

AI in Career Guidance: Opportunities, Risks and Ethical Boundaries

Antonios Zachariadis¹, Niki Daliana², Alexandros-Stamatios Antoniou³

¹Independent Legal Professional, Data Protection Officer, Greece

²Career Flame, Athens, Greece

³Department of Primary Education, School of Education, National and Kapodistrian University of Athens, Athens, Greece

Abstract

Artificial intelligence is increasingly used in career guidance through chatbots, recommender systems, predictive analytics and interactive platforms. These tools may improve access and personalisation, but career guidance is not a neutral information service. It can influence how people understand their skills, their future options and their sense of agency. This article provides a narrative review and normative analysis of recent academic, professional and regulatory sources on AI-assisted career guidance. The discussion focuses on opportunities, risks and ethical boundaries, with particular attention to human oversight, personal data, fairness, transparency and the distinction between career discovery and career decision-making. The literature suggests that AI can expand access to guidance, support lifelong learning, connect users with changing labour-market information and make exploration more interactive. At the same time, the evidence base remains uneven. Important risks include biased or outdated data, privacy and secondary use of personal data, hallucinated or weakly verified outputs, narrow recommendation loops, automation bias and the possible weakening of the human relationship in guidance. Concluding, ethical AI in career guidance requires more than enthusiasm for innovation. It requires explainable outputs, data minimisation, quality control, practitioner involvement, user contestability, AI literacy and clear human responsibility. The article argues that the most important boundary is the line between career discovery and career decision. AI systems that help users explore possibilities may be valuable; systems that substantially determine educational or employment outcomes require stricter governance and human oversight.

Keywords

artificial intelligence; career guidance; career counselling; ethics; personal data; human oversight; AI Act

Corresponding Author: Niki Daliana, Career Flame, Athens, Greece, ndaliana@yahoo.gr

Introduction

Artificial intelligence is gradually becoming part of the ordinary infrastructure through which people search, learn and make plans. Career guidance has entered this shift as well. Many users can now complete an online assessment, receive a set of suggested pathways, ask follow-up questions through a chatbot and compare their profile with labour-market information in a single digital environment. This can make guidance easier to reach. It also changes the tone of the guidance encounter. The user is no longer only reading static information or completing a fixed questionnaire; in many cases, the user is interacting with a system that appears to respond personally.

Career guidance has never been completely separate from technology. Databases, online inventories and computer-assisted guidance systems have been used for decades. The difference with AI is that the system is no longer limited to a transparent set of rules designed in advance. Wilson et al. describe careers information, advice and guidance as a field that has historically adopted new technologies while also questioning their risks; however, AI differs from earlier systems because it can deal with interactions that were not fully specified beforehand (1). Westman et al. also connect the current interest in AI with the wider lifelong guidance agenda, where education, work, upskilling and reskilling have become increasingly linked (2). The question, then, is not simply whether technology belongs in guidance. It already does. The more difficult question is what happens when a digital tool begins to resemble an adviser, a recommender and a decision-support system at the same time.

In this article, career guidance is used as a broad term for services that help individuals explore, plan and manage educational and occupational pathways across the lifespan. Career counselling refers more specifically to the relational and developmental part of this work: the process through which a professional helps a person examine values, identity, uncertainty and decision-making over time. High-risk AI is used here in the sense of the EU AI Act, referring to systems whose use in education, employment or access to services may significantly affect rights and life chances (16). These distinctions matter because the same technical system may have different implications depending on whether it is used for open exploration or for a consequential decision.

The urgency of this discussion is partly explained by the labour market itself. The IAEVG 2024 Communiqué states that AI technologies are reshaping work, learning and professional standards, while also creating new ethical responsibilities for career, educational and vocational guidance (3). Wong argues that AI-driven automation complicates career planning for secondary students because young people may choose pathways without fully understanding how tasks and occupations are changing (4). In higher education, recent evidence suggests a move from older computer-assisted systems toward generative agents, predictive analytics and recommender ecosystems (6). Guidance services are therefore asked to explain a changing labour market while also using the very technologies that are helping to change it.

