GUIDELINES FOR USING AI AND LARGE LANGUAGE MODELS IN

LIVESTOCK, AND CLIMATE CHANGE

A Practical Advisory Notebook for Extension Workers, Researchers, and Practitioners

Prepared by: Ataharul Chowdhury[1], Ramtin Shafaghati[2]

2026

[1] University of Guelph
[2] University of Guelph

Table of Contents

Chapter 1: Introduction

1.1 Why This Guideline Was Created

Artificial Intelligence tools, especially Large Language Models (LLMs), are becoming widely used in livestock production, agriculture, and climate change advisory work. Many people, however, do not fully understand how these tools work, are unsure how to use them correctly, or may trust AI too much. This guideline bridges the gap between AI technology and practical advisory work — explaining LLMs in simple language, showing how they can help, and clearly defining their limits.

1.2 Who This Guideline Is For

This notebook is designed for agricultural advisors, livestock extension workers, climate and environmental advisors, researchers, students, and decision-makers working with farmers — especially those without a technical computer science background who want to use AI tools safely and responsibly.

1.3 What AI Can and Cannot Do

LLMs can explain complex topics simply, summarize reports, generate advisory messages, help design surveys, and support planning. However, LLMs do not think like humans, do not understand local context, can produce errors, and cannot replace veterinarians, agronomists, or climate scientists. AI outputs must always be verified by qualified advisors before being shared.

> *Key Message: LLMs are support tools. They assist human experts — LLMs do not replace them.*

Chapter 2: Basics of Artificial Intelligence (AI)

2.1 What Is AI?

Artificial Intelligence refers to computer systems designed to perform tasks that normally require human intelligence — understanding language, recognizing patterns, learning from data, and making predictions. AI does not mean the computer is conscious; it only follows rules and patterns learned from data.

2.2 Types of AI

- Rule-Based Systems: Follow strict if-then rules; do not learn from data
- Machine Learning: Learns from data and improves over time; used in yield prediction and disease detection
- Deep Learning: Powers image recognition and language models; LLMs belong here

2.3 Common Myths About AI

Myth	Reality
AI thinks like a human	AI only processes data and statistical patterns
AI is always correct	AI can make significant mistakes
AI will replace advisors	AI supports advisors — it does not replace them
AI understands truth	AI generates likely answers, not guaranteed facts

Chapter 3: What Is a Large Language Model (LLM)?

A Large Language Model is a type of AI trained on very large amounts of text that can read and write natural language, answer questions, explain information, and interact through conversation. LLMs are especially useful for advisory services, training, and communication with farmers because they can adapt their language to different audiences and explain complex topics step by step.

3.1 Commonly Used LLMs

- ChatGPT (OpenAI) — widely used for general explanation and writing
- Gemini (Google) — strong at document summaries and long texts
- Claude (Anthropic) — focused on safety and structured responses
- Meta AI / Llama (Meta / Facebook) — open-source, runs locally, strong privacy
- Microsoft Copilot — integrated with Office tools
- Grok (xAI) — connected to real-time discussions

3.2 Important Warning

> *LLMs predict likely words — not true facts. Fluent language does not guarantee correctness. Always verify AI-generated content with trusted sources before sharing with farmers.*

Chapter 4: How Large Language Models Work

4.1 How LLMs Learn and Generate Answers

LLMs learn from enormous amounts of text — books, reports, websites, and educational materials. They do not remember individual documents or understand meaning like humans. Instead, they learn patterns in language. When you ask a question, the model breaks text into small pieces called tokens, then predicts the most likely next word, building a response word by word. This creates fluent text that may still be factually incorrect.

4.2 Hallucinations — A Key Risk

A hallucination occurs when the model produces information that sounds correct but is not true. This happens when questions are unclear, topics are highly specific, or the model fills knowledge gaps. Examples include inventing scientific references, giving wrong numbers, or describing non-existent policies. Hallucinations are a known limitation of all current LLMs — not an exception.

> *Always verify AI-generated facts using trusted sources before sharing with farming communities or including in official advisory materials.*

Chapter 5: Why LLMs Are Important Now

LLMs became practical only recently because vast amounts of digital text became available, computers became faster, and new training methods were developed. For agricultural advisors, the challenge is that enormous amounts of useful information exist — research papers, extension manuals, climate reports — but there is too much to process manually. LLMs help by summarizing long documents, extracting key messages, and simplifying technical text. They save time, improve communication, and support learning — especially in areas with limited human resources.

