

Grassroots Innovation, Organisational Inertia: Generative AI in Italian Career Guidance Organisations

Leonardo Evangelista

orientamento.it | l.evangelista@orientamento.it | ORCID: 0009-0009-7855-4023

Working paper, 2026

DOI: 10.5281/zenodo.20088215

Abstract

Background. Research on generative AI in professional practice has documented rapid individual adoption and significant transformation of task-level practices. Less attention has been paid to the organisational dimension: whether and how the organisations in which practitioners work are responding to, supporting, or governing this transformation.

Aim. This paper examines the organisational impact of generative AI adoption in Italian career guidance services, as reported by nine practitioners who were already using generative AI tools in their guidance practice, at varying levels of intensity.

Method. The analysis draws on nine semi-structured in-depth interviews with career guidance practitioners from Northern Italy, conducted as part of a broader mixed-methods empirical study (n=81 questionnaire respondents). All interviews were fully transcribed and analysed thematically.

Findings. Eight of nine practitioners report an absence of organisational-level change in their services. The pattern is consistent across three dimensions: no structured training for guidance staff, resistance or indifference among colleagues, and no shared guidelines, policies, or modified service delivery processes. Only two partial exceptions emerge. The individual innovation of practitioners is not matched by organisational response. This misalignment produces three significant consequences: privacy compliance risks arising from the use of personal accounts, high variability of practices within the same organisation, and a failure to exploit AI's potential for rethinking service delivery processes. The competences practitioners develop remain individual assets, not organisational knowledge.

Keywords: *generative artificial intelligence, career guidance, career counselling, organisational change, professional development, AI governance, grassroots innovation, Italy*

1. Introduction

The diffusion of generative AI tools across knowledge-intensive occupations is generating a well-documented pattern of individual adoption: practitioners learn to use new tools on their own initiative, transform their task-level practices, and develop new competences through self-directed experimentation. What is less well documented is what happens at the organisational level: whether the organisations in which these practitioners work are aware of this transformation, whether they

are supporting or governing it, and what the consequences are when individual innovation is not matched by organisational response.

This paper addresses that question in the context of Italian career guidance. Drawing on nine in-depth interviews with practitioners who have individually integrated generative AI tools into their practice at significant levels of intensity, it examines the organisational dimension of AI adoption: training, governance, colleague relations, privacy compliance, and the structural consequences of the gap between individual innovation and organisational inertia.

The Italian career guidance system provides a relevant context for this analysis. The broader Italian career guidance system includes public employment centres, private employment agencies, training agencies, and freelance career counsellors. However, the interview sample does not include practitioners working directly inside public employment centres: evidence on those contexts comes mainly from the questionnaire data. The variation in organisational context within the interview sample, covering private employment agencies, training agencies, HR firms, and freelance practice, nonetheless allows for a comparative analysis of how different institutional environments shape, or fail to shape, the organisational response to individual AI adoption.

2. Methodology

The data in this paper derive from nine semi-structured in-depth interviews conducted as part of a larger mixed-methods empirical study on generative AI integration in Italian career guidance practice (Evangelista, 2026). The broader study employed three instruments: an online questionnaire (n=81 practitioners), nine in-depth interviews, and a post-interview quantitative evaluation questionnaire administered to seven of the nine interviewees.

The nine interviewees were recruited through purposive sampling, with two selection criteria: active practice of career guidance counselling, and existing use of generative AI tools in professional practice. Participants were identified through multiple channels: some were practitioners already known to the author, others responded to a call for participation circulated via LinkedIn and email, and others were drawn from the questionnaire respondents. All were practitioners from Northern Italy. Interviews lasted between 45 and 90 minutes, were conducted remotely, audio-recorded with participants' consent, and fully transcribed.

All interviewees are referred to as Op1 through Op9. All quotations are translated from Italian by the author. Table 1 provides an overview of the organisational parameters relevant to this analysis.

Table 1. Overview of organisational context and AI governance parameters (nine practitioners).

