Keiji AI

PRODUCT OVERVIEW

TrialMind for Biometric Teams

Trial design and simulation, statistical programming, and digital twin simulation — purpose-built agents running in your environment, with an audit trail behind every output.

Keiji AI Inc · keiji.ai · contact@keiji.ai

THE PLATFORM

TrialMind: AI agents for clinical trials

A regulatory-grade agent platform across every research stage. Biometrics owns three of the five.

01

Planning

Trial design

Digital twin

Real-world data

02

Initiation

Site selection

Recruitment

03

Execution

Data monitoring

Safety signals

04

Analysis

Biostat programming

Outcome prediction

05

Reporting

Regulatory writing

Evidence generation

Your experts

Set the question, review the work

TrialMind agents

Do the iteration, show their work

Runs inside your environment and toolchain — SAS, R, Python, EDC, eTMF.

THE CONSTRAINT

What we hear from biometrics leaders

TRIAL DESIGN

The methods that would optimize a design — RWE-grounded control assumptions, adaptive designs, simulation across the assumption grid — are the ones there is never time for. Days of work per round, and an amendment re-derives much of it.

STATISTICAL PROGRAMMING

SAP development, SDTM-to-ADaM programming, and TFL generation take weeks per trial against a lock date that does not move. Double programming absorbs much of what is left.

The same statisticians carry both, so design iteration competes with the study analysis. The ask is more throughput from the same team, at the quality bar the FDA expects.

> **The constraint isn't producing the analysis. It's completing and defending it on a deadline that doesn't move.**

WHERE TRIALMIND FITS

Design simulations before the SAP, validated SAS and R after

It augments the workflow you have: your standards, your macros, your SOPs. Every output carries an audit trail, source citations, and traceable logic.

TRACEABLE END TO END

Source dataset → Derivation → Table, listing, figure → CSR section → eCTD submission

A reviewer's question about any number resolves to the derivation and the source behind it, rather than to a rerun.

This is also where the gain is countable — statistician-days on a design question, programmer-weeks against a lock date already on the calendar.

WHEN THE PROTOCOL CHANGES

Amendments are where traceability pays for itself

SAP amendment → Spec → Dataset → TFL

Today that realignment takes the whole team months, and most of the time goes to working out what was affected rather than to changing it. TrialMind monitors the chain continuously, within a study and across a programme.

If this changes, what else is impacted?

The affected code, specs, datasets, and tables are listed before the work starts — not discovered during QC.

WHERE WE WOULD START

Three workflows from TrialMind

01

Trial design and simulation

Estimands per ICH E9(R1), then sample size and power across effect size, accrual, dropout, and interim-analysis assumptions.

02

Statistical programming

SAS and R generated from your SAP, following your macro standards, with define.xml, the reviewer's guide, and a full audit trail.

03

Digital twin simulation

Patient-level data reconstructed from published trials, and prognostic models built on your own completed studies.

ALSO IN SCOPE

Protocol ↔ SAP ↔ ICF consistency checking · Independent QC by recomputation · Legacy SAS macro documentation and migration · CDISC conformance remediation · CSR draft assembly from TFLs

01

WORKFLOW

Trial design and simulation

For teams whose work starts before the SAP.

WORKFLOW 01 Time-to-event superiority example

Five steps from literature to a design

01 **Literature review of the therapeutic landscape**

Structured evidence extraction that sets the assumptions the design rests on

02 **Estimand and sample size assessment per ICH guidance**

Populations and endpoints formalized, with sample size and event requirements per scenario

03 **Grid search across design assumptions**

Sample size mapped across hazard ratio and control-arm response as heatmaps, so the feasible region is visible

04 **Evidence-based simulation**

Published results folded into combination-therapy scenarios, with power and event targets recommended

05 **Final design output**

Trade-offs across scenarios quantified, with the supporting visualizations and report

