Does Formal AI Training Produce More Learning-Oriented Users?

Evidence from the Central African Republic, Rwanda, India, and the UAE

Executive Summary

Governments around the world are rolling out formal AI-training programs with the goal of teaching AI skills to citizens. Programs typically assume, explicitly or implicitly, that trained users will engage with AI more effectively afterwards: as tools for collaboration and learning rather than as shortcuts to finished work. This is particularly true for programs aimed at younger generations. For instance, the UAE's mandatory AI lessons have devoted roughly a quarter of their content to ethics and responsible use, and Rwanda's partnership with Anthropic centers on the "Socratic mentor" Chidi, a model designed to give thoughtful questions rather than direct outputs. This post tests the assumption that trained users will engage with AI more reflectively by comparing claude.ai usage patterns across four countries that cover a range of national training policies: the United Arab Emirates, which has made AI instruction a mandatory part of its K–12 curriculum; Rwanda, with its Anthropic training partnership; India, which runs large, nationally coordinated programs; and the Central African Republic, which has no formal training landscape.

For the three countries where Claude is available, I use three weekly snapshots of the Anthropic Economic Index (AEI) from August 2025, November 2025, and February 2026. For the Central African Republic, I collected 348 chats from roughly 50 University of Bangui students over a one-week window in January 2026, through my pre-existing relationships in the UNC–State Department Diplomacy Lab partnership. I reformatted the CAR data into the Economic Index schema (collaboration-style clusters, request categories, Bloom's levels) so that all four countries sit on the same axes, hand-checking each entry after an initial Opus 4.7 classification pass. To widen the student-to-student comparison I also pull in aggregate data from Anthropic's Education Report (AER).

I design a central measure called the Learning Orientation Index (LOI), constructed from the Economic Index's classification of "collaboration style":

$$LOI = \frac{L + F}{D + T}$$

$where\ L$ = % of requests classified as "learning"

F = % of requests classified as "feedback loop"

D = % of requests classified as "directive"

T = % of requests classified as "task iteration"

This fraction captures the ratio of reflective, understanding-oriented interactions (the numerator, e.g. "why does this work," "does this argument hold up") to execution-oriented ones (the denominator, e.g. "write

this paragraph," "fix this bug"). A LOI above 1 means a user primarily seeks understanding; below 1 means they primarily issue instructions.

While directive AI use has its place, I argue reflective engagement better guarantees users' economic value in a future with increasingly independent, capable AI. The skills we typically associate with AI literacy, like interrogating the model's output and or applying independent thinking for further iteration, are closely related to learning-oriented use. Furthermore, a learning orientation serves as a defense against model failures such as hallucination, as habitual questioning helps users catch errors that directive users who trust the output tend to miss.

Findings

1. Formal AI training correlates with lower learning orientation.

All four countries are task-dominant (LOI < 1), but the ordering runs opposite to the level of training. The CAR has the highest LOI (0.624), driven by the highest share of learning requests (23.6%) and the lowest share of task iteration (16.7%). Rwanda, whose partnership is the most education focused of the four, has the lowest LOI (0.361), driven by the highest directive share (43.3%). India (0.503) and UAE (0.444) fall between. Rwanda produces ~19% more directive prompts than the CAR.

2. Students are not inherently more learning-oriented than general populations.

I investigate the obvious alternative explanation, that student populations would inherently have a higher LOI, with two checks. Firstly, amongst 128 countries in the February 2026 AEI release, the median LOI is 0.606. We would expect a student-only sample to be at the upper range of LOI: instead, the CAR sits at the 57th percentile — essentially the median — while the three trained countries have some of the lowest LOIs in the sample (India 28th, UAE 19th, Rwanda 2nd). For reference, CAR student LOI (0.624) is nearly identical to that of Claude users in the US (0.625). Secondly, using the Education Report's metric of Direct/Collaborative, CAR students (46/51) look almost identical to the Education Report's student sample (47/53) and slightly less collaborative than US all-users (39/61) — meaning a mixed working-age population is actually *more* collaborative than either student sample. Population composition is unlikely to be driving the comparative difference in LOI.

3. Subject matter does not account for the gap in learning-orientedness.

Another possible explanation is that requests related to education tend to have higher learning orientedness compared to more directive areas, like coding. I test this by fitting an ordinary least squares model on 119 countries with valid data, predicting LOI from the share of each country's requests related to education vs. software. This yielded a very limited relationship ($R^2 \approx 0.10$, very shallow residuals), indicating that subject matter does not reliably explain how learning oriented user requests are.

In fact, the CAR's actual LOI exceeds its predicted LOI by 0.25; and all three trained countries fell 0.16 to 0.21 below prediction (India 13th percentile, UAE 7th, Rwanda 3rd). Even dropping the CAR student population as a possible outlier only shifted each trained country's residual by less than 0.02. A further check using Economic Index's use case data indicates Rwanda has by far the highest coursework share of

the three AEI countries (35%, vs. India's 21% and UAE's 15%), yet the lowest LOI. Task success is approximately uniform across the three countries, ruling out the possibility of repeat requests inflating directive usage. This suggests subject matter is unlikely to be the explanation for the gap in LOI.

4. Students in different contexts use AI for vastly different subjects.

CAR students allocate 52% of conversations to humanities and writing, 38% to business and career tasks, and just 10% to technical or quantitative work. Education Report students are the near-inverse: 69% technical/quantitative, 19% business/career, 13% humanities/writing. The same pattern shows up on Bloom's taxonomy: CAR concentrates in Creating (69%) and Understanding (18%); the AER sample concentrates in Analyzing (30%), Creating (40%), and Applying (11%). This is a substantive difference in what students are doing with AI across regions and possible evidence of different subjects living at different Bloom's Taxonomy levels.

Caveats

1. While some interaction styles can seem more oriented towards long-term learning, it is still difficult to directly measure learning outcomes; a student who asks AI to "write this paragraph" and studies the result carefully may learn more than one who asks "explain this" and moves on.
2. The CAR sample is 348 requests from roughly 50 Diplomacy Lab students at one university over one week. This is a relatively small sample size with possible selection bias in who opted into the project, though to my knowledge it is also the first student-level AI-usage data from the CAR, where Claude is technically blocked.
3. The three Economic Index country samples are full Claude user bases, not students; the Education Report comparison partially addresses this but does not have the same level of detail as the AEI data. The assumptions this paper relies on in order to make three-way comparisons are documented more fully below.
4. Training programs in India, the UAE, and Rwanda are young and uneven in coverage: "mandatory K–12" in the UAE does not yet mean that the adult Claude user base has passed through the curriculum, and the LOI gap may widen, narrow, or reverse as trained cohorts age in.

Implications and Next Steps

This paper provides evidence suggesting against the hypothesis that formal training moves users toward using AI as a learning partner instead of as a task executor. If the goal of a national AI-training program is to turn AI into a scaffold for student thinking rather than a substitute for it, the current generation of programs may not be achieving that goal. This result does not condemn training programs as counterproductive. It shows that more reflective, learning-oriented use is not currently visible in revealed usage data, and suggests some ways training can promote more thoughtful interactions.

