Article

Constructing an Intelligent Admissions Model in Chinese Higher Education under the Strategy of Building a Leading Education Nation: Technical Architecture, Implementation Pathway, and Future Prospects

Chuanbao Peng[1,2], Tingting Zhang[1,*], Wei Liu[1], Liming Guo[1] and Chunxiao Wang[1]

[1] Kede College of Capital Normal University, Beijing, China
[2] City University of Macau, Macau, China
* Correspondence: Tingting Zhang, Kede College of Capital Normal University, Beijing, China; City University of Macau, Macau, China

Abstract: Intelligent admissions projects in universities usually stall for reasons that have little to do with technology: data access and decision authority. This paper applies one test to every component of an intelligent admissions model: can it be built and accepted now, within the institution's own authority and budget? Components that fail are sorted into three further grades, depending on whether the obstacle is data, authority, or both. Provincial behavior data, cross-database matching, household background fields, and gain-based early warning all fail the test. Four components pass: policy question answering, admissions-stage data governance, enrollment scenario estimation, and post-admission outreach. Fairness cannot be computed by a model. It has to be checked, through capacity rules, channel parity, and the right to appeal.

Keywords: intelligent admissions; technical architecture; feasibility assessment; institutional governance; educational equity

1. Introduction

1.1. Policy Context and Four Dilemmas

In 2025, China's gross enrollment rate in higher education reached 61.30%, with 48.7257 million students enrolled [1]. Expansion is no longer the main problem. The Outline of the Plan for Building a Leading Education Nation (2024–2035) asks admissions to match talent supply with industrial needs, to put quality before scale, and to win public trust in how opportunities are allocated [2]. All three tasks run through one admission process, bounded by rules, data, and authority, and none of them can bypass it. Whether that process is intelligent depends on three checkable facts: who owns the data, who may use it, and who maintains it.

Four problems block the process. Decisions are discontinuous: planning extends last year's totals, while bounded rationality says actors under time pressure settle for good-enough choices [3]. Information is asymmetric: candidates cannot verify program claims and fall back on reputation and scores, although timely and complete admission information changes the decisions of disadvantaged students [4,5]. Service stops at admission, so low-income and first-generation students keep facing information gaps that institutional support could narrow [5]. Records are not reused: reasons for non-enrollment rarely flow back, yet sound data governance requires records to be kept and used again [6]. The four problems reinforce one another, which is why adding systems has not helped.

1.2. Questions, Method, and Delimitations

Received: 12 September 2026
Revised: 23 October 2026
Accepted: 04 November 2026
Published: 10 November 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/).

Three questions follow. What is an intelligent admissions model? What architecture covers the whole cycle? And which components are deliverable now, which must be downgraded, which lack data, and which exceed institutional authority? This is a conceptual design-science study: the object is an artifact, and the standard is deliverability, acceptance, and accountability under constraints [7]. No design–implement–revise cycle was run, and the literature review is narrative rather than systematic. Three limits hold: decisions with legal effect stay with authorized bodies; traceability is not traded for accuracy; the result is a design that holds when conditions are met, not a blueprint.

2. The Model and Its Architecture

2.1. Conceptual Definition

An intelligent admissions model is an institutional arrangement in which data, algorithms, and routines support admission decisions under human oversight. Four points fix the definition. Technology is one element beside data rules, job duties, and appeal procedures. The scope covers the whole cycle: consultation, application, admission, registration, and feedback after enrollment. A system never decides eligibility, handles exceptions, or admits anyone. The model is judged by deliverability, verifiability, and accountability, not by the number of functions.

2.2. Theory and the Five-Layer Structure

Three lines of theory set the requirements. Signaling theory treats admission as two-sided information asymmetry, so services should lower verification costs [4]. Bounded rationality says candidates satisfice, so services should ease comparison rather than multiply options [3]. Design science ties success to usefulness under constraints [7]. Four commitments follow: every output cites its source text; structured items replace prose; services attach to process nodes that already exist; every output maps to a model version and a data version. If any commitment fails, the outcome is a failed design rather than a missing function.