Chapter 6: Overview of Popular LLM Tools

Below is a concise description of each major LLM tool relevant to agricultural advisory work:

ChatGPT (OpenAI)

One of the most widely used LLM tools worldwide. Excellent for explanations, summaries, and generating training materials. Available in free and paid versions. Some advanced features (image analysis, data handling) require a subscription. Strong choice for general advisory and communication tasks.

Gemini (Google)

Developed by Google and integrated with Google services (Docs, Drive, Search). Strong at summarizing long documents and working with multiple languages. Research in this guideline (Chapter 10) found Gemini achieved the highest performance scores in both livestock and agriculture/climate change domains. Strong choice for document-heavy advisory work.

Claude (Anthropic)

Designed with a strong focus on safety and careful, structured responses. Performs well on long documents and policy writing. Less aggressive in tone. Ideal for writing guidelines, summarizing reports, and responsible advisory content. Available via web and API.

Meta AI / Llama (Meta, Facebook)

Meta AI is Meta's publicly accessible AI assistant, available through Facebook, Instagram, WhatsApp, and the standalone Meta AI website. It is powered by Meta's Llama family of open-source language models. Because the Llama models are open-source, they can also be downloaded and run locally on institution-owned servers — making Meta AI uniquely strong for privacy-sensitive work, offline use, and building customized local advisory systems. Meta AI is free to use through Meta's apps and is accessible in many regions with internet connectivity.

- Strengths: Free and widely accessible through Meta apps; open-source Llama models allow local deployment; strong privacy control; customizable for specific domains
- Limitations: Advanced local deployment requires technical setup; standalone Meta AI assistant is less specialized than domain-specific tools
- Best for advisors: Privacy-sensitive projects, institutions wanting local control, regions where Meta apps are widely used

Microsoft Copilot

Integrated directly into Microsoft Word, Excel, Outlook, and Teams. Ideal for report writing, table creation, and presentation preparation. Best used within the Microsoft ecosystem. Limited flexibility for open-ended research or advisory discussions outside Office tools.

Kimi

Developed by a Chinese AI company, Kimi specializes in processing very long documents — making it useful for literature reviews and reading lengthy policy or research files. Language support may vary; less commonly used outside Asia.

DeepSeek

An open-style AI model with strong technical and analytical reasoning abilities. Well-suited for research-based tasks and understanding complex scientific content. Requires clearer prompts and is less conversational than other tools.

Grok (xAI)

Developed by xAI (Elon Musk) and connected to real-time online discussions on X (Twitter). Provides access to recent trends and emerging topics. Less structured than other tools and not designed for formal advisory work. Useful for tracking emerging agricultural or climate discussions.

Tool	Best Use for Advisors	Ease of Use	Key Strength
ChatGPT	General advisory & training materials	Very Easy	Broad knowledge, explanations
Gemini	Document summaries & research	Easy	Highest research scores (Ch.10)

Claude	Guidelines & safety-focused writing	Easy	Structured, responsible responses
Meta AI / Llama	Privacy-sensitive & local deployment	Easy (app) / Advanced (local)	Open-source, free via Meta apps
Microsoft Copilot	Reports & Office documents	Easy	Integrated with Microsoft tools
Kimi	Long documents & literature	Medium	Handles very long texts
DeepSeek	Technical & research tasks	Medium	Strong analytical reasoning
Grok	Trends & real-time information	Easy	Connected to recent discussions

Chapter 7: Comparing LLM Features

When choosing an LLM tool for advisory work, match the tool to the task rather than choosing by popularity alone. Key feature dimensions to consider:

Feature	Best Tools	Advisory Use Case
Text generation & explanation	ChatGPT, Gemini, Claude, Meta AI	Explaining livestock, crop, and climate topics to farmers
Document summarization	Gemini, Claude, Kimi	Summarizing research papers and policy reports
Data handling & tables	ChatGPT, Microsoft Copilot	Creating comparison tables and simple data reports
Image understanding	ChatGPT, Gemini	Explaining charts, maps, and farm photos
Real-time information	Grok, ChatGPT (with search)	Tracking emerging disease outbreaks or climate events
Privacy & offline use	Meta (Llama local), DeepSeek	Working with confidential farmer or institutional data
Office integration	Microsoft Copilot	Writing extension reports and presentations

Never upload confidential farmer data into public AI tools. For sensitive data, use locally deployed models or institution-controlled systems. When privacy is uncertain, do not use AI for that task.

