

AI Dubbing is becoming *the new distribution layer for global content.*

A Use Case Map analysis of 112,797 categorized professional dubbing projects on the Perso AI platform reveals industry-specific language patterns within Perso AI's global cohort of 4,023 professional creators across 80+ countries.

Released June 4, 2026

Author · Perso AI Data Team

Creative Commons Attribution 4.0

perso.ai/research/state-of-ai-dubbing-2026/

316,856

PROJECTS
16-MONTH

4,023

PRO
CREATORS

909

LANG
PAIRS

36×34

SOURCE ×
TARGET

96%

SHARE
RATE

80+

COUNTRIES
REACHED

Contents

PART I — THE THESIS

01	Executive Summary	04
	Thesis, headline statistics, and the spine of the report	
02	The 4-Layer Model	06
	Why AI Dubbing sits at a different layer of the AI media stack	

PART II — THE FINDINGS

03	The Use Case Map	08
	Cross-tabulation of industry × target language	
04	Finding 01 — Religion Has a Dual Hub	09
	English 25.6% ≈ Portuguese 25.2%	
05	Finding 02 — K-Content Spillover	10
	Korean as the structural #2 in sci/tech	
06	Finding 03 — Multi-Language Adoption Frontier	12
	Median 1, Top 1% (n=47) avg 15 languages	

PART III — DEEP-DIVES

07	Per-Industry Deep-Dive	14
	Six industries with distinct language geographies	
08	Per-Language Deep-Dive	15
	Six target markets with distinct industry concentrations	

PART IV — IMPLICATIONS

09	So What — Three Audience Implications	16
	For media, capital, and creators	
10	Why This Matters Now	18
	The reporting window	
11	Industry Voices	19

Five expert statements on AI localization

12	Predictions for 2027 Three directional shifts	20
----	---	----

PART V – CLOSING & METHODOLOGY

13	Three Things to Take Away Closing synthesis — actionable distillation of the three findings	21
14	Methodology Notes Two-lens approach, definitions, honest limitations	22
15	Appendix A — Top 1% Cohort Composition Anonymized data on the 47 power creators	24
16	Glossary Eight key terms defined as used in this report	26
17	Citation, Downloads & Contact CC BY 4.0 license and re-use information	27

CHAPTER 01

Executive Summary

This report frames AI Dubbing as a distinct layer of the AI media stack – the distribution layer – distinct from voice cloning and avatar generation. Across Perso AI's data, three counterintuitive findings emerge that connect to broader macro narratives in global media and creator economics.

THE THESIS

AI Dubbing is where finished video crosses language borders at the speed of distribution. It operates at the infrastructure layer, multi-polar and multi-vertical by design.

The Spine in Six Stats

METRIC	VALUE	SCOPE
Total projects	316,856	Jan 2025 – Apr 2026 (16 months)
Categorized projects (6-month)	112,797	Oct 2025 – Apr 2026
Professional creators	4,023	Defined: 6+ projects on the platform
Active language pairs	909	Source × target combinations
Source × target languages	36 × 34	Within Perso AI's data
Share rate	96%	Among all 316,856 projects

Key Findings (Preview)

Finding 01 — Religion Dual Hub. Within Perso AI's 6,229 religion projects, English (25.6%) and Portuguese (25.2%) form a near-equal dual hub. Brazilian Portuguese faith outreach reaches English-parity at scale.

Finding 02 — K-Content Spillover. Within Perso AI's 6,160 sci/tech projects, Korean is the #2 target language at 12.5% — ahead of Spanish. Patterns consistent with K-Content's broader cultural footprint extending into knowledge consumption.

Finding 03 — Multi-Language Frontier. Among Perso AI's 4,023 professional creators, the median dubs into 1 target language; the average is 2.43 (heavy-tail); the top 1% (n=47) averages 15. The infrastructure supports 33. The directional gap is the expansion-revenue opportunity.

Editorial Framing – Acknowledged Upfront

This report's 4-Layer Model framing (AI Dubbing at the distribution stage, distinct from voice cloning and avatar generation at the creation stage) is editorial, not an objective industry taxonomy. Voice cloning tools — including ElevenLabs Voice Lab — also offer dubbing features. We frame the categories around production stage and output type because we find this framing more useful for understanding where the AI media stack is heading, not as a settled classification.

