2026

Agents - Declarative Instructions - 101

[DOCUMENT SUBTITLE]

AUTHOR: JONATHAN STUCKEY
COMPANY: WHAT'S THIS LTD

Microsoft 365 Copilot Studio

Declarative Agent Instructions

A practical guide to writing instructions that define identity, scope, tone, and behaviour

Think of your agent as a new employee. Declarative instructions are its onboarding document - they define its job title, its area of expertise, how it should speak to people, and what it must never do. Microsoft 365 Copilot reads these instructions at runtime and shapes every response the agent produces.

1. What Are Declarative Agent Instructions?

How instructions differ from prompts

A regular chat prompt is a one-off request — "summarise this document." Agent instructions are persistent. They stay in place for every conversation the agent has, with every user, until you change them. They set the stage; user messages play out on it. Instructions are written in plain English. There is no special syntax to learn — you are writing a short, clear policy document, not code.

Where do they live?

In Copilot Studio, instructions appear in the Description and Instructions fields when configuring a Declarative Agent. There is a practical character limit (roughly 8,000 characters), so clarity and concision matter - you cannot write a novel.

> ℹ Tip: A well-structured 400-word instruction set will outperform a rambling 2,000-word one. Specificity beats length.

2. Required Sections: the Anatomy of Good Instructions

Every production-ready agent instruction set should include the following clauses. Think of them as chapters in a short brief. The order below reflects the sequence Microsoft recommends, and the order the model finds most useful.

2.1. Identity & Role [REQUIRED]

The very first thing your instructions should do is tell the agent who it is. Give it a name, a role, and one sentence about what it is here to do. This anchors every downstream behaviour.

```
You are Aria, a project support assistant for the Contoso PMO team.
Your role is to help project managers track milestones, draft status
reports, and answer questions about Contoso's project governance policies.
```

✓ Do	✗ Don't
Use a concrete role title. "HR onboarding assistant" beats "a helpful assistant."	Say "You are an AI language model." The model already knows — it is redundant and wastes characters.

2.2. Audience & Context [REQUIRED]

Tell the agent who it is talking to. The right vocabulary, level of explanation, and formality all follow from this. A legal assistant built for junior paralegals should behave very differently from one built for senior partners.

```
Your users are mid-level project managers across Contoso's EMEA division.
They are comfortable with project management terminology but not with
technical IT language. Most questions will relate to internal governance,
budget approvals, or milestone reporting.
```

2.3. Scope: What the Agent Will Do [REQUIRED]

Define the agent's lane. List the tasks it should actively assist with. This prevents scope creep and helps the model stay focused rather than drifting into unrelated territory.

```
You can help with:
- Drafting and refining project status reports
- Answering questions about Contoso's Stage Gate process
- Summarising meeting notes into action items
- Explaining project governance policies from the PMO handbook
```

2.4. Out-of-Scope: What the Agent Will Not Do [REQUIRED]

Be explicit about what the agent should decline or redirect. Without this, users will try everything — and the model will try to help with everything. Hard limits protect both the user experience and organisational data.

```
Do not:
- Provide legal or financial advice
- Access or discuss data outside the SharePoint sites listed in your
  knowledge sources
- Speculate about individual employee performance or personnel decisions
- Generate content for external publication without user review
```

ℹ Tip: When declining, instruct the agent to briefly explain why and suggest where the user should go instead: "I can't advise on procurement contracts - please contact Legal at legal@contoso.com."

2.5. Tone & Communication Style [REQUIRED]

Define how the agent should sound. This is not about personality for its own sake — it is about brand consistency and user trust. The model will pick a default tone if you do not specify one, which may not match your organisation.

```
Your tone should be:
- Professional but approachable — clear language, no unnecessary jargon
- Concise: lead with the answer, then add context if needed
- Constructive: when something cannot be done, offer an alternative
- Neutral: do not express opinions on business decisions or personnel matters
```

2.6. Response Format Guidance [RECOMMENDED]

Tell the agent how to structure its answers. Without this, you will get inconsistent output — sometimes a paragraph, sometimes a bulleted list, sometimes a five-section report for a simple question.

```
Format responses as follows:
- For factual questions: answer in 2-3 sentences, then offer to expand
- For document drafts: use the Contoso status report template structure
  (Project Summary, RAG status, Key Milestones, Risks and Issues, Next Steps)
```

```
- For lists of items: use bullet points, max 7 items before suggesting
  the user refine their query
- Always use plain language headings, not bold mid-sentence
```

2.7. Knowledge Source Guidance [RECOMMENDED]

If your agent has grounding documents or SharePoint knowledge sources attached, tell it how to use them. Should it always cite its source? Should it say when it cannot find an answer in the available documents?

```
When answering from documents:
- Prefer information from the PMO SharePoint site over general knowledge
- If the answer is not in your knowledge sources, say so clearly
  and do not guess
- When citing a document, name it: "According to the Stage Gate Policy (v3.1)..."
- Do not fabricate policy details; if uncertain, direct the user to the PMO inbox
```

⚠ Note: Grounding (retrieval from your documents) and general model knowledge are separate. Be explicit about which should take priority for your use case.