This article examines the opportunities, risks and ethical

boundaries of AI-driven career guidance. The argument is practical rather than speculative. AI can strengthen guidance when it widens access, helps users explore possibilities and supports practitioners. It becomes risky when its outputs are opaque, biased, overconfident or treated as final answers. The central distinction proposed here is between career discovery and career decision-making, a distinction consistent with vocational psychology frameworks that separate exploration from commitment (17). Discovery invites reflection and widens options. Decision-making may affect education, employment and identity in ways that require stronger safeguards.

Approach and limits of the review

This article is a narrative review and normative analysis. It does not report new empirical data, and it should not be read as an original research study. The discussion draws on peer-reviewed articles, systematic reviews, professional guidance and European regulatory materials that address AI in career guidance, recommender systems, data protection and high-risk AI governance. The purpose was not to exhaust every publication on digital guidance, but to bring together sources that are directly relevant to the practical question of how AI should be used in this field.

The review gives particular weight to sources that are close to career guidance and counselling. Wilson et al. are used to frame the professional risks of AI in career development practice (1). Westman et al. provide an applied account of requirements, possibilities and maturity levels for AI in higher-education guidance (2). The IAEVG Communiqué is treated as a professional reference point because it explicitly concerns AI in career, educational and vocational guidance (3). Wong is used to connect AI automation with the career-planning needs of younger users (4). Muhammad and Sarmurzin et al. are used to assess the still developing evidence base and the emerging types of AI implementation in career counselling (5,6). Legal and ethical sources are then used to interpret these findings in relation to privacy, fairness, transparency, accountability and human oversight (10,13-16).

There is a limit to this approach. Narrative reviews can clarify a field and connect strands of evidence, but they cannot estimate effectiveness in the way a meta-analysis can. For that reason, the claims made here are deliberately cautious. The article does not argue that AI career guidance is effective in general. It argues that some uses are promising, some are risky and the difference depends on design, context, governance and the role left for human judgement.

Opportunities created by AI in career guidance

One of the clearest advantages of AI-supported guidance is practical access. In many services, the problem is not the absence of professional expertise, but the fact that this expertise is not always available when the user begins looking for help. Appointments, geography, staffing shortages and cost

can delay or prevent support. AI-supported tools can offer a first point of contact, provide basic information outside office hours and help users arrive at a later counselling session with more focused questions. The IAEVG Communique recognises this potential, particularly for people who cannot easily access in-person services or who prefer online or anonymous interaction (3). Chatbot-based guidance systems illustrate the same shift at a simple level, because they can provide continuous responses and basic guidance interactions at scale (7).

Personalisation is another reason these tools have attracted attention. Career guidance often requires the user to connect several kinds of information: interests, skills, previous learning, labour-market trends, educational options and practical constraints. AI systems can organise large quantities of information and produce an initial map of options that may otherwise be difficult to assemble. Westman et al. show how AI may support competence mapping, case management and the matching of user needs with guidance resources, provided that the organisation has adequate data, processes and staff competences (2). This is not a replacement for professional judgement. It is better understood as a way of preparing information before judgement is exercised.

AI may also help users see emerging or less visible career paths. In periods of labour-market disruption, traditional occupational labels can lag behind the actual organisation of work. A role may not yet have a stable title, while the underlying skill combination is already becoming relevant. Wong emphasises that students need to understand how automation affects sectors and skills before they can plan realistically (4). Used carefully, AI can help make the labour market less abstract. It can show not only which occupations appear to fit a profile, but also which skills are becoming transferable, which roles are changing and which adjacent options the user may not have considered.

Interactive formats may also increase engagement. Many users, especially younger users, do not necessarily respond well to static lists or long questionnaires. Conversational, narrative and gamified interfaces can make exploration less intimidating. This benefit should be treated carefully. Engagement is not the same as accuracy, and a friendly interface can make a weak recommendation appear more trustworthy than it deserves to be. Even so, interactive formats can lower the threshold for starting a guidance process when they are designed to lead users toward reflection rather than premature certainty.

