

Chelsea AI Ventures

Scaling SME Business With Enterprise AI

*We help small and medium enterprises implement and scale
LLM technologies with measurable ROI*



www.chelseaai.co.uk

Contents

1	Company Overview	3
1.1	Our Mission	3
2	Our Validation-First Approach	3
3	Our Collaboration Models	5
3.1	Validation Partnership: Prove it works before you invest	5
3.2	Advisory Collaboration: We design, you build, we guide	5
3.3	Collaborative Implementation: Shared risk, shared success	5
4	The Open-Source AI Advantage	6
4.1	Cost Effectiveness	6
4.2	Customization & Control	6
4.3	Data Privacy & Compliance	6
4.4	Performance Optimization	6
5	Technology Expertise	6
5.1	Open-Source LLMs & Models	6
5.2	Machine Learning	6
5.3	NLP & Vector Databases	6
5.4	Front-End & Web Dev	6
5.5	Infrastructure	7
5.6	Development	7
5.7	High-Performance Computing	7
5.8	Computer Vision	7
6	Client Success Story	8
7	Contact Information	8

Company Overview

Chelsea AI Ventures Ltd specializes in making enterprise-grade AI accessible and effective for small and medium businesses (SMEs). We bridge the gap between complex AI technologies like Large Language Models (LLMs) and practical business applications, focusing on delivering measurable Return on Investment (ROI) for organizations that may have limited resources compared to large enterprises. Our approach prioritizes risk reduction, tangible results, and building stakeholder confidence from day one.

Our Mission

Our mission is to democratize access to advanced AI for SMEs, providing the same competitive advantages that larger enterprises enjoy, but with implementations tailored to SME budgets, timelines, and security constraints. We believe in validating ideas quickly and affordably, ensuring you only invest in solutions proven to work for your business.

Why Choose Chelsea AI?

- **Validation-First Approach:** We de-risk your investment by proving feasibility with small, fixed-cost projects.
- **Expert Leadership:** Founded by a bestselling author of AI/ML implementation books.
- **Open-Source Specialists:** Deep expertise in leveraging models like Llama 3 to reduce costs and enhance control.
- **SME Focus:** Enterprise-grade capabilities designed for SME budgets and operational realities.
- **Proven Results:** Client solutions delivering 70-90% cost savings and significant ROI improvements.

Our Validation-First Approach

We address real client concerns with structured validation. Each step is designed to reduce risk and build the stakeholder confidence needed to move forward. This process avoids costly, large-scale failures by focusing on tangible evidence.

"I think this could work, but I need real proof before asking for a big budget."

1. Identify Critical Assumptions: We start by isolating the core technical and business assumptions that will make or break your project. We focus on the 'unknown unknowns' that create the most risk.

"We need clear, measurable success metrics, not vague promises."

2. Define Go/No-Go Criteria: Together, we define specific, measurable criteria for success. This includes technical feasibility (e.g., latency, accuracy) and budget constraints, forming a clear basis for a go/no-go decision.

✓ **Example Criteria:** 90% of queries complete under 1000ms; 85% SQL generation accuracy; Fixed £18K budget, 4-week timeline.

"Let's do a short, fixed-price PoC on our infrastructure to avoid internal IT delays."

3. Execute Validation PoC: We build a self-contained proof-of-concept, typically in 3-4 weeks for a fixed cost. This validates the core assumptions with demonstrable evidence.

"Give us a clear answer. If it works, show us how. If not, tell us why."

4. Deliver Definitive Recommendation: You receive a clear go/no-go recommendation backed by the PoC results. If successful, you have the evidence needed for a larger investment. If not, you've avoided a costly failure.

"We want you as the brain of the project, guiding our team to build it right."

5. Guide Implementation (If Validated): For validated concepts, we transition into an advisory role, providing architectural oversight to ensure your team can build, own, and maintain the solution effectively.

Our Collaboration Models

We structure our collaboration to address the real challenges technical leaders face: budget approval, risk management, and stakeholder confidence.

Validation Partnership: Prove it works before you invest

"We think this could work, but need proof before we can get budget approval."

A self-contained proof-of-concept that validates your AI initiative's feasibility without internal dependencies or large commitments. This is ideal when you need proof an AI solution will work within your constraints before seeking larger budget approval.

What's Included

- Fixed-price validation (typically £15K-25K)
- 3-4 week timeline with weekly milestones
- Built on our infrastructure to avoid IT delays
- Clear go/no-go decision with detailed metrics

Advisory Collaboration: We design, you build, we guide

"We want you as the brain of the project - doing architecture and guiding our team."

