

DAILY ENGLISH READING

7 September 2026 - 30-minute active reading pack

TODAY'S QUESTION

Should CDISC knowledge live in prompts, tools, or reusable agent skills?

Difficulty: C1 | agent skills | CDISC/SAS | context engineering | ClinAgent

WHY SELECTED

Scientific Agent Skills packages recurring scientific procedure into versioned, human-readable modules instead of asking a base model to rediscover it every time. The library contains 163 skills across 16 areas. ClinAgent applies a similar pattern to clinical programming: Thin MCP, Thick Skills, with deterministic rules where needed.

READING ORDER

0-3 min Preview: why can working code still be scientifically wrong?
3-14 min Main: skill structure, progressive disclosure, context metrics, limits.
14-20 min ClinAgent: thick skills, thin tools, deterministic rule engines.
20-25 min CDISC TechniCon: structured standards, AI pipelines and traceability.
25-30 min Output: design one reusable SP agent skill.

SOURCES

Main: Scientific Agent Skills - arXiv v2, 2 Sep 2026
<https://arxiv.org/abs/2609.00065>
Support: ClinAgent - Biology Methods and Protocols
Support: 2026 CDISC TechniCon - 9 Sep 2026

BRIEF BACKGROUND

163 skill descriptions use 14,246 tokens, about 7.1% of a 200k reference window.
Full skill instructions total about 482k tokens; references total about 2.48M tokens.
Progressive disclosure leaves 83.7% of documentation unread unless a skill activates it.
Median example workflow uses 23.9% of the window; 29/46 overflow if all refs load.
Important limitation: the full library has no task-level validation or host-selection rate.
SP architecture: standards + narrow tools + versioned skills + validators + human review.

KEY VOCABULARY

procedural knowledge	程序性知识	how to perform a task correctly
agent skill	智能体技能	versioned task-specific procedure
progressive disclosure	渐进式披露	load details only when needed
resident context	常驻上下文	information always in model context
instruction file	指令文件	human-readable procedure file
domain logic	领域逻辑	CDISC/SAS-specific rules and decisions
stateless tool	无状态工具	narrow input-output operation
rule engine	规则引擎	executable deterministic rules
skill routing	技能路由	select the right procedure
context budget	上下文预算	limited space for model context
lexical separability	词汇可区分性	how distinct descriptions are
scope boundary	适用范围边界	what a skill covers or declines
versioned capability	版本化能力	controlled capability tied to a version
portable architecture	可移植架构	works across different agent hosts
scientifically defensible	科学上可辩护的	supported by procedure and evidence

USEFUL PHRASES

1. working code is not the same as a defensible analysis
2. encode recurring expert judgment once
3. load detailed guidance only when the task calls for it
4. keep data access thin and domain logic thick
5. make deterministic rules executable
6. separate procedural knowledge from the underlying model
7. version the instructions together with the rules
8. declare what the skill does not cover
9. route the task to the narrowest relevant skill
10. treat standards metadata as an authoritative source

COMPREHENSION QUESTIONS

1. Why can executable code still be scientifically indefensible?
2. How does progressive disclosure reduce context burden?
3. Why is a reusable skill different from a stateless tool?
4. What does ClinAgent's 72.1% derivation accuracy imply?
5. Why should standards, validators, study context and human review stay separate?

RETELLING PROMPTS

30 sec: define an agent skill and progressive disclosure.

45 sec: description -> routing -> skill -> refs/scripts -> tool -> result.

60 sec: divide responsibility among LLM, skill, validator, CDISC and reviewer.

5-MINUTE SPEAKING / OUTPUT TASK

Minute 1: choose ADAE, ADLB, log QC, TFL denominators, PK windows or terminology.

Minutes 2-3: define description, scope, declines, references, procedure, rules and tools.

Minute 4: add version, CDISC applicability, study config, tests and edge cases.

Minute 5: explain why the knowledge belongs in a reusable skill, not one giant prompt.

ONE SENTENCE TO KEEP

A strong SP agent should not rediscover professional procedure from model memory; recurring expertise should be versioned, reviewable, selectively loaded and paired with deterministic checks where correctness must be exact.