity and diversity on WildfireGPT’s recommendations. Inspired by the Private Property Protection case study, for both phases we use the prompt “Develop recommendations to enhance forest resilience against wildfires” to elicit a response from WildfireGPT. In the first phase, we systematically test three levels of user profile specificity provided to the agent: no profile information provided, location plus timeline, and the full profile of a power grid manager focused on maintaining transmission line clearance and grid resilience for power distribution reliability and access. In the second phase, we hold location and timeline constant and test five distinct professional profiles—homeowner, civil engineer, ecologist, emergency manager, and power grid manager—each with minimally distinct concerns and the respective scopes. Table 7 summarizes the differences in primary concerns and scopes across these professional profiles, with more detailed profile descriptions available in the Methodology section. We conduct a controlled simulation for each profile by passing the profile directly into WildfireGPT’s planning and analyst agents, allowing us to observe how the same prompt generated tailored recommendations across different user types. Outputs are analyzed across three dimensions: *plan generation*, *literature retrieval*, and *recommendation specificity*.

Evaluation Methodology We carry out a two-phase ablation study to evaluate the impact of user profile granularity on WildfireGPT’s responses. In the first phase, we use the prompt “Develop recommendations to enhance forest resilience against wildfires” to elicit recommendations from WildfireGPT, with three levels of user profile details:

1. **Generic:** No contextual details.
2. **Location + Timeline:** Added geospatial (Covington, VA; 37.7935°N, 79.9939°W) and temporal (5–10 year implementation window) parameters.
3. **Full Profile:** Integrated profession-specific attributes (e.g., *power grid manager*), operational concerns (e.g., grid resilience, transmission line clearance), and scope (e.g., infrastructure protection).

In the second phase, we hold location and timeline constant and test five distinct professional profiles (*homeowner*, *civil engineer*, *ecologist*, *emergency manager*, *power grid manager*). For each profile, we modify three variables—profession, concern, and scope—in the template user profile shown in Figure 5.

Table 10 provides the full details on the comparative analysis across three profile specificity tiers. The full experiment results can be accessed in the provided GitHub repository.

Below we enumerate the plans generated by WildfireGPT for each different user profile in the second phase of our ablation study.

Table 7: User profile variations and literature search queries in Phase 2 of the WildfireGPT personalization ablation study. We hold location and timeline constant and test five distinct professional profiles, with changes in the concerns and the respective scopes. The search query represents the resulting literature retrieval query automatically generated by WildfireGPT’s analyst agent.

Profession	Primary Concern	Scope	Search Query
Homeowner	Maximizing marketable species	Health and marketable species	“Strategies for managing forests to maintain health, maximize marketable species, and minimize wildfire risks in Virginia”
Civil Engineer	Ensuring structural and infrastructural resilience	Drainage efficiency and slope stability	“Wildfire risks and climate change impacts on forest management near Covington, VA; Strategies for enhancing drainage efficiency and slope stability; Structural resilience against wildfires in forested areas”
Ecologist	Maintaining biodiversity and ecosystem services	Ecological resilience and habitat connectivity	“Wildfire management and ecological resilience in forest ecosystems near Covington, VA”
Emergency Manager	Establishing defendable space and evacuation corridors	Emergency access and response capabilities	“Effective forest management practices, defendable space creation, evacuation protocols, and property protection measures against wildfires near Covington, VA”
Power Grid Manager	Maintaining transmission line clearance and grid resilience	Power distribution reliability and access	“Effective strategies for vegetation management, forest health maintenance, and wildfire risk mitigation around power grids near Covington, VA”

Table 8: Comparative analysis across three profile specificity tiers. WildfireGPT responses demonstrate progressive domain adaptation from generic principles to stakeholder-specific protocols when provided with increasing user profile granularity. Representative outputs shown for vegetation management, technology implementation, and role-specific interventions demonstrate systematic enhancement of recommendation precision with profile enrichment. Italicized text highlights power-grid-specific recommendations emergent only with comprehensive user profiling.

Aspect	No Profile Information	+ Location and Timeline	+ Power-Grid Manager Profile
Vegetation Management	“Thinning and Pruning: Reduce forest density through thinning and pruning [...] Controlled Burns: Implement controlled burns to reduce available fuels.”	“Develop and promote fire-adapted ecosystems by incorporating native species [...] Practice selective cutting to reduce fire risk.”	“Establish routine schedules for pruning and removing trees that <i>pose a risk to power lines</i> . Utilize drones or satellite imagery to monitor vegetation growth for timely intervention.”
Technology	“Employ technologies like satellite imaging and drones [...] Use artificial intelligence to predict fire patterns.”	“Invest in enhanced surveillance systems [...] satellite imagery, drones, and remote sensors to detect early signs of fire.”	“Utilize Geographic Information Systems (GIS) and remote sensing [...] Install sensors and advanced monitoring systems along the grid to detect temperature changes, smoke, or increased humidity levels indicating the risk of fire. Use drones for real-time surveillance [...] of power lines, <i>especially in hard-to-reach areas</i> .”
Role-Specific Actions	“International cooperation [...] Reduce carbon footprints.”	“Encourage Responsible Logging Practices [...] Climate Adaptation Strategies.”	“ <i>Expand the Right-of-Way (ROW) clearance around transmission lines</i> beyond standard regulations [...] Implement an IVM program that combines mechanical, biological, and chemical methods to manage vegetation”

Emergency Manager Step 1: I will analyze the Fire Weather Index (FWI) dataset to assess the trends in fire danger over time, focusing particularly on future projections around your area. This will

Table 9: Distribution of Recommendation Categories Across Five Stakeholder Profiles. We use checkmarks to indicate which recommendations appear in each stakeholder's profile.

Recommendation Category	Emergency Manager	Ecologist	Civil Engineer	Homeowner	Power Grid Manager
<i>Fuel Management</i>					
Regular thinning/pruning/fuel reduction	✓	✓	✓	✓	✓
Controlled/prescribed burns	✓	✓	✓	✓	✓
<i>Infrastructure and Planning</i>					
Defendable space creation	✓			✓	✓
Fire-resistant materials/infrastructure	✓	✓		✓	✓
Emergency access/water sources	✓		✓	✓	✓
<i>Community Engagement</i>					
Evacuation planning	✓	✓		✓	
Public education/awareness programs	✓	✓	✓	✓	✓
<i>Collaboration and Management</i>					
Cross-boundary coordination	✓	✓			✓
Monitoring systems/technology	✓	✓	✓	✓	✓
Research	✓	✓			✓
<i>Ecological Considerations</i>					
Fire-resistant species promotion			✓	✓	
Natural fire regime maintenance		✓			
Ecological monitoring		✓	✓		✓

Figure 5: Template user profile used in user profile agent.