Findings 1, 2, and 3 are scoped to Perso AI's professional creator cohort and do not claim to represent the entire AI dubbing market globally. Honest limitations — including user-acquisition bias and the exclusion of volume-based time-series — are documented in **Chapter 14 (Methodology Notes)**.

How the Three Findings Connect

AI dubbing is not a single global market running on one default workflow. The three findings below describe a distribution layer where industry-specific patterns coexist at scale. Religion concentrates in a Portuguese–English dual hub. Sci/tech extends into a Korean language frontier that mirrors K-Content's broader spillover. Across verticals, the most active creators on Perso AI dub into 15 target languages, while the typical creator stays at one. Read together, the findings describe a distribution layer that is multi-polar, multi-vertical, and multi-language at once.

CHAPTER 02

The 4-Layer Model

Mainstream coverage often groups AI Dubbing together with voice cloning and avatar generation. We propose framing them as different layers of the AI media stack, doing different work at different stages.

AI Dubbing's defining feature, in this framing, is that the output operates as a **distribution event** rather than a creation-stage asset. This 4-layer separation is editorial — voice cloning tools (including ElevenLabs Voice Lab) also offer dubbing features. Our category distinction emphasizes distribution-stage workflow over creation-stage assets, a framing we find useful for understanding where the AI media stack is heading.

LAYER	EXAMPLES	OUTPUT	PRODUCTION STAGE
Layer 1 Voice Cloning	ElevenLabs • Resemble • PlayHT	A synthetic voice. The asset is the voice itself.	CREATION
Layer 2 Avatar Generation	HeyGen • Synthesia • D-ID	A video featuring a synthetic person. The asset is the avatar.	CREATION
Layer 3 Text Translation	Google Translate • DeepL	Translated text. The asset is a file used in pre-distribution workflows.	PRE-DISTRIBUTION
Layer 4 AI Dubbing (this category)	Perso AI • category peers	A video deployed across multiple language markets simultaneously. The "asset" is a shipment.	* DISTRIBUTION

Voice clones and avatars are assets, produced once and reused many times. Translated subtitles sit as files in a production pipeline. A dubbed video is something different: it ships the moment it's produced.

EDITORIAL FRAMING • STATE OF AI DUBBING 2026

The Behavioral Fingerprint – 96% Share Rate

Share rate is the behavioral signal we use as a categorical fingerprint. **Among Perso AI's 316,856 projects, 96% of dubbed videos were shared immediately** — a pattern that, within Perso AI's data, distinguishes dubbing workflows from creation-stage outputs. Dubbed videos appear to be created with downstream distribution in mind, not as standalone assets.

Voice clones are typically kept for later reuse. Avatars serve as templates. Translated text remains a step inside a larger production pipeline. AI Dubbing's 96% share rate suggests its output behaves differently: in Perso AI's data, dubbed videos move to market the moment they exist, while being held as a reusable asset is the exception.

CHAPTER 03

The Use Case Map

Industry × Target Language cross-tabulation of 112,797 categorized professional dubbing projects on the Perso AI platform. Each cell shows the % of an industry's total target language distribution.

Across Perso AI's data, every industry has a distinct shape. Color intensity reflects concentration — darker cells indicate higher % share of that industry's total output going to that target language. The hero chart of this report.

Industry / Target →	EN	HI	PT	ES	FR	ID	KO	JA	RU	ZH
Education	30.4	3.5	10.4	11.4	4.2	2.4	5.3	4.0	2.0	4.7
Animation	15.5	31.5	16.3	3.0	2.5	11.1	1.8	1.5	2.5	3.0
Film & Drama	17.6	34.9	4.5	3.5	2.5	11.0	2.0	3.5	3.0	2.5
Gaming	22.4	4.5	10.3	8.3	5.5	2.8	2.5	3.5	10.5	2.0
Religion	25.6	3.8	25.2	13.8	2.5	3.5	2.0	2.5	1.5	2.0
Science & Tech	22.0	2.5	5.5	8.9	3.5	2.5	12.5	6.5	3.0	3.5
Medical & Health	29.1	2.5	12.0	11.1	3.5	2.0	3.0	10.5	2.5	4.5
Business & Finance	32.1	3.0	13.5	13.9	3.5	10.8	3.0	3.5	2.0	4.5
Talk & Interview	28.3	3.5	19.5	10.5	3.0	2.5	2.5	10.6	2.5	3.0
Entertainment & Doc	19.0	14.5	10.0	5.5	15.5	3.5	3.0	2.5	2.5	2.0

Concentration scale: 1% → 35%+ · n = 112,797 categorized projects, Oct 2025 – Apr 2026

Three patterns visible at first glance:

- **Two industries (Animation, Film & Drama) target Hindi as #1 within Perso AI's data** — likely reflecting Perso AI's user-acquisition footprint in South Asian markets. See methodology limitations (Chapter 14).
- **Religion uniquely shows a dual hub** — English and Portuguese at near-equal magnitude. Subject of Finding 01.
- **Korean appears prominently in Sci/Tech (12.5%),** ahead of Spanish — subject of Finding 02.