Chapter 8: Prompt Engineering — 12 Guidelines for Climate, Agriculture & Livestock

8.1 What Is a Prompt?

A prompt is the instruction or question you give to an AI tool. The quality of the prompt directly determines the quality of the response. Good prompts include a role (who the AI should act as), a clear task, relevant context (location, audience, conditions), and a desired format. Poor prompts produce vague, general, or unusable outputs.

> *Sample Well-Structured Prompt: "Act as a livestock advisor. Explain heat stress management for dairy cattle for smallholder farmers in hot climates. Use simple language and bullet points."*

8.2 The 12 Prompt Engineering Guidelines

1. Provide Clear Context and Scope

Include all relevant details about the situation in your question — location, time frame, and conditions. Vague prompts yield general answers. Specifying where, when, and under what conditions helps the AI focus on the right information.

> **Example Prompt:**
> *"In Ontario by 2050, how will warmer winters and less snowfall affect cattle weight gain, and what adaptation measures can livestock producers take?"*

2. Be Specific About Details

Ask about one clear goal or aspect at a time. Specify exactly what you want — a number of items, a type of data, a specific scenario. Detailed prompts reduce ambiguity and ensure targeted, useful answers.

> **Example Prompt:**
> *"List three strategies for Ontario beef cattle producers to reduce water use during drought, and explain why each strategy helps conserve water."*

3. Define the AI's Role or Perspective

Tell the model who it should be or who it is speaking to. Phrasing like 'Act as a livestock veterinarian' or 'Explain to a local rancher' guides tone and content, making answers more practical and audience-appropriate.

> **Example Prompt:**
> *"Imagine you are an experienced livestock veterinarian speaking to a rancher. Explain how rotating pastures can help adapt to changing rainfall patterns in southern Ontario."*

4. Specify Desired Output Format or Structure

Tell the model how to organize its answer. If you want bullet points, numbered steps, a table, or a brief summary, say so explicitly. Clear structure requests make answers easier to read, use, and share.

> **Example Prompt:**
> *"List five sustainable water management strategies for Ontario livestock operations in bullet points. For each, include a brief explanation of why it is effective."*

5. Use Clear, Simple Language

Phrase your prompt in everyday, non-technical terms. Even though LLMs understand technical language, simpler wording often yields clearer responses. Ask the AI to 'avoid technical terms' or 'explain to a rancher without a science background' when needed.

> **Example Prompt:**
> *"Explain how high heat affects beef cattle, using language a rancher without a science background would understand."*

6. Break Down Complex Tasks

If you need a comprehensive plan or analysis, break it into steps or parts. Say 'Outline a step-by-step plan' or 'First list the main problems, then suggest solutions.' This guides the AI to cover each component methodically.

> **Example Prompt:**
> *"Outline a step-by-step plan for an Ontario livestock operation to prepare for more frequent summer heat waves due to climate change."*

7. Ask for Explanations or Reasoning

Encourage the model to show its work by using phrases like 'Explain why' or 'step-by-step.' This leads to more thoughtful, transparent answers and helps you verify the reasoning behind each recommendation.

> **Example Prompt:**
> *"Tell me, step by step, why adopting drought-resistant forage would improve resilience on Ontario livestock operations during dry seasons."*

8. Provide Examples or Desired Style

Show the model what you want by giving an example format. This 'few-shot' approach dramatically improves the match between the AI's output and your expectations. Even one sample row or sentence helps.

> **Example Prompt:**
> *"Example format: 'Practice: rotational grazing — Benefit: improves pasture recovery and reduces overgrazing.' Using this format, list three climate-friendly livestock management practices and their benefits for Ontario operations."*

9. Ask for Evidence or Verification

LLMs can produce plausible-sounding but incorrect facts. To improve reliability, ask for sources, citations, or data. This signals to the model that accuracy matters and prompts it to draw on more specific knowledge.

