

Reconfiguring the Supervisor-Student Relationship in the Age of Artificial Intelligence: Challenges, Opportunities, and Governance Paths

Qing Zhao

School of Foreign Studies, Northwestern Polytechnical University, Xi'an 710129, Shaanxi, China

Abstract: *Generative artificial intelligence is profoundly reshaping the operational logic of postgraduate supervisor-student relationships. As AI becomes embedded in research practices, the traditional dyadic interaction between supervisor and student is transforming into a triadic structure of “supervisor-AI-student,” giving rise to multiple educational predicaments: the epistemic authority of supervisors is progressively diluted; students develop intellectual inertia and their capacity for independent critical thinking deteriorates; and the interpersonal emotional bonds between supervisor and student are steadily weakened. At the same time, however, AI also presents opportunities for reconstruction: automation of routine tasks liberates supervisors’ time for higher-order guidance; data-driven feedback enables precision-tiered mentoring; the traditional master-apprentice model is propelled toward a human-AI collaborative paradigm; and a new, open, and shared supervisory ecosystem begins to take shape. Grounded in the fundamental objective of fostering virtue through education, this paper develops a four-dimensional collaborative governance framework: clarifying the boundaries of AI technological intervention, refining supporting institutional mechanisms and supervisor evaluation systems, simultaneously cultivating AI digital literacy among both supervisors and students, and reconstructing a humanistic supervisory culture and community. The paper argues that the crux of reconfiguring the supervisor-student relationship in the age of AI lies not in technology adoption per se, but in striking a balance between enablement and regulation, adhering to the educational mission and ensuring that technology serves rather than dissolves the core values of the supervisory relationship.*

Keywords: Artificial Intelligence, Supervisor-Student Relationship, Postgraduate Education, Human-AI Collaboration.

1. Introduction

The supervisor-student relationship constitutes a foundational interpersonal bond in postgraduate education. It serves not only as the institutional vehicle for knowledge transmission and academic training, but also as a critical arena for shaping academic character and conducting value-oriented guidance. Over time, through sustained face-to-face interaction, supervisors have guided students in developing academic thinking and completing academic socialisation. This deep human-to-human connection has long been the irreplaceable cornerstone of high-level talent cultivation. The proliferation of generative artificial intelligence has fundamentally rewritten the operational landscape of traditional supervision. AI can independently perform literature organisation, text polishing, idea construction, and a substantial range of other research tasks, opening up external academic channels for postgraduate students that operate independently of their supervisors. The traditional dyadic master-apprentice model is progressively transforming into a triadic mediated interaction structure of “supervisor-AI-student.” While technology has undoubtedly enhanced the efficiency of research-related tasks, it has also given rise to structural tensions: students become susceptible to tool dependency and experience a weakening of independent critical thinking; the supervisor’s established epistemic authority is continuously eroded; and the depth of supervisory dialogue, rich with humanistic warmth, is steadily diminished. The changes brought about by technological transformation extend far beyond superficial adjustments at the instrumental level. They constitute a profound impact on the mode of supervisory interaction, the positioning of supervisor and student roles, and the core logic of education itself. How to delineate reasonable boundaries for AI intervention, how to redefine supervisory functions to accommodate human-AI symbiosis,

and how to safeguard the fundamental value of fostering virtue through education within the supervisory relationship — these are the central questions that postgraduate education urgently needs to address, and they form the point of departure for this paper.

2. Challenges and Impacts of Artificial Intelligence on Postgraduate Supervisor-Student Relationships

2.1 The Deconstruction of Supervisory Interaction Structure: From Dyadic to Triadic Tension

The traditional postgraduate supervisor-student relationship has exhibited a dyadic interaction structure of “supervisor-student” in its configuration. Within this structure, the supervisor serves as the dominant figure in knowledge provision and academic evaluation, while the student occupies the position of knowledge recipient and subject of academic socialisation. The supervisor enjoys unquestioned epistemic authority, and the student’s academic development is heavily dependent upon the supervisor’s judgement and feedback. The stability of this dyadic structure derives precisely from this functional complementarity, in which the supervisor’s guidance and the student’s learning are realised through their interaction.

The intervention of generative artificial intelligence has weakened this interaction. AI is no longer a passive assistive tool; it demonstrates considerable capabilities in knowledge provision, text generation, methodological advice, and even academic feedback. The supervisor ceases to be the sole object of students’ academic consultation and the only channel for obtaining improvement suggestions. Supervisory interaction consequently transforms from a linear supervisor-

student relationship into a triadic interaction involving supervisor, AI, and student.

