Review

Application Boundaries and Ethical Considerations of Generative AI-Assisted Teaching Design for Medicinal Chemistry

Zhanghong Fu[1,*]

[1] Hainan Vocational University of Science and Technology, Haikou, China
* Correspondence: Zhanghong Fu, Hainan Vocational University of Science and Technology, Haikou, China

Abstract: The explosive development of generative artificial intelligence has triggered profound transformations in higher pharmaceutical education. As a core course that heavily relies on spatial imagination, logical reasoning, and interdisciplinary knowledge integration, Medicinal Chemistry benefits greatly from generative AI in scenarios including molecular generation, retrosynthetic deduction, and personalized tutoring. Nevertheless, rapid technological advancement brings severe challenges such as knowledge hallucinations, cognitive outsourcing, and academic integrity risks. Based on the disciplinary characteristics of Medicinal Chemistry, this paper systematically sorts out the enabling scenarios of generative AI in instructional design. It analyzes the application boundaries concerning scientific validity, practical implementation, and affective interaction. Furthermore, it conducts a comprehensive ethical deliberation from four critical dimensions: academic integrity, algorithmic bias, data privacy, and accountability. To address these challenges, the study proposes a framework for evaluating AI-generated chemical structures and molecular pathways to ensure scientific accuracy. On this basis, a new human-machine collaborative and human-oriented teaching paradigm is constructed. This paradigm emphasizes the cultivation of students' critical thinking and ethical awareness alongside their technical proficiency. Ultimately, this research aims to provide robust theoretical references and practical pathways for curriculum teaching innovation and ethical governance of pharmaceutical majors in the intelligent era, ensuring that technological integration enhances rather than compromises educational quality and student development.

Keywords: generative ai; medicinal chemistry; application boundaries; ethics; teaching innovation; higher education

Received: 28 January 2026
Revised: 14 March 2026
Accepted: 24 March 2026
Published: 31 March 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Medicinal Chemistry constitutes a foundational and integrative discipline within pharmacy and pharmaceutical engineering curricula, characterized by its rigorous theoretical framework, strong emphasis on experimental validation, and intrinsic connections to organic chemistry, biochemistry, pharmacology, and computational biology. The course systematically explores the molecular architecture of bioactive compounds, analyzes their thermodynamic and kinetic behavior in physiological environments, investigates structure–activity and structure–property relationships through quantitative methods, examines metabolic fate—including phase I and phase II biotransformations—and introduces principles underlying rational drug design and optimization strategies [1]. Mastery of this subject demands that students cultivate advanced spatial reasoning skills to interpret conformational dynamics, chiral recognition, and non-covalent binding interactions; develop systematic analytical habits to correlate structural modifications with functional outcomes; and integrate multi-scale knowledge—from atomic orbitals to whole-organism pharmacokinetics. Conventional pedagogical approaches predominantly rely on static two-dimensional representations—

such as skeletal formulas, resonance structures, and hand-drawn reaction mechanisms—delivered via instructor-led expository lectures. While effective for conveying foundational concepts, these methods often prove insufficient for fostering intuitive comprehension of transient, three-dimensional molecular docking events, allosteric modulation, or the dynamic solvation effects that critically influence ligand–target affinity and selectivity.

The rapid advancement of generative artificial intelligence technologies—including large language models capable of domain-specific reasoning and deep generative models trained on vast chemical databases—has introduced transformative capabilities into pharmaceutical education. These systems can now propose novel scaffold-hopping candidates consistent with target pharmacophores, estimate logP, pKa, and membrane permeability with increasing accuracy, generate plausible synthetic pathways using retrosynthetic analysis algorithms, and even simulate protein–ligand binding poses under varying physiological conditions. Such functionalities significantly enrich instructional resources, enable personalized learning trajectories, support inquiry-based laboratory design, and facilitate real-time exploration of structure–function hypotheses previously constrained by time or instrumentation limitations [2]. Nevertheless, the integration of these tools into academic settings necessitates careful scrutiny: model outputs may contain undetected factual inaccuracies or chemically implausible suggestions due to training data biases or incomplete physical constraints; decision-making processes often lack transparency, hindering pedagogical traceability and student critical evaluation; and the handling of sensitive academic or institutional data raises legitimate concerns regarding confidentiality, intellectual property stewardship, and long-term archival integrity. Consequently, establishing clear methodological guidelines, implementing robust human-in-the-loop verification protocols, and embedding AI literacy within curriculum frameworks have emerged as essential priorities for sustaining academic rigor, safeguarding educational equity, and ensuring responsible innovation in pharmaceutical science education.

