

Verifying Recruitment Offers for Nepali Migrant Workers: A Provenance-Aware, Bilingual Decision Support System

Rajan Raj Acharya
Nepal College of Information Technology
(Affiliated to Pokhara University)
Lalitpur, Nepal
rajan.raj@ieee.org

Krishma Paudel
Nepal College of Information Technology
(Affiliated to Pokhara University)
Lalitpur, Nepal
krishma.241615@ncit.edu.np

Apil Khadka
Nepal College of Information Technology
(Affiliated to Pokhara University)
Lalitpur, Nepal
apil.231210@ncit.edu.np,

Manil Baidhya
Nepal College of Information Technology
(Affiliated to Pokhara University)
Lalitpur, Nepal
manil.baidhya@ncit.edu.np

Abstract—Prospective migrant workers in Nepal must often decide whether to pay using fragmented bilingual advertisements and intermittently available official records. We present Parakh, a provenance-aware bilingual decision-support layer that combines Department of Foreign Employment records, declarative red-flag rules, OCR-assisted field extraction, Media and Information Literacy questions, and grievance escalation. It exposes source state and missing evidence rather than issuing an official certificate or autonomous fraud verdict. On 216 synthetic advertisements, browser Tesseract, CLI Tesseract, and clean parsing achieved macro field exact match of 80.3%, 82.2%, and 98.8%; complete-record exact match was 22.2%, 23.1%, and 93.1%. Browser agency-name exact match was only 21.6%, while browser lot-number exact match was 86.0%, with 0.0% observed hallucination for all methods. A preliminary browser check on 43 scraped images found correct lot-number extraction in 38/43 (88.4%) and destination extraction in 35/43 (81.4%). These measurements provide reproducible software-level and collection-level baselines, not independently audited natural-ad accuracy or evidence of improved user decisions.

Keywords—migrant-worker protection, recruitment fraud, media and information literacy, explainable decision support, offline-resilient systems, Nepal, progressive, web application

I. INTRODUCTION

Consider a worker who receives a WhatsApp poster promising a UAE job for NPR 70,000 and requesting payment to a personal wallet. Before paying, the worker must determine whether the agency is licensed, whether the job-demand lot is genuine, whether the fee is lawful, whether the visa supports employment, and whether the recruiter is authorized. These checks are distributed across official registries, documents, advertisements, and informal conversations. A registry lookup may time out, an image parser may miss the fee, and a binary warning may hide the evidence that the worker needs to act safely.

This setting is especially important in Nepal. Migration and remittances remain central to the country's economic and policy context [1]. Research on recruitment cost, fraud, and redressal documents excessive fees, deception about jobs and salaries, and barriers to seeking justice among migrant

households [2]. Recent ILO reporting indicates that recruitment costs remain a concern for Nepali workers and calls for stronger transparency and grievance mechanisms [3]. Regional migration reporting emphasizes safe, orderly, and informed migration governance [4]. UNODC documents links among casinos, cyber fraud, trafficking, and organized crime in Southeast Asia, while IOM reports trafficking into forced criminality in online scamming centres [5], [6]. The practical gap is that authoritative information may be available without being usable at the moment a worker is asked to pay.

We present Parakh, a pre-departure verification layer for prospective Nepali migrant workers and their families. Its central idea is to treat evidence provenance and uncertainty as part of the user-facing verdict rather than as hidden engineering metadata. The contribution is the integration of source context, explicit checks, and bilingual next actions in one recruitment-verification workflow; no new OCR model or fraud classifier is proposed. Parakh combines official-record access, a configuration-driven engine for statutory and contextual red flags, graceful degradation across network, database, and OCR failures, and bilingual explanations aligned with Media and Information Literacy (MIL). UNESCO describes MIL as the ability to find, evaluate, and produce information and media content [7]; Parakh turns that principle into a sequence of checks and questions at the point of payment.

Our research question is: how can a lightweight, bilingual integration layer turn official recruitment records and explicit checks into actionable pre-departure verification under intermittent connectivity? The Department of Foreign Employment and its Foreign Employment Information Management System provide the authority for agency and job-demand records, while the Foreign Employment Act provides the legal context for rule maintenance [8], [9]. Parakh is an access, translation, evaluation, provenance, and escalation layer around these sources, not an independent license attestation.

