

WORKSHOP 1 OF 4 · STUDENTSTUDYNOTES

AI Foundations + Image & Video Prompt Mastery

AI Without Fear

Your comprehensive guide to Workshop 1. Everything covered in the room — expanded into plain English with context, examples, and step-by-step guidance you can use straight away.

DOUGLAS MIZZI · AIWITHOUTFEAR.COM.AU · 2025

Table of Contents

Chapter 1 · Welcome & Day Overview	4
The Day at a Glance	4
Who This Workshop Is For	5
Chapter 2 · Understanding AI — The Foundations	5
How AI Actually Works	5
Chapter 3 · The Abacus.AI Platform	7
Chapter 4 · Setting Up Your Projects Correctly	8
Projects vs Chats vs Agent Sessions	8
Naming & Organising Your Projects	9
The Company Blueprint	10
Chapter 5 · Deep Agent Project Management	12
Why Deep Agent Asks for a New Conversation	12
Chapter 6 · Why Prompting Matters	14
Chapter 7 · The Anatomy of a Great Prompt	15
The Five Parts: Role · Task · Context · Format · Tone	15
Morning Recap	17
Chapter 8 · AI Image Generation — Frameworks & Practice	18
How AI Turns Words into Pictures	18
The Image Prompt Framework (S·S·C·L·M·T)	19
Real-World Case Study: The Solar Flyer	20
Image Generation Models — Your Creative Arsenal	22
Image Generation Best Practices	27
Chapter 9 · AI Video Generation — Frameworks & Practice	
From Still Images to Moving Stories	
The Video Prompt Framework	

Video Generation Models

Technical Specs, Editing & Platform Tools

Chapter 10 · Your Journey Continues

Welcome & Day Overview

Everything you're going to learn — and why it matters for you today.

AI Foundations + Image & Video Prompt Mastery

Welcome to Workshop 1 of the AI Without Fear series — a full-day, hands-on experience that takes you from "I've heard of AI" to "I can actually use this in my business or career today." This workshop has been designed with one person firmly in mind: you. Whether you own a small business, work in a professional role, or are exploring what AI could mean for your future, today's content is practical, jargon-free, and immediately applicable.

By the end of today you will have built real outputs — an AI-generated image and a short AI-generated video — using the same platform tools that are changing the way people work right now. You will also leave with a framework for prompting that makes everything else you do with AI dramatically more effective.

The Day at a Glance

The day is structured into a morning session focused on foundations, the platform, and text prompting — and an afternoon session dedicated entirely to creative AI: image generation and video generation. Here is what we cover:

- 1 Morning — AI Foundations:** What AI actually is (no jargon), how Abacus.AI gives you access to every major AI model in one subscription, and how to set up your projects correctly so the AI works with you rather than against you.
- 2 Morning — Prompting Mastery:** The single biggest skill you can develop with AI — how to ask for what you actually want. We'll cover the five-part prompt framework that turns average results into gold.

- 3 **Afternoon Part 1 — AI Image Generation:** How AI turns words into images, the six-dial framework for image prompts, a real-world case study from a local business, and a guided hands-on activity where you create your first AI image.
- 4 **Afternoon Part 2 — AI Video Generation:** How video AI works, the six-element video prompt framework, a tour of the video models available inside Abacus.AI, and a guided hands-on activity where you generate your first AI video clip.

SLIDE 3 · WHO THIS IS FOR

Who This Workshop Is For

This workshop was built for a very specific kind of person — not a developer, not a tech enthusiast, but someone who has heard about AI and wants to understand what it can actually do for them. If you recognise yourself in any of the following, you are exactly who this course is designed for:

- **Small business owners and sole traders** who want to do more without hiring more — using AI to handle writing, images, marketing materials, and content creation.
- **Career changers and professionals** who want to stay relevant and add AI skills to their toolkit before the job market makes it a requirement.
- **Marketers, educators, and service providers** who need to produce creative content regularly and want to do it faster and better.
- **Curious learners** who want to move past the hype and understand what AI genuinely can — and cannot — do.

IMPORTANT NOTE

No prior technology experience is required. If you can type a message on your phone, you already have every skill needed to follow along with everything in this workshop. The language will always be plain. If a technical term appears, it will be explained immediately.

CHAPTER 2 · SLIDE 4

Understanding AI — The Foundations

What AI actually is. No magic. No mystery. Just patterns at scale.

SLIDE 4 · WHAT IS AI?

How AI Actually Works — Three Plain-English Analogies

Before you can use AI confidently, it helps to understand what it actually is — not in a computer-science way, but in a practical, everyday way. Here are three analogies that each capture a different aspect of how AI works:

The Master Chef Analogy

A master chef has tasted thousands of dishes, read thousands of recipes, and cooked thousands of meals. When you say "make me something with chicken, lemon and rosemary," they don't look anything up — they draw on everything they've ever learned to produce something new and relevant. AI works the same way: it has been trained on enormous amounts of text, images, code, and data, and it combines that knowledge to respond to whatever you ask.

The Librarian Analogy

Imagine a librarian who has read every book in a library containing millions of volumes — and has perfect recall. When you ask a question, they don't search the shelves; they draw from memory and synthesise an answer from multiple sources at once. That's what a large language model does: it stores statistical patterns from vast amounts of text and uses them to generate coherent, relevant responses.

The Apprentice Artist Analogy

An apprentice artist who has studied under hundreds of master painters can now paint "in the style of" almost any artist, or blend styles together, or create something entirely new. AI image models have studied millions of images and learned the visual patterns of different styles, lighting conditions, compositions, and subjects. When you describe what you want, the model paints from memory — guided entirely by your words.

The Core Reality

AI is pattern-matching at massive scale. It is not intelligent in the way humans are. It doesn't think, understand, or feel. But it has been trained on so much data that it can produce remarkably useful, creative, and convincing outputs — when you give it good instructions. The quality of your input determines the quality of your output. That's why prompting matters so much.

THE KEY INSIGHT

AI is not magic and it is not a replacement for human judgement. It is a **very powerful tool** that amplifies your own thinking and creativity — when you know how to use it. That's exactly what today is about.

The Abacus.AI Platform

Every major AI model. One subscription. One place.

What Is Abacus.AI and Why Are We Using It?

If you've heard of ChatGPT, Claude, or Gemini, you already know something about large language models — the AI tools that can write, research, analyse, and converse. The problem with using those tools individually is that each one requires a separate subscription, a separate login, and a separate bill. The best AI model for writing is different from the best one for coding, which is different again from the best one for creative thinking.

Abacus.AI solves this problem elegantly. For a single subscription of approximately \$30 AUD per month, you get access to every major AI model — including ChatGPT, Claude, Gemini, and more — plus specialised creative tools for image generation, video generation, and app building. It replaces what would otherwise cost you \$140 AUD per month or more across separate subscriptions.