The architecture has five layers and two pillars (Figure 1). Infrastructure supplies computing, storage, and authentication. The data layer handles collection, standardization, classification, and retention. The algorithm layer runs dialogue, matching, prediction, decision support, early warning, and evaluation. The application layer organizes scenarios for candidates, high schools, and internal units. Governance sets permissions, logging, review, and appeal, while standardization and operations run across all five. Governance decides which fields may be collected, which outputs need review, and which scenarios must not go live. Without it the system runs, but nobody can say who approved it or who answers for it.

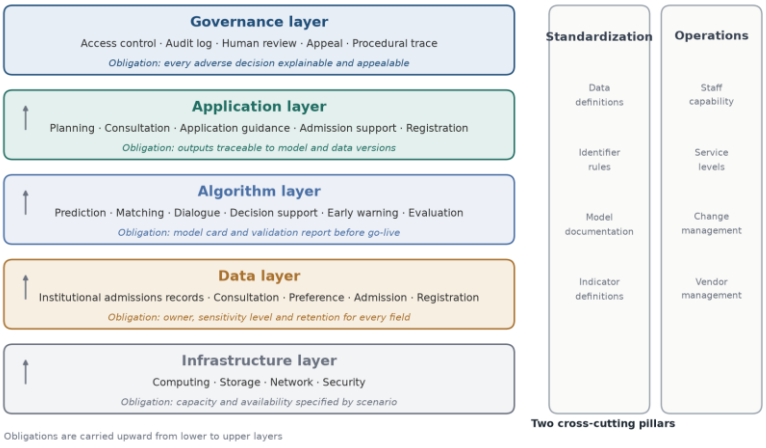


Figure 1. A Five-Layer Reference Architecture for Intelligent Admissions with Two Cross-Cutting Pillars (This Is an Ideal-Type Design. As Section 3 Shows, Only Part of the Data and Algorithm Layers Falls within What an Institution Can Deliver.)

3. Feasibility: A Four-Grade Scale

3.1. The Criterion

One test runs through this section: can the component be built and accepted now, within existing authority and budget, without external authorization? Anything that needs a rule change first, a provincial authorization, or ten years of accumulated data falls outside the deliverable range. The test yields four grades (Table 1).

Table 1. Feasibility Grades of Intelligent Admissions Components

Grade	Components	Basis	Disposition
1 Deliverable now	Policy Q and A; data governance; scenario estimation; post-admission outreach	Data and authority inside the institution	Minimum deliverable set (Table 2); independently verifiable
2 Workable, downgrade	Allocation simulation; program consideration sets; vendor contract controls	Placement parameters held by provinces	Narrow the claim; label outputs as reference
3 Data-starved	Gain-based early warning; effect evaluation; digital twin	No intervention control; broken series	Research agenda only
4 Beyond authority	Provincial platform; privacy computing; algorithm review	Built and authorized above institutional level	Removed from roadmaps; access preparation only

Note: Grades rest on data availability and authority, not technical maturity. One component may span two grades across institutions.

3.2. Data: Two Unavailable, Two Impermissible

Provincial candidate behavior data are closed to institutions. The 2025 Regulations make the institution responsible and the admissions office supervisory; provinces provide source distributions and transfer files of placed candidates only, and the 2026 notice keeps every stage under these rules [8,9]. An institution therefore holds its own records plus aggregate distributions; the preferences of candidates placed elsewhere are out of reach, so program-level forecasting must drop to interval-based estimates.

Household background fields do not exist. Rules such as "service coverage no lower than parity" assume background data. Admission records carry no income, parental occupation, or first-generation status; only a few categories exist, and provinces open them unevenly, so any coverage rule has nothing to measure. Cross-database matching is not permitted: placement data sit with provinces and academic records with the national student service, and their interfaces support query, not bulk matching. The same wall stands inside the institution, whose remit ends at registration. The first step is to link the five admissions-stage record types at the individual level.

