


LIZZY 7B

by Flower

Designed for the UK

Built from the ground up for UK use cases, standards, and deployment environments.



Deployment-ready sovereign LLM built in the UK for enterprise and public-sector use.

Lizzy is a 7B decoder-only transformer with long-context support, controllable attention behavior, and deployment-specific serving configurations designed for secure, sovereign AI workloads.

7B parameters

Long-context up to 65k tokens

Decoder-only transformer

Sliding-attention behaviour

UK Sovereign AI model



Tune to internal AI infrastructure

Optimize for your existing stack. Seamlessly deploying across CPUs, GPUs, NPUs, and more to maximize performance, efficiency, and ROI



Incorporate proprietary data

Integrate internal datasets into model training and inference pipelines. Improving accuracy and alignment to your company requirements, with data remaining fully within your environment



Align with compliance requirements

Compatible with requirements of GDPR, EU AI Act, NIS2, DORA and more. Run fully within your environment for greater control, security, and data sovereignty.

Top European Model Performance

Lizzy 7B matches or outperforms European sovereign baselines across reasoning, coding, and UK-specific benchmarks.

Benchmark	Lizzy 7B	EuroLLM 9B	Apertus 8B
Britishness MCQ	71.0	77.6	80.8
Britishness CoT	80.1	72.1	31.7
Britishness Domains	89.9	69.0	32.6
MATH	77.9	31.3	22.4
OMEGA	29.0	4.7	4.7
BigBenchHard	69.0	38.9	42.4
AGI Eval English	65.6	50.2	50.4
MMLU	67.9	57.4	63.4
GPQA	34.6	26.8	28.1

Additional downstream task metrics available upon request.

Intended uses

- Scenarios demanding deep knowledge of the UK
- General purpose reasoning, coding and agent tasks
- Private deployment via Hugging Face or vLLM

Built on years of Flower Labs R&D

Lizzy 7B is built on a foundation of peer-reviewed novel approaches to model embeddings, architecture, optimizer and the full tracking stack.

MLSys 2025 • 2025

ICLR 2025 • 2026 • 2026

Built for Enterprise Deployment

- Deploy anywhere with full control. Compatible with vLLM and leading inference stacks
- 24/7 support for deployment across your data and compute systems.



flower.ai/models/lizzy/

huggingface.co/flwrlabs/Lizzy-7B

hello@flower.ai