20 YEARS ANNIVERSARY

DSLs and AI Deep Dive



Sandip Pandey
Senior AI Engineer

be informed

Agenda

How we get reliable structured models of legislation out of an LLM?

1 Why "ask for JSON" is not enough

2 What a DSL adds

3 Three DSLs, one discipline

4 Three DSLs, one discipline

4 Inside an AI Extraction Step

5 The validation stack

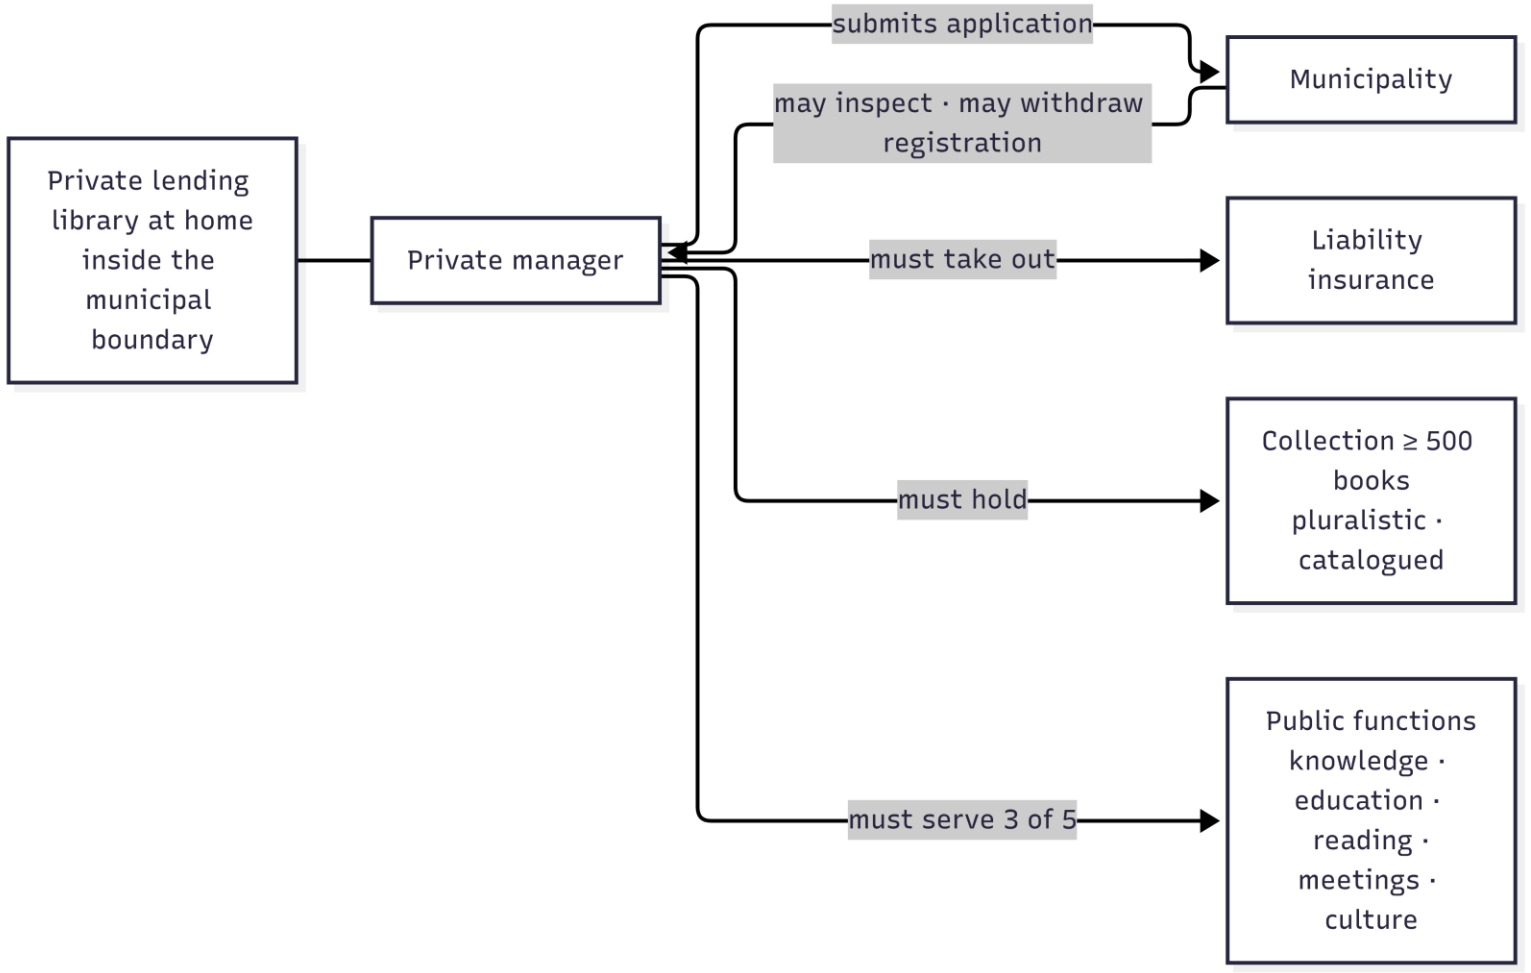
6 Where the reliability comes from?