With the emergence of the triadic structure, students acquire information sources and feedback channels independent of their supervisors, while supervisors face the new challenge of AI-generated content infiltrating the supervisory process. More critically, AI as a technological entity alters the behavioural logic of both parties: students reduce the frequency and depth of their engagement with supervisors because of AI's accessibility, while supervisors, unable to ascertain the extent of AI's involvement, struggle to make accurate judgements about students' actual abilities. Lai et al. (2025) [1] characterise this phenomenon as GenAI embedding itself as an active mediator within the supervisory system, thereby altering the division of roles and allocation of responsibilities among the various elements within the system. The originally closed supervisory interaction loop thus undergoes fragmentation and reorganisation, shifting from direct human-to-human dialogue to AI-mediated indirect interaction. The dyadic structure, at the functional level, can no longer sustain its original operational logic.

2.2 The Marginalisation of the Supervisor's Role: Erosion of Epistemic Authority and Supervisory Value

The legitimacy of the supervisor's role rests upon epistemic authority and supervisory value. Epistemic authority derives from the knowledge and judgement accumulated through years of deep engagement within a specific discipline; supervisory value is embodied in the supervisor's irreplaceable role in guiding students' academic development. The rise of generative artificial intelligence is simultaneously undermining both pillars.

At the level of epistemic authority, AI's capacity for knowledge provision is eroding the supervisor's status as the disciplinary authority. In traditional supervisory relationships, supervisors, by virtue of their deep academic accumulation, exert significant influence on students at the level of disciplinary frontiers, research methodologies, and theoretical frameworks. However, the knowledge coverage and response speed of generative AI surpass the supervisor's capabilities in certain respects. When students can obtain structurally complete academic responses generated by AI within seconds, the irreplaceability of the supervisor's "epistemic authority" is called into question. Henriksen et al. (2025) [2] observe that the emergence of generative AI "fundamentally shifts how educational knowledge is created, shared, and validated."

At the level of supervisory value, AI's capabilities are increasingly encompassing the supervisor's traditional responsibilities, including literature guidance, writing feedback, methodological training, and academic socialisation, ranging from preliminary editing and summarisation to structured academic recommendations. Consequently, the focus of the supervisor's work must shift away from these routine instructional tasks toward the deeper academic judgment and critical thinking training that AI cannot perform. Some scholars characterise this transformation as a hybrid supervisory model, in which AI undertakes structural and procedural tasks while human supervisors concentrate on judgement-intensive academic

work [3]. Yet this model itself suggests that the supervisor's traditional functions have been substantially curtailed. When a significant portion of foundational work in the supervisory process can be accomplished by AI, the functional positioning of the supervisor within the supervisory relationship faces pressure for redefinition, and the distinctiveness and irreplaceability of the supervisor's guiding value correspondingly diminish.

2.3 Erosion of Student Agency: Intellectual Inertia and Academic Capacity Deterioration

Postgraduate education takes the cultivation of independent thinking and academic autonomy as its fundamental objective. At the core of postgraduate-level learning lies the training of students' abilities to identify problems, analyse issues, and generate knowledge. However, the intervention of generative artificial intelligence is eroding student agency in two dimensions—thinking patterns and academic capacity—plunging students into the dual predicament of intellectual inertia and capacity deterioration.

At the level of thinking patterns, the convenience of AI has induced a tendency toward cognitive offloading among postgraduate students. Cognitive offloading refers to the behavioural pattern in which individuals transfer cognitive tasks that should be performed by themselves to external tools, thereby reducing mental effort. The all-encompassing functionality of generative AI renders it an exceedingly attractive cognitive substitute in students' academic work. When AI can handle tasks such as literature review, idea generation, and text polishing, users no longer need to personally perform those cognitive activities that would otherwise exercise their thinking, a phenomenon that may be termed "cognitive surrender." A survey of 1,781 university students found complex associations between the use of GenAI tools and students' critical thinking and autonomous learning abilities, with researchers calling for further careful evaluation of the boundaries of personalised assistance provided by such tools [4]. This concern is equally present among postgraduate students. A survey encompassing 58 postgraduate students from 21 universities revealed that over half (58%) of respondents believed that reliance on AI might weaken critical thinking or foster dependency [5]. These two studies, drawing on different populations, mutually confirm that the potential erosion of thinking capacity by AI is not an isolated anxiety but a widely shared consensus. Although postgraduate students actively use AI, they also perceive the latent risk of their thinking abilities being undermined through such use; this ambivalence precisely reflects that cognitive offloading has occurred to a certain extent.

