

Mapping the Mapmakers: Measuring Language and Source Bias in Global Cybersecurity Education Directories

Anonymous review draft for the 30th Colloquium for Information Systems Security Education (CISSE 2026)

Abstract—

International programme maps are shaped not only by educational provision but also by the systems used to observe it. We evaluate a reproducible, source-aware protocol across heterogeneous national registries and interfaces. Nine within-frame query-sensitivity experiments separate language, vocabulary, taxonomy, search, identity, and counting-unit effects. A tenth, cross-frame experiment in Kenya compares regulatory approval, 2026/27 placement, provider licensing, and institutional publication. Of 14 CUE-approved bachelor's programmes, 13 appeared in KUCCPS; KUCCPS also exposed three higher-education codes absent from the strict CUE extraction. For a Level 6 cybersecurity qualification, 10/18 KUCCPS provider codes matched the TVETA course page and only seven had a current TVETA licence, while six of 11 adjudicated discordances had current institutional publication. We contribute a data model and reporting standard that preserve regulatory identity, visible title, provider, temporal evidence, and source-specific state. The contribution is not another claim of global completeness, but an auditable method for measuring how observation choices change the population that becomes visible.

Index Terms— cybersecurity education; higher education; multilingual information retrieval; coverage bias; program mapping; educational observatory; missing data

1. INTRODUCTION

National cybersecurity capacity depends partly on opportunities to develop advanced knowledge and professional skills. Understanding where such opportunities exist matters to students, governments, universities, and workforce planners. Yet even the apparently simple question “How many university cybersecurity programs does a country have?” lacks a stable empirical answer.

The difficulty is not only definitional. National systems publish different objects: accredited titles, courses approved for international students, enrolled course instances, voluntary directory submissions, or current marketing pages. A title may be approved but no longer admit students; a current pathway may be hidden under a generic parent degree; one program may appear once nationally but operate at several campuses. Search language adds another selection mechanism.

Prior work includes curriculum comparisons, national inventories, regional directories, and international maps. SPARTA analysed 89 curricula and published an interactive worldwide Education Map [11]; Vykopal et al. examined 101 programmes across 24 countries, primarily through ranked universities [1]. CyberHEAD provides an important European directory [4]. Dillon and Tan compared supply and a cybersecurity-interest proxy across 11 Southeast Asian countries [3]. These precedents establish that international mapping and supply–need comparison are not themselves novel. Their frames, however, were designed for curriculum comparison, navigation, voluntary submission, regional coverage, or selected institutions rather than estimation of how discovery infrastructure changes programme observability.

This study asks: How do language, source architecture, temporal evidence, and unit of analysis alter the observed distribution of cybersecurity education? We report a five-country validation pilot, nine within-frame sensitivity experiments, a Kenyan cross-frame concordance study, and architecture probes in Sri Lanka, the Philippines, and South Africa. The public Global Cybersecurity Education Observatory (GCyEO) is treated as versioned research infrastructure, not evidence of a completed world census.

2. RELATED WORK AND RESEARCH GAP

Conference contribution

This study makes three contributions to cybersecurity education research. First, it shifts attention from the visible program list to the measurement system that produces it. That distinction matters because international comparisons can mistake language access, registry design, or stale institutional evidence for differences in educational capacity. Second, it operationalizes source-induced missingness through paired measures: verification estimates the current relevance of discovered records, while negative audits and external frames test whether relevant programs were never discovered. The two quantities are not interchangeable; a directory can contain mostly valid records and still omit a substantial share of provision. Third, it translates these findings into a reproducible observatory design that preserves original-language names, evidence dates, hierarchical relations, and alternative activity definitions. The empirical purpose of the five-country pilot is therefore not to rank national systems. It is to identify failure mechanisms that must be measured before

a global map can support claims about access, inequality, or workforce capacity.

Research on university cybersecurity education falls into four overlapping groups. First, curriculum studies compare knowledge areas, technical/non-technical balance, and experiential learning. Second, national inventories examine selected institutions or accreditation categories. Third, regional and international directories enumerate programs, often through voluntary submission or selective institutional frames. Fourth, national-capacity studies aggregate education, training, policy, and workforce indicators. Selected-master analyses [2], national maturity comparisons [5], the meta-discipline argument [6], and CSEC2017 curricular guidance [10] further establish the conceptual field without supplying a coverage-audit protocol.