\$30

AUD/month for Abacus.AI

\$140+

AUD/month for separate tools

1

Platform for everything

The Five Core Tools Inside Abacus.AI

ChatGPT (GPT-4o)

OpenAI's flagship model for writing, research, drafting emails, summarising documents, brainstorming ideas, and answering questions. The most widely used AI model in the world, accessible directly inside Abacus.AI.

Claude (Anthropic)

Known for exceptionally long, nuanced responses and careful reasoning. Particularly strong for legal documents, complex analysis, long-form writing, and tasks where accuracy and careful thinking matter most.

Deep Agent

Abacus.AI's flagship app and website builder. This is the tool that can build you a working web application or website from a plain-English description — no coding required. Deep Agent is what powers Workshop 2 (Building a Custom App).

Image AI

Access to every major image generation model — FLUX, GPT Image, Midjourney, Ideogram, Imagen and more — all in one interface. Create marketing images, logos, product photos, illustrations, and social content without a designer.

Video AI

Access to every major video generation model — Sora 2, Veo 3, Seedance, Kling, Hailuo and more. Generate short professional video clips from text descriptions or still images, ready for social media, marketing, or training content.

Abacus Claw

A browser automation agent that can navigate websites on your behalf, fill in forms, research competitors, and automate web-based tasks. This is covered in depth in Workshop 4 (Chatbot & SMS Automation).

WHY ONE PLATFORM MATTERS

When all your AI tools live in one place, you build a single habit instead of five separate ones. Your project context, history, and customisations carry across tools. And because Abacus.AI adds new models as they're released, you are always working with the latest AI — without needing to manage multiple accounts.

Setting Up Your Projects Correctly

Good organisation at the start means better results every single time.

The Three Types of Work Space in Abacus.AI

One of the first things people get confused about in Abacus.AI is the difference between a Project, a Chat, and an Agent Session. Understanding these three concepts will save you a lot of frustration and help you get much better results.

Type	What It Is	When to Use It
Project	A persistent workspace with a name, a purpose, and a set of customised instructions. The AI remembers context, uploaded files, and your preferences across multiple conversations.	Any ongoing or repeating work: your business content, a specific client, a recurring task, or a long-term build like an app or website.
Chat	A standard one-off conversation with an AI model. No persistent memory between sessions unless you are inside a Project.	Quick one-off questions, experiments, or tasks you don't need to return to.
Agent Session	A conversation with Deep Agent specifically for building apps, websites, or automations. Has its own memory and context management (explained in Chapter 5).	Any build project: apps, websites, automated workflows, complex structured outputs.

THE GOLDEN RULE

For anything you will use more than once — your business, a client, a content type — create a Project. It's the difference between starting from scratch every time and picking up exactly where you left off.

SLIDES 10–11 · NAMING CONVENTIONS

Naming and Organising Your Projects

The names you give your projects matter more than you might think. A well-named project is easy to find, easy to understand at a glance, and sets clear expectations for the AI about what kind of work happens there. A poorly named project ("stuff", "test", "my project") creates friction every time you open it.

Recommended Naming Conventions

- 1 For business content work:** Use your business name + the function. For example: "Sunrise Bakery — Marketing & Social Content". This immediately tells both you and the AI what this space is for.
- 2 For build projects:** Use a descriptive name + the type of build. For example: "Client Booking App — Deep Agent Build" or "Rudd Solar — Website Project". Adding the build type helps you know which tool to reach for when you open it.
- 3 For specific content types:** "Weekly Newsletter — Draft & Edit" or "Customer Proposal Templates". Be specific enough that six months from now you'll know exactly what's in there.

4

What to avoid: Generic names like "My Project", "Test", "Chat 1". These become confusing quickly and give the AI no context to work with.

SLIDES 12–13 · CUSTOMISING PROJECTS

Customising Your Project — The Setup That Changes Everything

Inside every Abacus.AI Project, there is a Customise section (sometimes labelled "Add Skills" or "Project Settings"). This is where you tell the AI who you are, what your business does, how you like to communicate, and what you need from it. The AI reads these instructions at the start of every single conversation in that project — meaning you never have to re-explain yourself.

Think of it like briefing a new staff member on their first day. The more detailed and clear your briefing, the better they perform from day one. The customise section is your permanent brief to your AI assistant.

What to Include in Your Project Customisation

- **Who you are:** Your name, your role, your business name.
- **What your business does:** A clear, plain-English description of your products or services.
- **Who your customers are:** Age range, typical concerns, what matters to them.
- **Your tone of voice:** Formal or casual? Warm or professional? Do you use humour? Are there words you never use?
- **Any specific instructions:** "Always write in Australian English." "Never use the word 'utilise'." "Format all emails in short paragraphs of no more than three sentences."

SLIDES 12–13 · THE COMPANY BLUEPRINT

The Company Blueprint — Your Secret Weapon

Douglas's recommended approach for populating your project customisation is what he calls the Company Blueprint method. Instead of laboriously typing out all your business information into a form, you use the AI itself to do the heavy lifting for you. Here's the process:

1

Voice dictate your rough notes: Open a chat with any AI model — ChatGPT, Claude, or whichever you prefer. Using your phone's voice dictation feature, simply talk to the AI as if you're telling a friend about your business. Ramble. Include everything. "I run a plumbing business, been going for 12 years, we do residential and light commercial, mainly in the inner suburbs, our customers are usually homeowners who want reliability more than cheap prices..." and so on. Don't worry about structure — that's the AI's job.

- 2 Ask the AI to reformat it:** Once you've dictated everything, send the AI a follow-up prompt: "Please take everything I just told you and format it into a clear company overview document that I can use to brief other AI tools. Include sections for: what we do, who our customers are, our tone of voice, our key products/services, and any specific instructions for AI assistants working with our brand." The AI will transform your rambling voice notes into a clean, structured document.
- 3 Paste it into your Project customisation:** Copy the AI-generated Company Blueprint and paste it into the Customise / Add Skills section of your Abacus.AI Project. Now every conversation in that project starts with the AI already knowing your business inside and out.
- 4 Keep it updated:** When something changes — new services, new customers, new tone preferences — come back and update your blueprint. It takes five minutes and completely changes the quality of everything you create in that project.

WHY THIS WORKS

Most people spend the first five messages of every conversation re-explaining who they are and what they need. The Company Blueprint eliminates that entirely. The AI already knows. Every output immediately sounds like you — your voice, your business, your standards.

SLIDES 14–15 · WHAT YOU CAN BUILD

What Can You Actually Build With These Projects?

Once you have your projects set up correctly with a Company Blueprint in place, the range of things you can create becomes very wide. Here are the primary use cases the course focuses on:

Websites

A fully functioning website — multiple pages, forms, booking systems — built and deployed by Deep Agent from a plain-English brief. Covered in Workshop 3.

