ALLEN + CLARKE

GETTING STARTED WITH AI

Practical things you can do today

GETTING STARTED WITH AI:
Practical things you can do today

Artificial intelligence represents a transformational opportunity for professionals who want to harness its potential whilst staying compliant with organisational policies and data protection requirements. This document provides practical guidance on getting started with AI, focusing on actionable steps that anyone can implement immediately. The content covers the evolution of AI from basic tools to powerful assistants, shares four core principles for maximising AI outputs, and demonstrates 16 real-world use cases that deliver immediate value.

Understanding the four stages of AI development

AI development follows four distinct stages: supporter (like autocorrect), assistant (ChatGPT), coworker (call centre agents), and colleague (autonomous AI like Jarvis from Iron Man). We're currently transitioning from the assistant stage into co-worker AI, where systems can handle more complex tasks with minimal human intervention.

This transition represents a fundamental shift in AI capabilities. Currently, assistant-level AI requires human direction for each task and provides responses to specific queries. Co-worker AI, however, can manage entire processes autonomously whilst maintaining human oversight for complex decisions. The practical implications are significant – instead of waiting on hold for 20 minutes to update your bank details, a co-worker AI can answer immediately, understand natural language requests, access your account information, and complete the update seamlessly.

The evolution continues towards colleague-level AI that will operate independently across multiple domains, but we're still years away from this reality. Understanding this progression helps organisations plan their AI adoption strategy and set realistic expectations for current capabilities.

Hallucination is a feature, not a bug – AI's tendency to generate convincing but potentially inaccurate information is actually what enables creativity and flexible thinking. However, this characteristic becomes more dangerous as we move towards co-worker AI, where errors can compound without human intervention.

"this is the worst AI will ever be – adoption is not optional,"

Four stages OF AI

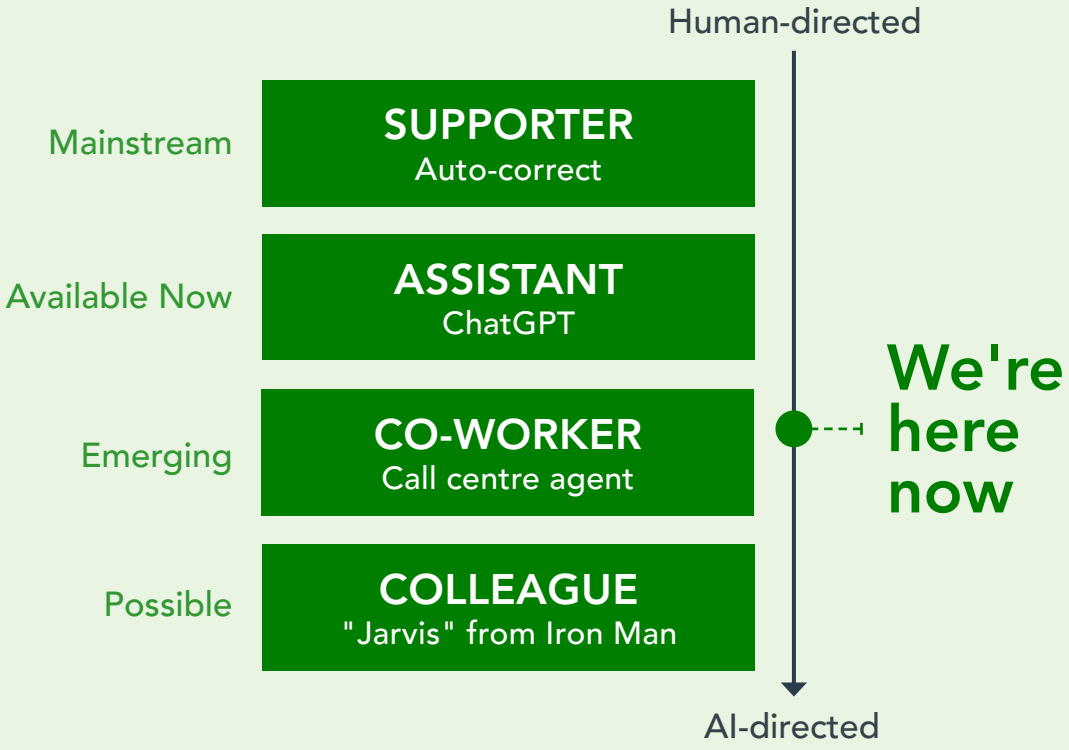

RECOMMENDATIONS:

