

Custom AI/ML models for your business — predict, automate, and generate in a fixed time & budget.

di~f Data Intelligence Factory

di-factory.biz | contact@di-factory.biz

di~f di-factory.co

Crafting AI/ML Models for Your Business Needs

di-factory | Data Intelligence Factory

Tailor-made machine learning. Fixed time. Fixed budget.

The Problem

- Businesses sit on **data goldmines** but lack the expertise to extract value
- Traditional ML projects are **expensive, slow, and unpredictable**
- Off-the-shelf solutions don't fit your specific business context
- Internal teams lack specialized ML engineering capacity
- Projects drag on for months with uncertain outcomes

Our Solution

- **Tailor-made ML/DL models** delivered in fixed time & budget
- No surprises — scope, timeline, and cost determined upfront
- From ideation to production in **2–4 weeks**
- You own everything: code, models, pipelines, IP

What We Build

Predict

- Sales, costs, inventory, churn, credit scoring

Automate

- Processes, fraud detection, quality control

Generate

- Content, recommendations, intelligent chatbots

Use Cases

-  **Predict** next sales / costs / inventory levels
-  **Credit scoring** & risk assessment
-  **Churn prediction** & retention triggers
-  **Intelligent chatbot** for customer service
-  **Logistics optimization** & route planning
-  **Product recommendations** engine
-  **Fraud detection** in real time
-   **Process automation** with ML decision-making

Our Pipeline — 5 Phases

PHASE	FOCUS	TIMELINE
1	Consultation & Solution Ideation	Days 1–2
2	Data Acquisition & Analysis	Days 3–5
3	Data Transform Pipeline	Week 2
4	ML Modeling & Transformation	Weeks 2–3
5	Deployment & Evaluation	Final days

Phase 1: Consultation & Solution Ideation

Days 1–2

- Design Thinking workshop with stakeholders
- Data landscape assessment
- Use case mapping & prioritization
- Success criteria definition
- Solution architecture proposal

Phase 2: Data Acquisition & Analysis

Days 3–5

- Data profiling & exploration
- Quality assessment & gap analysis
- Data governance review
- Source integration planning
- Baseline metrics established

Phase 3: Data Transform Pipeline

Week 2

- Data preprocessing & cleaning
- Feature engineering
- Feature Store setup
- Automated data pipelines
- Reproducible transformations

Phase 4: ML Modeling & Transformation

Weeks 2–3

- Algorithm selection & benchmarking
- Model training & hyperparameter tuning
- Validation & cross-validation
- ML Registry for versioning
- Performance evaluation against success criteria

Phase 5: Deployment & Evaluation

Final Days

- Production integration (API / batch / streaming)
- Monitoring & alerting setup
- Knowledge transfer sessions
- Documentation handoff
- Go-live support

The Squad

ROLE	RESPONSIBILITY
BIZ Engineer	Requirements, use cases, business alignment
Data Engineer	Pipelines, data quality, Feature Store
ML Engineer	Models, training, validation, registry
SW Engineer	APIs, integration, deployment, monitoring

Each specialized. Working in concert.

Timeline

- 2–4 weeks depending on complexity
- Determined **upfront** during Phase 1
- No surprises. No scope creep.
- Weekly progress demos
- Daily async updates

What You Get — Deliverables

-  Solution Brief & Architecture
-  Data Quality Report
-  Production Data Pipeline
-  Trained Model + Model Card
-  Monitoring Dashboard
-  Operations Guide
-  API Documentation
-  Knowledge Transfer Sessions
-  **Full IP Transfer** — you own everything

Industries We Serve

INDUSTRY	EXAMPLE USE CASES
 Healthcare	Patient risk scoring, readmission prediction
 Banking	Credit scoring, fraud detection
 Insurance	Claims prediction, risk assessment
 Fintech	Transaction anomaly detection, churn prediction
 Retail	Demand forecasting, product recommendations
 Manufacturing	Quality control, predictive maintenance
 Real Estate	Price prediction, market analysis
 Education	Student performance prediction, personalized learning

Why di-factory

-  **Fixed Time** — 2–4 weeks, guaranteed
-  **Fixed Budget** — no surprises, no overruns
-  **Lowest Investment** — AI-powered efficiency reduces cost
-  **Lowest Ops Cost** — open source, zero licensing fees
-  **Best of Both Worlds** — human expertise + AI acceleration

Technology

- **Open source** stack — no vendor lock-in
- **Cloud-native** architecture — scalable from day one
- **No licensing fees** — zero ongoing software costs
- **You own everything** — code, models, pipelines, IP
- Standard tools: Python, scikit-learn, PyTorch, TensorFlow, MLflow, Airflow, AWS/GCP

After Delivery

-  **2 weeks support** included at no extra cost
-  **Optional maintenance sprints** for enhancements
-  **Automated retraining pipelines** — models stay fresh
-  **Monitoring dashboards** for ongoing performance tracking
-  **Scale with us or independently** — no dependency

Next Steps

- 1 Choose an experiment — pick your highest-value use case
- 2 We run the pipeline — 2–4 weeks, fixed scope
- 3 Deliver & evaluate — production model + full documentation
- 4 You get a roadmap — what's next for your data strategy

Start small. Prove value. Scale with confidence.

Let's Talk

 di-factory.biz  contact@di-factory.biz

Fixed time. Fixed budget. Your AI model in production.

di-factory — Data Intelligence Factory