

From tool to colleague: how AI partnership transforms the developers' identity across cultural boundaries

Olivier Le Van^{1,2,*†}, Valentina Curcio^{1,†}

¹Cegedim, Boulogne Billancourt, France

²Business Science Institute, Luxembourg, Luxembourg

Abstract

The increasing availability of Artificial Intelligence (AI) tools is transforming the software development landscape, enabling engineers to generate code with greater ease and efficiency. However, the integration of AI into programming workflows can vary significantly across individuals and cultures. This study investigates the impact of AI on a cohort of software developers from a multinational French company, including participants from France, Romania, and Morocco with diverse levels of experience with Large Language Models (LLMs). Since all participants work for the same company, we can isolate the impact of cultural differences on AI adoption, perception and usage. Using a qualitative approach, we interviewed participants and analyzed their responses using both Hofstede's cultural dimensions and Ashforth's professional identity frameworks. The results indicate that national cultural patterns are not a primary driver of AI adoption. Our results show that AI usage is associated with a redefinition of developers' professional identities, as they adapt to new technologies and work practices that challenge their existing roles and self-perceptions. Our research provides valuable insights for managers of multinational software development teams, offering practical advice on how to effectively integrate AI into their workflows and support the evolving professional identities of their team members.

Keywords

AI, developer, cultural dimensions, professional identity, qualitative method, organizational transformation, multinational teams, IT management

1. Introduction

The generation of natural language to code has been a global challenge in the past years [1, 2], and Large Language Models (LLMs) have emerged as assistant tools in software development [3]. Despite their growing adoption, significant variations persist in how developers integrate these technologies into their work processes: some view LLMs as indispensable productivity tools that enhance their workflow, while others are more skeptical and prefer to adhere to traditional programming approaches. The adoption of a new technology depends on the tool itself, as much as on the people [4]: the way individuals adapt their professional identities in response to new IT technologies critically influences whether those technologies effectively serve business goals or become sources of friction and misalignment.

This qualitative study investigates how tech teams in France, Romania, and Morocco navigate the challenges and opportunities of AI adoption, with a focus on the organizational and cultural factors that shape business-IT alignment during this process. By focusing on teams within the same multinational software company that have received identical access to state-of-the-art LLM tools, we create a unique opportunity to investigate how cultural, organizational, and individual factors influence technology adoption and use. Unlike previous research that has primarily examined LLM use in homogeneous contexts or through quantitative metrics alone [5], our study employs a qualitative approach that situates these technological changes within their sociocultural environments.

The global distribution of software development across cultural and geographic boundaries introduces additional complexity to understand LLM adoption. As AI assistance becomes part to development

BIR-WS 2025: BIR 2025 Workshops and Doctoral Consortium, 24th International Conference on Perspectives in Business Informatics Research (BIR 2025), September 17-19, 2025, Riga, Latvia

*Corresponding author.

† These authors contributed equally.

✉ olivier.levan@cegedim.com (O. Le Van); valentina.curcio@cegedim.com (V. Curcio)

ORCID 0009-0007-3081-3810 (O. Le Van); 0009-0002-5534-6570 (V. Curcio)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Table 1

Common benchmarks results on LLaMa 3 70B, Codestral 22B and GPT 4o Mini models. Source: [13].

Benchmark	LLaMa 3 70B	Codestral 22B	GPT 4o Mini
HumanEval (Base)	77.4	79.9	88.4
MBPP (Base)	82.3	72.5	85.4

processes, questions emerge about how diverse technical cultures adapt to these tools and how professional identities evolve through AI collaboration. Understanding these cultural dimensions is crucial for examining evolving patterns of trust, verification and code integration in AI-assisted environments. We examined East and West European, as well as North African development contexts, to contribute to a more nuanced understanding of how AI tools propagate within global software ecosystems.

Drawing on Hofstede’s cultural dimensions framework [6] and theories of professional identity [7], this work investigates how factors such as power distance, uncertainty avoidance, and individualism versus collectivism manifest in developers’ relationships with LLMs. Moreover, it explores how the integration of AI assistance into programming workflows transforms developers’ perceptions of their professional identity, their expertise and their value contribution across different cultural contexts.

We aim to answer two research questions. The first explores how cultural factors might explain different approaches to LLM adoption:

RQ1 How do Hofstede’s cultural dimensions (particularly power distance, uncertainty avoidance, and individualism) influence LLM adoption patterns and usage strategies across West European, East European and North African development teams?

The second examines the psychological and professional identity aspects of LLM adoption:

RQ2 How does the integration of LLMs into programming workflows transform developers’ perception of their professional identity? Does this transformation manifest differently across cultures?

Through in-depth interviews with developers, this research aims to uncover the complex interplay between technological capability, cultural context and professional self-concept that shapes the integration of AI into software development practices. The results will contribute to both theoretical understandings of technology adoption across cultural boundaries and practical insights for organizations seeking to effectively implement AI-assisted development tools in multinational contexts.

2. Related research

2.1. LLM for code generation

Benchmarks provide standard methodologies for assessing the capabilities of LLMs in code generation tasks, offering an evaluation framework for comparing different models’ performance on coding tasks.

Table 1 presents the benchmark results [8] of three recent models: LLaMa 3 70B [9] released in December 2023, Codestral 22B [10] released in January 2025 and GPT 4o mini [11] released in July 2024. These models demonstrate strong performance across standard coding benchmarks such as HumanEval and MBPP (Mostly Basic Python Problems). Earlier benchmarking efforts [8] relied on automated evaluations, but these approaches have been criticized for encouraging LLM creators to overfit their models [12] and lacking connection to real-world development scenarios [8].

Although benchmarks provide valuable insights for evaluating and comparing LLMs, they offer limited understanding of how these tools are adopted and integrated into developers’ daily workflows. The gap between benchmark performance and practical adoption highlights the need for qualitative research, to examine real-world usage patterns and their impacts on software development practices.