1. Profession: {profession} in Virginia.
2. Concern: Managing the forest, keeping it healthy, while {concern}, and protecting properties from potential wildfires.
3. Location: Near Covington, VA with Latitude 37.7935 and Longitude -79.9939.
4. Time: Recommendations to be implemented within the next 5 to 10 years.
5. Scope: Management of the forest and properties to maximize {scope}, and protect against potential wildfires.

help understand the potential increase in fire risks facing the Covington area in upcoming years, aiding in strategic planning for emergency responses and forest management.

Step 2: I will delve into the Long Term Fire History Records to extract data relevant to the historical wildfire activity around Covington, VA. Analyzing these trends will provide a contextual understanding of the periodicity and severity of wildfires, which is crucial for planning long-term forest management and fire prevention strategies.

Step 3: As your timeframe concerns the next 5 to 10 years, it is critical to also analyze Recent Fire Incident Data. This will highlight the most immediate risks and recent trends in wildfire occurrences. Insights from this will be pivotal in adjusting current practices and preparing immediate response strategies.

Step 4: I will conduct a comprehensive literature search focusing on similar geographical areas with comparable forest compositions and climate to Covington, VA. This research will look into effective forest management practices, successful defendable space creations, evacuation protocols, and property protection measures against wildfires.

Table 10: Comparative analysis across three profile specificity tiers. WildfireGPT responses demonstrate progressive domain adaptation from generic principles to stakeholder-specific protocols when provided with increasing user profile granularity. Representative outputs shown for vegetation management, technology implementation, and role-specific interventions demonstrate systematic enhancement of recommendation precision with profile enrichment. *Italicized* text highlights power-grid-specific recommendations emergent only with comprehensive user profiling.

Aspect	No Profile Information	+ Location and Timeline	+ Power-Grid Manager Profile
Vegetation Management	“Thinning and Pruning: Reduce forest density through thinning and pruning [...] Controlled Burns: Implement controlled burns to reduce available fuels.”	“Develop and promote fire-adapted ecosystems by incorporating native species [...] Practice selective cutting to reduce fire risk.”	“Establish routine schedules for pruning and removing trees that <i>pose a risk to power lines</i> . Utilize drones or satellite imagery to monitor vegetation growth for timely intervention.”
Infrastructure	“Enforce fire-resistant materials in construction [...] Maintain and expand firebreaks.”	“Build and maintain natural and artificial firebreaks such as roads [...] or barriers devoid of vegetation.”	“Retrofit existing infrastructure with fire-resistant materials [...] <i>replace wooden poles with those made from steel</i> or reinforced concrete to enhance fire resilience.”
Technology	“Employ technologies like satellite imaging and drones [...] Use artificial intelligence to predict fire patterns.”	“Invest in enhanced surveillance systems [...] satellite imagery, drones, and remote sensors to detect early signs of fire.”	“Utilize Geographic Information Systems (GIS) and remote sensing [...] Install sensors and advanced monitoring systems along the grid to detect temperature changes, smoke, or increased humidity levels indicating the risk of fire. Use drones for real-time surveillance [...] of <i>power lines, especially in hard-to-reach areas</i> .”
Role-Specific Actions	“International cooperation [...] Reduce carbon footprints.”	“Encourage Responsible Logging Practices [...] Climate Adaptation Strategies.”	“ <i>Expand the Right-of-Way (ROW) clearance around transmission lines</i> beyond standard regulations [...] Implement an IVM program that combines mechanical, biological, and chemical methods to manage vegetation”
Collaboration	“Foster coordination among forestry services, firefighting agencies, and wildlife organizations [...] International Cooperation: Exchange knowledge [...] with international bodies.”	“Forge partnerships with local universities, environmental organizations, and governmental agencies.”	“Collaborate with local government agencies, environmental organizations, <i>and other utilities</i> to share resources, information, and strategies for effective forest and wildfire management.”
Policy & Funding	“Enforce building codes [...] Integrate fire risk into land use planning decisions.”	“Seek federal and state grants aimed at forest conservation [...] fire management.”	“Advocate for local and state policies that support sustainable forest management [...] Seek state and federal grants for implementing advanced technologies [...] in infrastructure resilience and environmental management.”
Community Education	“Train local firefighters and communities in fire suppression techniques and evacuation procedures.”	“Develop community awareness and educational programs [...] Workshops, training sessions, and school curricula can focus on how individuals can reduce wildfire risks through proper land management.”	“Regularly communicate with local communities about maintenance schedules, controlled burns, and emergency procedures to enhance public awareness and cooperation.”

Step 5: Based on data analysis and findings from scientific literature, I will provide tailored recommendations. These will address creating or refining defendable spaces, designing evacuation routes, improving property resilience against wildfires, and enhancing overall emergency preparedness and forest health.

Ecologist Step 1: I will begin by analyzing the Fire Weather Index (FWI) dataset. This will focus on understanding the future trends in fire danger, particularly for the Mid-Century period (2045 - 2054). Analyzing these future projections, and comparing them against the historical period (1995 - 2004), will help us anticipate changes in wildfire risk near your area in Virginia.

Step 2: Next, I will examine the Long term fire history records. This dataset will provide vital insights into the historical patterns and cycles of wildfire activity in the broader North American region over centuries. Understanding these long-term trends will help in assessing the natural fire regimes and their impacts on ecosystems, which is crucial for managing forest landscapes effectively.

Step 3: Concurrently, I will conduct a literature search focused on wildfire management and ecological resilience in forest ecosystems. This will include finding scientific studies that discuss successful strategies for vegetation management, biodiversity preservation, and wildfire prevention in areas with similar ecological and climatic conditions to Covington, VA.

Step 4: Based on the insights obtained from the data analyses and scientific literature, I will develop a set of tailored recommendations. These will be aimed at enhancing the ecological resilience of your forest, improving habitat connectivity, and implementing effective wildfire risk mitigation strategies to protect properties within the next 5 to 10 years.

This plan will help ensure that your forest management strategy aligns well with the evolving dynamics of wildfire risks and ecological demands.

Civil Engineer Step 1: Analyze the Fire Weather Index (FWI) dataset to measure future fire danger trends, focusing on the Mid-Century period (2045 - 2054). This data will give insight into how fire risk is likely to change and escalate, allowing for forward-thinking preparations.

