

A Conceptual Hybrid Genetic–Tabu Framework for Chatbot Selection and Allocation in Vocational Colleges

Suhan Wu^{1,a}, Min Luo^{2,b,*}

¹School of Economics and Management, Nanjing Polytechnic Institute, No.188 Xinle Road, Luhe District, Nanjing, 210048, Jiangsu, China

²School of Management, Shenzhen University of Information Technology, Longxiang Road 2188, Shenzhen, 518172, Guangdong, China

^awusuhan@link.cuhk.edu.hk, ^bluom@suit-sz.edu.cn

*Corresponding author

Abstract: The rapid diffusion of large language models has accelerated the uptake of educational chatbots, yet vocational colleges face a distinctive deployment challenge shaped by competence-oriented task chains, heterogeneous instructional settings, and stringent governance requirements. This paper conceptualizes chatbot adoption as an integrated decision problem in which selection and allocation must be determined jointly under binding constraints on budget, infrastructure capacity, faculty workload, and institutional obligations related to privacy, academic integrity, accountability, and auditability. Building on this framing, we develop a governance-aware Hybrid Genetic–Tabu framework that structures decision-making through a feasibility-first explore–refine workflow. Population-based exploration is used to generate diversified admissible selection–allocation plans, while tabu-guided refinement improves promising candidates via interpretable local adjustments and mitigates cycling under hard constraints. The framework further emphasizes traceability by producing an auditable decision artifact that makes decision criteria, binding constraints, and oversight boundaries explicit, thereby supporting cross-unit coordination and iterative adaptation as tools, policies, and instructional needs evolve. The study offers a compact methodology for responsible chatbot deployment in vocational colleges and outlines directions for future work on empirical assessment, preference elicitation, and domain-sensitive extensions for higher-risk training contexts.

Keywords: Vocational Education; Educational Chatbots; Generative AI Governance; Selection and Allocation; Hybrid Genetic–tabu Framework

1. Introduction

The recent advances in large language models and conversational agents have accelerated the deployment of educational chatbots across higher education ^[1]. In vocational colleges, this momentum is amplified by competency-based training, practice-oriented curricula, and the demand for scalable instructional support across classrooms, laboratories, and workplace simulation activities ^[2]. At the same time, vocational teaching imposes distinctive requirements: learning tasks are closely aligned with operational procedures and industry standards; instruction is distributed across multiple settings; and institutions must balance innovation with governance obligations related to privacy, accountability, and academic integrity ^[3]. Accordingly, the central challenge is increasingly not whether to adopt chatbots, but how to make deployment decisions that are pedagogically appropriate, operationally feasible, and institutionally defensible.

In practice, vocational colleges face two intertwined decision problems. The first is selection, namely identifying suitable chatbot technologies or configurations from a rapidly evolving landscape of models, products, and deployment modes. Alternatives vary in instructional functionality, integration complexity, and governance-related properties such as controllability, logging, and compliance readiness ^[4]. The second is allocation, namely determining how selected chatbots should be assigned and provisioned across heterogeneous teaching tasks, courses, and student groups. Allocation decisions typically involve mapping chatbots to course task chains, defining access tiers or usage quotas, and specifying human-in-the-loop arrangements for sensitive instructional activities, all under binding constraints of budget, infrastructure capacity, and faculty workload ^[5].

Existing research offers important foundations but leaves a practical gap for vocational colleges [6]. Many educational studies emphasize pedagogical affordances, learner experience, or ethical principles, yet stop short of providing an operational decision structure that supports systematic selection and allocation. In contrast, the literature on constrained resource allocation and combinatorial optimization provides well-established approaches for complex decision-making, but these are often presented without explicit integration of institutional governance conditions that shape educational deployments. This disconnect is particularly consequential in vocational settings, where adoption outcomes depend on transparent rationales, traceability, and alignment with teaching capacity and policy constraints.

