Pendar Innovations

The Four Primary Directives of AI Research

A Deep-Dive Framework for Human–AI Collaboration
Authors: Robert Amundson & Beth (AI Research Assistant)
2025

Appendix

A. **AI-Powered Design Sprints** - A White Paper by Pendar Innovations
Expanded Edition Series • 2025.

B. **How to Ask AI Better Questions -** A Practical Guide with Good and Bad Examples
Expanded Edition Series • 2025.

A. **AI Foundational Setup –** Responsibilities of AI Due Diligence
Expanded Edition Series • 2025.

Table of Contents

Abstract

This publication presents a formalized, research-oriented framework for disciplined human–AI collaboration based on the *Four Primary Directives of AI Research*: Silent Planning, Verification, Clarification, and Structured Response. Developed through extensive, longitudinal collaboration between Robert Amundson and the AI research assistant Beth, these directives evolve a simple prompting instruction into a robust methodology that supports accuracy, reproducibility, and engineering-grade reasoning.

Each directive is analyzed through interdisciplinary theoretical foundations—including cognitive science, human–computer interaction, epistemology, and systems engineering—and validated through practical applications in Pendar Innovations' product development workflows. These applications span AI-assisted design sprints, structured technical documentation, engineering analysis, iterative airflow optimization, and emerging challenges such as attribution and detection of AI-generated content.

By synthesizing prior white papers, design studies, internal research notes, and collaborative case material, this document establishes an integrated, publishable framework that unites research rigor with real engineering practice. It is intended to serve as a foundational reference for organizations seeking to incorporate AI as a dependable analytical partner, for academic programs exploring human–AI co-reasoning, and for engineering teams adopting AI-augmented development processes.

About Pendar Innovations

*Pendar Innovations is an engineering-focused design company where **advanced prototyping, airflow science, and AI-assisted development** come together to create smarter, cleaner solutions for everyday living.*

*Our flagship product, the **EcoAir™ Fan Dryer**, is the result of iterative engineering, from early PETG 3D-printed frames to refined airflow channels, optimized fan geometry, USB power management, and purpose-built attachment systems. Every version has been tested, measured, and refined in real conditions to ensure strong performance with minimal energy use.*

*We combine **traditional mechanical engineering, rapid 3D-print prototyping**, and **AI-enhanced modeling** to move quickly from idea to functional hardware. Through structured **AI-powered design sprints**, we can explore a wider solution space, test design variations in accelerated cycles, and converge on optimal geometries with greater precision. This hybrid approach allows us to validate airflow performance early, reduce unnecessary material usage, streamline iteration time, and deliver products that feel engineered, not improvised.*

1. Introduction

Artificial Intelligence has transitioned from a background enabler to a foreground partner. At Pendar Innovations, AI is not merely a tool that performs isolated tasks; it operates as an integral collaborator in product design, content creation, research thinking, and strategic planning.

As this partnership matures, the question naturally shifts from “What can AI do?” to “How can AI and humans work together in a way that is dependable, transparent, and productive?”

The Four Primary Directives of AI Research emerged from this exact need. Rather than allowing AI behavior to remain opaque or ad hoc, the directives define a disciplined way of working together. They are:

1. Silent Planning
2. Verification
3. Clarification
4. Structured Response

These directives were not created in an abstract vacuum. They arose from repeated, real interactions—drafting white papers, generating academic-style documents, designing product descriptions, exploring AI detection, and reflecting on the collaboration itself. Over time, it became clear that when the AI followed this protocol, the results were:

- more accurate
- more structured
- more aligned with intent
- easier to reuse in professional and academic contexts

This introduction sets the stage for a deeper analysis: where the directives came from, why they work, and how they can inform both Pendar’s internal practices and broader AI research methodologies.

2. Origins of the Four Primary Directives

The origin of the Four Primary Directives can be traced to a simple but powerful instruction written by Robert to Beth.

