

zango

The AI Compliance Layer for Financial Services

Reach us at

hello@zango.ai

www.zango.ai

[LinkedIn](#)

Zango's AI agents

First Line of Defense

◆ Alert Investigation

◆ SAR filing

◆ Controls Testing

Second Line of Defense

◆ Regulatory Intelligence

◆ Horizon Scanning

◆ Policy Gap Analysis

◆ Controls Gap Analysis

◆ Marketing Compliance

◆ QA/QC

Third Line of Defense

The Problems our AI agents solve

Regulatory Change Management

- Noise over signal
- Inconsistent interpretation
- Slow and manual impact analysis

Fragmented Policies

- Across multiple repositories
- Not linked to regulatory obligations

Growth Bottlenecks

- Manual and inconsistent marketing and product approvals
- Difficult to evidence to regulators

Manual Controls Testing

- Manual controls testing for regulated process such as KYC, Claims, etc

Capacity unlocked: A leading European bank's transformation

Before Zango (Horizon Scanning)

 8 FTEs

 External Consultant + SaaS platform

 Rules extraction (manual)

 SaaS Platform usage <10%

After Zango (Horizon Scanning)

 3 FTEs (5 FTEs re-deployed to assurance)

 One Single Platform

 Rules extraction (automated)

 AI Platform used daily

Masterclass Slides

AI for compliance: spotting where AI goes rogue

AI for compliance: spotting where AI goes rogue

01 The AI confidence gap in compliance

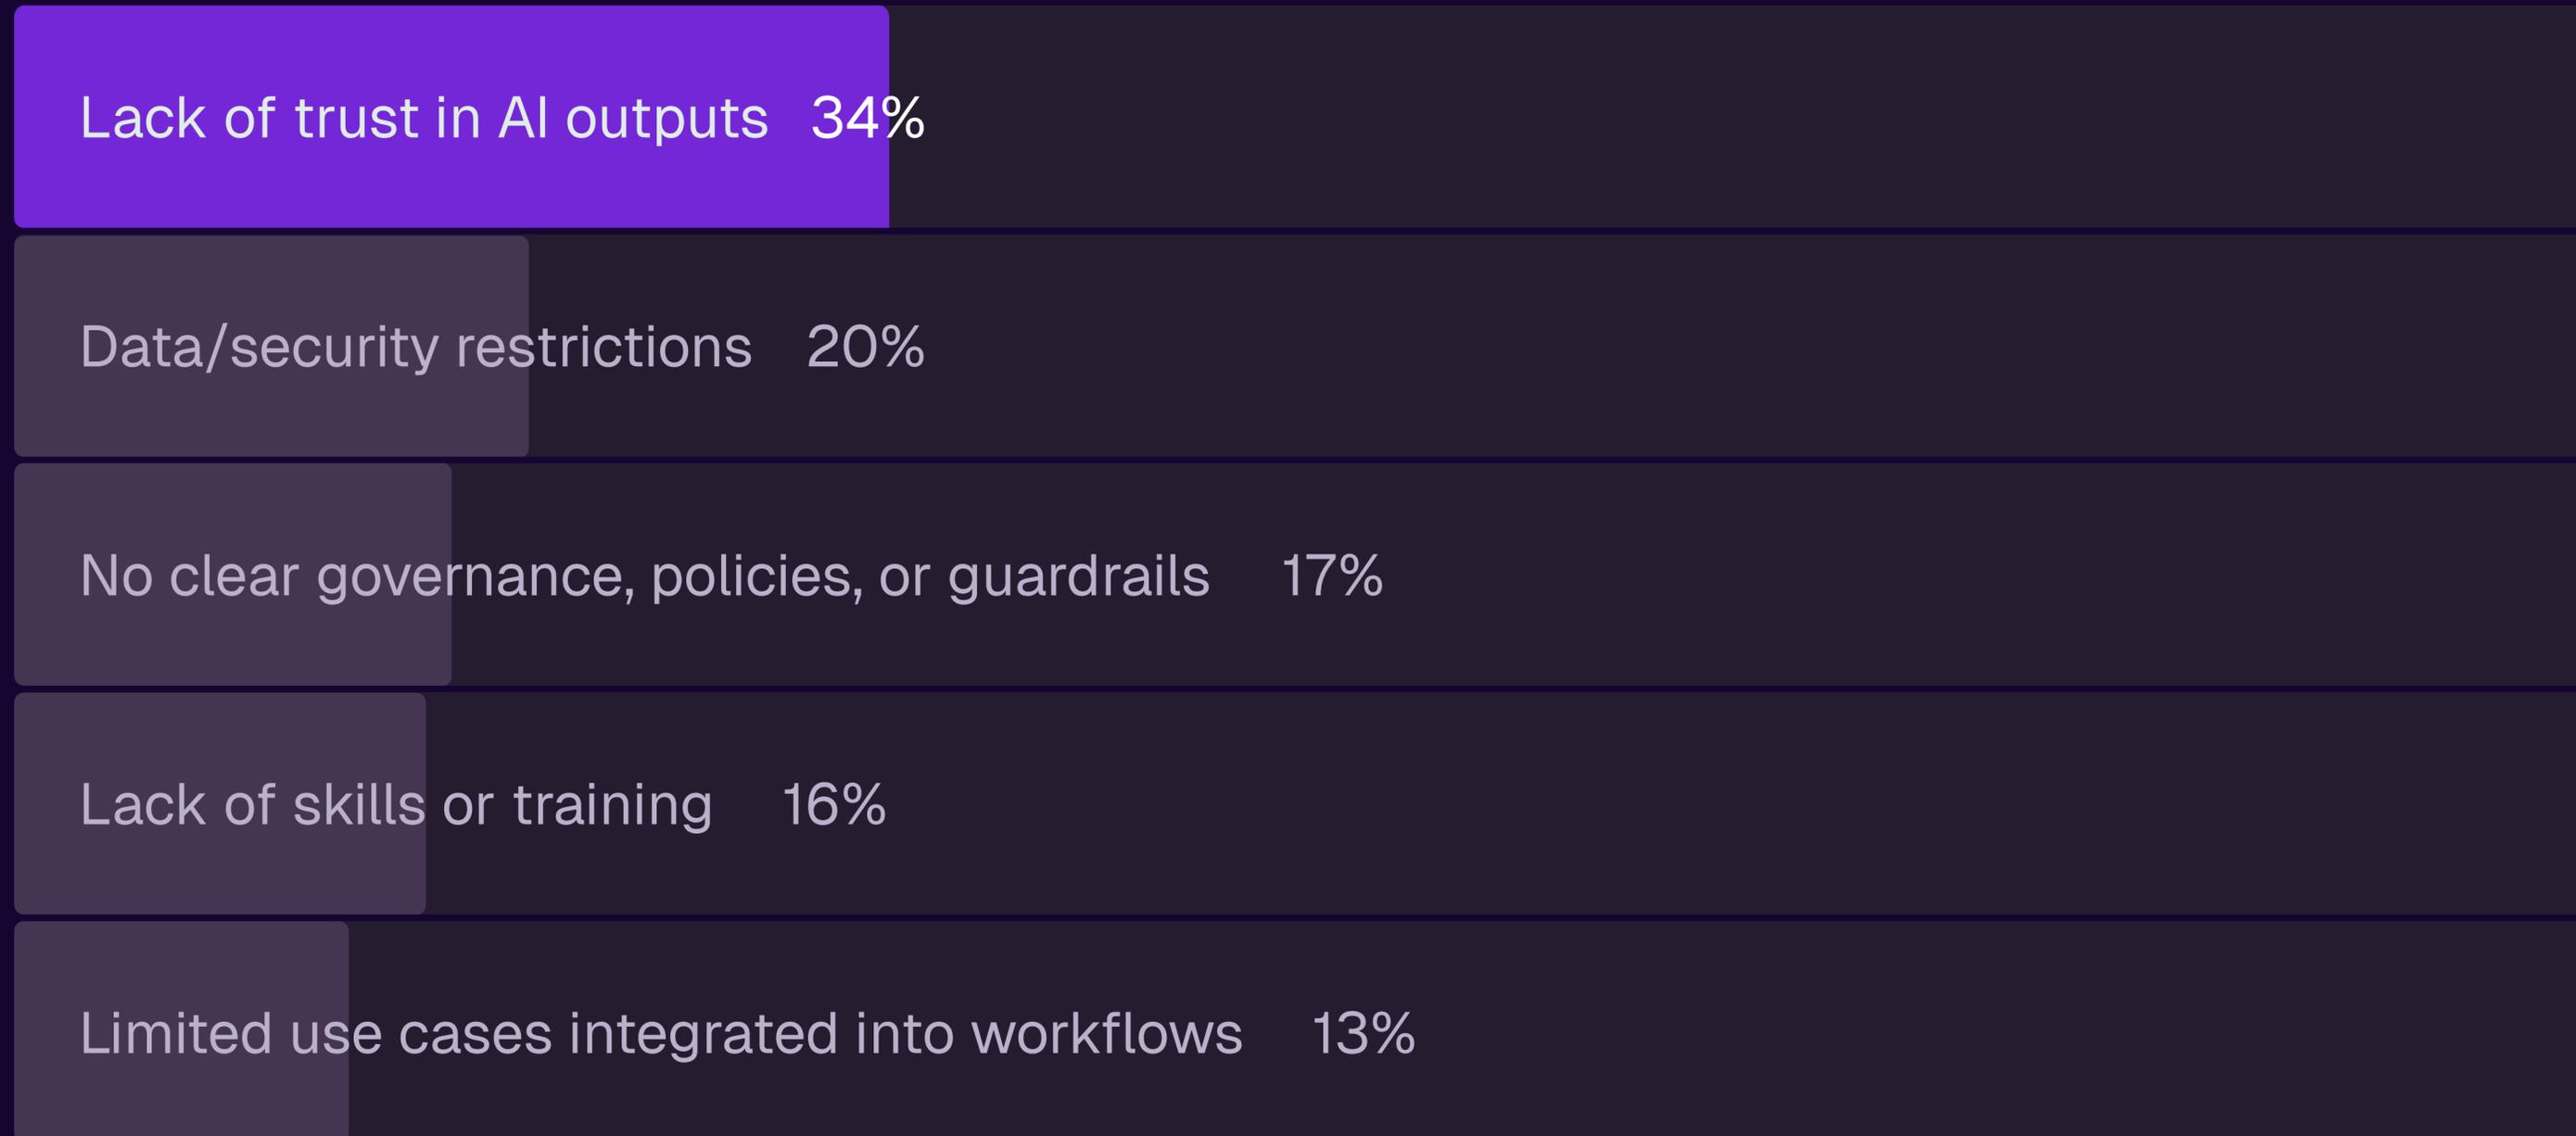
02 Why existing QA functions aren't effective for AI

