

Building Multi-Agent Copilot towards Autonomous Agricultural Data Management and Analysis

Yu Pan¹, Jianxin Sun^{2,4}, Hongfeng Yu^{2,4}, Joe Luck¹, Geng Bai⁵, Nipuna Chamara¹, Yufeng Ge¹, Tala Awada³

¹*Department of Biological System Engineering, University of Nebraska-Lincoln, Lincoln, NE, USA*

²*School of Computing, University of Nebraska-Lincoln, Lincoln, NE, USA*

³*School of Natural Resources, University of Nebraska-Lincoln, Lincoln, NE, USA*

⁴*Holland Computing Center, University of Nebraska-Lincoln, Lincoln, NE, USA*

⁵*Department of Biological and Agricultural Engineering, North Carolina State University, Raleigh, NC, USA*

Abstract—The ubiquity of sensors and IoT devices has led to an explosion in data availability in modern agriculture. The large volume and heterogeneity of the data, together with the complexity of data processing requirements, pose huge obstacles for achieving the principles of Findable, Accessible, Interoperable, and Reusable (FAIR). Current data management and analysis paradigms are to a large extent traditional, in which data collecting, curating, integration, loading, storing, sharing and analyzing still involve too much human effort and know-how. The experts, researchers and the farm operators need to understand the data and the whole process of data management pipeline to make full use of the data. The essential problem of the traditional paradigm is the lack of a layer of orchestrational intelligence which can understand, organize and coordinate the data processing utilities to maximize data management and analysis outcome. The emerging reasoning and tool mastering abilities of large language models (LLM) make it a potentially good fit to this position, which helps a shift from the traditional user-driven paradigm to AI-driven paradigm. In this paper, we propose and explore the idea of a LLM based copilot for autonomous agricultural data management and analysis. Based on our previously developed platform of Agricultural Data Management and Analytics (ADMA), we build a proof-of-concept multi-agent system called ADMA Copilot, which can understand user’s intent, makes plans for data processing pipeline and accomplishes tasks automatically, in which three agents: a LLM based controller, an input formatter and an output formatter collaborate together. Different from existing LLM based solutions, by defining a meta-program graph, our work decouples control flow and data flow to enhance the predictability of the behavior of the agents. Experiments demonstrate the intelligence, autonomy, efficacy, efficiency, extensibility, flexibility and privacy of our system. Comparison is also made between ours and existing systems to show the superiority and potential of our system.

Index Terms—Agricultural Data Management, Multi-Agent, Copilot, FAIR principles, Large Language Model, Autonomous, Paradigm shift

I. INTRODUCTION

The core topics of modern agriculture is to meet the escalating needs for food, fuel, feed, and fiber in the era of climate change, limited natural resources and intricate dynamics of global ecosystems. To solve this challenge, data-informed precision agriculture is not only required but also necessary, in which the deployment of sensors, Internet of Things (IoT)

devices [1], and advanced instrumentation [2] has enabled data collection in various modal, dimension, format, and resolution. While agricultural data has the potential to transform agricultural practice by providing valuable insights into crop productivity, soil quality, and sustainable farming techniques, in the meanwhile, the large volume and heterogeneity, together with the complexity of data processing requirements, pose huge obstacles for achieving the principles of Findable, Accessible, Interoperable, and Reusable (FAIR), which in turn result in in-efficient utilization and exploration of the collected data.

There are existing efforts trying to revolutionize current practices of agriculture data management [3]–[7], by implementing different data management platforms and technologies. Nevertheless, current practice of agricultural data management and analysis is to large extent traditional in the sense that the whole data management pipeline, including data collecting, curating, integration, loading, storing, sharing and analyzing still involve too much human effort and know-how. The experts, researchers or the farm operators make every tiny decisions from where to put a file, to which analysis tool to use. In this process, they need to understand the details of the data and tools at their hands, including location, semantics and format of data, and the location, semantics and the interface of the relevant data processing tools to make fully use of the data. This human-driven paradigm incurs high cost of training and makes reuse of the data and duplication of the data processing pipeline difficult. The essential reason of the challenges faced by the traditional practice of agricultural data management lies in the fact that there’s a lack of a layer of orchestrational intelligence which can understand, organize and coordinate each phase of the data processing pipeline to maximize the outcome.

In the meanwhile, Large Language Models (LLM) recently start to demonstrate reasoning and tool mastering capabilities. That is, augmented with a set of tools and given a task, LLM need to make plans for the task, select appropriate tools for each step, call the tools and observe the feedback from these tools. Also there can be a observe-and-take-action loop until the task can be accomplished. LLM’s reasoning capability makes it feasible to make plans for a task or make decisions about which tool will be used in next step. The application

