

Pilot Initiative Finds AI Can Speed Up a Key Step in Bay Area Sea Level Rise Planning

A San Rafael test shows how off-the-shelf AI tools can help turn stacks of local plans into standardized, citation-backed summaries.

As one of the 51 Bay Area jurisdictions required by California Senate Bill (SB) 272 to prepare a Regional Shoreline Adaptation Plan (RSAP), the City of San Rafael is already well ahead of the curve in its sea level rise preparedness efforts. But like many local governments, San Rafael faces a challenge familiar to planning staff everywhere: balancing limited time and resources against a growing list of planning requirements. Seeking opportunities to streamline some of the more administrative aspects of the RSAP process and allow staff to focus on higher-value planning work, the city partnered with the San Francisco Bay Conservation and Development Commission (BCDC) and the Metropolitan Transportation Commission's Data Strategy & Analytics (DSA) team to explore whether artificial intelligence could help.

The resulting pilot focused on one requirement in BCDC's Regional Shoreline Adaptation Plan framework: identifying and summarizing existing plans, studies, and reports related to coastal flooding and sea level rise. Known as Element B1, this task requires jurisdictions to review a wide range of documents, including general and land use plans, hazard and emergency plans, and climate resilience strategies. For a single city, that can mean working through thousands of pages spread across dozens of documents—a necessary but time-intensive effort.

Rather than replacing planners' expertise, the pilot explored whether AI could accelerate the initial document review process by producing organized, citation-backed summaries for staff to review, validate, and finalize. The effort demonstrated how AI can help reduce time spent on repetitive document analysis while maintaining planners' central role in interpreting results and making decisions.

Inside the San Rafael Pilot

Working with planning staff in San Rafael, the project team — University of San Francisco data science practicum students Jackson Dawson and Patrick Crouch — built a simple, repeatable workflow around widely available AI chatbots. They designed the workflow with AI quirks and limitations in mind, applying a handful of established techniques: chunking the job into smaller, sequential passes; writing strict, reusable prompts that fix the AI's role, rules, and exact output format; and requiring every claim to be grounded in the source material with a page citation.

Every design decision makes an impact on the result. Chunking keeps individual steps small enough for the AI to handle accurately. Strict prompts make results consistent from one run — and one user — to the next. Grounded citations keep the output tied to its source and easy to audit. Together they

turn a stack of documents into two reliable deliverables: plain-language summaries for each category, and a structured scorecard the team calls a "top sheet."

Importantly, the design keeps people in charge. Because the summaries will eventually appear in official government plans, the workflow requires a human review step, guided by a scoring rubric the team built to catch common AI mistakes — invented citations, overstated claims, or the wrong tone.

The team ran the full workflow from start to finish for San Rafael across all three categories, using the city's real planning documents. The result was a set of well-tuned summaries, ready for final review.

Just as important, the project includes a complete set of documentation so others can run the same process: a design document explaining the choices behind the workflow, a step-by-step user guide, the prompts, the evaluation rubric, and full example outputs for every run. Along the way, the team also tested its assumptions — comparing different AI tools and two ways of splitting up the work — and settled on the approach that produced the clearest, most reviewable results.

Next Steps

The pilot is a proof of concept, and it points to three clear next steps. The first is consistency: today the workflow relies on whatever version of a commercial AI tool is current, and those models change without notice. The team recommends moving to a standardized model so the same inputs reliably produce the same results. The second is automation: turning the manual, multi-step process into a single program where a jurisdiction loads its documents and the system handles the rest. The third is reach: running the workflow in more cities, each with different document types, to refine the prompts and prove it works beyond San Rafael.

The approach is not limited to sea level rise. Any task that means reading many long documents and distilling them into a consistent, reviewable format — across planning, compliance, or grant work — could use the same pattern: let AI handle the first draft and the indexing, and keep trained staff responsible for the final word. As more public agencies look for practical ways to use AI, the San Rafael pilot offers a grounded example of how to do it carefully — saving time on the busywork without handing over the judgment.