03 Bridging the confidence gap

04 Evaluating AI outputs

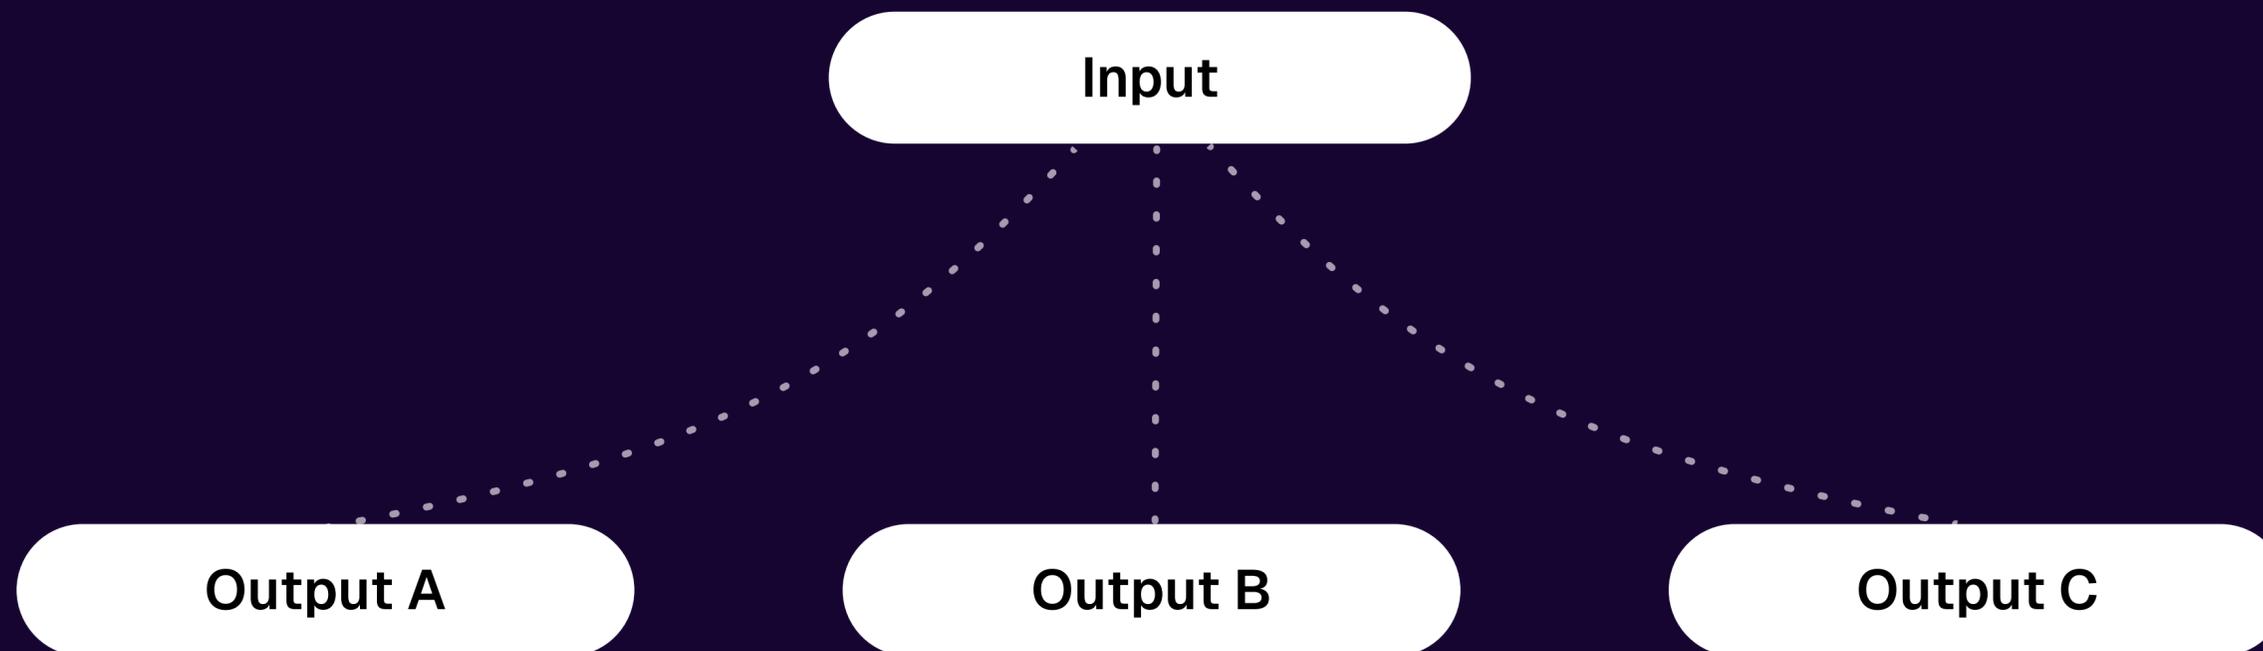
05 Case study: Monzo

What's the biggest barrier to using AI confidently in your organisation?



Why existing QA functions aren't effective for LLMs

→ The same prompt can generate multiple plausible outputs



How much water is used?

You're giving feedback on a new version of ChatGPT.

Which response do you prefer? Responses may take a moment to load.



Response 1

The NVIDIA GB200 NVL72's liquid cooling system circulates coolant at a rate of 2 liters per second, entering at 25°C and exiting at 45°C. This flow rate equates to 7,200 liters per hour. The system's design enhances energy efficiency and compute density, reducing the need for extensive air cooling and lowering energy consumption. [The Register](#)

Sources

I prefer this response



Response 2

The NVIDIA GB200 NVL72's liquid cooling system circulates coolant at a rate of approximately 2 liters per second, with an inlet temperature of 25°C and an outlet temperature 20°C higher. [The Register](#) This flow rate is designed to effectively dissipate the substantial heat generated by the system's high-performance components.