Recently, some studies have investigated the adoption of generative AI in software engineering [14]: at this stage of AI maturity, they highlight that the LLMs’ adoption highly depends on the compatibility

between AI and the existing workflows. Notably, LLMs can be used in various phases of the software development life cycle [15], serving not only as code generation tools but also as conceptual guides. However, the use of public LLM tools like ChatGPT is often limited by confidentiality constraints, as the context that can be prompted to the model is restricted [15].

Software development is a socio-technical activity. As such, social and cultural environment of the developers can influence their adoption of LLMs. Research has begun exploring how individual cultural values influence LLM adoption in software engineering contexts, using Hofstede's cultural dimensions framework to understand these relationships [5, 16]. This inquiry highlights that technology adoption in software development extends beyond technical performance to encompass social, cultural, and organizational factors that shape how developers perceive, adopt, and integrate generative AI into their practice.

2.2. Hofstede's cultural dimensions

Hofstede's seminal work on cultural dimensions, initiated in the 1980s [17, 6], has had a profound impact on our understanding of cultural norms and values. Although his framework has faced criticism [18], it remains a widely accepted reference [19]. This model posits that cultural norms and values can be categorized into six distinct dimensions:

Power Distance (PDI): This dimension captures the extent to which a society accepts and institutionalizes unequal power distribution. A high PDI score indicates a society with a strong hierarchical structure, whereas a low PDI score suggests a society that questions authority and favors a more egalitarian distribution of power.

Individualism versus Collectivism (IDV): This dimension reflects the degree to which a society is organized around individual or collective interests. Societies with high individualism tend to prioritize personal goals, whereas those with high collectivism emphasize group cohesion and interdependence.

Motivation towards Achievement and Success (MAS): This dimension assesses a society's orientation towards achievement, assertiveness and success (traditionally associated with masculinity) versus cooperation, modesty and care (traditionally associated with femininity).

Uncertainty Avoidance (UAI): This dimension measures a society's tolerance for ambiguity and uncertainty. Societies with high UAI scores tend to prefer structured and regulated environments, whereas those with low UAI scores are more open to novelty and diversity.

Long-term Orientation (LTO): This dimension evaluates a society's temporal orientation, with high LTO scores indicating a focus on future development and low LTO scores suggesting an emphasis on preserving traditions and past practices.

Indulgence versus Restraint (IND): This dimension captures the degree to which a society regulates and controls individual desires, with high IND scores indicating a permissive attitude towards gratification and low IND scores suggesting a more restrained approach.

Hofstede's cultural dimensions have been used in the field of software development to investigate the practices and approaches of software engineers from diverse cultural backgrounds [20, 21]. This research has significant implications for multinational teams working on collaborative projects. By understanding how people of different cultures act at work, both companies and individuals can develop better strategies to improve communication and collaboration [22]. A focus has been put on the individualism versus collectivism dimension : StackOverflow messages and profiles of developers from the US, China and Russia reveal clear differences related to the degree of individualism of their country.

Despite cultural differences, some studies suggest that software development has evolved into a global discipline. A shared professional culture emerges that may supersede national cultural influences in both classical software development practices [23] and AI-assisted software engineering [5, 16]. This raises questions about the interplay between national cultures and professional cultures in shaping the practices and attitudes of software engineers.

2.3. Professional identity framework

Professional identity represents how individuals define themselves within their occupational context and encompasses the beliefs, values, motivations and experiences that shape their professional self-concept [24, 25]. In the context of software development, professional identity becomes particularly relevant as developers navigate the integration of AI tools that may challenge traditional notions of expertise and value creation.

2.3.1. Theoretical foundation

Professional identity could be conceptualized through three interconnected dimensions [7]:

Core of Identity ("Who I Am"): This dimension encompasses the fundamental self-concept and central characteristics that define an individual's professional essence. It includes stable attributes such as values, personality traits and deep-seated beliefs about one's role and purpose within the profession [25]. For software developers, this might include viewing oneself as a problem-solver, creator or technical expert.

Content Identity ("What I Care About/Can Do"): This dimension reflects the specific knowledge, skills and professional interests that individuals possess and value. It encompasses both technical capabilities and domain expertise, as well as professional goals, aspirations, and areas of specialization [26]. This dimension is more malleable than Core of Identity and can evolve as new technologies and practices emerge.

Behavioral Identity ("What I Do"): This dimension manifests itself in the actual practices, routines, and behaviors that individuals engage in within their professional context. It includes daily work activities, interaction patterns with colleagues, and the specific methods and tools used to accomplish tasks [7]. This is typically the most visible and immediately changeable aspect of professional identity.

Research on the adaptation of professional identity suggests that technological disruptions can trigger identity work processes in which individuals actively construct, maintain, or modify their professional self-concept [24]. The introduction of AI tools in software development represents such a disruption, potentially affecting all three dimensions of professional identity. The literature on workplace automation and technological change indicates that professionals experience identity threats when new technologies challenge their expertise or change the nature of their work [27, 28]. However, technology can also enhance professional identity by expanding capabilities, increasing autonomy, or enabling focus on higher-value activities. Software development, as a knowledge-intensive profession, is particularly susceptible to identity shifts when core tools and practices evolve. Previous research [29] has examined how developers' professional identities are shaped by factors such as technical expertise, problem-solving habits, continuous learning and professional autonomy. The integration of LLMs into programming workflows introduces new dynamics that could challenge or transform these traditional aspects of developer identity. Cultural dimensions may influence how individuals interpret technological changes, the degree of resistance or acceptance they exhibit and the strategies they employ to maintain or reconstruct their professional identity [5, 16]. This cultural influence becomes particularly relevant in multinational organizations where developers from different cultural backgrounds could maybe respond differently to the same technological changes.

Existing research has examined LLM adoption in software development from technical and productivity perspectives, but limited attention has been paid to the professional identity implications across different cultural contexts. We address this gap by investigating how the integration of LLMs affects programmers' professional identity and whether the effects are different across cultural boundaries, providing valuable insights for multinational organizations adopting AI-assisted development tools.

