FREE PATTERN LIBRARY

Prompt Patterns for Business Teams

Eight reusable patterns for everyday business work. Each one tells you when to use it, gives you the template, shows a worked example — and names the most common way it fails, because knowing the failure mode is what separates a workflow from a party trick. These are the QA patterns I build into client systems.

Amy Sullivan · Buni LLC · sullivanamy75@gmail.com · luvbuniz.github.io/enterprise-portfolio1

EXTRACTION

1 · Meeting Notes → Action Items

When to use it: after any meeting where decisions and tasks are buried in a transcript or rough notes, and the team needs an executable list — not a recap.

TEMPLATE

```
You are turning meeting notes into an action-item list that someone who missed the meeting can execute.

Meeting notes:
[PASTE NOTES OR TRANSCRIPT]

Produce three sections:
1. Decisions made — one line each.
2. Action items — a table with columns: Task, Owner, Deadline, Depends on.
3. Open questions — anything discussed but not resolved.

Rules:
- Use only what is in the notes. If an owner or deadline was not stated, write "not assigned" — do not guess.
- After each action item, quote the phrase from the notes that justifies it.
```

EXAMPLE INPUT

A 40-minute project kickoff transcript: overlapping speakers, tangents, decisions made in passing ("yeah let's just go with the March date then").

OUTPUT SKETCH

3 decisions, a 7-row action table with 2 items flagged "not assigned," 3 open questions — each action item traceable to a quoted line.

MOST COMMON FAILURE MODE

The model invents owners and deadlines to make the table look complete. The "do not guess" rule and the required source quote exist specifically to block that.

WRITING

2 · Messy Doc → SOP

When to use it: when a process lives in someone's head, an email thread, or a rambling wiki page — and it needs to become a procedure a new hire can follow without asking anyone.

TEMPLATE

```
Convert the source material below into a standard operating procedure a brand-new team member could follow unaided.

Source material:
[PASTE DOC, THREAD, OR NOTES]

SOP format:
- Purpose (2 sentences max)
- Who does this, and when it's triggered
- Prerequisites (access, tools, information needed before starting)
- Numbered steps. Each step ends with a "Done when:" completion check.
- Escalation: what to do when a step fails, and who to contact

Rules:
- Do not invent steps. Where the source is silent or ambiguous, insert a line reading "GAP: [what's missing]" and continue.
- Keep the original terminology the team already uses.
```

EXAMPLE INPUT

A three-page email thread where a departing employee explains monthly invoice processing, out of order, with two contradictions.

OUTPUT SKETCH

A 12-step SOP with completion checks, plus 2 GAP flags — one where the thread never says who approves amounts over a threshold.

MOST COMMON FAILURE MODE

The model papers over holes with plausible invented steps, and the SOP reads complete while being wrong. Forcing explicit GAP markers turns missing knowledge into a visible to-do instead of a hidden landmine.

WRITING

3 · First Draft with Source Grounding

When to use it: drafting anything where accuracy matters — a client memo, a report section, a board update — and you have source documents the draft must stay faithful to.

TEMPLATE

```
Write a first draft grounded ONLY in the sources below.

Sources:
[S1] — [PASTE OR SUMMARIZE SOURCE 1]
[S2] — [PASTE OR SUMMARIZE SOURCE 2]

Draft to write: [WHAT IT IS — memo, report section, update]
Audience: [WHO READS IT]
Length: [TARGET LENGTH]

Rules:
- Every factual claim in the draft ends with its source tag, like [S1].
- Anything you believe is true but cannot tag goes in a separate "Needs verification" list at the end —
not in the draft.
- If sources conflict, say so in the draft rather than picking one silently.
```

EXAMPLE INPUT

Two internal research summaries about customer churn, plus "audience: VP of Customer Success, 500 words."

OUTPUT SKETCH

A 500-word memo with every claim tagged [S1] or [S2], and a 2-item "Needs verification" list of claims the model wanted to add but couldn't source.

MOST COMMON FAILURE MODE

The model blends its general knowledge into the draft so smoothly that sourced and unsourced claims read identically. The tag-or-quarantine rule keeps the two separable — which is the whole point.

REVIEW

4 · Rubric-Based Quality Review

When to use it: reviewing drafts or deliverables when you need consistent judgments across different reviewers, days, and moods — the same structure professional AI data-quality evaluation runs on: anchored scales, evidence quotes, per-criterion scores.

TEMPLATE

```
Review the artifact below against this rubric. Score each criterion independently before any overall judgment.

Rubric (score each 1–5):
- [CRITERION 1, e.g. Clarity]: 1 = [WHAT A 1 LOOKS LIKE], 3 = [WHAT A 3 LOOKS LIKE], 5 = [WHAT A 5 LOOKS LIKE]
- [CRITERION 2]: 1 = …, 3 = …, 5 = …
- [CRITERION 3]: 1 = …, 3 = …, 5 = …

Artifact to review:
[PASTE ARTIFACT]

For each criterion: the score, a one-sentence justification, and a direct quote from the artifact as evidence.
Then: the two highest-impact fixes, in order.
Do not give an overall score until every criterion is scored.
```

EXAMPLE INPUT

A sales one-pager, reviewed against a 5-criterion rubric (clarity, proof, audience fit, call to action, length discipline) with anchored scale points.

OUTPUT SKETCH

A score table with quoted evidence per criterion — proof scores 2 with the offending vague sentence quoted — then two prioritized fixes.

MOST COMMON FAILURE MODE

Grade inflation — everything scores 4 out of 5 and the review is useless. Anchoring what a 1, 3, and 5 actually look like, and demanding quoted evidence, is what forces the scores to mean something.