- Understand that this is the worst AI will ever be – it only improves from here
- Expect AI in customer service and administrative tasks, push yourself to adopt and trial these
- Recognise that human oversight remains critical as we move between stages

The Four Principles for Getting the Best AI Outputs

The four principles work together like a pyramid: Model Selection (choosing the right tool), Prompt (giving clear instructions), Contextual Data (providing helpful information), and Review and Refine (human oversight and improvement). These principles transform AI from a basic tool into a powerful assistant capable of producing professional-quality outputs.

Model Selection forms the foundation – different AI tools excel in different areas, and choosing the wrong model significantly limits your outputs regardless of prompt quality. However, a well-crafted prompt can overcome minor model limitations, making prompting the most critical skill to develop.

Contextual Data amplifies every other principle by providing the AI with relevant background information, examples, style guides, and domain-specific knowledge. This transforms generic outputs into tailored, organisation-specific content that aligns with your requirements and standards.

Review and Refine ensures quality and accuracy whilst allowing for iterative improvement. This stage prevents the delivery of hallucinated content and enables continuous refinement of both prompts and processes for repeated use cases.

The demonstration revealed the dramatic impact of following these principles systematically. A simple prompt in Perplexity produced a basic 500-word strategy outline, whilst a detailed prompt in Claude with proper contextual data generated a comprehensive 3,000-word strategy complete with specific funding targets, stakeholder engagement plans, and implementation timelines. The difference wasn't just in length but in depth, specificity, and practical utility.

Poor AI use v good AI use.
Click here to watch the demonstration.

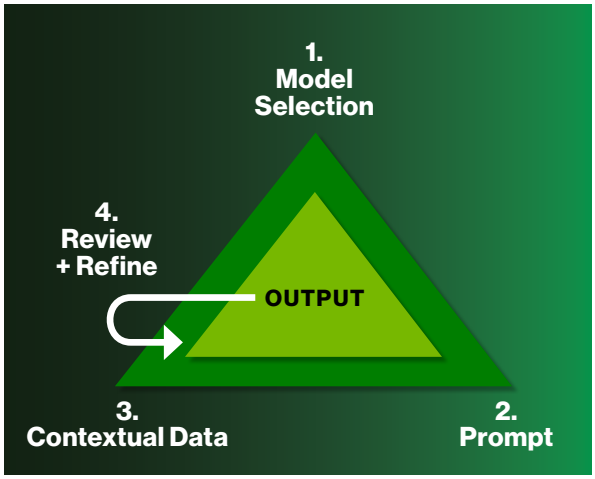

RECOMMENDATIONS:

- Always follow all four principles together for best results
- Start with model selection – choose the right tool for your specific task
- Never skip the review and refine stage, even for simple tasks

ESSENTIAL TIP:

Think of Generative AI as a less experienced junior colleague who needs clear instructions and context to perform well.

"A good prompt will generally overcome not using the best model. If you're using ChatGPT, Copilot, Google Gemini, Claude with a really good prompt, then the output will generally be pretty good."

Choosing the Right AI Models and Tools

As of April 2025, these are our favourite AI tools and use cases.