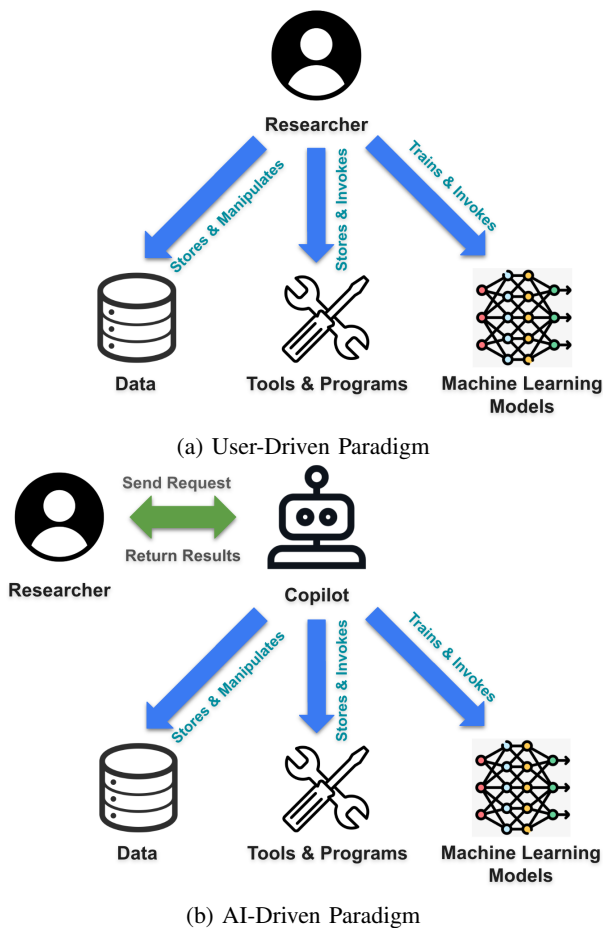


Fig. 1: Shift of Agricultural Data Management Paradigm

arises naturally where we explore the possibility to fit LLM as the aforementioned orchestrational layer which can coordinate and harness each component of agricultural data management and coordinate each phase of the data processing pipeline. Imagining the tools are not isolated with each other, but form a network where each tool can connect with each other through their inputs and outputs. The tool using intelligence of LLM can probably find new paths in the complicated tool networks and as a result gain new patterns and insight. This will bring us a new data management paradigm in which human effort can be reduced while the outcome of data analysis can be probably maximized. In one word, LLM based copilot will act as a controller in this new data management paradigm, which helps human users to manage data, tools or ML models. Furthermore, as new architectures of LLM-based data management framework being proposed and as new tools, IoT devices or sensors being continuously incorporated, the systems adopting the new paradigm will hopefully grow into a full-fledged autonomous data management system, which can take care of mundane operations with minimal human interference, and help human users to make decisions in data management tasks. Fig 1 illustrates the shift of agricultural data management.

In this paper, we propose and explore the idea of building a LLM-based copilot for autonomous agricultural data management system. We build a proof-of-concept multi-agent system, called ADMA copilot, which can try its best to understand user's intent through interactions in both natural language and GUI, and then make plan for data processing pipeline and accomplish the task automatically. The existing LLM-based tool using systems generate function calls together with arguments together, probably resulting in hallucinating problem where LLM fabricates the value of arguments without solid ground truth. Different from existing solutions, our system decouple the data flow and control flow, by introducing a meta-program graph which stores all the information about how to use a tool and what a variable stores. In this way, we solve the problem of hallucinating and increase the controllability and predictability of the copilot. Under the hood, the copilot has three agents collaborating together, one controller which decides the tool called in next step, one input formatter which modifies the value of variables based on user input, and one output formatter which generates UI, based on user's instruction. All the agents base their reasoning on the current status of the meta-program graph, which helps them to make collected decisions. We make comparison with existing agricultural data management platforms by defining a set of criteria to show the superiority and potential of our system in the long run.

We conduct extensive experiments to demonstrate the key features of the our agricultural data management copilot listed as follows:

- **Intelligence:** by introducing LLM based agents as program controller, input formatter and output formatter, the copilot demonstrates intelligence in terms of understanding human intent, solution planning and tool/function calling.
- **Autonomy:** data management tasks which implicitly contains several steps can be executed automatically with minimal human interference. The design principle of our system is not asking for user input or instruction unless necessary.
- **Efficacy:** multi-agent design decouples data flow from control flow, which will assure the predictability of the copilot and eradicate the phenomena of hallucination of LLM.
- **Efficiency:** when executing complicated data management tasks, the copilot is way more efficient than human being, thanks to streamlining the whole pipeline.
- **Extensibility:** the tool set utilized by the copilot can be extended, by augmenting the meta-program graph, which will guarantee the copilot can have versatile application in the future.
- **Flexibility:** the copilot can understand human input in a flexible way. In some cases, even the user's input is fuzzy and has mistake, the copilot can still get the intent of the user.
- **Privacy:** the copilot will protect user's data by asking the user to provide necessary credentials for various services.

Connected to our existing platform ADMA and various external tools, the data management copilot can not only help users to deal with tedious low-level data operations, but

also can be evolved to substitute human beings in decision making, which both help to make the shift of agricultural data management paradigm.

The key contributions of our work are summarized as follows:

- We design and implement Agricultural Data Management and Analytics Copilot (ADMA Copilot), which is a multi-agent copilot for autonomous agricultural data management and analysis. To our knowledge, it is the first of its kind, not only in agricultural data management, but in scientific data management as a whole.
- As a proof-of-concept product, ADMA Copilot can help to facilitate the shift of agricultural data management paradigm from traditionally human-driven paradigm to AI-driven paradigm.
- Extensive demos and evaluations demonstrate the intelligence, autonomy, efficacy, efficiency, extensibility, flexibility and privacy of our system. Comparison is also made between ours and existing systems to show the superiority and potential of our system.

The rest of the paper is organized as follows: Section II surveys existing work relevant to our research. Section III presents the design principles and frameworks of our copilot. In Section IV, we provide the thorough demos and evaluations of the current status of ADMA. Finally, in Section V, we conclude the paper by discussing the contributions and suggesting future directions.

II. RELATED WORK

A. Tool Learning with Large Language Models

Tool learning with Large Language Models (LLMs) has emerged recently as a promising paradigm for augmenting the capabilities of LLMs to accomplish complicated tasks, by utilizing LLMs reasoning abilities. Here the tool can be a program, a database, an API endpoint for a cloud based service, a sensor, hardware devices or anything of which the interface can be coded as plan text and thus can be recognized and utilized by LLMs.