A collaborative approach where we provide technical architecture and ongoing guidance while your team builds, retaining full ownership. This model is best for organizations wanting to build internal AI capabilities while ensuring technical success and mitigating risk.

What's Included

- Technical architecture design and validation
- Sprint-by-sprint guidance and reviews
- Knowledge transfer to your team
- Performance accountability without dependency

Collaborative Implementation: Shared risk, shared success

"The timeline has to work with the budget. We need something tangible with clear milestones."

Full development with payments tied to demonstrated progress, allowing you to validate value at each stage before continuing. This is suited for larger projects where you need to demonstrate progress incrementally to stakeholders and de-risk a significant investment.

What's Included

- 2-week milestone payments
- Working demos at each stage
- Ability to pause or pivot based on results
- Risk is shared between both parties

The Open-Source AI Advantage

Chelsea AI has deep expertise implementing both commercial and open-source LLMs. For many clients, particularly SMEs, open-source models like Meta’s Llama 3 offer significant advantages in cost, customization, and control.

Cost Effectiveness

Open-source LLMs can reduce operational costs by 70-90% compared to commercial API-based models, while maintaining comparable performance for many applications.

Customization & Control

Fine-tune models to your specific domain and use cases, creating AI that truly understands your business language. Maintain full control over your data without sharing it with third-party API providers.

Data Privacy & Compliance

Deploy models on-premises or in your private cloud, ensuring sensitive data never leaves your secure environment. This is ideal for industries with strict regulatory requirements.

Performance Optimization

We specialize in quantization and optimization techniques (e.g., GGUF, vLLM) that make large models run efficiently on standard hardware, ensuring maximum performance.

Our Open-Source LLM Expertise

Implementation Capabilities

- Local deployment of Llama 2 & 3 models
- Custom fine-tuning (LoRA)
- RAG pipelines for grounding
- Multi-modal capabilities

Business Outcomes

- Reduced operational costs (70-90%)
- Enhanced data security & control
- Customized AI for your industry
- Future-proof solutions

Technology Expertise

We leverage a comprehensive set of state-of-the-art AI and ML technologies to build robust, scalable, and effective solutions.

Open-Source LLMs & Models

- Meta’s Llama 2 & 3
- Mistral & Mixtral
- Local Model Deployment
- GGUF Quantization
- LoRA Fine-tuning
- Ollama & vLLM

- Hugging Face Ecosystem
- Custom Embeddings
- Production RAG Pipelines
- OpenAI & Anthropic APIs

Machine Learning

- PyTorch & TensorFlow
- scikit-learn & XGBoost
- Prophet & Causal Modeling

Front-End & Web Dev

- Flutter & Astro
- React & Next.js
- Responsive UI Design

NLP & Vector Databases

- LangChain & LlamaIndex

Infrastructure

- AWS, GCP, Azure ML
- Snowflake
- Docker & Kubernetes
- Serverless Architecture

Development

- Python (Django, FastAPI)
- Java & Kotlin
- C++

High-Performance Computing

- Distributed Systems
- Parallel Processing
- Multi-GPU Clusters

Computer Vision

- Deep Learning Vision
- Object Detection
- Multi-modal Models

Client Success Story

FinTech SME: Open-Source Cost Reduction

Challenge: A financial technology SME was using a commercial API for customer service automation. While effective, the costs were scaling rapidly with user growth, and they were unable to customize the model’s responses to their niche industry terminology.

Solution: Chelsea AI executed a **Validation Partnership** to prove the viability of a fine-tuned open-source model. We used Meta’s Llama 2, fine-tuned it on their internal documentation, and deployed it within their private cloud. The PoC was completed in 4 weeks.

Results:

- **85% reduction** in operational costs compared to the commercial API solution.
- **Significant improvement** in query accuracy for industry-specific terms.
- **Full data sovereignty**, as no sensitive data ever left their infrastructure.
- The successful PoC provided the evidence needed to fund a full-scale implementation.

"Chelsea AI helped us implement a fine-tuned Llama model that understands our industry terminology and handles our customer queries with remarkable accuracy. By hosting it on our own infrastructure, we've reduced costs by 85% compared to commercial API solutions." — Technical Director, FinTech SME

Contact Information

Chelsea AI Ventures Ltd	Connect With Us
Website: https://chelseaai.co.uk	✉ Email: info@chelseaai.co.uk
Registered in England & Wales	🌐 LinkedIn: https://www.linkedin.com/company/106959452
Company No. 16264968	📖: https://chelseaai.substack.com/
Location: London, UK	📰 Medium: https://medium.com/@auffarth
	🐙 GitHub: https://github.com/benman1

Ready to transform your business with AI?

Schedule a free 30-minute consultation to discuss your AI goals.

Book Your Consultation Today