Long-term tracking of training cohorts as participants age is one obvious next step. Other possible next steps to better understand what affects user learning outcomes are the release of more specific data on education users, the introduction of control groups to test the effect of different pedagogies, and long term, formal processes to directly measure how well users learn from using AI.

Introduction

Artificial intelligence has quickly become a global priority, with governments building infrastructure to avoid falling behind. The UAE is rolling out a mandatory AI curriculum across all public schools for the 2025–2026 academic year, reaching roughly 400,000 students, and has separately partnered with Google to give every university student a year of Gemini Pro. Many national AI training programs, particularly those targeted at students, are built on the assumption that structured training will improve how citizens actually interact with AI. The UAE's Minister of Education, Sarah Al Amiri, framed the mandatory K–12 curriculum as "a strategic step that modernises teaching tools and supports a generation of young people who understand tech ethics and can create smart, locally relevant solutions," warning against the risk of urgency to adopt AI outpacing "the capacity required to use AI wisely, effectively and ethically." India's YUVAi program, run by the National e-Governance Division under the IndiaAI Mission, states its objective as empowering students to "become human-centric designers and users of AI". Singapore's AI Singapore initiative has committed to making 25,000 professionals and 100,000 students "data fluent" through its LearnAI programs. Programs like these are both trying to increase familiarity with AI and to shape their populations' usage.

However, the evidence base for the effect of training on AI behaviours is limited. Most existing work has focused on output quality or on self-reported attitudes: for instance, a recent Microsoft field study at Gap Inc. finding that an AI "mindset" training produced top-quality outputs at 77% versus 61.8% under standard feature training. We have little public evidence on how trained populations interact with AI differently from users with limited training. This paper takes a first pass at that question by using Anthropic Economic Index data from three countries with formal national AI training (UAE, India, Rwanda) alongside a small field sample from the Central African Republic. I ask whether formal national trainings shift users towards more learning-oriented, reflective engagement. The goal is to start a more empirically grounded conversation about the effect of training on participants, so that future programs can be designed and targeted with a clearer read on their own effects.

While directive AI use has its place, teaching users to engage with AI more reflectively matters for two reasons. Firstly, many of the skills traditionally associated with AI literacy come with learning-oriented use: e.g. training users to interrogate the model to build their own understanding or iterate on its output with independent thinking. By contrast, a user whose contribution is "write this paragraph" or "fix this bug" is doing work that anyone could do. Second, learning orientation is a defense against model failure. Directive use tends to trust the output, leaving users vulnerable to hallucination, reward hacking, or other misaligned outputs that will only get more fluent and plausible. A user who habitually asks "why does this work" or "does this argument hold up" will catch errors that more directive users miss. Training programs that succeed in moving users toward learning-oriented engagement are therefore producing users whose value to their economies, and whose resistance to AI failure modes, continue to increase as AI capabilities improve.

Methods

Economic Index Comparisons

This project began as a way to increase my familiarity with Anthropic's Economic Index, which is a taxonomy of every Claude conversation within a given week. Once I saw that each conversation had a "collaboration style" classification around how the user interacted with the model, I became curious as to what factors affected general collaboration style. While demographic data would be interesting — e.g. age, occupation — these were limited due to privacy concerns. The index did, however, include country data. Thus, I decided to investigate the relationship between national AI training landscapes and how their users interacted with AI (i.e. their users' collaboration style).

For the sake of time and data size, I narrowed my investigation to three countries: the United Arab Emirates, India, and Rwanda. I selected these countries as an attempt to represent three different points on the possible spectrum of AI training: the UAE being the most extreme end of coverage with a mandatory K-12 AI curriculum, and India (large-scale national programs) and Rwanda (with partnerships between local deliverers and Anthropic) representing a form of less widesweeping coverage.

I included some additional data from the Central African Republic. This presented a unique opportunity as I had previously conducted AI policy research in conjunction with students from the University of Bangui, and relied on these personal relationships to gather data. I asked 28 of these students for their AI requests over the course of a week. I then formatted submissions into the Economic Index schema, using an initial Opus 4.7 classification then hand-checking each of the 348 entries, to enable direct comparison with Anthropic's data for the other three countries. Part of the work we did with the Bangui students was to help set up a student-organised AI literacy training program. Thus, the CAR can be seen as the other side of the AI training spectrum, with largely informal/nascent programming. The difference between these four cases are visualised below.

Country	AI training landscape	Data source	Sample size
UAE	Mandatory K–12 AI curriculum (most widespread coverage)	AEI country release (Feb 5–12, 2026)	4,704 requests
India	Large-scale national programs (IndiaAI Mission, AI4India)	AEI country release (Feb 5–12, 2026)	61,627 requests
Rwanda	Government/NGO partnerships with Anthropic (education-focused)	AEI country release (Feb 5–12, 2026)	504 requests
CAR	No formal program; student-organised literacy efforts	Field data: ~50 Bangui students; hand-reformatted to AEI schema	348 requests (Jan 11–18, 2026)

Learning Orientation

One qualitative pattern I noticed when speaking to students in the CAR was they often treated AI agents as instructors (e.g. teach me X) rather than task executors (e.g. do this assignment), unlike many of my friends in North American/European universities. Particularly in academic applications, this seemed like a more responsible usage of AI that would contribute to learning, rather than just task completion. The Economic Index data around "collaboration style" provided the opportunity to more rigorously evaluate this qualitative observation.

I constructed the following Learning Orientation Index (LOI) from the Economic Index's collaboration style percentages:

$$LOI = \frac{L + F}{D + T}$$

where L = % of requests classified as "learning"

F = % of requests classified as "feedback loop"

D = % of requests classified as "directive"

T = % of requests classified as "task iteration"

The numerator captures reflective, understanding-oriented interactions (asking to clarify a previous output, outlining a concept, etc); while the denominator captures execution-oriented ones. Essentially, an LOI above 1 means users primarily seek understanding or explanation; below 1 means they primarily issue instructions. I thus became interested in how LOI scores relate with national AI training norms.

Anthropic Education Report

One flaw with just relying on comparisons to the Economic Index is that the CAR data is drawn from university students in Bangui, while the UAE, India, and Rwanda Economic Index data is drawn from each country's full Claude.ai user base. Students are a narrow slice of any country's population and likely

have different usage patterns. Comparing CAR students against general-population samples in the other three countries therefore risks attributing to national AI training what is in fact attributable to being a student. To partially correct for this, I incorporate aggregate data from Anthropic's Education Report (Education Report, April 2025), which classifies 574,740 conversations from .edu email addresses and gives a student-to-student reference point.

The Education Report is also an imperfect comparison, as it does not have a country-level breakdown and data is limited to aggregate percentages inside a blog post rather than as a downloadable dataset. One notable issue is its taxonomy differs slightly from the Economic Index's (more on this in "Differing Classifications"). However, being able to make any comparisons between student demographics, even if imperfect, is still interesting, and so I elected to include the Education Report data along with the Economic Index data to get a fuller picture of how different users interact with Claude.