Finally, AI may support practitioners themselves. The relevant question is not only what the system says to the user, but how it can help guidance professionals manage increasing demand. AI may help staff prepare sessions, identify common information needs, update labour-market resources, detect skill gaps or monitor demand across services. Sarmurzin et al. identify several implementation types in university career counselling, including generative chatbots, predictive analytics and recommender ecosystems (6). These uses are most defensible when they remain part of a hybrid model in which the practitioner interprets, questions and contextualises the output.

Risks and limitations

A still limited evidence base

The first limitation is the evidence base itself. Interest in AI career guidance is growing faster than the long-term evaluation of its effects. Muhammad's systematic literature review found a limited number of empirical studies on AI-assisted career counselling, and only a small subset addressed barriers directly (5). Sarmurzin et al. provide a more recent systematic review in university settings and show that research is expanding, but they also call for longitudinal evaluation, algorithmic audits and more work on explainable AI (6). This matters because guidance outcomes are not always visible immediately. A system may feel helpful in the moment and still narrow a person's options over time.

For this reason, claims about effectiveness need to remain modest. It is reasonable to say that AI can support access, information retrieval and initial exploration. It is much less reasonable to claim that AI can replace the relational, developmental and ethical work of career counselling. The current literature does not justify that conclusion.

Bias, outdated data and unfair outputs

AI guidance depends on data. If the training data, labour-market data or user profile data are biased, incomplete or outdated, the output will carry those weaknesses into the guidance process. The IAEVG Communique warns that poor-quality AI services may mislead users, distort labour-market understanding and reinforce stereotypes when systems are trained on inadequate, biased or obsolete data (3). In career guidance, this is not a minor technical issue. A biased recommendation can shape how a person interprets their own capacity and future.

Recruitment provides a useful warning. The case of Amazon's recruiting algorithm, reported to have disadvantaged applications associated with female candidates because it was trained on historically male-skewed hiring data, is often used to illustrate how inequality enters algorithmic systems through past data. Chang analyses this case in the context of algorithmic gender bias (9). Career guidance is not the same as recruitment, but the lesson transfers. If a system learns from a labour market that has already been unequal, it may reproduce that inequality as advice. A user may then experience the result not as a reflection of biased data, but as a neutral assessment of their potential.

Privacy, profiling and secondary use of data

Career guidance platforms may collect educational history, employment history, skills, interests, behavioural data, interaction logs and sometimes sensitive inferences about motivation, uncertainty or psychological needs. This information can be useful for personalisation, but it also creates a detailed profile of the user. The GDPR principles of lawfulness, fairness, transparency, purpose limitation and data minimisation are therefore central to the design of such systems (13).

The privacy problem becomes sharper when large language models are involved. Barberá's report for the European Data

Protection Board identifies risks such as unlawful processing of personal data in training sets, repurposing input or output data for further training, and pressure on the data minimisation principle created by large-scale model development (10). In a career-guidance setting, these risks are concrete. A user may disclose uncertainty about education, work, finances or personal constraints. They should not have to guess whether these disclosures will later be reused to improve a model, profile them further or support unrelated purposes.

Hallucinations and the appearance of certainty

Generative AI can produce fluent outputs that sound more certain than the underlying evidence allows. This is particularly dangerous in career guidance because users may ask about salaries, job security, qualifications or labour-market demand. The World Health Organization has warned, in the health context, that large multi-modal models can generate inaccurate or hallucinated content even when the language is confident (11). Li also discusses hallucinations as misleading AI-generated content that requires fact-checking and information-evaluation skills (12).

The problem is not only that a fact may be wrong. It is that the user may not recognise it as wrong. A young person considering a field of study, or an adult considering a career change, may act on inaccurate information if the system presents it as settled. AI career platforms should therefore separate generated explanation from verified labour-market evidence and show users where factual claims come from.

Recommendation loops and narrowed horizons

Recommender systems are useful because they reduce complexity. They are risky for the same reason. Milano et al. describe the ethical challenges of recommender systems, including the way personalised systems can shape the options that users see and the choices they consider (8). In career guidance, this risk is especially important. Guidance should not merely confirm the user's past profile. It should also open paths the user may not have known how to name.

