Mapping cybersecurity education worldwide

A descriptive evidence atlas across 218 jurisdictions

Federico Pacheco · Universidad Tecnológica Nacional · ORCID 0000-0002-2803-1162

Working manuscript v4.0 · short edition · revised 2 November 2026

Global data cut-off 29 August 2026 · supplemental source checks through 2 November 2026 · descriptive mapping study

Abstract

This study maps the documented worldwide landscape of cybersecurity education using a frozen public-evidence atlas. We combine a global candidate layer with a separately curated five-country pilot, preserving source, credential and temporal distinctions. Across 218 evaluated jurisdictions, the global layer retains 7,167 candidate rows in 194 countries; four additional countries are represented only in the pilot, and 20 jurisdictions have no retained rows in the consulted frames. The pilot contains 616 records across Spain, Australia, Brazil, Canada and Morocco and overlaps the global layer in Australia; its total is not added to the global count. Documented education spans undergraduate and postgraduate awards, short-cycle and certificate routes, and embedded pathways. The five largest country frames contain 58.6% of candidate rows, while 115 positive frames contain nine or fewer. Regulatory URL fields alone are populated in 74.6% of rows. These patterns describe the evidence assembled, not national educational capacity. Bounded country checks demonstrate that language, source coverage, counting units and evidence dates affect the map. A separate 400-case audit did not identify a precise eligible-current share, so no adjusted global programme total is estimated. The contribution is a traceable descriptive baseline and jurisdiction atlas for locating documented education and identifying where comparable evidence remains insufficient.

Keywords: cybersecurity education; global mapping; education provision; public evidence; higher education; educational pathways

1 Introduction

Where is cybersecurity education documented around the world, what forms does it take, and how well does the available evidence describe it? These are practical questions for students locating educational routes, institutions situating their programmes, and researchers examining the international landscape. A useful map must report both the education found and the limits of what its sources can establish.

We present the Global Cybersecurity Education Observatory as a descriptive evidence atlas. The primary objective is to map the worldwide documented landscape using the records already assembled. The analysis describes geographic distribution, recorded educational forms and evidence coverage. It does not estimate a complete global stock of active programmes or assign national capacity rankings.

Prior work examines programme naming [1], master's education [2], national education and training initiatives [3], and regional higher-education directories such as CyberHEAD [4]. This study contributes a cross-jurisdiction descriptive baseline with inspectable record provenance, a separate deeper pilot and explicit distinctions between a missing record and missing education. It does not claim to be the first international inventory.

2 Data and descriptive methods

The global analysis uses the unchanged candidate and jurisdiction files frozen at 29 August 2026 [5]. A separate pilot file covers Spain, Australia, Brazil, Canada and Morocco [6]. Later audits through 2 November inform the

interpretation of these sources but are not silently added to the August counts. This is an exploratory secondary analysis of an existing assembled dataset, without fresh web searches or source acquisition for this revision.

Table 1. Evidence layers and their descriptive scope.

Evidence layer	Extent	Interpretation
Global candidates	7,167 rows in 194 countries	Heterogeneous retained evidence objects
Five-country pilot	616 records in 5 countries	Deeper curation; conditional checks have separate samples
Jurisdiction frame	218 evaluated jurisdictions	194 candidate-positive + 4 pilot-only + 20 without retained rows
Later source audits	September to 2 November 2026	Interpretive supplements; no refreshed global census

A candidate row is the primary descriptive unit. Depending on the national source, a row can describe an award, a specialization, a programme-provider or campus instance, an administrative reporting unit, or a scope-boundary lead. A programme family groups related identities; a delivery instance concerns a location, modality or offering. These units are retained separately. The 7,167 rows are not asserted to be 7,167 unique active degree programmes.

Sources include national regulatory registers, admissions and placement systems, statistical files, official directories and institutional catalogues. National and additional official languages were used where the existing collection required them. Country source descriptions and retrieval dates are retained in the jurisdiction atlas. Selection and search depth varied by country; the global collection was neither a probability sample nor an exhaustive common query applied to all jurisdictions [5–7].

We count rows by country code and by the regional labels already present in the jurisdiction file, trimming whitespace only. These source-region assignments are retained for reproducibility, including their grouping of Afghanistan and Pakistan with Middle East and North Africa. Exact scope and credential labels are tabulated without an unvalidated global recoding. Institutional-label counts use distinct nonempty strings within country; they do not resolve provider identities globally. URL-field presence is calculated separately from substantive source verification.

