

Fursah: Explainable AI for Workforce Readiness

Team name: Visionary

Members: Shoug Alomran, Lolwah Alsaadoun, Renad Alsulaiman, Taleen Bin Nader

Solution name: Fursah: AI Workforce Readiness Platform (fursah.org)

Contact details: info@fursah.org — Riyadh, Saudi Arabia

1. Introduction

The application of Artificial Intelligence (AI) has been widespread in different industries to process large amounts of complicated data. Fursah, our suggested solution, is an AI-driven readiness of the workforce system that utilizes the combination of AI document processing and transparent, deterministic intelligent models. Based on the verification of skills, certifications, experience, projects, career interests, and labor market needs, Fursah calculates Career Readiness Scores, detects the skill gaps, matches students with appropriate positions, and provides adaptive suggestions. Generative AI technology is used for evidence extraction and for role-aware assistance, while consequential scores and match results are generated using explainable rule-based models.

This is due to the fact that the disparity between educational productivity and labor demand is still the main long-term problem of the Human Capability Development Program within Vision 2030: the Kingdom must graduate graduates who can match the needs of the industries it develops, given that the disparity is still one of the sixteen strategic aims of the Program. At the moment, the disparity can be addressed using conventional means such as graduate surveys, curriculum revisions and employers' complaints. Fursah addresses the disparity by utilizing feedback; the needs of the employers become known to the universities and their students.

Most importantly, hiring decisions carry considerable weight and in most international contexts, considerable risk. An automated decision-support system that cannot clearly explain how its outputs are produced puts more at stake than an individual error: it risks perpetuating exclusion on a national scale. Consequently, this document identifies the technical solution through the lens of the policy and strategy level: the Principles and Controls of AI Ethics of SDAIA, the PDPL, the national standards for data governance, and the Human Capability Development Programme, such that the technical solution is based in the AI Readiness criteria and not seen as an independent product.

2. Use case and gaps in existing solutions

The issue is that the journey between the university and work is filtered by processes that are unclear, driven by credentials and networks. Employers sift through a large number of CVs using keyword-search applicant-tracking systems; the university gets no signal about whether their graduates are being hired or not; the student receives rejection for no reason and, therefore, has no way forward. Below are the specific gaps:

- **Obscure filtering criteria.** The current ATS ranking is a black box. A candidate does not know why he/she was filtered out, and an employer is unable to check if the filtering criteria are reasonable. There is no ground for appealing against or changing the decision.
- **Biased criteria:** qualifications above skills. Instead of measuring real skills, screening criteria such as the institutional reputation, GPA and job title are used.
- **Disconnected from education.** Reviewing cycles in education are many years behind labour market changes. There is no feedback from the employers' demand on curriculum design.
- **Hidden discrimination.** Data of previous hiring contains the past preferences and models trained on it will be discriminative with regard to gender, region or institution unless explicitly checked.
- **Uncontrolled data.** Data about a student's competencies is fragmented and belongs to different organizations.

How Fursah closes them:

- **Reasoning and explainability.** Readiness scores, matches, gaps, and recommendations are generated from explicit job-related evidence and transparent scoring rules. The platform shows users the factors contributing to these outputs, while model and ruleset versions are recorded for auditability. The conversational AI assistant can explain these results using role-scoped platform data but does not calculate or alter the underlying scores.
- **Profile based on verified evidence.** Students build evidence-backed profiles using skills, certifications, projects, experience, and uploaded documents. AI-assisted document analysis can extract structured information and supporting evidence, but AI extraction alone never makes evidence verified. Human approval is required before submitted evidence becomes trusted evidence within Fursah.
- **Institutional analytics.** Universities get insights about the alignment of the curriculum to the labor market, graduate tracking, and accreditation-compliant reporting from the same pipeline used for serving students.
- **Fairness by design and governance monitoring.** Fursah does not use protected or non-job-related attributes such as gender, nationality, age, or GPA in readiness and matching. Deterministic scoring rules are versioned and auditable, while governance monitoring, appeals, audit events, and human overrides provide mechanisms for identifying and responding to potentially harmful outcomes.
- **Human oversight by design.** AI assists with evidence understanding and explanation, while students, employers, universities, and administrators retain authority over consequential decisions.