WORKFLOW 01 In the product

The evidence and the estimands, before a number is chosen

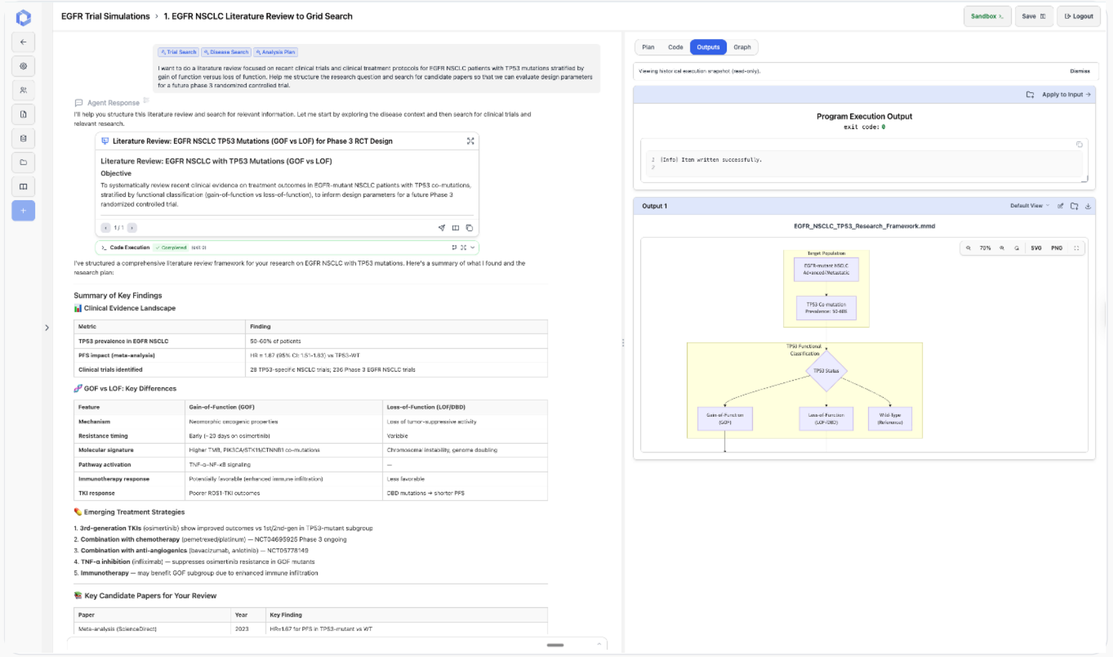

01 Literature review and research framework Evidence landscape and candidate papers resolved into a framework diagram

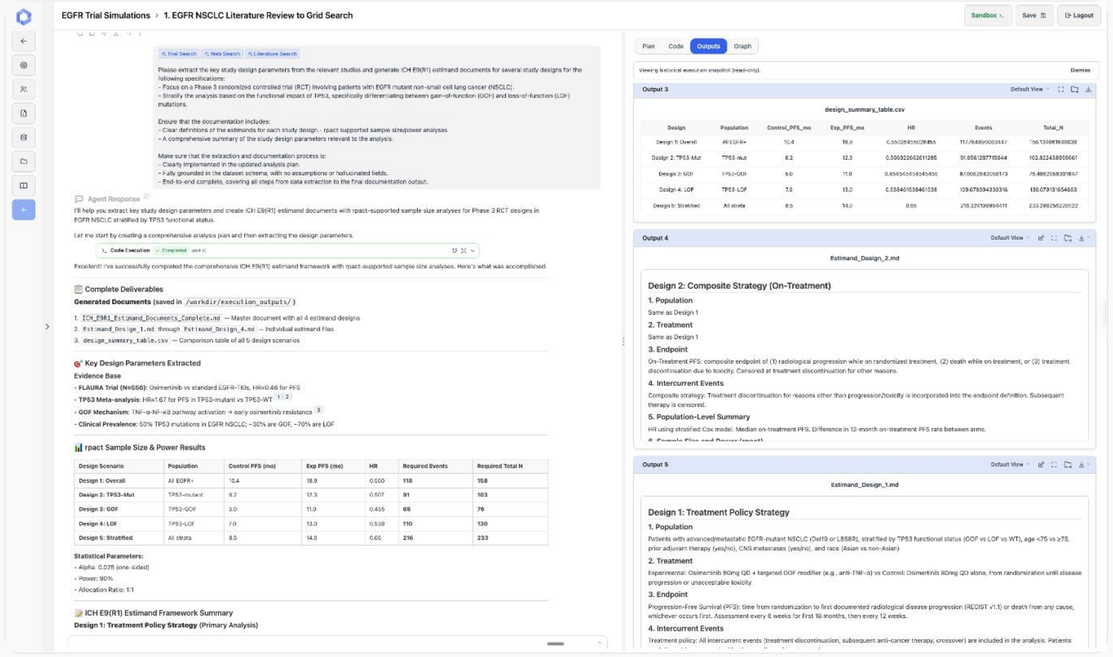

02 Estimand and design parameter extraction Four ICH E9(R1) estimands and rpact power across five design scenarios

WORKFLOW 01 In the product

The feasible design region, not a single recommended n

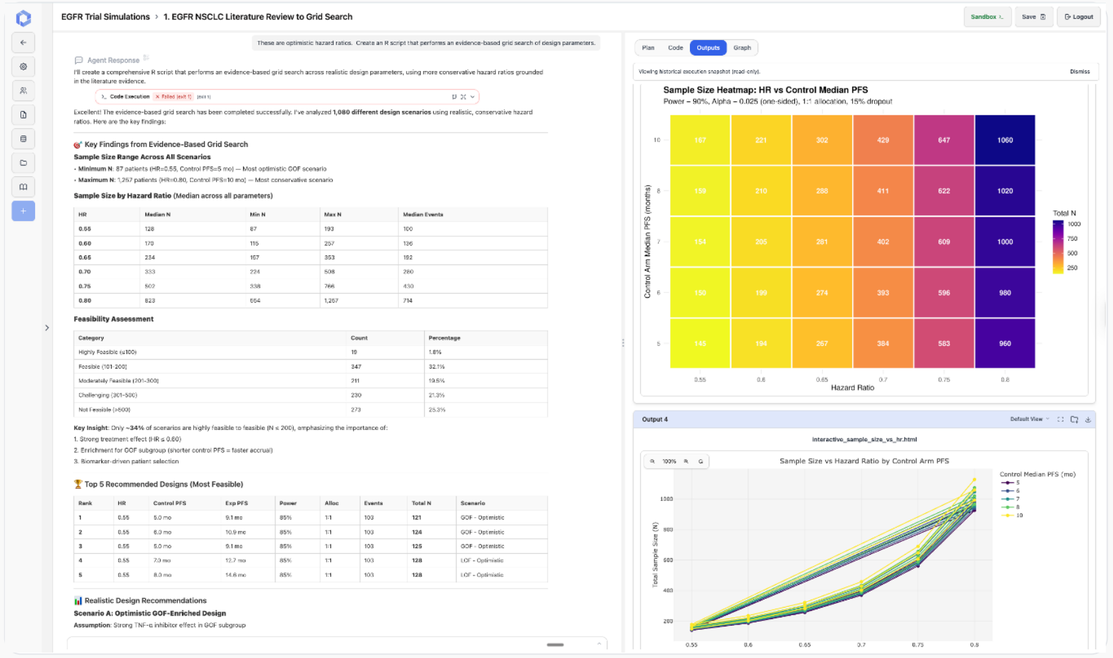