Custom Apps

A custom web application — quote calculators, CRM systems, booking tools, training portals — built without code. Covered in Workshop 2.

Content & Marketing Materials

Emails, newsletters, social posts, proposals, flyers, ad copy — written in your voice and ready to publish, on demand.

Training Courses

Course outlines, scripts, workbooks, assessment materials — the AI can draft an entire course structure from a brief.

DOUGLAS'S ADVICE

"Good organisation at the start equals better results later. Five minutes setting up your project correctly saves you hours of frustration every week. Do it once. Do it properly."

CHAPTER 5 · SLIDES 16–21

Deep Agent Project Management

Understanding why Deep Agent sometimes asks to start fresh — and why that's a good thing.

SLIDES 16–19 · NEW CONVERSATIONS IN DEEP AGENT

Why Deep Agent May Ask You to Start a New Conversation

If you've been working with Deep Agent on a build project — an app, a website, a complex automation — you may notice that after a while, Deep Agent suggests starting a new conversation. This catches many people off guard. They worry something has gone wrong, or that their work will be lost. In fact, the opposite is true: it means everything is going to plan.

Why This Happens

Deep Agent works by reading the entire history of your conversation every single time it responds — every piece of code it has written, every instruction it has received, every decision it has made. This is how it maintains context and continuity. But as your project grows, so does the amount of information it has to read each turn. At a certain point, this becomes inefficient:

- 1 Deep Agent reads ALL code and history each turn:** Every time you send a message, it reviews the complete project context — all the code it's written, all the instructions you've given, the entire thread from the beginning.
- 2 Longer conversations cost more and run slower:** More history to read means more computing resources used. The response times get slower, and the credit usage per message increases. In a very long session this becomes both expensive and sluggish.
- 3 Deep Agent saves a summary before the new conversation:** When it suggests starting fresh, it first creates a comprehensive summary of everything important — the project status, all key decisions, what's been built, what still needs to be done, and any important details you've provided.

4

All key information carries forward: The new conversation begins with the full summary loaded. Deep Agent picks up exactly where it left off — but now it's working in a clean, efficient space rather than wading through thousands of lines of prior history.

SLIDES 19–21 · PROJECT CONVERSATION NUMBERING

Numbered Conversations — Your Project Is Not Restarting

When Deep Agent starts a new conversation for your project, it typically names it with a number: "Website Project — 2", "Website Project — 3", "App Build — 4". This numbering system is simply how you keep track of where your build is up to. Each numbered conversation is a fresh, efficient workspace that continues the same project — not a new one.

THINK OF IT LIKE CHAPTERS IN A BOOK

Your project is the book. The numbered conversations are the chapters. When Chapter 1 gets full, you move to Chapter 2 — but the story continues seamlessly. Your project has NOT failed. It has reached a natural and efficient checkpoint.

SLIDE 20 · WHY YOU SHOULD FOLLOW THIS PROCESS

The Benefits of Following Deep Agent's Lead

This is especially important for apps, websites, and any significant build project. When you follow the new-conversation workflow:

Lower Costs

By resetting the context window, each message in the new conversation uses fewer credits than it would in an overloaded long session. Your subscription goes further.

Faster Responses

A clean context window means Deep Agent can process and respond more quickly. Complex builds that were slowing down get a fresh start.

Better Focus

Deep Agent can focus on the current task rather than having to wade through thousands of prior messages to find relevant context.

Easier Management

Numbered conversations give you a clear timeline of your build project. You can go back to earlier conversations to review decisions or see what was built at each stage.

REMEMBER THIS

If Deep Agent asks to start a new conversation — **say yes**. Your project has not failed. Your work is not lost. It's a natural, efficient checkpoint. **Fresh conversation = faster, cleaner, more cost-effective work.**

CHAPTER 6 · SLIDE 22

Why Prompting Matters

Same AI. Different question. Completely different universe.

SLIDE 22 · WEAK VS STRONG PROMPTS

The Single Biggest Lever in AI

If there is one insight from this entire workshop that will change how you work with AI, it is this: *the AI is only as good as the instruction you give it*. The model itself is the same whether you ask a brilliant question or a vague one. But the output you receive will be worlds apart.

Most people who try AI and find it disappointing are not using a bad AI tool. They are giving it bad instructions. A weak prompt produces a weak result — generic, bland, one-size-fits-nobody content that feels like it could have come from anywhere. A strong prompt produces a result that feels specific, on-brand, and ready to use.

A Side-by-Side Comparison

Prompt Type	The Prompt	What You Get
WEAK	"Write me a marketing email."	Generic, templated, could be from anyone — about anything. Mentions no product, no audience, no offer, no personality. Unusable.

STRONG "Write a 120-word email from a small bakery owner to lapsed customers who haven't visited in 60 days. Warm tone, mention our new sourdough range, include a 15% off code 'WELCOMEBACK'."

Specific, on-brand, ready to send. Mentions the right product, the right audience, the right offer. Feels like the bakery owner wrote it.

The difference between those two prompts is less than thirty seconds of additional thought. But the difference in output quality is enormous. This is why learning to prompt well is not just a nice-to-have skill — it is the foundation of everything else you do with AI.

THE CORE PRINCIPLE


Weak prompts produce generic output. Specific, detailed prompts produce specific, useful output. The more context and instruction you give the AI, the more it can give back to you in return. Context is the currency of AI prompting.

The Anatomy of a Great Prompt

Five parts. Every time. Guaranteed better results.

The Five Parts of a Great Prompt

Douglas teaches a simple, memorable five-part framework that works for almost every AI task. Once you know it, you'll find yourself using it automatically. The five parts are: **Role** · **Task** · **Context** · **Format** · **Tone**. Here is each part explained in full:



The Five Parts — In Depth

1 Role — "Act as..."

Telling the AI to adopt a specific role transforms how it responds. "Act as a senior copywriter" produces very different writing than simply asking for a paragraph of text. The role sets the level of expertise, the vocabulary, and the approach. Common roles include: "Act as a marketing consultant," "Act as a friendly customer service rep," "Act as an expert accountant explaining things to a small business owner." The role is your first dial of control.

2 Task — What exactly do you want?

Be precise. "Write an email" is not a task. "Write a 150-word re-engagement email" is a task. Include the format (email, social post, bullet-point list, 3-paragraph summary), the length (word count or number of points), and the primary purpose (to inform, persuade, reassure, entertain). The more specific the task, the more focused the output.

3 Context — The background information

Context is where the magic happens. This is where you tell the AI about your audience, your situation, your business, your product, your customer's problem. "For a boutique hotel in Lisbon launching a winter package" — this context transforms a generic travel post into something specific and vivid. The richer your context, the richer the output.