Consultation records cannot feed models. Articles 6, 14, and 24 of the Personal Information Protection Law require purpose limitation, explicit consent, and fairness in automated decisions with a right to refuse [10]. Statistical use stays within the original purpose; training, scoring, and profiling change it, and in practice consent cannot be collected because consultation concentrates in a short window at high volume. Profiling on consultation text stays off by default.

3.3. Algorithms: Workable, Downgraded, Data-Starved

Advisory response. The bottleneck is corpus upkeep and peak load. Statutes change yearly, so without a named owner and frozen versions the system cannot say which year's rule it cites. Demand concentrates around the application window at ordinary staffing levels, so tiering answers it: forms and retrieval for structured questions, a callback queue for judgment calls, refusal with a human channel for the rest.

Application guidance. The feasible level is low. The annual notice gives secondary schools the main venue and bans paid consulting by institutions; a published cross-institutional consideration set that costs a candidate placement becomes the institution's liability [9]. What remains is bounded by the institution's own rules: no generated lists, no probability claims, with the "ten prohibitions" codified since 2012 written into refusal rules [11].

Simulation and early warning. Placement parameters sit with provinces, so an institution can only simulate its own placed cohort under its own rules; a digital twin needs longitudinal data whose series break at program closures. Early warning needs treated and untreated groups, but outreach is never randomized, so a ranker learns who looks unlikely to enroll, not who would change if contacted. Output ranks cases for human review; purchased models are governed by contract: a reproducible test interface, acceptance on a test set, and an exit plan for migration.

3.4. Institutions: Three Incompatible Designs

Three governance designs fail the same test of fit. A named officer empowered to suspend operations has no place in the collective deliberation the notice protects; the compromise separates recommending suspension from deciding it [9]. Self-monitoring is biased, because the admissions office would grade its own work; monitoring belongs to the higher education authority. A task force lacks rank, since budget and procurement veto sit with the president, so the objective drops to three lists that the president's office can actually sign off.

4. Implementation Pathway

4.1. The Minimum Deliverable Set

The four Grade 1 components share one profile: data in institutional systems, authority inside the institution, no provincial approval, no background fields, no new posts. Table 2 lists their requirements, acceptance bases, and explicit non-goals. Policy question answering is the core deliverable: the technology is mature and needs no external data, governed by a frozen corpus with a named owner and refusal of undeterminable questions. Data governance makes "which consultation ended in registration" answerable for the first time. Scenario estimation turns planning from extending totals into weighing intervals. Outreach answers the third dilemma at almost no cost and uses no algorithm at all, which tests the definition in Section 2: whether something counts as intelligent does not depend on whether a model is present.

Table 2. Requirements, Acceptance Criteria, and Non-Goals of the Minimum Deliverable Set

Deliverable	Core requirements	Acceptance basis	Explicitly not done
Policy Q and A	Corpus limited to approved texts; every substantive answer quotes its source; named owner; versions frozen at each release	Sampled citation accuracy; days the corpus lags policy	Eligibility determination; exception handling; admission promises
Data governance	Five record types—consultation, application, preference, admission, registration—linked at individual level; field owners; access control	Linkage rate across the five types; completeness of field-owner list	Academic, aid, and student-affairs data, which need top-level authorization
Scenario estimation	Intervals and shortfall risk from three years of cutoff scores, admission counts, registration counts	Interval coverage; yearly forecast deviation recorded	Provincial popularity forecasting; point estimates

Post-admission outreach	Reminders at three nodes—payment, documents, registration; escalation; item-by-item reasons	Outreach delivery rate; completeness of non-enrollment reasons	Automated screening; differential treatment

Note: None of the four requires provincial authorization, household background fields, or new positions; each is separately acceptable.