A large corpus of exiting work has been proposed to explore different aspects and phases of Tool learning with LLMs [8]–[10]. Tool learning with LLMs is in fact an umbrella concept which encompasses a lot of subcategories, including augmenting LLMs with external knowledge [11]–[14], augmenting LLMs with mathematical reasoning ability [15], [16], augmenting LLMs with program executors [17], [18], and augmenting LLMs with domain specific tool sets [19], [20]. Generally speaking, there are four phases in tool learning with LLMs: task planning, tool selection, tool calling and response generation [8]. In task planning, a task is decomposed into several sub-tasks, by utilizing LLM’s reasoning capabilities [21]–[26]. In tool selection, the most appropriate tool is selected by LLM for each phase generated from task planning [27], [28]. In tool calling phase, the selected tool is called, based on given interface of the tool. Again LLM will be used to format the calling interface, that is, configure the order and

value of arguments for a method call [29]–[31]. Nowadays, large language models such as GPT-4 can output function calling, given the description of the tool and the user input [32]. Lastly, the output of the tools can be combined and formatted for output, by utilizing LLM’s language organization capability [33]. Each phase can utilize LLM and thus there will be several agents collaborating together for the whole process.

The existing works explore different frameworks to implement tool learning with LLMs, however, almost all the works don’t decouple data flow from control flow, which may result in hallucination problem and in turn fail a task. In our work, we generate a meta-program graph, as a representation of the tools and their affiliated data and use three agents: controller, input formatter and output formatter to deal with control flow and data flow respectively. In this way, we can increase the predictability of the result generated by the LLM-based agents and thus increase the success rate of the task solving process.

B. Agricultural Data Management Systems

The digital revolution in agriculture has paved the way for advanced data management platforms specifically designed to tackle the unique challenges and demands of agricultural research and production. Several notable platforms have emerged in related sectors.

CyVerse [34], initially developed for plant genomics, has expanded into a comprehensive platform, offering life scientists robust computational infrastructure to manage large datasets and conduct complex analyses. As a cloud-based solution, CyVerse supports collaborative efforts, enabling seamless sharing of data and tools. Its focus on scalability and interoperability has made it a crucial tool for modern agricultural data management. GARDIAN [5], the CGIAR’s (Consultative Group for International Agricultural Research) data discovery platform, allows users to find research datasets and publications from various CGIAR Centers and Programs. By aggregating data from multiple agricultural research projects, GARDIAN offers a holistic view of global agricultural research, fostering interdisciplinary collaboration. GEMS [4] integrates genomic, phenotypic, and environmental data, offering tools for analyzing and visualizing large-scale agricultural datasets, thereby linking genomics with environmental factors in agricultural research. TERRAREF [6] focuses on delivering high-resolution sensor data for plant research, using tools such as cameras and drones to gather detailed information on plant growth, health, and environmental conditions. This extensive dataset aids researchers in exploring the intricate relationships between plants and their environments. The evolution of Smart Farming into Agriculture 5.0 [35] emphasizes the critical role of data in optimizing farm management for sustainability and economic efficiency. Lastly, ADMA [7] introduces an innovative agricultural data management system designed according to FAIR principles, incorporating features such as intelligence, interactivity, scalability, flexibility, open-source pipeline management, and enhanced privacy and security.

Several platforms and studies emphasize the importance of integrated data management systems for addressing specific

agricultural topics or utilizing advanced technologies. For instance, Agricultural Remote Sensing Big Data [36] explores the management and application of remote sensing big data in agriculture, introducing a four-layer-twelve-level (FLTL) framework designed to handle remote sensing data for precision farming and local farm analyses. Data Warehouse and Decision Support on Integrated Crop Big Data [37] presents a thorough analysis of data warehousing and decision support systems, demonstrating their capability in leveraging integrated crop big data through the development of a continental-level Agricultural Data Warehouse (ADW). Big Data and Machine Learning (ML) With Hyperspectral Information in Agriculture [38] reviews critical research utilizing big data, machine learning, and deep learning, with a specific focus on processing hyperspectral and multispectral agricultural data. It highlights the potential of ensemble machine learning and scalable parallel discriminant analysis in managing agricultural big data. Lastly, Blockchain for Sustainable e-Agriculture [39] discusses the intersection of blockchain technology and IoT in e-agriculture, emphasizing how blockchain can enhance agricultural value chains, strengthen IoT networks, and ensure data validation, security, and privacy.

III. FRAMEWORK

To support the FAIR principles and the proposed key features such as intelligence, autonomy, efficacy, efficiency, extensibility, flexibility and privacy, we design and implement a proof-of-concept system called Agricultural Data Management and Analytics Copilot (ADMA Copilot).

A. Shift of Paradigm

Thanks to the emerging reasoning and tool learning ability of large language models, there will be a shift of agricultural data management paradigm. In the traditional paradigm, researchers need to understand every details of the data and tools (or ML models) at their hands, such as location, semantics and format of data, and the location, semantics and the interface of the tools (or ML models). When conducting data management, researchers need to manipulate the data, tools or ML models in person. So we can call the traditional paradigm as user-driven paradigm. On the other hand, in the new data management paradigm, an LLM-based copilot will take over to manipulate the data management pipeline, and the researchers will only send high-level requests to the copilot to initiate the pipeline and get involved when required. We can call the new paradigm as AI-driven paradigm. Fig 1 illustrate the shift of agricultural data management shift.

B. System Components

There are two types of components in our framework: internal components and external components. Internal components are the integral parts of the ADMA Copilot, including a copilot server, three LLM-based Agents, a meta-program graph, a data&tool registry. The ADMA Copilot talks with external components such as the original ADMA system, sensors, external cloud storage and computing services, external tools

& APIs, and finally, a web front-end. Figure 2 illustrates the the main components ADMA Copilot and their interactions. In The following subsections, we'll introduce each main components.