2. Enabling Scenarios of Generative AI in Medicinal Chemistry Teaching Design

Generative artificial intelligence, underpinned by sophisticated natural-language processing architectures and multimodal generative models, is fundamentally transforming pedagogical frameworks in medicinal chemistry education. Its integration enables dynamic lesson personalization, real-time generation of molecular visualization aids, automated synthesis pathway simulations, and context-aware feedback on student-generated chemical reasoning. Instructors can rapidly prototype interactive case studies involving drug-target interactions, ADMET prediction exercises, and structure-activity relationship explorations without requiring specialized programming expertise. Furthermore, generative AI supports scaffolding of complex concepts—such as bioisosteric replacement or metabolic soft spot identification—through adaptive question generation and stepwise explanatory narratives. This technological shift fosters deeper conceptual engagement, enhances accessibility for diverse learning styles, and promotes active knowledge construction over passive content transmission, thereby aligning with contemporary evidence-informed principles of science education [3].

2.1. Intelligent Deduction of Molecular Generation and Retrosynthetic Routes

In instructional modules focused on novel therapeutic agent design, generative artificial intelligence—particularly deep learning architectures such as Transformer-based molecular generation models—serves as a pedagogical catalyst for advancing students' understanding of structure–activity relationships and lead-compound optimization. Students begin by specifying target biological profiles, desired physicochemical properties, or pharmacophore constraints; the system then iteratively proposes chemically valid, synthetically accessible, and structurally diverse candidate molecules that align with those criteria. This approach fosters hypothesis-driven exploration rather than passive reception of pre-determined chemical entities. Within pharmaceutical synthesis

instruction, AI-powered retrosynthetic analysis platforms leverage advanced search strategies—including Monte Carlo tree search—to decompose complex target molecules into progressively simpler precursors, generating multiple logically coherent disconnection pathways in real time. Each proposed route undergoes systematic evaluation across quantitative metrics such as atom economy, step efficiency, functional group compatibility, thermal and catalytic requirements, and environmental impact indicators like E-factor. Consequently, learners transition from static memorization of canonical reaction sequences to dynamic, evidence-informed decision-making grounded in multi-parameter trade-off analysis, thereby cultivating critical synthetic intuition and process-oriented chemical reasoning essential for modern drug development practice [4].

2.2. Structure-Activity Relationship (SAR) Analysis and Target-receptor Docking Simulation

Structure–activity relationships represent a foundational pillar of medicinal chemistry education, enabling learners to systematically correlate molecular structural features with biological effects. Traditional pedagogical approaches typically rely on static tabular summaries or two-dimensional schematic comparisons, which often fail to convey the spatial complexity and dynamic nature of molecular recognition processes. To address this limitation, modern instructional design increasingly incorporates computational methodologies—particularly generative artificial intelligence integrated with molecular dynamics simulation frameworks—to construct immersive, interactive learning environments. Within such modules, students can actively manipulate chemical structures—for instance, by substituting hydrogen atoms with fluorine, extending alkyl chains, or modifying aromatic ring systems—and immediately observe corresponding changes in three-dimensional conformational behavior and binding orientation. The system dynamically generates and visualizes docking poses within receptor binding sites, explicitly highlighting key non-covalent interaction motifs including hydrogen bonding networks, van der Waals contacts, hydrophobic surface complementarity, and aromatic stacking arrangements. This real-time, visually grounded feedback loop supports deeper conceptual integration, significantly lowering cognitive load associated with abstracting microscale intermolecular forces. Furthermore, such tools foster experiential learning by transforming passive reception into active hypothesis testing, thereby strengthening analytical reasoning and structure-based design intuition essential for pharmaceutical sciences.