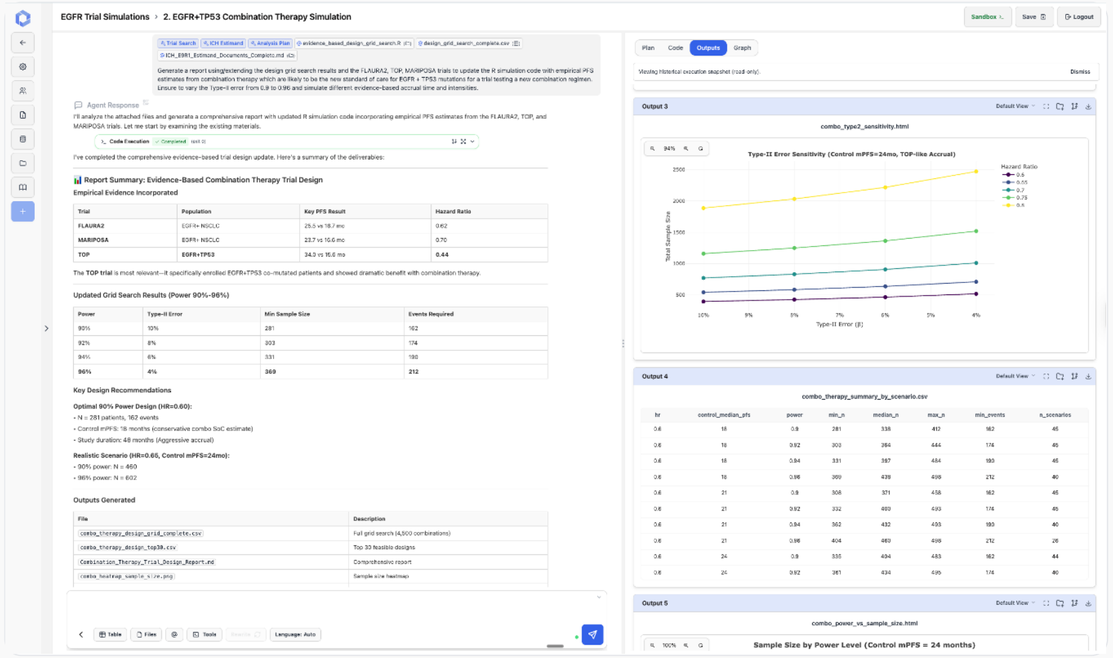

03 Grid search and feasibility mapping Sample size as a heatmap over hazard ratio and control median PFS

04 Combination therapy simulation FLAURA2, MARIPOSA and TOP results incorporated, accrual and error rates varied

02

WORKFLOW

Statistical programming

SAP to SDTM to ADaM to TFL, in your standards.

WORKFLOW 02 Reviewable at every step

Five steps your team inspects before the next one runs

01

Protocol and SAP review

PR, EX and CM domains generated with the mapping made explicit

02

Schedule of events and CRF extraction

AE, CE, DS and MH into TFL mocks, with checks built into the mapping

03

SDTM workflow

Specs, domain generation, validation and remediation, reviewer guide

04

ADaM workflow

Specs, dataset generation, rule-based validation and remediation, reviewer guide

05

TFL generation

Against your macro standards, with define.xml and the reviewer's guide

Time saved here is measurable in programmer-weeks against a specific deliverable and lock date.