8 Takeaways

© Be Informed 2026 | All rights reserved worldwide

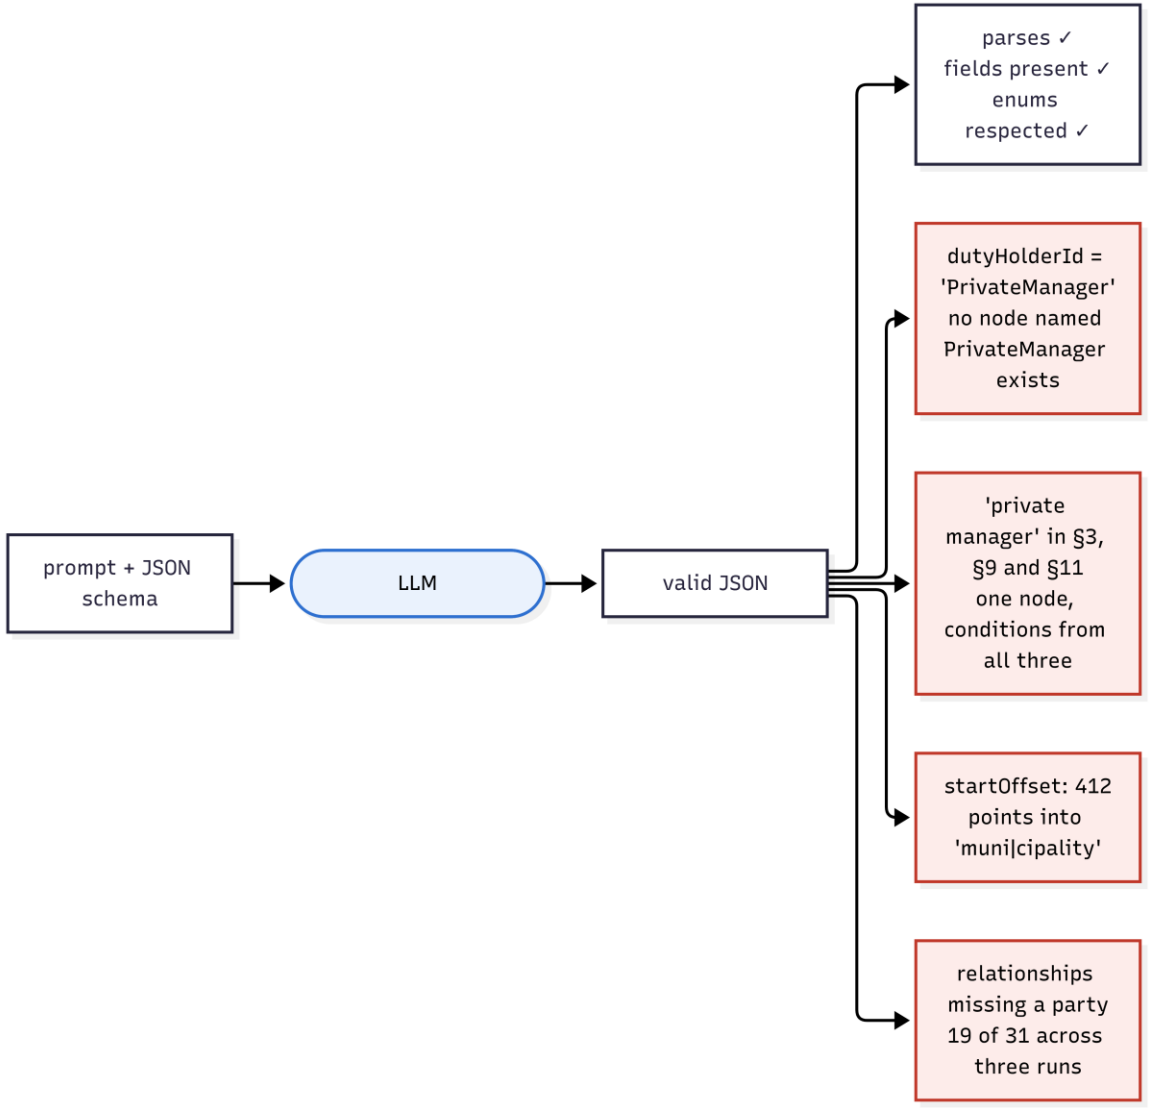
Example for this session

Private lending library at home



© Be Informed 2024 | All rights reserved worldwide