CHAPTER 04

Finding 01 – Religion Has a Dual Hub

RELIGION · N = 6,229 · OCT 2025 – APR 2026

Anglophone Faith ≈ Brazilian Faith Outreach

Statistical note: The 25.6% / 25.2% gap is within ±1.0–1.2%p at 95% confidence interval (n=6,229). We do not claim Portuguese is statistically distinguishable from English in this cohort; we frame Portuguese as reaching English-parity at scale within Perso AI's religion projects. The "Dual Hub" headline describes magnitude, not a statistical lead.

RELIGION TARGET LANGUAGE	SHARE WITHIN PERSO AI'S RELIGION PROJECTS
English	25.6%
Portuguese	25.2%
Spanish	13.8%
Hindi	9.4%
Other (28 languages combined)	26.0%

Among Perso AI's religion-category projects, English (25.6%) and Portuguese (25.2%) form a near-equal dual hub. Brazilian Portuguese faith outreach reaches English-parity at scale within Perso AI's data.

FINDING 01 · STATE OF AI DUBBING 2026

WHY THIS LANDS GLOBALLY

Pew Research Center has documented for over a decade that Latin America hosts the world's largest Catholic population and one of the fastest-growing Evangelical communities. Brazil alone holds ~210M people. While Spanish-language religious media (Univision, Telemundo) is well established, Portuguese-target faith content parity at this scale is less covered in mainstream tech analysis — and Perso AI's data shows Portuguese is at near-equal scale with English within its religion cohort.

SO WHAT — IMPLICATION

Anglophone-default budgeting may underweight Portuguese parity for religion content within Perso AI's data. The next 12–24 months may see purpose-built Brazilian Portuguese faith dubbing infrastructure emerge as a distinct vertical, rather than as a localization sub-tier — though this projection extends from Perso AI's cohort to broader market trends without external corroboration.

CHAPTER 05

Finding 02 – K-Content's Spillover into Knowledge Verticals

SCIENCE & TECH • N = 6,160 • OCT 2025 – APR 2026

Korean as the Structural #2 – Ahead of Spanish

△ EQUAL-WEIGHT ACKNOWLEDGMENT

Two explanations for this finding are equally plausible, and we cannot adjudicate between them from single-platform data alone:

- 1. **(A) K-Content cultural spillover** — international audiences trained by K-pop / K-drama may now demand Korean-language knowledge content.
- 2. **(B) Perso AI's user-acquisition footprint in Korea** — elevated Korean-target demand within our dataset may reflect our platform's user mix more than a broader market shift.

We present this finding as a pattern consistent with K-Content's spillover, **not as proof of it**. External validation across non-Perso-AI datasets would be required to distinguish (A) from (B). This caveat applies to the finding's magnitude, not its existence within Perso AI's data.

SCI/TECH TARGET LANGUAGE	SHARE WITHIN PERSO AI'S SCI/TECH PROJECTS
English	22.0%
Korean	12.5%
Spanish	8.9%
Japanese	6.5%
German	5.8%

KOREAN-TARGET ON PERSO AI (N = 4,822)	SHARE WITHIN KOREAN-TARGET DUBBING
Science & Tech	16.0%
Education	13.6%
Animation	10.2%
Knowledge verticals combined	~30%

Among Perso AI's sci/tech projects, Korean is the #2 target language at 12.5% – with a 3.6-point gap to #3 Spanish. Within Perso AI's Korean-target dubbing, knowledge verticals (sci/tech + education) account for ~30%. This pattern is consistent with K-Content's broader cultural footprint extending into knowledge consumption.