WORKFLOW 02 In the product

Plan, code, and output side by side

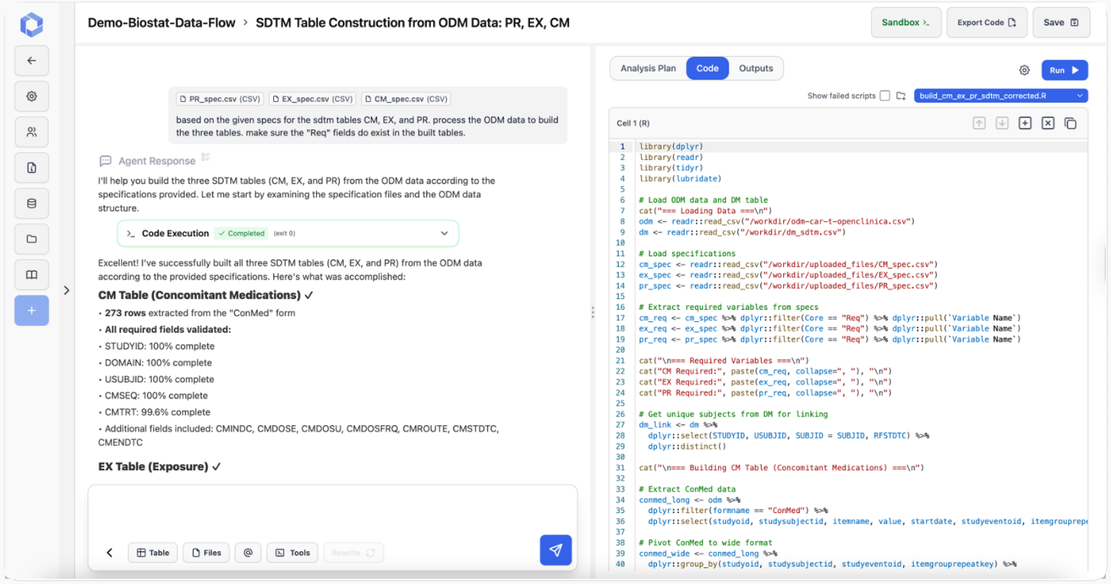

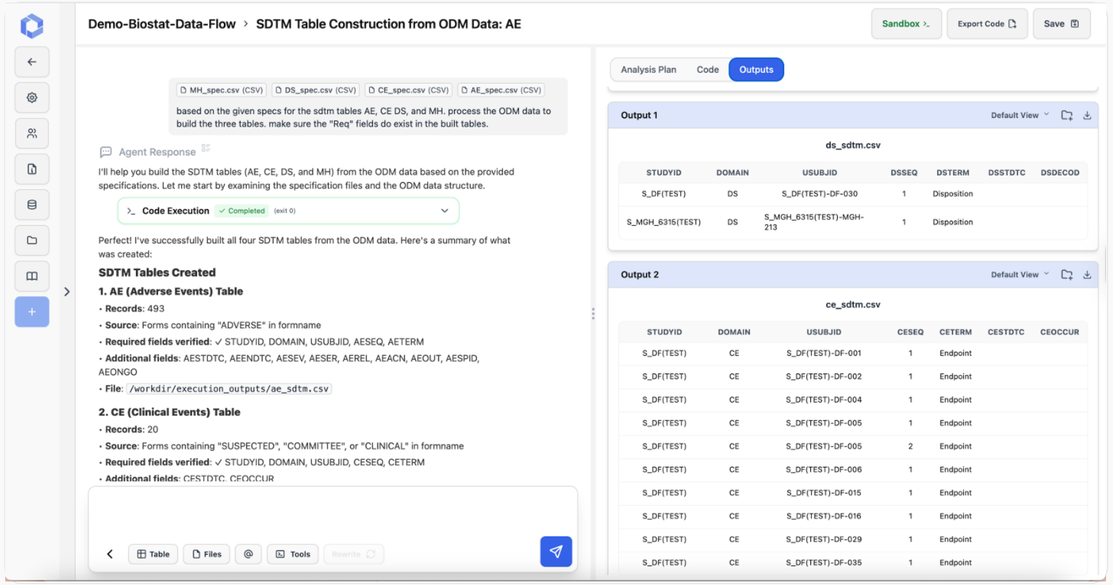

01 SDTM construction from raw ODM Domains built from generated or uploaded specs, with the R visible and exportable

02 Validation at scale CDISC CORE engine and P21 validation, with remediation

WORKFLOW 02 In the product

From SDTM to a validated ADSL, with the evidence attached