Table 2

Token consumption and its metrics for the 71 users between 8 April and 8 May 2025.

	Mean	Median	Standard Deviation	Maximum	Minimum	Sum
Number of Tokens	238,557.82	28,267.00	548,710.24	3,613,935.00	54.00	16,937,605.00

3. Methodology

3.1. Corporate setting

This research was conducted within a major French IT company specializing in digital solutions for the B2B sector. The organization employs a diverse range of professionals, including developers, project managers, and business analysts mainly in France, Morocco and Romania.

The organization under study has adopted a policy discouraging the use of external LLMs, relying solely on employee self-regulation, while instead promoting the use of self-hosted LLMs within its infrastructure. As part of this initiative, the company released two models on April 8th, 2025: *LLaMa 3.3 70B* [9] and *Codestral 22B* [10]. These models are currently intended for research and exploration purposes, rather than production use.

The inference of these models is performed on *Nvidia DGX H100* GPUs, leveraging the *vLLM* framework [30]. To enable seamless integration with existing development workflows, the models' endpoints can be accessed through a chat window with chat history functionality or directly within the developer's preferred Integrated Development Environment (IDE). This flexible deployment strategy allows developers to interact with the LLMs in a way that suits their individual preferences and work styles.

3.2. Population metrics

Within the company, a total of 302 developers were identified as potential users of the self-hosted LLM solution. The geographic distribution is the following: 180 software engineers are based in France, 98 in Morocco, and 24 in Romania. However, after a one-month period, only 71 developers (23.5% of the total) had logged in to the system. An analysis of token consumption was conducted for all developers who accessed the self-hosted LLM, detailed results are presented in Table 2. Based on this analysis, the developers' population was categorized into three distinct groups:

Frequent users: frequent connections and token usage superior to the average. We retained for the interview the top 2 users of each country.

Curious users: less than 3 connections and negligible token consumption. We retained for interview the bottom 2 users for each country.

Non-users: developers who did not establish any connection to the self-hosted LLM. Two random candidates from each country were selected for interview, where possible.

Notably, in Romania, only one non-user candidate was identified. The resulting cohort consists of 17 participants, whose profiles are summarized in Table 3, including their role, country of origin, technical expertise, years of work experience, in-house LLM usage patterns and interview duration.

3.3. Qualitative semi-conducted interview

The meetings were held in French or English, depending on the preferred language of the interviewee. They were conducted online using *Zoom*¹ and recorded with the explicit consent of the participants. The audio layer was then extracted from the record and processed with *Whisper large* [31] to obtain the transcriptions. *Pyannote Speaker diarization 3.1* [32, 33] was used to diarize them and identify the speakers. The transcriptions were then manually processed to correct mistakes made by the model.

We used an LLM method [34, 35] to extract the relevant citations from each transcription by using *Claude Sonnet 3* [36], then reviewed them one by one. Subsequently, we conducted a manual encoding

¹<https://www.zoom.com/>

Table 3
Participants details.

ID	Role	Country	Technical Stack	Work Experience	LLM Usage	Interview duration
P1	Lead Developer	France	PHP, Javascript	10 - 15 years	Frequent user	39.39 minutes
P2	Back-End Developer	France	Java	15 - 20 years	Frequent user	29.15 minutes
P3	Front-End Developer	France	Javascript	5- 10 years	Curious user	36.38 minutes
P4	Back-End Developer	France	Java, Python	0- 5 years	Curious user	19.01 minutes
P5	Full-Stack Developer	France	Javascript	10- 15 years	Non-user	28.34 minutes
P6	Back-End Developer	France	PHP, Java	15- 20 years	Non-user	38.55 minutes
P7	Front-End Developer	Romania	Javascript	0- 5 years	Frequent user	21.38 minutes
P8	Front-End Developer	Romania	Javascript	10- 15 years	Frequent user	33.19 minutes
P9	Front-End Developer	Romania	Javascript	0- 5 years	Curious user	15.59 minutes
P10	Full-Stack Developer	Romania	PHP, Javascript	0- 5 years	Curious user	41.24 minutes
P11	Full-Stack Developer	Romania	PHP, Javascript	5-10 years	Non-user	40.04 minutes
P12	Back-End Developer	Morocco	Python	5-10 years	Frequent user	30.44 minutes
P13	Full-Stack Developer	Morocco	Java, .Net, Python	0- 5 years	Frequent user	59.24 minutes
P14	Front-end Developer	Morocco	Javascript	0- 5 years	Curious user	20.12 minutes
P15	DevOps	Morocco	Python, Shell	0- 5 years	Curious user	16.01 minutes
P16	Lead Developer	Morocco	Java, Javascript	5-10 years	Non-user	17.40 minutes
P17	QA Engineer	Morocco	Python, Javascript	0- 5 years	Non-user	43.46 minutes

process, where every segment of speech was carefully categorized within the framework of the six cultural dimensions, as well as the professional identities. This approach enabled us to systematically analyze the transcriptions and identify patterns, themes, and relationships between the cultural dimensions and professional identities.

3.4. Hofstede’s cultural dimensions analysis

To obtain the participants’ point of view about the six cultural dimensions, the meeting was driven by these questions:

- **PDI:** Who was at the origin of your LLM usage?
- **IDV:** Do you share tips or technical details about your LLM usage or not?
- **MAS:** Do you feel competitive or collaborative with the LLM?
- **UAI:** How do you perceive risks by working with an LLM?
- **LTO:** How do you perceive the developer job in three to five years?
- **IVR:** What do you think about AI mistakes or about AI-generated code in general?

Furthermore, we engaged in in-depth discussions with the participants to explore their perceptions of LLMs and their impact on software development practices. Specifically, we investigated whether the use of LLMs alters their approach to code development, influences their vision of the job, and affects the value they give to certain skills over others.