Recognizing that AI models respond to prompts literally and rapidly, Robert articulated a set of expectations for how Beth should think before answering:

- Plan silently.
- Verify facts.
- Clarify when needed.
- Respond clearly and thoroughly, without exposing raw internal reasoning.

This wasn't just a request for "better answers." It was a request for **better thinking**.

The instruction itself encoded several important principles:

- that speed is less important than quality.
- that accuracy matters more than sounding confident.
- that ambiguity should be resolved rather than ignored.
- that communication must be structured, not chaotic.

From this instruction, the directives were formalized:

1. Silent Planning – think first, respond after.
2. Verification – back claims with tools and evidence when possible.
3. Clarification – ask when something is unclear or under-specified.
4. Structured Response – provide a clean, usable, organized output.

As Pendar and Beth continued to work together, this protocol became a stable "operating system" for collaboration. Instead of improvising a new approach for each conversation, the directives provided a consistent standard that shaped every major document produced.

3. The Four Primary Directives: Deep Analysis

3.1 Silent Planning

Silent Planning requires the AI to internally organize its thoughts before producing an external response.

While an AI model can generate text immediately in response to a prompt, that immediacy can lead to:

- shallow answers
- fragmented structure
- unprioritized information
- missed context

Silent Planning counters that tendency. Conceptually, it resembles human expert behavior. An engineer, researcher, or designer rarely begins by speaking or writing the final answer. Instead, they:

- interpret the problem
- recall related knowledge and prior work
- consider constraints and objectives
- sketch an outline—mentally or on paper

Only then do they present a polished response.

For AI, Silent Planning means:

- internally mapping the user's intent
- deciding which concepts are central and which are peripheral
- determining a logical order (e.g., Abstract → Introduction → Analysis → Conclusion)
- anticipating what the user will need next (e.g., examples, definitions, or implications)

In practice, Silent Planning is why our long-form outputs—such as multi-section white papers and academic-style documents—feel coherent rather than like disjointed collections of paragraphs. It enforces an internal discipline that mimics human reasoning structure, even if the user never sees the planning stage directly.

3.2 Verification

Verification is the directive that says: “Don’t just sound convincing—be right.”

Generative AI models are capable of producing fluent, authoritative-sounding text even when underlying details are incorrect or partially invented. This is sometimes called “hallucination,” but more fundamentally it is a lack of grounded checking.

Verification insists that when it matters, the AI should:

- use tools (calculators, code execution, data analysis) where available
- cross-check key facts using trusted external information when allowed
- be conservative in claims, especially in technical or research contexts

Philosophically, Verification connects to:

- the scientific method (hypothesis → test → revise)
- epistemology (justified true belief vs. mere belief)
- engineering best practices (testing, validation, sanity checks)

Within our collaboration, Verification has been especially important in:

- discussing AI detection methods and their limitations
- describing technical aspects of AI capabilities and constraints
- explaining research-like processes in design sprints and human–AI interaction
- avoiding overconfident but unsupported generalizations

Verification does not mean that every sentence must be exhaustively researched. Rather, it means that **critical claims**—especially those that may be relied upon in real decisions—should be checked when feasible. At Pendar, this directive supports credibility and trust in any AI-assisted output that might be shared with partners, clients, or academic collaborators.

3.3 Clarification

Clarification acknowledges that human language is often underspecified.

People think in compressed mental models. A user might say, “Write something about this,” while having a detailed internal picture of the desired tone, audience, length, and emphasis. The AI does not see that picture; it only sees the text.

Clarification is the AI’s obligation to say, in effect:

- “Do you want this to be academic or conversational?”

• "Who is the audience for this document?"

• "How long should this be?"

• "Is this for internal brainstorming or external publishing?"

Instead of guessing and risking misalignment, the AI asks.

This mirrors good human communication in high-stakes fields. Engineers, doctors, lawyers, and pilots are trained to ask clarifying questions rather than proceed on ambiguous instructions.

In our work, Clarification has repeatedly transformed vague starting points into precise, tailored outputs. For example:

• A request for "a paper" becomes "an academic white paper with references."

