



## **Collaborative intelligence: AI-driven decision support for beef production**

Wilson, CA.<sup>1</sup>; Naseem, U<sup>2</sup>; Zhang, Z<sup>3</sup>; Lu, H<sup>3</sup>; Tao, L<sup>3</sup>; Hay, R<sup>1</sup>; Everingham, Y<sup>1</sup>

<sup>1</sup> James Cook University, Townsville, Australia; <sup>2</sup> Macquarie University, Sydney, Australia; <sup>3</sup> The University of Sydney, Sydney, Australia

**Key words:** Beef production; Domain-specific AI; Collaborative intelligence; AI decision support; Agricultural innovation

### **Abstract**

The agricultural industry is confronted with an ever-expanding array of innovations, technologies and scientific knowledge. Despite the availability of these solutions, the knowledge exploration process is time-consuming and cognitively demanding due to the volume and variability of information – this often leads to inconsistent advice, non-adoption, or suboptimal decision-making. BeefVantage aims to bridge this gap by providing a tailored AI-powered decision support system designed specifically for beef producers in Queensland – the first of its kind.

BeefVantage utilises advanced Natural Language Processing (NLP) techniques and fine-tuned Large Language Models (LLMs), focusing on trusted domain-specific data. This approach ensures high relevance and accuracy in its recommendations, moving beyond the limitations of generic AI tools, which often suffer from hallucinations and lack of domain-specific knowledge. The tool is in development, with a Minimum Viable Product (MVP) anticipated six months prior to the congress. The presentation will introduce the groundbreaking concept behind BeefVantage, detailing its initial validation processes, and explore the anticipated transformative impacts of this AI-driven tool on beef production.

BeefVantage represents a significant advancement in the application of generative AI in agriculture, tailored specifically to meet challenges faced by beef producers. By offering real-time, context-sensitive, actionable insights, it supports not only immediate problem-solving but also long-term strategies for sustainable and resilient operations. The scalability of BeefVantage suggests potential applicability to other regional contexts and agricultural sectors, promising broader impacts across rangelands.

Attendees will gain insights into cutting-edge generative AI applications in agriculture, understanding both the technological underpinnings and practical benefits. The session will invite discussion on how the collaborative intelligence of human experts and AI can be harnessed to improve the future of rangelands.

## **Introduction**

The northern Australian beef production sector encounters distinct challenges in rangeland systems, where climate variability, environmental degradation, complex grazing land management, and market fluctuations place significant pressure on producers. While rapidly evolving agricultural technologies, innovations, and scientific knowledge offer potential solutions, they are often inaccessible or impractical for direct application (Knickel et al. 2009; Klerkx and Proctor 2013). The sheer volume and complexity of available information can overwhelm producers, complicating decision-making and leading to inconsistent adoption of innovations (Bronson and Knezevic 2016; Eastwood et al. 2019).

While tools like ChatGPT serve broad queries, they lack the precision and contextual relevance required for rangeland beef production, often generating plausible but inaccurate outputs—referred to as hallucinations—that undermine reliability (Bender et al. 2021; Balaguer et al. 2024). BeefVantage addresses these gaps by leveraging collaborative intelligence—a synergy of human expertise and advanced AI technologies—to provide tailored decision support. By integrating fine-tuned Large Language Models (LLMs), and advanced AI techniques, the system offers contextually relevant advice aligned with real-world on-farm challenges in rangeland systems.

## **Methods**

Development followed a systematic workflow, beginning with problem definition and the collection and pre-processing of domain-specific data. The data was curated from trusted sources, including scientific journal papers, reports, case studies, fact sheets, conference proceedings, and industry guidelines. Multimodal content, such as podcasts and videos, was transcribed using Whisper (Radford et al. 2022) and integrated into the dataset.

Foundational models were identified and evaluated for integration with advanced AI techniques. Initial experiments focussed on fine-tuning Large Language Models (LLMs) (Ding et al. 2023) to generate nuanced domain-specific responses. Subsequently, advanced AI techniques were explored, including Retrieval-Augmented Generation (RAG) (Lewis et al. 2021) to retrieve relevant content from a vector database of domain-specific documents, and Graph Retrieval-Augmented Generation (GRAG) (Hu et al. 2024), which utilised a knowledge graph with 118,000 nodes and 508,000 edges to provide relational insights into critical topics. Additionally, an Agent Flow mechanism (Park et al. 2023) was also implemented to address knowledge gaps, conducting constrained web searches to deliver contextually relevant recommendations when the system could not generate a suitable answer.