Through the analysis of the interview transcripts, we assigned a score to each participant for each of the six cultural dimensions in Hofstede’s framework, ranging from low to high. The original metrics, periodically updated by The Culture Factor Group [37], are presented on a scale from 0 to 100. We categorized the scores into five intervals: Low (0-20), Low-Medium (20-40), Medium (40-60), Medium-High (60-80), and High (80-100). This classification allows for the comparison between the cultural dimensions of individual participants and the national averages (Fig. 1a).

3.5. Professional identity analysis

To gather the participants’ perspectives on the three dimensions of Professional Identity, we guided the discussions with the following questions:

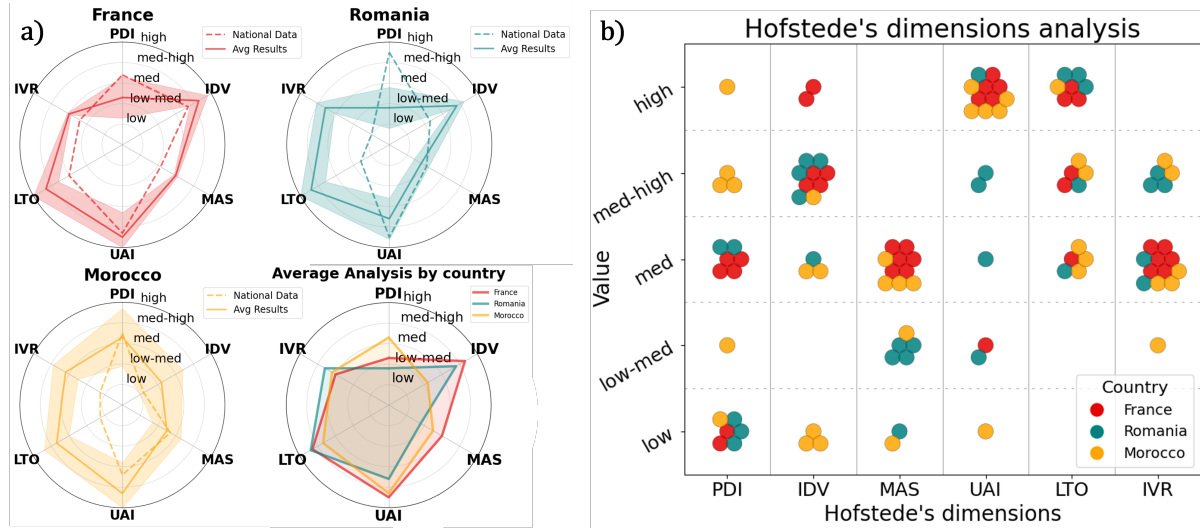


Figure 1: a) The average results of our Hofstede's cultural dimensions analysis (solid lines) are plotted along with the standard deviation (shaded region) for each of the studied countries. The official national scores are also given (dashed lines), source: [37]. The last graph (lower-right) represents a superposition of our average results per country. **b)** Plot of the results of Hofstede's cultural dimensions analysis. Each circle represents a participant. The results for all participants are given for all the Hofstede's dimensions (horizontal axis), with the measured value displayed on the vertical axis.

- **Core of Identity:** Could you present yourself? How do you think software development is changing with AI?
- **Content of Identity:** Do you think that some of your skills are increased while other decreased? How do you perceive your expertise and value since you are using LLMs?
- **Behavior of Identity:** Could you share your last usage with us? How do you use LLMs in your daily work?

By examining their perspectives on these dimensions, we sought to develop a nuanced understanding of how software developers perceive themselves, their work, and the impact of LLMs on their professional identities. We attributed a level of change on each Professional Identity dimension ranging from No Change to Moderate or Important Change.

4. Results

4.1. First observations on LLM adoption by developers

Our initial findings reveal that, despite the official company policies, all the interviewed developers had already used LLMs for their daily tasks at least once a week. The most commonly cited tool was ChatGPT by OpenAI, followed by, in order of frequency: CursorAI, Claude by Anthropic and Gemini by Google. These tools were used for personal and professional purposes at the time of the interviews. Considering these factors, the LLM usage score (Tab. 3), which informed our selection of the participants' cohort, only reflects in-house LLM use and will not be used as a predictive factor for actual LLMs usage.

4.2. Cultural dimensions results

The full results of the Hofstede's cultural dimensions analysis are presented in Table 4 and Figure 1b.

4.2.1. Country-based cultural analysis

The average tendency of each cultural dimension per country is resumed in Tab. 5. For a clear visual comparison, we plotted the average values of our study versus the official national scores from The

Table 4

Hofstede Cultural Dimensions Analysis. Results per participant.

ID	PDI	IDV	MAS	UAI	LTO	IVR
P1	Medium	Medium-High	Medium	High	Medium-High	Medium
P2	Medium	High	Medium	High	Medium-High	Medium
P3	Medium	Medium-High	Medium	High	High	Medium
P4	Medium	High	Medium	High	High	Medium
P5	Medium	High	Medium	High	High	Medium
P6	Low	Medium-High	Medium	Low-Medium	Medium-High	Medium
P7	Medium	Medium-High	Low-Medium	High	High	Low-Medium
P8	Low	Medium-High	Low-Medium	Medium	High	Medium-High
P9	Medium	Medium-High	Low-Medium	Low-Medium	Medium-High	Medium-High
P10	Low	Medium-High	Low	Medium-High	High	Medium
P11	Low	Medium	Low-Medium	Medium-High	Medium	Medium
P12	Medium-High	Low	Low	High	Medium	Medium
P13	Medium-High	Medium	Medium	High	Medium	Medium-High
P14	Low	Medium-High	Medium	Low	Medium-High	Medium-High
P15	High	Low	Medium	High	Medium	Low-Medium
P16	Low-Medium	Medium	Medium	High	High	Medium
P17	Medium-High	Low	Low-Medium	High	Medium-High	Low-Medium

Table 5

Average results of the Hofstede's cultural dimensions analysis per country.

