



The Complete ML Platform. From Data to Production.

AutoML, MLOps, Privacy Suite, Synthetic Data, ML Studio & AI-Powered Dataset Builder — one integrated platform that replaces 5 different tools.

320+	72+	5	15
API Endpoints	PII Types	Deploy Strategies	Job Types

coreplexml.io

— THE PROBLEM

ML teams are drowning in **too many tools**

Getting a model from experiment to production requires stitching together 5-8 separate tools — each with its own learning curve, pricing, and failure modes.

Fragmented Stack

Jupyter + MLflow + Airflow + Seldon +
Great Expectations + custom scripts.
Each team builds its own glue code.

Compliance Gaps

PII detection, data masking, and audit
trails bolted on as afterthoughts.
HIPAA/GDPR compliance becomes a
manual process.

Slow Time-to-Production

Months from experiment to
deployment. Models rot in notebooks
while data scientists wait for
engineering resources.

The result: 87% of ML models never reach production. The gap isn't talent — **it's tooling.**

— WHY COREPLEXML

Stop stitching. Start shipping.

Your ML team is wasting 60% of its time on operational overhead. One platform replaces the 6-7 disconnected tools slowing your team down.

BEFORE — 6-7 SEPARATE TOOLS

- ✗ 6-7 disconnected tools with separate APIs and credentials
- ✗ Months of integration work before the first model ships
- ✗ 80-100 hrs/mo of senior engineering time on infrastructure
- ✗ No lineage — impossible to trace a prediction back to its data

WITH COREPLEXML — ONE PLATFORM

- ✓ One login, one API, one audit trail for everything
- ✓ From CSV to production endpoint in under 10 minutes
- ✓ 80-100 hrs/mo returned to building and testing models
- ✓ Full lineage: dataset → experiment → model → deployment

80-100 hrs

Engineering time redirected from infra to model building **every month**

3-6 months → minutes

Time to production from raw data to live endpoint

6-7 → 1

Tools consolidated — more experiments, fewer context switches

— THE SOLUTION

One platform. Everything you need.

CorePlexML replaces your entire ML stack with six integrated modules — from data preparation to production monitoring.



AutoML

Multi-Engine AutoML (H2O + FLAML)



MLOps

Registry, canary, A/B, drift, auto-retrain



Privacy Suite

72+ PII types, HIPAA/GDPR/PCI-DSS/CCPA



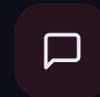
SynthGen

CTGAN, CopulaGAN, TVAE synthetic data



ML Studio

What-If analysis, SHAP, PDP, fairness



Dataset Builder

AI-powered data prep, no code needed

— WORKFLOW

Three steps to production ML

From raw data to deployed models with monitoring — in hours, not months.

1 **Upload Your Data**
CSV, Excel, JSON, XML — drag & drop your dataset and let the AI assistant prepare it. Automatic missing value handling, outlier detection, and type inference.

2 **Train Models**
AutoML selects the best algorithm, tunes hyperparameters, and builds ensemble models automatically. Multi-Engine AutoML (H2O + FLAML), including GBM, XGBoost, Deep Learning, and Stacked Ensembles.

3 **Deploy & Monitor**
One-click deployment with canary rollouts, drift detection, and auto-retraining. Real-time inference API with WebSocket streaming. A/B testing with automatic winner declaration.

— MULTI-ENGINE AUTOML

Two engines. One experiment. Best model wins.

Run H2O and FLAML in parallel on the same dataset. Each engine explores different optimization strategies independently. CorePlexML picks the best model across all engines.

H2O

15+ algorithms

XGBoost, GBM, Deep Learning, Random Forest, GLM, and Stacked Ensembles with Bayesian hyperparameter tuning.

FLAML

Cost-aware search

Microsoft FLAML finds optimal models within a time budget. Random Forest, Extra Trees, and Logistic Regression with scikit-learn.

Coming soon: AutoGluon Tabular, MLJAR-supervised, TPOT

Engine Access by Plan

PLAN	ENGINES	PARALLEL
Free	H2O only	—
Pro	H2O or FLAML	—
Team	H2O or FLAML	—
Enterprise	H2O + FLAML	✓ Up to 3