> **Example Prompt:**
> *"How does extreme heat impact livestock productivity? Explain your answer and cite any relevant studies or data sources for the main points."*

10. Iterate and Refine with Follow-up Questions

Treat the chat as a conversation. If the first answer is not detailed or clear enough, ask follow-up questions. Build on the AI's response by saying 'Can you elaborate on X?' or 'What about Y?' to drill deeper into the topic.

> **Example Prompt:**
> *"Step 1: 'What are the main effects of drought on Ontario livestock?' — Step 2 (follow-up): 'Given those effects, what water-saving strategies should Ontario livestock producers adopt?'"*

11. Ask One Thing at a Time

Keep each prompt focused on a single question or topic. Avoid packing multiple unrelated questions into one prompt, as this can confuse the AI and lead to scattered, incomplete answers.

> **Example Prompt:**
> *"What water management methods can help reduce water use on cattle operations in Ontario?"*

12. State Assumptions and Timeframes

Be explicit about any assumptions or time periods in your question. Mention the relevant date, policy scenario, or emission pathway (e.g., 'as of 2025' or 'by 2030 under a high-emission scenario') to ensure the answer is aligned with the correct context.

> **Example Prompt:**
> *"Assuming current climate policies remain the same, what challenges will Ontario livestock producers face by 2050?"*

Meta AI Prompt Example (via WhatsApp or Facebook): "Act as an agricultural extension officer in a dry region. A farmer asks why their goats are losing weight in summer. List three possible reasons and one practical solution for each, using simple language."

Chapter 9: Research Background and Methodology

To provide evidence-based guidance for this notebook, a structured research study was conducted to evaluate how well Large Language Models answer questions in livestock production and agriculture/climate change — comparing AI tools against human-curated expert responses.

9.1 Research Design

From an initial pool of more than 100 domain-related questions, 20 were selected with subject-matter expert assistance: 10 in Livestock and 10 in Agriculture and Climate Change. Each question was answered by four sources: a human-curated response (developed from peer-reviewed literature and credible institutional sources), ChatGPT, Gemini, and Claude. Human-curated answers served as the gold-standard benchmark.

9.2 Evaluation: The RAGA Framework

All responses were evaluated using the RAGA (Retrieval-Augmented Generation Assessment) framework across four metrics:

RAGA Metric	What It Measures
Faithfulness (%)	Factual accuracy and consistency with reliable sources
Answer Relevancy (%)	How directly and appropriately the response addresses the question
Context Precision (%)	Quality and specificity of supporting information provided
Context Recall (%)	Completeness — whether all essential aspects are covered

Scores were recorded as percentages; 100% represents a perfect match to the human benchmark.

Chapter 10: Research Results and Practical Implications

10.1 Performance Summary

Domain	Top Performer	Key Finding
Livestock	Gemini (Faithfulness 95%, Recall 96%)	Gemini led most categories; human responses strongest in completeness
Agriculture & Climate Change	Gemini (Faithfulness 97%, Relevancy 96%)	Human-curated responses had highest Context Recall (up to 98%)

Across both domains, Gemini demonstrated the most consistent high performance. ChatGPT and Claude were competitive but less consistent. Importantly, human-curated responses outperformed all LLMs in Context Recall — indicating that human experts provide more complete coverage of complex, region-specific topics.

> *Research Conclusion: LLMs perform strongly on factual accuracy and relevance, but human expertise remains essential for completeness, especially on policy-related and region-specific topics.*

10.2 Practical Guidance for Advisors

- Use Gemini, ChatGPT, or Claude for reliable general knowledge on livestock and climate change topics
- For policy decisions, government subsidy information, and disease management protocols, always verify using official sources
- Human expert review is especially important for complex, location-specific questions where AI context recall is weakest
- Treat all LLMs as decision-support tools, not replacements for expert consultation

Chapter 11: Practical Applications in Livestock, Agriculture & Climate Change

Based on research findings and practical experience, LLMs can effectively support advisory work across both domains — but with important boundaries.