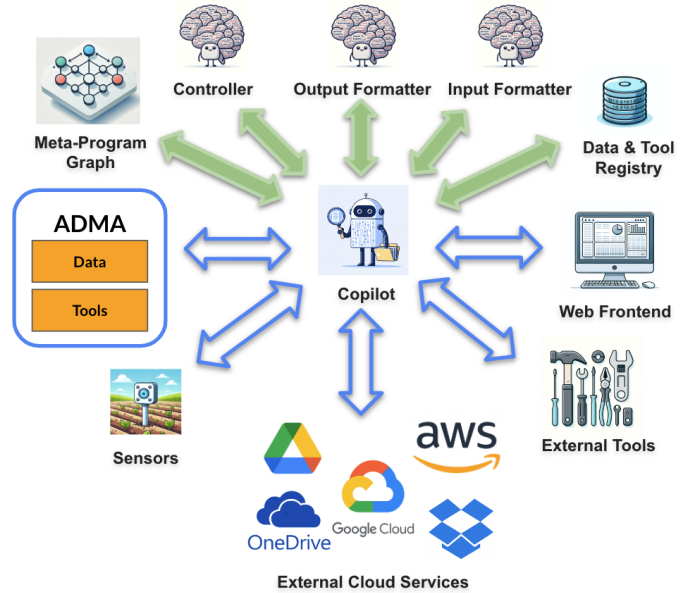


Fig. 2: Main Components of ADMA Copilot

1) *Copilot Server*: The copilot sever is the hub of the whole system, which is responsible for coordinating other components. The copilot server receives instructions from the web front-end and output the results to the front-end. When receiving a task, the copilot server will search in its data & tool registry for relevant tools to use and build a meta-program graph. The copilot server also talks to three LLM-based agents and meta-program graph, to decide the tools to call in each step and to format the input/output of the meta-program. When the tool to call is determined, the copilot server will reach out to the specified tool, execute the call, and retrieve the results from the tool.

The design principle of the copilot server is to decouple the programmatic part from the "intelligent" part, by only keeping the programmatic part within its own logic and off-loading the intelligent part to the three LLM-based agents. In this way, the logic of copilot server can be kept at its minimal.

2) *LLM-based Agents*: There are three LLM-based agents: a program controller, an input formatter and an output formatter, all of which will accept current task, meta-program graph and the execution history as input. The program controller is responsible for determining the control flow of the program, namely, which tool to execute in next step, by looking into the status of current meta-program graph and taking into account of the task to accomplish. The program controller has the highest requirements for the reasoning capability of the underlying LLM, because it needs to analyze current task and conduct the planning accordingly, and for some complicated tasks, the planning always requires high reasoning capability. The other two LLM-based agents, input formatter and output

IV. EVALUATION

In this section, we demonstrate the use cases of ADMA Copilot to prove its core features such as intelligence, autonomy, efficacy, efficiency, extensibility, flexibility and privacy.

A. Interface

ADMA Copilot has a conversational interface, where the user can type in the instruction into the chat box at the bottom of the web page. Then the results will be displayed above the chat box. All the chat history within a session will be preserved. In Figure 4, when the user asks "go to ADMA", then the main page of ADMA will be displayed. ADMA Copilot will also display the steps it has gone through.

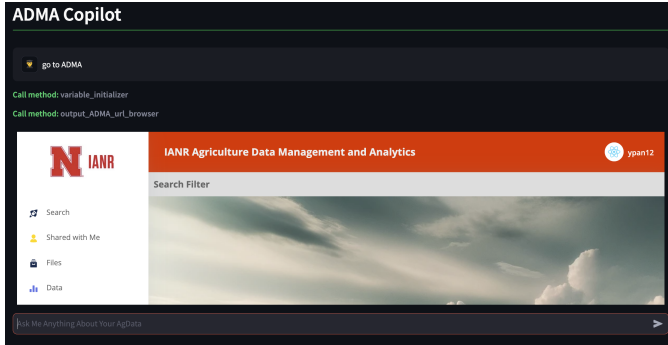


Fig. 4: A Demo for Conversational UI of ADMA Copilot

B. Intelligence

ADMA Copilot can understand user intent in natural language. Sometimes, even the user input a fuzzy or somewhat ambiguous instruction, the copilot will try its best to interpret the user's intent. Figure 5 illustrate a case in which the user type in the instruction: "I want to know how to use ADMA." Then the copilot will display the documentation page of the ADMA because this is the result which is most appropriate in this context.

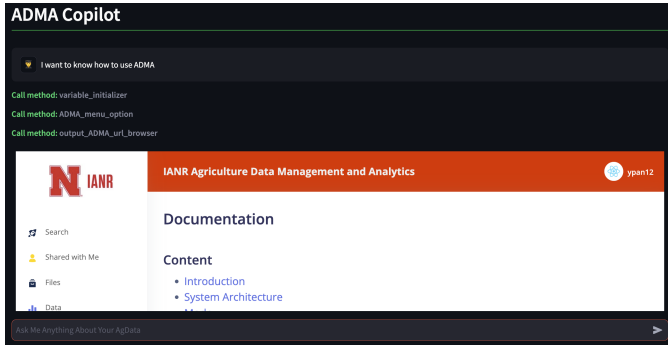
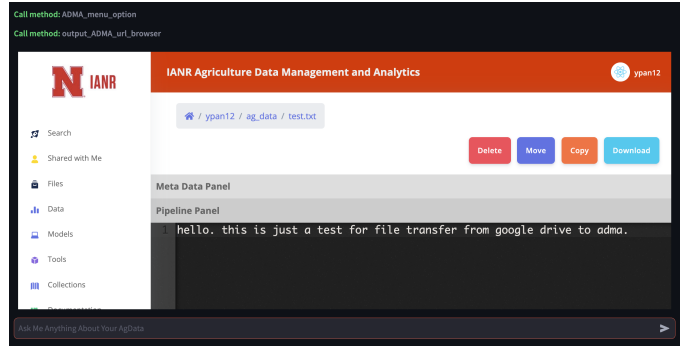