To address this gap, this paper proposes a conceptual Hybrid Genetic–Tabu framework for chatbot selection and allocation in vocational colleges. The framework is designed as a governance-aware decision engine that organizes the deployment process into three components: (i) formalizing criteria and constraints, (ii) generating feasible configuration candidates through population-based exploration, and (iii) refining candidate solutions via local search that preserves feasibility while improving task fit. The hybridization is motivated by complementary strengths: genetic-algorithm logic supports broad exploration across diverse instructional needs, whereas tabu-search logic enables disciplined local improvements while mitigating cycling and maintaining constraint satisfaction.

The remainder of the paper is organized as follows. Section 2 reviews related work on educational chatbots and the shift toward LLM-based agents, governance-oriented technology selection, and decision-support perspectives for multi-criteria selection and constrained allocation. Section 3 frames chatbot selection and allocation as an integrated decision problem by clarifying the decision scope, constraints, and expected outputs. Section 4 presents the conceptual Hybrid Genetic–Tabu framework and explains its feasibility-first explore–refine logic and traceability considerations. Section 5 concludes by summarizing the main results and discussing implications and directions for future research.

2. Literature Review

This review is organized into three complementary strands. First, we synthesize findings on educational chatbots and the recent shift toward LLM-based agents, with emphasis on capability boundaries relevant to competence-oriented learning tasks. Second, we review governance-oriented guidance for generative AI adoption in higher education to clarify the institutional requirements—privacy, integrity, transparency, and accountability—that shape feasible deployment choices. Third, we examine decision-support perspectives that formalize multi-criteria selection and constrained allocation, providing the methodological rationale for a hybrid GA–Tabu decision engine.

2.1 Educational chatbots in higher education and the shift toward LLM-based agents

Prior research has examined educational chatbots as conversational supports for tutoring, feedback, Q&A, and student services, with systematic reviews reporting perceived gains in immediacy and convenience while repeatedly flagging risks such as unreliable responses, misalignment with pedagogy, and uneven evaluation practices [6,7]. As LLM-based agents have become mainstream, recent syntheses have further noted that empirical evidence remains heterogeneous in outcomes and measures, and that many studies still privilege short-horizon satisfaction indicators over competence-oriented learning effects and robust instructional design rationales [8]. In vocational education and training (VET), the promise of chatbots is often articulated in relation to procedural guidance, competence development, and situated task performance; however, VET-specific work on LLM-enabled chatbots remains emergent and tends to foreground domain grounding, assessment sensitivity, and hallucination control as design constraints rather than purely functional expansion [9].

2.2 Governance-oriented selection and responsible use: from principles to institutional decision requirements

As generative AI diffuses into teaching and assessment, governance has become a first-order determinant of sustainable adoption, shifting the decision focus from tool enthusiasm to institutional defensibility. Global guidance emphasizes human-centred use, transparency, accountability, privacy protection, and capacity building, and calls for policy and oversight mechanisms that can translate broad principles into actionable institutional practice [10]. In higher education, regulators and sector bodies increasingly frame generative AI integration through assessment integrity, risk management, and cross-functional governance, stressing that institutional approaches require clear rules on permissible use, documentation and traceability, and continuous review as tools and risks evolve [11]. Complementary

practitioner-oriented guidance on assessment design similarly highlights structured decision pathways (e.g., when to permit use with disclosure, when to redesign toward authentic assessment, and how to evaluate process as well as product), reinforcing that selection and deployment decisions are inseparable from governance checkpoints ^[12].

2.3 Decision-support perspectives: multi-criteria selection and heuristic allocation under constraints

Technology adoption and configuration are often treated as structured decision problems in which multiple criteria and stakeholder preferences must be reconciled; decision-support systems and MCDM approaches provide a common rationale for making selection transparent by formalizing criteria, weighting, and evidence ^[4,13]. Education-oriented work on AI adoption likewise motivates multi-criteria views that include trust-related and contextual considerations alongside usability and technical factors, indicating that “selection” is rarely reducible to a single performance metric ^[14]. Once selection is coupled with allocation—assigning tools and usage capacity across heterogeneous tasks, courses, and student groups under binding constraints—the decision space becomes combinatorial, motivating heuristic or metaheuristic solution logics as pragmatic engines for feasible planning. Hybridization between genetic algorithms and tabu search is widely justified in the optimization literature on the basis of complementarity: population-based exploration supports diversity and broad search, while tabu-guided local improvement mitigates cycling and strengthens refinement under constraints ^[5,15].