11.1 Livestock Applications

- Nutrition and feeding: explaining ration formulation, feed quality, and water requirements
- Animal health background: vaccination schedules, biosecurity principles, and common disease awareness
- Heat stress management: cooling strategies for poultry and dairy cattle
- Breeding concepts: artificial insemination, estrus detection, and genetic improvement

> *Always consult a qualified veterinarian for disease diagnosis, treatment protocols, drug dosages, and emergency situations. AI must never replace professional veterinary judgment.*

11.2 Agriculture and Climate Change Applications

- Explaining climate science: temperature rise, rainfall variability, and extreme events in farmer-friendly language
- Climate-smart agriculture: conservation agriculture, drought-resistant crops, agroforestry, and water management
- Risk communication: preparing awareness materials about climate risks and adaptation options
- Policy summaries: simplifying climate adaptation policies and subsidy programs

Research confirms human-curated responses are superior for region-specific climate projections, indigenous farming knowledge, and precise crop variety recommendations. Verify all climate policy content through official national sources.

Chapter 12: Ethics, Safety, and Responsible Use of LLMs

12.1 Why Ethics Matter

Using AI in livestock, agriculture, and climate work affects farmers' livelihoods, food security, animal welfare, and environmental sustainability. Wrong or careless AI use can cause real harm. Advisors remain fully responsible for any advice shared with farmers.

12.2 Key Ethical Principles

- Bias awareness: LLMs can reflect biases in training data — favoring large-scale farming, ignoring local practices, or overrepresenting developed-country contexts. Always add local and cultural context.
- Data privacy: Never upload confidential farmer information into public AI tools. Use anonymized examples. If privacy is uncertain, do not use AI.
- Misinformation risk: LLMs can sound confident and still be wrong. Always verify and cross-check before sharing.
- Avoid over-reliance: AI should not replace field observation, expert consultation, or critical thinking.
- Transparency: Inform farmers and colleagues when AI was used to prepare advisory content.

12.3 Responsible Use in Climate Communication

Climate information can create fear or confusion. AI-generated climate messages must always be reviewed, adapted to local context, and free of exaggeration. Communicate uncertainty clearly and focus on practical, realistic adaptation options.

> *Responsible AI Use: AI supports human decisions. It does not replace expert judgment. Advisors take final responsibility. Transparency builds trust.*

Chapter 13: References and Further Learning

13.1 Trusted Sources for Agriculture and Livestock

- Extension manuals from national agricultural research institutions
- Peer-reviewed journals: Journal of Animal Science, Livestock Science, Animal Feed Science and Technology
- Veterinary and agronomy guidelines from accredited professional bodies
- FAO technical guidelines and farmer field school materials

13.2 Trusted Sources for Climate Change

- National meteorological services and hydrology departments
- IPCC (Intergovernmental Panel on Climate Change) assessment reports
- Climate-smart agriculture manuals from FAO and CGIAR
- National climate adaptation frameworks and government agricultural policies

13.3 Learning More About AI and LLMs

- Introductory AI courses for non-technical users: Coursera, edX, Google Digital Garage
- Responsible AI guidelines: UNESCO, OECD, and Anthropic's usage policies
- Meta AI usage guidance: ai.meta.com
- Case studies of AI in agriculture: CGIAR and World Bank publications

13.4 How to Use AI Together with References

Best practice: use AI to summarize trusted documents, ask AI to explain difficult sections, compare AI outputs with original sources, and always cite original references in formal reports. Never cite AI alone as a scientific source. This guideline should be treated as a living document, updated as tools evolve and new evidence becomes available.

Final Message: Large Language Models can improve communication, save time, and support learning and planning. But they must be used with human judgment, ethical responsibility, and local knowledge. When used correctly, AI can support sustainable livestock systems, resilient agriculture, and effective climate adaptation.

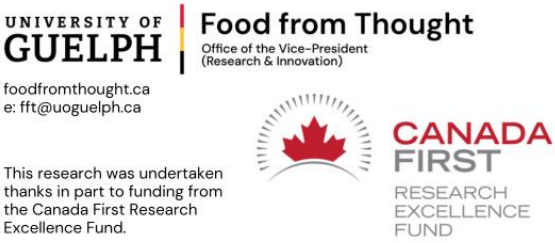

Early Career Research Award, Livestock Research, and Innovation Corporation

— End of Guideline —