TOOL	WHY WE LIKE IT	BEST USE CASES
ChatGPT	The most versatile and user-friendly option. Has a wide range of more advanced features for more complex tasks when you need them.	Ideation + brainstorming, writing, image generation, able to export outputs into Word or Excel, internet based research.
Perplexity	Better than Google at internet searches, including formatting search results and follow up questions.	Detailed internet research, formatting of search results, verifying facts, exploring new topics.
Claude	Writes like a human with the ability to set custom writing styles. Great at understanding nuance and long documents.	Drafting, editing, reviewing all types of content.
Copilot	Helps resolve privacy concerns by working within your existing Microsoft environment. Accesses SharePoint/OneDrive for context-aware tasks.	Project planning, task tracking, document prep.
Gemini	Great if you use Google suite for personal or work as it connects to G Suite. Has amazing screen share mode – your personal tech support.	Drafting and editing Google Docs, slides, etc. real-time support, summarising spreadsheets
Notebook LM	Accurately summarises and reviews documents with citations. Has a cool podcast mode so you can discuss the documents you've added.	Reviewing reports, research or contracts, extracting key points, building FAQs
Udio	A fun way to create your own songs with your own lyrics	Any moment you might need your own creative tune – a workshop or a webinar!
Originality	Best tool for detecting AI-generated content.	Verifying written content authenticity down to individual sentences
Whisper	Transcribes speech clearly, even from tricky recordings. Nails Te Reo and NZ accents.	Transcribing interviews, meetings, webinars, or voicemails

ChatGPT serves as the ideal starting point – functioning like the Arial font of AI, it's versatile, reliable, and handles most general tasks effectively. Its advanced features include image analysis, voice mode, and export capabilities to Word and Excel, making it suitable for diverse applications from brainstorming to content creation.

Perplexity revolutionises search capabilities by providing direct answers with proper citations rather than endless blue links. It excels at formatting search results into tables, comparing options, and following up with related questions. The ability to exclude specific sources (like Reddit) or focus on particular domains makes it superior to traditional search engines for professional research.

Claude demonstrates exceptional writing abilities with natural, human-like output and sophisticated understanding of context and nuance. It excels at long-form content, editing, and maintaining consistent tone throughout lengthy documents. Its ability to set custom writing styles makes it invaluable for organisations with specific brand voices.

Enterprise solutions like Copilot and Gemini offer significant advantages for organisations already using Microsoft or Google ecosystems. These tools access your existing files, emails, and documents as contextual data, providing more relevant outputs whilst maintaining better security controls. The integration eliminates the need to manually upload context, streamlining workflows considerably.

Specialised tools serve specific purposes: Whisper for transcription (particularly strong with New Zealand accents and Te Reo), NotebookLM for document analysis with podcast-style summaries, and Originality for detecting AI-generated content. Each tool represents best-in-class capabilities for particular use cases.

RECOMMENDATIONS:

- Start with ChatGPT for general tasks and learning
- Switch from Google to Perplexity for search – this simple habit change introduces you to AI-powered search
- Deploy enterprise versions for better data security and controls

PRACTICAL TIP:

Rather than trying to learn every AI tool, focus on mastering a few key models that suit your most common tasks. The session provided a comprehensive comparison showing why different tools excel in different areas.

"ChatGPT is kind of like the Arial font in Word. It's the market leader, it's normally the default, it's a great place to start from. However, there's lots of reasons that you might want to change that font."

Mastering the Art of Prompting

Effective prompts follow a specific structure: Persona (who the AI should be), Context (background information), Task (what you want done), and Output (how you want the result formatted). This framework ensures AI receives comprehensive instructions comparable to briefing a junior colleague on a complex assignment.

The quality differential between simple and detailed prompts is dramatic. Basic prompts like "create a funding strategy" yield generic, surface-level content lacking specificity or actionable insights. Conversely, detailed prompts that establish the AI as "an experienced government procurement advisor" working on"a comprehensive funding strategy report for a community-based environmental protection group in 2025" produce sophisticated, targeted outputs with specific funding recommendations, stakeholder engagement strategies, and implementation timelines.

Advanced prompting techniques include asking AI to generate clarifying questions before beginning tasks, which often reveals gaps in your initial brief and leads to more comprehensive outputs. Using AI to critique and improve your own prompts creates an iterative refinement process that dramatically enhances results over time.

Contextual integration within prompts involves incorporating specific requirements such as writing style, target audience characteristics, compliance considerations, and desired outcomes. Professional prompts often exceed 250 words and read like detailed project briefs, establishing clear success criteria and constraints.