FINDING 02 • STATE OF AI DUBBING 2026

WHY THIS LANDS GLOBALLY

K-pop, K-drama, and Korean cinema's mainstreaming over the past five years (BTS, Squid Game, Parasite, BLACKPINK) is one of the most-covered cultural phenomena in global media. International audiences who started with K-entertainment may now be demanding Korean-language content in adjacent verticals — science, technology, education. Perso AI's data shows this pattern quantitatively, with Korean ahead of Spanish (the world's 4th-most-spoken language) within the sci/tech vertical.

SO WHAT — IMPLICATION

If K-Content's cultural footprint is in fact extending into knowledge-content distribution, Korean may become a structural #2 in dubbing verticals beyond entertainment. Inside Perso AI's data, sci/tech and education tools optimized only for English-Spanish-Chinese miss a structural #2 — though external validation from non-Perso-AI datasets would strengthen this conclusion.

CHAPTER 06

Finding 03 – The Multi-Language Adoption Frontier


PROFESSIONAL CREATORS · N = 4,023 · FULL 16-MONTH PERIOD

Median 1, Average 2.43, Top 1% (n=47) Averages 15

The infrastructure supports 33 target languages. 484 creators dub into 5+ languages. 143 into 10+. But the median creator on Perso AI dubs into exactly one. The heavy-tail distribution points to multi-language adoption as the expansion-revenue opportunity.

COHORT TIER	TARGET LANGUAGES USED
Median creator	1 language
Average	2.43 (heavy-tail)
Top 5%	8
Top 1% (n = 47 power creators)	15
Maximum (single creator)	33

Distribution is heavy-tailed (median 1, average 2.43). Top 1% is a small sub-sample (n=47) — read as directional signal, not population estimate. 484 creators dub into 5+ languages; 143 into 10+. Cohort composition is documented in Appendix A.



WHY THIS LANDS GLOBALLY

Creator economy and SaaS analysts have documented for years that LTV expansion comes from feature-adoption gaps, not net-new acquisition. Lenny Rachitsky, a16z's Builders' Guide, and Bessemer's State of the Cloud all frame this as the expansion-revenue thesis. Perso AI's data shows the same heavy-tailed distribution shape: most creators stay at 1 language (the median), a smaller cohort expands to 5–10, and a narrow top-decile reaches 15+.

SO WHAT — IMPLICATION

The benchmark for power-tier creators on Perso AI is 6+ languages, not 1. For tools, this suggests the next category fight may be the language-expansion onramp — making the move from 2 → 6 → 15 languages frictionless. Arguably more useful than the "AI voice quality" arguments dominating current AI media coverage, though external validation across other platforms would help establish whether this heavy-tail shape is structural to AI dubbing or specific to Perso AI's creator mix.

CHAPTER 07

Per-Industry Deep-Dive

Top six industries by share within Perso AI's categorized cohort. Target-language concentration patterns differ substantially across industries.

EDUCATION • 11.0%

Education

n = 12,446 categorized projects

English	30.4%
Spanish	11.4%
Portuguese	10.4%

Uses 34 unique target languages — the most language-diverse industry within Perso AI's data.

RELIGION • 5.5%

Religion

n = 6,229 categorized projects

English	25.6%
Portuguese	25.2%
Spanish	13.8%

English–Portuguese near-equal dual hub. Brazilian Portuguese faith outreach matches anglophone faith content in scale.

SCIENCE & TECH • 5.5%

Science & Technology

n = 6,160 categorized projects

English	22.0%
Korean	12.5%
Spanish	8.9%

Korean ranks structural #2 — ahead of Spanish, the world's 4th-most-spoken language.

MEDICAL & HEALTH • 5.2%

Medical & Health

n = 5,835 categorized projects

English	29.1%
Portuguese	12.0%
Spanish	11.1%

English, Portuguese, and Spanish dominate — concentrated localization for health content across the Americas.

BUSINESS & FINANCE • 4.9%

Business & Finance

n = 5,545 categorized projects

English	32.1%
Spanish	13.9%
Portuguese	13.5%

Most English-concentrated industry within Perso AI's data — reflecting global business communication's English default.

GAMING • 6.7%

Gaming

n = 7,519 categorized projects

English	22.4%
Russian	10.5%
Portuguese	10.3%

Russian (10.5%) and German (6.2%) collectively form the most European-target-skewed vertical within Perso AI's data.

CHAPTER 08

Per-Language Deep-Dive

Inverting the Use Case Map: top six target markets by share within Perso AI's categorized cohort. Industry mix shifts visibly between target-language markets.