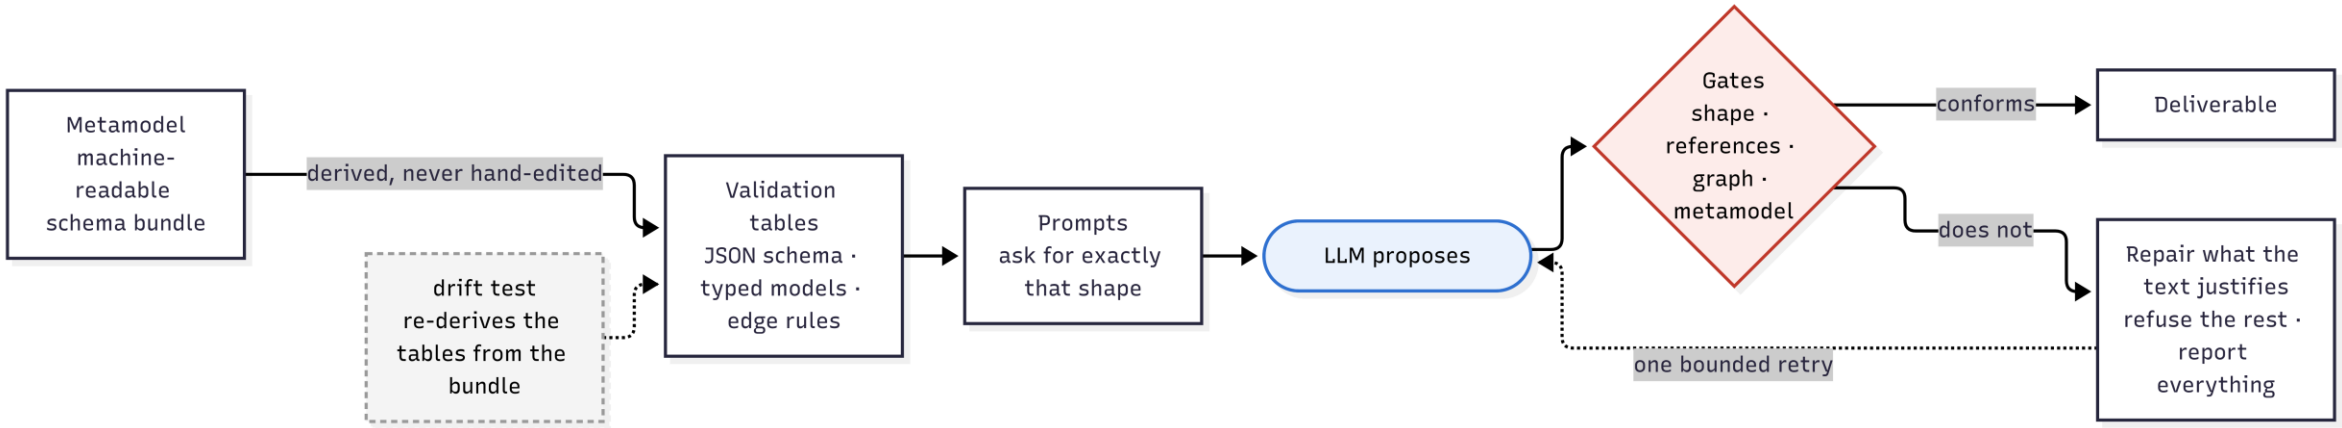
Why "ask for JSON" is not enough



be informed

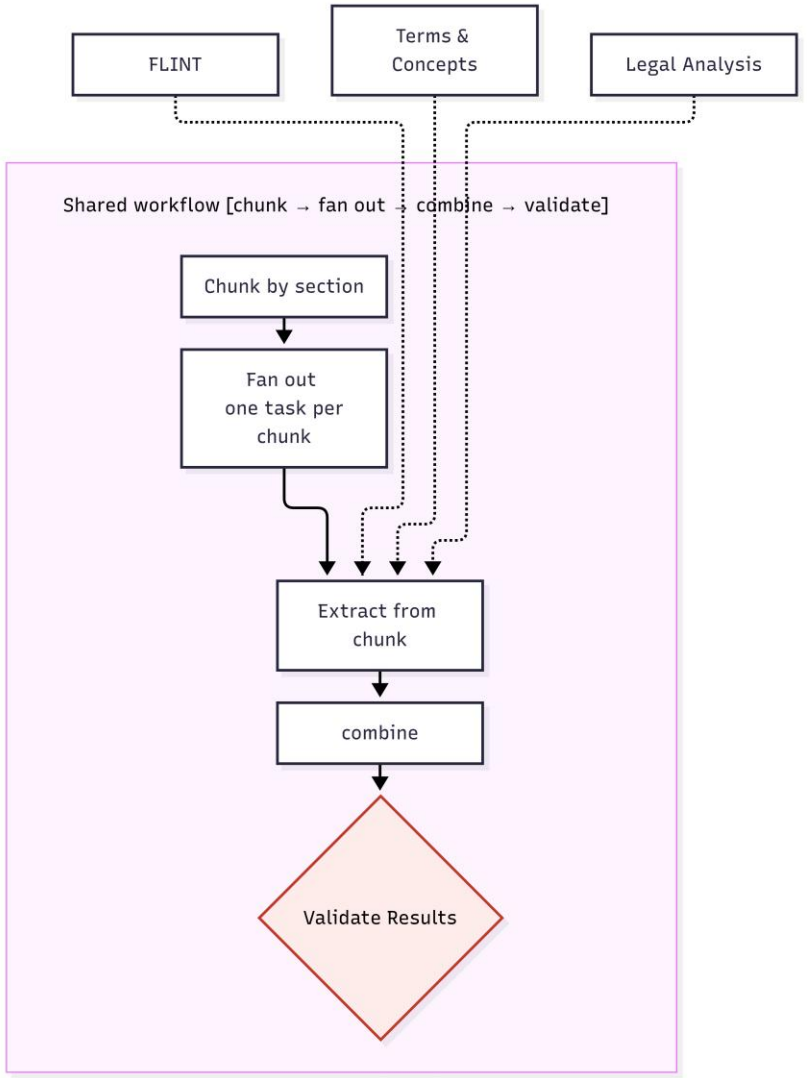
© Be Informed 2026 | All rights reserved worldwide