2.3. Personalized Learning-path Planning and Intelligent Q&A

Medicinal Chemistry constitutes a highly interdisciplinary domain characterized by intricate conceptual hierarchies, dynamic structure–activity relationships, and multifaceted biochemical mechanisms—such as enzyme inhibition kinetics, metabolic biotransformation cascades, and pharmacokinetic compartmental modeling. Learners exhibit substantial heterogeneity in prior knowledge, cognitive processing styles, and domain-specific reasoning abilities, necessitating pedagogical approaches that dynamically adapt to individual progression trajectories. Intelligent teaching assistants powered by large language models analyze longitudinal interaction data—including quiz responses, concept-mapping submissions, and reflective journal entries—to detect latent misconceptions, assess depth of conceptual integration, and quantify mastery thresholds across hierarchical knowledge units. Based on this diagnostic profiling, the system autonomously constructs adaptive learning pathways comprising sequenced micro-modules, spaced-repetition review cycles, and contextually scaffolded problem sets aligned with Bloom's taxonomy [5]. During asynchronous tutoring interactions, the assistant employs Socratic dialogue protocols grounded in cognitive apprenticeship theory: it poses calibrated, open-ended prompts that elicit metacognitive reflection, encourages hypothesis generation about reaction intermediates or binding conformations, and guides learners through iterative refinement of mechanistic reasoning—thereby fostering durable, transferable expertise rather than rote memorization of isolated facts.

3. Application Boundaries of Generative AI-aided Teaching

Despite its considerable potential to enhance instructional efficiency, automate routine tasks, and support personalized learning pathways, generative AI must be deployed within carefully defined pedagogical boundaries—particularly in foundational disciplines such as Medicinal Chemistry. The core purpose of education remains the holistic cultivation of human character, critical thinking, ethical reasoning, scientific integrity, and professional competence—not merely the transmission or replication of factual knowledge. In this context, AI tools should function strictly as pedagogical aids that augment, rather than replace, instructor-guided inquiry, student-centered experimentation, and reflective problem-solving [6]. Overreliance on AI-generated explanations risks undermining students' development of deep conceptual understanding, mechanistic reasoning, and the ability to evaluate chemical validity, stereochemical nuance, or pharmacokinetic plausibility. Furthermore, unchecked use may erode academic rigor, diminish opportunities for formative feedback, and obscure the epistemic responsibilities inherent in scientific practice. Therefore, clear institutional and curricular guardrails are essential to ensure alignment with educational objectives, disciplinary standards, and long-term learner development goals.

3.1. Boundary of Scientific Validity Triggered by Knowledge Hallucinations

Generative artificial intelligence systems operate fundamentally through statistical pattern recognition and probabilistic token prediction, rather than through domain-specific causal reasoning grounded in first principles of chemistry or pharmacology. This architectural limitation becomes especially pronounced when modeling intricate pharmaceutical-chemistry phenomena—such as reaction thermodynamics, stereoelectronic effects in drug–target binding, or multi-step enzymatic metabolic pathways. In such contexts, the models may generate outputs that appear chemically coherent at surface level yet violate fundamental constraints: for example, proposing molecular structures with impossible bond angles, non-integer valences, or ring systems exhibiting extreme Baeyer strain incompatible with known conformational stability; misrepresenting the kinetic and allosteric behavior of cytochrome P450 isoforms; or conflating competitive inhibition with time-dependent inactivation mechanisms. These knowledge hallucinations pose significant pedagogical risks when integrated into undergraduate or graduate curricula without rigorous expert oversight. Unchecked adoption could entrench conceptual misconceptions about molecular design logic, pharmacokinetic principles, or biochemical regulation. Therefore, subject-matter experts—particularly instructors and curriculum designers—must maintain authoritative responsibility for validating all AI-generated scientific content prior to instructional use, ensuring alignment with established chemical theory, empirical evidence, and regulatory science standards [7].