4.2. Division of Authority

A roadmap that includes the provincial platform, cross-institutional privacy computing, or algorithm review confuses categories: provinces build these, institutions only connect. Internal data governance and the advisory corpus are the institution's own to advance. Indicator definitions, data standards, and algorithm review are provincial, so the institution submits and documents. Application support is a public service that institutions supplement. Fairness audit has no dedicated provision, so institutions keep records and accept external sampling.

4.3. Sequencing and Acceptance

A ten-year, three-stage roadmap does not apply: pace follows funding, authorization, and data quality, not technology. Delivery means completing the four deliverables and passing acceptance; expansion starts when three of the four hold standard for two consecutive cycles; institutions with broken data may stay at the delivery stage for a long time. Access preparation—interface adaptation, format alignment, record keeping—costs little, has no completion date, and decides whether the institution can connect when provincial provision arrives (Figure 2).

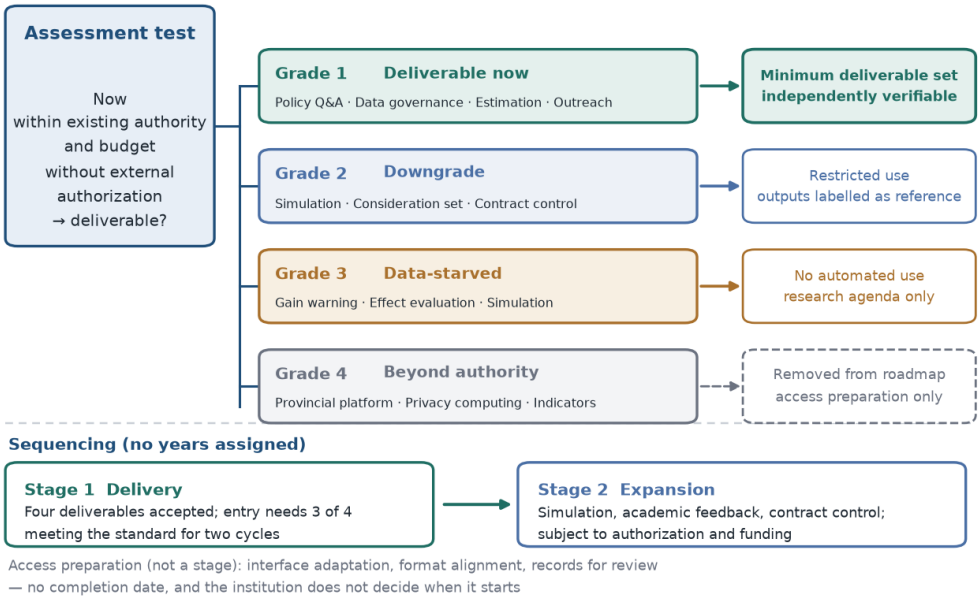


Figure 2. From Feasibility Assessment to a Minimum Deliverable Set (Note: No Years Are Given. Grade 1 Corresponds to the Delivery Stage and Grade 2 to Expansion, While Grades 3 and 4 Stay Outside the Institutional Roadmap.)

Acceptance standards must be fixed before go-live, or the built system writes them itself. Efficiency covers time to first response, plan completion rate, and cost per student; quality covers citation accuracy and refusal rate, reported with escalation logs. Fairness is checked, not computed: was reserved capacity delivered, were appeals accepted and corrected, how many days did the corpus lag policy. These checks prove that allocation was not driven by the model's preference, not that fairness was achieved.

5. Discussion and Conclusion

5.1. Fairness: From Ex Post Measurement to Ex Ante Rules

Personalization allocates attention by engagement, and engagement tracks family capital; advantaged students try harder while disadvantaged peers start behind in family, school, and effort [12]. The response is three rules rather than three quantities to measure. Capacity rules reserve part of service capacity, published before the season. Channel parity keeps personalized channels from becoming the only route to essential information. Appealability lets candidates demand explanation and review, and clustered appeals count as evidence of difference.