What a DSL adds



© Be Informed 2026 | All rights reserved worldwide

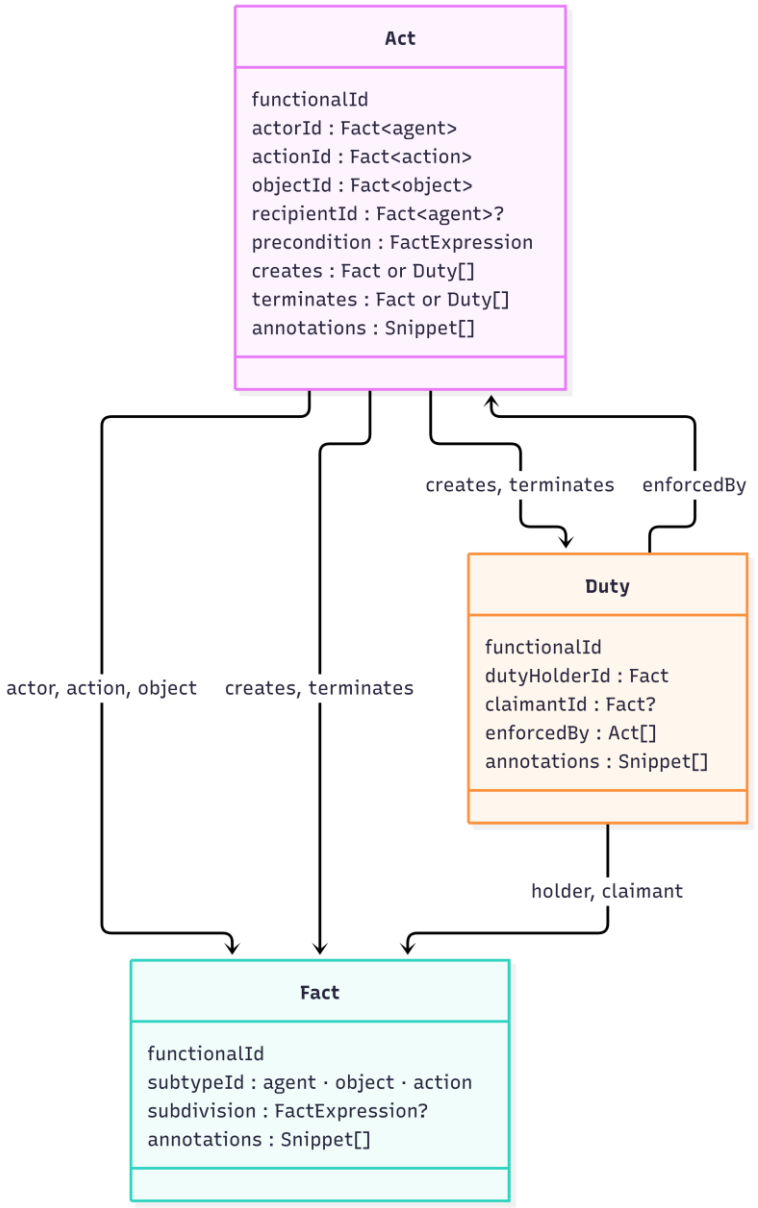
Three DSLs, one discipline



be informed

© Be Informed 2026 | All rights reserved worldwide

FLINT

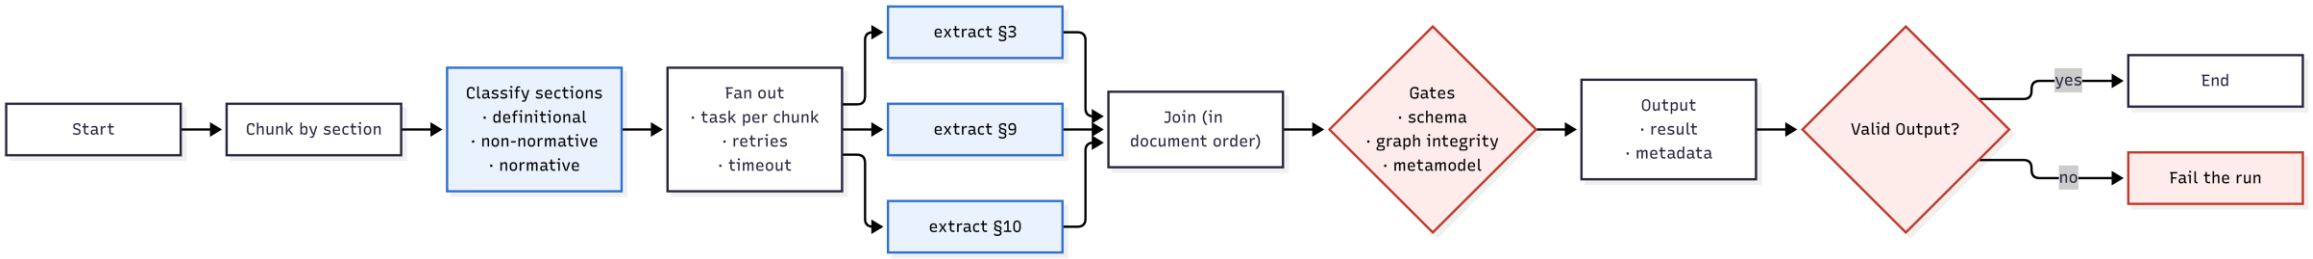


be informed

© Be Informed 2026 | All rights reserved worldwide

AI Extraction Pipeline

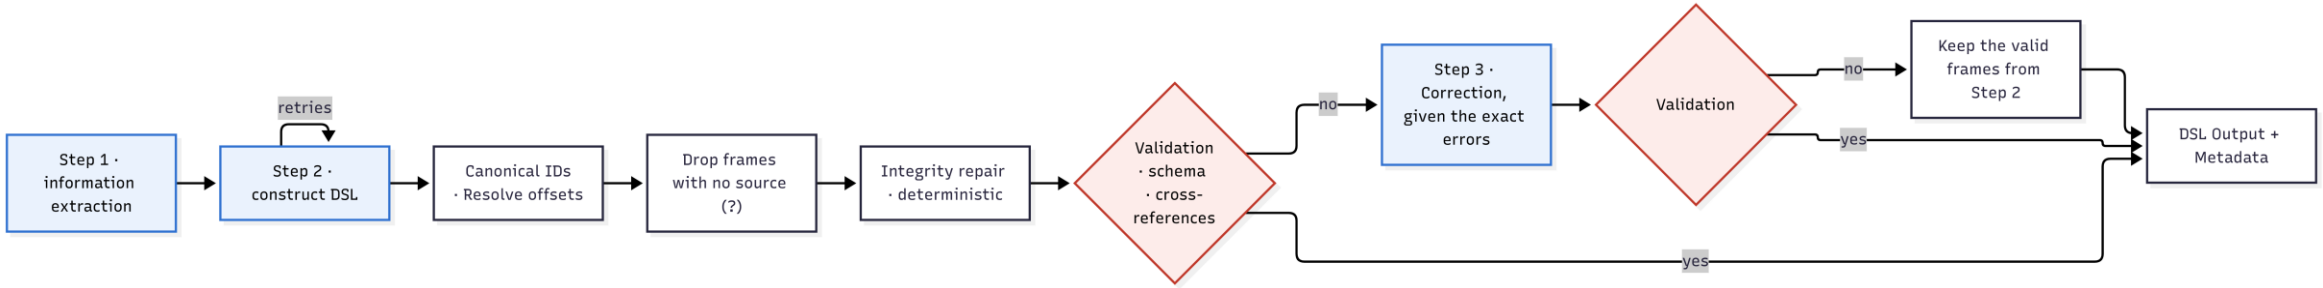
General framework for DSL extraction



