Layout-driven inspection recipe generation (RAC)



Bump recipe creation fell from **4 h to 20 min**. Large-area layout conversion fell from **314 to 42 s**.

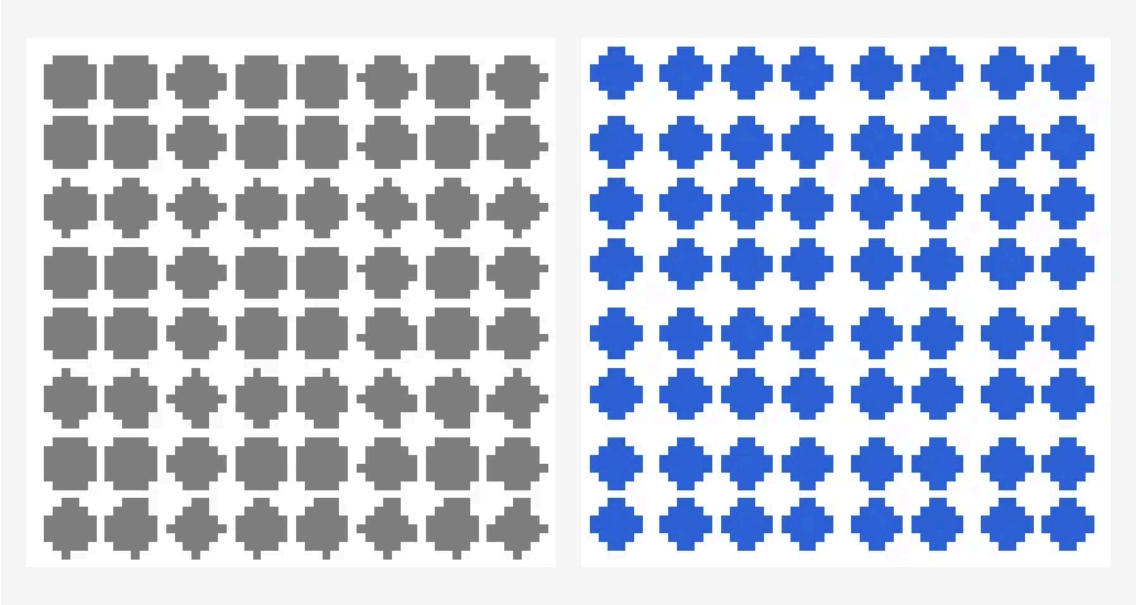

Left: naive conversion. Right: coordinate snapping + oversampling (synthetic pattern)

Problem	Each recipe took hours of manual teaching. Converting drawings to images distorted fine circular patterns and shifted their positions.
Implementation	Replaced min pooling with coordinate snapping, oversampling, box resampling, and thresholding (figure above). Verified against reference images with VIAs and bumps of several sizes; on production material (fine VIAs and bumps) position and shape matched, and it passed production validation.
Evaluation	Placed about 250,000 bumps with 4-point perspective correction, and later layers inherit the first layer's correction.
Status	Conversion passed production-material tests. The cross-site tool awaits deployment.

Evidence · internal evaluation records (H1 2026)

Training and validation for Shot Key detection

Rules and fewer scan lines cut field alignment from **140 to 23 s**, with marks in **12/12** slots. CNN false detections fell from **280 to 114** on 6,150 in-house images including augmented examples.