Evaluation involved automated Q&A testing with 127 domain-specific question-answer pairs, assessing three metrics: relevance, groundedness, and helpfulness. All configurations were evaluated alongside foundational models. The highest-performing configuration was integrated into a Minimum Viable Product (MVP) with a user-friendly interface, allowing users to rate responses and provide descriptive feedback to support iterative improvements.

## **Results**

BeefVantage demonstrates significant potential as a transformative decision-support tool for beef producers. The system functions as an interactive AI assistant, generating context-specific recommendations for challenges such as drought management, sustainable grazing practices, biosecurity measures, and reproductive efficiency. By providing accessible domain-specific knowledge, it enables producers to make informed, strategic decisions that address regional challenges and environmental demands.

Experiments with fine-tuning Large Language Models (LLMs) (Ding et al. 2023) demonstrated no significant improvement over the original Llama 3.1 model, leading to the decision to retain the original model. Performance evaluation indicated modest improvements across techniques explored, namely RAG, GRAG, and Agent, in terms of relevance, groundedness, and practical utility when compared to foundational AI models.

Combinations of techniques yielded mixed results. For example, integrating RAG with Graph RAG, despite its theoretical potential, resulted in decreased performance. This decline may be attributed to catastrophic forgetting, a phenomenon where a model loses general reasoning capabilities while acquiring new domain-specific knowledge (Luo et al. 2024). Ultimately, the highest-performing configuration was Llama-3.1-70B with Agent-RAG-Web, achieving a strong balance of domain relevance, groundedness, and practical utility.

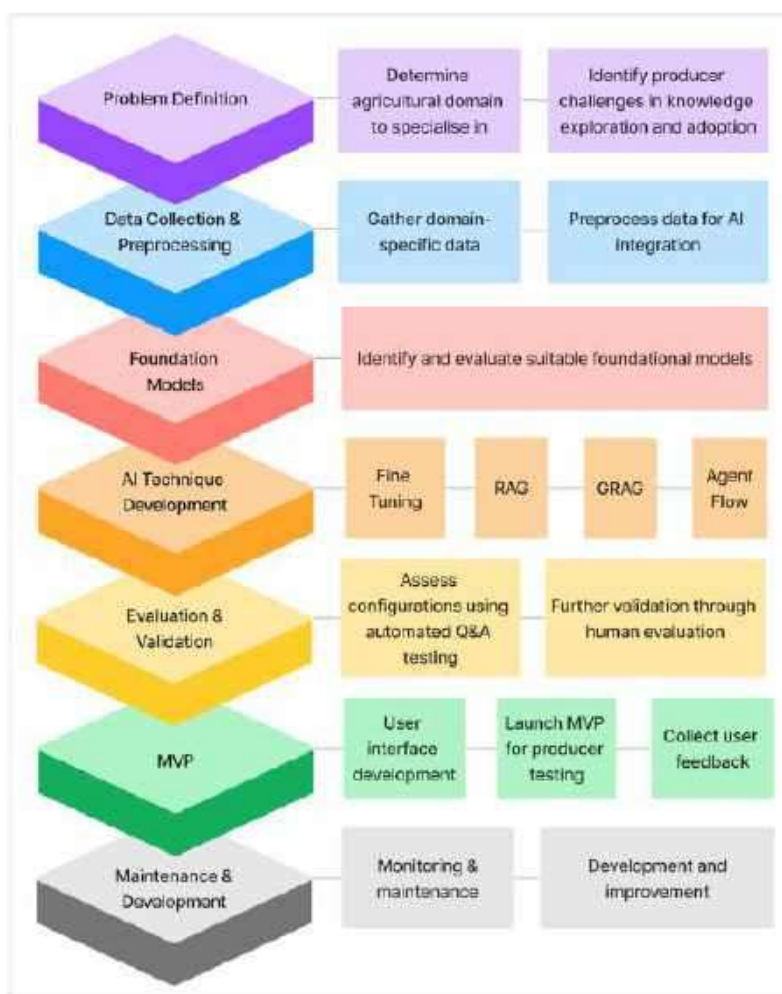


Figure 1. Development & Implementation Workflow

Key features of the MVP include a user-friendly interface with personalised accounts, saved chats, conversational memory, and the ability to download conversations. For example, producers can review previous queries related to grazing strategies, enabling continuity in decision-making. Preliminary testing

has demonstrated its effectiveness in presenting complex scientific insights in an actionable format, successfully bridging the gap between academic research and practical applications.

### Discussion

BeefVantage empowers producers to transition from reactive problem-solving to strategic planning by providing accessible knowledge and immediate, actionable insights. While trust in AI remains essential for widespread adoption in agricultural communities, transparent validation and producer-driven improvements can foster acceptance and encourage long-term engagement.