The closest precedents delimit the contribution. SPARTA already provides an international interactive map and detailed curriculum analysis [11]. Vykopal et al. analyse content and naming but not census coverage [1]. CyberHEAD is regional and partly submission-dependent [4]. Dillon and Tan restrict supply to QS-ranked universities [3]. None of these sources, as documented, estimates within-frame multilingual query sensitivity while also preserving regulatory identity, public title, delivery instance, temporal state, inaccessible-source outcomes, and negative coverage audits.

National studies add depth without solving global coverage: Crabb et al. review U.S. Centers of Academic Excellence [7], Caelli and Liu inventory Australian provision from university websites [8], and Catota et al. examine capacity and barriers across 13 Ecuadorian institutions [9]. None provides comparable multilingual discovery and explicit coverage audits across systems.

Accordingly, we do not claim the first international comparison or global map. The gap is measurement of observation bias: reproducible multilingual recovery, source recapture, frame-specific precision, temporal evidence, and explicit counting-unit algebra. The intended contribution is a protocol and reporting standard through which later inventories can state what their totals mean and what their sources could not observe.

3. RESEARCH QUESTIONS

- **RQ1:** How much candidate recovery is added by national and additional official languages relative to English-only discovery?
- **RQ2:** How often do regulatory or directory candidates have current institutional confirmation?
- **RQ3:** How often do external frames or negative institutional audits identify qualifying offerings omitted by the initial source?

- **RQ4:** Which mechanisms—language, title semantics, credential scope, temporal lag, interface synchronization, or hierarchy—explain observed omissions?
- **RQ5:** What data model and reporting rules are needed to avoid presenting source artifacts as educational inequalities?

4. METHODS

4.1 Design

We used a sequential multiple-frame design. Frames were selected purposively for variation in language, registry architecture, observable denominators, and preservation of staged retrieval; they are mechanism probes rather than a representative country sample. Within-frame query experiments required baseline hits and an adjudicated expanded denominator in the same national frame. This yielded nine experiments in Spain, Brazil, Canada, South Korea, Germany (two interventions), Poland, the Netherlands, and France. Kenya was analysed separately as a cross-frame concordance experiment; its denominators are not interpreted as query recall.

Country frames were source-specific. The core pilot used official or public frames in Spain, Australia, Brazil, Canada, and Morocco. The extension used AcademyInfo, Studiensus, POL-on/RAD-on, DUO/RIO, and Mon Master [12]–[14]. Kenya linked approval, placement, licensing, and institutional evidence [15]. Sri Lanka contrasted public admissions with recognized non-state degrees [16]; the Philippines audited national and regional CHED channels [17]; South Africa aligned SAQA qualification dates to DHET's private-provider register cut-off [18]. Each frame was dated and its represented object documented before screening.

4.2 Population and inclusion

The target population comprises credentials or formally named pathways issued by recognized higher-education institutions in which cybersecurity, information security, computer/network security, cyber defence, information assurance, digital forensics, or a validated local equivalent is an explicit focus.

We distinguished:

- `'dedicated_program'`: cyber is explicit in the principal credential title;
- `'named_pathway'`: cyber is a formal option, track, specialization, or pathway inside a broader program;
- `'broad_program_with_cyber_content'`: cyber appears only in modules; retained for sensitivity analysis but excluded from the main program count;
- non-degree and internal training: excluded from the principal population.

4.3 Multilingual discovery

Discovery followed a staged sequence: English; principal national language; other official or conceptually relevant local vocabulary; regulatory frame; institutional verification; negative audit. Original names, query terms, interface, access date, and first recovering stage were preserved. The comparative matrix reports baseline recall only within each experiment. It does not pool effects or rank countries because denominators, search functions, and source populations differ.

4.4 Verification and states

We did not treat “active” as a single variable. Evidence was classified across regulatory status, institutional publication, admission, and confidence. The principal verification outcome required a current institutional page, catalog, admission notice, timetable, or equivalent direct evidence. Regulatory-only and unresolved candidates were retained rather than forced into false-positive classifications.

4.5 Coverage audits

Candidate verification estimates precision conditional on discovery. Coverage was examined separately.

- Spain and Australia: ten institutions/providers without candidates were audited in each country.

