FarmWorks: Decentralized AI Agents for Personalized Solutions

Vladimir Baulin (ORCID: 0000-0003-2086-4271)
Alex Vyatkin (ORCID: 0000-0003-1306-4620)
Avel GUÉNIN—CARLUT (ORCID: 0000-0001-8239-7264)
Daniel Friedman (ORCID: 0000-0001-6232-9096)
John Bolt (ORCID: 0009-0002-0878-5546)
Stefan Falkenstein (ORCID: 0009-0008-5054-4969)
Parishrut Jassal (ORCID: 0009-0008-4351-805X)
Celio Trois (ORCID: 0000-0002-7386-9749)
Jonathan Minchin (ORCID: 0000-0001-8152-1796)

v1 ~ September 12, 2024.

Additional information (budget and more) and all contacts: blanket@activeinference.institute

Problem Statement

Climate change intensifies agricultural challenges, requiring more and more advanced technological solutions. Small farmers increasingly rely on technical assistance, which is becoming centralized, dominated by large agricultural corporations and governments imposing sophisticated pre-designated solutions. As AI proliferates within these centralized solutions, diseases mitigation methods, climate credits, government subsidies, and regulations risk monopolizing farmers' activities. This tendency, amplified by AI development, threatens to undermine farmers' autonomy and limit their ability to make independent decisions, converting them into consumers of centralized technological solutions.

Project Objectives

We propose to develop **FarmWorks** — a platform for human-AI interaction in agriculture that enables personalized, farm-scale solutions while resisting power concentration associated with centralized AI systems. FarmWorks addresses the above challenges by providing an open source decentralized AI-powered agricultural platform that empowers individual farmers with cutting-edge technology while preserving their autonomy and promoting sustainable practices. By integrating real-time data collection, edge computing, and Active Inference models, FarmWorks enables farmers to make informed decisions tailored to their specific contexts (for example, integrating humidity data and epidemiological models to assist farmers with remediative and anticipatory treatments for mold).
The platform's community-driven approach fosters knowledge sharing and collective innovation, while its ethical framework ensures that AI deployment aligns with human values and ecological sustainability. FarmWorks thus offers a comprehensive solution that not only enhances agricultural productivity but also promotes environmental stewardship and farmer empowerment in the face of global agricultural challenges.
Key components of Farmworks include:

1. Decentralized Data Collection and Processing
   - Customizable on-site sensor networks for real-time data gathering
   - Local data integration and AI model training to preserve privacy
2. AI-Assisted Decision Support
   - Virtual assistant for natural language interaction with farmers
   - Active Inference models for personalized, context-aware recommendations
3. Community-Driven Innovation Network
   - Social platform for farmers to share strategies and insights
   - Collective problem-solving and best practice identification
4. Ethical AI Deployment
   - Community-developed ethical framework

Transparent AI decision-making processes
Balancing AI assistance with human autonomy

5. Promoting Sustainable Agriculture
   Optimizing decentralized farm management through data-driven insights
   Supporting transition to regenerative practices

This approach addresses power concentration by centralized AI via:

Keeping data and decision-making power with individual farmers
Fostering a collaborative network and peer-to-peer technological support rather than relying on centralized solutions from governments and corporations
Ensuring transparency and ethical considerations in AI deployment
Empowering farmers with personalized AI tools while preserving their agency and creativity