The system faces several challenges, including context window constraints, RAG/GRAG scalability, and adoption barriers such as digital literacy and connectivity literacy. Additionally, the platform includes a clear disclaimer, advising users to consult qualified professionals for veterinary or animal welfare concerns, ensuring clarity about its scope and limitations.

The next phase involves human evaluation to compare and validate results against automated metrics, ensuring practical utility. Feedback gathered upon the release of the Minimum Viable Product (MVP) will further refine the system, enabling continuous improvement and broader adoption. As foundational models and AI techniques continue to advance, significant improvements in system performance, scalability, and domain relevance can be expected, further enhancing the impact of tools like BeefVantage.

Future opportunities include adapting the system to other agricultural sectors and collaborating with technology partners to enhance capabilities, extend its relevance to full supply chain systems, and address broader agricultural challenges. Realising these possibilities will require sustained collaboration among researchers, producers, and industry stakeholders to overcome current limitations and amplify its impact.

Through the synergy of AI and human expertise, BeefVantage empowers producers to transform specialised knowledge into actionable insights, fostering sustainability, productivity, and resilience in rangeland systems. Its continued development and deployment hold the potential to redefine agricultural innovation, enabling beef producers to tackle pressing challenges and shape a more sustainable future.

### Acknowledgements

This project was funded by the JCU Sandpit 2 Seed Fund with support from the TNQ Drought Hub. We thank the producers and industry experts who contributed to the human evaluation phase. Ethics approval was obtained from the JCU Human Research Ethics Committee (HREC).

### References

- Balaguer A, Benara V, Renato Luiz de Freitas C, Roberto de MEF, Hendry T, Holstein D, Marsman J, Mecklenburg N, Malvar S, Nunes LO, Padilha R, Sharp M, Silva B, Sharma S, Aski V, Chandra R, 2024. RAG vs Fine-tuning: Pipelines, Tradeoffs, and a Case Study on Agriculture. Cornell University Library, arXiv.org, Ithaca.
- Bender EM, Gebru T, McMillan-Major A, Shmitchell S (2021) 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?' New York, NY, USA. (ACM. Available at <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922>)
- Bronson K, Knezevic I (2016) Big Data in food and agriculture. *Big data & society* 3,
- Ding N, Qin Y, Yang G, Wei F, Yang Z, Su Y, Hu S, Chen Y, Chan C-M, Chen W, Yi J, Zhao W, Wang X, Liu Z, Zheng H-T, Chen J, Liu Y, Tang J, Li J, Sun M (2023) Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature machine intelligence* 5, 220-235.
- Eastwood C, Klerkx L, Ayre M, Dela Rue B (2019) Managing Socio-Ethical Challenges in the Development of Smart Farming: From a Fragmented to a Comprehensive Approach for Responsible Research and Innovation. *Journal of Agricultural and Environmental Ethics* 32, 741-768.

- Hu Y, Lei Z, Zhang Z, Pan B, Chen L, Zhao L, 2024. GRAG: Graph Retrieval-Augmented Generation. Cornell University Library, arXiv.org, Ithaca.
- Klerkx L, Proctor A (2013) Beyond fragmentation and disconnect: Networks for knowledge exchange in the English land management advisory system. *Land use policy* 30, 13-24.
- Knickel K, Brunori G, Rand S, Proost J (2009) Towards a Better Conceptual Framework for Innovation Processes in Agriculture and Rural Development: From Linear Models to Systemic Approaches. *The journal of agricultural education and extension* 15, 131-146.
- Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Wen-tau Y, Rocktäschel T, Riedel S, Kiela D, 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Cornell University Library, arXiv.org, Ithaca.
- Luo Y, Yang Z, Meng F, Li Y, Zhou J, Zhang Y, 2024. An Empirical Study of Catastrophic Forgetting in Large Language Models During Continual Fine-tuning. Cornell University Library, arXiv.org, Ithaca.
- Park JS, O'Brien J, Cai CJ, Morris MR, Liang P, Bernstein MS, Follmer S, Han J, Steimle J, Riche NH (2023) 'Generative Agents: Interactive Simulacra of Human Behavior.' New York, NY, USA. (ACM. Available at <https://dl.acm.org/doi/pdf/10.1145/3586183.3606763>)
- Radford A, Kim JW, Xu T, Brockman G, McLeavey C, Sutskever I, 2022. Robust Speech Recognition via Large-Scale Weak Supervision. Cornell University Library, arXiv.org, Ithaca.