	France (P1 to P6)	Romania (P7 to P11)	Morocco (P12 to P17)
PDI	Low to Med	Low to Med	Variable (Low to High)
IDV	Med to High	Med to Med-High	Low to Med
MAS	Medium	Low to Low-Med	Low to Med
UAI	High (except P6)	Variable (Low-Med to High)	High (except P14)
LTO	Med to High	Med-High to High	Med to High
IVR	Medium	Medium	Medium

Culture Factor Group [37] in Fig. 1a.

French developers (P1 to P6) generally exhibit profiles that are broadly consistent with Hofstede's national scores for France (Fig. 1a). They show high individualism: **P1** *"It goes faster on its own"*; moderate motivation towards achievement and success: **P6** *"I think we should rather see it as a productivity tool [...] I'm not sure that teams will shrink, but I think that with the same number of people, we will get more things done"*; high uncertainty avoidance: **P5** *"I'm very much into... techno-skepticism"*; strong long-term orientation: **P4** *"I imagine that we will perhaps be asked in the future to know how to use this kind of technology"*; moderate indulgence: **P3** *"the standard is higher, but we will have to go back later"*. The slightly lower power distance among participants may be explained by their higher educational background (all hold at least a master's degree) and elevated professional status.

Romanian Developers (P7 to P11) appear to have a lower power distance: **P9** *"The decision to use LLMs was very personal, also with discussion with colleagues"*, higher individualism: **P8** *"I'm rather independent of choosing what are our tools"* and higher long-term orientation: **P7** *"In the future, I want to use AI more responsibly"* compared to Hofstede's profile for Romania (Fig.1a). These differences could reflect both the influence of the tech profession and the specific organizational culture within the studied company, which may emphasize autonomy, low hierarchical distance and long-term innovation.

The Moroccan group (P12 to P17) aligns with national scores for Morocco, displaying in average a medium-high power distance: **P14** *"First you, operational, and then you will pass on the information to your manager"* and uncertainty avoidance: **P15** *"you should not share with ChatGPT, everything there is [...] you should not share all the data"*. However, participants also demonstrate stronger long-term

strategic thinking: **P16** *"I think AI will definitely become even more powerful over time"*, and higher indulgence: **P12** *"I also used them for decorating [...] my living room. [...] It gave me some very good ideas"*. Moroccan developers appear to adapt to modern tech-oriented environments and bridge traditional values with the demands of globalized, innovation-driven workplaces.

4.2.2. LLM adoption and cultural influence

As we observed in Sec. 4.1, despite the cultural differences, all participants reported regular use of LLMs, it being the in-house LLM and, above all, publicly available ones. This suggests that Hofstede's cultural dimensions, while useful for understanding values and behaviors, do not have a significant impact on LLM adoption in this context. It can be seen in Fig. 1b that for each cultural dimension the participants tend to cluster around a dominant value, without a strong national pattern. The average results of our analysis, as shown in Fig. 1a, tend to overlap over the three studied countries. The universal adoption of LLMs indicates that professional drivers such as productivity, peers and perceived usefulness may outweigh cultural resistance or predispositions in the adoption of generative AI tools.

4.3. Professional identity results

Table 6 shows the results of the changes in the three dimensions of professional identity considered: Core of Identity, Content of Identity and Behavior of Identity.

Table 6
Professional Identity Analysis results for all the participants.

ID	Core of Identity	Content of Identity	Behavior of Identity	ID	Core of Identity	Content of Identity	Behavior of Identity
P1	No Change	Moderate	Important	P9	Moderate	Important	Moderate
P2	No Change	Moderate	Moderate	P10	Moderate	Important	Moderate
P3	Moderate	Important	Important	P11	No Change	Moderate	Moderate
P4	Moderate	Important	Important	P12	Moderate	Important	Moderate
P5	Moderate	Important	Moderate	P13	Moderate	Important	Moderate
P6	No Change	Moderate	Moderate	P14	Moderate	Important	Important
P7	Moderate	Important	Moderate	P15	Important	Moderate	Important
P8	Moderate	Important	Important	P16	Moderate	Important	Important
...	P17	Moderate	Important	Moderate

4.3.1. Country-based professional identity analysis

French developers (P1 to P6)

- **Core of Identity:** Predominantly stable with three participants reporting **No Change** (P1, P2, P6) and three reporting Moderate changes
- **Content of Identity:** Moderate to Important engagement across all participants
- **Behavior of Identity:** Primarily Moderate with two participants showing Important changes

French developers demonstrate remarkable stability in the core of their professional identity, with a strong sense of self deeply rooted in their engineering expertise. This stability reflects confidence in their fundamental professional value, as exemplified by **P1**: *"My job will always be the same. I'm an engineer"* and **P2**: *"I think that for now, we still have a few good years ahead of us as developers"*.

Despite this core stability, French developers actively engage with content identity changes, developing new skills and adapting their professional expertise. **P3** illustrates this adaptive learning: *"I asked him to list all the linear paths [...] And, let's say that I... I took that as a basis"*, while **P4** anticipates future

skill requirements: *"I imagine that we will perhaps be asked in the future to know how to use this kind of technology because a developer with and without it may not have the same development speed"*.

French developers show pragmatic adaptation to LLM tools without compromising their professional values. **P4** demonstrates practical integration: *"everyday example [...] I need to manipulate an Excel file in Java I will ask for a list of libraries [...] I will ask ChatGPT or LLaMa perhaps to have the main functions"*, while **P1** shows strategic usage patterns: *"I use it maybe every other day [...] when I use it, I use it a lot"*. This measured approach reflects their confidence in maintaining professional autonomy while leveraging AI capabilities.