3.2. Cognitive Boundary of Spatial Stereochemical Perception

Drug bioactivity is fundamentally governed by three-dimensional molecular architecture, with particular dependence on stereochemical specificity—including absolute and relative chirality—as well as the thermally accessible conformational ensembles that molecules adopt under physiological conditions. While contemporary generative artificial intelligence models can produce visually compelling static 3D molecular renderings, these representations frequently omit critical biophysical dimensions: time-resolved conformational sampling, protein-ligand induced-fit adaptation, microenvironmental polarity gradients, and explicit or implicit solvation effects that modulate binding affinity and kinetic selectivity. Such omissions risk fostering a superficial understanding of molecular behavior, wherein students interpret structural images as definitive snapshots rather than probabilistic, context-dependent states [8, 9]. Overreliance on such static visualizations may inadvertently erode foundational spatial reasoning competencies—such as mental rotation, conformational interpolation, and topological inference—thereby impeding the development of robust mental models required for rational drug design. Consequently, pedagogical strategies must deliberately scaffold dynamic spatial cognition through interactive simulations, multi-conformer

ensemble analysis, and guided comparison of computational predictions against experimental structural biology data, ensuring learners remain active interpreters rather than passive consumers of molecular imagery.

3.3. Disconnection Boundary between In-silico Dry-lab Simulation and In-vitro/in-vivo Wet-lab Experiments

Medicinal chemistry is fundamentally an experimental discipline grounded in empirical validation, requiring rigorous integration of theoretical modeling with hands-on laboratory execution. Computer-aided drug design (CADD) constitutes an in-silico dry-lab methodology that relies on computational algorithms, molecular docking simulations, quantum mechanical calculations, and machine learning–driven property prediction to prioritize synthetic targets and optimize pharmacokinetic profiles. In contrast, wet-lab experimentation encompasses the full spectrum of physical laboratory work—including route scouting, multi-step organic synthesis, chromatographic purification, structural characterization via NMR and mass spectrometry, and comprehensive biological evaluation across cellular and animal models. Critically, AI-generated recommendations—such as novel synthetic pathways or structurally unprecedented lead candidates—often fail to account for practical constraints inherent in real-world laboratories: limited reagent availability, air- or moisture-sensitive intermediates, exothermic runaway reactions, scalability bottlenecks, regulatory restrictions on hazardous substances, and unanticipated off-target effects observed only in complex biological systems. Consequently, an overreliance on simulated outcomes without concurrent reinforcement of experimental technique risks producing graduates who possess strong algorithmic literacy but lack essential competencies in glassware handling, reaction monitoring, safety protocol adherence, troubleshooting under non-ideal conditions, and critical interpretation of discordant wet-lab data.

3.4. Boundary of Affective Interaction and Pharmaceutical Humanities Education

Teaching in pharmaceutical sciences integrates knowledge transmission with the cultivation of professional values and ethical dispositions. When examining historically significant developments—such as the pharmacovigilance lessons drawn from the thalidomide incident and the scientific, collaborative, and humanitarian dimensions of artemisinin's discovery—educators emphasize foundational principles including profound respect for human life, rigorous commitment to scientific integrity, and conscientious awareness of societal obligations inherent to pharmaceutical practice. These affective and value-laden dimensions are realized through dynamic, reciprocal teacher–student interactions characterized by dialogue, reflection, and shared inquiry [10]. Such pedagogical processes foster moral imagination, empathic reasoning, and contextual judgment—capacities that remain fundamentally distinct from algorithmic information retrieval. While generative artificial intelligence can efficiently supply factual summaries, historical timelines, or chemical data, it lacks the capacity for authentic emotional resonance, normative deliberation, or situated ethical interpretation. Human educators, by contrast, model reflective practice, respond contingently to learners' evolving concerns, and scaffold the internalization of humanistic commitments through lived example and sustained engagement. This irreplaceable relational dimension constitutes a core pillar of pharmaceutical humanities education, ensuring that scientific competence is consistently aligned with compassionate responsibility and social accountability.