3 Worldwide documented landscape

The atlas covers 218 evaluated jurisdictions: 194 contribute candidate rows, four contribute pilot records only, and 20 have no retained rows in the consulted frames. This establishes a broad documented footprint across multiple education systems while leaving completeness unmeasured. Figure 1 locates the retained evidence; jurisdiction_atlas.csv provides the full 218-row lookup, including territories not visible at this map scale.

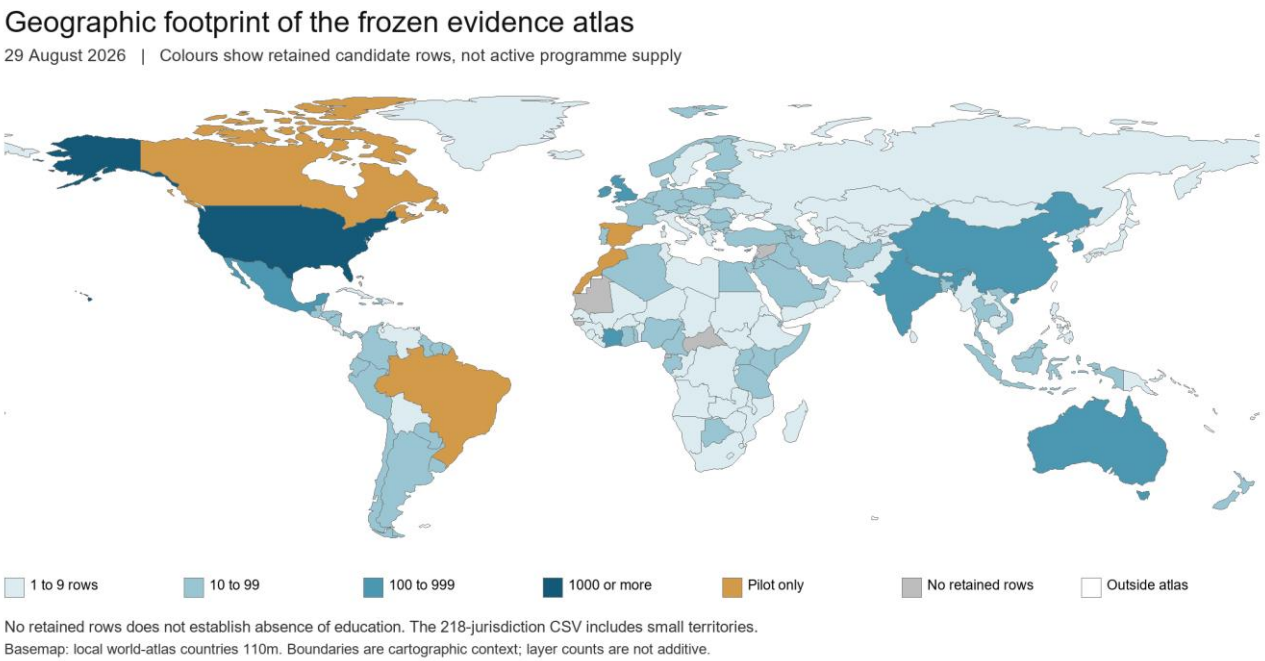


Figure 1. Geographic footprint of retained candidate evidence at the August cut-off. Colour bins count rows, not unique active programmes. Pilot-only countries are marked separately; Australia has candidate and pilot records and is coloured by candidate rows. Grey indicates no retained rows in the consulted frames. White indicates areas outside the evaluated atlas. Small territories are included in the accompanying table.

Table 2. Regional coverage in the assembled atlas. Columns count jurisdictions except candidate rows.

Source region	Evaluated	Candidate positive	Candidate rows	Pilot only	No rows
East Asia & Pacific	37	35	1,708	0	2
Europe & Central Asia	58	54	981	1	3
Latin America & Caribbean	43	34	766	1	8
Middle East, North Africa, Afghanistan & Pakistan	23	21	256	1	1
North America	3	1	2,602	1	1
South Asia	6	5	138	0	1
Sub-Saharan Africa	48	44	716	0	4

Table 2 reports distribution within the assembled evidence. For example, North America's 2,602 candidate rows all come from the United States, while Canada's 237 records belong to the separate pilot. A candidate-only regional comparison would therefore omit Canadian evidence that the project does contain. Region totals are not comparable estimates of regional educational provision.