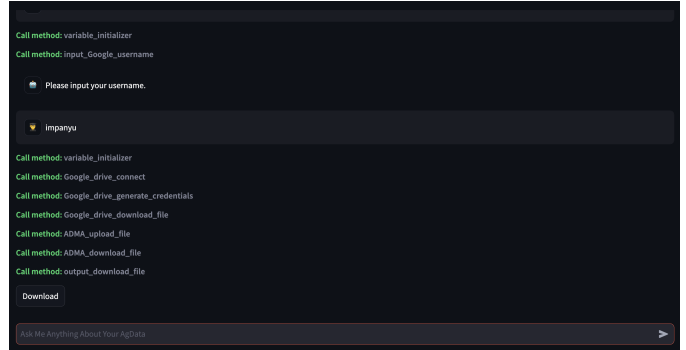
Fig. 5: A Demo for User Intent Understanding

C. Autonomy

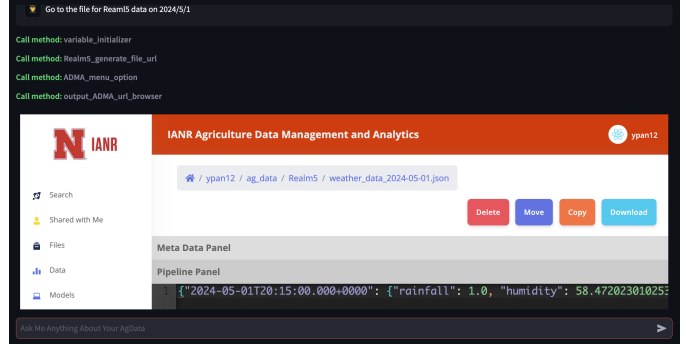
One of the key features of ADMA copilot is autonomy. It is designed to try its best to accomplish the tasks proposed by users autonomously, except extra user input is required. For



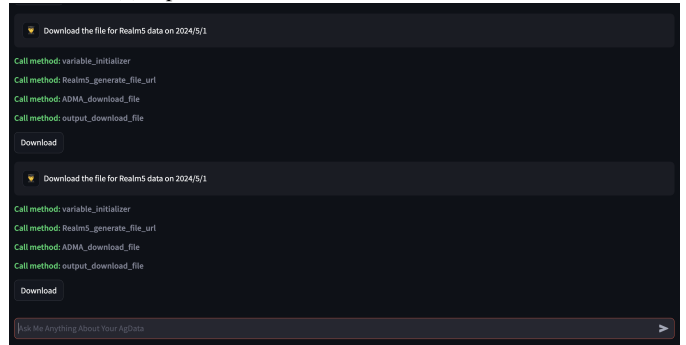
(a) Transfer File from Google Drive to ADMA and Open the Page



(b) Transfer File from Google Drive to ADMA and then Download



(c) Open the file for Realm5 Data on 2024/5/1




(d) Download the file for Realm5 Data on 2024/5/1

Fig. 6: Demo for Autonomous Task Completion

each task, the program controller of the copilot will route on the meta-program graph and complete several steps until the task is finally accomplished.

Figure 6 demonstrates some use cases of autonomous task



(a) Ask for User Input when Unsure



(b) Display the Result after Further Input

Fig. 7: Demo for Efficacy

completion. In Figure 6a, the user type in "Transfer the adma_test/test.txt on my google drive to my ADMA root folder, and open the uploaded file". The copilot decomposes the task into several steps, such as call the input formatter, authenticate the user for Google drive, download the file from Google Drive, upload the file to ADMA, then open the page of the uploaded file. In Figure 6b, the user type in "Transfer the adma_test/test.txt on my google drive to my ADMA root folder, and then download the uploaded file", similarly the copilot accomplish the task by first downloading the file from Google drive and then uploading it to ADMA, followed by popping up a download button for the user to click. In Figure 6c, the user types in "go to the file for Realm5 data on 2024/5/1", again the copilot decompose the task into several steps: initializing the date and menu name, generating the path for the file on ADMA, generating the url for the file on ADMA, displaying the web page for the file. In Figure 6d, the user asks "download the file for Realm5 data on 2024/5/1", the copilot first generates the url for the file on ADMA as previous task, and then call download API of ADMA to download the file, followed by popping up a download button.

Through these use cases, we can see ADMA Copilot has the basic capability to understand user's intent and complete tasks automatically, with minimal user interference.

D. Efficacy

The design of our copilot reduces the phenomena of hallucination of LLM to its minimal, by decoupling data flow from control flow. That is, when generating the calling interface of any tool, instead of asking LLM-based controller to do so, we use the value of the variable within meta-program graph. If there are not enough information, the copilot will ask for further input from the user, rather than fabricate some plausible value. This will make sure the predictability of the final results.

Figure 7 illustrates a case when the user inputs "plot temperature, humidity and wind speed" in 7a, since the copilot is not sure the time span of the data to draw, it will ask for further input: "Please input a data string for Realm5.". After the user provide the date "2024/5/1", the copilot will retrieve the relevant data and plot it to the user as in 7b.

E. Efficiency

Autonomy not only brings about convenience, but also efficiency. For some complicated tasks which contain several steps, the copilot demonstrates its efficiency compared with human users. Here we conduct a comparison between ADMA copilot and human user, in terms of task completion time for several data management tasks.




Fig. 8: Comparison of Task Completion Time between Copilot and Human Users.

In Figure 8, we choose 5 tasks as follows:

- "Go to directory 1 under root folder of ADMA",
- "Check the meta data of calculate_ndvi.py on ADMA",
- "Open the page on ADMA, the name of which contains the keyword soil",
- "Download the file for Realm5 data on 2024/5/1",
- "Transfer the adma_test/test.txt on my google drive to my ADMA root folder, and open the uploaded file"

The bars represent the average task completion time for copilot and human user, for each task. The green curve represents the number of steps the copilot has gone through for each task, which provides some clue about the complex level for a task. We can see for more complicated tasks, both the copilot and human user take longer time, however, the copilot outperforms human user in all tasks, and also the copilot keeps the completion time almost constant, regardless of the increasing complex level of tasks.