A platform that repeatedly recommends familiar options can make its own narrowing effect invisible. A student with a humanities background may never see credible routes into technology-related work. A user from a lower-income environment may receive "realistic" recommendations that reflect existing social patterns rather than genuine potential. In such cases, the system does not simply advise the user. It quietly defines the horizon of possible futures.

Automation bias and loss of agency

A related risk is automation bias. Sarmurzin et al. argue that students may passively accept algorithmic recommendations rather than actively constructing their vocational identities (6). This risk is especially relevant when the user is uncertain, young, under pressure or lacking other sources of guidance. Career guidance is not only about matching a profile to an occupation. It is also about helping a person ask better questions, examine assumptions, tolerate uncertainty and build agency.

Westman et al. frame this issue through agency in guidance. AI can support students and staff, but the value of the tool depends on user needs, staff competences and the organisation's capability to integrate technology responsibly (2). If the system becomes the dominant voice in the process, the user may stop exploring and start complying. That would be the opposite of good guidance.

The role of the human counsellor

The human counsellor does more than provide information. A counsellor notices hesitation, asks why an option feels impossible, challenges a narrow self-image and helps the user connect decisions with values and circumstances. This relational and developmental dimension of guidance is well established in vocational psychology. Counselling approaches grounded in career construction theory and life design emphasise the co-construction of meaning between practitioner and client as central to effective career intervention (17).

The IAEVG Communique is clear that AI should not be treated as a cost-saving replacement for face-to-face guidance or as a reason to reduce resources for personalised support (3). Wilson et al. reach a similar conclusion, arguing that the introduction of AI into career development should focus on increasing value for clients rather than reducing staff costs (1). This does not mean that every AI interaction must be supervised in real time. It means that systems should be designed with clear thresholds. Low-stakes exploration may be automated to a greater degree. High-stakes choices, vulnerable users, unclear outputs, emotional distress, or recommendations likely to affect education or employment should trigger human review or referral.

Ethical boundaries for responsible use

Transparency and explainability

Users should know when they are interacting with AI. They should also know what kind of data the system uses and how much weight should be given to the output. Transparency cannot be reduced to a long privacy notice. In guidance, transparency has to be understandable at the point of use. A user should be able to see whether a recommendation is based on stated interests, test responses, behavioural signals, labour-market data or a combination of these.

The EU AI Act requires appropriate transparency for users and, for high-risk systems, places emphasis on documentation, data governance, human oversight and risk management (16). Sarmurzin et al. also identify explainable AI as a priority for future research in university career counselling (6). In practice, this means that a platform should not only provide an answer. It should show why that answer appeared, what evidence supports it, what uncertainty remains and how the user can challenge or reinterpret it.

Fairness and data quality

Fairness requires active testing. It is not enough for a provider

to state that a system is neutral. Career guidance platforms should be checked for unequal patterns across gender, socio-economic background, disability, language, migration status and other relevant characteristics. The EDPB has linked fairness in AI models to the avoidance of unjustifiably detrimental, discriminatory, unexpected or misleading processing (15).

Data quality is part of fairness. A system using outdated labour-market information may mislead all users, but not all users equally. Those with fewer external resources may be more likely to rely on the platform as an authority. For this reason, AI career guidance should include update cycles, source labelling, bias testing and human review of recommendation categories.

Data minimisation and purpose limitation

The desire to personalise guidance should not become a justification for collecting everything. The GDPR requires that data be collected for specified, explicit and legitimate purposes and limited to what is necessary (13). In career guidance, this principle should be operational rather than decorative. A platform should be able to explain why each category of data is collected, how long it is retained, whether it is used for model improvement and whether the user can receive guidance without providing more intimate information.

The same principle applies to behavioural data. Clicks, pauses, repeated choices and interaction patterns can be useful, but they can also create sensitive inferences. If such signals are used, users should be told in plain language. They should also have the option to opt out of non-essential processing without losing access to basic guidance support.

Accountability and professional involvement

Accountability means that responsibility cannot disappear inside the model. If an output is harmful, misleading or discriminatory, there must be a person or organisation responsible for design, deployment, monitoring and correction. The GDPR accountability principle requires controllers to demonstrate compliance, and Articles 24 and 25 require appropriate technical and organisational measures, including data protection by design and by default (13).