The meta-approach of using AI to create better prompts proves highly effective – AI understands prompt structure and can suggest improvements, alternative phrasings, and additional contextual elements that enhance output quality.

The quality of your prompt directly determines the quality of your output. Simple prompts produce basic results, while detailed, structured prompts generate comprehensive, professional-quality work. The session showed how prompts can be 250+ words long and include specific requirements for tone, structure, and formatting.

4 prompting tips:

#1 – Brief your AI like you would a junior colleague: provide clear instructions and detailed context.

#2 – Use the flow: Persona Context Task Output.

#3 – Ask the AI to ask clarifying questions.

#4 – Use AI to help refine or rewrite your prompts.

Prompt example

Simple

Create a comprehensive strategy that's designed to get more funding for my community-based environmental protection group.

Detailed

Persona:

You are an experienced government procurement advisor with a strong track record of securing contracts and funding for community-based environmental protection groups in Aotearoa New Zealand.

Context:

You are preparing a comprehensive funding strategy report for a community-based environmental protection group in 2025. The report will be presented to the advisory board, which includes:

- **Ngahuia:** Chair, highly respected iwi leader focused on biodiversity.
- **Sam:** Treasurer, former Department of Conservation finance manager.
- **Priya:** Board member, sustainability strategist passionate about urban regeneration.

The political, environmental, and funding environment of New Zealand in 2025 must be considered, particularly:

- Current government environmental priorities
- Government funding trends for community and iwi-led environmental initiatives
- Stakeholder engagement practices appropriate for senior public sector and iwi leaders

Task:

Develop a comprehensive funding strategy tailored to securing additional support from the New Zealand Government in 2025, which must include:

1. Recommended Targets:

Identify and recommend government departments, agencies, and funding programmes most likely to fund biodiversity protection, urban regeneration, and community-led sustainability projects.

2. Stakeholder Engagement Strategy:

For each advisory board member (Ngahuia, Sam, and Priya), recommend specific roles they could play in engagement activities.

Suggest how their unique backgrounds and strengths can be leveraged to build credibility and partnerships with government stakeholders.

3. Proposal Alignment Guidance:

Provide tips for aligning funding proposals with current environmental, social, and partnership priorities of the New Zealand Government.

Highlight any specific policies, frameworks, or language that should be mirrored in proposals to increase alignment and approval chances.

Output:

Deliver a fully detailed strategy report outline (structured for a board document), written in plain, professional English, with clear and informative headings. Keep tone direct, confident, and action-oriented, following Allen + Clarke writing principles.

Report Structure:

Format the output as a Strategy Report Outline suitable for board approval. Include the following:

1. Executive summary
2. Key funding targets
3. Stakeholder engagement plan
4. Proposal alignment recommendations
5. Next steps

Also provide a bullet-point list summarising the key steps for board members at the start of the report.

ESSENTIAL TIPS:

- Brief your AI like you would a junior colleague with clear instructions and detailed context
- Always end complex prompts with "ask me clarifying questions before you complete the task"
- Use AI to review and improve your own prompts

RECOMMENDATIONS:

- Use the Persona → Context → Task → Output flow for all initial prompts
- Ask AI to ask you clarifying questions before completing tasks
- Use AI to help you create better prompts – it's quite meta but highly effective

Working Safely with Contextual Data

Contextual data encompasses any information that helps AI understand your task – previous reports, style guides, examples, datasets, organisational knowledge, or industry-specific background information. This data transforms generic AI outputs into tailored, relevant content that aligns with your specific requirements and standards.

The security implications are significant and require careful consideration. **Free-tier AI tools present the highest risk** as uploaded data will almost certainly be used for model training or stored indefinitely. Enterprise accounts provide the strongest protections through dedicated computing environments, data encryption, and explicit non-training agreements. Paid personal accounts offer intermediate protection, particularly when model training is disabled in settings.