4. Ethical Considerations of Introducing Generative AI into Medicinal Chemistry Teaching

The in-depth integration of generative artificial intelligence into medicinal chemistry education introduces multifaceted ethical challenges that demand rigorous, forward-looking reflection and proactive governance. These challenges encompass issues of academic integrity, such as the potential for overreliance on AI-generated molecular designs without sufficient critical evaluation of underlying chemical logic or pharmacological plausibility [9]. Pedagogical equity is also at stake, as disparities in access

to advanced computational infrastructure or AI literacy may widen achievement gaps among students from diverse institutional or socioeconomic backgrounds. Furthermore, transparency concerns arise regarding the provenance and training data of generative models—particularly when proprietary algorithms obscure decision-making pathways behind scaffold suggestions or property predictions. Instructors must carefully balance innovation with epistemic responsibility, ensuring that AI tools augment rather than replace foundational disciplinary reasoning, experimental validation, and ethical judgment in drug discovery contexts.

4.1. Academic-integrity Crisis and Cognitive-outsourcing Risks

Students increasingly rely on large-language models to generate drug-design reports, literature reviews, and other scholarly outputs for coursework and term assignments, raising serious concerns about academic integrity and intellectual development. This widespread practice constitutes a form of cognitive outsourcing, wherein learners delegate core epistemic tasks—such as information synthesis, critical evaluation of scientific evidence, hypothesis formulation, and logical argument construction—to artificial systems [11]. As a result, the distinction between authentic scholarly engagement and superficial content reproduction becomes increasingly ambiguous, undermining foundational academic values including originality, accountability, and intellectual ownership. Persistent dependence on AI-generated responses erodes essential metacognitive skills: students may gradually lose proficiency in conducting systematic literature searches across authoritative databases, interpreting primary research articles, identifying methodological limitations, and constructing evidence-based reasoning pathways. In pharmaceutical education specifically, this trend threatens the cultivation of domain-specific competencies required for rigorous drug discovery, safety assessment, and translational research. Educators therefore face an urgent pedagogical challenge: designing assessments that meaningfully integrate AI tools while preserving opportunities for deep learning, reflective practice, and authentic demonstration of subject mastery—without compromising ethical standards or long-term professional capability development.

4.2. Algorithmic Bias and Concealed Transmission of "Data Toxicity"

The output quality of generative artificial intelligence systems is fundamentally shaped by the composition, scope, and representativeness of their underlying training datasets. A notable limitation arises from the disproportionate reliance on English-language scientific literature and patent databases originating primarily from Western jurisdictions, which may inadvertently encode geographical, linguistic, and demographic imbalances [12]. These imbalances can manifest in downstream applications—for instance, in biomedical modeling tasks involving human genetic variation—where algorithmic outputs may fail to adequately reflect population-specific biological characteristics, such as allelic frequency distributions or pharmacokinetic enzyme polymorphisms across diverse ethnic groups. Consequently, recommendations generated for therapeutic development or molecular design may lack contextual validity when applied to broader global populations. Furthermore, historical datasets often contain outdated methodological descriptions, including synthetic procedures that employ hazardous reagents no longer aligned with contemporary safety standards, or experimental records referencing animal testing protocols that do not conform to current ethical review frameworks. When such content remains uncurated within training corpora, it risks being reproduced implicitly in model outputs, potentially influencing the knowledge formation of emerging researchers without appropriate critical contextualization or methodological safeguards.

