Five checks to run on any AI prosecution tool

A demo shows a tool at its best. These checks show what it does on your own work, with no technical setup.

> **How to run them.** Pick a recent office action your team has already answered. Run it through each tool you are evaluating, and have the associate who wrote the actual response score what comes back.

1 Quotes

Pick three claim limitations from the tool's chart at random and compare them with the pending claims. They should match word for word. A paraphrased limitation in an analysis is a paraphrase waiting to reach the prosecution history.

☐ Pass ☐ Fail Notes

2 Coverage

Every rejection in the office action, and every claim it rejects, should appear in the analysis. Count them against the office action, not against the tool's own summary.

☐ Pass ☐ Fail Notes

3 The examiner's art

Each argument should answer the art the examiner actually asserted against that claim, not a reference cited somewhere else in the action.

☐ Pass ☐ Fail Notes

4 A missing document

Remove one document, a cited reference for example, and run it again. The tool should ask for it, not proceed on assumptions about what it says.

☐ Pass ☐ Fail Notes

5 Time

Measure how long the associate spends getting the draft to a usable state. That number, not the demo, is what the tool is worth on a flat-fee matter.

Minutes to a usable draft

These are the checks we would run on any product in this category, including ours.