

Thomson Reuters doesn't just deploy AI.

We build it.



The story of **Thomson**, our own foundation model,
told in five stops.



1. The problem

General-purpose AI was built for work where being almost right is good enough.

#1 Ranking on the
hardest legal
benchmark tested — ahead
of every model evaluated.



2. The approach

We didn't rent intelligence, we built it: an open-source foundation trained on decades of our own proprietary legal, tax, and regulatory content.

100+ Legal-domain
benchmarks,
plus thousands of real legal-
specific test cases, used to
evaluate the model.



3. The proof

How it performs, backed by real benchmark data.

<10% Share of
Thomson
Reuters® proprietary content
used in training so far — the
rest is still ahead.



4. Still early

The model has only used a fraction of our available content so far, which reframes the proof point as a beginning, not a ceiling.

10,000 / 100
Documents/questions Tabular
Analysis in CoCounsel Legal
can process at once, today.



5. Where it lives today

It's already working inside a real product, not just a research demo.

To learn more about Thomson, visit
thomsonreuters.com/en/thomson-llm



Read the full technical report, with our experimental findings and details of the methods, data, and evaluations from training the Thomson model.