4.3. Intellectual-property Disputes and Data-privacy Leakage

In academic teaching and research settings, educators and learners routinely submit unpublished experimental data, novel molecular configurations, and preliminary scientific hypotheses to publicly accessible artificial intelligence platforms to obtain

algorithmic feedback for refinement and validation. Such inputs often constitute original intellectual contributions protected under applicable legal frameworks governing scientific innovation. However, many platform service agreements grant operators broad rights to utilize user-submitted content—including proprietary research materials—for purposes such as model retraining, parameter optimization, and system enhancement. This practice may result in the unconsented incorporation of confidential drug-design concepts into large-scale training corpora, thereby undermining ownership claims and introducing significant risks to intellectual property integrity. Concurrently, intelligent educational platforms routinely construct detailed learner profiles by aggregating granular behavioral metrics—including interaction patterns, response latency, navigation sequences, and error typologies. When such data undergo insufficient anonymization or pseudonymization—particularly where identifiers remain recoverable or cross-referenced across datasets—it poses tangible threats to the confidentiality and autonomy of student information. Robust data governance protocols, including purpose limitation, data minimization, and rigorous desensitization standards aligned with national information security regulations, are therefore essential to uphold both academic integrity and digital privacy rights in AI-augmented pedagogy.

4.4. Ambiguous Accountability and Engineering-ethics Dilemmas

Consider a scenario in which a student relies entirely on an AI-generated synthetic pathway for a pharmaceutical chemistry graduation project, without conducting independent verification of experimental safety parameters. Should adverse outcomes arise—such as unintended hazardous reactions during laboratory execution—questions of responsibility become complex and multifaceted. The student bears primary academic and procedural responsibility for experimental design and safety compliance; the academic supervisor holds oversight duties including risk assessment guidance and pedagogical supervision; while the AI system developer assumes accountability for algorithmic robustness, transparency of limitations, and appropriate documentation of uncertainty in chemical prediction outputs. In pharmaceutical education and research, where precision directly impacts human health, even minor inaccuracies in AI-provided reaction conditions or reagent compatibility assessments may compromise therapeutic safety profiles or compound stability. Existing regulatory frameworks and institutional policies do not yet establish unambiguous protocols for assigning legal or ethical responsibility across this tripartite relationship. This ambiguity risks normalizing passive reliance on automated tools, potentially weakening critical evaluation habits and diminishing the cultivation of professional judgment, rigorous safety consciousness, and foundational engineering-ethics competencies essential for future pharmaceutical practitioners.

5. Teaching-reconstruction Strategies for Crossing Boundaries and Mitigating Ethical Risks

Rather than adopting a passive stance of rejection toward generative artificial intelligence, educators are called upon to proactively reconstruct pedagogical frameworks and cultivate a robust, ethically grounded educational ecosystem in which technological innovation and humanistic values coexist in dynamic equilibrium. This reconstruction entails not merely adapting existing practices but reimagining the epistemological foundations of teaching and learning—shifting emphasis from information transmission to critical sense-making, from algorithmic compliance to principled judgment, and from isolated knowledge acquisition to integrated professional formation. Such an ecosystem must be deliberately designed to foreground human agency, uphold academic integrity, and reinforce domain-specific epistemic standards, particularly within high-stakes disciplines such as pharmaceutical sciences where errors may carry profound implications for public health and safety.

First, a human–machine collaborative, human-oriented teaching paradigm must be systematically established. Within this paradigm, artificial intelligence functions strictly

as a cognitive support tool—a 'co-pilot'—while the educator retains full pedagogical authority and ethical responsibility as the 'captain' guiding all instructional decisions. AI-generated content is intentionally positioned not as authoritative truth but as a heuristic stimulus: its outputs serve as provocations for deeper inquiry rather than substitutes for disciplined reasoning. Pedagogical innovations such as 'AI error-hunting' exercises exemplify this approach—teachers curate drug-design proposals generated by large language models that contain scientifically plausible yet substantively flawed assumptions, inconsistencies in pharmacokinetic modeling, or violations of structural bioactivity principles. Students then engage in rigorous, evidence-based critique using foundational knowledge in medicinal chemistry, pharmacology, and molecular modeling, thereby strengthening analytical rigor, domain-specific discernment, and metacognitive awareness of AI's representational limitations.

