

PRACTICAL GUIDE FOR AI AGENTS

Skills Instead of Tool Chaos

How teams translate recurring AI tasks
into reliable work standards.



Roger Basler de Roca

MSc Digital Business | PhD Candidate

Digital Entrepreneur, Author, and Top 100 Speaker

August 2026 Edition

skills.fragroger.ai

 ORIENTATION

From Principle to Application.

This eBook follows a simple question: How does a good AI response become a reliable work standard?

The answer begins with a precise task, moves through criteria and context, and only ends after a verified application. This may seem less spectacular than a new tool demo, but it is the difference between a lucky hit and an operating system.

01	The Starting Point	03
02	What a Skill Is	04
03	The Anatomy	05
04	The Operating Model	06
05	From Prompt to Standard	07
06	The Skills Library and Four Fields of Action	08-09
07	Implementation, Validation, Improvement	10-14
08	Principles, Sources, and About Roger	15-16

READING LOGIC

Understand the mechanics first. Then choose a task with high repetition. Test the skill on real cases. Only then is it worth making it a team standard.

THE STARTING POINT

AI rarely fails because of the model.

It fails due to vague tasks, missing criteria, and workflows that start from scratch with every request.

Most AI pilot projects don't fail because a model can't write text or read a table. They fail due to vague tasks. A team formulates a request from the gut, the result varies, no one measures quality, and the next time around, everything starts at zero again.

A skill is the counter-movement to this improvisation. It defines which task is to be solved, what information belongs to it, which steps are useful, and how to recognize a good result. This doesn't make every answer automatically correct. But the work becomes more traceable, verifiable, and repeatable.

THE CORE

The actual problem is not too little AI. It is too few good work standards.

A skill is not a substitute for expertise. It is a mechanism to provide expertise, experience, and process discipline in a form that an agent can use at the right moment.

Tool-Hopping

- Every task is reformulated.
- Quality depends on individuals.
- Errors remain experiential knowledge.
- Results are hard to compare.

Work Standard

- Task and goal are clearly defined.
- Context and criteria are visible.
- Test cases promote targeted improvement.
- Results become repeatable.

DEFINITION

A skill is a work instruction with memory.

It bundles metadata, procedures, and, if necessary, resources in a file and folder structure that an agent can use according to the task.

"A skill is a single Markdown file with instructions."

SKILLS LIBRARY | FRAGROGER.AI [1]

01

Metadata

Name and description say what the skill does and when it should trigger.

02

Instructions

The `SKILL.md` describes the workflow, inputs, results, and quality criteria.

03

Resources

Scripts, references, and templates provide detailed knowledge or deterministic steps.

The open Agent Skills specification describes a skill as a folder with at least one `SKILL.md`. The file connects YAML frontmatter with Markdown instructions. Supplementary directories for scripts, references, and assets are possible but not mandatory. [2]

IMPORTANT

A skill is not a magic add-on for a language model. It structures a good way of working. Whether and how an agent loads local skills depends on the respective environment. [4] [5]

STRUCTURE

Four answers before the agent starts.

What is the goal? When should the process take effect? Which steps are mandatory? How do you recognize a usable result?



1 · Goal

Describe the concrete task and the expected outcome.



2 · Trigger

Name words and situations that make the skill relevant.



3 · Procedure

Determine what the agent should check, do, and output.



4 · Validation

Define quality, limits, and approval points.

The `description` is more than just a label. It supports the discoverability of a skill. Therefore, it should clearly name the task, target audience, and typical triggers. The open specification requires at least `name` and `description`; other fields are optional. [2]

```
---
name: decision_letter
description: Creates a precise decision letter from meeting notes with the decision, justification, risks, and open points. Use it for meeting notes, steering committees, and project decisions.
---

# Working Method
1. Separate observed facts from assumptions.
2. Formulate the decision in one sentence.
3. List options, impacts, and open testing needs.
4. Only release when responsibility and the next date are visible.
```

DESIGN RULE

A skill should encapsulate a coherent unit of work. "Creating a decision letter from meeting notes" is clear. "Everything about transformation" is a fog machine with a file attachment.

PROGRESSIVE DISCLOSURE

Context only when it counts.

A well-structured skill provides the agent with the essentials first and deepens details only when the task actually requires them.



This approach is called progressive disclosure. The open specification distinguishes between metadata, the core instruction, and additional resources. ChatGPT/Codex and Claude also document a comparable pattern for the context-sensitive use of skills. [2] [4] [5]

This not only saves context. It forces better design. Whatever is always needed belongs in the core instruction. What only applies to special cases remains available as a targeted reference.

MANTRA

The quality of a skill is not shown by its length. It is shown by how quickly and reliably it translates a task into a verifiable result.

CLASSIFICATION

A prompt solves a case. A skill preserves a good methodology.

When a team adds the same source criteria to every search or demands the same dramaturgy for every presentation, the moment for a skill has come.

A prompt is a conversational contribution. It can be brilliant and still disappear after one session. A skill begins where a prompt is repeatedly corrected. The crucial difference lies in the specification: A prompt asks for a result. A skill additionally describes what the path should look like, when it makes sense, and what checks are required before submission.

