

Open Source AI for Public Good: Lessons from Falcon LLM for Labour Market Applications

A hand reaching out to a glowing digital sphere, symbolizing human-AI interaction. The background is a dark blue digital space with vertical lines and data points.

Technology Innovation Institute

Dr. Eleana Kafeza
Lead Researcher
Eleana.kafeza@tii.ae

Artificial Intelligence & Digital Science Research Center



Structure →



Program Development



Applied Research



Commercialization

Programs

<p>DATA</p>	<p>SUPER ALIGNMENT</p>	<p>EVALUATION</p>	<p>LLM VERTICALS</p>	<p>TRAINING & FRONTIER</p>	<p>HPC, DEPLOYMENT & OPTIMIZATION</p>	<p>GENAI FOR ROBOTICS</p>
-------------	------------------------	-------------------	----------------------	--------------------------------	---	---------------------------

KPI Overview – since January 2022

<p>15</p> <p>Nationalities</p>	<p>80</p> <p>Team members</p>	<p>19</p> <p>Emirati Scientists</p>	<p>93</p> <p>Publications</p>	<p>31</p> <p>Differentiated products (TRL4+)</p>	<p>1</p> <p>Board of Advisors</p>	<p>Partnerships</p>	<p>9</p> <p>Patents</p>
--------------------------------	-------------------------------	-------------------------------------	-------------------------------	--	-----------------------------------	---------------------	-------------------------

Leadership Team



Dr. Hakim Hacid
Chief Researcher



Hichame Khoulouf
Director, PMO



Dr. Ahmed Alhammad
Director

Advisory Board Member



Mounia Lalmas
Senior Director of Research, Spotify

Falcon Series : Overview



Falcon 3 series are the **most powerful small-scale** large language models, designed to meet a variety of needs and use cases. Available in four different sizes: **1B, 3B, 7B, and 10B**



Falcon 3 Multimodal are advanced LLMs capable of processing multiple data types including **text, images, video and audio**. These models offer **state-of-the-art** functionality while maintaining a **lightweight footprint**.



Falcon 3 Mamba 7B is an advanced state-space model designed for **efficient long-context processing**. It leverages a **pure Mamba architecture**, enabling **constant memory usage** and **faster token generation**.



Falcon-Arabic 7B sets a new benchmark for Arabic LLM. Built on the powerful **Falcon 3** architecture and optimized for **native and dialectal Arabic**.



Falcon-Edge redefines efficient AI with **ultra-lightweight** models designed for edge devices. Based on **BitNet architecture**, delivering strong performance with **minimal compute and memory needs**



Falcon-H1 blends classical Transformer attention with advanced State Space Models (SSMs) for exceptional **long-context reasoning** and efficiency. Six scalable sizes ranging from **0.5B-34B**.



Falcon-H1-Reasoning 7B blends **hybrid Mamba-Transformer** architecture with advanced reasoning optimization to deliver efficient long-context **logical and multi-step reasoning**.



Falcon-H1 Arabic sets a new benchmark for Arabic LLMs, built on H1 **hybrid architecture**. Optimized for **native and dialectal Arabic**. Available in **3B, 7B, and 34B**.

Key Focus Areas

- ✓ Natural Language Processing (NLP)
- ✓ Arabic & Multilingual Excellence
- ✓ Edge & Resource-Constrained Deployment
- ✓ Multimodal AI (Vision, Video, Audio)
- ✓ Hybrid and Next-Gen Architectures
- ✓ Scalability & Open Deployment



Falcon-H1-Tiny delivers **ultra-compact**, high-performance LMs built on the hybrid H1 architecture for on-device and edge AI. With **≤100M parameters**, it enables efficient **multilingual chat, reasoning, and coding** at minimal compute and memory cost.

What LLMs can do for Public Employment services and Labor Policy Design ?

- **Explainable AI for Labour-Market Governance and Policy Design**

LLMs with explainable reasoning can translate complex PES procedures—such as eligibility checks—into clear, step-by-step workflows.

- Decompose regulations into interpretable decision steps, (e.g. eligibility criteria for benefits) explain why specific outcomes apply, and provide traceable justifications aligned with legal and policy frameworks.

- **Administrative automation and case management support**

- Automate labour-market administration tasks such as CV parsing, vacancy analysis, document processing, and case-note summarisation, transforming unstructured information into structured and auditable records.
- More efficient case management, supports monitoring and evaluation activities, and enables evidence-based labour policy design, while reducing administrative burden on PES staff.