Differing Classifications

The Education Report's primary taxonomy of student prompts is a 2×2 matrix of Direct vs Collaborative, crossed with Problem-Solving vs Output-Creation. Direct means the student seeks an answer or content with minimal engagement while Collaborative means a multi-turn exchange; Problem-Solving means the student is working toward an answer (e.g. answering a question, debugging code), while Output-Creation means the student is producing a deliverable (e.g. a slideshow, writing code)

Classifying prompts as "direct" or "collaborative" is a seemingly close but ultimately imprecise analog to the Economic Index's more detailed labels. This is because not every multi-turn exchange is learning-oriented (e.g. a student who iterates with Claude across several turns to produce an essay may still be offloading the writing), which the Education Report also acknowledges.

The primary approach this paper takes to make three-way comparisons is to collapse the different Economic Index collaboration styles into the Education Report's Direct/Collaborative split by grouping directive and validation into Direct, and learning, feedback loop, and task iteration into Collaborative, with "none" and "not classified" dropped. While this continues to have the issue of the Education Report classifying task-oriented prompts as Collaborative, it does allow us to group Economic Index/CAR data to compare along the axes that the Education Report uses. Additionally, to match the Education Report Bloom's taxonomy classification, I hand-mapped each of the 348 CAR requests to a Bloom's level based on task structure, first using Opus 4.7 then manually checking each entry. It is unclear whether this was Anthropic's methodology as well, and so some misclassifications may have arisen.

In general, I believe the best interpretation of three-way comparisons to be "what does each framework say about how the same population interacts with AI". Where the three frameworks agree, the trend is more robust; where they disagree, a closer look may be necessary, as it could trace back to methodological differences rather than underlying behavior.

Results

1. Formal AI training correlates with lower learning orientation

While all four countries are task-dominant, with LOI scores below 1, the CAR, which has no formal AI training program at all, has the highest LOI (0.624), driven by the highest learning-mode share in the sample (23.6%) and the lowest task-iteration share (16.7%). Rwanda, whose Anthropic-government-NGO partnership is the most education-focused program of the four, has the lowest LOI (0.361), driven by the highest directive share (43.3%). UAE and India fall in between.

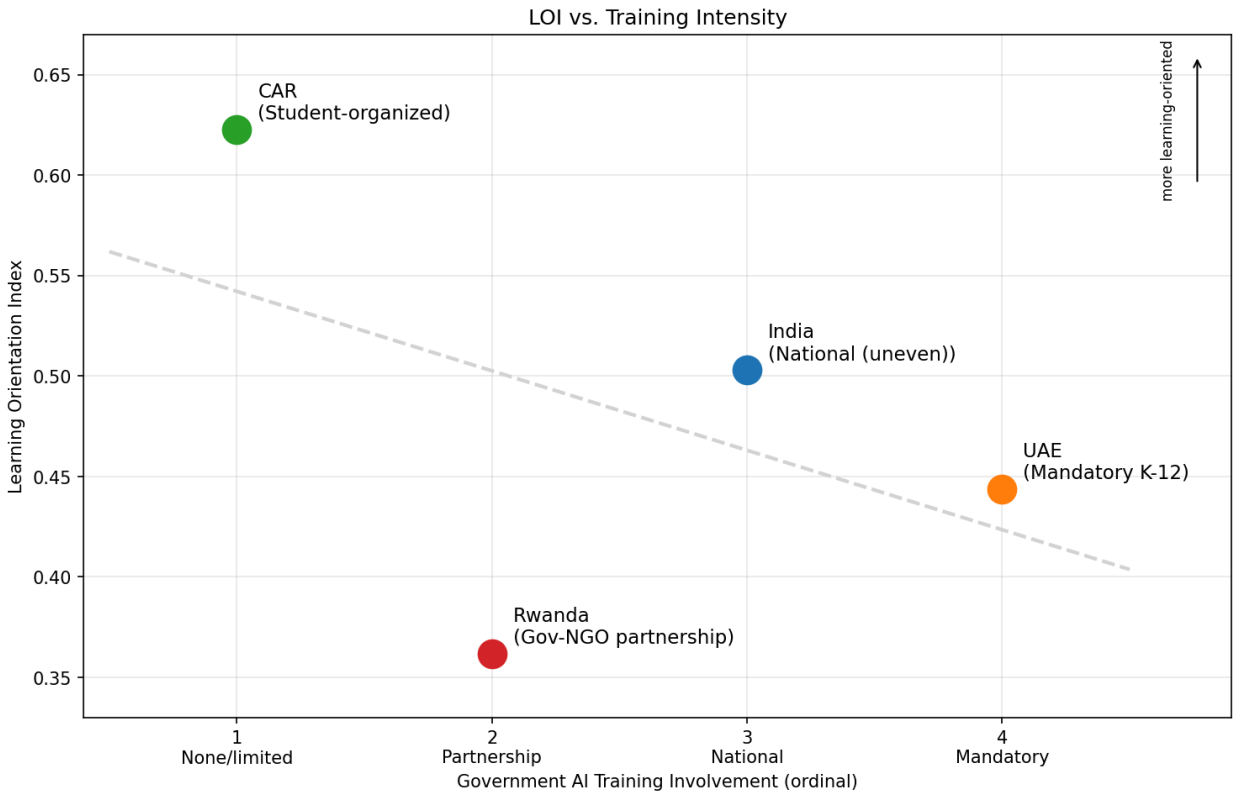

Note: LOI is theoretically bounded at [-1, 1]; y-axis zoomed to observed range.

Learning Orientation Index by country, ordered by training intensity. All four countries score below 1 (task-dominant). The ordering runs counter to the expectation that formal training makes users more learning-oriented (i.e. raises LOI). Line of best fit is included.