2x

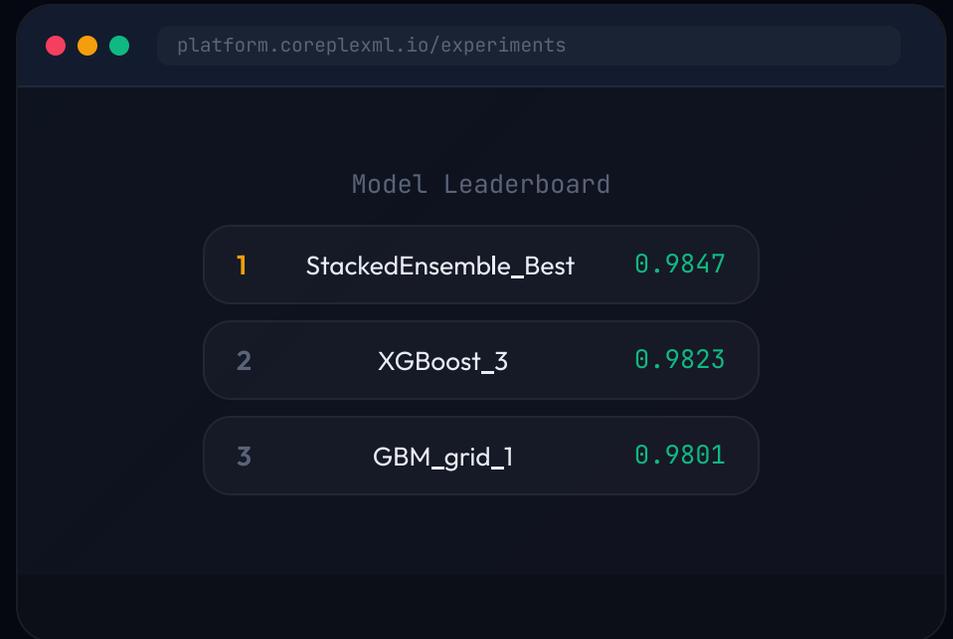
Model coverage with parallel multi-engine execution

— AUTOML

Configure once. Get production-ready models.

AutoML selects the best algorithms, tunes hyperparameters, and builds ensemble models. Then explore every metric, chart, and explainability view in detail.

- ✓ 4-step guided wizard — from dataset selection to training config
- ✓ Multi-Engine AutoML (H2O + FLAML)
- ✓ ROC curves, confusion matrix, precision-recall deep analysis
- ✓ SHAP values & variable importance for full explainability
- ✓ Cumulative gains, lift charts, and K-S statistics

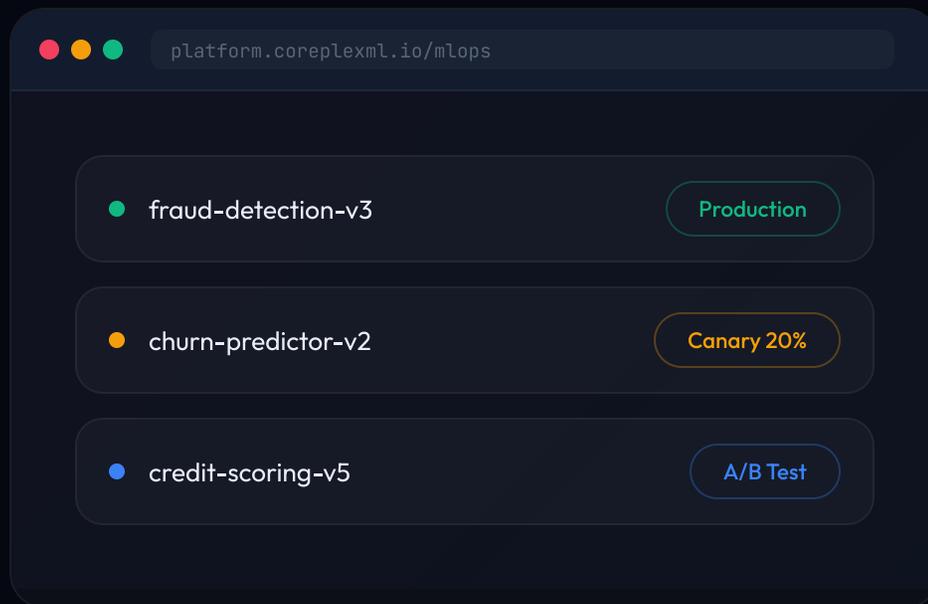


The screenshot shows a web browser window with the URL `platform.coreplexml.io/experiments`. The page displays a "Model Leaderboard" with three entries:

Rank	Model Name	Score
1	StackedEnsemble_Best	0.9847
2	XGBoost_3	0.9823
3	GBM_grid_1	0.9801

— MLOPS & MONITORING

Deploy, monitor, and iterate automatically.



From model registry to production deployment with canary rollouts, real-time monitoring, and automated retraining — all in one view.

- ✓ One-click deploy to staging & production environments
- ✓ Real-time inference API with WebSocket streaming
- ✓ A/B testing & canary deployments with traffic splitting
- ✓ Automatic drift detection with configurable alerting
- ✓ What-If Studio — compare scenarios vs baseline
- ✓ Automated retraining on schedule or drift trigger

— PRIVACY SUITE

72+ PII types. Compliance built in.

Detect, mask, redact, or encrypt sensitive data across your entire pipeline. Policy-based processing with complete audit trails.