- **Personalised career guidance and employability support**

- Support jobseekers with tailored career guidance, including skills identification, skills-gap analysis, and pathway recommendations.
- Leveraging multimodal labour-market data—such as work descriptions, photos of tasks performed, or tools used—to recognise informal, non-standard, and experiential work experience.
- Inform labour policy design, helping policymakers understand skills gaps, employment trends, and the impact of interventions, while enabling more inclusive and evidence-based employability support.

- **Evidence-informed labour policy design:**

Leverage LLMs to analyse labour-market data, administrative records, and stakeholder inputs, generating interpretable insights that support policy monitoring, evaluation, and the design of targeted, inclusive, and evidence-based interventions.

What LLMs can do for Public Employment services and Labor Policy Design ?

Training, reskilling, and upskilling guidance

Identify reskilling and upskilling needs based on individual profiles and labour-market trends, and recommend relevant training programs aligned with regional and sectoral demand.

Multilingual and inclusive access to PES services

Enable multilingual interaction for jobseekers and employers, improving accessibility for migrants and diverse populations and reducing language barriers in labour-market participation.

•Edge-based, privacy-preserving AI

Deploy lightweight LLMs (e.g. TinyFalcon) with

- Privacy preserving
- Low latency
- Reduced cloud dependency
- Small footprint →
 - Regional employment statistics
 - Secure input from stakeholders
 - No exposure of personal data for extracting trends, skill gaps, emerging policy risks

Ethical and Related Considerations

- **Ethical use in labour-market decision support**

Develop specific ethical frameworks, methodologies, and guidelines to ensure that LLM-based systems support employment services without undermining human autonomy, fairness, or the rights of jobseekers and employers.

- **Bias and non-discrimination in labour-market outcomes**

Implement safeguards to detect and mitigate bias in training data and model outputs, ensuring that recommendations do not disadvantage individuals or groups based on protected or proxy attributes (e.g. age, gender, ethnicity, disability, migration status).

- **Accountability and explainability**

PES operators and technology providers should ensure transparency regarding data sources, system capabilities, and limitations. LLM outputs must be explainable to counsellors, jobseekers, and oversight bodies, enabling contestability and accountability in decision-making.

- **Privacy and data protection of jobseekers and employers**

Jobseekers and employers must be clearly informed about how their personal data is processed. Especially when using multimodal inputs or AI-assisted profiling

Security Considerations for LLM-based Systems



Expanded attack surface introduced by LLMs

Software systems supporting Public Employment Services handle sensitive personal, employment, and institutional data. Security is therefore a fundamental requirement. LLMs introduce new security risks, including prompt injection, data leakage, model manipulation, and indirect system access through natural-language interfaces. These risks are amplified when LLMs are integrated into operational PES systems.

Protection of the infrastructure and data assets

Appropriate security tools and practices must be employed to detect and prevent attacks targeting PES software systems, including vulnerabilities in binaries, libraries, and third-party components. Security analysis tools (e.g. binary analysis and vulnerability detection frameworks such as *Binsleuth*) can support proactive protection of PES infrastructure.

Secure deployment environments for LLMs and agents

LLMs—particularly those supporting agentic reasoning and tool execution—must be deployed in secure, controlled environments. Sandboxing, strict permission management, and isolation of execution environments are essential to prevent unintended or malicious actions.

Some technical solutions for introducing ethics into LLMs

- ❑ Incorporate ethical guidelines and value-based reasoning directly during pre-training or fine-tuning phases to influence the model's behavior.
- ❑ Structured reasoning with retrieval support: Enable the model to perform structured reasoning over retrieved knowledge (RAG) while respecting ethical, legal, and regulatory boundaries in its outputs.
- ❑ Domain- and regulation-specific tailoring: Fine-tune the model to meet sectoral or regulatory requirements, ensuring decisions and recommendations are contextually appropriate.
- ❑ Ethical prompt design: Craft prompts that explicitly guide the model toward safe, ethical, and value-aligned outputs.
- ❑ Institutional data control: Keep sensitive reasoning and decision-support processes within organizational boundaries to protect privacy, ensure transparency, and enforce ethical oversight.

A futuristic digital scene with a human hand on the left and a white and blue robotic hand on the right. They are reaching towards each other, with a bright light at the point of contact. The background is dark blue with glowing lines, data points, and binary code. The text "Thank you" is centered in white.

Thank you