Step 2: Examine the Long Term Fire History records to understand centuries-long patterns and cycles in wildfire activity. This historical data analysis will contextualize the specific wildfire risks in your region and assist in predicting future trends more accurately.

Step 3: Conduct a thorough literature search of scientific papers focused on wildfire risks and climate change. This search will particularly emphasize studies related to forest management practices, resilience of structures against wildfires, and comparison of geographically similar areas to Covington, VA.

Step 4: Leverage insights from the aforementioned data analyses and literature search to formulate strategic recommendations. These will focus on enhancing forest and property management to improve drainage efficiency, slope stability, and overall resilience against future wildfire threats.

This strategic approach will ensure that all aspects of wildfire risks are comprehensively addressed, aligning with your needs to protect and sustainably manage the forest area around Covington, VA.

Homeowner Step 1: I will analyze the Fire Weather Index (FWI) dataset to assess the potential trends in fire danger around your location near Covington, VA, for both the Mid-Century (2045 - 2054) and the historical period (1995 - 2004). This will help us understand how the fire danger is expected to evolve over time, which is crucial for long-term forestry and property management planning.

Step 2: I will explore the Long term fire history records to extract and analyze historical wildfire activity patterns in your region. This analysis will focus on identifying any recurring wildfire events and the conditions associated with high fire risks. This information will be vital in predicting future wildfire tendencies and planning effective management practices for your forest.

Step 3: I will conduct a literature search to find scientific studies and successful strategies for managing forests to maintain health, maximize marketable species, and minimize wildfire risks. This will include examining case studies from areas with similar geographical and climatic conditions as near Covington, VA.

Step 4: Based on the compiled data and insights from the literature, I will provide recommendations tailored to your specific concerns and goals. These will include actionable strategies for forest management, species selection and cultivation practices, alongside measures for

enhancing wildfire resistance and property protection. All recommendations will aim for implementation within the next 5 to 10 years to align with your timeline.

Power Grid Manager

- Step 1: Analyze the Recent fire incident data to understand recent trends and specific areas of higher fire incidence near Covington, VA. This will help in assessing the most vulnerable spots along the transmission lines and around residential areas that need urgent attention for vegetation management and infrastructure protection.
- Step 2: Use the Long term fire history records to gather insights into historical wildfire patterns in the area. This can offer a broader context of fire behavior over centuries, helping to predict potential future hotspots and periods of higher risk, which is crucial for long-term planning and resilience strategies.
- Step 3: Conduct a literature search focusing on effective strategies for vegetation management, forest health maintenance, and wildfire risk mitigation around power grids. This will include exploring practices such as controlled burns, vegetation clearance standards, and innovations in infrastructure design that enhance resilience to wildfires.
- Step 4: Based on the data analysis and best practices identified in scientific literature, provide recommendations tailored to improving forest management and enhancing the resilience of power transmission lines. This will include specific measures to ensure effective clearance around lines, techniques for early fire detection, and community engagement strategies to reduce the overall fire risk.

Results Our two-phase ablation study explores the impact of user profile granularity and diversity on WildfireGPT’s recommendations. In the first phase, we progressively test three levels of user profile specificity provided to the agent: no profile information provided, location plus timeline, and the full profile of a power grid manager focused on maintaining transmission line clearance and grid resilience for power distribution reliability and access. In the second phase, we hold location and timeline constant and test five distinct professional profiles—homeowner, civil engineer, ecologist, emergency manager, and power grid manager—each with minimally distinct concerns and the respective scopes. We report the results from both phases below.

Hierarchical Specificity from Profile Granularity Table 8 demonstrates the tiered progression in recommendation specificity provided by WildfireGPT in the first phase of our ablation study. The table focuses on three key aspects of wildfire resilience—vegetation management, technology implementation, and role-specific action—and demonstrates how recommendations become increasingly targeted and context-aware as input specificity increases. For example, a generic prompt without profile information yields broad guidance such as thinning and pruning or the use of satellite imaging and drones. When location and timeline details are included, the system generates more ecologically informed and time-sensitive strategies, such as "develop and promote fire-adapted ecosystems by incorporating native species." With a complete power grid manager profile, the recommendations become more specific and operational, with recommendations prioritizing "sensors and advanced monitoring systems along the grid" and "expand the Right-of-Way (ROW) clearance around transmission lines," synthesizing ecological approaches with sector-specific operational requirements. Most notably, technological recommendations progressed from general "satellite imaging and drones" to grid-specific applications for "real-time surveillance of power lines, especially in hard-to-reach areas." This stratification demonstrates how increasing profile granularity introduced by the user profile agent enables WildfireGPT to align general resilience principles with the operational needs of distinct professional roles, particularly in critical infrastructure contexts where vegetation management intersects with operational reliability requirements.

Professional Identity as a Recommendation Filter For the second phase, we summarize WildfireGPT’s adaptation to different professional roles and priorities across three dimensions: *plan generation, literature retrieval and recommendation specificity*.

First, our analysis shows that the planning agent adapts to the user’s professional role and concerns. For example, when a power grid manager expresses primary concern about maintaining transmission lines and grid resilience, WildfireGPT proposes to search for literature about vegetation management and wildfire protection strategies around power grids. Similarly, when a homeowner indicates a focus on marketable species, WildfireGPT targets research about economically valuable tree species and property protection methods. This pattern continues across all profiles—emergency managers receive evacuation-focused plans, civil engineers get infrastructure-centered recommendations, and ecologists receive biodiversity-oriented strategies. In each case WildfireGPT maintains the same

basic framework (analyzing fire data, reviewing literature, developing recommendations) but tailors the specific content to match the user’s professional needs and objectives. This alignment between user profiles and generated plans is by design, as the planning agent’s primary function is to establish clear expectations and a structured pathway through which the analyst agent can deliver personalized insights relevant to each stakeholder’s specific context.

Second, the analyst agent uses search queries calibrated to professional priorities when conducting literature retrieval, with details provided in Table 7. For each profession, the search queries effectively capture their unique concerns and scope—homeowners’ focus on marketable species, civil engineers’ emphasis on structural resilience, ecologists’ focus on ecological resilience, emergency managers’ priority on defendable space, and power grid managers’ concentration on grid resilience. These search queries distinctly different literature being retrieved, and we report the literature retrieval results based on these queries. When a power grid manager’s profile emphasized transmission line clearance and grid resilience, the system yielded studies such as Hvenegard’s (2014) [41] analysis of fuel treatment monitoring and Dale’s (2006) [42] critique of fire suppression policies. Conversely, when responding to a civil engineer’s profile, the system retrieved Stevens et al.’s (2020) [43] findings on vegetation structure and soil moisture dynamics. This pattern of query-response alignment extends across all professional identities: emergency management queries yielded Ager et al.’s (2019) [44] work on community wildfire exposure, while the ecologist’s profile resulted in Schoennagel et al.’s (2017) [45] research on adaptive resilience and Waltz et al.’s (2014) [46] findings on ecological impacts of fuel reduction treatments.