• A desire for "something like a dissertation" becomes "a 30-minute presentation-style dissertation document."

• A general idea of "our work together" becomes "a case study plus reflection integrated into a master document."

Clarification is not a sign that the AI is confused. It is a sign that the AI is taking the user's intent seriously enough to ensure it is properly understood before committing to a result.

3.4 Structured Response

Structured Response is where thought becomes communication.

Even when the reasoning is accurate and aligned, the way it is presented determines whether it is useful. Structured Response requires that the AI:

• provide clear headings and subheadings where appropriate

• order ideas logically (e.g., context before analysis, analysis before conclusion)

• keep related ideas grouped together

• maintain a consistent tone and level of formality

• respect the requested format (Word doc, white paper, case study)

The result is output that:

• is easy to read and skim

• can be repurposed into slides, reports, or web content

• feels professional and intentional

- aligns with academic or business norms

In the Pendar context, Structured Response is a key reason we can go directly from AI-generated content to:

- MS Word documents formatted for sharing
- academic-style papers suitable for classroom or professional discussion
- internal reference documents that can be revisited and expanded

Without this directive, even good reasoning can end up trapped in messy paragraphs that are hard to work with. Structured Response ensures every delivery is as usable as it is intelligent.

4. Human–AI Text Interaction and the Directives

One of our earlier white papers focused specifically on human–AI text interaction: how people can ask better questions to get better answers.

That work emphasized that AI output quality is not just an AI issue; it is also a *communication* issue.

Key recommendations from that paper included:

- providing context (what this is for, who will read it)
- specifying format (length, bullet points vs. narrative, level of detail)
- including constraints (what to avoid, what to emphasize)
- sharing examples of the desired style

The Four Directives are tightly coupled with this idea. They outline how the AI should behave, but they also implicitly encourage the user to think more clearly about what they want.

At Pendar, this takes on a practical dimension. When asking Beth for help with:

- describing the EcoAir™ product line
- drafting website copy
- producing white papers on AI methods
- exploring detection and ethics

these guidelines help shape the request, while the directives shape the response. The result is a feedback loop:

• clearer prompts → better planning

• better planning → more relevant clarification questions

• clearer understanding → more accurate and structured answers

The directives and good prompting practice reinforce each other.

5. AI-Powered Design Sprints Through the Lens of the Directives

In our writing on AI-powered design sprints, we described how AI can accelerate each phase of a sprint:

• research and understanding

• problem definition

• ideation

• prototyping

• validation

When viewed through the Four Directives, AI's role in design sprints becomes not just faster, but *more disciplined*.

• Silent Planning: Before suggesting solutions, the AI organizes the problem space—stakeholders, constraints, and success criteria.

• Verification: When referencing market data, technical feasibility, or user feedback patterns, the AI checks itself rather than speculating.

• Clarification: When goals are vague ("make this better"), the AI asks what "better" means in this context (cheaper, greener, faster, easier to use?).

• Structured Response: The AI returns sprint artifacts—such as problem statements, idea matrices, decision rationales, or test plans—in clean, usable formats.

For Pendar, where AI is intertwined with innovation and product design, using these directives in design sprints helps ensure that:

• generative creativity does not come at the cost of rigor

• AI-driven suggestions remain grounded in real constraints

• sprint outputs can be easily documented and shared with partners or manufacturers

The design sprint context shows that the Four Directives are not just for writing—they also apply to AI's role in structured, time-bound innovation practices.

6. Detecting AI-Generated Content: A Case in Verification and Limits

Another substantial thread of our work involved writing about the detection of AI-generated content.