The largest candidate frames are the United States (2,602 rows), China (833), Côte d'Ivoire (311), South Korea (232) and Australia (222). Together they contribute 4,200 rows, or 58.6%. Among the 194 positive country frames, the median is seven rows and 115 have nine or fewer, jointly contributing 465 rows (6.5%) [7]. The atlas is consequently broad geographically but uneven in depth.

4 Educational forms and levels

The records describe a range of educational routes, including bachelor's and master's awards, associate and short-cycle qualifications, certificates and embedded pathways. These are recorded labels rather than a harmonized worldwide qualification taxonomy. The global file contains 461 distinct credential-level strings and 319 scope strings. We therefore report exact labels and preserve the full frequency dictionaries instead of merging ambiguous categories into apparently comparable degree totals.

Table 3. Most frequent exact credential labels. Percentages describe retained rows.

Exact credential label	Candidate rows	Share of 7,167
bachelor	1,636	22.8%
certificate	954	13.3%
master	951	13.3%
short_cycle	848	11.8%
associate	612	8.5%
unspecified	133	1.9%
All other credential labels	2,033	28.4%

The exact bachelor label occurs in 1,636 rows and master in 951. Certificate (954), short_cycle (848) and associate (612) demonstrate that the assembled landscape extends beyond bachelor's and master's titles. These categories do not establish internationally equivalent levels or active student places. Other strings can name the same broad level in local terminology or express unresolved level; the table therefore cannot estimate the global proportion of education at each level.

The exact scope label dedicated_program occurs in 5,453 rows (76.1%), scope_boundary in 854 and embedded_pathway in 198. Other labels account for 662 rows. These are collection labels, not verified cyber-degree counts; standalone awards and embedded content remain distinct.

5 Source coverage and bounded validation

Of the 7,167 candidate rows, 5,346 (74.6%) have a regulatory URL field without an institutional URL, 1,417 (19.8%) have both, and 404 (5.6%) have only an institutional URL. All rows populate at least one of these fields. Field presence documents a route to evidence; it does not establish that a link remains accessible, names the same award or proves current intake.

Table 4. The separate five-country pilot. Counts are not added to the global layer.

Pilot country	Retained records	Interpretation
Australia	170	Also represented in the candidate layer; overlap not deduplicated
Brazil	121	Pilot frame; conditional 40-case evidence check is separate
Canada	237	Pilot frame; Ontario coverage comparison is sector-bounded
Morocco	24	Pilot programme and campus records; no national completeness claim
Spain	64	Pilot catalogue; staged 57-offering comparison is a different object

Existing checks illustrate why this map should be read with source context [6,7]. In Spain, staged discovery expanded from five offerings found in English to 57 after Spanish and regulatory stages; language and source changed together. In Brazil, 18 of 40 sampled candidates had direct current-page evidence, 23 had evidence when group pages were accepted, and 36 had some institutional evidence. In Canada, 35 of 40 sampled guide records were strictly confirmed, yet an external Ontario public-college comparison found 14 active omissions among 52 candidates. These results describe their own samples and frames; they are not national validation or omission rates.

The separate stratified audit of 400 candidate rows produced three provisional inclusions, three provisional exclusions and 394 unresolved primary decisions [8]. It did not identify a precise eligible-current fraction and cannot estimate programmes absent from the candidate file. Later source work leaves only 14 cases without an identity signal, but identity is a different test from eligible current provision. Neither result justifies scaling the candidate count into a verified global programme estimate. Supplement S1 retains these outcomes and their limitations.

6 Interpretation and limitations

The atlas establishes that cybersecurity-related education is publicly documented across a wide range of jurisdictions and credential routes. It provides a starting point for locating named provision and examining source coverage. Sparse country evidence identifies a documentation gap in this atlas; it does not distinguish limited educational supply from limited discovery, public visibility or source granularity.

Several limits constrain any account of the state of education. Source populations and search intensity differ across countries; candidate units and qualification labels are heterogeneous; identity and currency are not established for all rows; and the pilot is deeper but nonrepresentative. Cross-border award, teaching and learner geography can differ. Joint programmes and multiple campuses create unresolved overlap. The study does not measure enrolment, teaching capacity, curriculum quality, affordability, access or labour-market outcomes. In particular, the field named enrollment includes completions in the US extraction and cannot support a global student total.

