
Most marketers using AI right now are improvising. They type something into a chat window, get something back, edit it heavily, and move on. It works — sort of. But it's inconsistent, slow, and heavily dependent on whoever happens to write the prompt that day.

Prompt frameworks fix this. They're structured approaches to writing prompts that consistently produce better, more usable outputs — because they force you to give the AI the right information in the right order before it starts generating.

This article covers ten of the most useful ones, what they're actually good for, and how to apply each one in a real marketing context.

A quick note: most of these frameworks are known by acronyms. Each letter stands for an element you include in your prompt. You don't need to label them rigidly — once you understand the logic, the structure becomes second nature.

1. **RTF — Role, Task, Format**

RTF is the starting point. It's the simplest structured framework and the one most marketers should default to for everyday, single-task prompts.

Role: Tell the AI what kind of expert it is. "You are a direct-response copywriter specializing in email marketing."

Task: Tell it exactly what to do. "Write a promotional email for a 20% discount on our bookkeeping software, targeting small business owners who haven't purchased yet."

Format: Tell it how to deliver the output. "Write it as a single email under 200 words, with a subject line, three short paragraphs, and a clear call to action."

RTF keeps prompts tight without making them complicated. It's the best framework

for volume work — when you're writing a batch of social posts, email subjects, product descriptions, or ad variations and you need consistent structure without a lot of overhead.

What it's best for: Quick tasks, batch content, one-off outputs.

2. **RISEN — Role, Instructions, Steps, End Goal, Narrowing**

RISEN is RTF with more depth. Where RTF handles simple tasks well, RISEN is built for anything complex — multi-phase campaigns, content series, strategic briefs, or any task where getting the process right matters as much as the output.

Role: Define the AI's expertise. Instructions: Explain the overall task. Steps: Break the task into a numbered sequence. End Goal: Define what success looks like.

Narrowing: Add the constraints — word count, tone, what to avoid, what to prioritize.

The Narrowing component is what makes RISEN particularly useful. It's where you prevent the AI from going broad when you need it to go specific, or from defaulting to safe when you need something with an edge.

Example for a marketing campaign brief:

“Role: You are a senior brand strategist. Instructions: Develop a campaign concept for a local gym targeting professionals aged 30-45 who have tried and quit gym memberships before. Steps: 1) Identify the core emotional barrier for this audience, 2) Define the campaign's central message, 3) Suggest three creative executions across social, email, and in-gym signage. End goal: A brief the creative team can execute from. Narrowing: No generic fitness messaging. Avoid motivational clichés. Keep the tone honest and direct, not cheerful.”

What it's best for: Campaign strategy, content briefs, multi-step deliverables, complex analysis.

3. **CRAFT — Context, Role, Action, Format, Tone**

CRAFT is arguably the most well-rounded general-purpose framework for marketing content. It covers everything RTF does and adds two elements that matter enormously for brand consistency: Context and Tone.

Context: What's the situation? What does the AI need to know before it starts? Role: What expert is it being? Action: What should it do? Format: How should it be structured? Tone: How should it sound?

See also [The One-Person AI Business Model Is Real](#)

The Tone element is what separates CRAFT from more stripped-down frameworks. For marketers who care about brand voice — and you should — specifying tone is one of the highest-leverage things you can do in any prompt. “Confident without being arrogant. Direct without being cold. Friendly without being casual” produces fundamentally different output than leaving tone undefined.

Example for a product launch announcement:

“Context: We’re launching a new project management tool for small construction businesses. Most competitors talk features. We focus on how it reduces missed deadlines and difficult client conversations. Role: You are a B2B copywriter who understands trade industries. Action: Write a launch email to our existing waitlist. Format: Subject line, three short paragraphs, one CTA. Tone: Straightforward and practical. Speak to people who work with their hands and don’t trust software that promises too much.”

What it’s best for: Brand content, email marketing, website copy, any output where voice matters.

4. **COSTAR — Context, Objective, Style, Tone, Audience, Response Format**

COSTAR takes the communication brief approach further. It was originally developed for use with large language models in enterprise contexts and has become one of the more widely adopted frameworks for content-heavy marketing tasks.