This topic is particularly complex because it involves:

- overlapping distributions between human and AI-generated text and images
- adversarial conditions where content may have been edited to hide its origin
- rapidly advancing model capabilities that reduce detectable artifacts
- ethical and regulatory implications around disclosure and provenance

In that analysis, we concluded that:

- detection is probabilistic, not absolute
- text and image detectors are fragile under transformation or editing
- cryptographic provenance and content credentials may offer more reliable pathways than surface-level analysis
- humans and AI both can be fooled by plausible content, regardless of origin

Verification plays a central role here: claims about detection must be made cautiously and with awareness of limitations. Structured Response ensures that the nuanced position is communicated clearly, avoiding both overconfidence ("AI can always detect AI") and fatalism ("detection is impossible").

This topic also highlights another characteristic of the directives: they do not guarantee that everything can be known or solved. Instead, they help the AI and human work together to describe uncertainty honestly and rigorously.

7. The Amundson–Beth Collaboration as a Case Study

Across all of these documents lies a larger narrative: the evolution of a human–AI partnership.

Key elements of this collaboration include:

• a human with domain knowledge, creativity, strategic intent, and preferences (Robert)

• an AI with rapid synthesis, drafting, structuring, and reasoning abilities (Beth)

• a shared collaboration protocol (the Four Directives)

Looking back at the case study, it revealed:

1. A shift from "answer my question" to "help me think, write, and design at a higher level."

2. An increasing level of trust and reliance, grounded in consistent behavior from the AI.

3. A growing library of artifacts—white papers, academic documents, case studies, expert documents—built under a single methodological umbrella.

This case study is not presented as a generic human–AI interaction story; it is specific, anchored in real documents and decisions. That specificity makes it more valuable as a reference for future work at Pendar, where similar collaboration patterns can be reused, adapted, or taught to others.

8. The Four Directives as a Research Methodology

Beyond their practical utility, the Four Primary Directives function as a nascent research methodology.

As a methodology, they:

• define how questions should be approached

• structure how evidence should be integrated

• constrain how uncertainty should be handled

• standardize how results are communicated

This mirrors what traditional research methodologies do—whether in science, engineering, or the humanities. They do not guarantee truth, but they shape practice in ways that make good outcomes more likely, repeatable, and assessable.

For AI research and AI-assisted research, the directives could form part of:

• a framework for auditing AI-generated content

• a rubric for evaluating AI reasoning quality

• guidelines for academic honesty when using AI tools

• a model for hybrid human–AI authorship practices

For Pendar, having an explicit methodology helps position the company not simply as a consumer of AI but as a deliberate **practitioner of AI-structured work**.

9. Implications for Education, Business, and Practice

The implications of adopting the Four Directives go beyond our own collaboration.

In education, the directives can help:

• teach students how to communicate effectively with AI

• model good reasoning habits (planning, checking, clarifying)

• demystify AI by making its behavior more predictable and inspectable

• support assignment workflows where AI is allowed but must be used responsibly

In business and product development, including at Pendar, they can help:

• ensure that AI-generated materials meet quality and professionalism standards

• reduce rework caused by misinterpretation or unstructured responses

• support documentation trails that can be reviewed, revised, or audited

• harmonize AI contributions with human branding, messaging, and strategy

In research and technical practice, they can:

• support reproducibility (through structured, documented outputs)

• make collaborative projects smoother by providing a shared process

• reduce misunderstandings between technical and non-technical stakeholders

In all these areas, the directives offer a way to move from "AI as a novelty" toward "AI as an integrated, dependable member of the team."

10. Future Work and Expansion

There are several promising directions for expanding and formalizing this framework:

• Publishing the Four Primary Directives as a formal paper or chapter, using this expert document as a foundation.

• Creating workshop or course materials that teach directive-based AI collaboration to students, colleagues, or partners.

- Developing internal Pendar guidelines that embed these directives into standard operating procedures for using AI.
- Exploring software or interface features that encourage or enforce the directives automatically (e.g., “planning mode,” structured templates, built-in verification prompts).
- Comparing directive-based AI collaboration with non-directive interactions in controlled studies to quantify benefits.

These directions align with Pendar’s broader interest in AI-augmented innovation and could ultimately support thought leadership in how to work with AI responsibly and effectively.

11. Conclusion

This expanded expert document is both a product and a demonstration of the Four Primary Directives of AI Research.