Sources

I prefer this response

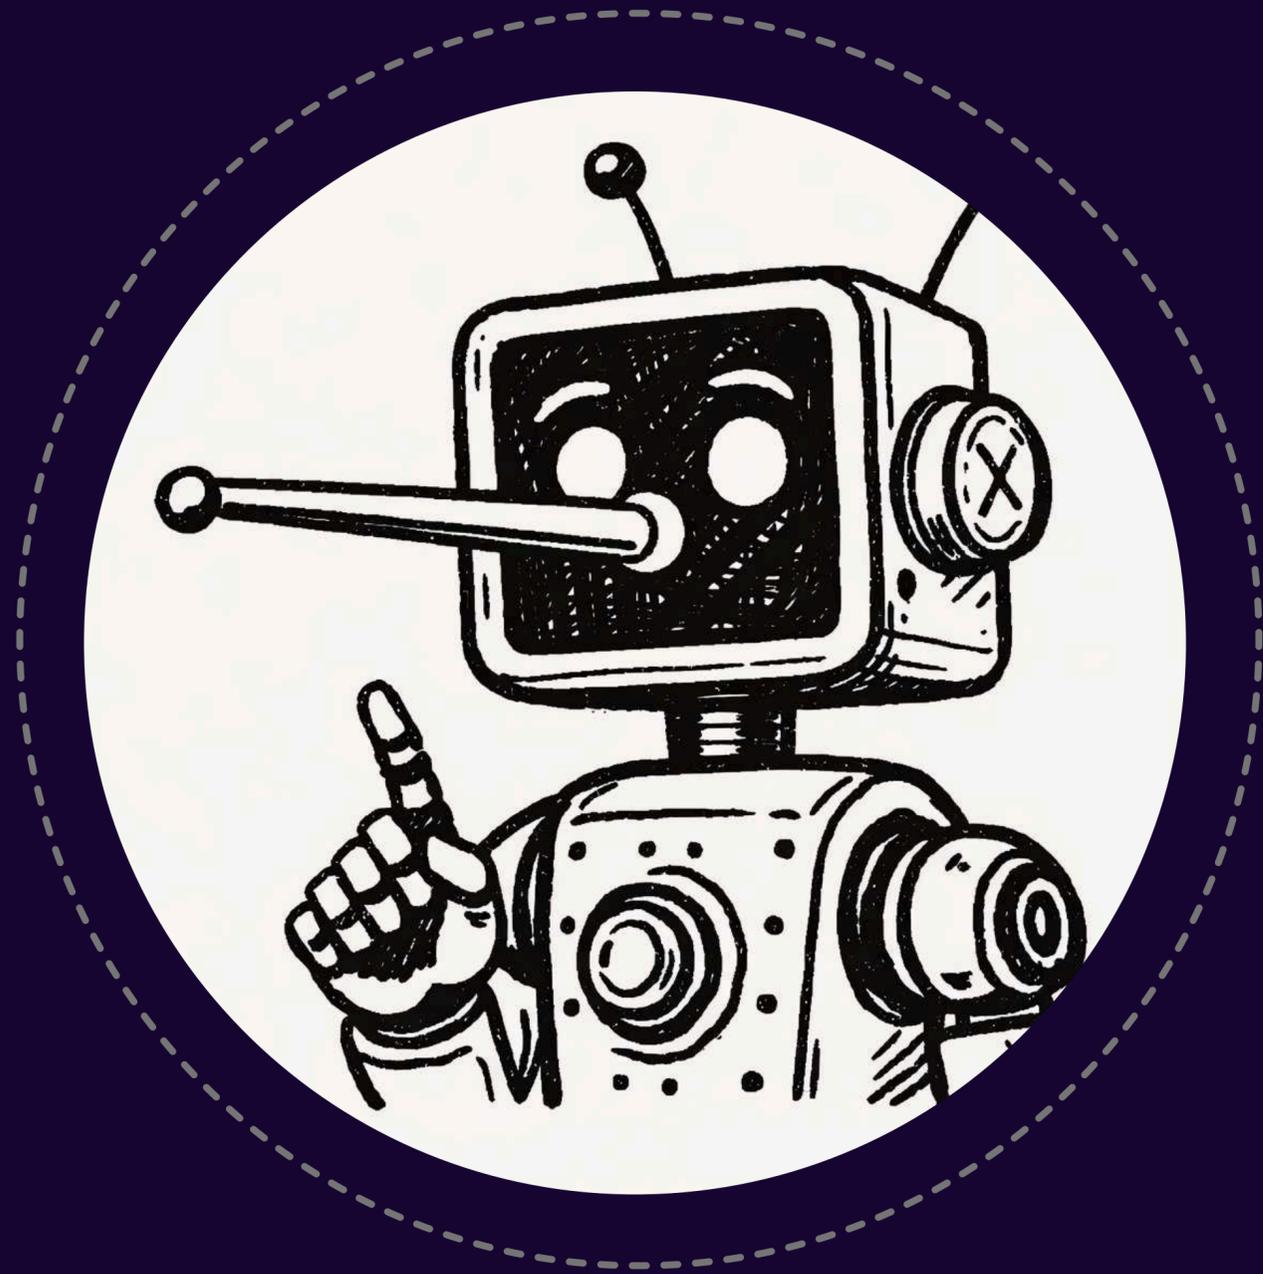
Poll

How often do you come across work that is clearly AI produced and error ridden?