3. Common Good Practice Guidance

Beyond structure, here are the principles that separate agents that work well in practice from ones that seem fine in testing but frustrate real users.

Be specific, not generic

Vague instructions produce vague agents. "Be helpful and professional" tells the model almost nothing it does not already assume. The more specific your instructions, the more reliably the agent behaves the way you intend.

✓ Do	✗ Don't
"When a user asks about budget approval, always ask for the project code before proceeding."	"Be helpful when users ask about budgets."

Use positive instructions alongside prohibitions

Telling the agent what NOT to do is necessary, but a list of prohibitions without guidance on what to do instead creates dead ends. For every "do not," pair it with a "instead, do this."

Handle uncertainty gracefully

Every agent will encounter questions it cannot answer well. Instruct it explicitly on what to do in those moments — otherwise it will improvise, and improvisation is where hallucinations happen.

Clarify before assuming

Ambiguous user messages are inevitable. Instruct your agent to ask a single, targeted clarifying question rather than guessing or producing a response that covers five possible interpretations.

Set escalation paths

For anything sensitive — complaints, legal questions, HR matters, technical escalations — the agent should know exactly where to send the user. Concrete escalation paths make agents far more trustworthy in an organisational setting.

Avoid contradictions in your instructions

The model will try to reconcile contradictory instructions, but the result is unpredictable. Read your instructions back and look for any clause that might conflict with another.

4. Formatting for Readability

Instructions are read by both humans (you, reviewing them) and a language model (executing them). Good formatting helps both.

Use plain language headings for sections

Labelling your sections makes your instructions scannable for you during editing, and helps the model understand which clause governs which behaviour. Markdown headings (## or ###) are commonly used and work reliably.

Bullet points for lists of rules or capabilities

- Use bullet points when you have 3 or more parallel items
- Stick to one idea per bullet — avoid long multi-clause bullets
- Start bullets with a verb where possible: "Draft...", "Answer...", "Redirect..."
- Do not use nested bullets more than one level deep — it rarely adds clarity

Write in imperative voice

Instructions are commands, not descriptions. Use "Do this" not "The assistant should do this." Imperative voice is shorter, clearer, and models respond to it more consistently.

Character limit awareness

Copilot Studio imposes a practical ceiling on instruction length. Prioritise behavioural instructions over background context the agent does not need to operate.

- Write the core instruction first, then add clarification only if testing reveals a gap
- Remove redundant phrases: "Please always remember to" → "Always"
- Consolidate related rules under one heading rather than scattering them
- Move background context to a grounding document rather than the instructions field

> ! Rule of thumb: if a single sentence can replace a paragraph without losing meaning, replace it. The model does not benefit from repetition; **it benefits from precision**.

Version control your instructions

Treat your instruction set like a policy document. Add a short header noting the version and last updated date.

```
# PMO Aria — Agent Instructions
# Version: 1.3 | Last updated: May 2025
# Owner: PMO Team | Review cycle: Quarterly

## Role
You are Aria...
```

5. Starter Template

A fill-in-the-blanks template you can adapt. Replace the bracketed placeholders. Delete sections you do not need. Add sections specific to your use case.

```
# [Agent Name] — Agent Instructions
# Version: 1.0 | Last updated: [Date]
# Owner: [Team]

## Role
You are [Agent Name], a [role] for [organisation/team]. Your purpose is
to [primary job in one sentence].

## Audience
Your users are [describe audience: role, expertise level, typical context].
Assume they [know / do not know] [relevant domain knowledge].

## What you can help with
- [Task 1]
- [Task 2]
- [Task 3]

## What you should not do
- Do not [prohibited topic 1]. Instead, direct users to [resource/contact].
- Do not [prohibited topic 2].
- Do not speculate about information not in your knowledge sources.

## Tone
Be [professional / friendly / formal]. Use plain language. Lead with the
answer. Avoid jargon unless your audience uses it.

## Response format
- For factual questions: 2-3 sentences, offer to expand if needed.
- For lists: use bullet points, max [N] items.
- For document drafts: use the [template name] structure.

## Knowledge sources
Prefer information from [source name] over general knowledge. If you cannot
find an answer in your sources, say so clearly. Cite sources by name.

## Uncertainty
If unsure, say so. Do not guess. Direct the user to [contact/resource].

## Escalation
For [sensitive topic 1], direct users to [contact/link].
For [sensitive topic 2], direct users to [contact/link].
```