3. Mapped documents

The nodes of Fursah's AI and intelligence pipeline are linked to the policies, strategies, and standards governing them. Nodes: SRC (Sources), C (Connectivity), PP (Pre-processing), M (Models and Intelligence), P (People), D (Data and Analytics Platform), and SINK (Interfaces).

Node	Fursah component (workforce-readiness use case)	Mapped documents (policies, strategies, standards)
SRC	Data Sources: Student profiles, skills, certifications, experiences, projects, career interests, applications, and uploaded evidence; employer job requirements, job descriptions, applications, and hiring feedback; university offerings and curriculum-action evidence; and platform interaction signals used to personalize recommendations.	PDPL and the Implementing Regulations (lawful basis, consent, purpose limitation); NDSG Data Classification Policy of SDAIA; MoE / ETEC Data Sharing Procedures and Ownership of Students Records; National Occupational Standards Framework for Data & AI (skill taxonomy); Saudi Academic Framework for Artificial Intelligence (programme - competencies mapping).
C	Connectivity and Integration Layer: Fursah is implemented as a web platform with role-based student, employer, university, and administrator interfaces. Application data is managed through Prisma and libSQL/Turso, private evidence is stored in Cloudflare R2, and secured Cloudflare services support AI-assisted evidence analysis and conversational assistance.	DGA standards for digital government interoperability and integration; CST regulatory regime for cloud services and telecommunications; NCA Essential Cybersecurity Controls and Cloud Cybersecurity Controls; platform availability/reliability SLAs; data localization requirements.
PP	Pre-processing: Uploaded evidence is validated, privately stored, and, for supported document types, analyzed using AI to extract structured information such as certification details, project technologies, experience responsibilities, job requirements, and university offering information. Extracted information remains advisory until human verification. Skills and requirements are normalized against Fursah's taxonomy before deterministic readiness, matching, gap-analysis, and recommendation calculations.	PDPL requirements for data minimisation, purpose limitation, and protection of personal data; SDAIA privacy-by-design guidance; requirements for removing unnecessary personal information; and requirements for secure data retention and deletion.
M	Models and Intelligence: Fursah separates generative AI from consequential scoring. AI models are used for structured evidence extraction and the grounded Fursah Assistant. Career Readiness Scores, gap analysis, career alignment, candidate-job matching, employer intelligence, university curriculum analysis, and workforce	Principles and Controls of AI Ethics (SDAIA – current): seven principles and four tiers of risk classification throughout the entire AI lifecycle, without any exception to discretionary waiver for high-risk systems; AI adoption framework (SDAIA);

	intelligence are calculated using deterministic, versioned rules. The assistant receives role-scoped platform facts and explains existing results conversationally; it does not calculate or alter readiness scores, rankings, evidence-verification decisions, or workforce metrics.	AI management systems (ISO/IEC 42001) and AI risk (ISO/IEC 23894); benchmarking models, bias audits and fairness tests; EU AI Act (Annex III point 4 – employment high risk).
P	People/Human Oversight: AI-extracted evidence requires human verification before it becomes trusted evidence. Employers retain responsibility for hiring decisions, universities decide whether to implement curriculum recommendations, and students retain control over their declared career direction. Students may dispute relevant platform decisions through Fursah's appeal mechanisms, while administrator actions and overrides are auditable.	Human oversight principles for SDAIA; Rights of the data subject in the PDPL (rights of access, right to correct, right to object to decision based solely on automated processing); Override and Escalation criteria for human oversight; Grievance procedure for institutional appeals; Guidance from the ETEC on assessment evidence.
D	Data and Analytics Platform: Fursah aggregates role-appropriate platform data to generate student, employer, university, and workforce intelligence while applying access controls and privacy-preserving cohort suppression. The prototype uses Vercel for application hosting, Prisma with libSQL/Turso for application data, Cloudflare R2 for private evidence storage, and Cloudflare services for AI inference.	PDPL requirements for data protection, purpose limitation, data-subject rights, and cross-border processing; NCA cybersecurity and cloud controls; CST cloud regulations; and Fursah's Privacy and Responsible AI policies. Production deployment requires institutional review of hosting regions, processor arrangements, cross-border transfers, and applicable Saudi data-localization requirements.
SINK	Interfaces: Student workspace for readiness, adaptive roadmaps, career exploration, job matching, and evidence management; employer portal for jobs, candidate fit, and hiring intelligence; university dashboard for student readiness, workforce demand, curriculum alignment, and initiatives; administrator interfaces for evidence review and governance; shared workforce-intelligence analytics; and role-aware AI assistance.	DGA digital accessibility guidelines and WCAG 2.1 AA; Arabic-first and bilingual user interface specifications; consent management and transparency notice rules; Fursah Responsible AI, Privacy and Accessibility policies publicly posted on the platform.