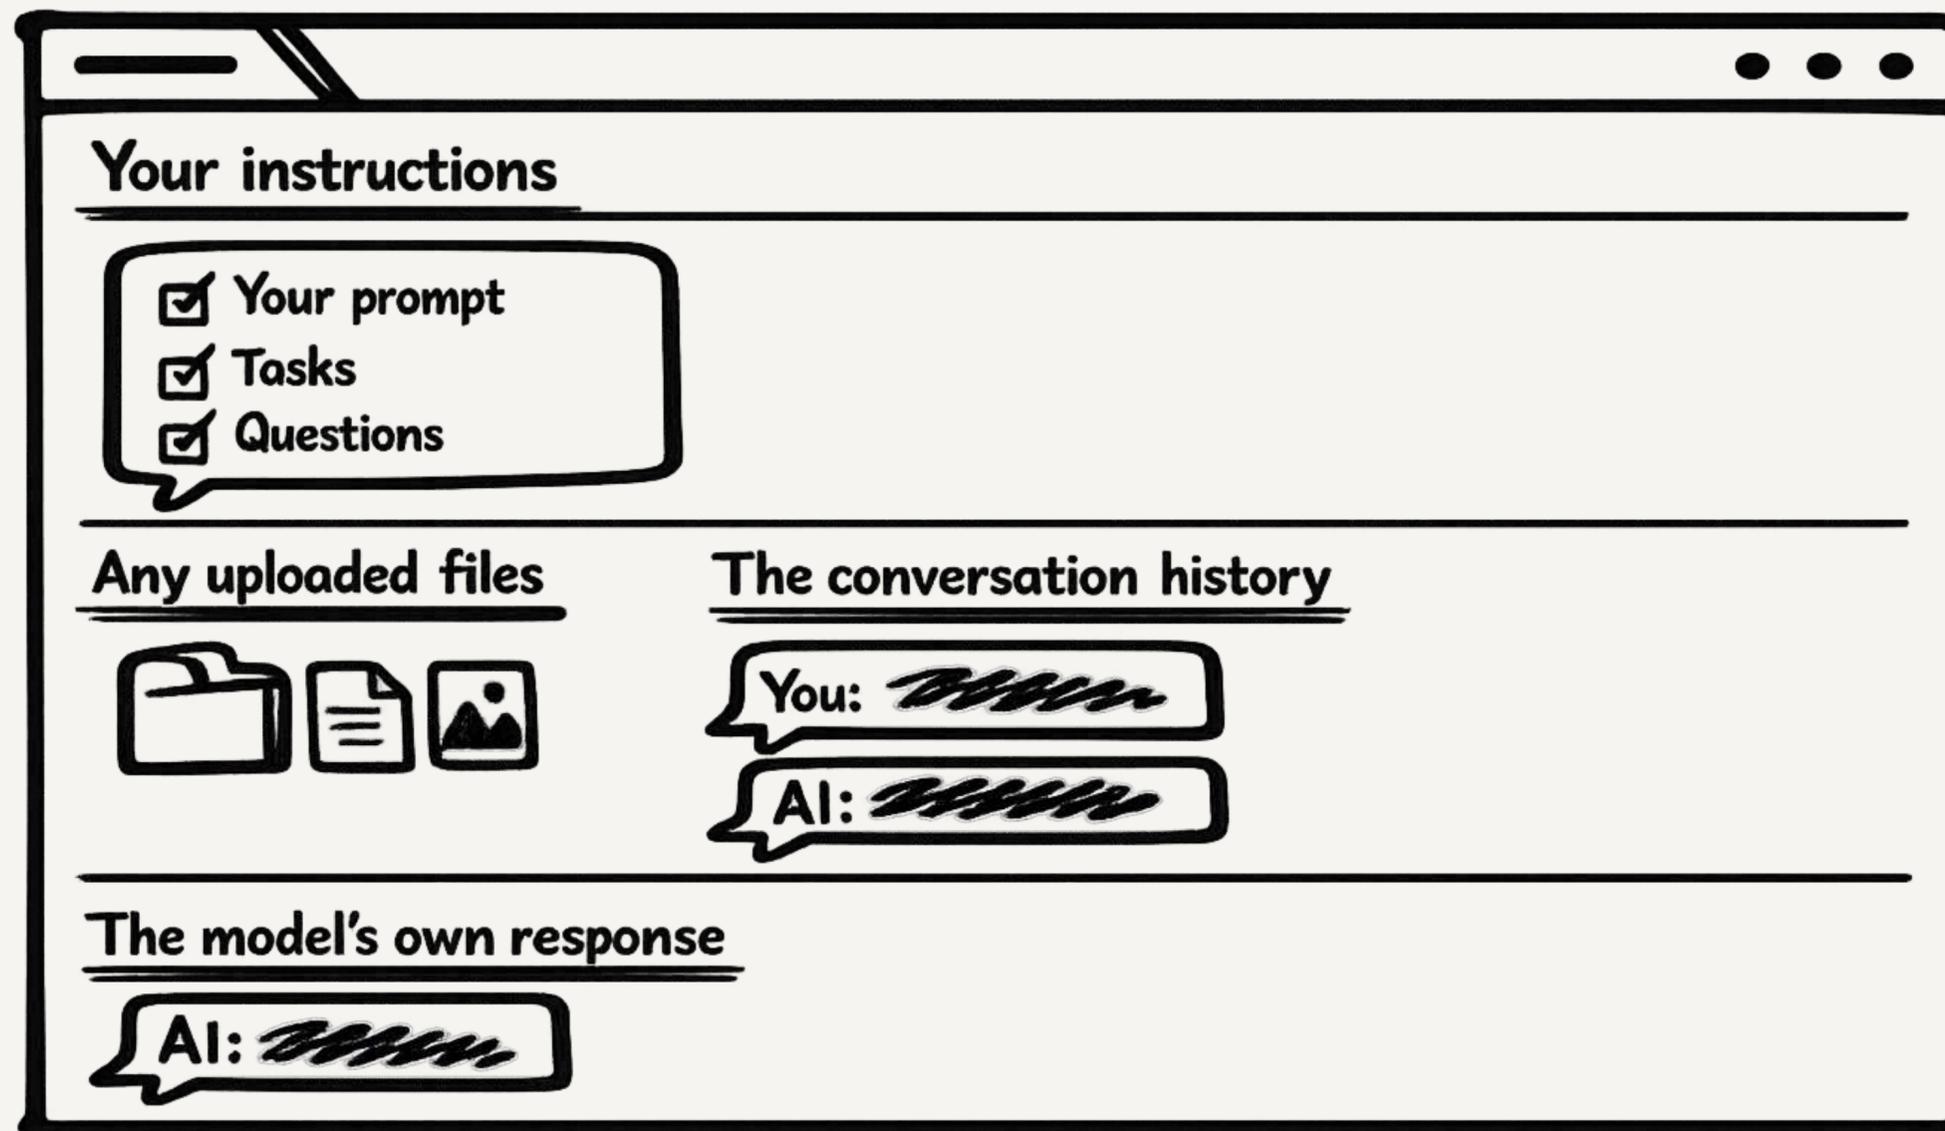
Bridging the confidence gap



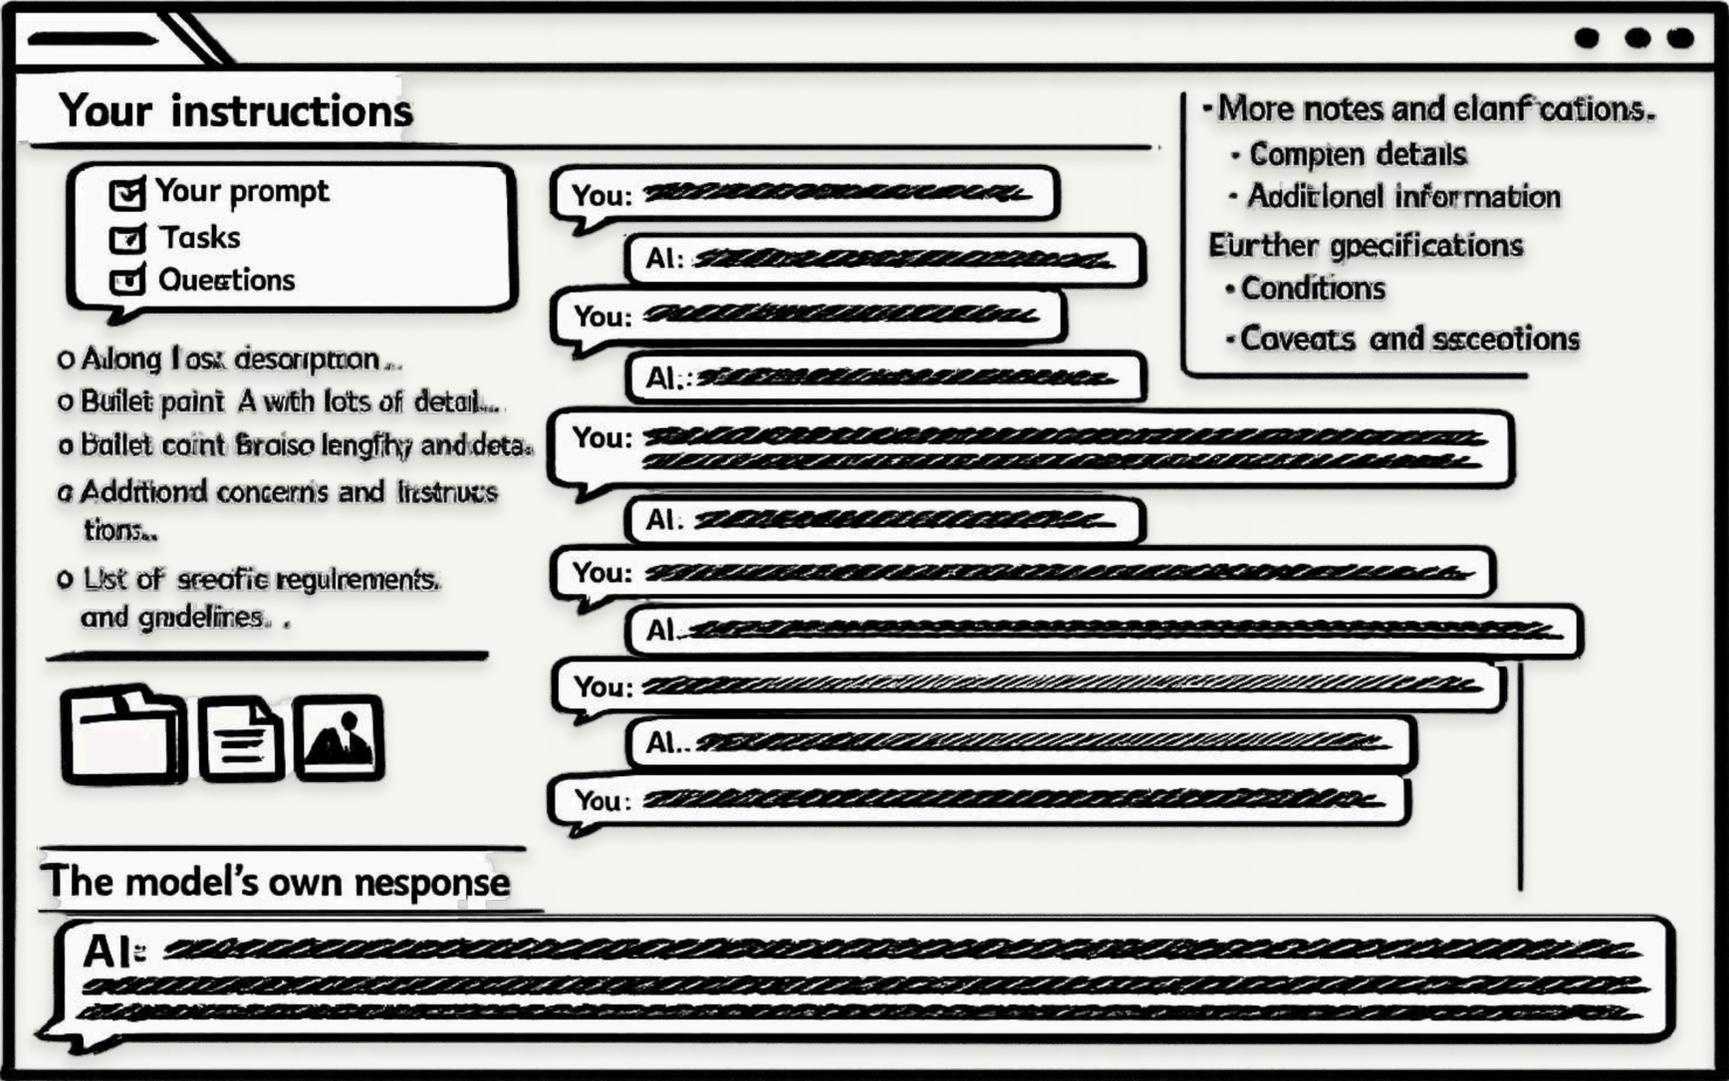
The context problem



The context window

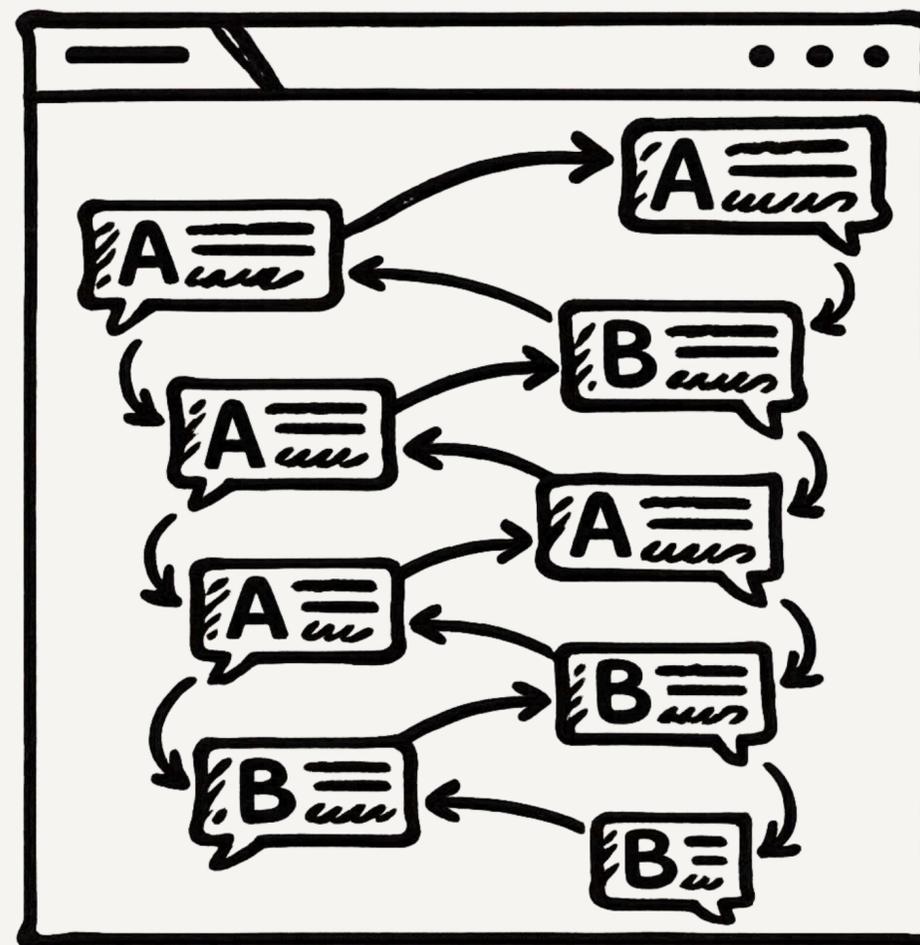


Source: Zango AI



Source: Zango AI

Long conversations



Prompt reset

If you had to complete this task perfectly using a single, well-structured prompt — what would that prompt be?

- 1. Rewrite my request as that improved, precise, high-quality single prompt.**
- 2. Identify:**
 - **Objective**
 - **Audience**
 - **Inputs assumed**
 - **Scope boundaries**
 - **Output format**
 - **Quality constraints**

Long documents



COMMISSION IMPLEMENTING REGULATION (EU) 2025/847 of 6 May 2025

laying down rules for the application of Regulation (EU) No 910/2014 of the European Parliament and of the Council as regards reactions to security breaches of European Digital Identity Wallets

THE EUROPEAN COMMISSION,

Having regard to the Treaty on the Functioning of the European Union,

Having regard to Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC (...), and in particular Article 5e(5) thereof,