	Context	Org. AI guidance	Structured training	Colleague resistance	Account type / privacy measures
Op1	APL	None	None	Not mentioned	Company account; no privacy measures mentioned
Op2	HR firm / Freelance	None; self-organised internal seminar	None	'Each person in their own small patch'	Personal; no privacy measures mentioned
Op3	Training agency / Freelance	None	None	Not mentioned	Personal; no privacy measures mentioned

	Context	Org. AI guidance	Structured training	Colleague resistance	Account type / privacy measures
Op4	GOL 4 / Trainer	None	None	Training delivered but not applied by managers	Personal; extreme anonymisation practiced
Op5	APL / Freelance	None; fear and suspicion noted	None	Not mentioned	Personal; no privacy measures mentioned
Op6	Training agency (GOL 4)	None	None	'I don't see the will to use AI in my colleagues'	Personal; removes sensitive data; disabled training
Op7	Training agencies / Freelance	Peer sharing encouraged by management	None (peer sharing only)	Colleagues less technological, supported by team	Personal; privacy setting managed
Op8	Training agencies / Freelance	None; missing ethics and guidelines noted	None	Not mentioned	Personal; removes sensitive data
Op9	APL	Company GPT account provided; training only for recruitment team	None for guidance staff	Generational gap (9 years to next youngest colleague)	Company account; removes sensitive data; starts new chat per client

3. The Absence of Organisational Impact

The central finding of this analysis is stark: eight of nine practitioners report that their organisations have made no structural changes in response to the individual adoption of generative AI tools. The pattern is consistent across the organisational contexts represented in the interview sample, mainly private employment agencies, training agencies, HR firms, and freelance practice. Evidence concerning public employment centres comes from the questionnaire data and from the broader discussion in the research volume.

Op5 offers the most direct summary:

"Absolutely nothing, nothing at all. In some cases I even sense some fear or suspicion about AI use."

Op1 situates the absence in the broader context of public employment services:

"As far as active labour market policies are concerned, it does not seem like it yet. We are using it too little, we have no proper training. It is being used at an intuitive level."

Op2 describes a context in which any organisational diffusion of AI practice has depended entirely on individual initiative:

"Not much, almost nothing. I had to push personally to run an internal seminar to explain to others how to use it."

Op2 also identifies the dominant organisational dynamic:

"Either the indications come from above, or everyone develops their own isolated practice."

This formulation captures the structural condition with precision: in the absence of organisational direction, AI adoption fragments into isolated individual practices, each practitioner developing their own approach independently of colleagues.

4. Three Dimensions of Organisational Absence

4.1 No structured training for guidance staff

None of the nine practitioners reports having received structured training on generative AI from their organisation. All have learned through self-directed experimentation, online resources, or informal peer exchange. Op1:

"We have no proper training. It is being used at an intuitive level."

Op9 describes a situation that illustrates the internal fragmentation within organisations: training was provided for the recruitment team, which uses AI intensively for candidate screening, but not for the guidance team:

"Training was done only for our colleagues in the personnel recruitment section, because they use it a lot. We did not receive training, even though we can use it."

This asymmetry reveals that the absence of training for guidance staff is not a general organisational indifference to AI, but a specific gap: guidance services are not prioritised in organisational AI strategies that do exist, even partially, in the same institutions.

All practitioners who have developed significant AI competences have done so as autodidacts. The passion for technology and personal curiosity, as several describe it, are the primary drivers of learning: individual dispositions, not organisational investment.

4.2 Colleague resistance and indifference

The absence of organisational training is compounded by resistance or indifference among colleagues. Op6:

"The main problem is the will to use AI that I do not see in my colleagues."

Op9 attributes part of this resistance to generational factors:

"I am the youngest here. The next youngest colleague is nine years older than me, so my approach is completely different from that of my colleagues in my division."

Op4 describes a frustrating dynamic in a training context where AI has been demonstrated but not adopted:

"The problem is that I explain to them how to use AI to simplify or better carry out certain management tasks, but then they do not do them and so they forget."

Op7 is the only practitioner who describes a positive collegial environment, where peer exchange on AI practice is actively cultivated:

"In the organisation where I work the good thing is that we are a team, there is a lot of exchange among us, we help those colleagues who are less technologically confident or wary of AI. We have periodic meetings, encouraged by management, where we talk about our work; in these meetings we sometimes also discuss how we use AI sites."