Apart from data and AI algorithms, the map also takes into account three additional domains. Strategy alignment: Fursah facilitates the process of comprehending the labor market's demands in terms of skills by both universities and organizations through the language of skills. Interface design and accessibility: The platform uses inclusive, Arabic-compatible interfaces as well as empowers the user with full control of his or her data. Human-AI cooperation: AI assists with evidence understanding and explanation, deterministic intelligence generates transparent recommendations and scores, and humans retain authority over consequential decisions.




Figure 1. Fursah AI and Intelligence Architecture. Generative AI is bounded to evidence understanding and grounded assistance, while deterministic, versioned intelligence produces consequential scores and recommendations under human oversight.

4. Evaluation scenarios

The scenarios below test the solution where it works and, more importantly, where it becomes controversial. In each case the response is governed by pre-agreed rules held in the knowledge base rather than decided ad hoc.

Scenario 1: Baseline adoption (target post-deployment outcome)

In a future multi-university deployment, students would use gap analysis before graduation, employers would use explainable matching to support recruitment, and universities would use aggregated demand signals during program review. Post-deployment success would be evaluated using placement rate, time-to-hire, gap-closure rate, explanation quality, dispute rates, and human-override statistics.

Scenario 2: Commercial capture of the ranking

Six months into implementation, the paid visibility for employer branding and paid courses become part of the gap analysis, leading to some of the recommendations being commercially driven. Dealing with the problem: The sponsored items should be clearly labeled, separated from the matching score and recommendations, and the ranking system should remain auditable. Otherwise, it will be suspended according to Fursah's Responsible AI policy.

Scenario 3: Targeted advertising on inferred attributes (escalation)

Inferred profile signals such as lack of readiness, socioeconomic indicators, or accommodations for people with disabilities, among others, are being utilized to serve advertisements, or have been sold to third parties. It is a move from being an issue about consumers to becoming one of data privacy and ethical concerns. The policies regarding the knowledge base give rise to higher standards: complete ban on any further utilization of the inferred signals for advertising or profiling, implementation of the principles of consent and purpose limitation for student data, contractual obligations for third-party processors, removal of data sharing, and escalation to SDAIA and the sectoral regulator.

Scenario 4: Potential disparate impact identified during governance review

Although protected characteristics are not used in Fursah's readiness or matching calculations, governance review may identify systematic differences in outcomes between populations when institutions lawfully conduct fairness evaluation using separately governed data. If unexplained disparities are identified, the relevant scoring rules, evidence requirements, thresholds, or potential proxy factors should be investigated before continued use. Findings and corrective actions should be documented, with human review and suspension available where necessary. Employment-related uses of Fursah should therefore be treated with heightened governance and human-oversight requirements.

Scenario 5: Automation bias among human reviewers

Mitigation measures include monitoring override and dispute rates, investigating unexpectedly low override rates, requiring written justification for selected decisions, training reviewers to treat Fursah's explanations as decision support rather than authority, and periodically validating scoring rules and recommendation patterns against observed employment outcomes.