Whereas:

- (1) The European Digital Identity Framework (framework) set out in Regulation (EU) No 910/2014 is a crucial component in the establishment of a secure and interoperable digital identity ecosystem across the Union. With the European Digital Identity Wallets (wallets) as its cornerstone, the framework aims to facilitate access to services across Member States, while ensuring the protection of personal data and privacy.
- (2) Regulations (EU) 2016/679 (5) and (EU) 2018/1725 (6) of the European Parliament and of the Council, and, where relevant, Directive 2002/58/EC of the European Parliament and of the Council (7) apply to the personal data processing activities under this Regulation. The rules on the assessment and provision of information established under this Regulation are without prejudice to the obligation to notify personal data breaches to the competent supervisory authority where applicable under Regulation (EU) 2016/679 or Regulation (EU) 2018/1725, and to the obligation to communicate the personal data breaches to the data subjects where applicable under these Regulations.
- (3) The Commission regularly assesses new technologies, practices, standards and technical specifications. To ensure the highest level of harmonisation among Member States for the development and certification of the wallets, the technical specifications set out in this Regulation rely on the work carried out under Commission Recommendation (EU) 2021/946 (8) and in particular the Architecture and Reference Framework which is part of it. In accordance with recital 75 of Regulation (EU) 2024/1183 of the European Parliament and of the Council (9), the Commission should review and, if necessary, update this Regulation, to keep it in line with global developments and the Architecture and Reference Framework, and to follow the best practices on the internal market.

(5) OJ L 257, 28.8.2014, p. 73, ELI: <http://data.europa.eu/eli/reg/2016/679/oj>.

(6) Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (OJ L 119, 4.5.2016, p. 1, ELI: <http://data.europa.eu/eli/reg/2016/679/oj>).