Advanced techniques for contextual data management include creating synthetic datasets using AI itself rather than uploading sensitive information. For example, using AI to research current funding opportunities and feeding that research back as context for strategy development, rather than uploading internal organisational data.

Publicly available information remains safe to use regardless of AI platform. Most models you use will have been trained on the entire internet and more. So, the simple rule of "if you can find it online" means it's generally safe for you to use as contextual data. For example: regulations, published research, industry guidelines, and public style guides can enhance outputs without security concerns. The New Zealand Government plain English guidelines, for instance, can help ensure all outputs use appropriate local language conventions.

Data classification becomes critical – establish clear boundaries between information suitable for free tools, paid platforms, and enterprise systems. Create organisation-wide policies that specify which data types require which security levels, enabling teams to use AI effectively whilst maintaining appropriate protections.

More context leads to better outputs, but you must be extremely cautious about what information you share. The session outlined the hierarchy of safety: enterprise accounts are safest, paid accounts with model training turned off are better than free accounts, and publicly available information is always safe to use.

CRITICAL TIP:

Establish clear organisational policies about what data can and cannot be shared with AI tools, and train your team on these boundaries.

RECOMMENDATIONS:

- Only use publicly available information with free AI tools
- Invest in enterprise accounts for organisational use
- Turn off model training in settings for all personal accounts
- Create synthetic contextual data using AI rather than uploading confidential information

The Importance of Review and Refine

The final principle stressed that **AI output always requires human oversight**. This stage involves checking for accuracy, improving clarity, and ensuring the output meets your specific needs. For simple tasks, this can be done within the same chat window, but complex or repeated tasks require systematic review processes.

AI can, and does, hallucinate – creating convincing but false information. The review and refine stage is your quality control, ensuring outputs are accurate, appropriate, and useful. This stage also includes iterating on outputs to improve them progressively.

RECOMMENDATIONS:

- Never deliver assume AI output is correct without human review, regardless of task complexity
- Build review checkpoints into your workflow by design
- Use multiple AI tools to cross-check important outputs
- Create template review processes for repeated tasks like meeting minutes

ESSENTIAL TIP:

For repeated tasks, invest time in perfecting your prompt and review process – the initial effort pays dividends over time.

Practical AI Use Cases for Immediate Implementation

Sixteen specific AI applications demonstrate how artificial intelligence can assist with diverse tasks from creative writing and meal planning to professional document analysis and IT support. Live demonstrations of the **MasterChef** (recipe creation from fridge photos) and **Workshop Wizard** (summarising whiteboard notes) use cases showed how AI handles real-world tasks with impressive accuracy.

These applications span both personal and professional contexts, enabling users to build confidence through low-risk, high-value activities. The key is starting with applications that address real problems in daily workflows, mastering them completely, then expanding to other use cases.

Work Applications provide immediate productivity gains: the **IT Help Desk** delivers step-by-step technical support with screen sharing capabilities, enabling users to resolve technical issues without lengthy support calls. **Meeting Minutes** functionality transforms transcripts into actionable summaries, eliminating manual note-taking and ensuring consistent documentation standards.

Research Assistant capabilities conduct in-depth research with proper citations, dramatically reducing the time required for comprehensive information gathering. The **Workshop Wizard** analyses photos of whiteboards and sticky notes, converting visual brainstorming outputs into structured digital formats.

Home Applications offer practical daily support: the **MasterChef** feature creates recipes from photos of available ingredients, addressing meal planning challenges and reducing food waste. **Storyteller** generates personalised bedtime stories for children, incorporating their names and interests for engaging content.

Trip Planner functionality organises family holidays based on preferences and budget constraints, whilst

Language Coach provides conversation practice in different languages. Additional applications include **Green-thumb Guru** for plant care guidance and **Budget-buddy** for financial analysis.

The breadth of applications demonstrates AI's versatility across different domains, from creative tasks to analytical work. Each use case represents an opportunity to experience AI capabilities firsthand whilst building practical skills.

IMPLEMENTATION TIP:

Start with one use case that addresses a real problem in your daily workflow, master it completely, then expand to other applications.