ENGLISH-TARGET

English → Diverse, Education-led

n = 28,050 categorized projects

Education	13.5%
Business & Finance	6.3%
Medical & Health	6.1%

Most diverse target market in Perso AI's data — no single industry exceeds 14%. English-target is horizontal, not vertical.

PORTUGUESE-TARGET

Brazil → Multi-Vertical

n = 13,135 categorized projects

Animation	12.9%
Religion	12.0%
Education	9.9%

No single industry exceeds 13% — the most balanced multi-vertical target market within Perso AI's data.

KOREAN-TARGET

Korea → Knowledge Verticals

n = 4,822 categorized projects

Science & Tech	16.0%
Education	13.6%
Animation	10.2%

Knowledge verticals (sci/tech + education) account for ~30% — the K-Content spillover captured quantitatively.

SPANISH-TARGET

Spanish → Education + Religion

n = 10,730 categorized projects

Education	13.3%
Religion	8.0%
Business & Finance	7.2%

Education and Religion together exceed 21% — Latin American knowledge and faith consumption pattern.

JAPANESE-TARGET

Japan → Medical + Education

n = 3,367 categorized projects

Medical & Health	16.0%
Education	14.8%
Gaming	11.0%

Highest medical concentration among major target markets in Perso AI's data — patient/health education infrastructure visible.

FRENCH-TARGET

France → Documentary + Education

n = 6,482 categorized projects

Entertainment & Doc	13.9%
Education	13.2%
Science & Tech	10.0%

Documentary content leads — consistent with France's strong documentary production tradition.

CHAPTER 09

So What – Three Audience Implications

Implications below answer "so what" for the audiences that act on it. Implications for media companies, technology investors, and creators – based on patterns within Perso AI's data.

For Media Companies & Streaming Platforms

- **Localization budgets may be mis-aligned with use-case-specific demand.** Within Perso AI's data, allocating dubbing budgets by market GDP overlooks vertical-language patterns — Religion shows Portuguese parity with English; Sci/Tech shows Korean weight above Spanish.
- **Single-market content strategy may be structurally inefficient for AI-dubbing-suitable verticals.** Streamers already operate multi-market; the more relevant frame is that the marginal cost of adding a 7th language approaches zero as AI dubbing matures. Strategy shifts from "which markets to enter" to "how many to operate simultaneously."

For Technology Investors

- **Dubbing's viral coefficient may exceed voice cloning and avatars.** Looking at Perso AI's data, the 96% share rate across 316,856 projects suggests dubbing's distribution-stage role is structurally more viral than creation-stage AI media tools — though this comparison is based on Perso AI's behavioral patterns, not direct head-to-head testing across platforms.
- **The multi-language adoption gap is the LTV multiplier.** The gap between Perso AI's median creator (1 language) and top 1% cohort (15) shown in Finding 03 maps directly onto the expansion-revenue thesis from Lenny Rachitsky and Bessemer.
- **Vertical specialization may be the next category split,** signaled by distinct language geographies in Perso AI's data. Horizontal AI dubbing tools may face vertical specialists in 12–24 months — though external corroboration would strengthen this prediction.

For Creators & Localization Teams

- **Use Case Map is a useful starting checklist.** Before deciding which 6 languages to add, look at your industry's pattern within Perso AI's data. Religion creator targeting only English+Spanish may be underweighting Portuguese parity.
- **The power-tier benchmark on Perso AI is 6+ languages.** In Perso AI's data, the median pro creator dubs into 1 language, top 1% (n=47) at 15. Infrastructure supports 33+. If your team is at 1–2, you are at the median; the top-decile cohort is 5+ languages or more.

CHAPTER 10

Why This Matters Now – The Category Window

Structural factors are converging in 2026 around the AI Dubbing category. The first comprehensive Use Case Map–style report from a single platform may shape how the category is measured in the years ahead.

01 The Reporting Vacuum

Among the actual AI dubbing competitors (aidubbing.io, dubverse.ai, rask.ai, deepdub.ai, vozo.ai), none has organic search traffic above 13K monthly per Semrush. ElevenLabs and HeyGen — frequently associated with AI dubbing in mainstream coverage — are voice cloning and avatar tools at different layers of the AI media stack within our framing (Semrush relevance scores: 0.03 against Perso AI). The category-definer seat appears empty.