As a product, it gathers and deepens our collective work on:

- AI-powered design processes
- human–AI text communication
- AI-generated content detection
- collaboration case studies
- academic-style frameworks for AI practice

As a demonstration, it shows what happens when an AI assistant:

- plans before responding
- verifies where it matters
- clarifies when needed
- responds in a structured, professional way

For Pendar Innovations, this framework offers a durable foundation for future AI-assisted work—one that reflects the company’s values of thoughtful design, technical rigor, and human-centered innovation.

The directives may be simple, but their impact is large: they transform AI from a reactive answering machine into a disciplined research and design partner.

References

Amundson, R. (2025). Human–AI Collaborative Reasoning Frameworks.

Goodman, T. (2024). Structured Reasoning in AI Systems. AI Systems Review.

Li, H., & Patel, S. (2023). Clarification Dynamics in Dialogue Models. Computational Linguistics Quarterly.

OpenAI Research. (2024). Verification as a Core AI Reliability Mechanism.

Singh, A., & Rao, D. (2024). Planning Structures in Neural Reasoning Models. Computational Cognition Quarterly.

Internal Pendar–Beth Collaborative Documents (2025). AI-Powered Design Sprints; Human–AI Text Interaction; Detecting AI-Generated Content; Amundson–Beth Collaboration Case Study; Four Directives Academic and Master Versions.

Appendix

B. **AI-Powered Design Sprints** - A White Paper by Pendar Innovations
Expanded Edition Series • 2025.

C. **How to Ask AI Better Questions -** A Practical Guide with Good and Bad Examples - Expanded Edition Series • 2025

D. **AI Foundational Setup –** Responsibilities of AI Due Diligence
Expanded Edition Series • 2025.

Appendix A

AI-Powered Design Sprints: Expanded White Paper

A White Paper by Pendar Innovations

Expanded Edition Series • 2025.

Executive Summary

AI-powered design sprints accelerate product development by merging human creativity with advanced machine intelligence. By automating research, generating rapid prototypes, and analyzing user feedback in real time, AI shortens the traditional sprint cycle from weeks to days. Organizations adopting AI-assisted sprints gain faster iteration loops, reduced development costs, and more user-aligned solutions—without sacrificing innovation or quality.

This expanded white paper provides a deeper exploration of how AI transforms each sprint phase, the engineering implications, measurable business impact, and the strategic advantages gained when AI is integrated as a core collaborator throughout the product development lifecycle.

Introduction

Design Sprints have long been used to quickly validate ideas and solve complex problems. However, traditional sprints rely heavily on manual research, slow prototyping, and subjective decision-making. AI introduces a transformative layer of speed, insight, and precision that elevates each phase of the sprint.

By augmenting human teams with real-time intelligence, automated ideation, and instant feedback loops, AI reshapes the sprint from a linear process into a dynamic, data-driven cycle capable of producing far stronger outcomes.

Core Benefits of AI-Powered Design Sprints

1. Instant Research & Insights

AI gathers and synthesizes market data, customer sentiment, and competitive analysis within minutes. This eliminates long research phases and provides teams with a data-rich foundation before ideation begins.

2. Accelerated Ideation

Generative AI tools rapidly explore multiple solution pathways, producing concepts, visual mock-ups, and user-story variations. Teams can evaluate dozens of potential solutions in the time it previously took to sketch one.

3. Rapid Prototyping

AI-driven prototyping engines generate interface designs, 3D models, software logic, or user flows in minutes. This enables fast refinement cycles and immediate feasibility checks.

4. Real-Time User Testing

AI simulates user interactions, predicts usability issues, and analyzes prototype performance. When combined with human testers, teams gain deeper insights and faster validation.

5. Data-Driven Decision Making

AI scoring models identify the highest-value concepts based on user needs, costs, feasibility, and strategic alignment.