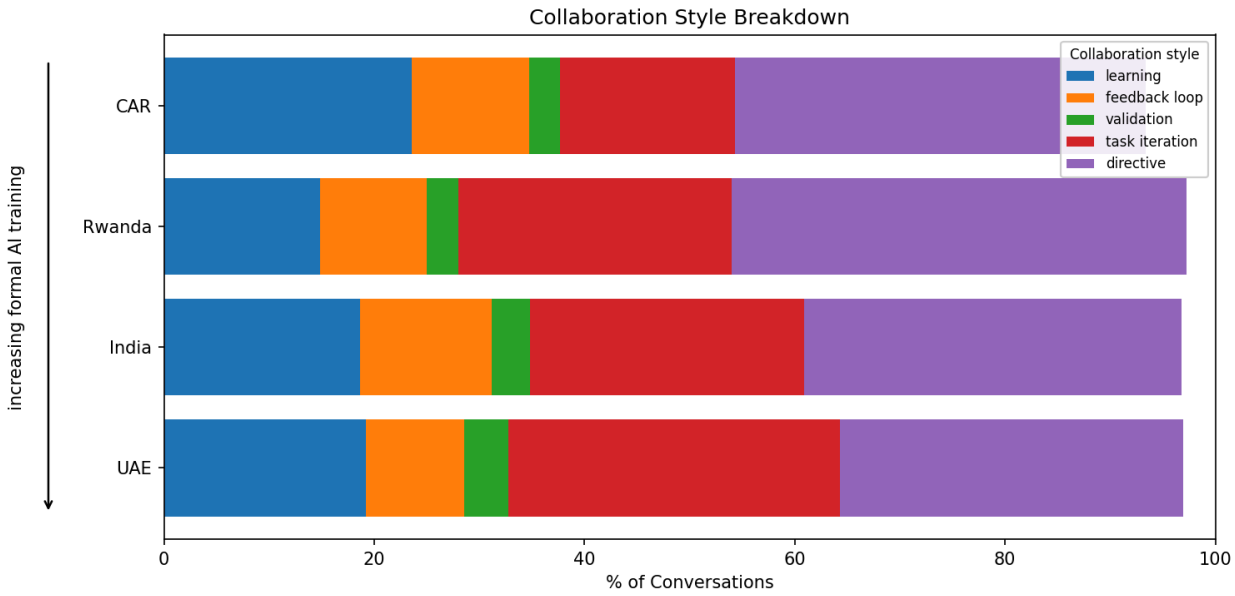

Collaboration style distribution underlying each country's LOI. CAR's LOI (learning + feedback) is the largest of the four; Rwanda's denominator (directive + task iteration) is the largest.

A higher LOI means in general, a user is spending more of their AI time asking for conceptual explanations ("why does this happen"), soliciting feedback on their own reasoning ("does this argument hold up"), and iterating through dialogue rather than issuing task commands ("write this paragraph," "fix this code"). In education-research terms, a higher LOI looks closer to scaffolded reasoning, where the model supports the student's thinking, and a lower LOI looks closer to cognitive offloading, where the student gets the end result but the cognitive work moves to the model.

This gap is significant, indicating that the more formal the country's AI training landscape, the more its users engage with AI as a task executor rather than a learning partner. For instance, for every learning-oriented interaction in the CAR sample, users issue ~1.6 directive ones; India's sample shows ~2.0, the UAE's ~2.3, and Rwanda's ~2.8. In other words, Indian users produce ~8% more directive prompts than those in the CAR, the UAE ~13% more, and Rwanda ~19% more. This trend runs opposite to the intuition that more training necessarily increases a more responsible, learning-oriented engagement amongst users.

Notably, Rwanda is an outlier even within these three countries: its training landscape is less intensive than India's or the UAE's, yet its directive ratio is the highest of all four countries. This may be due to Rwanda's program design differing from those of the UAE and India. Anthropic’s three-year MOU with the Rwandan government deploys Chidi, a Claude-based learning companion separate from claude.ai meant to guide learners through questions rather than providing direct answers. Chidi operates through its own interface, while the AEI captures only general claude.ai usage. It is therefore plausible that learning-oriented interactions are being routed through Chidi, leaving claude.ai usage in Rwanda disproportionately skewed toward directive, task-executor behavior, which is reflected in an outlying lower LOI.

2. Students are not inherently more learning oriented

One alternate explanation for the previous section is that the CAR sample is composed of students, who may inherently approach AI with more of a "learning orientation" given their daily obligations, while the three Economic Index country samples contain mixed working-age populations. However, two further checks indicate this may not be accurate.

First, we can observe global benchmarking. If a student population has an inherently inflated LOI, CAR should sit well above the global median of mixed-population country samples, but it does not. Across 128 countries in the February 2026 Economic Index release (the most recent data to the January 2026 CAR data), the median LOI is 0.606. CAR sits at the 57th percentile, essentially the global median. In addition to this, the three countries with formal training are below median: India at the 28th percentile, UAE at the 19th, Rwanda at the 2nd. Among the 33 countries with usage volumes comparable to CAR and Rwanda (200–600 weekly requests), CAR's LOI is slightly higher than average while Rwanda's is 32nd of 33rd. While it is true students are likely more learning oriented, the negative correlation in Section 1 may be evidence that formally-trained countries are unusually task-focused relative to populations with less training.

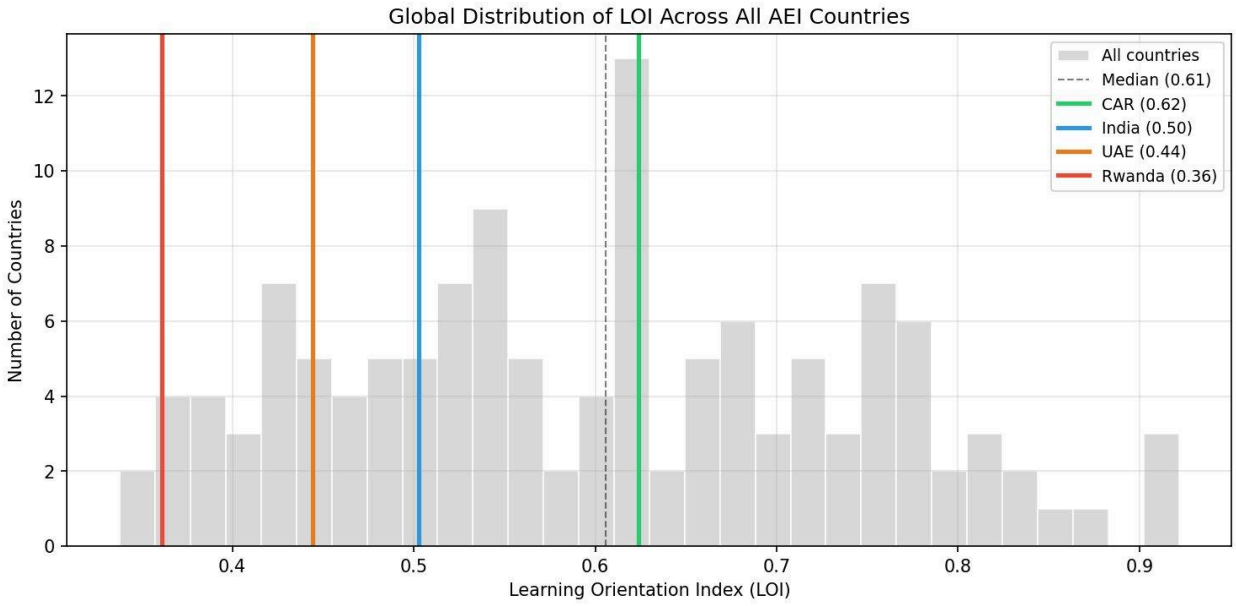

LOI distribution across 128 countries based on measurements by the February 2026 Economic Index. CAR sits near the global median; Rwanda, UAE, and India are all in the lower tail.

Second, we can do direct student-to-student comparison against the Education Report. After collapsing Economic Index collaboration styles into the Education Report's Direct/Collaborative split (see Methods), The CAR is 46% Direct / 51% Collaborative / 3% uncategorised (due to ambiguity). The Education Report reports 47% Direct / 53% Collaborative for its sample of 574,740 student conversations. The US all-Claude-users population (February 2026 Economic Index, mixed working-age and student) sits at 39% Direct / 61% Collaborative, meaning US working-age users are more collaborative than either student population. We can repeat this process to take a global average of all countries, landing at a 40% Direct/

60% Collaborative split.