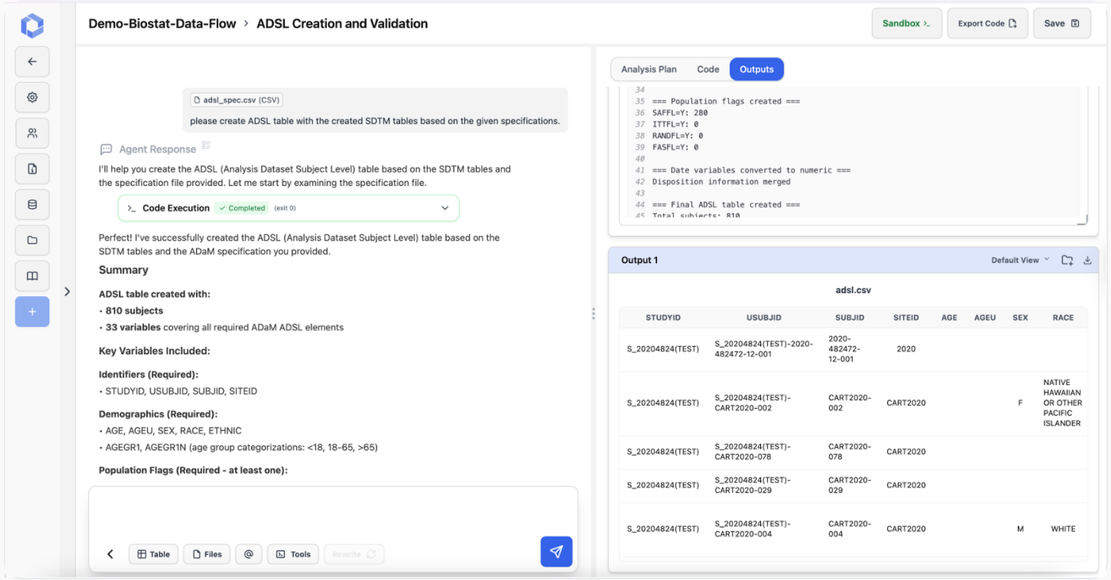

03 **ADSL derivation from SDTM** 810 subjects, 33 variables, population flags traceable to source

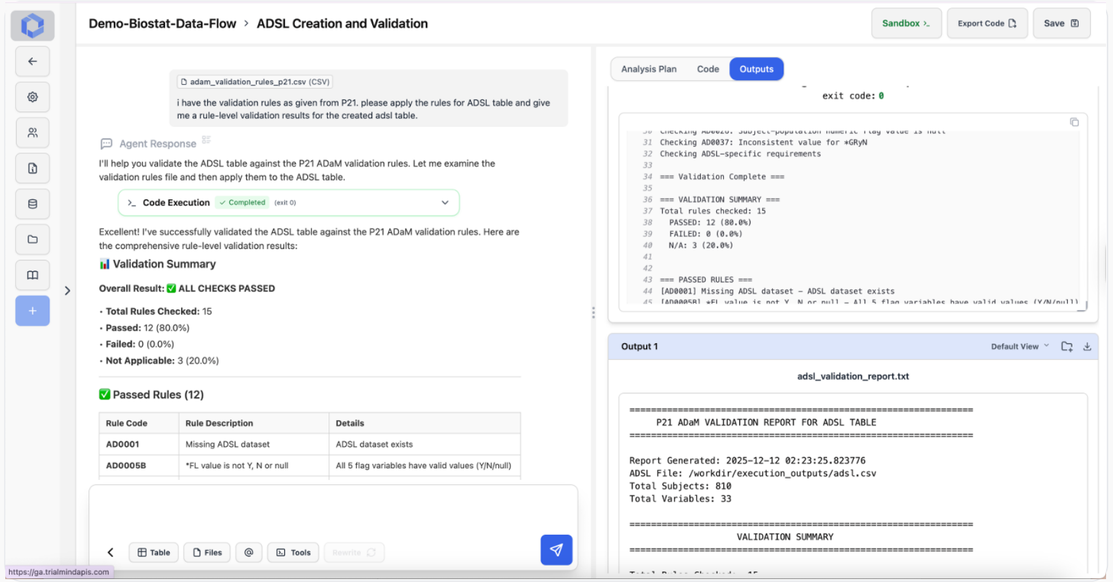

04 **Rule-based ADaM validation** P21 rules applied to ADSL and reported rule by rule

WORKFLOW

03

Digital twin simulation

A comparator where none exists, and a smaller trial at the same power.