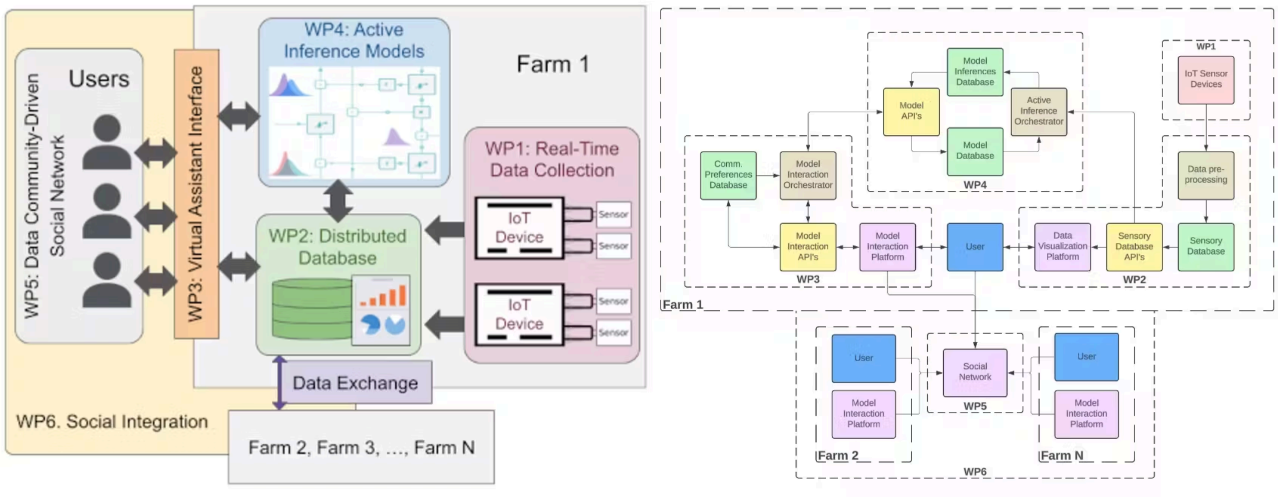

FarmWorks can serve as a model for decentralized AI applications in other fields, demonstrating how to harness AI's benefits while mitigating risks of power concentration.

Implementation

FarmWorks establishes a platform for human-AI interaction in agriculture, where local data drive personalized solutions tailored to each individual farm. This approach offers multiple strategic options and counteracts the risk of power concentration inherent in centralized AI systems, ensuring that farmers not only retain control and autonomy over their agricultural practices, but become co-creators of personalized solutions that can be shared to other farmers. Further, this approach can be replicated across various fields, such as decentralized science and strategies to mitigate the accumulation of power in broader contexts.
We present FarmWorks with its setup of sensors and the community of interconnected farmers with the creation of agents for predicting diseases, crop harvest predictions, etc. It is important to realize that this proposal can be replicated in other fields, showcasing the broader potential of this approach in terms of digital trust of AI-empowered humans.
We aim to augment this decision power of farmers in the setting of decentralized AI and farm-scale sensemaking and decision-making. We aim to support farmer practices through data-guided decisions and consensual knowledge sharing (rather than reliance on centralized AI systems). A key objective is to align farmer expectations and requirements with an advanced recommendation system based on Active Inference principles. This system takes into consideration harvest predictions, meteorological factors, and scientific disease models, creating an iterative feedback loop provided by Active Inference. The FarmWorks development and implementation will be divided into six Work Packages (WP), described below, whose iterations are shown in Figure 1.

WP1: Real-Time Data Collection Module

WP1 is focused on the integration and deployment of a fully customizable network of on-site sensors and cameras that collect critical environmental and crop data and adapting to specific field conditions. This network is based on developments of open-source projects Vin-Q and ROMI. The data collected includes soil moisture, temperature, pH, conductivity, humidity. The system is designed to operate under any conditions, continuously collecting and transmitting vital agricultural information to a local server autonomously, independent of connectivity or electric power. The sensors are connected to the local data server via communication technologies including 4G, 5G and LoRa networks working simultaneously. The flexibility of these communication protocols allows the system to function effectively even in remote areas with limited connectivity. LoRa, in particular, offers long-range, low-power communication, making it ideal for continuous data transmission across large fields or challenging terrains. Furthermore, the sensors and gateways are powered with solar panels to avoid dependence on power cuts. By focusing on decentralization in in-field processing, WP1 reduces dependency on centralized solutions such as cellular networks and closed-source AI models, providing monitoring system sthat continuously collects data in the fields.
Deliverables: 1. Design of affordable network of customized sensors that can communicate with a gateway via 4G-5G connection and LoRa network. 2. A fully operational real-time data monitoring and alerting system that integrates with IoT edge devices and securely transmits and processes data from diverse sources. 3. Installation in Vin-Q farmers pilot fields (20 fields in Catalonia region). **Key persons**: Jonathan Minchin, Vladimir Baulin