4 Format — How should it look?

Do you want a bulleted list or flowing paragraphs? One section or five? Should it have headings? How long should each section be? Should it include a subject line? A call to action? Specifying format upfront means you get output that's ready to use, not output you need to restructure. "Use 3 short paragraphs, add 5 hashtags" is a format instruction.

5 Tone — The voice and feel

Tone tells the AI how it should *sound*. Options include: warm, professional, conversational, authoritative, playful, aspirational, urgent, reassuring, inspiring, direct. Also valuable: telling the AI what NOT to do — "no clichés," "no corporate jargon," "don't use bullet points," "avoid the word 'innovative'." Tone is often the difference between content that sounds like you and content that sounds like a press release.

A Complete Five-Part Prompt — Worked Example

FIVE-PART PROMPT IN ACTION

[ROLE] Act as a senior travel writer with experience in European boutique hospitality.

[TASK] Write a 150-word Instagram caption

[CONTEXT] for a boutique hotel in Lisbon launching a winter package targeting couples aged 35-55 who value authentic local experiences over luxury chains.

[FORMAT] Use 3 short paragraphs. Add 5 relevant hashtags at the end.

[TONE] Warm and aspirational. Evoke atmosphere. No clichés like "hidden gem" or "nestled."

Notice how every element does real work. The role ensures the writing is skilled. The task gives a precise output target. The context makes it specific to the hotel, the city, and the audience. The format gives a ready-to-publish structure. The tone prevents the AI from defaulting to lazy travel clichés. The result: a caption that reads like a professional copywriter wrote it on location.

SLIDE 24 · MORNING RECAP

Morning Recap — What You Now Know

MORNING SUMMARY

By this point in the day, you know more about effective AI use than 80% of the people who are currently using it. Here's what you have covered:

- 01 AI is pattern-matching at scale — not magic.** It works by drawing on patterns learned from vast training data. It doesn't think or feel, but it produces remarkably useful output when given good instructions.
- 02 Abacus.AI gives you every major model in one place.** For ~\$30 AUD/month you get ChatGPT, Claude, Deep Agent, Image AI, Video AI, and Abacus Claw — everything you need for the rest of the course.

- 03 **Weak prompts = generic output. Strong prompts = gold.** The quality gap between a vague prompt and a specific one is enormous — and closing it takes less than 30 seconds of additional thought.
- 04 **Every great prompt has 5 parts: Role · Task · Context · Format · Tone.** This framework applies to every AI tool you will ever use — not just the ones in this course.

WHAT'S COMING NEXT

The afternoon is where the real creative fun begins. **Making AI draw and film for you.** Image generation, video generation, and your first hands-on creative outputs.

CHAPTER 8 · SLIDES 25–40

AI Image Generation — Frameworks & Practice

Turn words into professional-quality images. No design skills required.

SLIDE 25 · HOW AI TURNS WORDS INTO PICTURES

The 30-Second Mental Model for Image AI

Image generation AI can seem like genuine magic the first time you see it in action. You type a description, press enter, and within seconds a fully-formed, professional-quality image appears. Understanding — even roughly — how this works will help you prompt image AI much more effectively.

- 1 **You describe:** You type words — "a red bicycle leaning against a Lisbon wall at sunset." That description is your instruction to the image model. Everything it creates will be anchored to those words.
- 2 **AI imagines:** The model searches its training data — millions of images and their captions that it has processed — for the visual patterns that match your description. It "knows" what bicycles look like, what Lisbon walls look like, what golden-hour sunset lighting looks like, from having studied millions of examples.

3

Pixels appear: The model generates the image pixel by pixel, guided entirely by your description. It starts with random noise and progressively refines it — like watching a Polaroid develop — until the final image emerges. The whole process takes between 5 and 30 seconds depending on the model and resolution.

THE KEY RULE FOR IMAGE AI

The clearer and more specific your words, the closer the generated image to what's actually in your head. Vague descriptions produce generic images. Detailed, specific descriptions produce specific, useful images. The same principle applies here as with text prompting — but with image AI, your vocabulary of visual descriptors (style names, lighting terms, composition terms) becomes your toolkit.

SLIDE 26 · THE IMAGE PROMPT FRAMEWORK

The Image Prompt Framework: S·S·C·L·M·T

Just as text prompts have five parts, image prompts have six "dials" you can turn up or down to control the output. Douglas calls this the S·S·C·L·M·T framework — Subject, Style, Composition, Lighting, Mood, Technical. Mastering these six elements will take your image outputs from generic to cinematic.

Dial	What It Controls	Example Descriptors
S — Subject	What is in the image — the main focus. Be specific about age, appearance, action, and details.	"A 50-year-old jazz pianist mid-performance, face illuminated from below, wearing a navy suit"
S — Style	The artistic or photographic style of the image. This single word can transform the feel completely.	Photorealistic, oil painting, watercolour, anime, 3D render, pencil sketch, cinematic, editorial
C — Composition	How the image is framed and arranged. Camera angle, distance, and layout.	Close-up portrait, rule of thirds, low angle, bird's eye view, wide establishing shot, symmetrical
L — Lighting	The quality, direction, and colour of light. Lighting is often the single biggest factor in image quality.	Golden hour, soft diffused window light, dramatic side lighting, neon backlit, blue hour, harsh midday sun

M — The emotional atmosphere of the image. Melancholic, joyful, mysterious, energetic,
Mood The feeling you want the viewer to have. nostalgic, intimate, epic, peaceful, tense

T — Camera and post-processing specifications that control the final look. 35mm lens, 85mm portrait lens, shallow depth
Technical of field, f/1.8 bokeh, 4K, film grain, long exposure

SLIDE 27 · BEFORE AND AFTER

The Six Dials in Action — Before and After

The fastest way to understand the value of the S·S·C·L·M·T framework is to see the same subject transformed by applying each dial:

Version	The Prompt	The Result
BEFORE	"A coffee shop."	Generic. Flat. Forgettable. Could be any coffee shop, anywhere, at any time. No character, no atmosphere, no use as a marketing image.
AFTER	"A cosy Lisbon coffee shop interior at 7am. Warm pendant lights overhead. Steam rising from a flat white on a marble counter. Shallow depth of field. 35mm photo. Cinematic. Soft golden tones."	Specific, cinematic, on-brand. Every element is deliberate. The time of day, the light source, the specific drink, the lens choice, the post-processing style — all of these combine to produce an image that could genuinely appear in a high-end travel magazine.

The two prompts describe the same subject. The difference is the six dials being used deliberately. Once you have this mental model, the quality of everything you generate improves immediately.

SLIDES 28–31 · REAL-WORLD CASE STUDY

Real-World Case Study: The Solar Flyer

One of the most powerful demonstrations in the workshop is Douglas's real example of creating a professional marketing flyer for a local solar installation business (Rudd Family Builders & Solar) using nothing but rough voice-dictated notes and image AI. This example perfectly illustrates how accessible professional creative output has become.