At the level of academic capacity, excessive use of AI may lead to substantive deterioration of postgraduate students' foundational academic abilities. The core of traditional postgraduate training lies in cultivating independent scientific thinking through cognitive struggle characterised by repeated trial and error, iterative revision, and persistent search for direction from failure. This process, though arduous, constitutes an indispensable pathway for academic growth. When AI can handle tasks such as literature review, methodological design, and draft composition, students bypass this developmental process. Critical reading,

reasoning, and writing skills resemble muscles: they weaken through disuse, a phenomenon that may be described as “intellectual atrophy.” [6] A more profound issue is that senior scholars trained before the proliferation of AI developed their writing and analytical skills through repeated manual practice, whereas contemporary postgraduate students, with unrestricted access to AI platforms, may never need to independently exercise these skills. This disparity in the pathways of academic capacity development between two generations of scholars implies a substantial risk of foundational academic capacity deterioration among contemporary postgraduate students.

2.4 The Dilution of Supervisory Emotional Bonds: From Human Warmth to Machine Efficiency

The supervisor-student relationship is not merely a relationship of academic guidance, but also one of emotional connection. Throughout the prolonged process of academic training, postgraduate students require not only the supervisor’s intellectual and methodological direction, but also emotional support, spiritual encouragement, and psychological reassurance. Supervisors often simultaneously assume the multiple roles of academic mentor, psychological supporter, and career development advisor. The intervention of generative artificial intelligence, however, is placing the emotional bonds between supervisor and student at risk of dilution.

The first manifestation of this dilution is the “displacement” of emotional support. Academic research is a highly individualised, complex, and frequently emotionally charged endeavour. Boyd and Harding (2025) found that GenAI plays a role in providing “surrogate emotional support” in academic research [7]. This means that when postgraduate students encounter setbacks, experience confusion, or feel the need to confide in someone during their research, they may turn to AI rather than their supervisor for emotional consolation. AI provides immediate, non-judgemental responses, and students need not fear exposing their uncertainties or vulnerabilities. This convenience renders AI a “safe” emotional outlet, yet simultaneously means that supervisors miss opportunities to forge emotional connections with students at the very moments when they are most in need of support.

The second manifestation is the “dehumanisation” of supervisory interaction. Effective postgraduate supervision fundamentally relies on trust, dialogue, guidance, and emotional support. When students become accustomed to obtaining immediate feedback and emotional responses from AI, their previously frequent, face-to-face interactions with supervisors may be compressed into transactional, functionally oriented exchanges. The supervisor ceases to be the student’s primary choice for emotional confiding, and the dialogue between supervisor and student degrades from human communication to mere information exchange. When supervisory interaction is reduced to standardised question-and-answer and feedback, the supervisory relationship degenerates from a human bond rich with warmth into a functional relationship oriented solely toward efficiency. The supervisor’s role is thus downgraded from “mentor” to “service provider”; the former is concerned with the holistic development of the person, while the latter is merely

responsible for delivering specific services.

A deeper problem lies in the often imperceptible nature of this emotional dilution. Students may not realize that they are using AI to replace the supervisor’s emotional function, and supervisors may similarly fail to notice that students are seeking emotional support from alternative sources. When neither party recognises that the emotional bond is being eroded, repair becomes impossible, and the emotional ties between supervisor and student prove difficult to restore. Students cease to open their hearts to supervisors, and supervisors find it increasingly difficult to enter the inner world of their students. The supervisory relationship gradually degenerates from a human connection into a mechanism of machine efficiency, and the most essential affective dimension of its educative function is thus eroded in silence.

3. Opportunities for Ai-Empowered Reconstruction of Postgraduate Supervisor - Student Relationships

3.1 Enhancement of Supervisory Efficiency: From Experience-Driven to Data-Driven

The introduction of generative artificial intelligence provides a new technological pathway for enhancing the efficiency of postgraduate supervision. Traditional supervisory decisions have relied on supervisors’ personal experience and subjective judgement, whereas AI, through the large-scale processing of academic data and student behavioural data, shifts supervisory activities from experience-dominated to data-supported approaches. This efficiency enhancement is primarily manifested in three dimensions.

First, the automation of routine tasks liberates supervisors’ time for guidance. Conventional tasks such as preliminary literature summarization, correction of writing grammar and formatting, standardization of references, and initial organization of research data previously required substantial time investment from supervisors to complete item by item. Although necessary, these tasks involve limited technical complexity and are highly repetitive. Generative AI can accomplish the above tasks at speeds far exceeding human capacity, with output quality that is generally acceptable. Supervisors are thus relieved of trivial matters and can reallocate the saved time to higher-order supervisory activities such as academic argumentation, research design, and critical dialogue.