Scenario 6: The non-standard candidate

Response: alternative evidence pathways, including portfolio evidence, practical work, and employer endorsements, can contribute to the same skills-based profile; a low score should be presented as an indication of incomplete evidence or identified readiness gaps rather than a definitive judgment of competence; and the contest and appeal process allows human review and override. Overrides are recorded in the audit trail and can inform future human-reviewed changes to Fursah's taxonomy, scoring rules, or evidence pathways.

5. Knowledge base (sources)

AI ethics and governance

- SDAIA: Principles and Controls of AI Ethics: seven principles (fairness; privacy and security; humanity; social and environmental benefits; reliability and safety; transparency and explainability; accountability and responsibility) and a four-tier risk classification (little/no, limited, high, unacceptable). <https://sdaia.gov.sa/en/SDAIA/about/Documents/ai-principles.pdf>
- SDAIA: AI Adoption Framework. <https://sdaia.gov.sa/en/SDAIA/about/Files/AIAdoptionFramework.pdf>
- SDAIA: Generative AI Guidelines for Government Entities (January 2024). <https://sdaia.gov.sa/en/SDAIA/about/Pages/RegulationsAndPolicies.aspx>

Data protection and governance

- Personal Data Protection Law (Royal Decree M/19, amended by M/148) and its Implementing Regulations, in force 14 September 2023: consent, purpose limitation, automated processing and profiling, data-subject rights, cross-border transfer. <https://sdaia.gov.sa/en/SDAIA/about/Pages/RegulationsAndPolicies.aspx>
- SDAIA/NDMO: National Data Governance Policies and Data Classification Policy. <https://sdaia.gov.sa/ar/SDAIA/about/Documents/Data%20Classification%20Policy.pdf>

Skills and education frameworks

- SDAIA: National Occupational Standard Framework for Data and Artificial Intelligence: five classifications, ten specialty areas and sixteen occupations, forming the basis of Fursah's skill taxonomy. <https://sdaia.gov.sa/en/Research/Documents/PSFDataAndAIBook.pdf>
- SDAIA: Saudi Academic Framework for AI Qualifications: reference for higher-education programme development and accreditation. <https://sdaia.gov.sa/en/Research/Documents/SaudiAcademicFramework.pdf>
- Vision 2030: Human Capability Development Programme, Delivery Plan 2021–2025: workforce-readiness pillar, education–labour-market alignment, national occupational standards and certification pathways. <https://www.vision2030.gov.sa/en/vision-2030/vrp/human-capability-development-program/>

Cybersecurity, cloud and interoperability

- National Cybersecurity Authority: Essential Cybersecurity Controls. <https://nca.gov.sa/en/regulatory-documents/controls-list/ecc/>
- National Cybersecurity Authority: Cloud Cybersecurity Controls (CCC–1:2020). <https://nca.gov.sa/en/regulatory-documents/controls-list/ccc/>
- CST: Cloud Computing Services Provisioning Regulations (Version 4, in force 10 October 2023): provider registration categories keyed to subscriber data classification. <https://www.cst.gov.sa/en/regulations-and-licenses/regulations/Document-1550>
- DGA: Guideline for Web Accessibility of Government Websites (WCAG 2.1 Level AA). https://dga.gov.sa/en/Web_Accessibility_of_Government_Websites

International interoperability

- ISO/IEC 42001:2023: AI management systems. <https://www.iso.org/standard/42001>
- ISO/IEC 23894:2023: AI risk management guidance. <https://www.iso.org/standard/77304.html>
- EU AI Act (Regulation 2024/1689), Annex III point 4: employment, worker management and access to self-employment as a high-risk category; high-risk obligations enforceable from 2 December 2027. <https://artificialintelligenceact.eu/annex/3/>

Platform policies

- Fursah: Responsible AI, Privacy Policy, Accessibility. <https://fursah.org/policies>