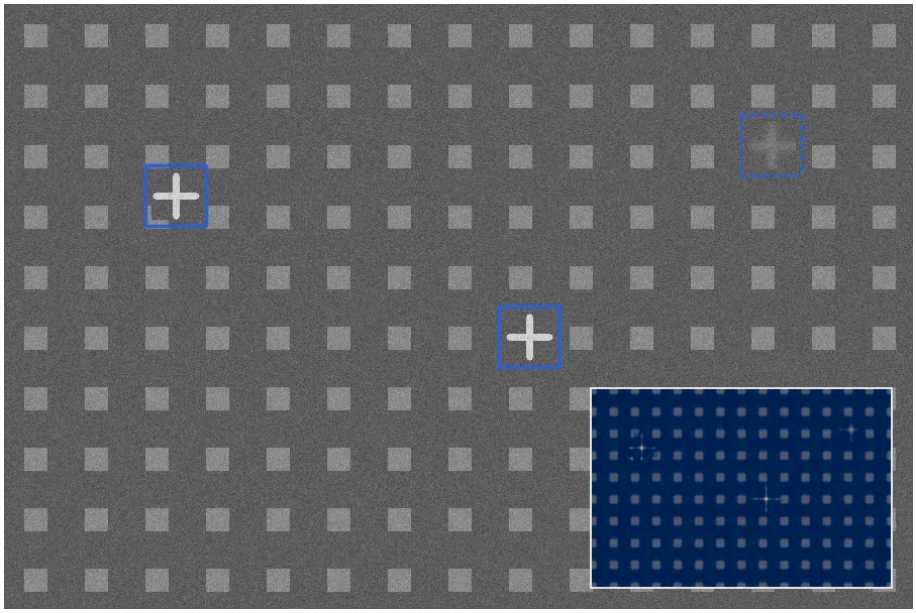

Rule detections (solid) and a faint mark missed by score (dashed); bottom right is the correlation map. Synthetic image

Problem	Template-based align was slow and needed marks re-registered at every new site. Rule parameters did not fit every image, causing false detections.
Implementation	Found same-day leakage and introduced a fixed date-based evaluation split and regression checks. Built a tool that learns the rule-detection tolerances people used to tune from labeled field images, and exports them as the profile file the tool reads.
Evaluation	Mark detections on 1,030 unaugmented images rose from 996 to 998. This difference was not statistically significant.
Status	Rules passed field tests. The CNN runs in shadow alongside the existing decision.

Evidence · internal evaluation records (H1 2026), training tool documentation

Weekly engineering reports with a local LLM

In a two-week synthetic replay of the public app, inputs through the export request fell from **25 to 16**.

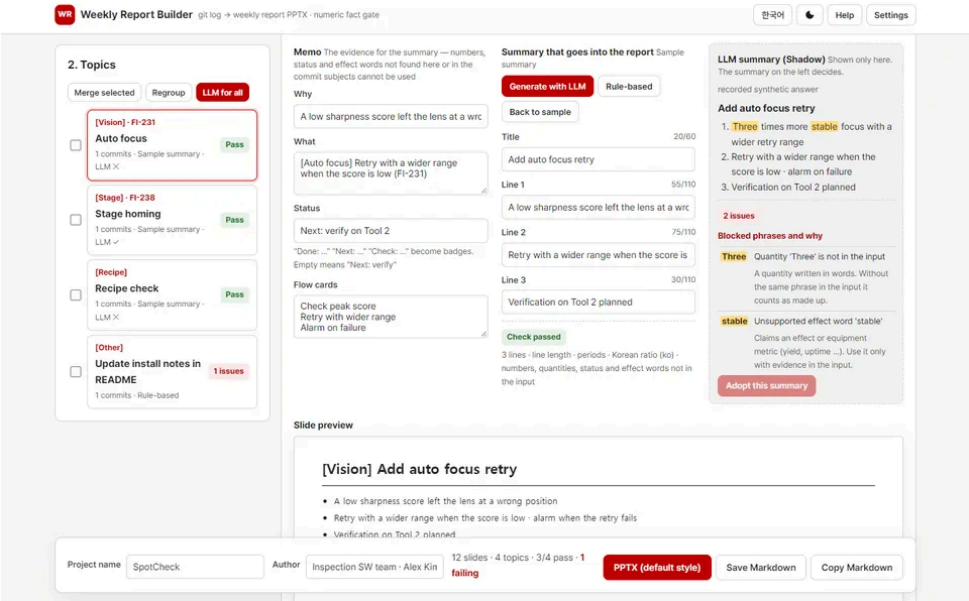

The public app. It rebuilds the in-house tool's flow so anyone can run it on their own git log.

Problem	Every week, writing the report meant rereading the week's commits and notes and copying them by hand into the company PPT template.
Implementation	Rules check line count, length and unsupported numbers or status claims such as "done." A failed generated draft is retried once before falling back to rule-based text.
Evaluation	The available records do not track drafting or review time, the share of AI draft sentences edited, or repeat usage by author.
Status	The internal tool had 10 team users. One weekly run produced a report from 5 manually entered topics.

Evidence · November 2026 team usage, internal report build record and public app benchmarks