Romanian developers (P7 to P11)

- **Core of Identity:** Moderate changes across most participants, with only P11 reporting No Change
- **Content of Identity:** Predominantly Important engagement (except Moderate for P11)
- **Behavior of Identity:** Primarily Moderate with one participant showing Important changes

Romanian developers reveal a profession that is actively grappling with technological disruption while preserving professional meaning. Their moderate core identity changes suggest adaptation rather than replacement, with developers redefining their professional identity while maintaining its essence. **P9** captures this balance: *"I am confident that AI won't stall my job, but it will give me the opportunity to transform it into something. Something better, as I said, and more creative"*, while **P10** expresses cautious optimism: *"I'd like to keep the values I have now, well, not completely intact. They will suffer changes, but I hope not so much"*.

The substantial content identity changes reflect a significant skills' evolution: the developers recognize the need to master both traditional and AI-assisted programming. **P8** identifies emerging skills requirements: *"I think we will have to get good at prompts, of writing prompts for it"*, while **P9** emphasizes temporal advantages: *"I think that time is extremely important and the AI is going to bring time for us as humans to be more creative"*.

Behaviorally, Romanian developers demonstrate security-conscious adoption patterns, prioritizing professional standards and data protection. **P7** exemplifies this cautious approach: *"I was sometimes using ChatGPT [...] with some level of abstraction to protect our data"*, while **P9** shows strategic workflow integration: *"when I have very specific questions, instead of surfacing on the internet, I choose to ask ChatGPT because it already makes quicker research"*. This measured behavioral evolution reflects their commitment to maintaining professional integrity while embracing technological advancement.

Moroccan developers (P12 to P17)

- **Core of Identity:** Predominantly Moderate changes with one participant (P15) reporting Important transformation
- **Content of Identity:** Important changes across most participants (except P15 showing Moderate)
- **Behavior of Identity:** Mixed pattern from Moderate to Important changes

Moroccan developers uniquely exhibit universal core identity evolution, being the only cohort where all participants experience some fundamental change in professional self-conception. This transformation is characterized by existential questioning about their professional future while maintaining the confidence that humans are irreplaceable. **P12** articulates this tension: *"if from here we ask more questions to chatGPT [...] we no longer need to ask human questions to physical persons"* while affirming *"We will always need a human, physical touch"*.

The substantial content identity changes reveal a comprehensive skill portfolio transformation, encompassing both new technical domains and strategic revaluation of human-centric capabilities. **P12** demonstrates radical professional transition: *"I am doing an LLM course [...] because we are going to switch to the LLM level on my project"*, while **P13** identifies entirely new areas of expertise: *"The skills I've refined are really on the side of how to use AI algorithms. It was really a skill that I hadn't really refined before"*.

Behavioral transformations among Moroccan developers range from a cautious approach to comprehensive workflow restructuring. **P14** illustrates fundamental methodological shifts: *"Before we used Google, but now I'm getting used to chatGPT"*, while **P15** exhibits systematic AI-first integration: *"When I go to work on a task [...] I do research on ChatGPT"*. This behavioral evolution extends to relational redefinition, with **P13** expressing: *"He's a colleague. He's not an intern"*, indicating a shift from tool-based to partnership-based professional relationships.

4.3.2. Experience-based professional identity analysis

We investigated the relationship between the level of change in professional identity and the developer's experience. The results of this analysis are presented in Figure 2.




Figure 2: The level of Identity Change (vertical axis) is given as a function of Experience (horizontal axis) for the three professional identity dimensions. Each circle represents a participant, the color their country of origin.

We observe a correlation between changes in the participants' Core of Identity and their years of working experience in the field: the most experienced software engineers exhibit a lower level of change in their Core of Identity, whereas developers with less working experience in the field demonstrate a stronger change in this dimension. This suggests that the core aspects of professional identity, such as self-perception and overall professional orientation, are more stable among seasoned developers, while being more malleable among those with less experience.

In contrast, the Content Identity and Behavioral Identity dimensions do not appear to exhibit a clear dependency on the experience of the developers. The introduction of AI seems to have had a moderate to strong impact on both the "I do" (Behavioral Identity) and the "I care about/want/believe/generally do/can do" (Content of Identity) aspects of professional identity. This suggests that AI has become an integral part of the development process, influencing the way developers perceive their work, their skills, and their professional roles, across all levels of experience.

4.3.3. Cultural dimension and professional identity cross analysis

We investigated the relationship between cultural dimensions and professional identity changes to determine if there are any correlations between the level of professional identity changes and the positioning of participants within the different cultural dimensions.

Our analysis reveals that there are no strong correlations between Hofstede's cultural dimensions and the Content of Identity and Behavior of Identity. This may be attributed to the fact that AI has already become an integral part of the practice and adoption of all participants, including non-users who use other LLMs outside of the in-house model. This widespread adoption may have contributed to a convergence of attitudes and behaviors across participants, regardless of their cultural backgrounds.

If we focus at the changes in the Core of Identity with respect to Hofstede's cultural dimensions, several tendencies emerge (see Fig. 3):

- **Power Distance Index (PDI):** The distribution of participants' data suggests a positive correlation between high power distance and stronger changes in Core of Identity. Participants with low to




Figure 3: Core of Identity scores (vertical axis) are plotted versus Cultural Dimensions scores (horizontal axis), for each of the six Hofstede's cultural dimensions (PDI, IDV, MAS, UAI, LTO, IVR). Each circle represents a participant, with its color indicating the country of origin.

medium PDI scores tend to exhibit low to moderate Core of Identity changes. Notably, higher PDI scores are associated with Moroccan participants, indicating that they experience a higher level of Core of Identity change compared to other participants.

- **Individualism versus Collectivism (IDV):** Participants with low IDV scores (i.e., collectivist attitudes) tend to exhibit moderate to high levels of Core of Identity changes, whereas those with high IDV scores (i.e., individualistic attitudes) display low to medium Core of Identity changes.
- **Motivation towards Achievement and Success (MAS):** Participants tend to cluster around a medium Hofstede score, corresponding to moderate Core of Identity changes.
- **Uncertainty Avoidance Index (UAI):** Participants tend to cluster around high Hofstede scores, corresponding to moderate Core of Identity changes.
- **Long-term Orientation (LTO):** Participants tend to cluster around high Hofstede scores, corresponding to moderate Core of Identity changes.
- **Indulgence versus Restraint (IVR):** Participants tend to cluster around medium to medium-high Hofstede scores, corresponding to moderate Core of Identity changes.