Context: Background on the situation. Objective: What you want to achieve, not just

what you want produced. Style: The writing style — conversational, journalistic, formal, punchy. Tone: The emotional register — reassuring, urgent, playful, authoritative. Audience: Who specifically is reading this. Response Format: The structure and length of the output.

The distinction between Style and Tone is worth understanding. Style is how the writing is constructed — sentence length, vocabulary, rhythm. Tone is the emotional quality underneath it. A piece can be written in a conversational style but carry an urgent tone, or a formal style with a reassuring tone. COSTAR lets you specify both, which matters for things like crisis communications, sensitive customer messages, or content aimed at audiences with specific professional expectations.

What it's best for: Content strategy, brand campaigns, multi-channel content that needs strict voice alignment.

5. **RACE — Role, Action, Context, Execute**

RACE is a leaner framework that works well when you need something fast but more precise than RTF. The Execute component is what distinguishes it — rather than just describing format, you're giving the AI a direct instruction to produce the output with specific parameters already built in.

Role: What expert is responding. Action: What they're doing. Context: The situation and relevant background. Execute: The instruction to produce — including format, length, and any final constraints.

RACE is particularly effective for research and analysis tasks in marketing: competitive analysis, customer persona development, positioning reviews, market summaries. The framework pushes you to front-load the context before the execution instruction, which produces more grounded, specific outputs than leading with the task.

What it's best for: Research tasks, competitive analysis, audience insights, strategic summaries.

6. **BAB — Before, After, Bridge**

BAB comes directly from direct-response copywriting and is one of the most natural frameworks for marketers because it mirrors how good persuasive writing already works.

Before: Describe the situation as it currently is — the problem, the frustration, the pain point your audience knows firsthand. **After:** Paint the picture of how things look once the problem is solved — the outcome, the relief, the improved state. **Bridge:** Connect the two by introducing the product, service, or strategy that creates the transformation.

This framework taps into something fundamental about how people respond to marketing: they don't buy products, they buy movement away from a problem and toward a better situation. BAB makes that structure explicit for AI so it produces copy that follows the same emotional logic.

See also [How I Build AI Workflows Instead of One-Off Prompts](#)

Example:

“**Before:** Small business owners are spending hours each week manually chasing unpaid invoices and still missing some. **After:** Imagine having every invoice tracked automatically, late payments flagged before they become a problem, and getting paid on average two weeks faster. **Bridge:** Write copy introducing our invoicing software as the solution. Under 100 words, warm but confident tone.”

What it's best for: Sales copy, landing pages, case studies, product pages, campaign messaging built around transformation.

7. **AIDA — Attention, Interest, Desire, Action**

AIDA is not new — it's been a cornerstone of advertising since the late 1800s. But it translates directly into a prompt framework and remains one of the most reliable structures for conversion-focused marketing content.

Attention: Hook the reader immediately. Stop the scroll. **Interest:** Build curiosity or relevance. Give them a reason to keep reading. **Desire:** Create want — through benefits, outcomes, social proof, or emotional appeal. **Action:** Tell them exactly what to do next.

As a prompt framework, AIDA works by instructing the AI to build each element intentionally rather than producing generic copy that checks no particular box well.

Example:

“Write a Facebook ad for a local pet grooming business using the AIDA structure. Attention: Something that stops a dog owner mid-scroll. Interest: What makes this grooming service different from the chain option down the road. Desire: How their dog will look and feel, and how that reflects on the owner. Action: Book online, first appointment 20% off. Keep it under 80 words. Friendly tone.”

What it’s best for: Ad copy, landing pages, email sequences, promotional content with a clear conversion goal.

8. Chain-of-Thought

Chain-of-Thought is different from the frameworks above because it’s less about structure and more about process. Rather than telling the AI what to produce, you instruct it to think through the problem step by step before arriving at its answer.

The practical instruction is simple: add “Think through this step by step before answering” or “Walk me through your reasoning before giving your recommendation” to any prompt.

Why does this work? AI models produce better reasoning when they’re prompted to show their work. The act of generating intermediate steps improves the final output — particularly for any task involving analysis, strategic judgment, or decision-making.