F. Extensibility

To accommodate more tools, ADMA Copilot is designed to be able to incorporate new tools which will be incorporated in the future. Given a documentation of any new API or function interface, the only thing we need to do is to plug in the tool into the meta-program graph and inform the copilot how to use the new tools. Then all things left will be handled to the

copilot without any changes as before. The copilot will make calls to the new tools when required. For instance, when we design a new user interface which can draw a map for the fields on our extension farm, we can just plug in the description of the UI tool into meta-program graph, and then when the user inputs "I want to see the map for the field named 1863N", the output will be like in Figure 9

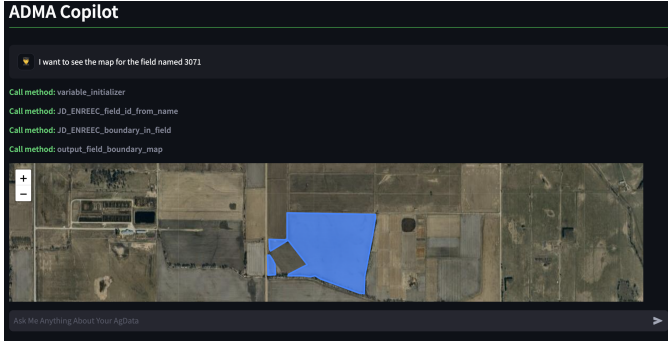


Fig. 9: Demo of Drawing a Map for a Field after Plugging in the Map UI

G. Flexibility

Thanks to the reasoning capability of the underlying LLM, our copilot demonstrates some degree of flexibility in terms of understanding the user’s intent based on a wide spectrum of user input. The user can be free from precisely phrasing the instruction, and the user input can be fuzzy or even ambiguous. Nevertheless, the copilot will try its best to interpret the user’s instruction and execute the task accordingly.


In Figure 10a, the user inputs "download the file adma_test/test.txt on Google drive", then the copilot will download the specified file and display a download button on the interface. In Figure 10b, the user types in "I need to check the file adma_test/test.txt on Google drive", the copilot can get the same intent as previous input and still display a download button on the interface.

H. Privacy and Authentication


ADMA Copilot will protect the user’s data and privacy, by ask for user’s credentials to authenticate when the user need to get access to protected services or data. In Figure 11a, the user want to "list my Google drive root folder", then the copilot decides it is a protected resource and ask for the user to go through the authentication process for Google drive. In Figure 11b, the user types in "download the Realm5 data on 2024/5/1 on ADMA", then the copilot decides this is also a protected data and asks the user to input ADMA token.

I. Comparison

By adopting similar evaluation frameworks as in [7], we compare ADMA Copilot with the existing agriculture platforms. The evaluation criteria and corresponding explanation are listed in Table I. These criteria cover a broad range of dimensions, from the basic requirements for a typical

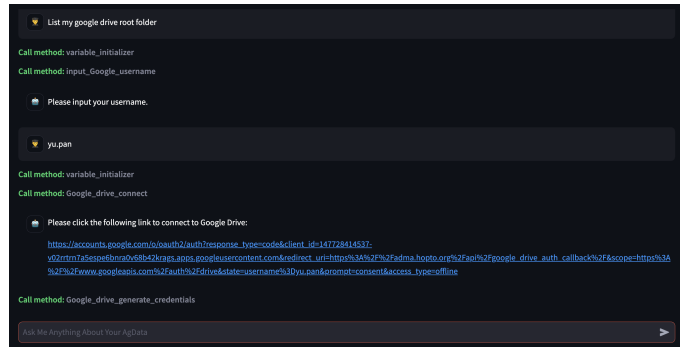


(a) Download A File from Google Drive

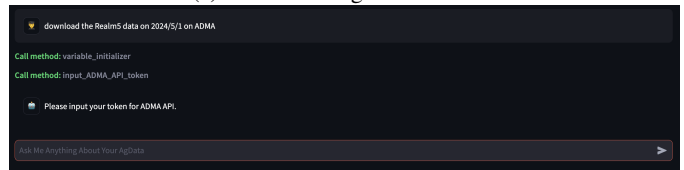


(b) Check A File from Google Drive

Fig. 10: Demo of Flexible User Instruction



(a) Ask for Google Credentials



(b) Ask for ADMA Token

Fig. 11: Demo of Privacy Protection

data management platform to more innovative features, which also take into consideration the principles of FAIR. The eleven dimensions of criteria are, namely, privacy and security, CRUD repository, interoperability, extensibility, data model, data granularity, analytical tools, tractability, ML Models, and Intelligence and Autonomy.

We make qualitative comparison between ADMA Copilot with existing agricultural data management platforms including CyVerse, GARDIAN, GEMS, TERRAREF and ADMA. Table II shows the comparison result. We mark × when a platform satisfies certain criteria and leave the cell empty otherwise.

ADMA Copilot meets all the criteria under consideration. While this comparison is qualitative and may introduce bias due to the selection and specific definition of evaluation

criteria, the results suggest that our platform excels in the chosen dimensions. This indicates that the design principles of our system address a broader range of potential user needs and interests, positioning it well for future requirements.

V. CONCLUSION

In this paper, we propose and explore the idea of a LLM based multi-agent copilot for autonomous agricultural data management. Based on our previous developed platform: Agricultural Data Management and Analytics, we build a system called ADMA Copilot, which can understand user's intent, makes plan for data processing pipeline and accomplishes the task automatically, in which three agents: a LLM based controller, an input formatter and an output formatter collaborate together. Different from existing LLM based solutions, by defining a meta-program graph, our work decouples control flow and data flow to enhance the predictability of the behaviour of the agents. As a proof-of-concept product, ADMA Copilot can help to facilitate the shift of agricultural data management paradigm from traditionally human-driven paradigm to AI-driven paradigm. Experiments demonstrates the intelligence, autonomy, efficacy, efficiency, extensibility, flexibility and privacy of our system. Comparison is also made between ours and existing systems to show the superiority and potential of our system.