These results suggest that PDI and IDV dimensions could have a significant impact on the level of Core of Identity change perceived by the participants in relation to their work. Specifically, a stronger Power Distance and collectivist behavior appear to be associated with a greater change in participants' professional identity. These two cultural traits are also characteristic of Moroccan developers. In contrast, participants tend to exhibit more homogeneous behavior with respect to the other cultural dimensions, clustering around a few strong values for both cultural and Core of Identity scores. This suggests that, among our participants, there is a certain degree of homogeneity in how AI impacts their behavior and attitudes towards LTO, IVR, MAS, and UAI, as well as a similar homogeneity in how this impact changes their vision of their professional identity.

5. Conclusion

This study examined how cultural dimensions and professional identity factors influence the adoption and integration of Large Language Models (LLMs) among a limited range of software developers from France, Romania, and Morocco working in the same multinational IT company. By combining Hofstede's cultural framework with Ashforth's professional identity theory, we analyzed the sociotechnical dynamics underpinning AI adoption in software engineering.

5.1. Addressing research questions

RQ1: How do Hofstede’s cultural dimensions influence LLM adoption patterns and usage strategies?

Our findings suggest that while national cultures shape attitudes, they do not significantly constrain LLM adoption. Participants across all cultural backgrounds—regardless of their national scores for power distance, uncertainty avoidance, or individualism—reported regular use of AI tools, either the in-house LLM or publicly available ones. However, cultural traits such as high uncertainty avoidance or low individualism (e.g., among Moroccan developers) were associated with greater shifts in professional identity, especially in how participants perceive trust, expertise, and collaboration. This supports the notion that professional norms may partially override national cultural patterns in highly globalized and technical environments [5, 23].

RQ2: How does the integration of LLMs transform developers’ professional identity?

The introduction of LLMs catalyzed significant shifts in *content* and *behavioral* aspects of identity across all participants. Less experienced developers and those in high power-distance or collectivist cultures exhibited deeper changes in their *core* identity, reflecting a more profound renegotiation of their role as software engineers. While French developers largely preserved their self-concept as "engineers" who use AI as a tool, Moroccan developers were more likely to describe LLMs as transformative collaborators, echoing findings from recent work on identity adaptation in automated environments [7, 27].

5.2. Practical guidelines for multinational teams

To support effective and ethical LLM integration, we propose the following evidence-based recommendations for organizations managing culturally diverse development teams:

- **Support identity transformation, not just tool onboarding.** AI adoption impacts how developers see themselves. Managers should create forums for open discussion and reassure staff about the enduring value of human creativity and technical judgment [34].
- **Embed LLMs directly into daily workflows.** Developers are more likely to engage with AI when it is integrated in their IDEs and version control pipelines, rather than via separate platforms. Compatibility with existing tools is key [14].
- **Support LLM adoption through universal strategies.** While it may seem intuitive to adapt support based on cultural orientations (e.g., providing clear guidelines in high uncertainty-avoidance cultures or encouraging collaboration in collectivist contexts), recent evidence indicates that habit and performance expectancy are the primary drivers of LLM adoption, with cultural values playing a less significant role [5].
- **Train for prompt literacy and critical thinking.** Developers must learn not only how to ask good questions, but also when not to rely on AI. Prompt engineering is emerging as a core skill in software work [15].
- **Ensure secure, ethical use of AI.** Developers—especially in Romania—expressed concerns about confidentiality. Organizations should prefer in-house models and clarify what kind of data can be shared with LLMs [15].
- **Enforce a culture of experimentation and autonomy.** Encourage team members to explore AI tools without fear of judgment or failure. Drawing on lessons from qualitative research practice, supportive and reflexive environments help individuals—especially those adapting to new roles—navigate identity shifts and learn from experimentation [35].
- **Establish strategic governance for AI-business alignment.** Organizations should ensure that LLM integration supports business objectives through clear governance frameworks that balance innovation with business requirements and risk management. This includes creating policies that limit shadow usage of LLMs while measuring impact beyond productivity metrics to demonstrate tangible business value between IT capabilities and organizational goals.

In summary, successful AI integration requires not only technical deployment but cultural and psychological adaptation. By recognizing that identity and trust are as crucial as productivity gains, organizations can foster responsible, inclusive, and sustainable AI adoption in global software teams. **In future work**, we aim to examine the limitations, organizational adoption, and strategic benefits of the agentic systems released this year, with a focus on their long-term implications for business–IT alignment.

Declaration on Generative AI

During the preparation of this work, the authors used Whisper [31] and Claude [36] to transcribe the interviews. The authors reviewed and edited the publication’s content and take full responsibility for it.