HIPAA

Protected health information: names, MRNs, dates of birth, diagnoses

GDPR

Personal data: emails, phone numbers, addresses, national IDs

PCI-DSS

Payment data: credit card numbers, CVVs, expiration dates

CCPA

California consumer data: SSNs, driver's licenses, biometric data

Automatic Detection

72+ PII types detected automatically using NLP and pattern matching. No manual tagging required.

Flexible Actions

Mask, redact, hash, encrypt, or replace with synthetic data. Configure per-column policies with rule priorities.

Audit Trail

Every transformation logged with timestamp, user, policy, and before/after snapshots. Export-ready for compliance audits.

— SYNTHGEN

Generate millions of privacy-safe records.

Create statistically faithful synthetic datasets that preserve distributions without exposing real data. Perfect for development, testing, and data augmentation.

- ✓ 4 engines: CTGAN, CopulaGAN, TVAE, Gaussian Copula
- ✓ Statistical quality scoring — distribution fidelity metrics
- ✓ Privacy scoring — membership inference protection
- ✓ Generate 100 to 1M+ records with configurable epochs
- ✓ Column-level correlation preservation

Quality Score

94.7%

Distribution fidelity across 42 columns

Privacy Score

99.2%

Membership inference resistance

— ML STUDIO

What-If analysis & model comparison workspace

Interactive workspace for model exploration. Create scenarios, compare against baselines, and understand model behavior — no code required.

- ✓ SHAP values & variable importance for full explainability
- ✓ Partial Dependence Plots (PDP) for feature effects
- ✓ Scenario comparison with delta analysis
- ✓ 20 diagnostic tabs: ROC, confusion matrix, fairness, lift
- ✓ Schema-driven forms — adjust inputs, see predictions change

Diagnostic Tabs

20

Compare, Classify, Explain, Analyze, Advanced, Operations

Loan Underwriting

Test approval thresholds

Insurance Pricing

Simulate risk scenarios

Customer Churn

Test retention levers

Fraud Detection

Analyze decision bounds

— DATASET BUILDER

Tell the AI what you need. It builds your dataset.

Conversational AI interface powered by LLMs. Describe your goals in natural language and the AI executes the transformations — full script transparency.

You:

"Fill missing Age values with the median, drop the Cabin column, and one-hot encode Embarked."

AI:

"Done. Age filled (median=28.0), Cabin dropped, Embarked one-hot encoded into 3 columns. 891 rows, 0 missing values."

6-Step Pipeline

- 1 Load & Explore
- 2 Clean (missing values, outliers)
- 3 Transform (types, dates, text)
- 4 Encode & Scale
- 5 Feature Selection
- 6 Export (CSV / Parquet)

CSV/Excel/JSON/XML

Input formats

15+

Transform types

100%

Script transparency

— ARCHITECTURE & SECURITY

Built for enterprise teams

RBAC & SSO/SAML

Owner, editor, and viewer roles per project. SSO with SAML 2.0, OIDC, Google and GitHub OAuth.

GPU Scheduling

Automatic GPU allocation for training and inference. CUDA acceleration with intelligent CPU fallback.

A/B Testing

Split traffic between model variants. Bayesian and Frequentist analysis with automatic winner declaration.

Model Registry

Semantic versioning with stage transitions. Model cards, lineage tracking, and tag-based organization.

Audit Logging

Every action logged with user, timestamp, and context. Security events, access logs, and compliance trails.

Multi-Tenancy

Project-level isolation with PostgreSQL Row-Level Security. Per-tenant billing with usage metering.

PostgreSQL
Single source of truth

FastAPI
320+ REST endpoints

H2O.ai
ML engine

Docker
Container deployment

Built for your industry

Enterprise-grade ML solutions tailored for regulated and data-intensive industries.

Financial Services

Fraud detection, credit scoring, AML monitoring

AutoML Fraud Models

PCI-DSS Compliance

Drift-Triggered Retrain

99.2%

Fraud detection

67%

False positive reduction

<15 min

Model retrain time

Healthcare

Patient outcomes, clinical analytics, HIPAA-compliant pipelines

HIPAA Privacy Profiles

Synthetic Patient Records

Explainable Predictions

100%

HIPAA compliance

10x

Dataset augmentation

SHAP

Full explainability

Retail & E-Commerce

Demand forecasting, personalization, churn prevention

Demand Models

Canary Deploys

What-If Pricing

Government

Citizen data analytics with full anonymization and compliance

Multi-Compliance

Synthetic Data Sharing

Policy Impact Modeling

 **Insurance** — Claims prediction, underwriting, fraud

 **Manufacturing** — Predictive maintenance, quality

 **Telecom** — Anomaly detection, churn

 **Energy** — Load forecasting, compliance

— BY ROLE

Built for your team

Purpose-built workflows for every role in your ML organization.