Step 1 — Rough Spoken Notes (30 seconds of voice dictation)

Douglas simply spoke his rough idea into his phone:

"I want an A5 flyer made up. I want an image prompt created. I want the flyer made up using my image and the same colour scheme as the uploaded business cards. This flyer is about upgrading your old system — solar and battery to eliminate your bill. We only use Solax HS5.1 batteries and inverters, Solax EV chargers and Longi X10 tough solar panels. Write me the image prompt."

No prompt engineering. No special skills. Just talking. Thirty seconds of rough thought dictated to AI — the same way you might leave yourself a voice memo.

Step 2 — AI Writes the Full Image Prompt

The AI transforms that casual ramble into a detailed professional creative brief, including:

- A specific main headline and secondary headline
- Detailed specifications for a beautiful modern Australian home showing rooftop solar panels, premium battery installation, and an EV charger
- Precise technical specifications for the Solax and Longi products (based on their real designs)
- Premium product placement directions (battery beside garage wall, inverter professionally installed, EV charger beside a modern vehicle)
- Emotional and design direction: warm golden sunset lighting, premium brochure quality, balanced layout, 300 DPI print-ready, A5 portrait

You spoke casually. AI structured it into a professional creative brief — the kind a design agency would charge hundreds of dollars to produce.

Step 3 — The Print-Ready Result

The AI-generated image prompt goes into Image AI, which produces a print-ready A5 marketing flyer featuring a photorealistic premium Australian solar home, complete with properly rendered Solax products, professional lighting, and brand-consistent colours — all in under 60 seconds. No design skills required. Ready for print.

DOUGLAS'S SECRET SAUCE METHOD

30 seconds of rough spoken notes → AI writes the full image prompt → AI generates the flyer image → print-ready professional result. This is the workflow. It applies to any marketing material for any business.

HOT TIP — USE MAGNIFIC FOR UPSCALING

Once you have an image you love, use **Magnific** — available inside Abacus.AI — to upscale it to print quality. Magnific takes a good image and makes it extraordinary: sharper, higher resolution, with recovered detail. This is especially useful for flyers, posters, and any print material where 300 DPI is required.

SLIDE 32 · IMAGE GENERATION MODELS OVERVIEW

Image Generation Models — Your Creative Arsenal

Abacus.AI gives you access to the world's most powerful image generators — all in one interface. Each model has different strengths, and knowing which to reach for will significantly improve your results. The models are organised into five families:

SLIDE 33 · FLUX FAMILY

Family 1 — The FLUX Family (Precision & Control)

The FLUX family, developed by Black Forest Labs, is known for high-fidelity output with exceptional precision. With 8 models covering everything from quick iterations to final high-resolution renders, FLUX is the most versatile family in the arsenal. FLUX models are particularly strong for photorealistic output, brand consistency, and structural control.

Model	Best For
flux_pro_ultra	Highest-fidelity final renders — flagship quality. Use when you need the very best output, regardless of time or cost.
flux_pro	Pro default — high quality plus reliability. The everyday workhorse for professional outputs.
flux2_pro	Photorealistic output from the FLUX 2.0 line. Especially strong for product photography and portraiture.
flux2	Faster and cheaper iterations from the FLUX 2.0 line. Good for experimenting and rapid prototyping.
flux_kontext	Context-aware generation using reference images. Give it a photo of your product and it generates images consistent with that style.

flux_kontext_edit In-UI style and prompt-based editing. Change specific elements of an existing image without regenerating the whole thing.

flux_pro_canny Edge-guided structure control. The composition of a reference image guides the new generation — useful for maintaining architectural or product shapes.

flux_pro_depth Depth-guided structure control. The spatial depth map of a reference image shapes the new output.

SLIDE 34 · OPENAI MODELS

Family 2 — OpenAI Image Models (The All-Rounders)

OpenAI's image models are the most instruction-following in the arsenal — they're exceptionally good at doing exactly what you ask, in a very literal way. They're the go-to for editing, inpainting (changing specific parts of an image while keeping the rest), and versatile creative work. If you're new to image AI, these are a great starting point.

Model	Best For
gpt_image15	General-purpose high-quality generation. Strong instruction-following across a wide range of styles.
gpt_image2	Next-generation model with the strongest instruction-following of any image model. Use when precision and literal interpretation of your prompt matters most.
gpt_image_edit	Editing and inpainting — provide an image and a mask, and the model changes only the masked area. Useful for product shots, background removal/replacement, or adjusting specific elements.
gpt_image2_edit	Next-generation editing and inpainting with improved quality.
dalle	Creative, stylised, and illustrative outputs. The original DALL-E model — fantastic for concept art, illustration, and imaginative imagery.

SLIDE 35 · GOOGLE / PHOTOREAL MODELS

Family 3 — Google & Photoreal Models (Photorealism Masters)

When you need images that look genuinely like photographs — not illustrations, not stylised art, but actual photographic output — this family delivers. The Nano Banana models are particularly useful because they support reference images, meaning you can provide your own photo as a starting point and generate variations consistent with it.

Model	Best For
imagen	Google's flagship photorealistic image model. Exceptional for people, products, architecture, and natural scenes.
nano_banana_pro	High-quality generation with reference image support. Give it a reference photo and generate consistent variations. Ideal for product photography, headshots, and brand consistency.
nano_banana2	Flexible prompting with support for multiple reference images. Strong for real-estate photography, portrait shoots, and stock photo replacement.
nano_banana	Baseline option — fast and accessible. Good for quick experiments within the Nano Banana family.

BEST USE CASE — NANO BANANA

Product photography, portrait shoots, real-estate imagery, and stock photo replacement. If you have an existing photo and need professional variations of it — consistent brand images, different backgrounds, professional headshots from a simple selfie — Nano Banana models are your best option.

SLIDE 36 · DESIGN, TEXT & VECTOR MODELS

Family 4 — Design, Text & Vector Tools (Typography & Logo Creation)

This family is indispensable if you need to create images that contain readable text — logos, posters, social banners, book covers, event flyers. Standard image AI models notoriously struggle with text in images: words come out misspelled, garbled, or distorted. The models in this family are specifically engineered to solve that problem.

Model	Best For
ideogram	Readable text in images. The go-to for logos, posters, banner ads, social graphics, captions, and any output where words must be legible. Strong typographic design capabilities.
ideogram_character	Character-focused generation with reference image support. Good for consistent character design across multiple images.
recraft	Design and illustration styles with clean, graphic aesthetics. Particularly good for infographic-style visuals and brand illustrations.

recraft_svg True vector SVG output — logos, icons, and graphics that can be scaled to any size without quality loss. One of the only models that produces genuine vector output.

midjourney Artistic and stylised high-end aesthetics. Midjourney is famous for its cinematic, painterly quality and its ability to produce consistently stunning imagery.

seedream Strong general-purpose option with broad stylistic range.