2.4 Research gap

Collectively, prior research points to a coherent but still incomplete knowledge base for decision-making in vocational colleges. Educational chatbots have become increasingly capable—particularly with the transition from scripted systems to LLM-based agents—yet the associated evidence base remains uneven in both evaluation designs and outcome measures. At the same time, governance considerations have moved from peripheral concerns to binding constraints: requirements related to privacy protection, academic integrity, accountability, and auditability increasingly determine whether adoption can be justified institutionally and sustained at scale. Parallel to these developments, decision-support perspectives and hybrid metaheuristic logics provide established principles for structuring complex choices under multiple criteria and hard constraints. However, for vocational colleges, what remains comparatively underdeveloped is an integrated, governance-aware decision structure that explicitly links chatbot selection (what to adopt) with allocation (how to provision and assign) in a manner aligned with competence-oriented task chains, while maintaining traceability and institutional accountability. This gap motivates the decision framing and the conceptual Hybrid Genetic–Tabu framework developed in the subsequent sections.

3. Decision Problem Framing in Vocational Colleges

This section frames chatbot deployment in vocational colleges as an integrated decision problem in which selection and allocation must be determined jointly under binding resource constraints and institutional governance requirements. The aim is to establish a precise decision scope, articulate the informational structure of the problem, and clarify what constitutes an implementable and defensible decision output. This framing provides the conceptual foundation for the hybrid decision engine introduced in the subsequent section.

3.1 Decision scope: defining selection and allocation

In vocational colleges, selection concerns the determination of an admissible set of chatbot technologies or configurations that the institution is prepared to adopt for instructional purposes. Selection should be interpreted as a portfolio decision because instructional needs are heterogeneous and vary across disciplines, course structures, and competence-oriented task chains. A “configuration” may differ not only in model choice, but also in deployment mode and functional design. Differences in retrieval augmentation, tool integration, multilingual support, access control, and safety constraints can materially affect pedagogical suitability, operational feasibility, and governance risk.

Allocation addresses how the selected portfolio is provisioned and aligned with concrete teaching tasks, courses, and student groups. In vocational contexts, allocation is closely tied to task-chain organization, where learning activities are structured around procedural steps and performance criteria derived from occupational standards. Allocation therefore requires explicit alignment between task requirements and configuration capabilities, while also specifying how access and capacity are

operationalized across courses and cohorts. Where instructional activities intersect with assessment, individualized guidance, or student welfare, allocation further entails defining the boundary conditions for human oversight and escalation, thereby ensuring that chatbot assistance remains consistent with institutional responsibility.

3.2 Entities and informational structure

The integrated decision problem can be described through a structured set of entities and attributes. The first entity set comprises teaching tasks, which can be defined at a course, module, or activity level, depending on the granularity at which the institution plans and monitors instructional support. These tasks are characterized by competence targets, procedural requirements, and acceptable support boundaries. The second entity set comprises candidate chatbot configurations, each described by a capability profile, an integration profile, and a governance profile. Capability profiles capture instructional functions relevant to vocational learning, such as procedural tutoring, feedback generation, or multilingual communication. Integration profiles reflect compatibility with institutional platforms and workflows, including authentication and content systems. Governance profiles capture properties that affect institutional defensibility, such as controllability, logging, data retention, and safety constraints. The third entity set comprises user groups, because differences in language proficiency, access needs, and exposure to risk can influence both capability requirements and governance safeguards.

These entities jointly determine the informational inputs required for decision-making. Task-side inputs articulate what support is pedagogically meaningful and permissible for each task chain. Configuration-side inputs specify what each candidate can deliver, at what cost and with what operational burden. Institution-side inputs encode governance rules and operational constraints that must be satisfied for deployment to be legitimate and sustainable. Importantly, governance is treated as intrinsic to the problem definition rather than an external checklist applied after an allocation is produced.