For marketers, Chain-of-Thought is most valuable when you’re using AI for something that requires judgment rather than just generation: evaluating a campaign concept, analyzing why a launch underperformed, recommending a positioning strategy, or stress-testing a marketing argument.

Example: “We ran a promotional email campaign last month with a 12% open rate and a 0.8% click rate. Our list is 4,000 subscribers, mostly previous customers. Think through step by step what might explain these results, what benchmarks suggest about performance, and what the most likely improvements would be.”

What it's best for: Strategic analysis, problem-solving, diagnosis, decision support, any task where reasoning quality matters.

9. **TAG — Task, Action, Goal**

TAG strips everything back to the essentials and is worth knowing specifically because of its speed. For experienced prompt writers, there are moments when a longer framework creates more friction than value. TAG handles those moments.

Task: What kind of task is this. Action: What the AI should do specifically. Goal: What outcome you're working toward.

The power of TAG is that it forces clarity on the goal — not just the output, but the underlying objective. This small addition changes how AI approaches the task. “Write three Instagram captions” and “Write three Instagram captions with the goal of driving profile visits from people who don't follow us yet” produce meaningfully different results, even though the action is the same.

See also [The Psychology Behind Prompts That Actually Work](#)

What it's best for: Quick tasks where you need speed but still want intentional output, not random generation.

10. **APE — Action, Purpose, Expectation**

APE is a clean, beginner-friendly framework that's particularly effective for marketers who are newer to prompting or who want something easy to teach across a team.

Action: What the AI should do. Purpose: Why — the reason behind the task, the context that explains the goal. Expectation: What good output looks like — format, length, quality bar.

The Purpose element is what makes APE more useful than a bare instruction. “Write a LinkedIn post about our new service” tells the AI what to do. “Write a LinkedIn post about our new service — the purpose is to generate enquiries from operations

managers at companies with 10 to 50 employees, not just to announce the launch” tells it why. The why reshapes the what.

Example: “Action: Write a follow-up email sequence (three emails) for prospects who attended our webinar but didn’t book a call. Purpose: Move them from interested but uncertain to confident enough to schedule a conversation — they know our product is relevant, they just haven’t prioritized it yet. Expectation: Each email under 150 words, no hard sell, progressively more direct. Tone should feel like a helpful colleague, not a salesperson.”

What it’s best for: Teams learning to prompt consistently, everyday marketing tasks, building internal prompt standards.

How to Choose the Right Framework

You don’t need to memorize all ten and decide which one to use every time you open a chat window. A simpler mental model:

For quick, single-task content, use RTF or TAG. Fast, structured, effective.

For anything involving brand voice or specific audience, add CRAFT or COSTAR. The extra elements pay off when consistency matters.

For strategic or multi-step work, use RISEN or RACE. They’re built for complexity.

For conversion-focused copy, use BAB or AIDA. They’re grounded in how persuasion actually works.

For analysis and reasoning tasks, use Chain-of-Thought. Output quality improves measurably when AI shows its work.

For teams building shared prompt standards, APE is the easiest to teach and apply consistently.

The best marketers don’t pick one framework and apply it to everything. They know which tool fits which job — and they’ve built enough familiarity with each one that switching between them is instinctive.

Start With One, Then Build

The most practical advice: pick the one framework that fits the task you do most often with AI right now. Learn it properly. Build a few working prompts around it for your specific business and audience. Save those prompts.

Then add another framework when you hit a task the first one doesn't handle well.

That's how a prompt library gets built — not by collecting frameworks in theory, but by applying them to real tasks until you have a bank of working prompts that produce reliable results.

If you want to skip the trial-and-error stage and start with prompts that already apply these frameworks to real marketing use cases, that's exactly what our prompt library is built for.

Browse the full [library](#)

Frameworks don't replace good thinking. They give good thinking a structure it can work inside. And in marketing, structure is what turns one good prompt into a repeatable system — which is what separates businesses that occasionally get something useful from AI and businesses that consistently do.

Share this:

- [Share on Facebook \(Opens in new window\) Facebook](#)
- [Share on X \(Opens in new window\) X](#)