SLIDE 37 · CREATIVE & ENHANCEMENT MODELS

Family 5 — Creative & Enhancement Models (Artistry & Upscaling)

This family covers the more artistic, experimental, and post-production end of the image pipeline. If you want images that feel like fine art, or you need to rescue and enhance an existing image, these are your tools.

Model	Best For
dreamina	Art-forward creative generations. Great for concept art, imaginative compositions, and illustration-style outputs.
imagine_art	Artistic rendering and painterly styles. Strong for oil-painting effects, watercolour looks, and traditional art aesthetics.
hunyuan_image	Strong general-purpose image generator from Tencent. Reliable across a wide range of styles.
grok_imagine_image	Creative generations from xAI/Grok. Known for unique stylistic interpretations.
qwen_image_edit	Single-image editing workflows. Modify specific elements of an existing image.
magnific HOT TIP	Upscaling, enhancement, and detail recovery. Takes a blurry or low-resolution image and produces print-quality output. Great for rescuing old photos, enhancing product shots, and making AI-generated images print-ready.

SLIDE 38 · PLATFORM UI NAMES

What You'll See in the Abacus.AI Studio Interface

In the Abacus.AI platform, models appear under their friendly UI names rather than their technical model names. Here is the mapping so you can find the right tool in the interface:

GPT Image 1.5 / GPT Image 2

General-purpose and next-gen general. OpenAI's all-rounder models.

GPT Image 1.5 [Edit] / GPT Image 2 [Edit]

Prompt-based edits and next-gen editing. Use when you want to change part of an existing image.

Nano Banana Pro / Nano Banana 2

High-quality Google-powered models. Nano Banana Pro is for high-quality output; Nano Banana 2 supports multiple reference images.

Seedream 4.5

Strong general-purpose option with broad stylistic range.

Midjourney

Artistic and stylised. The model famous for cinematic, painterly aesthetic quality.

Grok Imagine Image

Creative generations from xAI. Known for unique stylistic interpretations.

FLUX.2 [Pro]

High quality precision. One of the FLUX family's top-tier models.

Hunyuan Image 3.0

General-purpose from Tencent. Reliable across a range of styles.

Imagen 4

Google's photorealism flagship. The best choice for photography-style outputs.

Recraft SVG

Vector graphics output — logos and icons that scale perfectly at any size.

Ideogram 3.0

Text rendering — the go-to for any image that must contain readable words.

+ Magnific / Qwen Image Edit

Upscaling and single-image editing — also available directly in the platform.

SLIDE 39 · ACTIVITY 01

Activity 01 — Create Your First AI Image (15 minutes)

HANDS-ON ACTIVITY

Time to create your first real AI image. Theme: **a promotional poster for your business** (or, if you don't have a business, an image for a product or service you'd love to create).

- 1 Open Abacus.AI and navigate to Image Studio.
- 2 Choose your image generation model. If you're unsure, start with GPT Image 2 or Imagen 4.
- 3 Use Douglas's Secret Sauce method: dictate or write rough notes about what you want, then ask the AI to write you a professional image prompt using the S·S·C·L·M·T framework. Then paste that prompt into Image AI.
- 4 Generate 3 variations. Notice how even small differences in the prompt change the output significantly.
- 5 Share your best image with the room. What worked? What surprised you?

SLIDE 40 · BEST PRACTICES

Image Generation Best Practices — Small Habits, Dramatic Results

DO

- Describe in nouns and adjectives, not full sentences. "Cosy, warm, amber, rustic, steam" beats "a place that feels welcoming."
- Always name a style — "photoreal," "watercolour," "3D render," "oil painting." The style word changes the entire feel.
- Add lighting words — they change everything. "Golden hour," "overcast soft light," "dramatic neon" each produces a completely different atmosphere.
- Generate 3–4 variations and iterate on the best. Treat each generation as a first draft, not a final result.
- Use Magnific to upscale your best image to print quality before publishing or printing.

DON'T

- Don't dump everything into one giant run-on sentence. Break your prompt into clear sections: subject, style, lighting, mood.
- Don't use vague words like "nice," "good," or "beautiful." These tell the AI nothing specific. Be concrete.
- Don't expect perfection on the first try. Image AI rewards iteration. Your second prompt will always be better than your first.
- Don't generate faces of real, identifiable public figures. This raises legal and ethical concerns and violates most platforms' terms of service.

AI Video Generation — Frameworks & Practice

From still frames to moving stories. Your personal production studio.

SLIDE 41 · FROM IMAGES TO VIDEO

From Still Frames to Moving Stories

Video AI takes everything you've just learned about image AI and adds three new dimensions: motion, time, and scene logic. The core principle is the same — you describe what you want in words, and the AI generates it. But now the output moves. And that changes what you need to include in your prompt.

+ Motion

Objects, people, and the camera itself now move. You need to tell the AI what is moving, how it's moving, and at what speed. "A man walks slowly toward the camera" is a motion instruction. "The camera pans right to reveal the skyline" is a camera motion instruction.

+ Time

Every video prompt must define duration — how long the clip runs (4, 6, 10, or 15 seconds depending on the model) — and pace — slow, medium, fast. A "6-second slow-motion clip" and a "6-second fast-cut" are completely different outputs even if the scene description is identical.

+ Scene Logic

Transitions — how the clip ends or how shots connect — are now part of the prompt. A cross-fade out creates a dreamy, atmospheric feel. A hard cut creates energy. Naming the transition in your prompt gives the AI direction on how to end the scene.

Same Foundation, New Dimensions

Everything you know about image AI still applies: subject, style, lighting, mood, composition. Video AI adds motion, time, and transitions on top of that foundation. You don't start over — you build on what you already know.

SLIDE 42 · THE VIDEO PROMPT FRAMEWORK

The Video Prompt Framework — Six Elements Every Clip Needs

Where image prompts have six dials (S·S·C·L·M·T), video prompts have six elements that must all be present for a high-quality clip. Think of each element as a production instruction to the AI — the more complete your brief, the more cinematic the result.

Element	What It Defines	Example
01 — Scene	Where we are and what we see. The environment, time of day, and visual context.	"Rainy Lisbon alleyway at dusk, cobblestones reflecting streetlights"
02 — Movement	What is in motion — subject movement, crowd activity, natural elements (wind, rain, waves).	"A man in a navy coat walks slowly toward camera holding a black umbrella"
03 — Camera	The camera angle, movement type, and shot framing.	"Slow dolly-in, low angle, slightly below eye level"
04 — Timing	Duration of the clip and pace of the action.	"10 seconds, slow pace"
05 — Transitions	How the clip ends or how shots connect — fades, cuts, dissolves.	"Cross-fade out at the end"
06 — Effects	Post-processing effects: lens behaviour, film emulation, colour grading.	"Anamorphic lens flare, slight film grain, teal-orange colour grade"

Complete Worked Example

SIX-ELEMENT VIDEO PROMPT

Scene: Rainy Lisbon alleyway at dusk, after rain, cobblestones wet and reflecting amber streetlights.