AI USE CASES

AT HOME	AT WORK
Master chef – take photos of your pantry and fridge and ask ChatGPT to create a recipe with cooking instructions based on the ingredients you have.	I.T. help desk – can't figure out how to do something on your phone or computer? Use Gemini, share your screen and ask it to help you solve the issue. It will 'look' at your screen and talk you through it
Storyteller – Have your AI assistant create a unique bedtime story based on your kids favourite adventure. Ask it to influence your family to make it extra personal	**Research Assistant** – use Elicit or Litmaps for in-depth research with citations
Custom Podcast – got a topic you want to deep dive into? Use Notebook LM's podcast mode to create an entertaining and in-depth podcast which you can interrupt and ask questions	**Meeting Minutes** – take your meeting transcripts and instantly create useful summaries and actions
Musician – Use Udio to create a custom song for your party or webinar — a fun, personalised experience.	**Ideation** – generate and critique ideas to develop the best potential solutions
Trip planner – use your assistant to plan your next family day out or holiday. Have it rank activities based on your preferences, budget and location	**Tough talker** – Role-play difficult conversations using ChatGPT's voice model — great for preparation and confidence-building.
Language coach – have ChatGPT role-play daily dialogues so you can finally order tapas in decent Spanish.	**Data wrangler** – use ChatGPT to summarise and categorise large datasets
Green-thumb guru – Snap a sick houseplant and let ChatGPT-Vision diagnose the issue and suggest care steps.	**Content Reviewer** – use Claude to critique and enhance your writing
Budget-buddy – provide Claude a CSV of your account statement (anonymised) and get a detailed breakdown of where the dollars wander and how you could improve your finances.	**Workshop Wizard** – use ChatGPT to categorise, write up and summarise your workshop outputs from images of your whiteboard or post-it notes

Be wary of privacy, IP, and data protection and always use AI in accordance with your organisations policies.

Data Protection and Privacy Considerations

Organisational AI policies are essential before any team member begins using AI tools professionally. The biggest risk isn't AI itself, but people using free-tier tools without proper guidance or controls. Clear policies, training, and appropriate tool selection can mitigate most risks whilst enabling teams to benefit from AI capabilities.

Get started with your AI policy using this template

Policy frameworks must address multiple dimensions of AI usage, from data classification to tool selection and user training. Organisations need explicit guidelines about what information can be shared with different AI platforms, establishing clear boundaries between free, paid, and enterprise tools.

Risk management requires proactive approaches rather than reactive restrictions. Teams who receive proper guidance and approved tools are far less likely to create security incidents than those who use AI tools without organisational support. The challenge lies in balancing innovation with protection.

Training programmes should cover practical scenarios rather than abstract policies. Users need specific examples of what constitutes appropriate versus inappropriate AI usage, with clear trigger points that indicate when additional caution is required.

Enterprise solutions provide superior security controls through dedicated computing environments, data encryption, and explicit non-training agreements. However, smaller organisations can achieve adequate protection through careful tool selection and user education.

Data governance principles apply equally to AI tools as to other technology platforms. The same considerations around personal information, confidential data, and intellectual property that govern traditional systems should inform AI usage policies.

KEY REQUIREMENTS:

- Establish clear organisational AI policies before implementation
- Use enterprise-tier tools for professional work when possible
- Provide team training on safe AI usage practices
- Implement express permission processes for sensitive data use

RECOMMENDATIONS:

- Assume people are already using AI – provide guidance rather than prohibition
- Create trigger points that help users understand when extra caution is needed
- Consider enterprise accounts for better data controls and security
- Regular review of AI usage and policy updates as technology evolves

CRITICAL CONSIDERATION:

Balance innovation with protection – overly restrictive policies may push AI use underground, creating greater risks.

"If you're not telling them what they can do and training them what's OK for your organisation, people are going to be going to the free tier of ChatGPT and putting whatever in there. That's what we want to get away from."