Second, data-driven monitoring of student progress enhances the precision of supervisory intervention. Under the traditional model, supervisors’ grasp of students’ academic status relied on periodic meetings and stage-based reviews of deliverables, with information acquisition characterised by delay and incompleteness. AI-driven learning analytics can continuously track behavioural indicators such as students’ literature reading progress, writing output frequency, and data analysis completion rates, generating visualizable progress reports. Supervisors are thus able to monitor each student’s research status in real time and intervene when necessary, with both the precision and timeliness of intervention

substantially enhanced.

Third, the systematisation of personalised support expands the coverage of supervisory guidance. Each student possesses a distinct knowledge structure, research capacity, and learning rhythm, and ideal guidance should be tailored to individual differences. However, supervisors' time and energy are limited, making it difficult to provide in-depth personalised guidance to multiple students simultaneously. AI can analyse different students' academic output characteristics and competency gaps, recommending differentiated literature resources, methodological tools, and training tasks accordingly, and generating initial personalised learning plans for each student. Supervisors can then review, adjust, and deepen the foundation plans provided by AI, thereby achieving differentiated support for more students at a lower marginal cost.

3.2 Enhancement of Supervisory Quality: From Generalized Supervision to Precision Empowerment

The enhancement of postgraduate supervision quality is also reflected in the precision and depth of guidance. Under the traditional model, supervisors' guidance to students has been based largely on general academic experience and impressions formed from students' stage-based deliverables, with instructional content carrying a strong generality that makes it difficult to respond precisely to each student's specific needs and real-time status. The intervention of generative artificial intelligence enables supervisors to achieve precise engagement across various dimensions of students' academic training, propelling supervision from generalized approaches toward precision empowerment.

First, the personalization of instructional content enables more adequate responses to students' differentiated needs. Each student possesses a distinct knowledge foundation, research direction, and competency gaps. Under the traditional supervisory model, it is difficult for supervisors to tailor differentiated training plans for each student. Supervisors typically can only form general judgements based on limited interactions and submitted work, offering broadly applicable recommendations. Generative AI can conduct continuous diagnostic analysis of students' academic texts, research data, and knowledge structures, precisely identifying individual shortcomings in areas such as argumentative logic, methodological application, and literature coverage, and accordingly generating targeted improvement suggestions and learning resource recommendations. Supervisors, building upon the diagnostic foundation provided by AI, can then review and deepen these insights, enabling supervisory recommendations to precisely correspond to each student's actual needs.

Second, the instantiation of feedback mechanisms enables students to obtain sustained academic support. Under traditional supervision, students often wait days or even weeks for feedback on draft manuscripts and research designs, with long feedback cycles and limited frequency. This delayed feedback model constrains students' iterative efficiency in that problems in writing may remain undiscovered until entire chapters are completed. Generative AI can provide real-time feedback during students' work processes, whether the

literature review framework is complete, whether the data analysis methods are appropriate, whether the chain of argumentation is coherent, and these issues can be identified immediately with AI assistance. By the time students submit their work to supervisors, they already possess a preliminarily diagnosed version, allowing supervisors' feedback to bypass foundational issues and focus directly on deeper academic judgement.

Third, the expansion of interdisciplinary knowledge enables both supervisors and students to transcend disciplinary boundaries. Under the traditional supervisory model, the supervisor's knowledge boundaries constitute the upper limit of students' academic horizons. When research involves areas unfamiliar to the supervisor, whether frontier methodologies or interdisciplinary theoretical resources, the supervisor struggles to provide effective guidance. Generative AI can rapidly map interdisciplinary research trajectories, introduce key concepts and core literature in relevant fields, and provide references for methodological choices. Supervisors can quickly enter relatively unfamiliar knowledge domains with AI assistance, and students can independently access interdisciplinary resources beyond their supervisors' guidance. Supervisory activities are no longer constrained by the supervisor's personal knowledge boundaries, and the possibilities for interdisciplinary research are thus expanded.

Fourth, the refinement of academic writing training enables sustained improvement in students' expressive capabilities. Academic writing is a core component of postgraduate training and one of the most time-consuming aspects of traditional supervision. Supervisors often need to repeatedly annotate logical structure, academic tone, argumentative approach, and language expression within the same draft. Generative AI can provide fine-grained diagnostic feedback and revision suggestions across various dimensions of writing: from syntactic expression to paragraph cohesion, from the formality of academic tone to the clarity of argumentative structure. Students gain concrete directions for improvement with each writing iteration, rather than receiving vague feedback such as "not well written." More importantly, AI can demonstrate before-and-after comparisons, enabling students to understand why particular revisions are superior, thereby gradually internalising the norms of academic writing through sustained training.