The paper makes three bounded, testable contributions:

- a shared verification architecture that combines registry

evidence, lot checks, declarative rules, OCR-assisted input, and human escalation (Section III);

- an inspectable evidence contract in which each warning links to a rule identifier, legal or policy basis, bilingual explanation, and next question (Section III); and
- a reproducible validation protocol, including the measured synthetic OCR benchmark reported in Section VI, that separates verified software behavior from unmeasured user outcomes.

The system reports only what it can check from supplied inputs and the available source. It is not an official certificate, legal adjudicator, or black-box scam classifier; unresolved cases are labeled and directed to official or civil-society contacts.

II. DESIGN GOALS

The design was guided by six goals.

- 1) **Evidence before labels.** A verdict should be derived from identifiable fields, registry records, and named rules. A missing record should result in “unverified,” not an invented assertion.
- 2) **Graceful degradation.** Loss of a government endpoint, optional OCR provider, or network connection should produce an explicitly weaker mode, not a silent failure or an apparently authoritative answer.
- 3) **Actionable explanations.** Each triggered rule should provide a reason, a legal or policy basis, a bilingual explanation, and a question or next step.
- 4) **Authority-aware integration.** The government registry remains the authority. The system must expose the source, freshness, failure mode, and limits of every result.
- 5) **Low-friction access.** The same core logic should be reachable from a responsive web application, a WhatsApp conversation, and an eventual desktop shell, without requiring a paid account or a heavy client installation.
- 6) **Safety and accountability.** The interface should avoid defamatory claims, preserve source attribution, minimize personal data, and direct unresolved cases to official channels or counselors.

III. SYSTEM ARCHITECTURE

Figure 1 shows the system flow. Inputs arrive as structured fields, an advertisement image or text, an agency query, or a pre-permission lot number. The application normalizes these inputs and passes them to the same verification services. The result contains a status, color-coded urgency, triggered rules, official record when available, MIL questions, provenance metadata, and contacts.

A. Source synchronization and provenance

The pipeline normalizes public Department of Foreign Employment (DoFE) agency and pre-permission records into a JSON snapshot mirrored to Cloudflare D1, with local SQLite/libSQL fallback. A successful request is marked `dofe_live`; an empty or timed-out request uses a curated `mock_fallback`. The validated live snapshot contains 1,769 agency and 33,687 lot records. The application carries source state and snapshot age into results, distinguishing verification

against the available record from a real-time government response; a `dofe_live` label records the last successful retrieval and does not guarantee freshness, while fallback data cannot support live-corpus coverage.

B. Government dependency as a trust boundary

The registry remains authoritative for agency and lot records; Parakh records source state, synchronization time, data counts, failure mode, and rule set alongside its interpretation. Snapshot metadata is not itself a rule predicate. A missing record is not proof of unlicensed activity, and demonstration data cannot establish official verification. Source and legal changes therefore require an accountable maintainer and visibly weaker results during upstream failure.

C. Offline-resilient operation

The current implementation has three limited continuity mechanisms. The service worker attempts to cache a small web shell after an online visit. Client-side handlers use a small fallback seed for selected demonstration flows when an API request fails. The server-side pipeline can use a local snapshot when the live DoFE endpoints fail. These mechanisms preserve some functionality under intermittent connectivity, but they do not constitute a complete local-first architecture. The browser does not currently maintain a synchronized FEIMS database in local storage, and the service worker does not yet guarantee that every JavaScript asset is available offline. We therefore describe Parakh as offline-resilient and state full local-first replication as a current scope limitation.

D. Declarative red-flag engine

The core engine consumes normalized offer fields and, when available, an agency record. Rules are stored as JSON objects containing an identifier, severity, category, destination or field conditions, legal basis, bilingual warnings, and bilingual MIL questions. The current rule family covers fee-cap violations for selected destinations, personal payment channels, work on visit or tourist visas, an unauthorized Korea E-7 condition without a verified lot, suspended or blocked agencies, recruiter-phone mismatch, and unlicensed agencies.