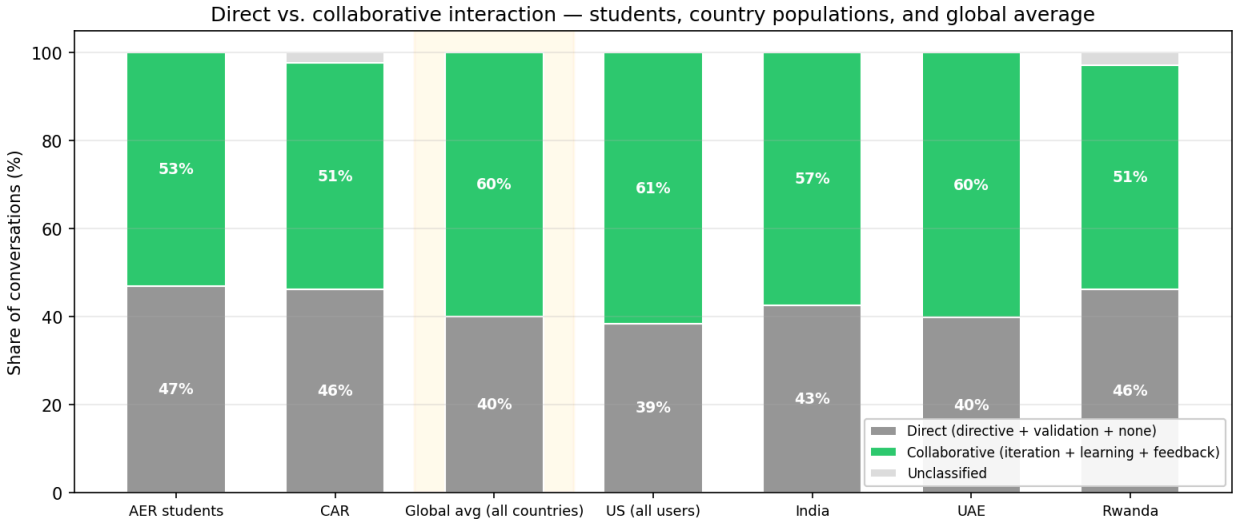

Direct vs. collaborative share across CAR students, select country-level populations, and the global country-level average.

Most strikingly, CAR student LOI (0.624) is almost identical to US all-Claude-users LOI (0.625) and UK all-Claude-users LOI (0.617).

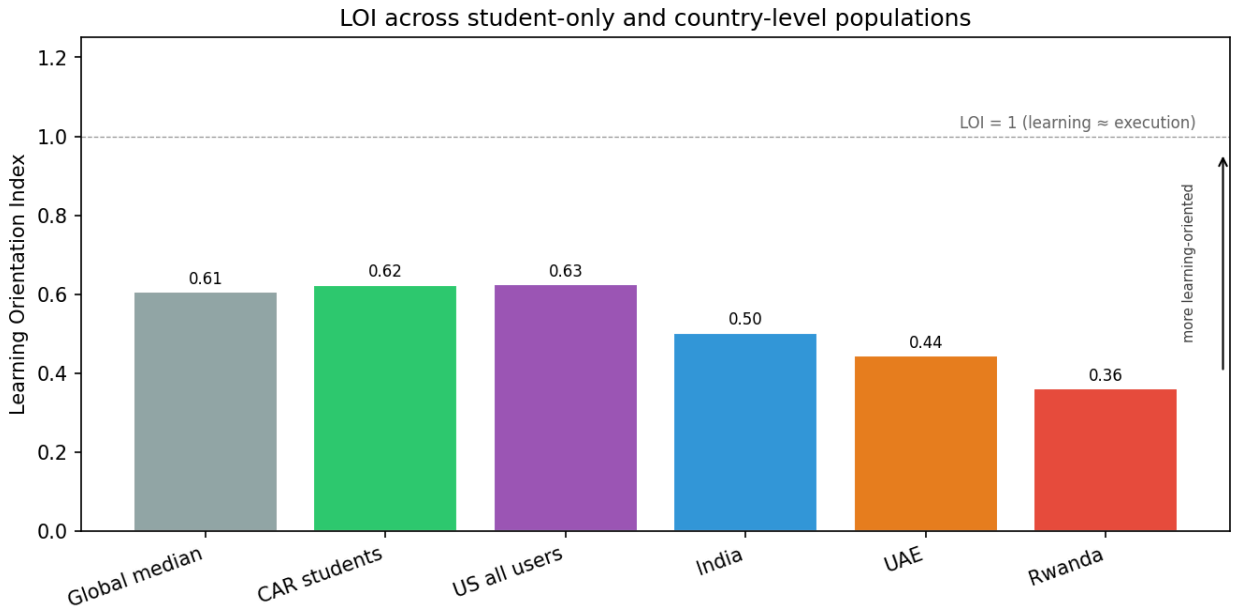

LOI across Education Report students, CAR students, US all-Claude-users, and the three Economic Index training-country populations. CAR students and US all-users cluster around the global median; the trained countries all fall below it.

US all-users are the most collaborative (61%), India, Rwanda, and UAE sit only a few points lower, and both student populations are lower still (51–53%). The LOI of US all-users and both student samples are slightly above the global median while the three formally-trained countries fall well below it. Reading the two charts side by side may cause confusion: how can these three countries be more “collaborative” than students, but have lower LOIs?

The answer recurs back to the Education Report's classification system. "Collaborative" groups learning-oriented prompts together with task-iteration prompts, meaning all-population samples may look more collaborative than students — even if broadly less learning-oriented — because they iterate more. This indicates that the collaborative turns of trained countries are disproportionately task iteration rather than learning, as evidenced by their low LOIs.

More generally, this data shows students are not inherently more learning-oriented or collaborative than general populations: both student samples are low on the collaborative axis, and the CAR student sample lands at essentially the same LOI as all US users. Population composition is therefore unlikely to be the primary driver of the difference in LOI between well-trained countries and countries with limited training.

3. Education-related requests are not inherently more learning oriented

Another alternative explanation is that subject matter drives LOI. Perhaps education-heavy countries ask more learning-style questions simply because tutoring is a learning-shaped task, and the countries with lower LOI happen to ask more task-heavy questions. To test this, I classified each country's request categories as education-related or software-related (which would intuitively be more task-based e.g. asking Claude to debug code) by keyword matching. I then fit a simple ordinary least squares (OLS) model predicting LOI from those two shares across all 119 countries with valid data.

One notable null finding, observing the residuals, is that subject matter is not a strong predictor of LOI. Education share carries a Pearson correlation with LOI of $r = -0.14$, software share $r = +0.20$. These estimated slopes are extremely small: a country would need to shift its education share by 15 percentage points to move its predicted LOI by 0.10, and the corresponding shift in software share would be 28 percentage points. One contributing factor to this potentially surprising result is that even within subjects, a wide range of collaboration styles are possible. "Education-related" could indicate a curious adult learning a new topic, but it also could mean a tutor preparing a lesson handout or a student working through homework. The limited residual relationship indicates that the subjects users are engaging in have limited effect on LOI compared to how users interact with the models.