3.3 Paradigm Shift in Supervisory Relationships: From Master-Apprentice Model to Human-AI Symbiosis

Postgraduate supervision is not merely a teaching activity, but also a deeply cognitive and social relationship. The core paradigm of the traditional supervisory relationship can be characterised as the master-apprentice model, in which postgraduate students, through prolonged engagement with their supervisors and close observation, emulation, and dialogue, progressively acquire the thinking modes, working habits, and professional ethics of academic research. This paradigm exhibits three fundamental features: the supervisor serves as a model of knowledge and methodology, with students learning through imitation and repeated practice; the supervisory process relies on face-to-face, long-term, and immersive interaction; and the cultivation of academic character is inseparable from the transmission of knowledge

and skills. The deep intervention of generative artificial intelligence is challenging the foundations of this paradigm, propelling the supervisory relationship toward a new paradigm of “human-AI symbiosis.”

The first dimension of the human-AI symbiosis paradigm is the epistemological transformation. In the traditional supervisory relationship, the supervisor constitutes the highest authority in academic judgement, and the student’s task is to learn the supervisor’s standards of judgement and gradually internalise them as their own capabilities. With generative AI entering academic activities as an epistemic partner, academic judgement is no longer a unidirectional output from the supervisor. Students can form preliminary judgements through dialogue with AI, then discuss and refine them with their supervisors; supervisors can also draw upon the analytical frameworks provided by AI to reflect upon and adjust their own judgements. The formation of academic understanding progressively transforms into a collaborative model of “student exploration-AI assistance-supervisor validation.”

The second dimension of the human-AI symbiosis paradigm is the transformation of interaction patterns. In the traditional master-apprentice model, supervisory interaction is centred on the supervisor: the supervisor arranges meetings, dominates discussions, and makes judgements. Students occupy a relatively passive position within this structure. When AI becomes an academic partner that students can access and engage with at any time, students acquire information sources and cognitive support independent of their supervisors. They can complete preliminary literature reviews and idea construction before meeting with supervisors, bringing more developed questions and proposals to supervisory discussions. Supervisory interaction thus shifts from a model of “supervisor questions-student answers” to a model of “student exploration-collaborative discussion.” Students transform from passive recipients into more active participants, while supervisors evolve from dominant figures into collaborators and gatekeepers.

The third dimension of the human-AI symbiosis paradigm is the transformation of relational structure. The traditional master-apprentice model constitutes a direct human-to-human relationship, in which emotional connection and academic transmission operate along the same track. With AI intervention, the relational structure evolves into a mediated form of human-AI-human. The risk of this transformation lies in the dilution of emotional bonds, yet it also brings new possibilities: students can engage in free preliminary exploration with AI in an environment free from the fear of judgment, and after gaining confidence, enter into deeper dialogue with supervisors. AI thus serves as a “buffer” between supervisor and student, lowering the psychological threshold for students to expose their inadequacies to supervisors.

3.4 Reconfiguration of the Supervisory Community: From Individual Mentoring to an Ecosystem

The basic unit of the traditional supervisory relationship is the

dyadic combination of supervisor and student, with supervisory activities revolving around this core relationship. The supervisor constitutes the student’s primary, and sometimes sole, source of academic support. While this individual mentoring model ensures relational closeness, it also carries inherent limitations: the student’s academic support network is over-reliant on the supervisor, and the supervisor’s knowledge boundaries, time, energy, and supervisory style together form the ceiling of the student’s academic development. The intervention of generative artificial intelligence, however, is propelling the supervisory relationship from one-to-one individual mentoring toward an open ecosystem involving multiple participants.

The first manifestation is the diversification of knowledge sources. Generative AI provides knowledge channels independent of the supervisor. Students can access interdisciplinary theoretical resources, information on frontier methodologies, and literature from diverse academic perspectives through AI, no longer relying entirely on the supervisor’s knowledge reservoir. At the same time, AI can help supervisors rapidly enter relatively unfamiliar knowledge domains, enabling them to transcend their own disciplinary boundaries. When both supervisors and students can obtain knowledge increments from AI, supervisory dialogue is no longer confined to the limited intersection of their existing knowledge, but is instead built upon their respectively expanded intellectual territories, thereby broadening both the depth and scope of their exchanges.

The second manifestation is the opening of interactive networks. The intervention of generative AI opens the academic research network in multiple directions. AI expands students’ direct access to external academic resources: interdisciplinary literature, frontier research developments, and discussions within the international academic community. These resources previously required supervisor mediation to enter students’ field of vision, but students can now access them independently and bring them into dialogue with supervisors. The direction of interaction likewise shifts to two-way or even multi-way exchanges between students and supervisors following collaborative exploration with AI. The boundaries of the supervisory community are thus opened, evolving from a single, confined system into an open system that continuously exchanges information with the wider world of knowledge.