Table 1. Source architecture, validation design, and permitted inference

Country	Initial frame	Verification	Coverage check	Permitted inference
Spain	QEDU / official titles	20 candidates; combined Wave 1: 34/40 strict active	Directed audit: 10 universities	Detect omission mechanisms; no national omission rate
Australia	CRICOS course records	20 candidates; combined Wave 1: 34/40 strict active	Directed audit: 10 providers	Detect scope and semantic omissions; no national omission rate
Brazil	INEP census; 121 candidates	18/40 strict active	Directed audit: 10 institutions	Estimate candidate verification only; audit is mechanism evidence
Canada	Federal bilingual guide; 237 records	35/40 strict active	Ontario public-college universe: 52 candidates	Frame-specific Ontario omission proportion; not all Canada
Morocco	Accreditation plus public triangulation; 24 records	19/24 strict active	Directed audit: 10 academic units	Observed recapture and mechanisms; no national omission rate



Figure 1. Versioned observation cycle used by the Global Cybersecurity Education Observatory.

5. RESULTS

5.1 Multilingual recovery

In Spain, the protocol's English-only stage identified five offerings; the subsequent Spanish-language and regulatory stages produced 57 offerings (54 unique registrarial titles). The English stage therefore recovered 8.8% of the final observed

- Brazil: ten institutions stratified by region and control were audited.
- Canada: the federal guide was compared with the complete Ontario public-college program file (52 cyber candidates), and jurisdictions with zero federal records were separately reviewed.
- Morocco: ten academic units across seven regions were searched to identify omitted dedicated and embedded provision.

The Ontario file supports a frame-specific omission proportion for its public-college universe. Smaller directed audits only falsify assumed completeness and identify mechanisms; their denominators are not national program populations. We neither pool nor rank these proportions.

4.6 Analysis

We report counts, conditional proportions, Wilson intervals where appropriate, marginal language recovery, and recapture gains. Because national frames and audit designs differ, country proportions are not pooled and are not interpreted as rankings. Capture–recapture population estimates were rejected because source independence and homogeneous capture probabilities were not defensible.

offering set, while adding Spanish multiplied protocol recovery by 11.4. This contrast is conditional on the staged queries, sources, and cut-off date; it is not a general estimate of English-language web discoverability.

In Brazil, the protocol's English vocabulary returned no qualifying candidates from the official 2024 higher-education census. Portuguese matching produced 6,167 offer/dimension

rows, representing 123 unique course codes. Two public-security false positives were removed, leaving 121 cyber candidates at 90 institutions. Under this fixed census and query design, the language intervention changed observed recovery from zero to 121 candidates; the result should not be generalized to unconstrained English web searching.

Canada's federal guide contained 235 English-interface and 236 French-interface records. Their union contained 237: 234 shared and three interface-exclusive. All three exclusive records were confirmed institutionally. Thus, the marginal language gain was small in quantity but perfect in observed relevance.

Across the expanded matrix, baseline recovery was 5/57 in Spain (8.8%), 0/121 in Brazil (0.0%), 235/237 for Canada's English interface (99.2%), 26/118 for South Korea's literal terminology (22.0%), 37/71 for Germany's official subject taxonomy (52.1%), 48/71 for its German query set (67.6%), 3/28 for Poland's broad base-word query (10.7%), 9/17 for the Netherlands' cyber-literal title query (52.9%), and 7/36 for English-titled French Mon Master families (19.4%). The range is not an international distribution: it diagnoses different selection mechanisms within each source.

5.2 Conditional candidate verification

Across the combined Spain–Australia verification sample, 34/40 candidates (85.0%) were institutionally current under strict nominal matching; 35/40 (87.5%) when name drift was accepted.

Brazil showed a different temporal pattern. Strict direct confirmation was 18/40 (45.0%; Wilson 95% CI 30.7–60.2). Including active group pages increased confirmation to 23/40 (57.5%; 42.2–71.5). Nevertheless, 36/40 (90.0%) had some institutional evidence, including historical pages and catalogs. The discrepancy therefore reflected regulatory persistence and public visibility more than wholesale fabrication.

In Canada, 35/40 guide records were strictly active (87.5%; 73.9–94.5), rising to 36/40 (90.0%; 76.9–96.0) with group pages and 37/40 (92.5%; 80.1–97.4) with catalog evidence.

Morocco's expanded frame contained 24 program-campus records. Nineteen were directly active (79.2%); twenty (83.3%) when a current university-group page was included. Four remained accreditation-only. No case was classified as a confirmed false positive.

5.3 Coverage failures under non-equivalent audit designs

The Wave 1 directed negative audit found at least one qualifying omission in 4/10 Spanish universities and 2/10 Australian providers. These detection proportions describe the audited samples rather than national omission prevalence. Mechanisms differed: permanent education, proprietary credentials, microcredentials, hidden pathways, semantic title

variation, and the domestic/international scope of the Australian source.