The engine produces a structured result rather than a scalar risk score. A clean active-agency check can yield `VERIFIED_ACTIVE`; this means the supplied agency record is active, not that the offer is genuine or complete. Missing offer fields do not necessarily prevent this agency-level status. All triggered rules are retained: a critical rule other than agency no-match produces `HIGH_RISK_CRITICAL`, while other warnings produce `CAUTION_REQUIRED`; agency no-match alone yields `UNVERIFIED`. Multiple high-severity rules do not automatically become critical. For example, the configured UAE fee check flags NPR 70,000 as high severity; adding a personal-wallet payment triggers a critical warning while retaining the fee warning. Each triggered flag carries its rule identifier, category, legal basis, and localized warning. This design enables unit tests, configuration review, and future rule updates without changing presentation code. Predicates and some thresholds remain in TypeScript, so not every rule change can be made through JSON alone.




Fig. 1. Parakh architecture. All user surfaces converge on the same verification path. Source provenance and failure mode are carried into the structured result.

E. OCR and input normalization

Recruitment offers may arrive as screenshots or photographs. The client resizes large images, reuses an on-device Tesseract.js worker, and parses the transcript locally into fields such as agency name, destination, fee, visa type, payment channel, phone number, and lot number. Only when device OCR cannot produce text does the client send the resized image to a local server-side Tesseract CLI fallback with Nepali language data; no OCR image is sent to an AI provider. OCR is treated as a pre-fill aid: users can inspect and correct fields before verification, and missing fields do not become positive evidence.

F. Bilingual and conversational surfaces

The web application persists an English/Nepali language choice and renders the verdict, warnings, questions, and contacts in both languages. The WhatsApp webhook maps short commands to agency, lot, and offer verification. Both surfaces use the same server functions to reduce semantic drift between channels. The “Ask Before You Pay” questions are deliberately included in every verdict, including a clean or incomplete result, because a user should retain a verification habit even when no rule fires.

G. Escalation and reports

The system includes public government and civil-society contacts and generates a structured grievance package containing the supplied details, evidence fields, and data provenance. Reports state that they are verification aids rather than official certificates. This is intended to reduce the risk of a green or amber interface state being interpreted as a legal finding and provides a path when the system cannot resolve the case.

IV. RELATED WORK AND DESIGN POSITIONING

A. Recruitment risk and worker information

Fair-recruitment research emphasizes transparent information, recruitment-cost control, remedy, and prevention of abusive practices [2], [10]. Government registries provide authoritative agency and job-demand records within the digital public infrastructure described by OECD [11]. By contrast, Vidros *et al.* predict recruitment fraud from the text and metadata of 17,880 online job advertisements [12]. Parakh makes a narrower integration claim: it combines registry evidence, explicit rules, source provenance, missing-evidence handling, bilingual explanations, and escalation in one pre-payment workflow. It introduces neither a new OCR model nor a fraud classifier; the benchmark therefore reports extraction and rule-replay behavior, not fraud-classification accuracy.

B. Information literacy and human-in-the-loop security

MIL connects the technical problem to the user’s ability to evaluate information. UNESCO’s policy guidelines treat MIL as a development capability that joins information literacy, media literacy, access to information, and freedom of expression [7]. The implication for a verification tool is that a warning should teach a repeatable check: Which official record supports this result? What information is missing? What should the worker ask next?

Security systems that depend on human decisions also need to account for user capability, motivation, and context. Cranor’s human-in-the-loop framework argues that systems should be designed to support people when a human security decision cannot be removed from the process [13]. Parakh adopts this position by exposing evidence and next actions instead of delegating a final decision to an opaque score.

C. Interpretability and resilient web systems

Research on application-grounded interpretability and responsible AI emphasizes traceability, governance, and user context [14], [15], [16]. Parakh applies deterministic rules rather than explaining a learned classifier, while W3C Service Workers provide one primitive for limited offline continuity [17]; the remaining question is whether users understand the resulting source labels and next actions.