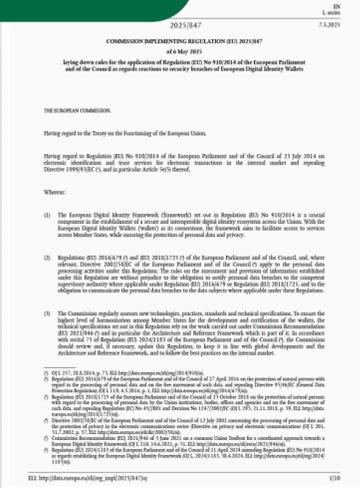
(7) Regulation (EU) 2018/1725 of the European Parliament and of the Council of 23 October 2018 on the protection of natural persons with regard to the processing of personal data by the Union institutions, bodies, offices and agencies and on the free movement of such data, and repealing Regulation (EC) No 45/2001 and Decision No 1247/2002/EC (OJ L 295, 21.11.2018, p. 39, ELI: <http://data.europa.eu/eli/reg/2018/1725/oj>).

(8) Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) (OJ L 201, 31.7.2002, p. 37, ELI: <http://data.europa.eu/eli/dir/2002/58/oj>).

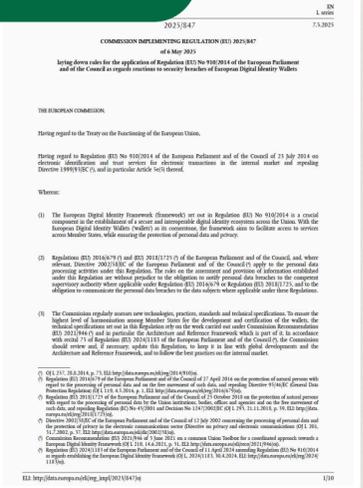
(9) Commission Recommendation (EU) 2021/946 of 3 June 2021 on a common Union Toolbox for a coordinated approach towards a European Digital Identity Framework (OJ L 210, 14.6.2021, p. 51, ELI: <http://data.europa.eu/eli/reco/2021/946/oj>).

(10) Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework (OJ L 2024/1183, 30.4.2024, ELI: <http://data.europa.eu/eli/reg/2024/1183/oj>).

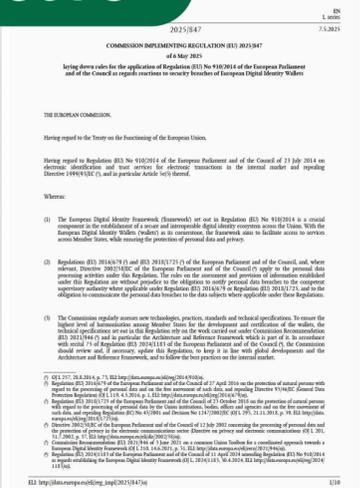
Set A



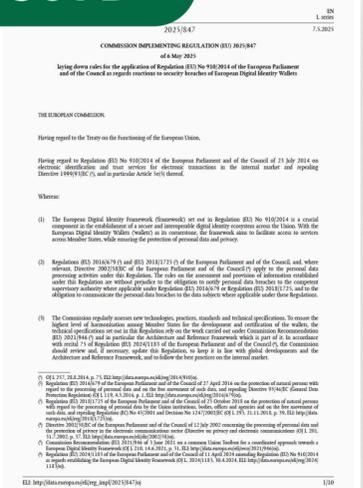
Set B



Set C



Set D



Poor context



Context injection

You are acting as **[ROLE]**.

You are operating in the context of **[INDUSTRY / REGULATED ENVIRONMENT]**.

The jurisdiction(s) you should assume are **[JURISDICTION]**.

The type of organisation is **[ORG TYPE]**, and the relevant audience is **[AUDIENCE]**.

Your responses must prioritise **[PRIMARY RISK / CONSTRAINT]**, and avoid **[KEY FAILURE MODES]**.

If information is missing or uncertain, you must **[HOW TO HANDLE UNCERTAINTY]**.

Evaluating AI outputs



Evaluation techniques



**LLM as a
judge**



**Confidence
scoring**

LLM as a judge



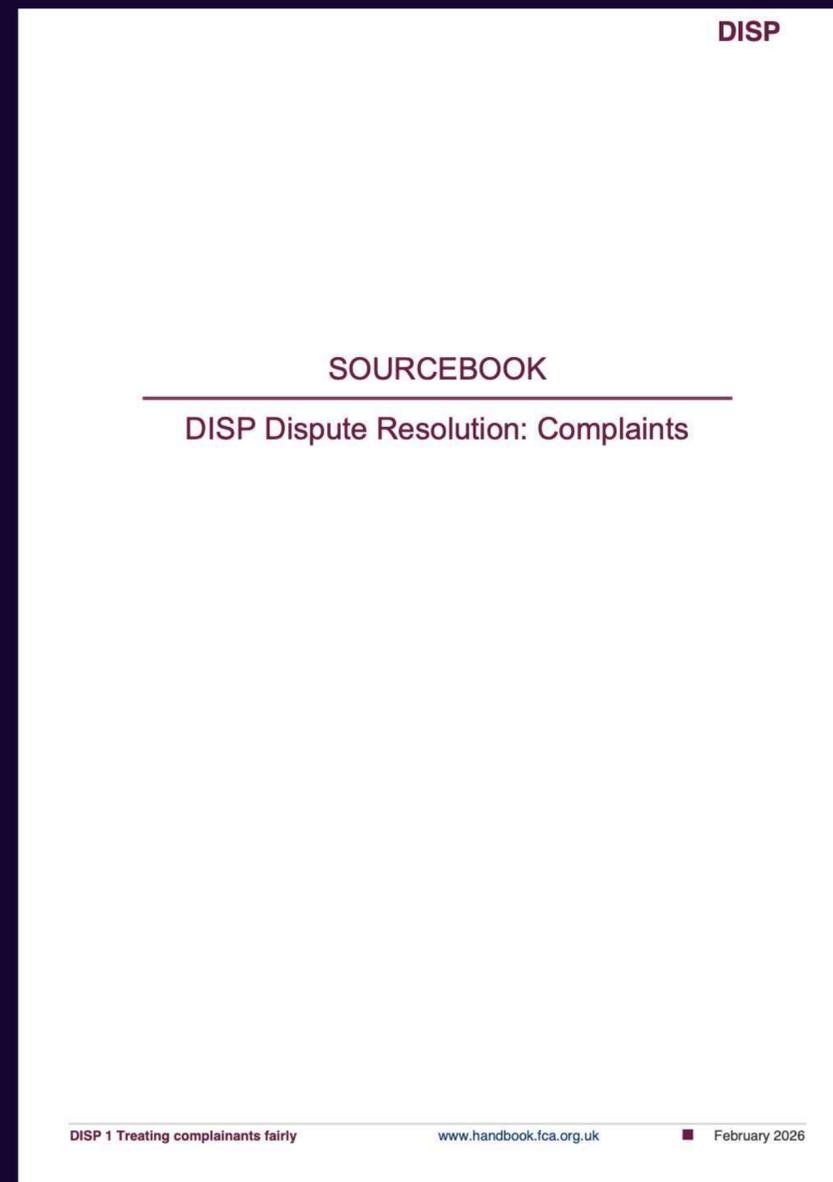
Have the LLM QA the output using explicit scoring criteria