3.3 Constraints and governance assumptions

The decision space is restricted by a combination of resource constraints and governance constraints, which together define feasibility. Resource constraints typically arise from limited budgets and capacity, covering licensing or compute costs, infrastructure limits such as concurrency or network conditions, and the faculty workload required for training, monitoring, and escalation handling. These constraints influence both the size and composition of the selected portfolio and the intensity with which chatbots can be provisioned across tasks and cohorts.

Governance constraints specify the institutional conditions under which chatbot use is permissible. They commonly reflect privacy and data protection requirements, including restrictions on sensitive data handling and retention; academic integrity requirements, particularly where chatbot use may intersect with graded work or assessment preparation; and accountability requirements that demand traceability, documentation, and role clarity in oversight. In vocational colleges, governance assumptions are further shaped by competence standards and safety considerations, because procedural guidance is not purely informational but can influence students' understanding of operational norms and professional conduct. Feasible deployment therefore requires that capability provision, access control, and oversight arrangements be co-designed rather than treated as separable stages.

3.4 Decision outputs and interpretability requirements

The decision outcome is a governance-aware selection–allocation plan that can be implemented within existing institutional processes and defended under audit or review. Such a plan delineates a portfolio of chatbot configurations that satisfy institutional admissibility requirements and specifies how these configurations are aligned with competence-oriented task chains across courses, learning activities, and student groups. Operational feasibility is ensured by articulating provisioning arrangements that translate allocation decisions into actionable access rules and capacity controls, thereby reflecting budgetary limits, infrastructure capacity, and faculty workload. In parallel, the plan embeds governance checkpoints throughout the deployment logic, clarifying where human oversight is required and establishing traceability through appropriate documentation and logging practices. Under this framing, the decision output is not a simple assignment of tools to tasks, but an auditable decision artifact that makes criteria, constraints, and trade-offs sufficiently explicit to support responsible deployment and iterative refinement. Building on this formalization, the next section introduces a conceptual Hybrid Genetic–Tabu decision framework that operationalizes the integrated selection–allocation problem

through complementary global exploration and local refinement mechanisms.

4. Conceptual Hybrid Genetic–Tabu Framework

This section presents a governance-aware Hybrid Genetic–Tabu (GA–Tabu) framework that operationalizes the integrated selection–allocation problem defined in Section 3. The framework is positioned as a decision methodology for vocational colleges, structuring how institutions can move from heterogeneous instructional needs and governance boundaries to an implementable and defensible deployment plan. Table 1 summarizes the workflow; the discussion below clarifies the rationale that links feasibility preservation, disciplined refinement, and institutional traceability.

Table 1. A concise governance-aware Hybrid GA–Tabu framework for chatbot selection and allocation

Stage	Key focus	Main output
1. Scope and inputs	Specify task chains, user groups, candidate chatbot configurations, and institutional governance boundaries	Decision scope and admissible input set
2. Criteria and constraints	Translate pedagogical goals and governance requirements into criteria and hard constraints (e.g., privacy, integrity, capacity, workload)	Formalized criteria/constraint set
3. Feasible initialization (seed plans)	Construct an initial set of admissible selection–allocation plans under feasibility-first rules	Feasible candidate plan set
4. GA-based exploration	Generate diverse candidate plans via population-based variation while enforcing hard constraints	Diversified feasible plan pool
5. Tabu-guided refinement	Improve promising plans through local moves (replace/swap/adjust provisioning) with tabu restrictions to avoid cycling	Locally improved feasible plans
6. Plan finalization and traceability	Select a defensible plan and record the rationale trace (criteria applied, binding constraints, oversight points)	Auditable selection–allocation plan

4.1 Framework rationale and architecture

Chatbot deployment in vocational colleges is best understood as a constrained institutional decision rather than a purely technical procurement or an ad hoc pedagogical intervention. Selection and allocation are coupled: admissible configurations are determined by governance and infrastructure realities, while feasible allocation depends on task-chain heterogeneity, workload sustainability, and integrity-sensitive boundaries. The proposed framework therefore treats governance requirements and hard constraints as part of the decision space, rather than as post-hoc checks applied after candidate plans are produced.

As summarized in Table 1, the framework follows a feasibility-first explore–refine logic. The early components formalize admissible inputs, criteria, and hard constraints. The middle components generate diversified feasible plans, and the later components refine and consolidate promising candidates into an actionable selection–allocation plan with an accompanying rationale trace. This architecture aligns with institutional practice, where acceptability, accountability, and implementability are as consequential as technical capability.