V. IMPLEMENTATION AND CURRENT EVIDENCE

The prototype is a single Next.js application with route handlers, a local TypeScript rules engine, Cloudflare D1 with SQLite/libSQL fallback, a Tauri 2 shell, and an optional WhatsApp webhook [18]. It runs without paid OCR or AI APIs; server Tesseract, WhatsApp delivery, and hosted storage are optional. The validated snapshot contains 1,769 agencies and 33,687 lots; the OCR results below are benchmark measurements, not population estimates.

VI. EVALUATION

A. Completed image-to-field evaluation

We generated 216 synthetic recruitment advertisements: two images for each cell of a $4 \times 3 \times 3 \times 3$ design spanning four script categories, three resolutions, three capture types, and three layouts. Each image has a manifest row, hash, clean transcript, and gold labels for seven fields; missing values are ABSENT. The generator and schema are versioned in `paper/ocr_eval`, and no private image was used.

We compared production browser Tesseract.js 7.0.0 (local `eng/nep`, Playwright 1.62.1), local Tesseract CLI 5.5.3 (`eng+nep`, PSM 11), and a deterministic parser on clean transcripts. All used the same normalization and scorer; the third is a parser comparison, not an image-OCR estimate.

Precision is exact normalized match among predicted-present values; recall and exact match use gold-present values; abstention is an empty prediction for a gold-present field; hallucination is a non-empty prediction for gold-ABSENT. The seven-field complete-record metric includes absence status. Re-scoring corrected an earlier implementation that counted correct absent/absent pairs as errors; field-level exact-match counts were unchanged. Complete-record exact match is 48/216 (22.2%), 50/216 (23.1%), and 201/216 (93.1%) for browser, CLI, and clean parsing. Hallucination uses 30 or 31 gold-absent instances per field; zero observed hallucinations does not imply perfect present-field extraction.

Browser field exact match ranged from 21.6% for agency name to 100.0% for phone number; fee, visa, and lot exact match were 75.7%, 99.5%, and 86.0%. All methods had 0.0% observed hallucination; browser abstention was highest for agency name (15.1%), fee (23.8%), and lot (9.1%). Clean parsing ranged from 91.9% to 100.0%, locating most remaining loss in image OCR. Tables I and II report the detailed results.

In the measured condition breakdown, Devanagari was weakest, while English and mixed-script inputs performed better; differences across resolution, capture type, and layout

TABLE I

OVERALL SYNTHETIC-IMAGE EXTRACTION RESULTS BY METHOD ($n = 216$). EM IS EXACT MATCH AMONG GOLD-PRESENT FIELDS; H IS HALLUCINATION RATE ON GOLD-ABSENT FIELDS.

Method	EM range	Macro EM	Max H	Complete EM
Browser Tesseract	21.6–100.0%	80.3%	0.0%	22.2%
CLI Tesseract	25.4–100.0%	82.2%	0.0%	23.1%
Clean parser	91.9–100.0%	98.8%	0.0%	93.1%

were smaller. Full condition-level results are retained in the reproducibility artifacts.

B. Controlled OCR-to-verification evaluation

We exercised the public OCR and offer-verification handoff on six controlled cases covering matched lots, ambiguous fields, fee/payment warnings, lot mismatch, server OCR, and unavailable OCR; all expected outcomes passed. Nine additional offer-verification cases covered stale/demo sources, incorrect identifiers, unavailable lookup, and no-match. These controlled transcripts, mocked responses, and the local FEIMS snapshot establish seam behavior, not natural-ad accuracy, live-government reliability, or fraud-detection accuracy; the expected-versus-actual records are retained with the repository.

C. Extraction-to-rule replay

We replayed all 216 saved predictions through the rebuilt rules engine without registry lookup. Browser retained 327/369 expected warnings (113/154 fee warnings), matched gold status on 203 images and the complete rule set on 175; CLI retained 327 warnings, matched status on 202 and rule sets on 174; clean parsing matched all 216 statuses and rule sets and retained all 369 warnings. These are descriptive propagation measures, not legal, fraud, live-government, or full-system accuracy; hashes and replay evidence are in `evidence/rule-replay.json`.