The third manifestation is the equalisation of support resource distribution. Variations in supervisors’ guiding styles and levels of commitment have resulted in uneven academic support for students. Generative AI promotes the equalisation of support resources at two levels. First, the foundational academic support provided by AI, including literature review, writing feedback, and methodological advice, is accessible to all users, and the quality of service does not fluctuate significantly due to individual differences among supervisors. Second, the tasks undertaken by AI free up supervisors’ time, enabling those who were previously under-invested to redirect more energy toward high-quality guidance, thereby narrowing the gap in supervisory effectiveness at the foundational level to a certain extent.

4. Governance Paths for Postgraduate Supervisor - Student Relationships in the Age of Artificial Intelligence

4.1 Technological Enablement: Delineating Boundaries for AI Intervention and Standardising Rules for Human-AI Collaborative Use

Technology serves as the foundational vehicle for achieving human-AI collaborative supervision. Reasonably defining the scope of AI application can alleviate the interactive tensions arising from the triadic supervisory structure, and fundamentally prevent issues such as student intellectual inertia, the concealment of academic misconduct, and the weakening of supervisor-student trust.

First, it is essential to establish tiered boundaries for AI usage scenarios and clearly delineate the division of labour between humans and machines. A differentiated control mechanism should be implemented according to the nature of research activities: routine tasks such as literature organisation, formatting, and grammar proofreading may be delegated to AI, thereby freeing both supervisors and students to devote their energy to higher-order guidance and critical inquiry. For tasks including research frameworks and writing approaches, AI may only provide reference materials, which students must subsequently verify point by point and substantially revise before use. Critical junctures that embody academic originality, such as topic design, core argumentation, close reading of primary sources, and the generation of innovative ideas, must be strictly off-limits to AI, precluding any substitution of independent thinking by technological tools. Additional constraints should be imposed for humanities disciplines, with strict controls on AI-generated foreign-language citations and theoretical interpretations, so as to prevent AI hallucinations from causing fabricated references or distorted theoretical readings.

Second, a full-process traceability mechanism for AI-generated outputs should be established to ensure that research processes remain verifiable. A written declaration system for AI use should be implemented, requiring students to truthfully document the scenarios, tool types, and revision processes involving AI in their research proposals, progress reports, and theses. Concurrently, AI dialogue logs and multiple draft versions should be retained to serve as critical evidence for supervisors in evaluating students' genuine research capabilities. AIGC text detection tools should be embedded within institutional systems to automatically identify large-scale AI-generated writing and evasive paraphrasing, thereby reducing supervisors' verification costs and shifting supervisory evaluation from a results-oriented to a process-oriented approach.

Finally, institutionally approved, compliant AI research tools should be developed to mitigate the risks associated with external platforms. Universities should create dedicated, lightweight AI-assisted systems that offer only compliant functions such as literature retrieval and corpus analysis, while blocking high-risk modules such as one-click full-text generation, with the systems automatically retaining records of supervisor-student usage. These systems should incorporate a supervisory collaboration portal, enabling

supervisors to view students' AI usage trajectories in real time and intervene proactively when excessive tool reliance or uncritical acceptance of erroneous content is detected. In addition, built-in sensitive text filtering and algorithmic bias alert mechanisms should be implemented to provide automatic warnings for one-sided narratives or inaccurate historical materials, thereby balancing technological efficiency with the guidance of academic values and fostering a constructive triadic "supervisor-AI-student" collaborative mentoring model.

4.2 Institutional Adaptation: Refining Supporting Regulations and Reforming Supervisor Evaluation Systems

Institutions constitute the rigid safeguard for regulating supervisory behaviour and reconciling triadic supervisory tensions in the age of AI. Existing postgraduate management and supervisor evaluation systems are largely designed for the traditional dyadic master-apprentice model and are ill-equipped to address emerging issues such as ethical ambiguity surrounding AI, the transformation of supervisory roles, and the alienation of supervisory interaction. Institutional renewal is therefore required across three dimensions: usage standards, evaluation mechanisms, and conflict early warning systems.

First, it is necessary to promulgate institutional regulations on the use of AI academic tools, clearly defining behavioural red lines and corresponding disciplinary standards. Detailed requirements for AI use at each stage of postgraduate research should be established, with explicit classification of the levels of academic misconduct corresponding to AI-generated writing, fabricated references, and undisclosed tool usage. Differentiated treatment should be stipulated between unintentional reliance on erroneous AI-generated content and deliberate fabrication. Concurrently, supervisors' managerial responsibilities in verifying AI-generated outputs should be clearly defined, and AI output verification should be incorporated into their routine supervisory duties as a mandatory task, ensuring that both supervisors and students have clear rules to follow.