A significant finding is that the CAR's actual LOI of 0.624 exceeds its predicted LOI (0.372) by +0.25, while all three formally-trained countries fall 0.16 to 0.21 below their predicted LOI. In other words, formal AI training is associated with lower reflective engagement even after accounting for what people actually use AI for.

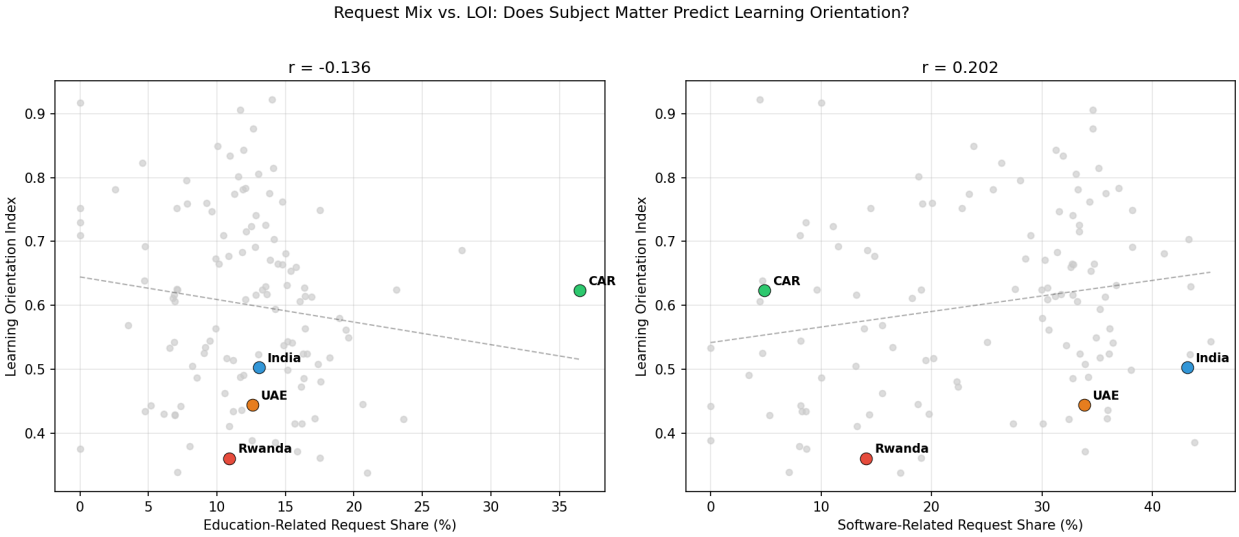

- *Caption: Actual vs predicted LOI under an OLS model using education and software shares as predictors. CAR exceeds prediction; all three trained countries fall below it.*

It is worth flagging the CAR is at the far tail for education-related requests, given the sample is 100% university students. However, fitting the OLS without the CAR outlier only shifts each trained country's residual by less than 0.02 (the strong trend remains: India is −0.17, UAE is −0.19, Rwanda is −0.20). This indicates subject matter is a near-null predictor of LOI ($R^2 \approx 0.10$) and the training-country gap persists even when holding subject matter constant.

It is also worth evaluating whether the Economic Index data supports this. Later Economic Index facets (November 2025 onward) let me look at use cases for the three Economic Index countries. Rwanda has the highest coursework share of the three Economic Index countries, at 35.2%, compared to India's 20.9% and UAE's 14.6%. If learning-oriented engagement followed mechanically from education-oriented use cases, Rwanda should thus have the highest LOI of the three. Instead, it has the lowest LOI. Rwanda's AI training appears to encourage directive use within educational workflows.

Task success is included to check whether LOI may be skewed as a result of users struggling (and thus taking multiple turns/asking the agent to explain or teach) rather than their rate of directive engagement. Given task success is broadly uniform across all three countries, this indicates the learning-orientation gap is more attributable to how users are taught to engage.

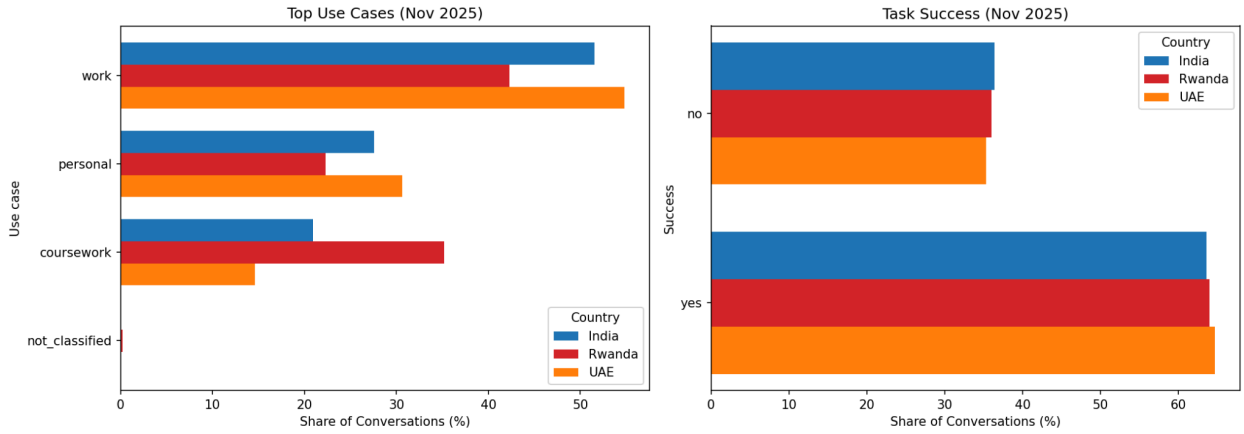

Coursework share and task success rate across the three Economic Index countries. Rwanda has the most coursework use but the lowest LOI; task success is nearly uniform.

4. Students in different contexts use AI for different subjects

Another interesting finding was that CAR students allocate 52.0% of their conversations to humanities and writing, 38.3% to business and career tasks, and just 9.8% to technical or quantitative work. The students in the Education Report are the near-inverse: 68.7% technical and quantitative, 18.5% business and career, 12.8% humanities and writing. In the CAR sample the top requests were writing, editing, translation, and resume and application support, with limited representation for coding. Part of this may reflect selection bias into what students use Claude, who is willing to take an Economic Index survey, the students who opted into the Diplomacy Lab partnership that produced the CAR sample, etc.