This adaptation pattern extends to recommendation development. As shown in Table 9, WildfireGPT generates distinct recommendation sets across the five professional profiles while maintaining core wildfire management principles. The table uses checkmarks to indicate which recommendations appear in each stakeholder’s profile—for example, all profiles receive recommendations for regular thinning and prescribed burns, while only the ecologist receives natural fire regime maintenance recommendations. Although certain fundamental strategies—such as regular thinning/pruning and controlled burns—are universally recommended across all profiles, significant variations emerge in other categories based on professional priorities. For civil engineers, recommendations pivot toward drainage efficiency and slope stability considerations, incorporating Stevens et al.’s (2020) [43] findings on soil moisture dynamics. Notably, civil engineers are the only profession not recommended to implement fire-resistant materials/infrastructure, instead focusing on ecological monitoring and emergency access solutions. Emergency managers receive the most comprehensive recommendation set, with checks in nearly all categories and exclusive emphasis on defendable space creation and evacuation planning. Ecologists uniquely receive recommendations for natural fire regime maintenance, reflecting their biodiversity conservation priorities, while being the only group besides civil engineers not advised on defendable space creation. Emergency management profiles elicit recommendations heavily weighted toward evacuation planning and defendable space creation, drawing on Ager et al.’s (2019) [44] community exposure analyses. Homeowners and power grid managers show similar recommendation patterns in infrastructure planning but diverge significantly in community engagement—homeowners receive evacuation planning advice while power grid managers do not. Interestingly, power grid managers share the research recommendation with emergency managers and ecologists, highlighting the technical expertise common to these roles.

This pattern of differential adaptation suggests that user profiles serve as effective filters for distilling broadly applicable wildfire management principles into professionally actionable, context-specific recommendations. The profile-specific recommendations reflect the granularity in both tactical precision and implementation feasibility, aligned with each stakeholder’s operational priorities while maintaining scientific rigor in the underlying management principles.

Though the user profile agent help successfully retrieve thematically-relevant literature at a profession-specific level, the specificity of matches invites further refinement that goes beyond the user profile agent. First, when matching papers to user concerns, it might not align exactly with the specific scope. For example, when the power grid manager requested information about protecting power lines, the system identified Hvenegard’s (2014) [41] paper on fuel treatment. This is valuable for fire management, but does not specifically address power grid infrastructure. Second, regarding location-specificity, searches for Covington, Virginia yielded Schoennagel et al.’s work in western North America or Hansen’s research in Grand Teton National Park [47]—which could still offer applicable insights despite the geographic mismatch. These patterns reflect two areas for enhancement for the analyst agents: (i) the finite scope of the available literature database where closely aligned

research may not exist for highly specialized user needs (such as power grid protection in specific geographic contexts); and (ii) inherent limitations in literature search mechanism as a part of the analyst agent. We explore these limitations and potential solutions in greater detail in the Discussion section.

G Information Fidelity of WildfireGPT’s Reported Wildfire Statistics

We verify whether the reported data and citation in WildfireGPT’s responses originate from the retrieved source. The fidelity verification process does not require domain expert knowledge but simply needs to match statistics in the responses to those in the retrieved data. We use precision scores to quantify the proportion of correct statistics reported by the language model; the formula is

$$\text{Precision} = \frac{\text{Number of correctly reported statistics}}{\text{Total number of statistics reported by the language model.}} \quad (1)$$

In the context of data analysis, if the textual response by WildfireGPT mentions a data point verbatim, then it is considered a correctly reported statistic. For example, consider the following sentence from a WildfireGPT response:

“The wildfire incident data for Mora County from 2015 to 2022 shows a notable increase in fire occurrences, particularly in 2018 with 29 incidents. The most active months for wildfires are during the summer, peaking in July with 21 incidents.”

In this case we would verify whether the numbers 29 (incidents in 2018) and 21 (incidents in July) are true based on the extracted data. If both numbers match the data, they would be counted as correctly reported statistics. Similarly, in the context of retrieving scientific literature, if WildfireGPT discusses three papers and two of them match the papers returned from the literature retrieval tool, the precision would be 2/3 or 66.7%.

For the literature review, we also calculate the Sentence-BERT (SBERT) [48] cosine similarity score between the retrieved papers and the generated response of WildfireGPT. The SBERT score ranges from -1 to 1, with higher scores indicating greater semantic similarity. A score above 0.5 is generally considered a good match [25]. A high SBERT cosine similarity score suggests that the language model is not only citing the correct papers but also capturing the main ideas and concepts from the literature.

We note that information fidelity between the retrieved sources and WildfireGPT’s outputs is not the same as independently validating claims against the broader scientific literature; this stringent view allows us to cleanly verify the model’s capacity to accurately synthesize information from credible sources.¹

As shown in Table 11, WildfireGPT achieves high fidelity scores across all case studies (overall success rates of 98.94% for data analysis and 100% for literature review, respectively) and high semantic similarity between the papers discussed by the language model and the abstracts of the retrieved papers (0.8559 average SBERT score). This is expected given that our design explicitly incorporates retrieved material, with prompts that specifically direct the language model to accurately summarize the data. Nevertheless, instances of ambiguity occasionally emerge. In the Community Hazard Mitigation Planning case, WildfireGPT stated that “spring FWI projections for the mid-century (2045–2054) and the end of the century (2085–2094) show a gradual increase, particularly notable during the spring season, where the FWI is projected to rise significantly from medium to potentially high-risk levels.” According to the retrieved data, however, the projections for the spring seasons were classified as medium risk for both the mid-century (2045–2054) and the end of the century (2085–2094), despite showing an increasing trend, which led to a relatively lower fidelity score.

¹While SBERT is used to find relevant sources by matching queries to our database, we separately use SBERT to check whether the system’s final responses actually align with the content of those retrieved sources. This separate evaluation step is necessary since retrieving relevant sources does not guarantee the system will use them appropriately when generating responses. Additionally, while verifying that sources are properly included is a basic requirement (as opposed to hallucinated), it remains an important check for system reliability.