Dimension	Single Prompt	Skill
Purpose	One-time response	Recurring workflow
Context	Reformulated in conversation	Structured as working knowledge
Trigger	Manual	Explicit or by matching description, depending on the platform [4]
Validation	Often implicit	Integratable as a quality criterion
Development	Hard to see	Versionable and testable

Repetition without a standard is just routine waste.

THE CONSEQUENCE OF WORKFLOWS WITHOUT VERIFIABLE QUALITY

Skills do not guarantee correctness. However, they make quality explicit. And what is explicit can be tested, improved, and accounted for.

PRACTICE LIBRARY

Think from the task perspective.

The Skills Library at FragRoger.ai is not oriented towards model names, but towards recurring tasks. That is the right perspective.

No one buys a screwdriver to own a screwdriver. The intention is to turn a screw safely and cleanly. Accordingly, the library groups its skills into current fields of work such as Content, SEO, Governance, Prompting, Documents, Quality, Speaking, Visualization, Finance, Communication, Research, and Productivity. [1]

Content

SEO

Governance

Prompting

Documents

Quality

Speaking

Visualization

Finance

Communication

Research

Productivity

Text Humanizer

Revises artificial-sounding texts according to clear structural and formatting rules.

CONTENT

GEO Optimizer

Structures articles and technical texts for AI extractability, search intent, EEAT signals, and technical requirements.

SEO

AI Guidelines Check

Categorizes use cases, data categories, and relevant risk questions as a basis for a professional review.

GOVERNANCE

Workflow Visualizer

Translates process descriptions into triggers, steps, decisions, tools, and outputs.

VISUALIZATION

ENTRY POINT

Don't choose the most spectacular skill. Choose that recurring task that occurs frequently, costs noticeable time, and can be checked against clear criteria.

PRIORITIZATION

Four fields for the start.

The Skills Library offers many starting points. For the beginning, however, four fields are particularly effective because they combine frequent work with clear controls.

01 · QUALITY AND GOVERNANCE

Not prettier, but more resilient.

CrapCheck verifies claims and sources. AI Guidelines Check structures data, legal, and risk questions. Both do not replace a professional release but make checkpoints visible. [1]

02 · COMMUNICATION AND CONTENT

Not more text, but more intent.

Text Humanizer, LinkedIn Post Writer, Hook Writer, and Better Emails standardize tonality, argumentation logic, and call to action. [1]

03 · KNOWLEDGE AND RESEARCH

Not more hits, but better classification.

Quick Research, Morning Briefing, and GEO Optimizer combine information gathering, evaluation, and visibility. Good research also shows uncertainty. [1]

04 · OPERATIONAL PRODUCTIVITY

Not maximum automation, but controlled.

Receipt Scanning, Invoice Generator, and Budget Dashboard fit clearly structured routines. What is contentious remains a human decision. [1]

What is standardized can be controlled and automated. What is contentious remains a human decision.

PRIORITIZATION RULE FOR IMPLEMENTATION

DEVELOPMENT MODEL

The six-sided reality check.

Before a skill is written or installed, its use should fulfill a clear mandate. The Skill Canvas translates this mandate into six questions.

1. Trigger When should the skill take effect? "When meeting notes require a decision."	2. Goal What result is expected? "One-page decision letter."	3. Context What must be known? "Project goal, options, budget limit."
4. Procedure Which steps are mandatory? "Separate facts, compare options, mark open points."	5. Quality Check What is being checked? "Decision, responsible parties, risks, and next date present."	6. Limit What must the skill not decide? "No legal release, no budget commitment."

The most important part is the last field. Good skills make it clear where an agent must stop and involve a professional. A skill without a limit is not a sign of autonomy, but of poor risk management.


CONTROL QUESTION

Would a new team member be able to work reliably with this skill without anyone having to budget for their mind-reading?

FIRST TEST RUN

Prove first. Then scale.

The technical entry point is small. Operational discipline determines whether it becomes a reliable routine.



The Skills Library describes a simple start: download `SKILL.md`, place it in a suitable folder like `skills//`, connect it to the workspace of a compatible agent, and use a suitable trigger. [1]

The specific storage and activation differ depending on the agent. ChatGPT and Codex, for example, document personal, project-related, and system-wide locations for skills. [4] Therefore, always check the documentation of your chosen environment.

PRACTICE RULE

The first run is not a showcase. It is a test case.

Compare the result with your previous way of working. Measure not only speed but also rework, types of errors, and acceptance by those who continue to work with the result.

GOVERNANCE

Trust is not a permission level.

A skill is a package of instructions and, under certain circumstances, scripts or other files. Therefore, the same basic rule applies as with any software component: check before productive data is involved.

1

Source

Only adopt skills from trustworthy sources. Check all files, not just the description.

2

Data

Clearly define which data may be processed and which does not belong in external systems.

3

Result

Make technical correctness, source logic, and open uncertainties visible.

4

Approval

For high-impact cases, build in a human review step and clear responsibility.