This account is notable precisely because it is exceptional. The conditions that Op7 describes, management encouragement, a collaborative team culture, regular structured exchange, are entirely absent in the accounts of the other eight practitioners.

4.3 No guidelines, policies, or process modifications

The third and most significant dimension of organisational absence is the complete lack of shared governance frameworks. No practitioner reports the existence of organisational protocols for AI use, guidelines on data protection in AI-assisted sessions, criteria for deciding when to use AI with clients, or modifications to service delivery processes to incorporate AI systematically.

Op8 names this absence explicitly:

"What is missing is an ethical framework, that is, guidelines for the use of AI tools."

The absence of guidelines has a direct consequence for privacy compliance, examined in the following section.

5. The Partial Exceptions

Two cases offer partial, and instructive, exceptions to the pattern of organisational absence.

Op9's organisation provided a company GPT account for all staff, actively promoted by management. This is a significant organisational step: it addresses the privacy compliance problem associated with personal accounts and signals institutional recognition of AI as a legitimate professional tool. However, the initiative stopped there:

"The manager pushed for all divisions to have their own GPT profile, providing shared accounts for practitioners. However, this initiative was limited to providing the tool without training for guidance staff."

The provision of a tool without training, guidelines, or process modification is organisational recognition without organisational integration. It creates the infrastructure for AI use without the conditions for its responsible and effective deployment.

Op7's case is different in character. Rather than a top-down provision of tools, what Op7 describes is a management-encouraged culture of peer learning: periodic team meetings in which AI practice is discussed alongside other professional topics. This is not a structured training programme, but it creates a form of distributed organisational learning that is entirely absent elsewhere in the sample. It is, however, entirely dependent on the existence of a collaborative team culture and supportive management, conditions that cannot be assumed across diverse organisational settings.

6. Three Consequences of the Individual-Organisation Gap

6.1 Privacy compliance risks

The most immediate and legally significant consequence of organisational inertia is the privacy compliance risk created by the widespread use of personal or non-business accounts. Only two practitioners report using business or company accounts: Op1, whose organisation has a company subscription, and Op9, whose organisation proactively provided a shared GPT account. For the remaining practitioners, account status is personal or unclear. The broader study identifies possible GDPR non-compliance for six of the nine practitioners.

The practitioners themselves are aware of the risk and have developed individual mitigation strategies. Op4 practices systematic and demanding anonymisation:

"I am very attentive to data protection. I have disabled the option to use responses for training in both GPT and Claude. I have learned that in handwritten documents it is not enough to blacken out words I do not want it to read: you have to actually delete the sensitive data. I remove everything that could identify the person: for example, if the person has worked for Coca Cola there is no identification problem, but if they have worked in a small specific company in a small town in a specific period, the person becomes identifiable and I cut that information."

Op6 adopts a lighter approach:

"There is always the privacy problem, so I do not enter sensitive data. If I am looking for occupations compatible with certain medical conditions, I enter the condition and nothing else, without additional data."

Op8:

"I am very careful about privacy: when I load personal data I avoid entering sensitive data."

Op9:

"I copy and paste my notes removing name, surname, and all sensitive data, because I think this is an ethical problem: I do not know where this data goes."

These individual strategies are thoughtful and well-intentioned. They are also, by definition, inconsistent: each practitioner makes their own judgements about what constitutes adequate anonymisation, with no shared standards, no verification, and no organisational accountability. The risk is not eliminated; it is distributed across nine individual risk-management approaches of varying rigour.

6.2 High variability of practices

The second consequence of organisational inertia is the high variability of AI practices within the same organisations and, more broadly, across the guidance sector. In the absence of shared guidelines, each practitioner develops their own approach to every dimension of AI use: which tools to use, how to introduce them to clients, what tasks to use them for, how to manage privacy, whether and how to teach AI to clients.