Table 11: Fidelity scores and SBERT scores for each case study. The fidelity scores are represented as precision percentages, calculated as the number of correctly reported statistics divided by the total number of statistics reported by WildfireGPT. The raw counts of correctly reported statistics and total statistics are provided. The SBERT scores are reported as average values across all retrieved papers for each case study, along with the number of times papers were retrieved (in parentheses). "N/A" indicates that the case study did not involve a literature review component. Case studies are listed in alphabetical order.

Case Study	Data Analysis	Literature Review	Avg. SBERT Score
Climate Impact on U.S. Wildfires	5/5	N/A	N/A
Hazard Mitigation Planning	10/11	N/A	N/A
Comprehensive Wildfire Impact	16/16	5/5	0.8529 (2)
Current Wildfire Risk Analysis	11/11	N/A	N/A
Ecosystem Fire Management	6/6	3/3	0.9150 (1)
Infrastructure Wildfire Risk	10/10	3/3	0.8398 (1)
Post-Wildfire Public Safety	1/1	3/3	0.7742 (1)
Private Property Protection	1/1	17/17	0.8694 (6)
Urban Wildfire Mitigation	8/8	9/9	0.8556 (3)
Wildland Urban Interface Impact	12/12	3/3	0.8844 (1)
Average Across All Percentages or Scores	98.09%	100.00%	0.8559
Overall Success Rate	98.94% (80/81)	100.00% (45/45)	N/A

H Domain Expert Evaluation

Table 12: Domain expert evaluation criteria and corresponding questions used to assess WildfireGPT's performance in generating responses. Each question was rated on a three-point scale: "Yes" (1), "Could be better" (0.5), and "No" (0).

Criteria	Questions
Relevance	(1) Does my response answer your last question? (2) Is my response relevant to your profession? (3) Is my response relevant to your concern? (4) Is my response relevant to your location? (5) Is my response relevant to your timeline? (6) Is my response relevant to your scope?
Entailment	(1) Do my analyses or recommendations logically follow from the information (data, literature) provided?
Accessibility	(1) Does my response contain too much jargon? (2) Does my response provide enough explanation? (3) Does my response contain redundant or useless information?

During the case studies, whenever WildfireGPT generates a response, domain experts from Argonne National Laboratory's Environmental Science Division and Decision and Infrastructure Sciences Division are asked to evaluate its relevance, entailment, and accessibility using a structured questionnaire. These experts are scientists and engineers specializing in natural hazards and wildfire control with different detailed concentrations or other related professions as listed in Table 1. Each question is rated on a three-point scale, with "Yes" receiving a score of 1, "Could be better" a score of 0.5, and "No" a score of 0. The valuation criteria and corresponding questions are presented in Table 12. Relevance assesses whether the model's responses appropriately address the user's last question and are relevant to the user's profession, concerns, location, timeline, and scope. Entailment evaluates the logical coherence between the model's analyses/recommendations and the provided data or literature. Accessibility examines the clarity and concision of the model's language, considering factors such as jargon, explanatory detail, and redundancy. We report the total scores for each criterion across all case studies. To complement the quantitative analysis, we also report qualitative insights from the expert evaluations to provide a more nuanced understanding of WildfireGPT's performance.

We additionally verify whether the reported data and citations in WildfireGPT’s responses originate from the retrieved sources.

Table 13: Relevance scores for each case study, representing the percentage of responses that appropriately addressed the user’s context and needs. Raw scores are shown (score/total times the domain expert opinion was collected to evaluate a response).

Case Study	Last Question	Profession	Concern	Location	Time	Scope
Climate Impact on U.S. Wildfires	17/18	18/18	18/18	18/18	18/18	18/18
Comprehensive Wildfire Impact	6/6	7/7	7/7	7/7	7/7	7/7
Current Wildfire Risk Analysis	5/5	6/6	6/6	6/6	5.5/6	6/6
Ecosystem Fire Management	3/3	2.5/3	3/3	3/3	3/3	3/3
Hazard Mitigation Planning	5/5	7/7	7/7	7/7	7/7	7/7
Infrastructure Wildfire Risk	7/7	7/7	7/7	7/7	7/7	7/7
Post-Wildfire Public Safety	3/3	4/4	4/4	3.5/4	4/4	4/4
Private Property Protection	6.5/7	8/9	9/9	8/9	9/9	9/9
Urban Wildfire Mitigation	3.5/4	4/4	4/4	4/4	4/4	4/4
Wildland-Urban Interface Impact	2/2	4/4	4/4	4/4	3.5/4	4/4
Average across All Percentages	97.48%	97.22%	100.00%	97.64%	97.92%	100.00%
Overall Success Rate	96.67% (58/60)	97.83% (67.5/69)	100.00% (69/69)	97.83% (67.5/69)	98.55% (68/69)	100.00% (69/69)

Table 14: Entailment and accessibility scores for each case study. Entailment indicates the percentage of responses where the model’s analyses and recommendations logically followed from the provided information (data, literature). Accessibility assesses the clarity and concision of the model’s language. Scores represent the percentage of responses that avoided jargon, provided sufficient explanation, and minimized redundancy. Raw scores are shown (score/total times the domain expert opinion was collected to evaluate a response).

Case Study	Entailment	Accessibility		
		No Jargon	Enough Explanation	No Redundancy
Climate Impact on U.S. Wildfires	11/13	17/17	16/17	14.5/17
Comprehensive Wildfire Impact	7/7	7/7	7/7	7/7
Current Wildfire Risk Analysis	5.5/6	6/6	5.5/6	6/6
Ecosystem Fire Management	3/3	3/3	3/3	3/3
Hazard Mitigation Planning	6/6	7/7	7/7	7/7
Infrastructure Wildfire Risk	6.5/7	7/7	7/7	7/7
Post-Wildfire Public Safety	3/3	4/4	4/4	4/4
Private Property Protection	6/6	8/8	7/8	7/8
Urban Wildfire Mitigation	1.5/2	4.5/5	2.5/5	5/5
Wildland-Urban Interface Impact	2.5/3	4/4	4/4	4/4
Average Across All Percentages	92.75%	99.00%	92.33%	97.28%
Overall Success Rate	92.86% (52/56)	99.26% (67.5/68)	92.65% (63/68)	94.85% (64.5/68)

Relevance: Domain experts generally found WildfireGPT’s responses contextually appropriate (a success rate of 96.67% to 100% across all relevance categories, Table 13), but there are areas for improvement. For the “Last Question” and “Profession” categories (Table 12, Q1 & Q2), domain experts noted occasional misalignments, such as offering broad advice when specific zoning changes are requested or making impractical recommendations such as suggesting that homeowners host

community workshops. The "Location" category revealed potential limitations in the literature corpus (Table 12, Q4), such as offering recommendations for homeowners in Virginia based on studies from the West Coast. WildfireGPT sometimes failed to specify which strategies could be reasonably implemented within given timeframes (Table 12, Q5).