ROLE	PAIN POINT	HOW COREPLEXML HELPS	RESULT
Developer API-first ML integration	ML integration complexity — different APIs for each tool	REST API & Python SDK — one unified interface for everything	Programmatic MLOps with CI/CD-ready SDK and 320+ endpoints
ML Engineer Full lifecycle MLOps	Experiment sprawl and deployment risk across environments	Model registry, canary/blue-green deploys, auto-retraining	Safe rollouts with drift detection and automated pipelines
Data Scientist AutoML & explainability	Algorithm selection fatigue and explainability gaps	50+ algorithm AutoML with built-in SHAP, PDP, and fairness	More experiments, less infra — focus on model quality
Business Analyst No-code ML tools	Technical barrier to scenario testing and reporting	Visual What-If Studio and automated PDF reports	Test business scenarios and generate reports without code
Compliance Officer Privacy & audit trails	Manual compliance audits across multiple regulations	72+ PII detection, HIPAA/GDPR/PCI-DSS/CCPA profiles	Full audit trails, one-click compliance reports

— COMPARISON

Replace your entire ML stack

One platform. All capabilities. A fraction of the cost.

CAPABILITY	COREPLEXML	DATAROBOT	H2O.AI	ALTERYX
AutoML	✓	✓	✓	✓
MLOps & Model Registry	✓	✓	—	Partial
Privacy Suite (72+ PII)	✓	—	—	—
Synthetic Data Generation	✓	—	—	—
What-If Analysis	✓	✓	—	✓
A/B Testing	✓	Partial	—	—
Conversational Builder	✓	—	—	—
Open Source	✓	—	✓	—
Starting Price	Free	\$\$\$\$\$	\$\$\$\$	\$\$\$\$

— PRICING

Start free.

Scale as you grow.

No credit card required. Upgrade when you need more. Cancel anytime.

Free

\$0

 /mo

Explore the platform. No credit card required.

WHAT'S INCLUDED

- ✓ 1 project, 5 models
- ✓ 10 jobs / day
- ✓ AutoML (all algorithms)
- ✓ 1 deployment
- ✓ API (rate-limited)
- ✓ ML Studio
- ✓ Dataset Builder (1/day)

Pro

\$49

 /mo

Full platform access for production ML.

EVERYTHING IN FREE, PLUS

- ✓ Unlimited projects & models
- ✓ Unlimited jobs & API
- ✓ Unlimited deployments
- ✓ Privacy Suite (72+ PII)
- ✓ SynthGen
- ✓ Canary & A/B testing
- ✓ Auto-retraining
- ✓ Priority support

Team

5 seats included

\$199

 /mo

Shared workspace for growing ML teams.

EVERYTHING IN PRO, PLUS

- ✓ 5 seats (up to 10)
- ✓ Team projects
- ✓ Centralized billing
- ✓ RBAC
- ✓ Team dashboard

Enterprise

Unlimited seats

Custom

Advanced security and compliance.

EVERYTHING IN TEAM, PLUS

- ✓ Unlimited seats
- ✓ SSO / SAML / OIDC
- ✓ On-premise deployment
- ✓ Custom SLA
- ✓ Dedicated account mgr
- ✓ GPU scheduling

Annual billing saves ~17%. Pro: \$490/yr · Team: \$1,990/yr

Grow at **your own pace**

Start with zero risk. Scale when you're ready.

①

Free[Start here](#)

Sign up in 30 seconds. Upload data, train your first model, deploy to staging. No credit card, no time limit.

②

Pro Trial[15 days free](#)

Unlock everything: unlimited models, Privacy Suite, SynthGen, canary deployments. Full access, no charge until trial ends.

③

Pro / Team[Production](#)

Choose individual Pro (\$49/mo) or Team (\$199/mo for 5 seats). Annual billing saves 17%. Upgrade or downgrade anytime.

④

Enterprise[Scale](#)

Unlimited seats, SSO/SAML, on-premise deployment, custom SLA, dedicated account manager. Talk to sales for custom pricing.

— IMPACT

Why teams choose CorePlexML

10x

Faster Deployment

From months to hours. One-click deployment with built-in monitoring eliminates the MLOps bottleneck.

5+

Tools Replaced

One platform instead of Jupyter + MLflow + Seldon + Great Expectations + custom privacy tools.

\$0

To Get Started

Generous free tier with real capabilities. No credit card. No time limit. Start shipping models today.

100%

Built-in compliance (HIPAA, GDPR, PCI-DSS, CCPA)

320+

REST API endpoints — fully programmable platform



Ready to simplify your ML workflow?

Join the data teams shipping models faster with CorePlexML.
Start free, no credit card required.

[Start Free Today](#)

[Talk to Sales](#)

Website
coreplexml.io

Platform
platform.coreplexml.io

Email
sales@coreplexml.io

Docs
docs.coreplexml.io