The research documents substantial variation on all these dimensions. Some practitioners use AI in every session; others in fewer than 5%. Some teach AI systematically to all digitally able clients; others rarely or never. Some use AI for the full range of eight functions identified in the broader study; others confine it to document production. This variation is not problematic in itself: different practitioners in different contexts serving different client populations will reasonably make different choices. What is problematic is when variation is driven by the absence of shared standards rather than by legitimate contextual adaptation.

The practical consequence is that two clients of the same organisation, assigned to different practitioners, may receive substantially different services: one incorporating AI as a structural tool across the full guidance process, the other receiving no AI assistance at all.

6.3 Failure to exploit AI's organisational potential

The third consequence is, in a sense, an opportunity cost rather than a harm: the failure of organisations to exploit AI's potential for rethinking service delivery at the process level. The individual practitioners in this study have demonstrated that AI can substantially accelerate core guidance tasks, improve the quality of outputs, extend the practitioner's informational reach, and enable new forms of support that were previously impossible or impractical.

None of these improvements has been translated into organisational process redesign. No organisation in the sample has used AI adoption as an occasion to rethink caseload management, session structures, inter-session client support, referral protocols, or quality assurance processes. The efficiency gains are real but remain individual: they benefit each practitioner's own practice without generating collective organisational improvement.

The competences that practitioners develop through self-directed AI learning are similarly confined to individual possession. They are not codified, not transmitted systematically to colleagues, and not protected against the risk of loss through staff turnover.

7. Discussion

The findings document a pattern that is familiar from the broader literature on technology adoption in public and professional services: individual innovators transform their practice ahead of organisational response, creating a structural lag between the frontier of individual practice and the baseline of organisational provision.

What makes the career guidance case particularly significant is the character of the tools being adopted. Generative AI tools are not auxiliary productivity aids that can be added to existing workflows without affecting their structure. As the other papers in this series document, they reconfigure the practitioner's role, modify the dynamics of guidance sessions, expand the scope of client autonomy, and create new ethical responsibilities. These are changes that require organisational attention, not merely individual management.

The privacy compliance issue is the most urgent. The use of personal accounts for processing client personal data is not a minor procedural irregularity: it is a potential GDPR violation affecting vulnerable clients who have shared sensitive personal information in a professional context. The individual anonymisation strategies practitioners have developed are creative responses to an organisational failure, not adequate substitutes for proper governance.

The colleague resistance documented across the sample points to a professional development challenge that individual initiative cannot address. Practitioners who have developed significant AI competences describe colleagues who have not, and the gap is widening as the competence frontrunners continue to develop while their organisations provide no mechanism for broader diffusion. This dynamic, if sustained, will produce a two-tier guidance workforce within the same institutions: AI-capable practitioners and those who are not, with significant implications for service equity.

The two partial exceptions, Op9's company account and Op7's peer learning culture, are instructive precisely because they are partial. They show that organisational responses are possible, and that even limited organisational action (providing a compliant account, encouraging peer exchange) can create conditions for more responsible and effective individual practice. They also show that these conditions do not emerge spontaneously: they require deliberate organisational choice.

8. Connection to the Broader Study and Further Materials

This paper draws on material from: Evangelista, L. (2026). *Tra Mercurio e Virgilio: l'integrazione dell'intelligenza artificiale generativa nella consulenza di orientamento. Un'indagine empirica sugli operatori italiani*. Amazon KDP. <https://doi.org/10.5281/zenodo.19855951>

The full research volume (in Italian) is available at: <https://www.orientamento.it/intelligenza-artificiale-e-orientamento-professionale-una-ricerca-sugli-operatori-italiani/>

Additional materials related to this study, including an executive report, translated practitioner testimonies, and papers on the eight functions of AI in guidance, AI use within the guidance session, and the disintermediation risk, are freely available at: <https://www.orientamento.it/generative-ai-in-career-guidance-practice-evidence-from-italian-practitioners/>

The author welcomes correspondence from researchers working on related questions.

References

Evangelista, L. (2026). Tra Mercurio e Virgilio: l'integrazione dell'intelligenza artificiale generativa nella consulenza di orientamento. Un'indagine empirica sugli operatori italiani. Amazon KDP. <https://doi.org/10.5281/zenodo.19855951>