Ontario's official public-college universe contained 52 cyber candidates. Fourteen were confirmed active omissions from the federal guide (26.9% lower bound). Omissions included recent launches, future first cohorts, degree or master's offerings, Francophone provision, and specialized titles.

The Canadian guide contained no record in four jurisdictions. Institutional review found related embedded or non-degree provision in three, demonstrating title dependence rather than false negatives under the strict dedicated-program definition.

Brazil's directed audit found one omission among ten institutions: a current introductory Cisco cybersecurity course. It was non-degree and therefore outside the main population; the mechanism was credential scope, not language. The 1/10 result is evidence that the audit can expose a scope boundary, not an estimator of Brazilian coverage.

Morocco's directed audit found two dedicated omissions among ten academic units: a Bachelor in Cybersecurity of Intelligent Connected Systems at EST Fès and a program in Port and Multimodal Management, Digitalization and Cybersecurity at ENCG Tanger. Adding three broad programs with substantive cyber modules raised the number of audited units with an educational omission to five, although these broad programs remain outside the strict principal count. Because the units were selected to probe mechanisms rather than sampled probabilistically, confidence intervals and national prevalence interpretations are inappropriate here.

5.4 Hierarchy and structural visibility

The Netherlands yielded 15 regulated-provider families, 17 visible-title families, and 21 delivery instances. Poland's broad query returned 504 records but only three strict families; expanded vocabulary produced 28, giving 0.6% precision and 10.7% recall. In France, crosswalks reconciled all 31 prior families despite raw identifiers suggesting 13 entries and eight exits. Kenya exposed a different mechanism: 13/14 CUE-approved bachelor's programmes appeared in KUCCPS, while three KUCCPS higher-education codes were absent from the strict CUE extraction. For Cyber Security TVET Level 6, 10/18 KUCCPS provider codes matched TVETA and seven had current licences. Of 11 adjudicated TVET discordances, six had current institutional publication, one only historical publication, and four remained unresolved.

These cases expose separable mechanisms. Language, interfaces, taxonomies, registry structure, annual identifiers, and regulatory purpose all change what becomes visible. Sri Lanka's 2025/26 public-admissions frame yielded zero strict titles while a ministry list recovered five recognized non-state degrees; three had exact current institutional matches. In the Philippines, a national directory exposed 2,734 institutional rows, but three observable regional programme channels

yielded only six strict delivery records. In South Africa, six private HEQSF qualifications became five after aligning SAQA registration dates to DHET's 28 February 2026 cut-off; three matched exactly and two did not. Treating the post-cut-off sixth record as missing would have manufactured a discordance. These are frame contrasts and lower bounds, not national prevalence estimates.

6. SENSITIVITY ANALYSES

6.1 Definition of “active”

Activity thresholds remain consequential. Kenya provides direct evidence: regulatory approval, placement, provider licensing, institutional publication, and dated admission did not align record-for-record. Absence from one frame cannot be converted to closure, while approval cannot be interpreted as current intake. The observatory therefore publishes state-specific evidence instead of one active flag.

6.2 Dedicated versus embedded provision

Dedicated credentials and named specialisations answer different questions. They are retained in separate scope classes and connected to their parent programmes. Broad curricular content is excluded from the principal programme population unless a formally named pathway is evidenced.

6.3 Language union

The nine experiments identify five non-poolable mechanism classes. Even Canada's 1.3% interface-union gain represented three confirmed records, while literal or English baselines missed most of the adjudicated frame in Brazil, Spain, South Korea, Poland, and France. Yield alone cannot determine whether a stage is methodologically necessary.

6.4 Counting unit

Counting-unit sensitivity should be reported alongside every total. In the Dutch case, 15 regulated-provider families, 17 visible titles, and 21 delivery instances are all correct answers to different questions. The data model must preserve parent code, institutional name, mode, location, and crosswalk relation rather than force premature deduplication.

7. DISCUSSION

7.1 Counts are properties of observation systems

Observed programme supply is co-produced by education and observation architecture. Regulatory, admissions, placement, licensing, regional, voluntary, and institutional sources do not enumerate equivalent populations. Kenya refutes one-source-status-equals-activity; Sri Lanka demonstrates a frame-specific zero; Philippine decentralization produces a lower bound; and South Africa shows that unsynchronised cut-offs can create false discordance.