ACKNOWLEDGMENT

This research is funded by partnership on Data Innovation (PDI), Agricultural Research Services (ARS), United States Department of Agriculture (USDA) and the Nitrogen Research for Agriculture Transformation and Enhancement (NRATE) project, ARS, USDA. This work was completed utilizing the Holland Computing Center of the University of Nebraska, which receives support from the UNL Office of Research and Economic Development and the Nebraska Research Initiative.

REFERENCES

- [1] N. Chamara, M. D. Islam, G. F. Bai, Y. Shi, and Y. Ge, "Ag-iot for crop and environment monitoring: Past, present, and future," *Agricultural systems*, vol. 203, p. 103497, 2022.
- [2] G. Bai, Y. Ge, D. Scoby, B. Leavitt, V. Stoerger, N. Kirchgessner, S. Irmak, G. Graef, J. Schnable, and T. Awada, "Nu-spidercam: A large-scale, cable-driven, integrated sensing and robotic system for advanced phenotyping, remote sensing, and agronomic research," *Computers and Electronics in Agriculture*, vol. 160, pp. 71–81, 2019.
- [3] T. L. Swetnam, P. B. Antin, R. Bartelme, A. Bucksch, D. Camhy, G. Chism, I. Choi, A. M. Cooksey, M. Cosi, C. Cowen *et al.*, "Cyverse: Cyberinfrastructure for open science," *bioRxiv*, pp. 2023–06, 2023.
- [4] S. D. Senay, G. C. Shurson, C. Cardona, and K. A. Silverstein, "Big data, data privacy, and plant and animal disease research using gems," *Agronomy Journal*, vol. 114, no. 5, pp. 2644–2652, 2022.
- [5] I. Jouini, D. Leibovici, M. Kurian, E. Bongcam-Rudloff, J. Kieti, and S. Sellitti, "Evaluation of cgjar platform for big data in agriculture," 2021.
- [6] D. LeBauer, M. A. Burnette, J. Demieville, N. Fahlgren, A. N. French, R. Garnett, Z. Hu, K. Huynh, R. Kooper, Z. Li *et al.*, "Terra-ref, an open reference data set from high resolution genomics, phenomics, and imaging sensors," 2020.
- [7] Y. Pan, J. Sun, H. Yu, G. Bai, Y. Ge, J. Luck, and T. Awada, "Transforming agriculture with intelligent data management and insights," in *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 2023, pp. 3489–3498.
- [8] C. Qu, S. Dai, X. Wei, H. Cai, S. Wang, D. Yin, J. Xu, and J.-R. Wen, "Tool learning with large language models: A survey," *arXiv preprint arXiv:2405.17935*, 2024.
- [9] Z. Shen, "Llm with tools: A survey," *arXiv preprint arXiv:2409.18807*, 2024.
- [10] G. Mialon, R. Dessi, M. Lomeli, C. Nalmpantis, R. Pasunuru, R. Raileanu, B. Rozière, T. Schick, J. Dwivedi-Yu, A. Celikyilmaz *et al.*, "Augmented language models: a survey," *arXiv preprint arXiv:2302.07842*, 2023.
- [11] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders *et al.*, "Webgpt: Browser-assisted question-answering with human feedback," *arXiv preprint arXiv:2112.09332*, 2021.
- [12] Y. Zhuang, Y. Yu, K. Wang, H. Sun, and C. Zhang, "Toolqa: A dataset for llm question answering with external tools," *Advances in Neural Information Processing Systems*, vol. 36, pp. 50 117–50 143, 2023.
- [13] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [14] W. Fan, Y. Ding, L. Ning, S. Wang, H. Li, D. Yin, T.-S. Chua, and Q. Li, "A survey on rag meeting llms: Towards retrieval-augmented large language models," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024*, pp. 6491–6501.
- [15] J. He-Yueya, G. Poesia, R. E. Wang, and N. D. Goodman, "Solving math word problems by combining language models with symbolic solvers," *arXiv preprint arXiv:2304.09102*, 2023.
- [16] Z. Gou, Z. Shao, Y. Gong, Y. Shen, Y. Yang, M. Huang, N. Duan, and W. Chen, "Tora: A tool-integrated reasoning agent for mathematical problem solving," *arXiv preprint arXiv:2309.17452*, 2023.
- [17] L. Gao, A. Madaan, S. Zhou, U. Alon, P. Liu, Y. Yang, J. Callan, and G. Neubig, "Pal: Program-aided language models," in *International Conference on Machine Learning*. PMLR, 2023, pp. 10 764–10 799.
- [18] X. Wang, Y. Chen, L. Yuan, Y. Zhang, Y. Li, H. Peng, and H. Ji, "Executable code actions elicit better llm agents," *arXiv preprint arXiv:2402.01030*, 2024.
- [19] A. M. Bran, S. Cox, O. Schilter, C. Baldassari, A. D. White, and P. Schwaller, "Chemcrow: Augmenting large-language models with chemistry tools," *arXiv preprint arXiv:2304.05376*, 2023.
- [20] W. Zhang, L. Zhao, H. Xia, S. Sun, J. Sun, M. Qin, X. Li, Y. Zhao, Y. Zhao, X. Cai *et al.*, "A multimodal foundation agent for financial trading: Tool-augmented, diversified, and generalist," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024*, pp. 4314–4325.
- [21] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [22] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, and Y. Zhuang, "Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [23] Z. Shi, S. Gao, X. Chen, Y. Feng, L. Yan, H. Shi, D. Yin, Z. Chen, S. Verberne, and Z. Ren, "Chain of tools: Large language model is an automatic multi-tool learner," *arXiv preprint arXiv:2405.16533*, 2024.
- [24] Y. Zhuang, X. Chen, T. Yu, S. Mitra, V. Bursztyrn, R. A. Rossi, S. Sarkhel, and C. Zhang, "Toolchain*: Efficient action space navigation in large language models with a* search," *arXiv preprint arXiv:2310.13227*, 2023.
- [25] Y. Ge, W. Hua, K. Mei, J. Tan, S. Xu, Z. Li, Y. Zhang *et al.*, "Openagi: When llm meets domain experts," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [26] Y. Qin, S. Liang, Y. Ye, K. Zhu, L. Yan, Y. Lu, Y. Lin, X. Cong, X. Tang, B. Qian *et al.*, "Toollm: Facilitating large language models to master 16000+ real-world apis," *arXiv preprint arXiv:2307.16789*, 2023.
- [27] S. Qiao, H. Gui, C. Lv, Q. Jia, H. Chen, and N. Zhang, "Making language models better tool learners with execution feedback," *arXiv preprint arXiv:2305.13068*, 2023.
- [28] X. Liu, Z. Peng, X. Yi, X. Xie, L. Xiang, Y. Liu, and D. Xu, "Toolnet: Connecting large language models with massive tools via tool graph," *arXiv preprint arXiv:2403.00839*, 2024.
- [29] Y. Song, W. Xiong, D. Zhu, W. Wu, H. Qian, M. Song, H. Huang, C. Li, K. Wang, R. Yao *et al.*, "Restgpt: Connecting large language models with real-world restful apis," *arXiv preprint arXiv:2306.06624*, 2023.