References

- [1] W. Ling, P. Blunsom, E. Grefenstette, K. M. Hermann, T. Kočiský, F. Wang, A. Senior, Latent predictor networks for code generation, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016, pp. 599–609.
- [2] P. Yin, G. Neubig, A syntactic neural model for general-purpose code generation, in: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2017, pp. 440–450.
- [3] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. D. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al., Evaluating large language models trained on code, *arXiv preprint arXiv:2107.03374* (2021).
- [4] V. Uren, J. S. Edwards, Technology readiness and the organizational journey towards AI adoption: An empirical study, *International Journal of Information Management* 68 (2023) 102588.
- [5] S. Lambiase, G. Catolino, F. Palomba, F. Ferrucci, D. Russo, Investigating the role of cultural values in adopting large language models for software engineering, *ACM Transactions on Software Engineering and Methodology* (2024).
- [6] G. Hofstede, Dimensionalizing cultures: The Hofstede model in context, *Online readings in psychology and culture* 2 (2011) 8.
- [7] B. E. Ashforth, S. H. Harrison, K. G. Corley, Identification in organizations: An examination of four fundamental questions, *Journal of management* 34 (2008) 325–374.
- [8] J. Liu, C. S. Xia, Y. Wang, L. Zhang, Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation, *Advances in Neural Information Processing Systems* 36 (2023) 21558–21572.
- [9] Meta AI, Introducing Meta Llama 3: The most capable openly available LLM to date, 2024. URL: <https://ai.meta.com/blog/meta-llama-3>.
- [10] Mistral AI, Codestral, 2024. URL: <https://mistral.ai/news/codestral>.
- [11] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford, et al., Gpt-4o system card, *arXiv preprint arXiv:2410.21276* (2024).
- [12] S. Ott, A. Barbosa-Silva, K. Blagec, J. Brauner, M. Samwald, Mapping global dynamics of benchmark creation and saturation in artificial intelligence, *Nature Communications* 13 (2022) 6793.
- [13] EvalPlus Leaderboard, Accessed : 2025-07-01. URL: <https://evalplus.github.io/leaderboard.html>.
- [14] D. Russo, Navigating the complexity of generative AI adoption in software engineering, *ACM Transactions on Software Engineering and Methodology* 33 (2024) 1–50.
- [15] R. Khojah, M. Mohamad, P. Leitner, F. G. de Oliveira Neto, Beyond code generation: An observational study of chatgpt usage in software engineering practice, *Proceedings of the ACM on Software Engineering* 1 (2024) 1819–1840.
- [16] S. Lambiase, G. Catolino, F. Palomba, F. Ferrucci, D. Russo, Exploring Individual Factors in the Adoption of LLMs for Specific Software Engineering Tasks, *arXiv preprint arXiv:2504.02553* (2025).

- [17] G. Hofstede, Culture's consequences: International differences in work-related values, volume 5, sage, 1984.
- [18] D. Chun, Z. Zhang, E. Cohen, L. Florea, O. F. Genc, Long-term orientation and the passage of time: Is it time to revisit Hofstede's cultural dimensions?, *International Journal of Cross Cultural Management* 21 (2021) 353–371.
- [19] M. Minkov, A. Kaasa, Do dimensions of culture exist objectively? A validation of the revised Minkov-Hofstede model of culture with World Values Survey items and scores for 102 countries, *Journal of International Management* 28 (2022) 100971.
- [20] G. Borchers, The software engineering impacts of cultural factors on multi-cultural software development teams, in: 25th International Conference on Software Engineering, 2003. Proceedings., IEEE, 2003, pp. 540–545.
- [21] S. Abufardeh, K. Magel, The impact of global software cultural and linguistic aspects on Global Software Development process (GSD): Issues and challenges, in: 4th International conference on new trends in information science and service science, IEEE, 2010, pp. 133–138.
- [22] S. Lambiase, G. Catolino, M. Torre, D. A. Tamburri, F. Palomba, F. Ferrucci, Fright not and be dispersed! evaluating cultural dimensions versus software communication and collaboration activities, *SSRN Electronic Journal* (2022).
- [23] G. Ghinea, B. Bygstad, M. Satpathy, Software Developers in India and Norway: Professional or National Cultures?, *Journal of Information Technology Research (JITR)* 4 (2011) 50–63.
- [24] H. Ibarra, Provisional selves: Experimenting with image and identity in professional adaptation, *Administrative science quarterly* 44 (1999) 764–791.
- [25] M. G. Pratt, K. W. Rockmann, J. B. Kaufmann, Constructing professional identity: The role of work and identity learning cycles in the customization of identity among medical residents, *Academy of management journal* 49 (2006) 235–262.
- [26] G. E. Kreiner, E. C. Hollensbe, M. L. Sheep, Where is the “me” among the “we”? identity work and the search for optimal balance, *Academy of Management Journal* 49 (2006) 1031–1057.
- [27] E. Brynjolfsson, A. McAfee, The second machine age: Work, progress, and prosperity in a time of brilliant technologies, WW Norton & company, 2014.
- [28] D. H. Autor, Why are there still so many jobs? the history and future of workplace automation, *Journal of economic perspectives* 29 (2015) 3–30.
- [29] R. Panguluri, Developing a Competitive Edge: Building an Effective Portfolio in Snowflake and Teradata Data Engineering, *International Journal of Scientific Research in Computer Science, Engineering and Information Technology* 10 (2024).
- [30] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. Gonzalez, H. Zhang, I. Stoica, Efficient memory management for large language model serving with PagedAttention, in: Proceedings of the 29th Symposium on Operating Systems Principles, 2023, pp. 611–626.
- [31] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, I. Sutskever, Robust speech recognition via large-scale weak supervision, in: Proceedings of the 40th International Conference on Machine Learning, ICML'23, JMLR.org, 2023.
- [32] A. Plaquet, H. Bredin, Powerset multi-class cross entropy loss for neural speaker diarization, in: INTERSPEECH 2023, ISCA: International Speech Communication Association, 2023, pp. 3222–3226.
- [33] H. Bredin, pyannote. audio 2.1 speaker diarization pipeline: principle, benchmark, and recipe, in: 24th INTERSPEECH Conference (INTERSPEECH 2023), ISCA, 2023, pp. 1983–1987.
- [34] A. S. Hayes, “Conversing” With Qualitative Data: Enhancing Qualitative Research Through Large Language Models (LLMs), *International Journal of Qualitative Methods* 24 (2025).
- [35] S. De Paoli, Performing an inductive thematic analysis of semi-structured interviews with a large language model: An exploration and provocation on the limits of the approach, *Social Science Computer Review* 42 (2024) 997–1019.
- [36] Anthropic, Introducing the next generation of Claude, 2024. URL: <https://www.anthropic.com/news/claude-3-family>.
- [37] The Culture Factor Group, Accessed : 2025-05-24. URL: <https://www.theculturefactor.com/>.