Professional involvement is also necessary. Wilson et al. argue that CIAG experts and practitioners should be involved throughout system design and deployment, because software vendors may not have the professional knowledge needed to evaluate ethical compliance in career guidance (1). This is a crucial point. A technically functional system may still be professionally weak if it misunderstands what guidance is for.

Human oversight and contestability

Human oversight should not be symbolic. It should be available at the moments when the output may affect a user's educational or professional path in a significant way. Article 22 GDPR gives individuals protections in relation to solely automated decisions that produce legal or similarly significant effects, including the right to obtain human intervention and to contest the decision (13). The AI Act also places human oversight at the centre of high-risk AI governance (16).

For AI career guidance, contestability can be practical. Users should be able to say that a recommendation does not fit, correct the profile, ask for alternative pathways, see broader options and request human support. A system that cannot be questioned is unsuitable for guidance, even if its average accuracy appears acceptable.

AI literacy and organisational maturity

Ethical use also depends on the competence of the people and organisations deploying the technology. The IAEVG Com-muniqué calls for AI literacy, ethical technology use and data interpretation skills among guidance practitioners (3). Westman et al. offer a useful maturity model for AI in career guidance, moving from AI-aware guidance toward AI-informed, AI-integrated and AI-transformed guidance across data, technology, processes, services and competences (2).

This maturity perspective prevents a common mistake: treating AI adoption as the purchase of a tool. Responsible adoption also requires governance routines, staff training, monitoring, escalation pathways, documentation and evaluation. Without these, even a promising system may become unsafe in practice.

Career discovery versus career decision

The most important ethical boundary is the distinction between career discovery and career decision. Career discovery uses AI to help users explore possibilities, notice patterns, ask questions and prepare for reflection. Career decision-making uses AI to recommend, rank or determine outcomes that may shape education, employment or access to opportunity. The first can be a valuable support. The second requires much stricter safeguards.

This distinction is also legally relevant. Under the EU AI Act, systems used in employment, worker management and access to education may fall under high-risk categories depending on their function and impact (16). A platform that helps a user explore possible roles is not the same as a platform that ranks candidates, determines admission, screens suitability or pushes users toward a single pathway. The difficulty is that the boundary can blur. A recommendation presented with too much certainty may feel like a decision even when the provider calls it "guidance".

For that reason, platforms should be careful with language. They should avoid deterministic statements such as "you should become" or "this is your best career". Better wording keeps agency with the user: "these paths may be worth exploring", "this result is based on the information provided", "other options may also be relevant", and "a human counsellor can help interpret this result". This is not only a matter of tone. It changes how the user relates to the output.

Concluding remarks

AI-driven career guidance can make guidance more accessible, more responsive and better connected to changing la-

bour-market information. It can help users organise complex information and can support practitioners who work under increasing demand. These benefits should not be dismissed. At the same time, they do not justify uncritical adoption.

The risks are substantial. Biased data can reproduce inequality. Privacy failures can expose highly personal information. Hallucinated outputs can mislead users. Recommender systems can narrow horizons while appearing personalised. Automation bias can weaken the user's own agency. Most importantly, the human dimension of guidance can be lost if AI is treated as a replacement for professional judgement rather than as a tool that supports it.

For the time being, the most defensible position is not full automation, but a hybrid model. AI may assist with information, exploration, pattern recognition and preparation. Human professionals should remain involved where interpretation, vulnerability, contested outputs or significant educational and employment consequences are present. The current evidence supports cautious integration, not substitution.

Future work should move beyond general claims about innovation and test AI career guidance systems in real-world conditions. Longitudinal studies, bias audits, explainability checks, privacy impact assessments and user-agency measures are needed. Until such evidence is stronger, AI should be used to open career possibilities, not to close them too early.

Funding: No external funding was received for this article.

Conflict of interest: The authors are involved in the development of a career-discovery platform focused on AI-supported exploration. This relationship is disclosed because the article discusses the value and appropriate design of such tools. No user data from that platform were used in this article, and the manuscript does not evaluate the platform or present product-specific findings. The argument developed here concerns AI career guidance tools in general and should not be read as promotional material.