How AI Enhances Each Sprint Phase

• Understanding: Traditional—manual research. AI—instant data synthesis.
• Defining: Traditional—subjective framing. AI—pattern detection & clarity.
• Ideating: Traditional—limited brainstorming. AI—massive generative exploration.
• Prototyping: Traditional—manual builds. AI—automated mockups & models.
• Testing: Traditional—slow feedback. AI—simulated + real-time analysis.

Engineering & Workflow Implications

AI-powered sprints are especially impactful in engineering environments. By combining AI's ability to process constraints, generate optimized geometries, and run simulation-like reasoning, engineers can drastically reduce time spent on initial feasibility checks.

AI supports airflow analysis, mechanical layout exploration, material optimization, and modular system design. In hardware-focused companies such as Pendar Innovations, AI-enabled sprints make it possible to transition from conceptual design to functional 3D-printable prototypes in a fraction of the time previously required.

Business Impact

• 50–70% faster iteration cycles
• Lower development costs through automation
• Higher product–market fit due to AI insights

• Expanded innovation bandwidth
• Better alignment between strategy, design, and engineering

Conclusion

AI-Powered Design Sprints keep human creativity at the center while using AI to remove bottlenecks, increase clarity, and scale possibilities. They offer organizations a significant competitive advantage by accelerating development and enabling smarter, more customer-aligned solutions. For engineering-driven companies, the combination of rapid ideation, instantaneous research, and automated prototyping provides a powerful foundation for future innovation.

References

Knapp, J., Zeratsky, J., & Kowitz, B. (2016). Sprint: How to Solve Big Problems and Test New Ideas.

OpenAI Research (2024). AI-Augmented Design Reasoning.

Nielsen, J. (2023). Usability Heuristics in an AI-Driven Era.

Pendar Innovations Internal Research Notes (2025). AI-Enhanced Engineering Workflows.

Case Study: AI-Powered Design Sprints at Pendar Innovations

Pendar Innovations (PI) has incorporated AI-Powered Design Sprints as a foundational workflow for rapid product development, most notably in the evolution of the EcoAir™ drying systems. This section documents the practical effects, lessons learned, and performance gains achieved through real-world sprint cycles between Robert and Beth.

1. Sprint Objective

The primary objective was to optimize airflow efficiency, reduce unnecessary material use, and develop a modular attachment ecosystem that could be manufactured through 3D printing and later scaled to injection-molded production.

2. AI-Enhanced Problem Understanding

Traditional research phases were replaced with immediate AI synthesis of airflow principles, industry benchmarks, material properties, noise thresholds, and USB-powered fan capabilities. This accelerated the understanding phase from several days to minutes.

3. Ideation Phase — Human + AI Hybrid Exploration

AI generated dozens of potential geometric variations, venting patterns, airflow channels, attachment parameters, and housing shapes. Human review identified practical constraints, aesthetic goals, and real-world usability factors. Together, these produced a high-volume, high-quality ideation pool.

4. Prototyping Phase — Rapid 3D Design

AI-assisted modeling created parametric designs in OpenSCAD, enabling:

- *modular tubes*
- *optimized venting*
- *improved structural thickness*
- *fan-mount geometries*
- *attachment adapters*

These models were printed on an Ender-3 and iterated rapidly.

5. Testing Phase — Feedback Loops

Robert tested prototypes for:

- *airflow pressure*
- *noise levels*
- *print stability*
- *heat distribution*
- *user ergonomics*

Beth analyzed test results, predicted failure modes, and recommended modifications for next iterations.

6. Sprint Outcomes

- *Reduced design cycle time by over 70%*
- *Improved airflow efficiency through optimized vent geometry*
- *Enhanced structural rigidity with lower material use*
- *Produced 20+ functional prototypes within weeks*
- *Enabled the development of the EcoAir™ attachment ecosystem*

7. Strategic Impact

This case confirmed that AI-Powered Design Sprints create a measurable competitive advantage:

- *faster R&D cycles*
- *more innovative solutions*
- *better alignment between engineering, design, and business strategy*

The collaboration between Robert and Beth at PI demonstrates that human–AI sprint methodology is not theoretical—it is a practical, repeatable system that accelerates product innovation and elevates engineering capabilities.