Second, the supervisor evaluation system should be restructured to align with the new educational positioning in the AI era. Quantitative metrics such as publication output and project counts should be de-emphasised, while greater weight should be assigned to qualitative indicators, including the cultivation of virtue, the development of critical thinking, and the guidance of academic character. Supervisors' performance in identifying students' AI dependency, guiding students' critical use of tools, and conducting in-depth offline academic exchanges should be incorporated into annual assessments. This restructuring should redirect supervisors away from the outdated identity of "knowledge transmitters" toward roles as intellectual guides and shapers of academic character, thereby alleviating the identity anxiety induced by the impact of AI.

Third, a dynamic early warning and communication mediation system for supervisor-student relationships should be established. Drawing on institutional AI usage records and academic performance data, a risk monitoring system should

be developed to identify supervisor-student pairs exhibiting long-term excessive AI reliance or significantly reduced communication frequency. In such cases, student affairs professionals and graduate school administrators should intervene proactively through counselling and dialogue. The communication and mediation mechanism for supervisor-student conflicts should also be improved, with regular faculty-student forums convened to gather information on practical challenges encountered in human-AI collaborative supervision, and to dynamically refine management regulations. This closed-loop governance system, encompassing regulatory constraints, process monitoring, and timely intervention, provides institutional support for the stable operation of triadic supervisory relationships.

4.3 Capacity Building: Enhancing AI Digital Literacy for Both Supervisors and Postgraduate Students

The sound functioning of human-AI collaborative supervision hinges critically upon the AI digital literacy of both supervisors and students. Deficiencies in either party's capabilities will amplify the various contradictions inherent in the triadic supervisory structure. It is therefore necessary to construct a tiered cultivation system targeting both supervisors and students, achieving simultaneous capability enhancement in both directions.

For supervisors, regular thematic training programmes should be offered, with emphasis on two core competency areas. The first is the capacity to critically evaluate AI-generated outputs: training should cover text detection methods for various generative tools, instruction on how to identify AI-fabricated citations, logical fallacies, and inherent algorithmic biases, thereby addressing supervisors' difficulty in distinguishing students' original contributions from AI-generated text. The second is the capacity for novel supervisory guidance: training should provide guidance strategies for the triadic interaction model, directing supervisors to reduce routine knowledge transmission and focus instead on higher-order educational activities such as the cultivation of research thinking and academic character, while alleviating the role anxiety induced by the erosion of their epistemic authority. Regular supervisor forums should be organised to share practical experiences in human-AI collaborative supervision, establishing a peer-learning community for mutual support.

For postgraduate students, systematic education should be delivered through orientation courses on academic literacy and thematic seminars. On the one hand, training should cover compliant usage standards for AI tools, clarify the boundaries of human-AI division of labour at each research stage, explain common deficiencies such as AI hallucinations and narrative biases, and cultivate students' habit of verifying AI-generated content point by point, thereby breaking their blind reliance on AI. On the other hand, training in independent critical thinking and originality should be strengthened through practical exercises such as close reading of primary sources and textual analysis, consolidating foundational research skills in the humanities and subjectively reducing the likelihood of academic misconduct.

In addition, collaborative practical activities involving both supervisors and students should be organised, with joint

exercises centred on academic writing and literature research. Through the shared use of AI tools, supervisors and students can develop a common understanding, refine their modes of communication, and bridge the gaps in their perceptions and usage habits regarding AI. Such initiatives will cultivate a community of teachers and learners equipped with both digital discernment and academic steadiness, thereby laying a solid capability foundation for the human-AI symbiotic supervisory relationship.

4.4 Value Shaping: Reconstructing a Warm and Humanistic Supervisory Community

Technical regulations and institutional frameworks can only establish behavioural baselines. To counteract the erosion of instrumental rationality brought about by AI and to restore the diluted emotional bonds between supervisors and students, it is also necessary to rely on value-based cultural reconstruction to reshape the triadic supervisory relationship and return to the fundamental objective of postgraduate education: fostering virtue through education.

First, the scholarly culture of the traditional master-apprentice model should be preserved, and in-depth offline interaction scenarios free from AI should be provided. Regular offline activities such as reading groups, academic seminars, and fieldwork should be organised to reduce the model in which supervisors and students communicate solely through online tools for fragmented exchanges. In face-to-face interactions, supervisors can provide value guidance, impart scholarly character, and convey research ideals and academic responsibility—thereby compensating for the emotional nourishment and spiritual leadership that AI cannot offer, and rebuilding the sincere and trusting humanistic bonds between supervisors and students.

Second, the core responsibility of supervisors in value leadership should be clarified, while the anxiety over diminished authority induced by technological tools should be alleviated. Supervisors should be guided to move beyond the singular role of knowledge transmission, integrating academic integrity, critical spirit, and commitment to the nation and society into their daily guidance. They should proactively explain the limitations and potential risks of AI to students, helping them establish a correct research perspective in which AI serves their purposes. Through sustained communication, an inclusive and egalitarian collaborative supervisory atmosphere should be fostered.