Second, assessment systems must undergo comprehensive transformation to align with the realities of AI-augmented learning. A decisive shift is required from static, outcome-focused summative evaluation toward dynamic, process-oriented value-added assessment that captures intellectual growth, methodological sophistication, and ethical deliberation. Standardized tests emphasizing easily retrievable factual knowledge should be significantly reduced; instead, authentic assessments—including supervised hands-on laboratory experiments conducted without digital assistance, time-bound molecular-structure deduction tasks under offline conditions, and structured oral defenses of AI-assisted drug-design projects—must constitute core evaluation components [13]. Crucially, students are required to submit comprehensive prompt logs documenting their iterative engagement with AI tools: detailing how domain-specific questions were formulated with precision, how algorithmic outputs were critically appraised against empirical evidence and theoretical frameworks, and how AI suggestions were selectively adapted, modified, or rejected through reasoned integration with professional expertise. Proficiency in purposeful, transparent, and ethically reflective AI utilization is explicitly embedded in rubrics and weighted accordingly in final grading.

Third, AI literacy and pharmaceutical engineering ethics education must be elevated from peripheral topics to central, scaffolded curriculum pillars. Discipline-integrated micro-modules on the ethical dimensions of AI-aided drug discovery should be embedded throughout core courses such as Medicinal Chemistry, Pharmacokinetics, and Pharmaceutical Analysis. Instructors provide transparent explanations of large model architectures, training data provenance, known failure modes—including hallucination tendencies in chemical space prediction—and inherent privacy vulnerabilities associated with cloud-based inference. Institutionally endorsed, course-specific AI usage guidelines are codified and disseminated, clearly demarcating permitted applications (e.g., literature summarization under supervision), restricted uses (e.g., preliminary SAR analysis with mandatory experimental validation), and strictly prohibited actions (e.g., autonomous generation of regulatory documentation). Drawing upon documented historical cases—including adverse drug reactions linked to insufficient computational validation—the curriculum reinforces foundational tenets: no simulation, however sophisticated, can replace empirical verification; no algorithmic recommendation absolves human professionals of moral and legal accountability; and reverence for human life remains the non-negotiable cornerstone of pharmaceutical practice.

Fourth, institutions must prioritize the deployment of locally hosted, controllable, domain-specific large models tailored to pharmaceutical education and research [12]. Universities are encouraged to establish strategic tripartite collaborations among academic departments, pharmaceutical enterprises, and trusted domestic technology partners to jointly develop and maintain fine-tuned medicinal-chemistry-oriented language models. These models are trained exclusively on rigorously curated, high-fidelity corpora—including national pharmacopoeias, peer-reviewed journals indexed in authoritative databases, Good Manufacturing Practice (GMP) regulations, and validated compound-property datasets—ensuring terminological precision, contextual fidelity, and alignment with regulatory science standards. Domain specialization significantly reduces

the incidence of knowledge hallucinations, mitigates biases arising from overgeneralized training data, and enhances reliability in chemical nomenclature, reaction mechanism prediction, and ADMET profiling. Local deployment architecture ensures that all student interactions, research queries, and proprietary compound data remain fully contained within secure, on-campus infrastructure—eliminating external data exposure, safeguarding intellectual property, and fulfilling stringent requirements for data sovereignty and personal information protection under national cybersecurity and education informatization policies.