WP2. Decentralized Data Integration and Processing Module

The FarmWorks' platform foundation is built on the automatic real-time data collection, processing, and analysis of data from field sources. Each farm stores locally collected data from several boards and it is equipped to handle in-field analytics capable of running AI and machine learning models. Apart from avoiding centralized power accumulation, such setup reduces the need for centralized data storage and processing, by localizing data processing.
Decentralized data-processing modules will be created based on the Vin-Q platform, as it offers consistent integration of various open source tools: Apache Superset, Zeppelin, Kafka, Spark, integrating data from different sources such as traditional databases (e.g. PostgreSQL, MySQL, SQLite) as well as distributed (e.g. Apache Druid, Hive). This data is curated and organized into a Distributed Database, which forms the foundation for advanced analysis *in-situ*. It is also designed in such a way that a curated database that is needed for decision support is composed from preregistered datasets.
The module operates on a decentralized model, where data collection and processing will be done in an instance running locally at each farm, enhancing privacy and security. This decentralized approach allows AI agents to read on local data and contribute updates to a global model without sharing raw data. The decentralized architecture ensures that all training data remains on each farm, with only model updates that can be optionally shared with other users to combine in a larger model. This method not only protects data privacy but also improves the efficiency and accuracy of the models by leveraging local, highly personalized context-specific information tailored by the intentions and the needs of the user.
The module includes a visual interface that enables users to interact with, explore, and visualize data through custom dashboards. These dashboards serve as a centralized hub for accessing real-time and historical information about farm conditions, soil health, and more. Users can quickly identify trends, anomalies, and potential issues using visual tools like graphs, calendar, charts, and maps. The interface allows farmers to make data-driven decisions about irrigation, fertilization, and other practices, ultimately enhancing sustainability and productivity. This module will provide data availability for customized AI agents to train and infer the decisions based on highly adaptable Active Inference models.
Deliverables: 1. Integration of real-time data from various sources, such as IoT sensors, weather updates, and user inputs (month 1 - 10). 2. Make data available for consumption by bots and Active Inference models (month 10 - 13). 3. A user-friendly AI-powered interface with customizable dashboards that provide

real-time and historical data visualization, supporting informed decision-making in regenerative agriculture (month 13-24). **Key persons:** Celio Trois, Vladimir Baulin.

WP3. Virtual Assistant for AI-human Interaction

The platform integrates virtual assistants designed to facilitate seamless interaction between farmers and AI, making complex data and insights more accessible and provide a control of farmer of the agents seemingly. The virtual assistant functions realized as a chat bot (e.g. Telegram or other platform) with an intuitive interface that allows farmers to engage with the system in a natural, conversational manner. Through the farmer information available in the personalized database that keeps all the records of the farmer fields , optionally using embeddings of LLMs, this assistant is trained on the specific data of each farm, including soil analysis and sensor information, enabling it to provide essential feedback loop between the system and give personalized recommendations and insights that will change with conditions.
The virtual assistant optionally leverages Open Source Large Language Models (LLMs) installed locally in the platform with the function to interpret natural language queries, allowing farmers to easily interrogate the extensive personalized Knowledge Database. It tailors responses based on the unique conditions of each farm, offering evidence-based practices that have proven successful in similar contexts. Beyond answering questions, the assistant can help farmers visualize current field conditions, predict future developments, and understand the impacts of climate change and weather events on soil health through interacting with AI agents processing the data of the farmer field.
Additionally, the virtual assistant supports the setup of monitoring systems, enabling farmers to establish custom alerts based on specific conditions, such as weather projections or soil moisture levels. This proactive approach allows for timely interventions, helping farmers optimize their practices and make informed decisions while maintaining a sense of digital trust and full control over localized data, without inadvertently sharing their decisions or data with government authorities, big tech firms, or other stakeholders.
Deliverables: 1. A virtual assistant that enables natural language interaction with the AI-driven platform implemented in Telegram, providing personalized insights, recommendations, and monitoring capabilities to support farmers in making informed, data-driven decisions. 2. Integration with decentralized platform and AI agents through API. **Key persons**: Design: Jonathan Minchin, Vladimir Baulin, Technical implementation of the front end of interaction: Stefan Falkenstein, Daniel Friedman

WP4. Active Inference Models

The project seeks to develop generative state space models that leverage Active Inference to elevate the interaction between IoT data, AI agents and farmers. These will include predictive and causal wine-grape epidemiological models, models for recommending interventions to the farmer, and models of other farms. Active Inference models leverage variational free energy and expected free energy minimization to conduct approximate online Bayesian inference and planning. In the context of **FarmWorks**, this means that the AI agents will be designed to perpetually learn from and predict data gathered through soil sensors, weather stations, interactions with the farmer, and interactions with other farm's models via the Social Data Network. In particular, the interaction between farmers and their AI agents forms a feedback loop that is central to the platform's operation. As farmers dynamically evaluate the recommendations provided by the AI and take action, the outcomes of these actions are fed back into the system, further refining the AI's models. This continuous loop of action, observation, and adaptation ensures that the platform becomes more accurate and personalized over time, aligning closely with the specific needs and conditions of each farm. Deep learning and reinforcement learning models may be effective, however can have high computational cost and low interpretability, exacerbating the risks of centralized AI. In contrast, Active Inference generative models implemented in RxInfer.jl can be computationally efficient, interpretable, and operated in a decentralized sense/decision-making setting.
To support the implementation of these Active Inference models, the platform will use and extend RxInfer.jl, an open-source probabilistic programming interface for conducting variational inference. RxInfer is specifically tailored for fast, online inference on edge devices, making it a critical enabling technology for