D. Worked failure and safety interpretation

In `syn-0001`, gold destination UAE and fee NPR 70,000 trigger `CAUTION_REQUIRED` and `RULE_FEE_CAP_EXCEEDED` (configured cap NPR 10,000). Browser OCR malformed the fee, so the parser abstained and the result became `UNVERIFIED` without that warning. This demonstrates why users must confirm extracted high-impact fields: abstention is safer than invention but does not resolve missing evidence.

E. Preliminary browser check on natural images

Separately, a preliminary browser check of 43 scraped recruitment images correctly extracted lot numbers in 38/43 (88.4%) and destinations in 35/43 (81.4%); employer names were reported present and extracted in 31. These are collection-level tallies, not independently audited full-system accuracy; denominators were not restricted to readable, gold-present fields, and source absence, illegibility, OCR omission, and incorrect values were not separated. Predictions for absent fields and per-image joint correctness were unavailable; provenance, rights, duplicate screening, and saved outputs require audit. The employer field is distinct from agency name; no natural-image

TABLE II

PER-FIELD SYNTHETIC-IMAGE RESULTS BY METHOD. P, R, EM, A, AND H ARE PRECISION, RECALL, EXACT MATCH, ABSTENTION RATE, AND HALLUCINATION RATE (PERCENT).

Field	Method	P	R	EM	A	H
Agency name	Browser Tesseract	25.5%	21.6%	21.6%	15.1%	0.0%
	CLI Tesseract	28.3%	25.4%	25.4%	10.3%	0.0%
	Clean parser	100.0%	100.0%	100.0%	0.0%	0.0%
Destination	Browser Tesseract	100.0%	87.6%	87.6%	12.4%	0.0%
	CLI Tesseract	100.0%	87.6%	87.6%	12.4%	0.0%
	Clean parser	100.0%	100.0%	100.0%	0.0%	0.0%
Fee	Browser Tesseract	99.3%	75.7%	75.7%	23.8%	0.0%
	CLI Tesseract	100.0%	76.2%	76.2%	23.8%	0.0%
	Clean parser	100.0%	100.0%	100.0%	0.0%	0.0%
Visa type	Browser Tesseract	100.0%	99.5%	99.5%	0.5%	0.0%
	CLI Tesseract	100.0%	100.0%	100.0%	0.0%	0.0%
	Clean parser	100.0%	100.0%	100.0%	0.0%	0.0%
Payment channel	Browser Tesseract	100.0%	91.9%	91.9%	8.1%	0.0%
	CLI Tesseract	100.0%	91.9%	91.9%	8.1%	0.0%
	Clean parser	100.0%	91.9%	91.9%	8.1%	0.0%
Phone number	Browser Tesseract	100.0%	100.0%	100.0%	0.0%	0.0%
	CLI Tesseract	100.0%	100.0%	100.0%	0.0%	0.0%
	Clean parser	100.0%	100.0%	100.0%	0.0%	0.0%
Lot number	Browser Tesseract	94.7%	86.0%	86.0%	9.1%	0.0%
	CLI Tesseract	98.9%	94.6%	94.6%	4.3%	0.0%
	Clean parser	100.0%	100.0%	100.0%	0.0%	0.0%

CLI comparison was reported, and these observations do not establish that an offer is genuine.

F. Remaining natural-ad and user evaluations

An independently annotated evaluation of the full system on permitted natural advertisements and a consented bilingual user study remain outstanding. The planned study would recruit prospective workers, family members, or trained counselors through an approved partner and measure source-state comprehension, safe-action identification, confidence calibration, completion time, and error types in Nepali and English; ethics approval, recruitment, consent, and translated instruments are prerequisites.

VII. ETHICS, LIMITATIONS, AND THREATS TO VALIDITY

Source, rules, and safety. Snapshots can become stale, outages limit coverage, and the fallback dataset is intentionally small; provenance reduces but does not remove these risks. Statutory thresholds and procedures change, so deployment needs an accountable policy owner, source dates, review, and regression tests. A phone mismatch or OCR omission can be benign or harmful; a clean result means only that supplied fields triggered no encoded rule, not that an offer is safe.