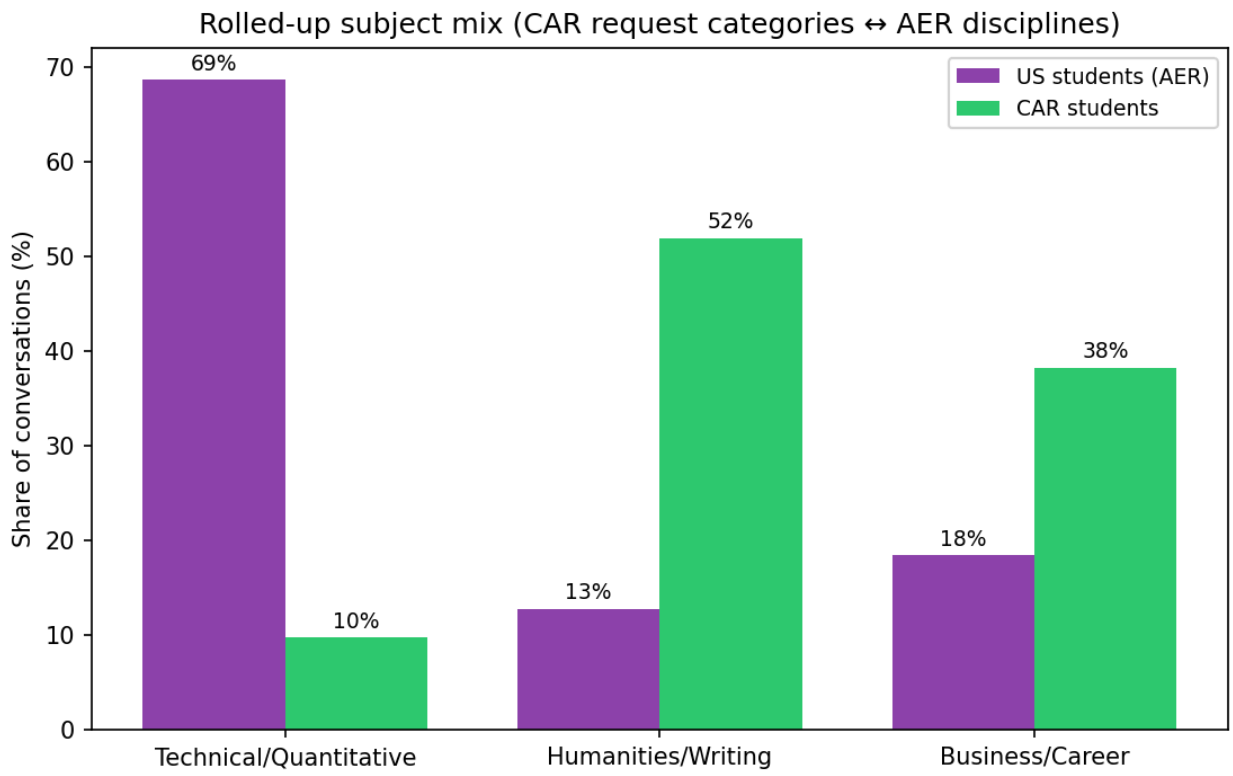

Subject mix across CAR students and Education Report students. The cleanest difference in the dataset: CAR tilts heavily toward humanities and writing; the Education Report tilts heavily toward technical and quantitative work.

Connecting to earlier results, this may indicate that humanities and writing tasks are inherently more learning-oriented; an equally plausible reading is that CAR students simply face different daily demands or have different norms around AI usage. Further research into how different majors use AI agents could be a natural next step based on this data.

This difference in subjects is also reflected when mapping CAR student requests to Bloom's taxonomy levels (see Methods). CAR students concentrate in Creating (68.7%) and Understanding (18.0%); Education Report students concentrate in Analyzing (30.2%), Creating (39.8%), and Applying (10.9%).

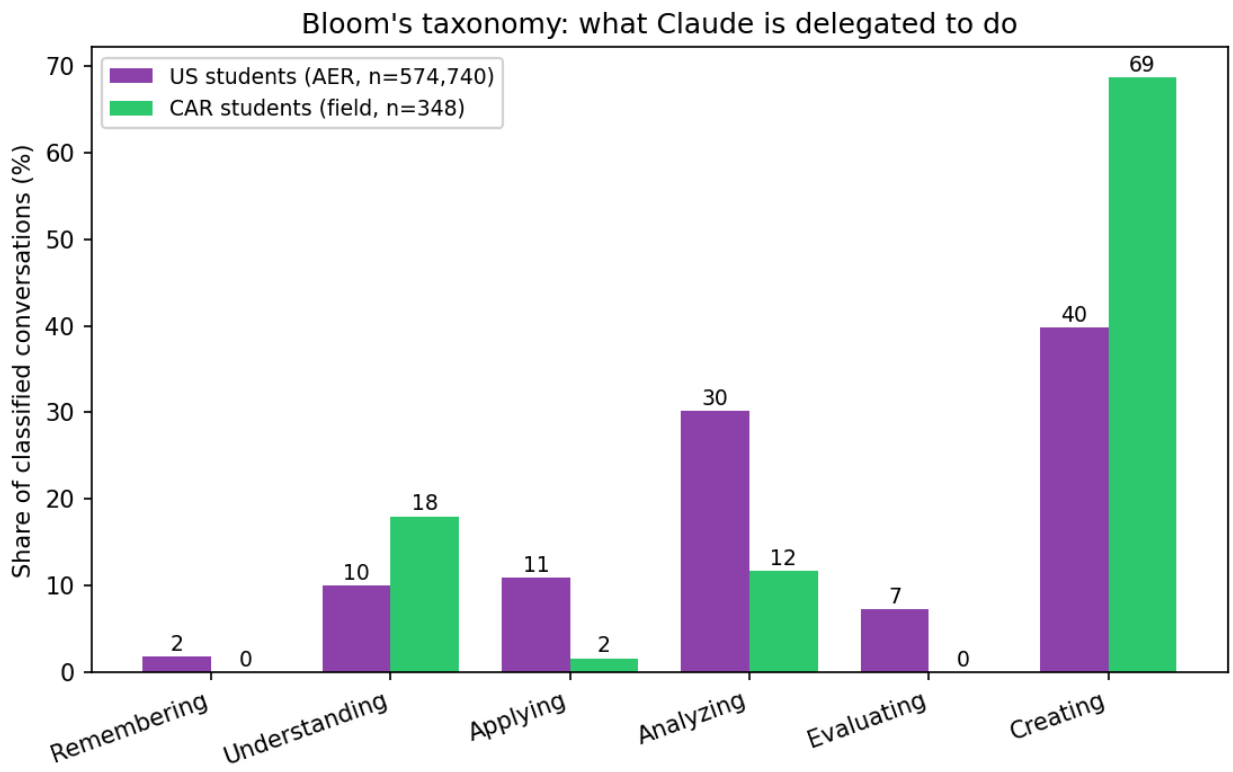

Distribution across Bloom's taxonomy levels for CAR students and Education Report students. CAR concentrates in Creating and Understanding; Education Report concentrates in Analyzing and Creating.

This is unsurprising, given code and quantitative work skew towards analyzing and applying; writing, editing, and translation lend themselves more to creating and understanding. It is worth cautioning that Anthropic constructs the Education Report's Bloom's distribution using its own classification pipeline, and so my CAR distribution may not be strictly commensurable.

Implications & Further Directions

Key findings:

1. **Formal AI training correlates with lower learning orientation**
   A higher degree of formal AI training corresponded with a lower relative Learning Orientation Index score, with Rwandan users forming a particularly low outlier. A lower LOI indicates users tend toward using AI as a task executor rather than a learning partner.

2. **Students are not inherently more learning oriented**
   CAR students cluster around the global median LOI rather than above it, meaning they sit essentially identical to US all-Claude-users (0.625). Population composition is therefore unlikely

to explain the low LOIs of well-trained countries.