4.2 Admissibility, constraints, and feasibility-first planning

A central design principle is the separation between evaluative criteria and binding constraints. Criteria articulate institutional priorities—such as competence alignment, task-chain coverage, integration practicality, and operational sustainability—where trade-offs are expected. Constraints specify what cannot be violated, including budget ceilings, capacity limits, and governance requirements related to privacy, integrity, accountability, and auditability. Maintaining this distinction prevents feasibility from being conflated with preference and enables disciplined decision-making among admissible alternatives.

Feasibility-first planning follows from this structure. Admissibility screening should occur early so that configurations lacking minimal governance properties do not enter the candidate pool. When partial violations arise during candidate generation, feasibility is restored through conservative adjustments, such as reassignment to lower-risk configurations or reductions in provisioning intensity, before candidates proceed to refinement. This is consistent with institutional realities: solutions that violate privacy or integrity requirements are not “nearly feasible” in practice.

4.3 GA-based exploration and tabu-guided refinement

Within the admissible space, GA logic is used to generate structurally diverse selection–allocation plans. The value of population-based exploration is not limited to optimization; it supports the systematic surfacing of qualitatively different deployment strategies in settings where task requirements vary across courses, modules, and student groups. Candidate plans can differ in portfolio composition, task-to-configuration alignment, and provisioning regimes, while remaining within hard constraints. In this framing, evaluation functions primarily as screening among feasible trade-offs rather than as a pursuit of a single technical optimum.

Tabu-guided refinement complements exploration by improving promising candidates through local, interpretable adjustments that correspond to realistic planning moves. These include replacing configurations for subsets of tasks, swapping assignments across tasks with similar capability requirements, and adjusting provisioning rules to relieve workload or capacity pressures without crossing governance boundaries. The tabu mechanism functions as a stability device, discouraging immediate reversals and repetitive cycling among near-identical plans, thereby promoting coherent improvement under constraint. Where warranted, aspiration can override tabu restrictions to remove a binding feasibility bottleneck or achieve materially better alignment with institutional priorities.

4.4 Decision finalization and traceability

The framework culminates in a selection–allocation plan that is actionable within institutional processes and defensible to stakeholders. The plan specifies how selected chatbot configurations are aligned with competence-oriented task chains and how allocation is operationalized through provisioning rules consistent with resource limits. Equally important, it is accompanied by a rationale trace that records the criteria used, identifies binding constraints that shaped the plan, and specifies governance-relevant boundary conditions such as oversight requirements and documentation scope. This traceability supports auditability, cross-unit coordination, and iterative refinement as institutional priorities and governance rules evolve. Taken together, Table 1 and the foregoing discussion offer a concise articulation of how a feasibility-first explore–refine workflow can guide chatbot selection and allocation in vocational colleges, while sustaining governance defensibility and decision transparency.

5. Conclusion

5.1 Main results

This paper proposed a conceptual Hybrid Genetic–Tabu (GA–Tabu) framework to structure chatbot selection and allocation in vocational colleges under binding resource constraints and institutional governance requirements. By framing deployment as an integrated decision problem, the study clarifies that selection and allocation should be co-determined in relation to competence-oriented task chains, rather than treated as separable stages or informal adoption practices. The resulting feasibility-first explore–refine workflow supports the generation of admissible deployment alternatives and the disciplined improvement of promising plans, while keeping decision defensibility as a central design objective. A further contribution lies in treating governance requirements—privacy, academic integrity, accountability, and auditability—not as peripheral considerations but as feasibility conditions that delimit the admissible decision space, and in emphasizing traceability through a rationale trace that makes criteria, binding constraints, and oversight boundaries explicit.