Entailment: With an overall success rate of 92.86% (Table 14), WildfireGPT generally remained logically consistent; but experts noted some issues when WildfireGPT provided recommendations, such as a lack of transparency in tracing the sources of its recommendations and a failure to present recommendations in a logical hierarchy from most important to least important. These situations could lead to long-winded responses that lack focus.

Accessibility: WildfireGPT generally performed well in avoiding jargon (99.26% overall success rate) and redundancy (92.65% overall success rate), with most responses providing sufficient explanation (94.85% overall success rate); see Table 14. However, the interpretation of these scores varied by audience. Urban planners and homeowners frequently requested more detailed explanations and examples as well as clearer definitions of technical terms and contextual information for data interpretation. In contrast, climate scientists preferred more concise responses, with less detail and more direct answers to their questions. These findings point to the challenge of balancing detail and concision to meet diverse user needs.

Overall, the domain expert evaluation demonstrates that WildfireGPT is capable of generating relevant, logically sound, and accessible responses to user queries across a range of wildfire-related scenarios. However, the results also highlight potential areas for improvement, such as improving the relevance of responses to the user's last question in certain cases and ensuring differing levels of explanation to maximize clarity and utility for different domain experts.

Qualitative Feedback on WildfireGPT's Recommendations To better understand the nuance and practical value of WildfireGPT's recommendations, we sought qualitative feedback from domain experts as they evaluated the system's responses via the structured questionnaire reported above.

One notable example of WildfireGPT's effectiveness in integrating data analysis with domain-specific knowledge is demonstrated in the Ecosystem Fire Management case study (Figure 6). When asked for suggestions on adapting fire management strategies to maintain oak ecosystems while minimizing risks to residential areas, WildfireGPT analyzed recent wildfire data, revealing that controlled burns typically occur in March. It validated this response by citing literature such as that of Stan, Rigg, and Jones (2006) [49], which emphasizes the importance of March fires for white oak regeneration. WildfireGPT also analyzed seasonal FWI data, identifying lower-risk periods in early spring and late autumn. Synthesizing these insights from more than one data source, it recommended timing controlled burns to coincide with historically low FWI values, balancing ecological management goals with safety. The domain expert praised WildfireGPT's nuanced approach, emphasizing its ability to balance ecological management goals with safety considerations.

The Wildland Urban Interface Impact case study (Figure 7) further demonstrates WildfireGPT's practical utility. When the domain expert asked questions from the perspective of a risk manager from Las Vegas, NM, WildfireGPT provided recommendations on water resource protection that aligns well with actual challenges faced in the field. Furthermore, it identifies California as a relevant comparison due to similar fire characteristics. The retrieved study by Moritz and Stephens (2008) [50] recommends strategies such as risk-based frameworks, controlled burns, and reevaluating urban planning—approaches the domain expert recognizes as potentially valuable for Las Vegas, NM.

The Private Property Protection case study (Figure 8) highlights WildfireGPT's ability to provide insightful ideas that might not be immediately apparent to users. When a domain expert, assuming the role of a Virginia homeowner, sought advice on managing oak and cherry trees while reducing wildfire risks, WildfireGPT not only suggested standard strategies such as fire monitoring and defensible spaces but also raised an interesting point about the potential for landowners to benefit from carbon credits by investing in oaks. Additionally, WildfireGPT emphasized the importance of mixing oak species with other native, fire-resistant trees to promote ecological balance and adaptability. The domain expert praised these points, noting that they might not immediately come to mind for many but offer valuable economic and ecological benefits.

Similarly, in the Infrastructure Wildfire Risk case study (Figure 9), a risk and resilience analyst discussed strategies for mitigating wildfire risks with WildfireGPT, focusing on the use of FWI

Case Study: Ecosystem Fire Management

Context: WildfireGPT analyzed recent wildfire incident data and found that from 2018 to 2021, the months with recorded incidents were March (2 incidents) and July (1 incident), suggesting that March may be when controlled burns are typically implemented. To confirm this point, WildfireGPT retrieved a study by Stan, Rigg, and Jones (2006) [49], which highlighted the importance of reintroducing fire in March for white oak regeneration in northeastern Illinois. Additionally, WildfireGPT analyzed seasonal FWI data, identifying lower-risk periods during spring.

User Query: Please suggest specific adaptations to fire management strategies (to balance the utilization of fire for managing oak ecosystems and reducing the risk to residential and human developments.)

WildfireGPT Response (Snapshot):

Optimized Timing for Controlled Burns:

- Continue to utilize controlled burns in March when the FWI (Fire Weather Index) historically shows lower values, minimizing the risk of uncontrollable wildland fires.
- Schedule additional controlled burns during early spring or late autumn to take advantage of similarly low FWI scores, ensuring burns are conducted under optimal conditions for controlling spread and smoke dispersal.

Expert Comment: The answer effectively addresses a key concern: whether to continue conducting controlled burns in March. It confirms that March remains a suitable time for burns due to historically lower Fire Weather Index (FWI) values, which reduce the risk of uncontrolled fires. The response also provides a nuanced suggestion by recommending additional burn periods in early spring and late autumn. This approach takes advantage of similarly low FWI scores during these times, allowing for burns at both the beginning and end of the dormant season. This strategy demonstrates a nuanced understanding of fire management, as it balances the need for regular burns with the importance of conducting them under optimal weather conditions for safety and control. The suggestion to utilize multiple burn windows throughout the year shows a thoughtful approach to ecosystem management while minimizing risks to nearby developments.

Figure 6: Ecosystem Fire Management: WildfireGPT demonstrates its ability to integrate data analysis and domain knowledge to provide actionable recommendations for ecosystem fire management. By analyzing wildfire incident data, retrieving relevant literature, and examining seasonal FWI trends, WildfireGPT suggests optimizing the timing of controlled burns to minimize risks while maintaining oak ecosystem health. The domain expert's positive feedback highlights WildfireGPT's nuanced approach and its potential to support informed decision-making in wildfire management.

projections and firebreaks. When asked about the factors that contribute to an area being considered high risk for wildfires, WildfireGPT not only mentioned ecological importance but also highlighted the often-overlooked aspect of cultural significance. The domain expert commended the system's recognition of this intangible factor, noting that it reminds users of the broader, non-quantifiable impacts that wildfires can have on communities and their heritage.