Finally, a culture of community characterised by symbiosis and mutual advancement among supervisors and students should be cultivated. Through supervisory salons and teacher-student growth-sharing sessions, universities should guide both parties to engage in joint discussions on the boundaries and scales of human-AI collaboration, thereby dissolving the estrangement caused by divergent views on AI use. It is essential to advocate consensus on the pursuit of truth and authentic engagement, balancing research efficiency with educational warmth. In this manner, the triadic supervisory structure ceases to be a source of adversarial tension between human and machine, but rather becomes an integrated organic whole in which technology, supervisors, and students collaborate in education, providing spiritual support for the

long-term and stable development of supervisory relationships in the age of AI.

5. Conclusion

The transformation of postgraduate supervisory relationships brought about by generative AI is, in essence, a continuous contestation between instrumental rationality and humanistic educational values. The triadic interaction configuration cannot recede to the traditional pure dyadic master-apprentice model. Neither outright rejection nor laissez-faire permissiveness regarding AI use can adequately meet the demands of high-level talent cultivation in the new era. Technology itself is neither good nor evil; the ultimate shape of human-AI collaborative supervisory relationships depends on the regulatory boundaries and value consensus jointly established by universities, supervisors, and students. The core of the supervisory relationship has never been the efficiency of knowledge transmission, but rather the cultivation of academic character, the nurturing of critical thinking, and the guidance offered through shared intellectual and spiritual companionship. These humanistic essences cannot be replaced by algorithms. Delineating boundaries for AI use, refining supporting institutional frameworks, cultivating the digital discernment of both supervisors and students, and reconstructing an offline scholarly community: these multiple governance measures converge on a single ultimate objective: safeguarding the fundamental educational mission of fostering virtue through education. Looking ahead, postgraduate education cannot passively respond to technological impacts; it must actively explore new human-AI symbiotic paradigms for education. Only by subordinating technological application to humanistic values, and by striking a dynamic balance between tool-enabled empowerment and value-based commitment, can AI serve as an auxiliary vehicle for mutual growth in supervision and learning rather than a barrier that severs the supervisor-student bond. Only then can a truly new supervisory community be constructed, one that embodies both efficiency and warmth.

Acknowledgments

This research was supported by the 2024 Annual Northwestern Polytechnical University Degree and Postgraduate Education Research Fund Project (2024JM014) titled "Research on the Cultivation Paths of Postgraduate Students' Awareness of Social Responsibility from the Perspective of the Holistic Approach to National Security".

References

- [1] Lai, S., Liu, S., Dai, Y., Lim, C. P., & Liu, A. (2025). The impacts and tensions of generative AI on doctoral students' supervisory and peer dynamics: An activity theory analysis. *Australasian Journal of Educational Technology*, 41(5), 1-17. <https://doi.org/10.14742/ajet.9916>
- [2] Henriksen, D., Mishra, P., Woo, L., & Oster, N. (2025). The education doctorate in the context of generative artificial intelligence: Epistemic shifts and challenges to practical wisdom. *Impacting Education: Journal on Transforming Professional Practice*, 10(1), 73-79. <https://doi.org/10.5195/ie.2025.485>
- [3] Omodan, B. I. (2025). Redefining the role of supervisors in the era of artificial intelligence: Implications for hybrid postgraduate research governance. *Cogent Education*, 12(1), 2536534. <https://doi.org/10.1080/2331186X.2025.2536534>
- [4] Qi, J., Xu, Y. R., Liu, J. A., & Xue, K. (2025). The impact of generative artificial intelligence tools on college students' critical thinking and autonomous learning ability. *Frontiers of Education in China*, 20(4), 398-412. <https://doi.org/10.3868/s110-020-025-0021-0>
- [5] Governors State University. (2026, April 10). Graduate students' perceptions of generative AI's impact on critical thinking development [Poster presentation]. Research Days 2026, University Library, Governors State University, University Park, IL, United States. https://opus.govst.edu/research_day/2026/fri/3/
- [6] Sobo, E. J., et al. (2025, December 17). Faculty are first responders in averting an epidemic of intellectual atrophy. *Inside Higher Ed*. <https://www.insidehighered.com/opinion/views/2025/12/17/faculty-are-key-preventing-ai-overreliance-opinion>
- [7] Boyd, P., & Harding, D. (2025). Generative AI: Reconfiguring supervision and doctoral research. *Buildings & Cities*, 6(1), 294-309. <https://doi.org/10.5334/bc.560>