The map is a dated baseline. Global descriptive counts refer to August 2026, not to a synchronized November 2026 census. Later observations through 2 November are source-specific checks [9,10]; they establish no general programme opening or closure and do not constitute a completed longitudinal panel. Temporal stability and growth in worldwide provision are therefore not assessed.

7 Conclusion

Using existing public evidence, GCyEO maps a worldwide documented landscape across 218 evaluated jurisdictions. Its global layer contains 7,167 candidate rows in 194 countries, complemented by a separate five-country pilot of 616 records. Education is documented through degree, short-cycle, certificate and embedded routes, but country coverage and record semantics remain uneven. The resulting atlas supports discovery and contextual description; it does not estimate a complete worldwide stock of active programmes. Publishing the geographic lookup, original labels and source boundaries makes that baseline usable without concealing its uncertainty.

Data availability and disclosure

This local revision accompanies jurisdiction_atlas.csv, regional_summary.csv, credential_labels.csv, scope_labels.csv, descriptive_metrics.json, a basemap crosswalk and an artifact manifest [12]. The builder uses local files only. Original data and frozen primary decisions are unchanged. Source links refer to third-party material and do not grant redistribution rights. The earlier v3.3 preprint has a recorded Zenodo DOI [11]; this v4.0 manuscript has not been deposited or submitted by this revision.

AI assisted source coding in earlier stages and the restructuring and generation of this manuscript. Earlier archived agreement results are not treated as an accuracy reference or representative human validation. Supplement S1 identifies the targeted earlier human audit and the later no-additional-human-review constraint. This revision performed no new case review by humans or external models. The author remains responsible for the analysis and manuscript.

References

[1] J. Vykopal, V. Švábenský, M. T. Lopez, and P. Čeleda, Cybersecurity Study Programs: What's in a Name?, SIGCSE TS 2025, pp. 1169–1175. doi: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 75, pp. 24–35, 2018. doi:10.1016/j.cose.2018.01.015.

[3] R. Shillair et al., Cybersecurity education, awareness raising, and training initiatives, Computers & Security 119, 102756, 2022. doi:10.1016/j.cose.2022.102756.

[4] European Union Agency for Cybersecurity, Cybersecurity Higher Education Database CyberHEAD. https://www.enisa.europa.eu/tools/cyberhead-cybersecurity-higher-education-database.

[5] GCyEO, frozen global candidate and country layers, cut-off 29 August 2026. app/data/wave3-candidates.json and app/data/countries.json. Exact input SHA-256 values and derived jurisdiction table accompany this manuscript in artifact_manifest.json and jurisdiction_atlas.csv.

[6] GCyEO, five-country pilot layer, app/data/programs.json, 616 retained records. Sampling provenance: research/outputs/GCyEO_BRA_CAN_sampling_provenance_v0.1.md. The data layer and the conditional verification samples are separate denominators.

[7] GCyEO, canonical 24-contrast registry and frozen headline metrics, 23 September 2026. research/outputs/GCyEO_v3_experiment_registry.json and GCyEO_v3_headline_metrics.json. Underlying source paths, hashes and identification limits are retained in the registry.

[8] GCyEO, probability audit execution and subsequent identity reconciliation, 2–29 October 2026. research/next_study_v3_4/probability_400_execution_report_2026-10-02.md and probability_400_unified_source_identity_ledger_v98.csv. Primary and post-result outcomes are distinct.

[9] GCyEO, cumulative November source-panel contrast, 2 November 2026. research/next_study_v3_4/manuscript_v3_4_180_source_index_2026-11-02.md and monthly_observation_report_2026-11-02.md. This follow-up does not update the frozen global counts.

[10] GCyEO, detailed measurement manuscript v3.4.180, full edition, 2 November 2026. research/outputs/GCyEO_manuscript_v3_4_180_full.docx. Preserved as the historical technical companion; its measurement-focused framing is superseded for the new descriptive manuscript.

[11] GCyEO, measurement preprint v3.3 and reproducibility package, 23 September 2026. doi:10.5281/zenodo.22925125. The existing deposit does not contain or publish this v4.0 revision.

[12] GCyEO, global mapping v4.0 descriptive outputs and analytical supplement, 2 November 2026. research/outputs/global_mapping_v4_0/. Offline reconstruction from unchanged local inputs; no new source acquisition.