FarmWorks's decentralized AI. Furthermore, the flexibility of probabilistic programming will allow for the development and composition of a wide variety of proven models, from classic regression, causal modeling, and time series analysis, to modern deep learning methods — all grounded in the data being collected from sensors and farmer interactions.
Deliverables: 1. Development and deployment of the RxInfer.jl generative model library, which will be adapted to the specific needs of agricultural decision-making. 2. Integration with data platform, 3. Integration with Virtual assistant. 4. Interaction between various agents trained on different models or different management styles (connection to social network based on data). Model Library: RxInfer.jl
Key persons: John Bolt, Daniel Friedman, Alex Vyatkin, Stefan Falkenstein

WP5. Social Network based on Data and Community-Driven Innovation

The platform extends its capabilities by fostering a new type of social network integrating data exchange in existing networks of trust. Such social network enables farmers to share their personalized strategies and insights within a community in form of well-defined protocols supported by extensive data, while also keeping a sense of agency on their farming practices and the manner in which they share these practices. This network facilitates a powerful exchange of knowledge, where individual successes can be collectively analyzed, refined, and applied to solve common challenges. By participating in this community, users contribute to a system of swarm intelligence, where collaborative problem-solving becomes a shared endeavor, amplifying the effectiveness of individual actions through collective input.
Our main challenge in that regard is to overcome adoption issues related to the perception of the platform, as potential users find it too removed or even antagonistic from their existing practices. To address this, we propose to use a very simple QR code or NFC-based system, requiring embodied interaction, to transfer permission for data access. Hence, the exchange of data access permission is likely to be perceived as an extension of (rather than a threat to) existing social practices and network of trust, while the system handles the technical data exchange in a transparent manner. Furthermore, this incentivizes farmers to seek peer-to-peer interaction with new people, therefore extending prior networks of trust ; and ensures innovation through the uneven diffusion of data. This system will however require a careful design process to ensure ergonomy, accessibility, and adequacy to local social practices.
A core requirement for the adequacy of this system is to enable a coherent, but flexible coupling between peers' practices, hence enabling both the diffusion and selection of innovative techniques. In other words, the users should be encouraged to keep an attitude of curious attention to the output of FarmWorks, as well as the practices of other farmers, rather than either ignoring or blindly following them. This design principle will likely necessitate a careful design process based on extensive observation of and feedback from potential users. However, we propose to articulate our basic design concept from existing research on social organization and self-identity under Active Inference (see WP6). Overall, the process of community validation will be designed to help propagate and test practices, hence building over time a collective databank of techniques that are proven to work under similar conditions.
Deliverables: 1. The design of social interaction between farmers through AI agents, 2. Creation of a social network platform automatized by AI agents that supports data and solution sharing, enabling swarm intelligence, collaborative problem-solving, and the identification of best practices through community-driven innovation. Key persons: Avel Guénin-Carlut, Jon Bolt

WP6. Social Integration of AI: Design, Ethics, & Sociology

The design of an agency-augmenting, autonomous AI system subtly but critically reframes the line of reflection developed to describe and regulate the use of centralized AI systems. FarmWorks is indeed to be understood as a resource augmenting the power of human cognition, rather than an authoritative system replacing it. Concretely implementing this model while minimizing the potential for abuse raises an array of design, ethical and sociological questions, which we aim to investigate during the project - both to inform our development of FarmWorks, and to leverage this process into a more general insight in the social integration of AI.

The first question we need to tackle is at the very foundation of our project: it is to motivate our decision to take agency augmentation as the core expected quality of our AI system instead of more classical consequentialist or deontological considerations. Although we believe that there exist fundamental reasons to consider the respect of agency as a core tenet of ethics (notably to be found in virtue and enactive approaches to ethics), our main argument here centers on the historical sociology of power. As outlined in the introduction, the concentration of power usually goes with the domination of perspectives aligned with the logic of power (Scott's "view from above"). Following this line of reasoning, we propose to produce academic work on the "longue durée" dynamics of the social integration of AI, addressing how agency-augmenting systems may (or may not) help us prevent some of its potential harmful effects. Furthermore, this project requires a demonstration of whether and how it is possible to implement AI system in human social organization in an agency-respecting manner. To do so, we plan to rely on the Active Inference literature on social organization and collective intelligence. A core intuition of this literature is that human social behavior follows contextual prediction of one's self-model, which is notably scaffolded by social communication and the observation of peer behavior. A key parameter determining the level of human agency in a given social context is the rigidity of the coupling between self- and peer-prediction. Hence, the core design principle we propose to follow is to keep a coherent but flexible coupling between humans users' self-model and FarmWorks, as well as between human users. We propose leverage this design principle to design the interface of FarmWorks (see WP5), while using this process to articulate a more general framework for agency-respecting AI.
Deliverables: 1. Comprehension research of the "longue durée" prospective history of social integration of AI. 2. Comprehensive research studies to analyze the effectiveness of this decentralized ethical model, providing insights and guidelines for fostering transparent and fair coexistence between humans and AI agents in various societal contexts. 3. Implementation of tools and protocols that ensure clear visibility into AI decision-making processes and data usage, including user-friendly dashboards and audit trails.
Key persons: Avel Guénin-Carlut, Parishrut Jassal