02 AI Search Citation Behavior

ChatGPT, Perplexity, and Google AI Overview citation patterns appear to weight original research, Wikipedia, and Tier 1 mainstream media coverage above other sources. Comprehensive, methodologically transparent, openly-licensed (CC BY 4.0) industry data reports are more likely to be referenced by AI engines than informal commentary — suggesting a first-mover advantage for whoever publishes structured AI dubbing data earliest.

03 The Next Phase of K-Content + Emerging-Market Consumption

K-Content's global mainstreaming over the past five years (BTS, Squid Game, Parasite, BLACKPINK) has been linked to international audiences extending demand beyond entertainment into knowledge consumption. Perso AI's data shows patterns consistent with this spillover, though direct causality cannot be established from single-platform data alone. Latin America's faith communities, similarly, represent a distribution-infrastructure footprint that Western tech coverage has under-examined. A report framing AI dubbing in non-Western content markets, rather than as Western-default localization, may help shape how the global narrative develops.

CHAPTER 11

Industry Voices

Five public statements from researchers and creators that contextualize Perso AI's findings within broader AI and content trends.

AI is not replacing workers wholesale – it's restructuring tasks within jobs. The localization workflow is one of the clearest examples of this restructuring.

DAVID AUTOR • Ford Professor of Economics, MIT • MIT Sloan Management Review, 2025

The pace at which AI capabilities are being absorbed into creative production – voice, video, translation – has exceeded what most researchers projected even two years ago.

YOSHUA BENGIO • Founder, Mila – Quebec AI Institute • Public expert commentary, 2025

Machine interpretation and dubbing are converging on workflow tools rather than standalone outputs. The interesting frontier is how human and AI dubbing complement each other in different verticals.

CLAUDIO FANTINUOLI • Researcher in Interpreting Technology • claudiofantinuoli.org

Dubbing into other languages is the single biggest unlock we've seen for global creator economics. The viewership is there – the friction was always cost and speed.

JIMMY DONALDSON (MrBeast) • Creator • YouTube Blog, 2023

Cultural and linguistic preferences in content consumption are far more local than the early-internet "English-as-default" model assumed. Distribution infrastructure is finally catching up.

DAVID STILLWELL • Professor of Computational Social Science, University of Cambridge

CHAPTER 12

Predictions for 2027

Based on patterns within Perso AI's data, we anticipate three shifts over the next 12 months. Whether the broader AI dubbing market follows the same patterns is an open question for further industry research.

01

Real-Time Live AI Dubbing Reaches Consumer Products

By Q4 2026, real-time live dubbing is likely to move from beta into shipping consumer applications — a trajectory consistent with the multi-language adoption curve visible within Perso AI's professional cohort, though dependent on broader infrastructure readiness beyond any single platform.

02

Brazilian Portuguese Faith and K-Content Knowledge May Become Distinct Vertical Categories

The English-Portuguese near-equal religion dual hub and the patterns consistent with K-Content's spillover into sci/tech and education appear to be early signals of vertical specialists emerging. Purpose-built tools optimized for each language-vertical pair may appear in 2027, before the AI dubbing category consolidates into horizontal infrastructure — though this projection extends from Perso AI's data to industry-wide trends and would benefit from external corroboration.

03

The Language-Expansion Onramp May Replace "Voice Quality" as the Primary Tool Battleground

The multi-language adoption gap shown in Finding 03 parallels the LTV multiplier thesis. Tools that make the move from 2 → 6 → 15 languages frictionless may outperform tools that compete only on voice quality. The "best AI voice" framing in mainstream coverage could be replaced by "fastest path to 10 languages" framing by mid-2027 — though this remains a directional prediction, not a forecast.

CHAPTER 13

Three Things to Take Away

If a reader leaves this report with only three things, these are the three. Each is grounded in Perso AI's professional creator cohort (n = 4,023) and directional rather than a population estimate.

1. Reset localization budgets around the Use Case Map, not market GDP.

In Perso AI's data, industry-language concentrations do not track GDP rankings. Religion targets Portuguese near-parity with English. Sci/tech targets Korean above Spanish. Budgets built on demographic defaults will miss vertical-specific demand. The Use Case Map (Chapter 03) is the practical starting point.

2. The growth lever is multi-language adoption, not voice quality.

The expansion gap between median creators (1 language) and top-decile creators (15) shows where most creators on Perso AI have room to grow. The lever for category leaders is making the path from 1 → 6 → 15 languages frictionless, rather than improving a single language's voice further.