Quality Control and Avoiding Common Pitfalls

AI tools have inherent biases from their training data, and users must be aware of these limitations whilst implementing appropriate quality controls. Successful AI implementation requires understanding both capabilities and limitations, with quality control involving multiple verification methods, appropriate human oversight, and recognition that AI is a powerful assistant, not a replacement for human expertise.

Content detection tools like Originality.ai can identify AI-generated text when transparency or verification is required. However, the focus should be on quality and accuracy rather than simply detecting AI involvement.

Skills atrophy represents a long-term risk if users become overly dependent on AI tools without maintaining their own expertise. The goal should be using AI to enhance productivity whilst preserving human capabilities and decision-making authority.

Quality Control Methods:

- Use multiple AI tools to cross-verify important information
- Implement systematic fact-checking for critical outputs
- Create human checkpoints throughout workflows
- Use tools like Originality.ai to detect AI-generated content when needed

Common Issues to Avoid:

- Over-reliance leading to skill atrophy
- Accepting AI outputs without verification
- Using inappropriate tools for specific tasks
- Sharing sensitive information through free platforms

Conclusion

AI represents a transformational opportunity for individuals and organisations willing to approach it systematically and safely. The four principles – model selection, quality prompting, contextual data, and review and refine – provide a reliable framework for achieving consistently valuable outputs. Success depends on starting with simple, low-risk applications and building expertise progressively.

Current AI capabilities offer immediate practical value for professionals across diverse industries and roles. The technology is ready for implementation today, provided users follow proper security protocols and maintain appropriate human oversight.

Early adoption provides competitive advantages as AI capabilities continue to expand rapidly. Organisations and individuals who master fundamental principles now will be better positioned to leverage future developments and innovations.

Bottom line: This is the worst AI will ever be. Early adopters who master these fundamental principles now will have significant advantages as AI capabilities continue to expand. The technology is ready for practical implementation today, provided you follow proper security protocols and maintain appropriate human oversight.

The most important next step is simply to start – choose one use case from the 16 demonstrated, ensure it complies with your organisational policies, and begin building your practical AI experience today.

Key Tips Summary

1. **Understand the 4 Stages of AI** – This is the worst AI will ever be; it's only going to become more capable
2. **Get an Account** – Enterprise is best, but at the very least use paid accounts with model training turned off
3. **Ask AI How to Use AI** – Use AI itself to help you learn by prompting with "Help me create a prompt to..." or "What can I use you for?"
4. **Start Playing** – The best way to learn AI is by using it; begin with simple tasks and you'll gain confidence quickly
5. **Replace Google with Perplexity** – A simple switch that introduces you to AI-powered search with faster, more relevant results and sources
6. **Use the Right Tool for the Job** – Start with ChatGPT for general tasks, Perplexity for research, Claude for writing, and Gemini for anything Google-integrated
7. **Follow the 4 Principles for Better Outputs** – Model selection, quality prompting, contextual data, and review + refine work as a system, not individual magic wands
8. **Prompt Like a Pro** – Use the structure: Persona → Context → Task → Output, and ask for clarifying questions to improve results
9. **Feed it Context, But Keep it Safe** – More context equals better outputs, but avoid using sensitive, personal, or confidential information unless you're using paid/enterprise tools
10. **Always Review + Refine** – Treat AI like a junior colleague: useful and fast, but still needs your oversight for quality and accuracy
11. **There's an AI for That** – From cooking dinner to summarising whiteboard photos, explore real-world use cases to unlock AI's practical power
12. **Check Your Organisation's AI Policy** – Before implementing any AI tools professionally, ensure compliance with your organisation's guidelines and data protection requirements
13. **Start with Low-Risk, High-Value Tasks** – Begin with applications that won't cause problems if they go wrong, but provide clear benefits to your workflow
14. **Build Human Checkpoints by Design** – Structure your AI workflows to include regular human oversight and quality control points
15. **Use AI to Improve AI** – Employ different AI tools to review and refine outputs from other AI tools for better quality control
16. **Focus on Augmentation, Not Replacement** – Use AI to enhance your expertise and productivity while maintaining your role as the subject matter expert and decision maker