5.2. The Limits of Human Oversight

Keeping decisions with humans is not sentiment. Admission has legal effect and must be reproducible from records; a model that cannot explain itself in language a candidate could appeal against fails regardless of accuracy. Effective oversight needs four conditions: reviewers can see inputs and confidence values, have time and standing to object, objections are recorded and analyzed, and objecting costs nothing career-wise. At peak load, "escalate to a human" becomes "escalate to a queue" with a committed response time.

5.3. Future Prospects and Limitations

Components graded infeasible are not permanently unsolvable; each lacks one precondition. Data: the blocker is liability for data leaving the institution, not interfaces; once legislation settles cross-boundary liability, aggregate queries on a "usable but invisible" basis become possible. Algorithms: gain-based warning needs randomized outreach; a province could register staged outreach in advance within a special program and create comparable groups. Accountability: fairness monitoring collapsed because indicator fields do not exist; statistical compliance checks backed by external audit could bring ex post measurement back. All three preconditions are advanced by the state and the provinces, so the future prospect for an institution comes down to one item: stay connectable.

Three judgments close the paper. An intelligent admissions model is an institutional arrangement, and its typical failure is accumulating technology without allocating authority. Architecture lowers the cost of reform without replacing it. Only a small, technically modest part is deliverable now, and its value lies in verifiability, not novelty. The design rests on literature synthesis without implementation testing, so performance and cost remain empirical questions.

Acknowledgments: The author thanks the admissions staff whose practical questions shaped the feasibility assessment presented here.

Funding:

Disclosure statement:

Author contributions:

References

1. Ministry of Education of the People's Republic of China, "Statistical Bulletin on National Education Development in 2025," *Ministry of Education*, Beijing, 2026. [Online]. Available: https://www.moe.gov.cn/jyb_sjzl/sjzl_fztjgb/202607/t20260706_1442870.html
2. Central Committee of the Communist Party of China and State Council, *Outline of the Plan for Building a Leading Education Nation (2024–2035)*, People's Publishing House, Beijing, 2025.
3. H. A. Simon, "A Behavioral Model of Rational Choice," *The Quarterly Journal of Economics*, vol. 69, no. 1, pp. 99–118, 1955.
4. F. L. Li, S. P. Xu, and L. Wang, "Can the Diploma Itself Raise Individual Income? An Empirical Test of the Sheepskin Effect in China's Labor Market," *Exploring Education Development*, vol. 39, no. 5, pp. 39–45, 2019.
5. Q. X. Wu, "The Differences in Students' Access to College Admission Information and Crossing the 'Information Gap'," *Educational Research*, vol. 45, no. 6, pp. 53–62, 2024.

6. G. X. Chen and C. L. Wu, "The Composition, Problems and Solutions of the Data Governance System in Chinese Universities: An Activity Theory Perspective," *Journal of Higher Education Management*, vol. 17, no. 3, pp. 63–75, 2023.
7. A. R. Hevner, S. T. March, J. Park, et al., "Design Science in Information Systems Research," *MIS Quarterly*, vol. 28, no. 1, pp. 75–105, 2004.
8. Ministry of Education of the People's Republic of China, *Regulations on the Admission Work of Regular Higher Education Institutions*, *Ministry of Education*, Beijing, 2025.
9. Ministry of Education of the People's Republic of China, "Notice on the Admission Work of Regular Higher Education Institutions in 2026 (Jiaoji [2026] No. 1)," *Ministry of Education*, Beijing, 2026.
10. Standing Committee of the National People's Congress, *Personal Information Protection Law of the People's Republic of China*, Beijing, 2021.
11. C. L. Zhang, "The Ministry of Education's 'Ten Prohibitions' Safeguard Fairness and Justice in College Admissions," *Guangming Daily*, p. 03, Jun. 29, 2012.
12. H. X. Qin, F. Fang, and J. Zhou, "The New College Entrance Examination Reform and Access to High-quality Higher Education Opportunities," *Educational Research*, vol. 45, no. 6, pp. 28–41, 2024.

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.