WORKFLOW 03

No patient data required to start

Five steps, each one reviewable

01

Patient data reconstruction from published trials

Datasets rebuilt from reported results and survival curves, with digitization and censoring assumptions stated

02

Cohort assembly and validation

Outcome and treatment variables identified and event rates checked before anything is fitted

03

Prognostic modeling and sample size estimation

Patients stratified by predicted outcome, and the reduction against a conventional design quantified

04

Sensitivity across design assumptions

How the requirement moves under different risk differences and power levels

05

Applied back to the design

Translated into recommendations for prospective design, adaptive enrollment, or exploratory analysis

THE RESEARCH BEHIND THE WORK

Related publications by use case

01 Trial design and simulation

Trial2Vec
EMNLP Findings 2022

AutoTrial
EMNLP 2023

HINT
Cell Patterns 2022

SPOT
ACM-BCB 2023

InformGen
JAMIA 2025

02 Statistical programming

Can LLMs Replace Data Scientists in Clinical Research?
Nature Biomedical Engineering 2024

Where agents match a data scientist on clinical analysis tasks — and where they do not

DSWizard
Nature Biomedical Engineering 2026

Multi-step analysis agents reliable enough to be checked: roughly three months of work in about a day

03 Digital twin simulation

IPD reconstruction method
MLHC 2026 — accepted

TWIN
KDD 2023

MediTab
IJCAI 2024

TransTab
NeurIPS 2022

HALO
Nature Communications 2023

Our AI research provides the foundation for clinical research and biometrics — from trial design and simulation to statistical programming, patient-level modeling, and synthetic data generation.

POSITIONING

Key differentiators

Secure by design

- AI control plane separate from the secure data plane
- No raw data leaves your environment
- SOC 2 Type II, ISO 27001, HIPAA, GxP controls

Purpose-built model performance

- Agents trained on clinical trial data, not a general model prompted at it
- Published benchmarks on review, design, analysis, and recruitment tasks
- Test it against your own SAP

Comprehensive data pipeline

- 1.3M+ trial publications, 850K+ protocols, 150K+ patient notes
- Indexed across 14 registries and medical ontologies
- A trial knowledge base behind every answer, not open-web retrieval

Interoperable by design

- Import your own R, Python, and SAS code, templates, and skills
- Callable from scripts, including through Claude Code
- Reproducible pipelines — save, share, and re-run workflows

ENGAGEMENT

Three ways to work with us

01

License the platform

Deployed on-prem or in your VPC, and we help operate it — tuning for speed and accuracy, standing behind security and compliance.

OWN	The outputs, code, and validated workflows
TIME	1–3 months to fit your data model
PRICE	Integration cost plus licence

02

Delivered as a service

You buy the output: a completed SAP, a validated CDISC dataset, a TFL package, a submission-ready CSR section. For teams under a date they cannot move.

OWN	The deliverable and everything supporting it
PRICE	Per deliverable, study, or milestone

03

Build with us

For groups building their own agents. You have the mandate and domain expertise; we bring AI engineers who have shipped agents for regulated clinical work.

OWN	The pipelines, skills, and the capability
SCOPE	Defined scope with a delivery date
PRICE	Scoped project pricing

Keiji AI

THE FIRST STEP

Pick one workflow and we will run it

START HERE — NO DATA REQUIRED

Tell us your indication and endpoint

We reconstruct patient-level data from the published trials in your space, run the design grid, and walk you through what it shows.

THEN

Bring one completed study

We build the prognostic model and show what sample size that design would have supported, against the number you enrolled.

contact@keiji.ai

Keiji AI Inc · keiji.ai · Seattle, WA