SECURITY LOGIC

Anthropic points out that skills should be treated like software from trustworthy sources. Possible risks include unexpected network requests, tool misuse, and data leakage. [5]

A skill does not automatically expand permissions. However, it can access files, tools, or systems within the existing environment. That's why the source of the skill is part of the security architecture.

For legal, medical, finance, or personnel decisions, the responsible professional remains in the lead. Agents support the process. They do not automatically assume its liability.

IMPLEMENTATION

One month. One task.

A resilient start.

The most sensible way into the skill economy is not the big rollout. It is a clearly defined test with a frequent, measurable task.

WEEK 1

Choose Friction

List ten recurring tasks. Mark frequency, effort, and error costs. Decide on a task with high repetition and a clear quality definition.

WEEK 2

Write First Version

Describe triggers, inputs, steps, output, and limits. Keep the core instruction lean and understandable.

WEEK 3

Test Real Cases

Test on five to ten cases. Check time, rework, error types, source quality, and acceptance by subsequent recipients.

WEEK 4

Stabilize

Remove unnecessary steps, add pitfalls, and define responsibility for version, test cases, and approvals.

MEASUREMENT QUESTION OF THE WEEK

Days 1-7: How frequent, costly, and error-prone is the task?

Days 8-14: Is the result, process, and limit understandable for third parties?

Days 15-21: Does rework decrease with the same or better quality?

Days 22-30: Who maintains the version, test cases, and approvals?

The fastest draft is worthless if it only drives quickly into the next correction loop.

WHY SPEED WITHOUT QUALITY MEASUREMENT IS A POOR KPI

The best-practice documentation of the open specification recommends developing skills from real work, corrections, and concrete test cases. [3]

WORKSHOP

Don't start with a tool, start with a friction.

Take the decision letter after meetings. The problem is not that notes are missing. It is that decisions, responsibilities, and open risks blur within them.

INITIAL PROBLEM

Complete notes, unclear consequences.

A meeting ends with many statements. What exactly was decided, who does what by when, and which assumption still needs to be checked, however, remains open.

SKILL MANDATE

One-page decision letter.

The skill separates facts from assumptions, formulates a clear decision, makes risks visible, and names responsible parties as well as next dates.

The first draft may remain short. It doesn't need a 30-page manual as long as the goal, trigger, procedure, and quality check are clear. Extensive details can be specifically outsourced to reference files. [2] [3]

PROMPT CARD · START MANDATE

"Help me turn the recurring task '[task]' into a lean AI skill. First, ask about target audience, trigger, required context, mandatory work steps, desired output, quality criteria, and clear limits. Then create a `SKILL.md` with precise `name` and `description` specifications. Mark assumptions and output five realistic test cases."

The prompt serves as a thinking framework. It does not replace the validation of the finished skill in real work situations.

NEXT STEP

Find suitable starting points and field-tested examples in the **Skills Library**. Choose a task. Test a skill. Measure quality. Then scale.

THREE PRINCIPLES

Skills are the operating instructions for recurring AI work.

01

Impact Over Scope

A skill is valuable if it makes a recurring job easier, more reliable, and more verifiable. Not if it contains as many buzzwords as possible.

02

Precision Over Length

The best instruction provides exactly the context and quality logic that a good first run needs. Details follow specifically.

03

Testing Over Scaling

A skill without test cases, limits, and responsible maintenance remains a pretty document. Only application turns it into a standard.

The economic leverage of AI does not lie in the number of tools installed. It lies in the ability to make good decisions repeatable.

CONCLUSION OF THIS EBOOK

START NOW

Discover suitable starting points in the [Skills Library](#).

Choose a task. Test a skill. Measure the quality. Then scale. More AI is not a plan. A better way of working is.

SOURCES & LICENSE

Thinking further with substance.

This eBook is based on the Skills Library of FragRoger.ai as well as the open Agent Skills specification and the documented implementation at ChatGPT/Codex and Claude. Platform information refers to the documented status as of August 30, 2026.

[1] Skills Library | FragRoger.ai.

<https://skills.fragroger.ai/>

[2] Agent Skills. Specification.

<https://agentskills.io/specification>

[3] Agent Skills. Best practices for skill creators.

<https://agentskills.io/skill-creation/best-practices>

[4] OpenAI. Build skills.

<https://learn.chatgpt.com/docs/build-skills>

[5] Anthropic. Agent Skills.

<https://platform.claude.com/docs/en/agents-and-tools/agent-skills/overview>

LICENSE NOTICE

The Skills Library specifies CC BY-NC-ND 4.0 and describes use for personal use. Check the current license information on the website as well as the details of the respective file before each use. [1]

ABOUT ROGER

25+ years in

digital worlds

Roger Basler de Roca is a digital entrepreneur, author, and Top 100 speaker. He deals with artificial intelligence, algorithms, and human sovereignty in the digital age.

His work combines technology, education, and entrepreneurial thinking. At the center is the question of how people and organizations not only use digitalization and AI but also shape it in a self-determined way.

LEARN MORE

rogerbasler.ch

fragroger.ai

skills.fragroger.ai

Created as a thinking offer. Not as a substitute for professional examination, data security concept, or entrepreneurial responsibility.