7.2 Precision does not imply coverage

Canada illustrates why verifying found records is insufficient: guide entries were usually current, yet the Ontario comparison exposed substantial omissions. Conversely, Brazil's broad regulatory enumeration had lower current web confirmation. A defensible observatory must estimate precision and coverage separately.

7.3 Multilinguality is necessary but insufficient

Multilingual discovery is necessary but insufficient. Spain and Brazil show language substitution; Canada shows interface asymmetry; South Korea shows conceptual vocabulary; Germany shows taxonomy mismatch; Poland shows morphology–search interaction; and the Netherlands shows naming plus unit structure. Reproducibility requires recording the mechanism, not merely stating that local-language searches were performed.

7.4 Implications for inequality research

Raw national totals should not be normalized into rankings until coverage states and units are explicit. Otherwise, source availability, registry purpose, or query behavior can be misreported as educational inequality. The present data support within-frame sensitivity statements and architecture comparisons, not country performance scores.

8. PROPOSED OBSERVATORY MODEL

We propose four linked entities: organization, campus, program, and pathway/credential. Annual offerings and modalities are temporal variants. The public interface should expose two layers:

1. named dedicated credentials and pathways;
2. formally documented embedded or curricular provision, reported separately.

Each record retains original and translated name, country, provider, regulatory or parent identifier, visible-title family, delivery instance, source and access date, scope, temporal evidence, retrieval mechanism, and change history. Country pages state frame population, query profile, audit design, supported inference, unresolved exclusions, and known inaccessible channels. The public interface exposes downloadable records and caveats; a versioned snapshot anchors each paper release.

9. THREATS TO VALIDITY AND LIMITATIONS

The work is an incremental observatory, not a probability sample or completed world census. Directed audits identify mechanisms but usually cannot estimate national omission prevalence; Ontario is exhaustive only for one provincial public-college frame. National sources differ in purpose and update cycles. The nine sensitivity estimates must not be pooled. Sri Lankan, Philippine, and South African counts are

architecture probes with unequal frames, and the public candidate layer is not a globally comparable programme total.

Coding and adjudication were conducted by the primary research workflow. Institutional spot checks and deterministic transformations strengthen traceability but do not replace independent coding. A frozen 95-case multilingual second-coder instrument has been prepared, but no agreement result is reported because independent coding is not yet complete. Reliability claims are therefore limited to deterministic reproducibility and evidence audit. Query vocabularies may still omit local concepts, and private or non-funded provision can fall outside official frames.

10. ETHICS, LICENSING, AND DATA AVAILABILITY

The study uses public institutional and governmental information and does not recruit participants or process sensitive personal data. It reports programme-level metadata rather than applicant-level records. Whether an institutional ethics determination is required remains subject to the submitting author's institution and must be documented before final submission. Published data should contain factual metadata, classifications, URLs, and short evidence notes rather than reproducing protected curricular content. Source licenses and redistribution conditions must be reviewed before public release of raw regulatory extracts.

The protocol, dictionaries, query logs, processed data, audit forms, source snapshots where licensing permits, and sensitivity matrix are versioned. The public observatory provides navigation and downloads. A frozen conference snapshot and persistent repository identifier are required before the camera-ready release; until then, the live site is supporting infrastructure rather than an archival record. Counts are accompanied by provenance and interpretation warnings.

11. CONCLUSION

Existing projects already map cybersecurity programmes internationally. GCyEO asks whether the mapped population is stable under changes in language, source, temporal evidence, identity, and counting unit. It is not: Poland combined 0.6% precision with 10.7% recall; the Dutch registry yielded 15, 17, or 21 valid totals; French apparent churn disappeared after reconciliation; Kenya separated programme states; Sri Lanka produced a frame-specific zero; Philippine decentralization produced a lower bound; and South African cut-off alignment prevented a false discordance.

A defensible observatory must therefore publish the measurement system with the map. Each figure should name its denominator, observational unit, inclusion rule, temporal state, source architecture, and tested query profile. This turns absence into a graded evidentiary state, makes source bias

auditable, and creates a safer foundation for later comparisons of educational access.