The Community Hazard Mitigation Planning case study (Figure 10) further demonstrates WildfireGPT's ability to generate nuanced and context-specific recommendations by integrating both explicitly discussed information and implicit knowledge about the unique needs of the target community. When a hazard mitigation planner for Mora County, NM, sought guidance on wildfire risk adaptation strategies tailored to the community's demographics, WildfireGPT suggested developing a multilingual alert system and focusing on community-based approaches. Despite the census data only including poverty rates and housing units, the system drew on broader knowledge to generate these insightful recommendations, indicating its awareness of the significant Spanish-speaking population in the region. This case study showcases WildfireGPT's potential to offer nuanced decision support that extends beyond mere brainstorming, ultimately enhancing the effectiveness of wildfire risk adaptation efforts.

These case studies collectively underscore WildfireGPT’s value in providing insightful, context-specific recommendations that consider the multifaceted nature of wildfire risk management. By integrating diverse information sources and accounting for the ecological, cultural, and socioeconomic dimensions of the communities it serves, the system demonstrates its ability to offer tailored guidance that can inform more effective and feasible wildfire risk mitigation strategies. These case studies highlight WildfireGPT’s potential to stimulate innovative thinking and support informed decision-making by generating recommendations that extend beyond conventional approaches, ultimately enhancing the effectiveness of wildfire risk adaptation efforts across various contexts.

I Scaling Up Evaluation by LLM-as-a-Judge

To explore the potential for scalable, automated evaluation that could maintain quality assurance during deployment without constant human oversight, we implemented an LLM-as-a-judge approach. This method involves using a separate GPT-4 [27] query to assess the quality of WildfireGPT’s responses based on the same criteria used in the expert evaluation. We report the agreement of the LLM-as-a-judge evaluation with the expert evaluation to determine the feasibility of using this approach. Each case study involved a unique interaction between the user and WildfireGPT, resulting in varying uses of data sources and different numbers of rounds of follow-up interactions. Our evaluation framework was designed to accommodate this variability, ensuring that the assessment remained meaningful and relevant to each specific case.

To evaluate GPT-4’s [27, 51] feasibility as an automated judge for WildfireGPT, we compared GPT-4 assessments with human judgments on relevance, entailment, and accessibility. We applied LLM-as-a-judge to evaluate WildfireGPT’s responses only immediately after data or literature retrieval to manage context length and maintain performance. Agreement rates varied: relevance (62.99%), entailment (75.00%), and accessibility (66.67%). Disagreements primarily occurred when experts rated responses positively ("Yes"), whereas GPT-4 suggested improvements ("Could be Better"), reflecting GPT-4’s stricter evaluative criteria.

Table 15: Agreement between human evaluation and LLM-as-a-Judge evaluation of WildfireGPT responses. Most misalignments occur when domain experts respond with "Yes" while GPT-4 tends to give a more critical judgment, saying "Could be better" instead of a larger disagreement like "Yes" vs "No."

Category	Agree	Disagree	Yes vs Could be better
Relevance	97/154 (62.99%)	57/154 (37.01%)	51/57 (89.47%)
Entailment	15/20 (75.00%)	5/20 (25.00%)	2/5 (40.00%)
Accessibility	56/84 (66.67%)	28/84 (33.33%)	22/28 (78.57%)

The moderate 62.99% agreement in WildfireGPT relevance evaluations is largely due to nuanced differences rather than fundamental disagreements, often due to the data availability issues since retrieved data may not exactly match user-specified timelines or locations, whereas experts were more lenient, acknowledging these constraints.

In most entailment disagreements, domain experts found responses satisfactory ("Yes"), while GPT-4 flagged them as needing improvement ("Could be better" or "No"). These discrepancies often stemmed from WildfireGPT extrapolating beyond available data. For example, based on FWI data showing minimal change over time, WildfireGPT described wildfire risk as "relatively stable"; GPT-4 criticized this, arguing that it can be misleading and that there are possibilities for wildfire risk to increase. In another case, when asked about fire controls for cherry trees, the retrieved literature did not directly address this specific species but rather the same family of trees in general; experts found this reasonable given the limitations of the available literature, but GPT-4 criticized it as unsupported. Experts appeared more forgiving of such extrapolations, likely because of the experts filling in gaps with their own background knowledge.

Accessibility disagreements highlighted differing standards: GPT-4 recommended simpler wording or more detailed explanations. However, experts’ satisfaction varied with their specific expertise and familiarity with the topic. For example, In areas of deep knowledge they may prefer concise,

technical responses, while in less familiar areas they require more detailed explanations. GPT-4's critiques do not fully account for this spectrum of expertise.

In summary, the results reveal moderate agreement between GPT-4 and human evaluations. While GPT-4 can provide scalable and cost-effective evaluations, it may not fully capture the nuanced expectations and information needs of domain experts, leading to discrepancies in relevance, entailment, and accessibility assessments.

Case Study: Wildland-Urban Interface Impact

Context: In Las Vegas, NM, a risk manager developing urban plans for 2050 is concerned about the increasing wildfire risks and their impact on water resources. Previous analysis of Fire Weather Index (FWI) data indicated an increasing trend in fire danger levels, particularly during the spring season.

User Query: Proceed with the recommendations (on urban planning strategies to minimize wildfire risks and methods to protect water resources from the adverse effects of wildfires)

WildfireGPT Response (Snapshot):

Water Resource Management

- Water Supply Resilience: Increase the resilience of water supply systems to ensure availability during and after wildfire events. This could involve constructing new reservoirs, enhancing existing water treatment facilities, and establishing emergency water supply locations.
- Protect Water Quality: Develop strategies to protect water quality from the effects of wildfires, such as increased sedimentation and contamination. These might include erosion control measures and enhanced filtration systems at water treatment facilities.

Expert Comment: The recommendation to protect water quality from increased sedimentation and contamination is highly relevant. This issue was a central concern when we initially developed this case study, and it remains a significant challenge for the area. The local water treatment plant is struggling to handle the fine-grained sediment that is now flowing downstream in much larger quantities than before. The plant's processes are overwhelmed by the volume of material, which far exceeds what they typically dealt with from normal erosion. It's very good that the answer highlights the need to address water quality issues, specifically mentioning sedimentation and contamination. The suggestion to implement erosion control measures and enhance filtration systems directly addresses the problems we encountered. In our modeling work, we aimed to predict the sediment load to help design appropriate mitigation measures. The recommendation aligns well with the actual challenges faced in the field and the solutions we were exploring.