TABLE I: Dimensions of Evaluation Criteria

Dimension	Explanation
Privacy&Security	The system protects data privacy and security, including but not limited to: 1. the ability to separate data and metadata; 2. the ability to share data or tools with other users safely; 3. the ability to differentiate public and private data; 4. the ability to run user-defined programs safely; 5. unified authentication module; 6. data encryption.
CRUD Repository	Data can be created (uploaded), read (rendered), updated, and deleted from/to/on the repository.
Interoperability	Data from external sources of different formats or various modalities can be added or incorporated into the system and managed by the system without difficulties.
Extensibility	New capabilities and functionalities can be added to the system.
Data Model	Extract, transform and load procedure employed to populate a relational, graph, object, or vector database governed by an explicit data model.
Data Granularity	Ability to index and query for sub-file level entities such as individual records within datasets.
Analytical Tools	Provides tools for users to clean text, numeric, and geospatial outliers, as well as conduct unique analysis.
Trackability	Ability to keep track of different operations on each file, maintain and visualize the pipeline record accordingly.
ML Models	Provide ready-to-use or support training and hosting of agricultural models.
Intelligence	The system utilizes Artificial Intelligence to facilitate data management, analysis, and interaction.
Autonomy	The system demonstrate certain degree of autonomy, by accomplishing data management pipeline with minimal interference or guiding from human users.

TABLE II: Comparison of Agriculture Data Management Platforms across a Variety of Criteria

Platform	Privacy& Security	CRUD Repository	Interoperability	Extensibility	Data Model	Data Granularity	Analytical Tools	Tractability	ML Models	Intelligence	Autonomy
Cyverse	x	x	x	x		x	x		x		
GARDIAN	x	x	x	x			x				
GEMS	x	x	x	x			x				
TERRAREF	x	x	x	x	x	x	x				
ADMA	x	x	x	x	x	x	x	x	x	x	
ADMA Copilot	x	x	x	x	x	x	x	x	x	x	x

[30] R. Yang, L. Song, Y. Li, S. Zhao, Y. Ge, X. Li, and Y. Shan, "Gpt4tools: Teaching large language model to use tools via self-instruction," *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[31] L. Li, Y. Chai, S. Wang, Y. Sun, H. Tian, N. Zhang, and H. Wu, "Tool-augmented reward modeling," *arXiv preprint arXiv:2310.01045*, 2023.

[32] OpenAI, "Openai function calling guide," 2023. [Online]. Available: <https://platform.openai.com/docs/guides/function-calling>

[33] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom, "Toolformer: Language models can teach themselves to use tools," *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[34] N. Merchant, E. Lyons, S. Goff, M. Vaughn, D. Ware, D. Micklos, and P. Antin, "The iplant collaborative: cyberinfrastructure for enabling data to discovery for the life sciences," *PLoS biology*, vol. 14, no. 1, p. e1002342, 2016.

[35] V. Saiz-Rubio and F. Rovira-Más, "From smart farming towards agriculture 5.0: A review on crop data management," *Agronomy*, vol. 10, no. 2, p. 207, 2020.

[36] Y. Huang, Z.-x. Chen, Y. Tao, X.-z. Huang, and X.-f. Gu, "Agricultural remote sensing big data: Management and applications," *Journal of Integrative Agriculture*, vol. 17, no. 9, pp. 1915–1931, 2018.

[37] V. M. Ngo, N.-A. Le-Khac, and M.-T. Kechadi, "Data warehouse and decision support on integrated crop big data," *International Journal of Business Process Integration and Management*, vol. 10, no. 1, pp. 17–28, 2020.

[38] K. L.-M. Ang and J. K. P. Seng, "Big data and machine learning with hyperspectral information in agriculture," *IEEE Access*, vol. 9, pp. 36 699–36 718, 2021.

[39] K. Dey and U. Shekhawat, "Blockchain for sustainable e-agriculture: Literature review, architecture for data management, and implications," *Journal of Cleaner Production*, vol. 316, p. 128254, 2021.

[40] John Deere, "John deere: Agricultural and construction equipment," 2024. [Online]. Available: <https://www.deere.com/en/index.html>

[41] RealmFive, "Realmfive: Agriculture iot solutions," 2024. [Online]. Available: <https://realmfive.com/>