Proposed Methodology

Real-Time Data Integration (WP1+WP2)

The platform will integrate real-time data from various sources, such as sensors, weather updates, and user inputs. This data will be transmitted to the platform, where it will be processed and analyzed using Active Inference models developed by the Institute.

User Interfaces (WP3)

Users will have access to intuitive interfaces that allow them to:

Chat/informal interfaces and Data/Analytic/API options.
Customize and tune AI models based on their specific needs and preferences.
Set personal preferences and ethical boundaries for AI assistance.
Visualize and interpret AI decision-making processes for transparency.
Provide feedback and corrections to improve AI performance and alignment.

Active Inference Models (WP4)

The Active Inference models will utilize variational free energy minimization to generate adaptive, context-aware AI agents. These models will be designed to:

Using RxInfer.jl to make generative models which process real-time data from various sources.
Generate dynamic, personalized decision-making scenarios using discourse graphs.
Leverage relevant context to compute near-optimal policies through expected free energy-minimizing algorithms.

Active Inference-driven Dynamic Decision Database (WP2 + WP4) with real-time sensor data (WP1)

The platform will generate a dynamic database of possible decisions using discourse graphs that will:

Contain personalized scenarios and solutions based on real-time data and user inputs.
Be continuously updated through iterative learning from real-time data.
Enable users to choose from a range of optimal solutions guided by the Active Inference models.

Community-Driven Innovation (WP5)

Users can share their personalized strategies within the community, allowing for swarm intelligence, collaborative problem-solving, Valorization of effective strategies by the community.
Development of a social network based on data and solutions, fostering collective intelligence.

Social and Research outcomes (WP6)

Developing qualitative and social science research to contextualize and communicate the work being done, and how it relates to broader efforts to ameliorate risks of centralized AI.

Implementation Plan and Timeline

Year 1: Foundation and Core Development

Months 1-3: Project setup, team assembly, and detailed planning.
Months 3-9: Development of the core Active Inference models and real-time data integration.
Months 9-12: Initial implementation of user interfaces and dynamic decision database.

Year 2: Refinement and Community Building

Months 13-18: Iterative improvement of core components based on testing and feedback.
Months 19-24: Development of community-driven innovation tools and engagement strategies.
Continuous: Growing the user and developer community through workshops, and online events.

Evaluation and Success Metrics (examples)

1. **Technical Metrics**: Performance and scalability of the Active Inference models, Quality and Quantity of models, Accuracy and consistency of data integration and decision-making processes.
2. **Community Engagement**: Number of active users and contributors to the platform, Diversity of participants in the community-driven innovation network.
3. **Adoption and Impact**: Number of farms/users adopting the platform and the variety of applications, Measurable improvements in productivity and sustainability.
4. **Public Perception**: Media coverage and public discourse around decentralized AI, Surveys assessing user trust and satisfaction with the platform.
5. **Policy Influence**: Citations or references to the project in policy discussions and regulatory frameworks, Adoption of project-inspired principles in AI governance initiatives.

Risk Assessment and Mitigation Strategies

1. **Technical Challenges in Integrating Real-Time Data with Active Inference generative models**:
   Mitigation: Leverage our expertise in sensor technology and data integration, and collaborate with leading researchers in these fields.
2. **Privacy and Security Vulnerabilities**:
   Mitigation: Implement rigorous security audits, bug bounty programs, and monitoring
3. **Resistance from Established AI Industry Players**:
   Mitigation: Engage in constructive dialogue, emphasize complementary rather than competitive positioning, and demonstrate unique value propositions.

4. **Slow Adoption Due to Technical Complexity**:
   Mitigation: Prioritize user experience design, provide comprehensive documentation and tutorials, and offer support through community forums.
5. **Ethical Dilemmas in AI Decision-Making**:
   Mitigation: Establish an ethics advisory board, implement transparent decision-making processes, and engage in ongoing public consultations.