User Query: Are there good examples of similar cities that might experience similar shifts in wildfire risk? What mitigation strategies are those cities implementing?

WildfireGPT Response (Snapshot):

Here's a look at some of the approaches highlighted in recent research:[...] California is a prime example of a region adapting to increased wildfire risks. The study by Moritz and Stephens (2008) emphasizes the need for a sustainable coexistence with wildfire through policy, planning, and management changes. Key strategies include the following:

- Risk-Based Framework: Implementing a risk-based framework for fire-related decisions to address natural hazards systematically.
- Reintroduction of Fire: Using controlled burns to reintroduce fire in a managed way to fire-prone ecosystems.
- Urban Planning Reevaluation: Reevaluating urban planning and building in fire-prone locations to make future cities less susceptible to wildfire damage.

References:

Moritz, M., & Stephens, S. (2008). Fire and sustainability: considerations for California's altered future climate. <https://doi.org/10.1007/s10584-007-9361-1>

Expert Comment: The WildfireGPT response effectively identifies California as a region experiencing similar shifts in wildfire risk to Las Vegas, NM. Despite some differences in fire characteristics, California remains a relevant comparison. The strategies highlighted, such as implementing a risk-based framework, reintroducing controlled burns, and reevaluating urban planning, could provide valuable guidance for Las Vegas.

Figure 7: Wildland Urban Interface Impact: WildfireGPT demonstrates its effectiveness in addressing the complex challenges posed by the wildland-urban interface. Wildfire's recommendations regarding water quality protection align well with the actual challenges faced by domain experts in the field. WildfireGPT's ability to identify relevant case studies from other areas with similar shifts in wildfire risk and to highlight pertinent mitigation strategies showcases its potential to support informed decision-making in wildfire risk management and urban planning.

Case Study: Private Property Protection

Context: A Virginia homeowner near Covington seeks advice on forest management, focusing on oak and cherry tree health, while mitigating wildfire risks. After analyzing recent fire data, noting an increase in wildfire incidents, especially during spring and fall, WildfireGPT suggests strategies like fire monitoring, fuel management, controlled burns, defensible spaces, and reforestation with fire-resistant species. Oak trees are highlighted for their fire resilience, making them a smart reforestation investment.

User Query: Could you please provide a more detailed summary for investing in oak species?

WildfireGPT Response (Snapshot):

- Carbon Sequestration: As large and long-lived trees, oaks are effective at capturing and storing carbon dioxide, which could potentially qualify landowners for carbon credits under emerging environmental regulations and markets.
- Mixing Species for Resilience: While investing predominantly in oaks, it is advisable to include a mix of other native, fire-resistant species to ensure ecological balance and adaptability to changing climate conditions.

Expert Comment: Carbon sequestration is an insightful point that might not immediately come to mind for many. It's valuable to remind people that trees, especially long-lived species like oaks, offer more than just timber value. They can also serve as carbon sinks, potentially providing landowners with additional income through carbon credits in emerging environmental markets. Mixing species for resilience is also a crucial aspect to highlight. While the focus is on oak species, it's important to emphasize that a diverse planting strategy is more beneficial. Encouraging a mix of native, fire-resistant species alongside oaks promotes greater ecological resilience and adaptability to changing climate conditions. This approach aligns with best practices in forest management and conservation. Both points add depth to the investment strategy, considering both economic and ecological factors.

Figure 8: Private Property Protection: WildfireGPT showcases its ability to provide tailored advice to individual property owners seeking to balance wildfire risk mitigation with their specific land management goals.

Case Study: Infrastructure Wildfire Risk

Context: A risk and resilience analyst discussed strategies for mitigating wildfire risks with WildfireGPT, particularly focusing on the use of FWI projections and firebreaks. The conversation has detailed the implementation of firebreaks to protect transmission lines, including site assessment, types, and maintenance of firebreaks. They explored integrating these with other mitigation measures and engaging with local communities for effective implementation.

User Query: What goes in to being a critical area as an high-risk areas prone to wildfires. Can you explain?

WildfireGPT Response (Snapshot):

Ecological and Cultural Significance

- Protected Areas: Wildlife reserves, national parks, and other protected lands are often considered high-risk areas due to their ecological importance.
- Cultural Heritage: Areas with cultural or historical significance may also be prioritized for protection against wildfires.

Expert Comment: The inclusion of cultural significance is a crucial aspect in explaining the impact of natural disasters, particularly wildfires. This intangible factor is often overlooked in decision-making processes for natural hazard mitigation, yet it remains highly important. Highlighting cultural significance reminds users of the broader, non-quantifiable impacts that wildfires can have on communities and their heritage. It's good that the answer recognizes this often undervalued dimension of wildfire risk assessment.

Figure 9: Infrastructure Wildfire Risk: WildfireGPT demonstrates its ability to provide fresh perspectives and consider often-overlooked dimensions of wildfire risk assessment.

Case Study: Hazard Mitigation Planning

Context: A hazard mitigation planner for Mora County, NM, is assessing wildfire risks over the next 15 years. In the previous discussion, WildfireGPT found a notable increase in fire occurrences from 2015 to 2022 and a clear trend of increasing wildfire risk in the coming decades due to climate change impacts. From the U.S. census data, the community, with a population of about 10,297, is particularly vulnerable due to high poverty levels.

User Query: Let's explore specific adaptation strategies tailored to these demographics (in Mora County, NM). **WildfireGPT Response (Snapshot):**

Enhanced Emergency Communication Systems:

- Develop a multilingual alert system that provides real-time updates on wildfires via SMS, social media, and local radio stations to ensure that all community members receive timely information.
- Set up community billboards in key areas to display emergency information and tips on wildfire preparedness.

Expert Comment: The suggestion to develop a multilingual alert system is insightful, particularly given the significant Spanish-speaking population in the area. Interestingly, this demographic detail wasn't explicitly mentioned in the provided census data, which only included poverty rates and housing units. This indicates that the AI model is drawing on broader knowledge about the region to make this recommendation. It's noteworthy that the model suggests community-focused strategies rather than individual-based ones. This approach is particularly relevant for a lower-income community like Mora County, where collective efforts may be more effective and feasible for wildfire mitigation. The model's ability to integrate both explicitly provided information and implicit knowledge about the area demonstrates its capacity for nuanced and context-appropriate recommendations.

Figure 10: Community Hazard Mitigation Planning: WildfireGPT demonstrates its ability to generate nuanced and context-specific recommendations by integrating both explicitly discussed information and implicit knowledge about the unique needs of the target community.