6. Conclusion

Generative AI has emerged as a transformative force in medicinal chemistry education, offering unprecedented capabilities for simulating molecular interactions, generating synthetic pathways, visualizing three-dimensional pharmacophores, and personalizing learning experiences through adaptive feedback systems. However, its integration must be guided by deliberate pedagogical intentionality—not merely technical feasibility. The deployment of such tools requires careful calibration to ensure alignment with core educational objectives: cultivating scientific rigor, nurturing ethical discernment, and reinforcing the humanistic foundations of pharmaceutical practice. Defining clear application boundaries is not an act of technological restriction but a necessary safeguard to preserve academic integrity, prevent epistemic overreliance on algorithmic outputs, and uphold the irreplaceable role of instructor-mediated interpretation and contextualization. In forward-looking pharmaceutical curricula, educators must assume an active stewardship role—designing scaffolded AI-assisted activities that foreground hypothesis generation, error analysis, and value-laden decision-making in drug discovery contexts. This includes explicit instruction on model limitations, data provenance awareness, bias detection in chemical prediction outputs, and reflective discourse on professional accountability. Furthermore, assessment frameworks must evolve beyond output validation to evaluate metacognitive engagement, ethical reasoning, and integrative synthesis across computational, experimental, and clinical domains. Ultimately, the most effective future paradigm envisions generative AI not as an autonomous knowledge provider, but as a collaborative cognitive partner—one that amplifies, rather than supplants, the educator's capacity to foster deep disciplinary understanding, empathetic patient-centered orientation, and unwavering commitment to public health responsibility.

References

1. L. T. Ramos and E. Casas, "Applications, trends, and perspectives of large language models in education: A literature review," *Frontiers in Education*, 2025.
2. X. Chen, J. Zhang, Q. Zhao, L. Ding, Z. Wu, Z. Jia, and D. He, "Application and teaching of computer molecular simulation embedded technology and artificial intelligence in drug research and development," *Open Life Sciences*, vol. 18, no. 1, p. 20220675, 2023.
3. R. AlAli and Y. Wardat, "Opportunities and challenges of integrating generative artificial intelligence in education," *International Journal of Religion*, vol. 5, no. 7, pp. 784–793, 2024.
4. I. M. García-López and L. Trujillo-Liñán, "Ethical and regulatory challenges of Generative AI in education: a systematic review," in *Frontiers in Education*, vol. 10, p. 1565938, Frontiers Media SA, Jun. 2025.
5. A. Gangwal, A. Ansari, I. Ahmad, A. K. Azad, V. Kumarasamy, V. Subramaniyan, and L. S. Wong, "Generative artificial intelligence in drug discovery: basic framework, recent advances, challenges, and opportunities," *Frontiers in Pharmacology*, vol. 15, p. 1331062, 2024.
6. W. P. Walters and M. Murcko, "Assessing the impact of generative AI on medicinal chemistry," *Nature Biotechnology*, vol. 38, no. 2, pp. 143–145, 2020.
7. X. Tang, H. Dai, E. Knight, F. Wu, Y. Li, T. Li, and M. Gerstein, "A survey of generative AI for de novo drug design: new frontiers in molecule and protein generation," *Briefings in Bioinformatics*, vol. 25, no. 4, p. bbae338, 2024.
8. Q. Vanhaelen, Y. C. Lin, and A. Zhavoronkov, "The advent of generative chemistry," *ACS Medicinal Chemistry Letters*, vol. 11, no. 8, p. 1496, 2020.
9. X. Tong, X. Liu, X. Tan, X. Li, J. Jiang, Z. Xiong, and M. Zheng, "Generative models for de novo drug design," *Journal of Medicinal Chemistry*, vol. 64, no. 19, pp. 14011–14027, 2021.

10. X. Zeng, F. Wang, Y. Luo, S. G. Kang, J. Tang, F. C. Lightstone, and F. Cheng, "Deep generative molecular design reshapes drug discovery," *Cell Reports Medicine*, vol. 3, no. 12, 2022.
11. K. Swanson, G. Liu, D. B. Catacutan, A. Arnold, J. Zou, and J. M. Stokes, "Generative AI for designing and validating easily synthesizable and structurally novel antibiotics," *Nature Machine Intelligence*, vol. 6, no. 3, pp. 338–353, 2024.
12. A. Gangwal and A. Lavecchia, "Artificial intelligence in natural product drug discovery: current applications and future perspectives," *Journal of Medicinal Chemistry*, vol. 68, no. 4, pp. 3948, 2025.
13. Y. Bian and X. Q. Xie, "Generative chemistry: drug discovery with deep learning generative models," *Journal of Molecular Modeling*, vol. 27, no. 3, p. 71, 2021.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Publisher and/or the editor(s). Publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.