



JIGYASA BY MINDFORGEAI

Machine Learning in the Age of AI

Equations. Data. Models. Decisions.

Akash Chatake



BABYX OPENER

Learning, interaction and appearance

SENSE

What enters the system?

UPDATE

What changes with experience?


BabyX First Words



Watch the linked lab film, then distinguish a convincing response from evidence of understanding.

ROADMAP

One experiment, from equation to prediction



The same learning loop reappears in larger ML systems.

DIFFERENTIAL EQUATION

A cooling model with one unknown

$$dT/dt = -k(T - 25)$$

$$T(t) = 25 + 55 \exp(-kt)$$

25°C

Room temperature

80°C


Initial temperature

k = ?

Rate learned from data

SYNTHETIC OBSERVATIONS

Observed data constrains the equation



7 observations

The equation is fixed.
The rate k is still unknown.

Separate training observations from later held-out checks.

PARAMETER ESTIMATION

Fitting is training this model




Loss = mean of (prediction – observation)²

The computer changes k repeatedly; the learner specifies the model and objective.

TRAINING RESULT

The fitted curve follows the observations




$k \approx 0.03979$

per minute
A parameter learned from
seven observations.

Training error alone does not establish future performance.

LEARNING HISTORY

Training reduces the stated loss




175.46 → 0.094

Training MSE in °C²
Same observations, better
parameter.

An improving training curve is evidence of optimisation; an unseen check asks a different question.

HELD-OUT CHECK

A check on two later observations



0.133°C

Mean absolute error on two held-out synthetic observations.

A small synthetic check demonstrates the method; real use needs more varied data.

INFERENCE

A prediction at a new time

$t = 45$ minutes

$$25 + 55 \exp(-0.03978592 \times 45)$$

$$\approx 34.18^{\circ}\text{C}$$

Inference substitutes a new input into the fitted equation; changing k would be additional training.

LIMITS

The assumptions travel with the prediction

What if the cup changes?


The fitted rate may change.
A good fit does not protect against a different physical setup.

Physics, data quality and evaluation are part of the same engineering problem.



CAMPUS-ENERGY CASE

Data and exploratory analysis



600 cases

Units · missingness · ranges
Timing · duplicates ·
provenance

The data dictionary comes before the model; transformations are fitted on training data only.

LEARNING TASKS

The task determines the output

Regression

How many kWh?

Classification

High demand or normal?

Clustering

Which patterns are similar?


Forecasting

What comes next?

Generation creates text or media and requires its own evaluation.

SUPERVISED MODELS

Linear and logistic regression



$$y = w_0 + w_1x$$

Linear: a numeric estimate


Logistic: a class probability

Both learn weights from labels.


A simple baseline makes later improvements meaningful.

SUPERVISED MODELS

KNN and SVM use geometry differently



KNN · nearby labelled cases



SVM · a separating margin

Distance and margins depend on representation, units and kernel choices.

DECISION RULES

Trees and ensembles




A tree splits feature space

Tuning and leakage can still produce an invalid result, even with a strong algorithm.

Forest: combine trees
XGBoost: correct error in stages


UNSUPERVISED LEARNING

K-means groups unlabelled examples



The method groups by distance; a person explains whether those groups are useful.


LSTM carries state through a sequence



Gated memory passes between ordered steps

Sequence construction, history length and future-data leakage matter more than the model name.

Transformers relate tokens through attention



A token representation uses information from other tokens.

The full system also includes training data, objectives, tuning, evaluation and deployment.

A map of the model names

KNN / SVM

Supervised methods

K-means

Unsupervised clustering

XGBoost

Boosted-tree implementation

LSTM / transformer


Neural architectures

GPT / Gemini

Foundation-model families

CAMPUS-ENERGY VALIDATION

Evaluation starts with a fair comparison



Same cases

Same metric
Only advance inputs
A separate held-out test

These synthetic results do not establish real-campus performance.

LEAKAGE INVESTIGATION

The perfect score exposes a timing error

0.0

kWh validation MAE

End-of-day meter

= the outcome copied after the day

Remove the feature, retrain and audit availability at prediction time.

CLASSIFICATION DECISIONS

A label and an alert threshold

125 kWh

Observed boundary
Defines the target class.

0.30 probability

Decision boundary
Triggers an illustrative alert.

Precision, recall and error costs explain the decision better than accuracy alone.

REAL-LIFE EXAMPLES

Applications are decisions in context

Retail	Stock forecasts → replenishment
Factory	Vibration anomaly → technician review
Healthcare	Risk estimate → validated clinical review
Support	Triage and draft → verified response



How a foundation model is developed



Training objective, data, tuning and system tools all influence observed behaviour.

AS OF 17 OCTOBER 2026

Current model examples

GPT-6 Astra

Gemini 4 Argon

A product name or version number is not an intelligence metric. Read its model card and limits.

BENCHMARKS

What a model score actually measures

90 / 100

on which questions?

A score needs context

Dataset and task
Metric and denominator
Date and test conditions

Benchmark success does not certify a different local workflow.

RELIABILITY

Hallucinations and factual checks


“The paper proves 100% accuracy.”

Can we open the paper and find that claim?

Source inspection → evidence → review

Fluency and confidence in wording do not establish factual correctness.

Agentic AI is a controlled tool loop



Permissions · state · checks · stop rules · human approval


A written plan is one component of a working agent system.

Engineering work in the age of AI

AI can assist with tasks

Draft code and reports
Inspect data and suggest tests
Repeat bounded workflows

Engineers frame the problem, verify the evidence and remain responsible for deployment.



OUR ECOSYSTEM

Project demonstrations

MindForgeAI

13-project catalogue

Apollo ABS-100

Reference engineering assembly

GFIS Upgrade


Simulation and evidence workflow

MLOps Academy

Versioning, deployment and monitoring

STUDENT ROADMAP

Your next experiment



One small model, an honest test and a clear limitation

Continue with the notebook, printed notes and Workshop 01.