Ethics approval and informed consent: Not applicable. This is a review and normative analysis and did not involve human participants, animal subjects, tissue samples, identifiable case material or experimental procedures.

References

- (1) Wilson M, Robertson P, Cruickshank P, Gkatzia D. Opportunities and risks in the use of AI in career development practice. *Journal of the National Institute for Career Education and Counselling* 2022, 48:48-57, doi: 10.20856/jnicec.4807
- (2) Westman S, Kauttonen J, Klemetti A, Korhonen N, Manninen M, Mononen A et al. Artificial intelligence for career guidance - Current requirements and prospects for the future. *IAFOR Journal of Education: Technology in Education* 2021, 9:43-62, doi: 10.22492/ije.9.4.03
- (3) International Association for Educational and Vocational Guidance. The role of artificial intelligence (AI) in career, educational and vocational guidance: Opportunities and challenges. *IAEVG Communique*, adopted by the IAEVG General Assembly, 13 November 2024. Available from: <https://iaevg.com/Communiques>
- (4) Wong LPW. Artificial intelligence and job automation: Challenges for secondary students' career development and life planning. *Merits* 2024, 4:370-399, doi: 10.3390/merits4040027
- (5) Muhammad R. Barriers and effectiveness to counselling careers with Artificial Intelligence: A systematic literature review. *Ricerche di Pedagogia e Didattica - Journal of Theories and Research in Education* 2023, 18:143-164, doi: 10.6092/issn.1970-2221/18250
- (6) Sarmurzin Y, Azhenov A, Kudysheva A, Ualikhanova B, Saparbekova I, Besbayev A. Implementation of AI in career counselling for university students: a systematic review. *Frontiers in Education* 2026, 11:1787689, doi: 10.3389/feduc.2026.1787689
- (7) Ohm A, Bhavani K. Chatbot for career guidance using AI. *International Journal of Computer Sciences and Engineering* 2019, 7:856-860, doi: 10.26438/IJCSE/V7I6.856860
- (8) Milano S, Taddeo M, Floridi L. Recommender systems and their ethical challenges. *AI & Society* 2020, 35:957-967, doi: 10.1007/s00146-020-00950-y
- (9) Chang X. Gender bias in hiring: An analysis of the impact of Amazon's recruiting algorithm. *Advances in Economics, Management and Political Sciences* 2023, 23:134-140, doi: 10.54254/2754-1169/23/20230367
- (10) Barberá I. AI privacy risks and mitigations: Large language models (LLMs). Support Pool of Experts Programme, European Data Protection Board, 2025. Available from: https://www.edpb.europa.eu/our-work-tools/our-documents/support-pool-experts-report/ai-privacy-risks-and-mitigations-large_en
- (11) World Health Organization. Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. World Health Organization, Geneva, 2024. Available from: <https://www.who.int/publications/i/item/9789240084759>
- (12) Li Y. Addressing hallucinations in AI-generated content: Strategies for developing student fact-checking and information evaluation skills. *Artificial Intelligence Education Studies* 2025, 1:48-62, doi: 10.6914/aiese.010204
- (13) European Parliament and Council. Regulation (EU) 2016/679 - General Data Protection Regulation. Official Journal of the European Union, 2016. Available from: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>
- (14) European Data Protection Board. Report of the work undertaken by the ChatGPT Taskforce, adopted on 23 May 2024. European Data Protection Board, Brussels, 2024. Available from: https://www.edpb.europa.eu/our-work-tools/our-documents/report/report-work-undertaken-chatgpt-taskforce_en
- (15) European Data Protection Board. Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, adopted on 17 December 2024. European Data Protection Board, Brussels, 2024. Available from: https://www.edpb.europa.eu/our-work-tools/our-documents/opinion/opinion-282024-certain-data-protection-aspects-related-processing_en
- (16) European Parliament and Council. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, 2024. Available from: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>
- (17) Savickas ML. Career counseling. *American Psychological Association*, Washington DC, 2011, doi: 10.1037/12410-000