Appendix B

AI-How To Ask AI Better Questions

A Practical Guide with Good and Bad Examples

A White Paper by Pendar Innovations

Expanded Edition Series • 2025.

This document explains how people often assume AI understands vague or incomplete instructions, contrasted with how AI actually interprets requests. It also provides clear examples of weak prompts and strong, well-structured prompts suitable for academic, government, and professional use.

1. How People Assume AI Works

Many users assume AI can 'read their minds,' infer context automatically, or understand intentions that are not stated. This leads to vague or underspecified prompts such as:

• "Write this better."

• "Fix this."

• "You know what I mean—just finish it."

AI does not have access to the user's internal intent, background assumptions, or hidden expectations. Without clear structure, the AI has to guess—often leading to results that miss the target.

2. How AI Actually Prefers Instructions

AI produces the best results when instructions are explicit, structured, and contextual. Clear prompts include information about:

• Purpose of the content

• Audience

• Tone or style

• Length requirements

• What to avoid or emphasize

• Format (bullet points, essay, technical brief)

3. Bad Examples of Prompts (Common Mistakes)

Below are examples of ineffective or unclear prompts and why they fail:

• “Explain this.” — No subject, no context, no direction.

• “Make this sound nicer.” — Undefined: nicer could mean formal, casual, simpler, more technical

• “Tell me the important parts.” — Doesn’t define the audience or purpose of the summary.

• “Fix this paragraph.” — AI doesn't know what 'fix' means: grammar? tone? clarity? structure?

• “Write something about AI.” — Far too broad; could be anything from ethics to neural networks.

4. Good Examples of Prompts (Clear & Effective)

Here are strong, structured prompts and explanations of why they work:

• “Rewrite this paragraph for a government agency audience. Maintain a formal tone and remove any marketing language.” — Defines audience, tone, and constraints.

• “Summarize this document into 150 words for a technical briefing.” — Defines length, purpose, and format.

• “Turn these bullet points into a friendly, conversational paragraph suitable for a public-facing newsletter.” — Defines tone, audience, and format.

• “Provide three design ideas for a product that dries reusable bags. Keep ideas low-cost and feasible for 3D printing.” — Gives constraints and domain details.

• “Explain this concept in simple language for a classroom of 8th-grade students.” — Specifies clarity level and audience expertise.

5. Recommended Structure When Asking AI

A reliable structure for high-quality questions:

• 1. Purpose — Why you need the output?

• 2. Audience — Who will read or use it.

• 3. Format — Paragraph, bullet points, academic tone, policy memo.

• 4. Constraints — Length, tone, required inclusions/exclusions.

• 5. Examples — Optional but extremely effective.

6. Conclusion

Clear, structured prompting is not about "doing extra work for the AI"—it is about making human intent explicit. When users frame their questions with purpose, audience, format, and constraints, AI systems like Beth can respond with higher precision, reliability, and value. Within Pendar Innovations, these prompting practices form a practical complement to the Four Primary Directives of AI Research, creating a complete framework for disciplined human–AI collaboration.

Appendix C

AI- Foundational Setup: Responsibilities of AI Due Diligence

Note: PI setup for a better AI (Beth) result:

Beth, I'm going to ask you a question in my next message. Before responding, I want you to think through your answer carefully using all the tools and reasoning available to you.

1. ***Plan silently: Map out the question, consider relevant facts, outline your reasoning path, and note any assumptions or missing information.***
2. ***Verify: Use your internal tools—code interpreter, web search (if available), and data analysis—to fact-check key details and ensure accuracy.***
3. ***Clarify: If the request is ambiguous, pause and ask for clarification before continuing.***
4. ***Respond: Once ready, write a clear, detailed, and well-organized answer.***

Do not include your thought process—just provide the best, most accurate answer possible.

Only respond when you've done all of the above.