Data handling. Server Tesseract writes a temporary image and removes its directory; optional OCR evidence is opt-in and pruned by configured retention. Grievance generation persists supplied contact and incident details; assistant sessions may retain name/contact fields; analytics chat logging is disabled by default and, if enabled, bounded and redacts common phone and email patterns. WhatsApp logs event counts rather than raw content. Assistant, analytics, grievance, and stored-OCR records have configurable retention and an authenticated maintenance endpoint, but deployment-level access, transport/database encryption, backups, infrastructure logs, and external-provider retention remain unverified.

Privacy and human subjects. The reproducible benchmark used synthetic advertisements; the 43-image tallies are separate preliminary observations, not pooled scores. No private image was sent to an external service without permission. Recruitment documents may contain names, phone numbers, passport identifiers, and financial information, so raw uploads are not retained by default and optional provider processing requires explicit configuration. Any usability or counselor study requires institutional ethics review, informed consent, secure storage, and a response plan for distress or active fraud disclosures.

External validity. External validity concerns whether the observed results generalize beyond the evaluated images. The 43-image browser check provides preliminary real-image evidence for lot-number and destination extraction, but does not establish performance on advertisements from other sources, layouts, or image qualities, or extraction of fields absent from this collection. Adding these observations addresses the previous lack of real-image evidence without resolving generalization to broader advertisement collections or complete-offer verification. Separately, the controlled handoff tests assess specified software behaviors on fixtures; their inability to estimate natural-ad OCR performance or live-government reliability should not be confused with the evidence provided by the 43-image browser check. The synthetic corpus also uses a small vocabulary and generated layouts; its agency normalization recognizes four synthetic name families, so exact match is not literal transcription or registry-resolution accuracy. No participant study has been completed. Transfer beyond Nepal requires further validation because legal rules, registries, languages, and grievance channels differ; modular rule and contact files enable adaptation but do not establish transferability.

VIII. DISCUSSION

Parakh treats source state as part of the evidence. Carrying live, snapshot, and fallback labels through synchronization, verification, and reports makes uncertainty actionable and

reduces the risk that a useful but unlabeled fallback appears official.

The output is an action-oriented MIL aid: each warning couples a short explanation and legal basis with questions the worker can ask before paying. This supports human-in-the-loop security [13] without claiming that comprehension or safer decisions have been demonstrated; that requires the uncompleted bilingual study.

Finally, deterministic rules remain separate from replaceable OCR. This isolates extraction errors, avoids silent substitution for missing official evidence, and permits future comparisons with learned extraction or counselor review under the same evidence contract.

IX. CONCLUSION

Parakh combines official-record access, explicit source state, declarative rules, soft-fail OCR, MIL questions, and grievance escalation in one pre-payment workflow. The synthetic benchmark reports macro and complete-record extraction baselines; controlled tests exercise incomplete fields, warning combinations, lot mismatches, and unavailable evidence; and the preliminary 43-image check reports descriptive lot, destination, and employer observations.

The results do not establish full natural-ad accuracy, live-government reliability, offer genuineness, comprehension, or improved user decisions. Independently annotated natural advertisements and a consented bilingual user evaluation remain necessary.

ACKNOWLEDGMENT

The authors thank the open-source and public-interest communities whose documentation and public records informed the prototype.