WRITING

5 · Persona-Based Stakeholder Rewrite

When to use it: one piece of content needs to land with different audiences — the executive, the front-line team, the customer — without maintaining three documents by hand.

TEMPLATE

```
Rewrite the source content for the reader profile below.

Reader profile:
- Role: [WHO THEY ARE]
- What they care about: [THEIR PRIORITIES]
- Time budget: [e.g. 30 seconds / 5 minutes]
- What they already know: [CONTEXT THEY HAVE]
- Decision they need to make after reading: [THE DECISION]

Source content:
[PASTE SOURCE]

Rules:
- Facts stay identical to the source. Only selection, order, and framing change.
- Format for the reader: if the time budget is short, lead with the decision and keep it under [N] lines.
- End with exactly what you need from them, stated once.
```

EXAMPLE INPUT

A detailed engineering status update, rewritten for a COO with a 30-second budget who must decide whether to approve two more weeks.

OUTPUT SKETCH

A 4-line summary: current state, the one risk that matters, the ask ("approve 2-week extension"), and the cost of saying no.

MOST COMMON FAILURE MODE

"Simplifying" quietly changes the claims — numbers get rounded into fiction, hedged findings become confident ones. The "facts stay identical" rule is the guardrail; spot-check the numbers anyway.

EXTRACTION

6 · Structured Data Extraction from Documents

When to use it: pulling the same fields out of many documents — contracts, invoices, applications, intake forms — where the output feeds a spreadsheet or system and errors compound downstream.

TEMPLATE

```
Extract the following fields from the document below.

Fields:
- [FIELD 1] — [TYPE / ALLOWED VALUES, e.g. date as YYYY-MM-DD]
- [FIELD 2] — [TYPE]
- [FIELD 3] — [TYPE]

Document:
[PASTE DOCUMENT]

Output a table with three columns: Field, Extracted value, Verbatim quote from the document supporting it.

Rules:
- If a field is not stated in the document, the value is "NOT FOUND" — never infer or compute it.
- The quote column must be word-for-word from the document. No quote, no value.
```

EXAMPLE INPUT

A vendor contract; fields: parties, effective date, term length, renewal type, termination notice period, payment terms.

OUTPUT SKETCH

A 6-row table, each value paired with its supporting quote; renewal type comes back NOT FOUND because the contract never states it — caught in seconds.

MOST COMMON FAILURE MODE

Silent normalization: the model converts "Net 30 from invoice date" into an invented calendar date, or guesses a missing field from context. The verbatim-quote column makes every value auditable by a human in seconds.

DECISION

7 · Multi-Option Decision Memo

When to use it: a real decision with two or more viable options — vendor choice, build vs. buy, timing — where stakeholders need a document to react to instead of a meeting to relitigate.

TEMPLATE

```
Write a one-page decision memo.

Decision to make: [THE DECISION + DEADLINE]
Constraints: [BUDGET, TIMELINE, POLICY — anything non-negotiable]
Criteria, weighted: [e.g. Cost 40%, Time-to-value 30%, Risk 30%]

Options:
[OPTION A — WHAT IT IS]
[OPTION B — WHAT IT IS]
[OPTION C — WHAT IT IS]

Structure:
1. For each option: what it costs, what it gets us, its biggest risk, and what we'd have to believe for it
to be the right choice.
2. A comparison table scoring each option against the weighted criteria.
3. Recommendation with reasoning — including the strongest case AGAINST the recommended option.
4. "What would change this recommendation" — 2–3 concrete tripwires.

Build the table before stating the recommendation.
```

EXAMPLE INPUT

Three project-management vendors, a hard budget cap, migration deadline next quarter, criteria weighted toward time-to-value.

OUTPUT SKETCH

A one-page memo: weighted table, a recommendation that survives its own counter-argument, and two tripwires ("if migration exceeds 6 weeks, revisit").

MOST COMMON FAILURE MODE

The model picks a winner early — often mirroring a preference you hinted at — and back-fills the analysis to justify it. Requiring the table first and a genuine case against the winner keeps the memo honest.

WRITING

8 · Email Tone Calibration

When to use it: the email is drafted but the stakes make tone the hard part — an escalation, a decline, chasing an overdue payment, delivering bad news to a good client.

TEMPLATE

```
Calibrate the tone of this email without changing what it asks for.

Draft:
[PASTE DRAFT]

Context:
- Recipient and relationship: [WHO + HISTORY, e.g. "client of 3 years, second late invoice"]
- What I need to happen: [THE CONCRETE OUTCOME]
- Tone target: warmth [1–5], directness [1–5]

Produce two versions:
A) Minimal edit — keep my structure and voice, adjust only what misses the tone target.
B) Full rewrite — same facts and same ask, rebuilt for the tone target.

After each version, list what you changed and why, in one line per change.
The ask must survive in both versions, stated clearly, exactly once.
```

EXAMPLE INPUT

A blunt overdue-invoice email to a long-term client; target: warmth 4, directness 4 — friendly, but the payment date doesn't move.

OUTPUT SKETCH

Version A softens two sentences and keeps the deadline; version B restructures around the relationship first — both still name the date and the amount.

MOST COMMON FAILURE MODE

Over-correction into corporate mush: apologetic, hedged, and suddenly not asking for anything. Pinning "what I need to happen" as a hard requirement anchors the rewrite against dissolving the ask.

Want these tuned to your team's documents, tone, and failure modes? **sullivanamy75@gmail.com** — subject "Workflow Audit"

Case studies & live systems: **luvbuniz.github.io/enterprise-portfolio1**