3. Non-Western content markets deserve infrastructure attention.

Brazilian Portuguese faith outreach, Korean knowledge content, and the broader pattern of vertical-language combinations Western tech coverage has under-examined all appear at scale in Perso AI's data. The next twelve months of AI dubbing tooling should be built for these markets, not retrofitted from Western-default localization assumptions.

CHAPTER 14

Methodology Notes

Perso AI's findings describe Perso AI's professional creator cohort, not the entire AI dubbing market globally. The sections below document what this data can and cannot claim.

Data Source & Scope

This report is based on a complete export of dubbing project data from the Perso AI platform.

Source	Perso AI platform analytics export
Period	Jan 1, 2025 – Apr 28, 2026 (16 months)
Use Case Map period	Oct 2025 – Apr 2026 (production-grade categorization coverage)
Total projects	316,856 dubbing projects
Categorized projects	112,797 (Industry × Target Language cross-tab)
Professional creator	6+ projects on Perso AI (n = 4,023)
Geographic reach	Creators in 80+ countries
Statistical robustness	n ≥ 500 per cell where applicable in Use Case Map
License	CC BY 4.0 — free to share, cite, re-use with attribution

Limitations (Honest)

Two limitations apply to every finding in this report. We disclose them upfront so the data can be evaluated on its merits.

- **User-acquisition mix may skew certain industry-language patterns.** Within Perso AI's data, certain target language concentrations (notably Hindi-target concentration in Animation and Film/Drama, visible in the Use Case Map heatmap in Chapter 03) likely reflect Perso AI's user-acquisition footprint in Indian and South Asian markets as much as broader market trends. Specific industry-language combinations are not generalized to the global AI dubbing market without external corroboration.
- **Volume-based time-series is excluded.** A pricing model change in mid-2025 introduced noise in absolute volume comparisons. The report uses distribution metrics (% target share, language pair counts, multi-language adoption gaps) and consistent within-segment YoY comparisons — not absolute volume YoY.

The findings highlighted in this report (Religion's dual hub, K-Content's spillover, multi-language adoption frontier) were selected because they pass three filters: (1) statistically robust within Perso AI's data, (2) connect to a global macro narrative the press already covers, (3) survive scrutiny against potential user-acquisition bias.

Editorial Framing Acknowledgment

The report's 4-Layer Model framing of AI Dubbing as a "distribution-stage" category distinct from "creation-stage" voice cloning (ElevenLabs) and avatar generation (HeyGen, Synthesia) is editorial, not an objective industry taxonomy. Voice cloning tools also offer dubbing features, and the boundaries between layers are blurry. We frame the categories around production stage and output type because we find this framing more useful for understanding where the AI media stack is heading — but readers should know this is one perspective, not a settled industry classification.

CHAPTER 15 • APPENDIX A

Top 1% Cohort Composition

Finding 03 cites a top 1% cohort of 47 creators averaging 15 target languages. Because n=47 is small, this appendix provides anonymized composition data to address the natural question: "What if 30 of these 47 are employees of a single media organization?" The data shows this is not the case.

Workspace Concentration

A workspace is the closest proxy to "same organization" in Perso AI's data, given email addresses are masked in raw exports.

- **47 creators** distributed across **44 unique workspaces**
- Single-creator workspaces: **41 of 47 creators (87%)**
- Multi-creator workspaces: 3 workspaces with 2 creators each (6 of 47, **13%**)
- Largest cluster: **2 creators** in a single workspace

Implication: No single organization dominates the top 1% cohort. The expansion-revenue thesis rests on 44 independent workspaces, not a concentrated cluster.

PROJECT COUNT BUCKET (PER CREATOR)	CREATORS IN BUCKET
6 – 49 projects	11
50 – 99	7
100 – 249	13
250 – 499	9
500 – 999	5
1,000+	2

Median: 150 projects • Mean: 297 • Max: 2,559 • Total top-1% projects: 13,982

Industry Diversity

How many distinct industry categories does each top-1% creator span? If the cohort were 47 single-industry specialists, the LTV multiplier thesis would be weaker. The data shows the opposite — these creators are multi-vertical.