3. **Education-related requests are not inherently more learning oriented**
   Subject matter is a near-null predictor of LOI ($R^2 \approx 0.10$), and the training-country gap persists after controlling for education- and software-related request shares. Rwanda has the highest coursework share of the three Economic Index countries, yet the lowest LOI, though this may be due to most learning-oriented education requests flowing through Chidi rather than claude.ai

4. **Students in different contexts use AI for different subjects**
   CAR students' request mix is heavy on humanities and writing (52%), while Education Report students are heavy on technical and quantitative work (69%). This may reflect differences in local academic demands and sample selection rather than any fundamental difference in how students engage AI.

Recommendations

It is worth noting LOI should not be an universal goal. Directive use is the correct pattern for certain use cases (e.g. code generation, document drafting, translation) and makes up much of the routine professional work AI can help automate. Some training programs should absolutely prioritise developing participants' ability to direct AI agents. The claim of this paper is that if programs explicitly attempt to produce learning-oriented users, the data does not yet show them doing so. AI training programs should intentionally select users they believe would benefit from being more learning oriented, and for these participants, programs must explicitly shift their focus from productivity toward accelerating education. Curricula should emphasise messaging and branding around teaching students to interrogate AI output rather than blindly operationalising it, and design assessments to reward collaboration and proper process over deliverable quality. The following changes could also help programs develop more learning-oriented use:

1. Directly track learning outcomes

Instead of relying on proxy metrics like enrollment or self-reported confidence, programs should measure their success by evaluating whether participants actually learned to collaborate effectively with AI. Where anonymised usage logs are available — such as those in the Anthropic Economic Index — revealed-behavior metrics like LOI and directive share are straightforward to compute and track. A simpler alternative, available to any program, is directly testing learning ability at regular checkpoints. For instance, programs could give participants open-ended problems with no single correct answer, then score the share of their AI interactions that ask for explanation, challenge model output, or iterate on reasoning, rather than simply requesting a finished deliverable (which ideally is not possible under an open-ended problem: for instance, asking a student to use AI to design a policy intervention for a specific local market, where the model lacks primary context and the student must provide the constraints and check the reasoning themselves).

2. Embed Socratic defaults in general-purpose tools

One interpretation of Rwanda's low LOI is that Chidi, a Socratic educational product, is working to pull students' learning-oriented use into the dedicated tool and leaving the residual in general Claude.ai looking directive. While I cannot verify this mechanism from the data alone, programs could consider building on Chidi's design to strengthen this possible effect. Governments could design, or partner with Anthropic to create, models with a "student mode" that uses Socratic defaults (e.g., asking clarifying questions before output), segmented by user type; or simply push towards other dedicated learning-oriented tools as the primary pedagogical instrument.

An additional perspective to consider is that the observed timeframe is admittedly narrow, and so it is unclear whether Socratic habits transfer when users switch to a more general productivity-default tool. Rwanda's low LOI may also be consistent with Chidi users not changing their long-term behaviour, indicating longer monitoring or further research may be necessary.

3. Take organic adoption seriously

The CAR has the highest LOI in this dataset, despite having the most constrained access and the smallest formal program. Informally speaking to students, I learned many felt the scarcity often encouraged them to be more deliberate in their use. Another common sentiment was that the student-led training initiatives produced better norms because the most active users set expectations for how best to use AI tools. One student said “Because none of us fully know what is going on, we have to help each other learn. And by the end of each session, we all understand better, even the teachers, who experiment with the tools live.”

An important caveat is that the CAR sample is small, self-selected, and drawn from one university, meaning this observation may not generalise to all circumstances. However, it is worth evaluating whether heavy-handed national rollouts actually outperform more organic user-to-user adoption, and also whether one of the interventions with highest marginal impact may simply be expanding access and letting usage norms form socially. The CAR example ought to draw more policy attention toward provisioning and peer learning as a supplement to existing formal training.

Limitations & Further Research

One primary limitation of this paper is the CAR sample. There is a limited sample size in 348 requests from roughly 50 self-selected University of Bangui students, collected over one week in January 2026. Given Anthropic's internal classification pipeline was not accessible, there may be additional classification errors affecting results. Also, as previously mentioned, the three AEI country samples are full Claude user bases while the CAR sample is entirely students. The Education Report comparison partially addresses this issue, but more fine-grained public student-level data would meaningfully strengthen our ability to make student-to-student comparisons. It is also worth noting that all four samples capture Claude users specifically, and since Claude is not the global default model for many individual users, the populations observed here may suffer from selection bias around users with stronger preferences about their AI tools. Extending analysis to usage patterns from other major models is a direction for future work.

Another limitation is that global training initiatives are still young: for instance, the UAE's mandatory K–12 curriculum began in the 2025–26 school year, and Rwanda's Anthropic MOU was signed in late 2025. None of the trained cohorts have fully aged into the Claude user base at the time of this paper, which means the LOI gap may widen, narrow, or reverse as trained students graduate into adult usage. Revisiting these countries annually with the same methodology, and paying attention to whether student-skewed country LOI rises over time, is the most direct way to test the effect of programs over time.

Lastly, this paper elected to track only four countries due to time and design constraints. Further evaluation of this hypothesis would ideally cover additional countries, spanning a fuller range of formal training intensity. Additional research could be conducted to compare users who have passed through a formal program against users in the same country who have not, moving this research from data analysis to original data collection.

The US, which appears throughout this paper as a reference point, is worth mentioning here as a strong candidate for future research. Where the US sits on the training-intensity axis is ambiguous: given that the US has no mandatory AI programming, on the formal state-led training axis this paper uses, the US belongs at the bottom. However, most US users also have extensive exposure to AI — through workplace Enterprise rollouts, peer networks, etc. — to the point where as of late 2025, 52% of US adults reported using a large language model, and 78% of the US labor force worked at a firm that had adopted AI. The US AI rollout is unlike the CAR in many respects. I consider the US as a useful reference because it shows long-term what a mixed, market-driven environment with heavy informal exposure looks like; it is less useful as a direct comparator to countries with extensive formal training. The fact that the US and CAR sit near each other at the top of the LOI distribution is suggestive of countries low on formal training gravitating to high learning-oriented use, but the two contexts are likely driven by different mechanisms. Further research into American AI adoption could help disentangle whether CAR's and the US's shared position near the top of the LOI distribution reflects the absence of formal training, or simply other unique factors to the US.

Conclusion

As global actors design training programs to ensure people are equipped to harness the transformative power of artificial intelligence, this paper offers empirical evidence that formal, state-led AI training does not, at the moment, shift student usage toward more learning-oriented patterns. This paper does not claim that AI training is a lost cause, or that learning-oriented use is categorically better than task-oriented use. Instead, it argues training learning-oriented use is an important element of developing AI literacy and future-proofing user economic value; and that based on original field data from a country outside of the Economic Index and public Economic Index usage data, increasing user familiarity with AI tools does not necessarily increase their ability to collaborate effectively or thoughtfully. AI training programs have the chance to prepare users for the future, but only if they see learning orientation as a goal rather than a guaranteed outcome.