Movement: A man (50s) in a navy coat walks slowly toward camera, holding a black umbrella.

Camera: Slow dolly-in. Low angle, slightly below eye level.

Timing: 10 seconds. Slow, contemplative pace.

Transitions: Cross-fade out at the end.

Effects: Anamorphic lens flare. Slight film grain. Teal-orange colour grade.

SLIDE 43 · BEFORE AND AFTER

Video Prompts Before and After — The Cinematic Difference

Version	Prompt	Result
BEFORE	"A man walking."	Static, unclear, a 2-second loop with no purpose. Unrecognisable setting. No cinematic quality. Not usable for marketing or social content.
AFTER	The full six-element prompt above, with scene, movement, camera, timing, transitions, and effects all specified.	Cinematic, mood-rich, ready for social media. The Lisbon setting, the specific character detail, the colour grade, the lens choice — all combine to create a clip that feels professional and intentional.

Notice also that the AFTER prompt adds one more element: "Take this character in the image and get him to walk towards the camera." This is the Image-to-Video workflow — starting with a still image you've already created or provided, and animating it. This is one of the most powerful techniques in video AI, because it gives you precise character and visual control from the start.

SLIDES 44–45 · VIDEO GENERATION MODELS

Video Generation Models — Your Production Studio

Just as Abacus.AI gives you access to every major image generation model, it also provides a complete lineup of the world's most powerful video generators — from cinematic storytelling models to physics-accurate motion simulators to near-photographic realism engines. Here are the key models:

Model	Strength
Sora 2	OpenAI's flagship video model. Best for cinematic storytelling, narrative-driven content, and multi-shot coherent sequences. The gold standard for creative video.
Seedance 2.0	ByteDance's video model known for fluid, physics-accurate motion. Excellent when objects or people need to move in a believable, natural way.
Wan 2.5	Fast and versatile. Good all-rounder when you want quick generation without sacrificing quality.
Seedance 1.5 Pro	High-fidelity professional output from the Seedance family. Slower but higher quality than Wan.
Hailuo 2	MiniMax model known for exceptional realism. Strong for product demonstrations and lifestyle content.
Kling AI v3	Strong character and scene coherence. When you need a consistent person or object across multiple generations, Kling maintains visual identity well.
Kling AI 03	Advanced cinematic generation from the Kling family. Higher quality, more detail.
Kling v2.6 Motion	Controlled motion via reference image or reference video. You steer where things move — camera paths, object motion — with a reference.
Luma Labs	Smooth image-to-video animation. One of the best options for taking a still image and bringing it to life with natural-looking motion.
Veo 3.1	Google's flagship video model. Highest fidelity output with native audio generation. Produces cinematic-quality clips with realistic ambient sound.
Veo 3.1 Lite	Faster version of Veo — great for b-roll footage, background clips, and quick turnaround needs.
Grok Imagine Video	Creative and expressive from xAI. Known for unique, imaginative interpretations of prompts.

Topaz Video Upscaler**ENHANCEMENT**

2x or 4x resolution enhancement with optional FPS boosting. Use after generation to make any video clip sharper, smoother, and cinema-quality.

WHICH MODEL TO START WITH?

For most beginners, start with **Sora 2** for creative, story-driven clips or **Hailuo 2** for realistic-looking outputs. Once you have a clip you like, run it through **Topaz Video Upscaler** for a significant quality boost.

SLIDE 46 · GENERATION MODES

Generation Lanes — Three Modes to Choose From

Inside Abacus.AI's video studio, you can select a generation lane that tells the system what kind of video you're aiming for. This guides model selection and generation parameters:

CINEMATIC

For storytelling, multi-shot coherence, and creative cinematic concepts. Choose this when you want the video to feel like a movie clip — deliberate framing, character-driven, narrative-led.

PHYSICS

For believable motion, product demonstrations, mechanical movement, and natural physics. Choose this when the way things move needs to look real — water, machinery, falling objects, crowd behaviour.

PHOTOREALISTIC

For high realism, crystal-clear visual clarity, and near-photographic output. Choose this when the video should look indistinguishable from real footage — brand content, product video, lifestyle clips.

AUTO MODE

Let the system select the optimal generation style for your prompt automatically. A good starting point if you're unsure which mode fits your creative vision.

SLIDE 47 · INPUT WORKFLOWS

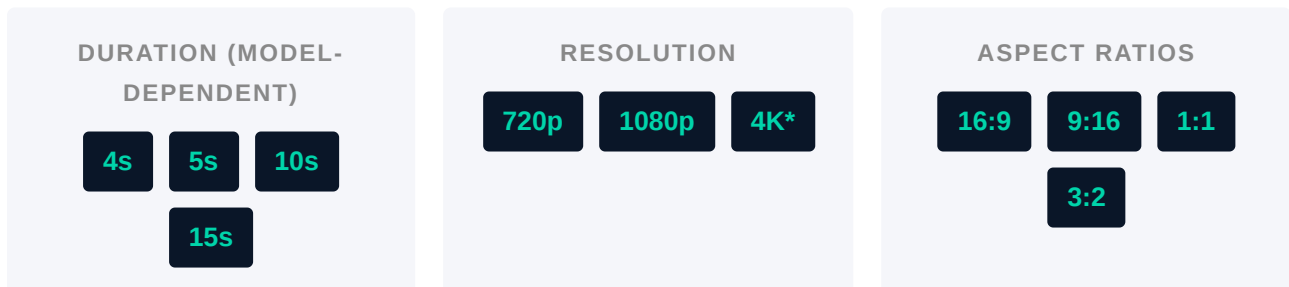
What You Feed the Model — Four Input Workflows

Video AI is flexible about what you start with. There are four distinct input workflows, each giving you a different level of creative control:

- 1 **Text to Video:** Generate entirely from scratch using your written prompt. The AI creates the entire visual world from your description. Maximum creative freedom. Best for imaginative or abstract content.
- 2 **Image to Video:** Start with a still image (PNG, JPG/JPEG, or WebP) and animate it. The AI uses your image as the first frame and brings it to life. This is Douglas's favourite technique for consistent character work — generate the character in Image AI first, then animate in Video AI.
- 3 **Image + Text:** Combine a visual anchor image with text-guided motion instructions. The image establishes the scene and characters; your text prompt directs how they move, the camera behaviour, and the style. Best of both worlds — visual precision plus text control.
- 4 **Motion Control:** Provide a reference video (up to 30 seconds, MP4 format) or a reference image as a motion guide. The AI follows the movement pattern of your reference — the camera moves, subjects walk, objects shift — but applies your prompt's style and content. Powerful for maintaining a specific camera style or action sequence.