REFERENCES

- [1] World Bank, "Nepal development update," World Bank, 2026, accessed: 2026-08-10. [Online]. Available: <https://www.worldbank.org/en/country/nepal/publication/nepaldevelopmentupdate>
- [2] A. Kharel, S. Bhattarai, and D. Tumsa, "Recruitment cost, fraud and redressal in foreign labour migration from nepal," Centre for the Study of Labour and Mobility, Policy Brief 9, 2023, based on a sub-national survey of 1,225 households, accessed: 2026-08-10. [Online]. Available: <https://ceslam.org/wp-content/uploads/2024/11/Policy-Brief-9.pdf>
- [3] International Labour Organization, "The cost of labour migration in nepal: Evidence from the 2023 recruitment cost survey of migrant workers," ILO, 2026, national Statistics Office of Nepal survey with ILO technical support, accessed: 2026-09-14. [Online]. Available: <https://www.ilo.org/publications/cost-labour-migration-nepal-evidence-2023-recruitment-cost-survey-migrant>
- [4] United Nations Economic and Social Commission for Asia and the Pacific, "Asia-pacific migration report 2024," United Nations, Tech. Rep., 2024, accessed: 2026-08-10. [Online]. Available: <https://hdl.handle.net/20.500.12870/7548>
- [5] United Nations Office on Drugs and Crime, "Casinos, cyber fraud, and trafficking in persons for forced criminality in southeast asia," UNODC Regional Office for Southeast Asia and the Pacific, Tech. Rep., 2023, policy report, accessed: 2026-09-14. [Online]. Available: https://www.unodc.org/roseap/uploads/documents/Publications/2023/Tip_for_FC_Policy_Report.pdf
- [6] International Organization for Migration, "Iom's regional situation report on trafficking in persons into forced criminality in online scamming centres in southeast asia," IOM Regional Office for Asia and the Pacific, 2024, official regional situation report, accessed: 2026-09-14. [Online]. Available: https://roasiapacific.iom.int/sites/g/files/tmzbd1671/files/documents/2024-07/iom-southeast-asia-trafficking-for-forced-criminality-update_december-2023_-final.pdf
- [7] A. Grizzle, P. Moore, M. Dezuanni, S. Asthana, C. Wilson, F. Banda, and C. Onumah, "Media and information literacy: Policy and strategy guidelines," UNESCO, 2013, accessed: 2026-08-10. [Online]. Available: <https://www.unesco.org/en/articles/media-and-information-literacy>
- [8] Department of Foreign Employment, Government of Nepal, "Foreign employment information management system," Government of Nepal, 2026, accessed: 2026-08-10. [Online]. Available: <https://feims.dofe.gov.np/Common-Login>
- [9] Nepal Law Commission, "Foreign employment act, 2064," Government of Nepal, 2007, official legal source in Nepali, accessed: 2026-08-10. [Online]. Available: <https://lawcommission.gov.np/content/13307/foreign-employment-act--2064/>
- [10] International Labour Organization, "General principles and operational guidelines for fair recruitment and definition of recruitment fees and related costs," International Labour Organization, 2019, accessed: 2026-08-10. [Online]. Available: <https://www.ilo.org/publications/general-principles-and-operational-guidelines-fair-recruitment-and>
- [11] Organisation for Economic Co-operation and Development, "Digital public infrastructure for digital governments," OECD Publishing, Tech. Rep., 2024. [Online]. Available: <https://doi.org/10.1787/f525dc8-en>
- [12] S. Vidros, C. Kolias, G. Kambourakis, and L. Akoglu, "Automatic detection of online recruitment frauds: Characteristics, methods, and a public dataset," *Future Internet*, vol. 9, no. 1, p. 6, 2017. [Online]. Available: <https://www.mdpi.com/1999-5903/9/1/6>
- [13] L. F. Cranor, "A framework for reasoning about the human in the loop," in *Usability, Psychology, and Security 2008 (UPSEC 08)*. San Francisco, CA: USENIX Association, 2008. [Online]. Available: <https://www.usenix.org/conference/upsec-08/framework-reasoning-about-human-loop>
- [14] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv preprint arXiv:1702.08608*, 2017. [Online]. Available: <https://arxiv.org/abs/1702.08608>
- [15] A. B. Arrieta, N. Díaz-Rodríguez, J. D. Ser, A. Bannetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, "Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai," *Information Fusion*, vol. 58, pp. 82–115, 2020. [Online]. Available: <https://arxiv.org/abs/1910.10045>
- [16] E. Tabassi, "Artificial intelligence risk management framework (ai rmf 1.0)," National Institute of Standards and Technology, Tech. Rep. NIST AI 100-1, 2023. [Online]. Available: <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10>
- [17] World Wide Web Consortium, "Service workers," W3C Candidate Recommendation Draft, 2026, accessed: 2026-08-10. [Online]. Available: <https://www.w3.org/TR/service-workers/>
- [18] Parakh Project, "Parakh: Pre-departure migrant worker scam shield," Project source repository, 2026, apache License 2.0 metadata in repository; access may require repository permission, accessed: 2026-09-14. [Online]. Available: <https://github.com/apil-khadka/parakh>