REFERENCES

- [1] J. Vykopal, V. Švábenský, M. T. Lopez, and P. Čeleda, "Cybersecurity Study Programs: What's in a Name?," in *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1*, 2025, pp. 1169–1175. <https://doi.org/10.1145/3641554.3701976>
- [2] K. Cabaj, D. Domingos, Z. Kotulski, and A. Respício, "Cybersecurity education: Evolution of the discipline and analysis of master programs," *Computers & Security*, vol. 75, pp. 24–35, 2018. <https://doi.org/10.1016/j.cose.2018.01.015>
- [3] R. Dillon and K.-L. Tan, "Cybersecurity Workforce Landscape, Education, and Industry Growth Prospects in Southeast Asia," *Journal of Tropical Futures*, vol. 1, no. 2, pp. 172–181, 2024. <https://doi.org/10.1177/27538931231176903>
- [4] European Union Agency for Cybersecurity, "Cybersecurity Higher Education Database (CyberHEAD)." <https://www.enisa.europa.eu/tools/cyberhead-cybersecurity-higher-education-database>
- [5] R. Shillair, P. Esteve-González, W. H. Dutton, S. Creese, E. Nagyfejeo, and B. von Solms, "Cybersecurity education, awareness raising, and training initiatives: National level evidence-based results, challenges, and promise," *Computers & Security*, vol. 119, Art. no. 102756, 2022. <https://doi.org/10.1016/j.cose.2022.102756>
- [6] A. Parrish et al., "Global perspectives on cybersecurity education for 2030: a case for a meta-discipline," in *Proceedings Companion of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education*, 2018, pp. 36–54. <https://doi.org/10.1145/3293881.3295778>
- [7] J. Crabb, C. Hundhausen, and A. Gebremedhin, "A Critical Review of Cybersecurity Education in the United States," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, 2024, pp. 241–247. <https://doi.org/10.1145/3626252.3630757>
- [8] W. Caelli and V. Liu, "Cybersecurity Education at Formal University Level: An Australian Perspective," *Journal of The Colloquium for Information Systems Security Education*, vol. 5, no. 1, pp. 26–44, 2018. <https://journal.cisse.info/index.php/cisse/article/view/78>
- [9] F. E. Catota, M. G. Morgan, and D. C. Sicker, "Cybersecurity education in a developing nation: the Ecuadorian environment," *Journal of Cybersecurity*, vol. 5, no. 1, 2019. <https://doi.org/10.1093/cybsec/tyz001>
- [10] Joint Task Force on Cybersecurity Education, *Cybersecurity Curricula 2017*, Version 1.0. ACM, IEEE Computer Society, AIS SIGSEC, and IFIP WG 11.8, 2018. https://cybered.hosting.acm.org/wp-content/uploads/2018/02/newcover_csec2017.pdf
- [11] SPARTA, 'Cybersecurity Training and Awareness: Education Map,' and D9.2, 'Curricula Descriptions,' 2020. <https://www.sparta.eu/training/>
- [12] Dienst Uitvoering Onderwijs, 'HO Opleidingsoverzicht,' Registratie Instellingen en Opleidingen, 2026. https://onderwijsdata.duo.nl/datasets/ho_opleidingsoverzicht
- [13] RAD-on, 'POL-on Courses API documentation,' 2026. https://radon.nauka.gov.pl/api/katalog-udostepniania-danych/dane-polon/polon/reports_course
- [14] Ministère de l'Enseignement supérieur et de la Recherche, 'Mon Master: données du portail d'information — parcours de formation,' 2024–2025. <https://data.enseignementsup-recherche.gouv.fr/explore/dataset/fr-esr-tmm-donnees-du-portail-dinformation-trouver-mon-master-parcours-de-format/>
- [15] Commission for University Education, Kenya Universities and Colleges Central Placement Service, and Technical and Vocational Education and Training Authority, official programme, 2026/27 placement, and provider-licensing records, accessed Aug. 2026. <https://www.cue.or.ke/>; <https://students.kuccps.net/>; <https://www.tveta.go.ke/>
- [16] University Grants Commission and Ministry of Education, Sri Lanka, public-university admissions handbook 2025/26 and recognized-degree records, accessed Aug. 2026. <https://www.ugc.ac.lk/>; <https://mohe.gov.lk/>
- [17] Commission on Higher Education, Philippines, HEIDA directory and regional programme inventories, accessed Aug. 2026. <https://ched.gov.ph/>
- [18] South African Qualifications Authority and Department of Higher Education and Training, qualification records and Register of Private Higher Education Institutions, updated 28 Feb. 2026, accessed Aug. 2026. <https://regqs.sqa.org.za/>; <https://www.dhet.gov.za/>

GENERATIVE AI DISCLOSURE

Automated tools supported drafting and formatting. Sources, classifications, results, and interpretations were author-verified; no generative system served as an evidentiary source.