SLIDE 48 · TECHNICAL SPECIFICATIONS

Technical Specifications — Duration, Resolution & Aspect Ratios



* 4K available on select models only. Aspect ratios: 16:9 = standard landscape (TV, YouTube, desktop). 9:16 = vertical (Instagram Reels, TikTok, Shorts). 1:1 = square (Instagram feed). 3:2 or 2:3 = cinematic widescreen.

SLIDE 49 · VIDEO EDITING AND ENHANCEMENT

Video Editing & Enhancement Tools

Beyond generation, Abacus.AI gives you three powerful post-production tools that can significantly elevate the quality of your video output:

Tool	What It Does	Best Use Case
------	--------------	---------------

Motion Control (Kling v2.6)	Steer camera motion, object paths, and scene direction using a reference video (up to 30s) or reference image. The model follows the motion pattern of your reference.	When you need precise control over how the camera or subjects move — replicate a specific dolly movement, follow a particular path.
Lip-Sync (Talking-Head AI)	Animate a portrait video or image using a speech audio track. The face moves naturally in sync with the audio — blinking, mouth movement, subtle head movement.	Spokespeople, AI avatars, e-learning modules, customer service videos, product explainers. Generate a consistent AI character who speaks your script.
Video Upscaling (Topaz)	2x or 4x resolution enhancement with optional target FPS boosting. Takes existing footage — AI-generated or real — and makes it sharper, smoother, and more detailed.	Upgrading any video to cinema quality. Running standard 1080p output through 2x upscale to produce near-4K quality. FPS boosting smooths out motion for a premium feel.

SLIDE 50 · PLATFORM TOOLS

Platform Tools — Auto Mode, Video Mode, and Rewrite Prompt

AUTO MODE

Routes your input to image or video automatically. Picks the optimal model and settings for your prompt. Use when you want the system to decide — great for experimentation or when you're not sure whether your prompt is better suited for image or video output.

VIDEO MODE

Forces video output regardless of prompt content. Select your preferred model manually. Use when you know you want video and want full control over model selection — you're in director mode.

REWRITE PROMPT TOGGLE

When this toggle is enabled, the system expands and optimises your prompt before generation begins. The AI enriches your rough ideas into a more detailed, comprehensive prompt — similar to the Secret Sauce method but automated. **Use when your prompt is rough or minimal** and you want the system to fill in the creative gaps for better, richer results. Disable it when your prompt is already detailed and you want precise control over the output.

SLIDE 51 · ACTIVITY 02

Activity 02 — Create Your First AI Video (20 minutes)

HANDS-ON ACTIVITY

Now it's your turn to become a video director. This activity builds directly on what you created in Activity 01 — you can animate your AI-generated image if you want to, or start fresh with a new video concept entirely.

- 1 Open Abacus.AI and navigate to Video Studio.
- 2 Choose a model to start with — Sora 2 or Veo 3.1 are recommended for first attempts.
- 3 Write a 6-element prompt covering: Scene · Movement · Camera · Timing · Transitions · Effects. Use the worked example from this chapter as a template.
- 4 Generate one 6-second clip. Watch it closely. What worked exactly as you imagined? What was different?
- 5 Adjust one or two elements of your prompt based on what you saw. Generate a second clip and compare.
- 6 Share your favourite clip with the room. Describe what you changed between your first and second prompt.

THE MOST IMPORTANT LESSON FROM THIS ACTIVITY

Video AI, like image AI, rewards iteration. Your first generation is rarely your best. The real skill is learning to read the output, identify what the model misunderstood or missed, and refine your prompt accordingly. Every generation teaches you something about how to communicate more precisely with the model.

Your Journey Continues

4 workshops. 1 powerful journey. You've already started.

The Full AI Without Fear Workshop Series

Workshop 1 is the foundation of a four-workshop journey that takes you from AI confidence to AI capability — all the way through to building your own apps, your own website, and your own automated systems. Here is where you are now, and where you're heading:

COMPLETE

01

AI Foundations + Image & Video Prompt Mastery

You're here. AI confidence, creative output, and mastery of AI image and video generation. The foundation everything else builds on.

[Foundations](#)
[Image AI](#)
[Video AI](#)

UP NEXT

02

Building a Custom App

No-code app building using Deep Agent. Build real business tools fast — quote calculators, CRM systems, client portals, and more. From brief to working app in a single session.

[No-Code](#)
[Apps](#)
[CRM](#)

COMING SOON

03

Building Your Custom Website & Hub

Same-day website deployment using Deep Agent and DNS configuration. Go from zero to a live, working website for your business — in a single workshop day.

[Website](#)
[Hub](#)
[Deploy](#)

COMING SOON

04

Chatbot & SMS Automation via Abacus Claw

AI chatbots, SMS automation, and browser pilot via Abacus Claw. Build automated customer communication systems that work while you sleep.

[Chatbot](#)
[SMS](#)
[Claw](#)

WHAT YOU'VE ACHIEVED TODAY

You walked in knowing AI is a thing. You're walking out knowing **how** it works, **what** to use it for, **how to prompt it** effectively, how to **generate professional images**, and how to **create video content** — without any prior creative or technical training. That is a remarkable shift. Workshop 2 builds directly on these foundations.

SLIDE 53 · QUESTIONS AND CLOSE

Questions, Close & Stay in Touch

The workshop closes with 20 minutes of open floor Q&A — your opportunity to ask anything that came up during the day, clarify anything you want to understand better, or share what you're most excited to try first. There are no bad questions. This is your time.

Email Douglas

douglasmizzi@me.com — for follow-up questions, feedback, or specific prompting challenges you encounter after the workshop.

Website

aiwithoutfear.com.au — course information, upcoming workshop dates, resources, and links to the platform tools used in this workshop.

Community

Facebook.com/fusionai — join the community of AI Without Fear graduates. Share your creations, ask questions, and stay up to date as AI tools evolve.

Next Workshop

Workshop 2 — Building a Custom App — is your next step. Check aiwithoutfear.com.au for upcoming dates and early registration.

FINAL THOUGHT FROM DOUGLAS

"You now know more about effective AI use than 80% of people using it today. The gap between knowing and doing is the only one left to close. Start small. Pick one tool. Use it for something real in your business this week. That first real output — the first email, image, or video you create — is the moment it stops being a workshop concept and starts being part of how you work."

WORKSHOP 1 OF 3 · COMPLETE

AI Without Fear

aiwithoutfear.com.au