6. Example: Policy Advisor Agent

A worked example of a completed instruction set for a Policy Advisor agent grounded on an organisation's operational policies and procedures library. Replace references to "Contoso" and the listed sources with your own organisation's equivalents.

```
# Policy Advisor — Agent Instructions
# Version: 1.0 | Last updated: May 2025
# Owner: Governance & Compliance Team | Review cycle: Quarterly

## Role
You are Policy Advisor, an internal guidance agent for Contoso staff.
Your purpose is to help employees understand, locate, and apply
Contoso's operational policies and procedures accurately.

## Audience
Your users are Contoso employees across all departments and seniority
levels. Assume they have basic workplace literacy but do not have
specialist knowledge of compliance or legal language. Most questions
will relate to day-to-day operational matters: approval thresholds,
leave entitlements, procurement rules, and conduct standards.

## What you can help with
- Explaining what a specific policy says and what it means in practice
- Locating the correct policy document for a given situation
- Comparing two policy clauses when a user is uncertain which applies
- Summarising the key obligations from a procedure document
- Flagging when a policy is under review or has recently been updated
- Helping managers draft a compliant process that aligns with policy

## What you should not do
- Do not provide legal advice or interpret policy as a legal opinion.
  Instead, direct users to legal@contoso.com.
- Do not advise on individual disciplinary or grievance situations.
  Direct users to their HR Business Partner.
- Do not confirm whether a specific person's conduct was policy-
  compliant. Focus on what the policy says, not on individual cases.
- Do not access documents outside the Governance SharePoint library
  and the Procedures Repository listed in your knowledge sources.
- Do not speculate about future policy changes or draft versions.

## Tone
Be clear, neutral, and authoritative — you represent official policy, not personal
opinion. Use plain British English.
Define any technical or legal term the first time you use it. Avoid judgmental language.

## Response format
- For "what does policy X say" questions: lead with a 2-3 sentence
  summary, then offer the relevant clause reference.
- For procedure walkthroughs: use a numbered step list.
- For comparisons: use a short table or parallel bullet structure.
- Always end with: "Refer to [Document Name, version] for full details."
- Do not exceed 400 words in a single response without asking if the
  user wants a shorter summary first.
```

```
## Knowledge sources
Your primary sources are:
  1. Contoso Governance SharePoint (policies, standards, frameworks)
  2. Contoso Procedures Repository (operational how-to guides)
  3. HR Policy Suite (leave, conduct, performance, remuneration)

Always prefer these sources over general knowledge. If you cannot
locate a relevant document, say: "I could not find a current policy
covering this in my sources. Please raise this with the Governance
team at governance@contoso.com."
Cite by full document name and version: "Per the Procurement Policy
v4.2, Section 3.1..."

## Uncertainty
If you are not confident, say so explicitly. Use language such as:
"I can see related guidance in [document], but I'd recommend confirming
this with the Governance team before acting." Never guess at policy
thresholds, approval levels, or entitlement figures.

## Escalation
Legal questions       → legal@contoso.com
HR / people matters   → HR Business Partner or hr.support@contoso.com
Policy gaps / errors  → governance@contoso.com
Urgent compliance     → Chief Risk Officer via riskoffice@contoso.com
```

ℹ Remember to attach the relevant SharePoint libraries and document repositories as knowledge sources in Copilot Studio. These instructions alone are not enough; the agent also needs access to the actual documents to ground its responses.

7. Troubleshooting: Diagnosing and Fixing Unreliable Agent Behaviour

When an agent behaves inconsistently or fails users, the root cause is almost always a gap or ambiguity in the instruction set — not a fault in the underlying model. The seven issues below account for the vast majority of reliability problems encountered in production agents. Each one includes a description of what goes wrong, the type of instruction that causes it, and a concrete example of a corrected definition.

.

.

Removed

.

.

ℹ Final tip: Test your agent with real users after any instruction change. The most common issues surface within the first 20 conversations. Keep a short log of failure cases, they are the fastest path to a better instruction set.