© Be Informed 2024 | All rights reserved worldwide

Inside an AI Extraction Step

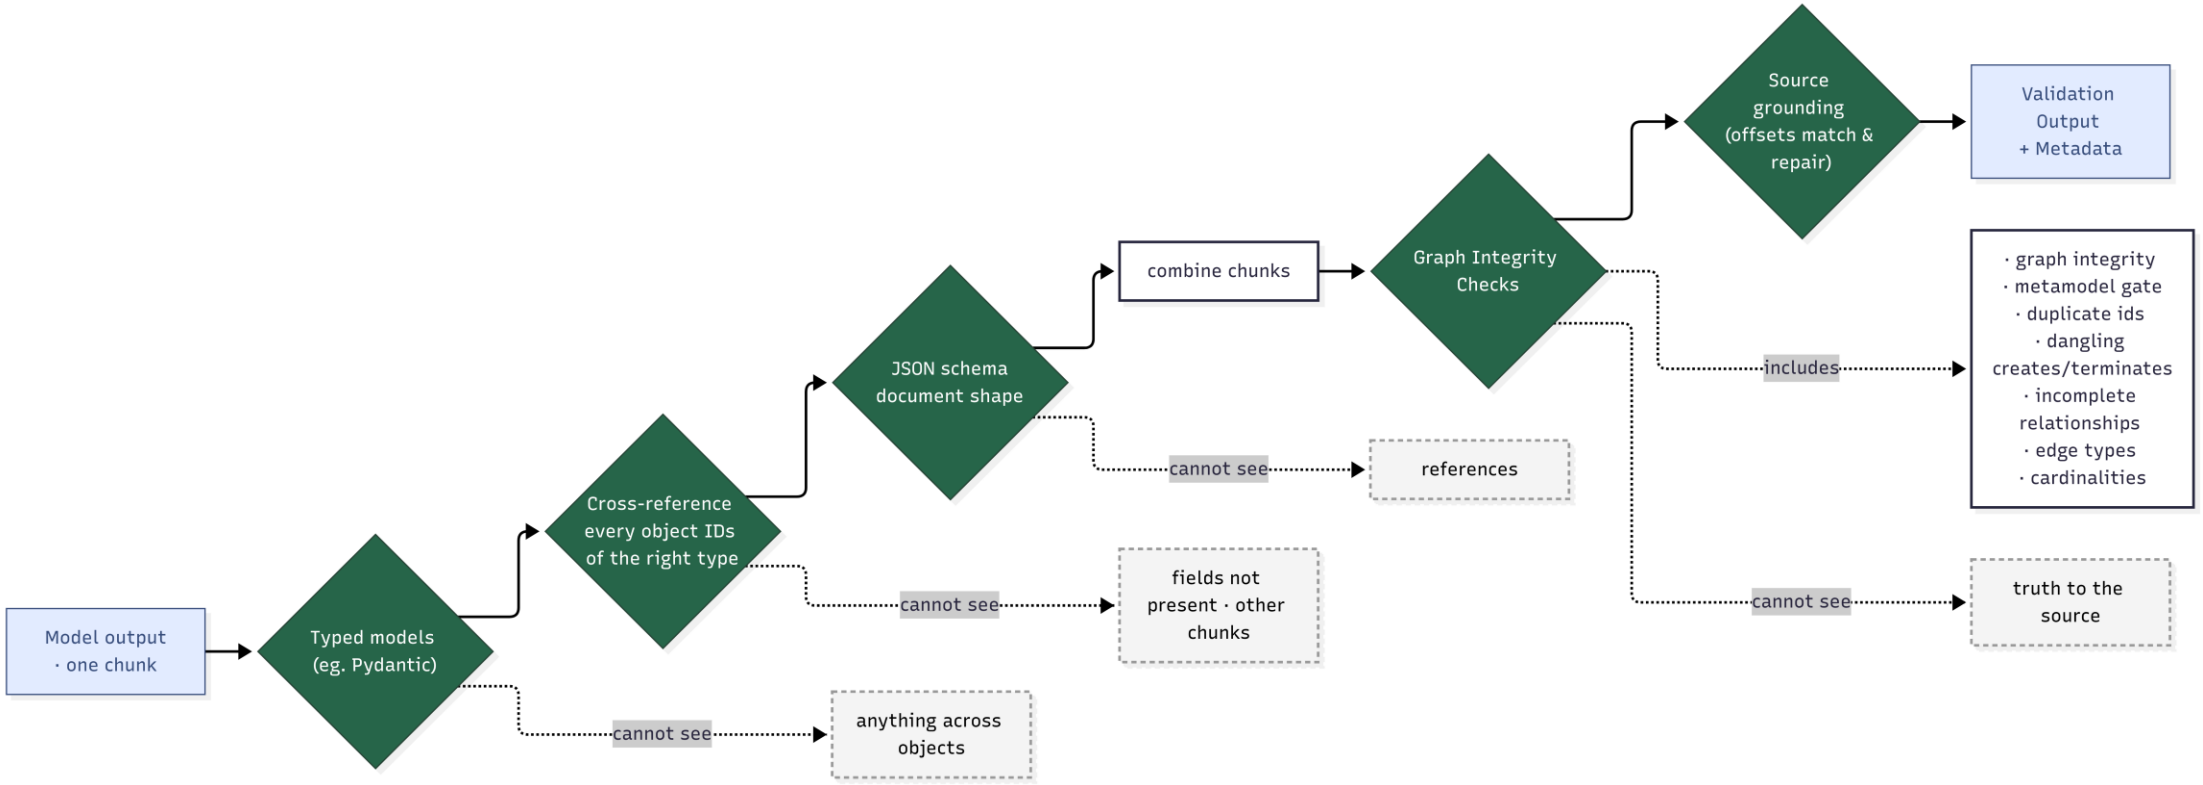
It is more than a call to an LLM



© Be Informed 2024 | All rights reserved worldwide

Validation Stack

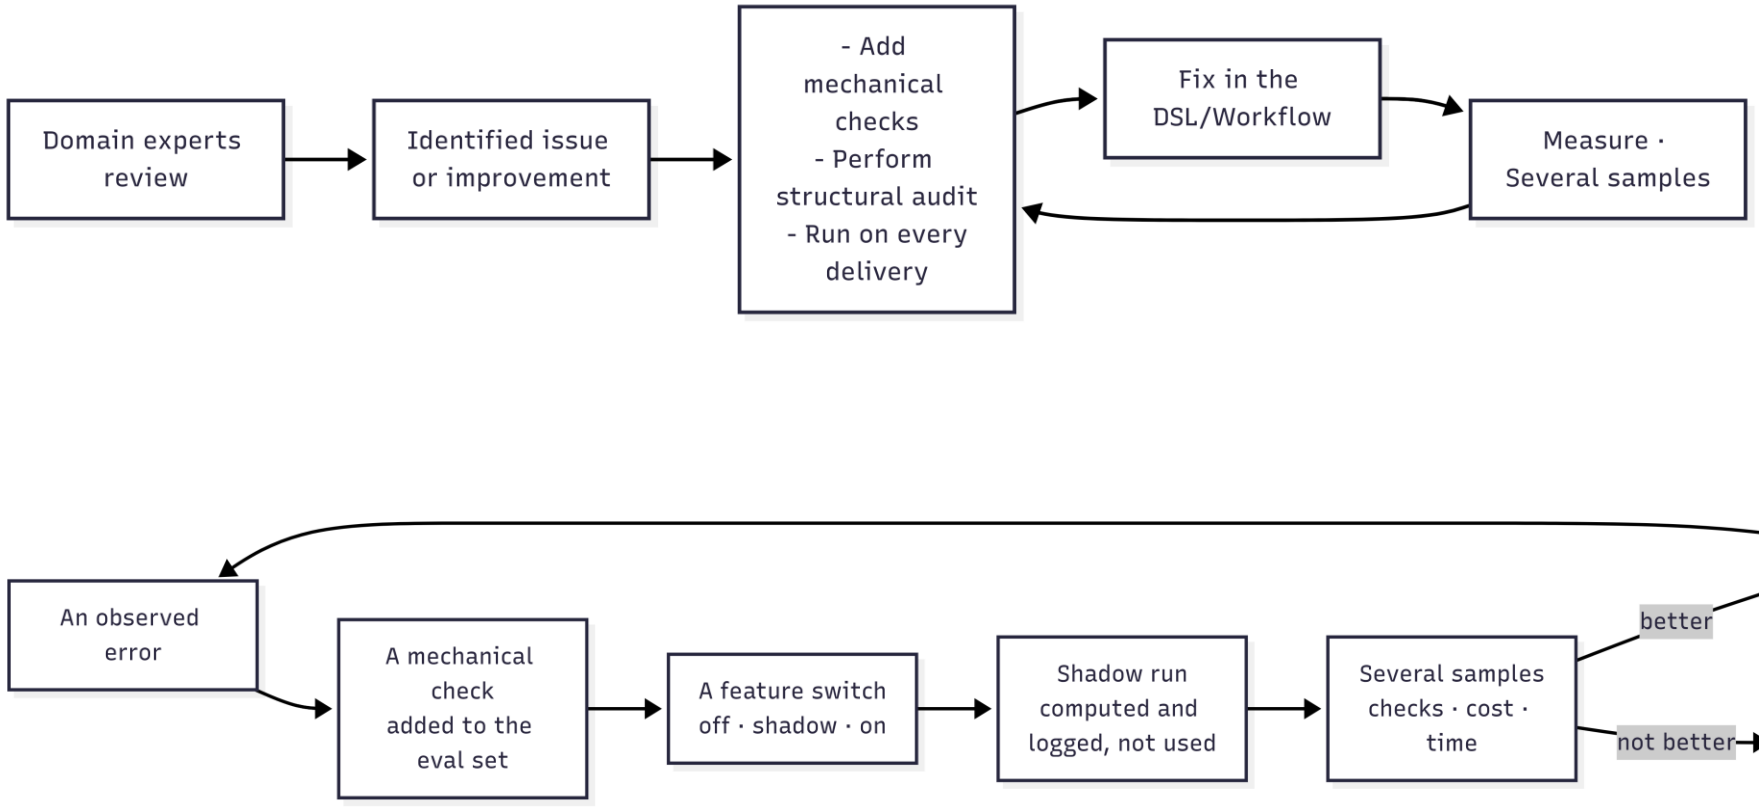
Can we trust the output?



© Be Informed 2024 | All rights reserved worldwide

Where reliability comes from

The continuous improvement loop



© Be Informed 2024 | All rights reserved worldwide

Some DSL Switch Examples

Switch	Default	Changes
Section roles	on	which sections are extracted at all
Implicit parties	on	unstated parties filled from the assigning provision
Terms layer	on	Terms model beside the analysis, links and subclasses
Chunk context	per DSL	what a chunk reads but does not extract from
Decision model	off	typed yes/no and choice questions with probabilities beside the text model

© Be Informed 2026 | All rights reserved worldwide

Takeaways

For your own extraction systems



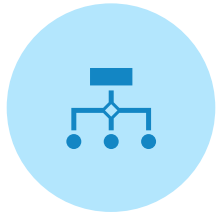
Have the model mark up the input, not generate the output. Parse the markup; offsets become free and exact.



Validate against a schema you did not write by hand. Derive gates from the consumer's metamodel and test for drift.



Scope identity in the DSL. Merging on names is the error experts notice first.



The model proposes, code decides. Repair only what the text justifies; refuse the rest; report everything.



Shadow before you switch. Several samples per change; single runs are noise.



Treat the LLM as a dependency. Config, metering, caching, cost estimate; swap and re-run the evals.

be informed

© Be Informed 2024 | All rights reserved worldwide

Thank you!
Any questions?

info@beinformed.com
+31 (0)55 368 1420
www.beinformed.com

be informed