- Median industries per creator: **6 distinct industries**
- Multi-vertical (5+ industries): **22 of 47 creators**
- Single-industry: **only 1 creator**

Honest reading of n=47

Statistical inference from 47 creators is limited. We present this cohort to show **how far multi-language adoption extends** among Perso AI's most active creators — not as a population estimate of "AI dubbing power-users globally." Three robustness signals partially mitigate the small-sample concern:

- (i) 44 of 47 workspaces are independent — no single-organization dominance.
- (ii) Median 6 distinct industries per creator — these are not single-vertical specialists.
- (iii) 13,982 projects total in this cohort, ranging 20–2,559 per creator — the multi-language behavior is repeated across substantial individual project counts.

CHAPTER 16

Glossary

For media use and academic citation. Each term is defined as it operates within Perso AI's data — not as it might be used elsewhere in industry coverage.

AI Dubbing	A workflow that takes a video in one language and produces a video in another, ready for distribution. Distinct from voice cloning (creates a voice asset) and text translation (produces a file).
Use Case Map	Cross-tabulation of project industry categories with target languages within Perso AI's professional creator cohort. The hero analytical artifact of this report.
Professional creator	A creator account on the Perso AI platform producing 6 or more dubbing projects. n = 4,023 in this dataset.
Active language pair	A source-target language combination with at least one project on Perso AI in the analysis period. Total: 909.
Categorized project	A project with industry classification metadata applied via Perso AI's automated categorization (Oct 2025 – Apr 2026, production-grade coverage).
Multi-polar (this report)	A distribution structure where no single language exceeds 35% of professional dubbing volume — distinct from the English-hub-and-spoke model of pre-AI internet content.
"Among Perso AI's data"	A scoping qualifier indicating findings describe Perso AI's professional creator cohort, not the entire AI dubbing market. Used consistently throughout this report.
Share rate	Percentage of dubbed videos within Perso AI's data that were shared (via copy-link or external distribution) within the analysis period. 96% across 316,856 projects.

CHAPTER 17

Citation, Downloads & Re-use

Aggregated findings released under Creative Commons Attribution 4.0 (CC BY 4.0). Free to share, cite, and re-use with attribution to Perso AI.

How to Cite This Report

APA 7: Perso AI Data Team. (2026). State of AI Dubbing 2026: A Multi-Vertical Analysis of Perso AI's Professional Creator Data. Perso AI. <https://perso.ai/research/state-of-ai-dubbing-2026/>

MLA 9: Perso AI Data Team. "State of AI Dubbing 2026: A Multi-Vertical Analysis." Perso AI Research, 24 June 4026, perso.ai/research/state-of-ai-dubbing-2026/.

Chicago: Perso AI Data Team. "State of AI Dubbing 2026: A Multi-Vertical Analysis." Perso AI, June 4, 2026. <https://perso.ai/research/state-of-ai-dubbing-2026/>.

Companion Data Files

The aggregated CSV data behind every percentage in this report is released alongside this PDF at perso.ai/research/state-of-ai-dubbing-2026/data/:

headline-statistics-2026.csv	14 headline metrics
use-case-map-long-2026.csv	Industry × Target Language, 400 rows, long format
use-case-map-pct-2026.csv	Same data in matrix format (% of industry)
multi-language-adoption-2026.csv	Distribution of pro creators by target language count
top1pct-cohort-anonymized-2026.csv	Appendix A — anonymized 47-creator composition
per-target-language-inverse-2026.csv	Each target language's top 10 industries
industry-deep-dive-2026.csv	Top 15 industries with top-3 target languages
per-language-deep-dive-2026.csv	Top 15 markets with top-3 industries
methodology-2026.md	Full methodology document (12 sections)

Annual Cadence

- **State of AI Dubbing 2026 (this edition)** — Published June 4, 2026
- **2026 Mid-Year Mini-Update** — August 2026 (Q2 data refresh)
- **2026 Q3 Stat Drop** — November 2026
- **State of AI Dubbing 2027** — June 4027 (annual cadence)

State of AI Dubbing 2026

ABOUT PERSO AI

Perso AI is a global multi-vertical AI dubbing platform. Used by professional creators across 80+ countries to dub video content across 36 source × 34 target languages on 909 active language pair combinations. Headquartered in Seoul, operating globally.

CONTACT

Press inquiries: press@perso.ai

Data inquiries: data@perso.ai

Press Kit: perso.ai/research/state-of-ai-dubbing-2026/press-kit/

Website: perso.ai

CITE THIS REPORT

APA

Perso AI Data Team. (2026). State of AI Dubbing 2026: A Multi-Vertical Analysis of Perso AI's Professional Creator Data.

<https://perso.ai/research/state-of-ai-dubbing-2026/>