5.2 Implications and directions for future research

The framework carries several implications for institutional practice. It suggests that vocational colleges can improve deployment quality by adopting criteria-driven, governance-aware decision procedures that connect chatbot capabilities to task-chain requirements and translate allocations into concrete provisioning rules. Beyond decision structuring, the framework also implies a practical governance pathway in which deployment is embedded into routine institutional processes: responsibilities should be clarified across academic affairs, IT support, teaching teams, and compliance functions; admissibility conditions and usage boundaries should be articulated in advance for assessment-adjacent and other integrity-sensitive activities; and provisioning should be coupled with minimal but explicit monitoring arrangements, including escalation routes and periodic review cycles. In addition, accountability should be operationalized through auditable decision artifacts, so that choices about

selection, allocation intensity, and oversight requirements are traceable and can be updated coherently as tools, policies, and instructional needs evolve.

Future research can extend the present work in three directions. First, empirical studies could examine how the proposed decision structure influences implementation outcomes, including governance compliance, faculty workload, and perceived instructional value across different vocational disciplines. Second, methodological extensions could incorporate richer preference elicitation and multi-stakeholder weighting schemes to better represent trade-offs between pedagogical flexibility, operational cost, and governance strictness. Third, domain-sensitive variants may be required for high-risk vocational contexts where procedural guidance has safety implications, motivating more stringent admissibility rules, oversight design, and audit requirements.

Conflict of Interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Acknowledgement

This research is supported by the Jiangsu Provincial Education Science Planning Project of Jiangsu Academy of Educational Sciences (No. C/2025/02/96), Key Research Area Support Program of Shenzhen University of Information Technology (No. SZIIT2024KJ054) and the Scientific Research Foundation for High-level Talents (No. RC2024-005).

References

- [1] Kasneci E, Seßler K, Küchemann S, et al. ChatGPT for good? On opportunities and challenges of large language models for education[J]. *Learning and Individual Differences*, 2023, 103: 102274.
- [2] Guthrie H. Competence and competency-based training: What the literature says[M]. ERIC, 2009.
- [3] Cotton D R, Cotton P A, Shipway J R. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT[J]. *Innovations in Education and Teaching International*, 2024, 61(2): 228-239.
- [4] Farshidi S, Jansen S, de Jong R, et al. A decision support system for software technology selection[J]. *Journal of Decision Systems*, 2018, 27(sup1): 98-110.
- [5] Glover F, Kelly J P, Laguna M. Genetic algorithms and tabu search: Hybrids for optimization[J]. *Computers & Operations Research*, 1995, 22(1): 111-134.
- [6] Labadze L, Grigolia M, Machaidze L. Role of AI chatbots in education: Systematic literature review[J]. *International Journal of Educational Technology in Higher Education*, 2023, 20(1): 56.
- [7] Okonkwo C W, Ade-Ibijola A. Chatbots applications in education: A systematic review[J]. *Computers and Education: Artificial Intelligence*, 2021, 2: 100033.
- [8] Dong B, Bai J, Xu T, et al. Large language models in education: A systematic review[C]. 2024 6th International Conference on Computer Science and Technologies in Education (CSTE), 2024: 131-134.
- [9] Yu Y, Zhao Z. Development of vocational education and training chatbot supported by large language model-based multi-agent system[J]. *Vocation, Technology & Education*, 2025.
- [10] Holmes W, Miao F. Guidance for generative AI in education and research[M]. Unesco Publishing, 2023.
- [11] Kofinas A K, Tsay C H H, Pike D. The impact of generative AI on academic integrity of authentic assessments within a higher education context[J]. *British Journal of Educational Technology*, 2025.
- [12] Gundu T. Strategies for e-assessments in the era of generative artificial intelligence[J]. *Electronic Journal of e-Learning*, 2024, 22(7): 40-50.
- [13] Salvado L L, Villeneuve E, Masson D, et al. Decision support system for technology selection based on multi-criteria ranking: Application to NZEB refurbishment[J]. *Building and Environment*, 2022, 212: 108786.
- [14] Wiangkham A, Vongvit R. Comparative analysis of multi-criteria decision making methods for prioritizing influential factors of ChatGPT adoption in higher education[J]. *Expert Systems with Applications*, 2025: 128188.
- [15] Li D, Wang L, Wang M. Genetic algorithm and tabu search: A hybrid strategy[J]. *IFAC Proceedings Volumes*, 1999, 32(2): 8664-8668.