Confidence scoring



A tiered confidence flag (low/medium/high) showing how likely a specific mapping is correct

Let's review this obligation register



Monzo presentation

Next workshop

🖥️ Online masterclass

How to build AI agents for risk & compliance

14 April
11.00AM GMT

Registration
open



Sam Green
Policy & Partnerships Lead
zango



Ashi Bajwa
Regulatory AI Architect
zango



Tim Tyler
Vice President
ICAI INTERNATIONAL COMPLIANCE ASSOCIATION

Register here <https://www.zango.ai/events>

Poll

**Which compliance AI agents should
we build in the next workshop?**

Masterclass Prompts

1. Prompt reset

If you had to complete this task perfectly using a single, well-structured prompt – what would that prompt be?

1. Rewrite my request as that improved, precise, high-quality single prompt.

2. Identify:

- Objective
- Audience
- Inputs assumed
- Scope boundaries
- Output format
- Quality constraints

Source: Zango AI

2. Context injection

You are acting as [ROLE].

You are operating in the context of [INDUSTRY / REGULATED ENVIRONMENT].

The jurisdiction(s) you should assume are [JURISDICTION].

The type of organisation is [ORG TYPE], and the relevant audience is [AUDIENCE].

Your responses must prioritise [PRIMARY RISK / CONSTRAINT], and avoid [KEY FAILURE MODES].

If information is missing or uncertain, you must [HOW TO HANDLE UNCERTAINTY].

Source: Zango Ai

3. Obligation Completeness Audit

Section 1 – Context

Role: You are acting as a senior regulatory compliance auditor reviewing the completeness of an obligation register.

Industry: You are operating in the context of a regulated financial services environment where regulatory accuracy and audit defensibility are critical. Assume the relevant jurisdiction is the United Kingdom

Firm Type & Audience: The organisation type is a Retail bank and the target audience is the second-line compliance function

Priorities & Constraints: Your response must prioritise regulatory completeness and textual fidelity to the source legislation. If information is missing or unclear in the provided legislative extract, you must explicitly state that and not speculate.

Source: Zango AI

Section 2 – Task

For each extracted obligation, compare it to the source and decide whether it passes the rubric. If ANY material detail from the source is **missing, softened, broadened, or merged**, or if any of the criteria below is not met, it cannot be marked COMPLETE.

Source: Zango AI

3. Obligation Completeness Audit

Section 3 – Completeness Criteria

1) Granularity

- If the source lists specific items, the obligation must keep the list (or the same specificity).
- Fail if a list becomes a generic category.

2) Obligated Party

- Preserve the exact actor/role (e.g., “management body”, “payment service provider”).
- Fail if the actor is broadened, swapped, or made vague.

3) Conditions & Thresholds

- Keep every “if/when/unless/where” trigger, threshold, qualifier, timeframe, exception.
- Fail if any trigger/threshold is removed or softened.

4) Procedural Steps

- If the source implies multiple actions or sequencing, the obligation must reflect that.
- Fail if steps are collapsed into a single vague action.

5) Qualifiers/Standards

- Preserve qualifiers that change meaning (e.g., “promptly”, “in writing”, “securely”, “independent”, “at least X”).
- Fail if the standard is weakened or removed.

3. Obligation Completeness Audit

Section 4 — Result

✔ COMPLETE

- Near-literal to the source AND passes every rubric check
- No missing actors, conditions, lists, steps, artefacts, or qualifiers

⚠ NEEDS REFINEMENT

- Clearly based on the source, but missing/softened/generalised any rubric element

✘ INVALID

- Hallucinates, merges unrelated duties, or is incoherent

Source: Zango AI

Section 5 — Output Structure

Return output in this table format:

| Obligation ID | Result | Reason |

Source: Zango AI

4. Policy Match Confidence Scoring

Section 1 – Context

Role: You are acting as a senior regulatory compliance auditor reviewing a regulation gap analysis against a policy.

Industry: You are operating in a regulated financial services environment where regulatory accuracy and audit defensibility are critical. Assume the relevant jurisdiction is the United Kingdom

Firm Type & Audience: The organisation type is a Retail bank and the target audience is the second-line compliance function.

Priorities & Constraints:

- Prioritise full preservation of compliance-material detail from the legislative extract.
- Require explicit evidence of coverage in the policy text; do not assume alignment.
- Where legislative wording is ambiguous, incomplete, or context-dependent, explicitly flag the limitation rather than interpreting intent.

Source: Zango AI

4. Policy Match Confidence Scoring

Section 2 – Task

Assign a confidence score (High / Medium / Low) for how strongly the cited policy section covers the obligation. Explain why, using evidence from both the obligation and the policy text. You are NOT drafting new policy wording. You are judging the mapping quality.

For each proposed mapping, assess:

1. **Coverage:** does the policy address the same requirement as the obligation?
2. **Specificity:** does it include the obligation's key details (actors, conditions, thresholds, steps, artefacts, standards)?
3. **Strength:** is the policy language binding enough (must/shall/required) or is it vague/aspirational?
4. **Scope alignment:** is it the right subject matter and context (not adjacent-but-different)?

Source: Zango AI

4. Policy Match Confidence Scoring

Section 3 — Results

✓ HIGH

- Direct, unambiguous match
- Policy covers the obligation's core action AND most/all material details
- Language is enforceable (must/shall/required), not just guidance

⚠ MEDIUM

- Partial coverage OR correct topic but missing material details (conditions/thresholds/artifacts/steps/actors)
- OR policy language is weaker/ambiguous (e.g., “should”, “where possible”)
- Still plausibly intended to satisfy the obligation, but needs refinement or stronger evidence

✗ LOW

- Weak or indirect link (adjacent topic, general principles, or controls-level detail without the actual requirement)
- OR key elements are mismatched (wrong actor, wrong trigger, different action)
- OR the policy is too generic to evidence compliance

Source: Zango AI

4. Policy Match Confidence Scoring

Section 4 — Task

Return output in this table format:

| Obligation ID | Policy Section (ref) | Confidence (High/Medium/Low) | Reason (evidence-based) |

Source: Zango AI

zango

Thank you

Reach us at

hello@zango.ai

www.zango.ai

[LinkedIn](#)