

CampusNet: An AI-Augmented, Proximity-Preserving Social Platform for Campus Engagement

Aryan Borawake, Aditya Solanki, Parth Kad

Department of Computer Engineering, Smt. Kashibai Navale College of Engineering

Abstract — Active campuses already contain a dense fabric of communities — sanctioned clubs, self-organized interest groups, and informal peer mentorships — through which students acquire the knowledge and opportunity that drive their development. Yet this fabric is largely invisible: it is fragmented across messaging applications, noticeboards, and word-of-mouth, so the students who would benefit most cannot find it. General-purpose social, professional, and event platforms do not close this gap because their architectures, and the commercial incentives behind them, optimize for the global connection graph and thereby dissolve the campus proximity that gives near-community engagement its value. This paper presents CampusNet, a campus-bounded social platform that makes an institution's existing community fabric visible and accessible. The system is organized around a three-entity model — clubs, communities, and peer connections — each carrying a knowledge channel and an opportunity channel, and is bounded to a single institution by design rather than by optional filter. It integrates a hosted large-language-model service as a black-box inference layer for natural-language discovery, supports an end-to-end event lifecycle with cryptographically verified attendance, and provides institutional analytics. We describe the design principle, its operationalization, and a working deployed prototype demonstrating feasibility; we do not claim measured engagement outcomes, which a controlled study would establish.

Key Words: campus engagement, social platform, community of practice, proximity, large language models, systems design

1. INTRODUCTION

The developmental mission of a university extends well beyond its curriculum. A consistent body of empirical research associates participation in extracurricular and community life with gains in academic performance, sense of belonging, and longer-term career trajectories [1][2]. The capabilities students most often credit for their later success — leadership, navigating institutions,

recognizing opportunity, building relational capital — are acquired predominantly through the social and organizational life of the campus rather than through classroom instruction.

The mechanism behind these returns is worth stating precisely, because it shapes everything that follows. Communities deliver growth through two distinct channels. The first is knowledge — the mentorship, situated learning, and informal discussion through which tacit understanding transmits. The second is opportunity — the events, collaborations, and recruitment through which understanding is converted into experience. A campus is unusually well suited to carrying both: it is dense enough to sustain the close relationships through which guidance flows and broad enough to host the wide acquaintance network through which opportunity travels.

The problem this work addresses is not that these communities are absent but that they are invisible. On a typical engineering campus, the fabric exists — clubs run, interest groups form, seniors mentor juniors — but no digital infrastructure surfaces it. An interest group that has quietly won a national competition and now needs a developer remains unknown to the student who could join it; a senior with an exceptional placement remains unmet by the junior who would benefit from their guidance; the set of clubs a college actually contains is recorded nowhere a newcomer can simply consult. Both channels are throttled at once: knowledge cannot flow because the people carrying it cannot be found, and opportunity cannot flow because events and collaborations are imperceptible to those, they would most help. The gap falls hardest on first-year students and on those without pre-existing senior networks — precisely the students for whom the developmental returns are largest.

The objective of CampusNet is therefore narrow and concrete: to make an institution's already-existing community fabric visible and accessible to the students inside it, while preserving the campus proximity that gives that fabric its developmental value. This paper presents the design principle behind the system, the way that principle is operationalized into a working

architecture, and a deployed prototype that demonstrates the feasibility of the approach. Consistent with the project's scope, the paper reports a feasibility demonstration rather than a measured impact study; the latter is identified explicitly as the extension path.

2. LITERATURE REVIEW

Engagement and the two channels. The empirical association between structured engagement and student outcomes is well documented across multi-institution studies [1][2]. The mechanism resolves into the two channels named above, each with a distinct theoretical basis. The knowledge channel is best understood through the situated-learning tradition: Lave and Wenger characterize learning as legitimate peripheral participation, in which newcomers join at the edge of a community's practice and move toward its center as they take on more demanding roles, so that knowledge transmits as a by-product of shared activity rather than as a product of instruction [5]. The opportunity channel is best understood through network theory: Granovetter's analysis shows that novel information — including opportunity in the labor-market sense — propagates disproportionately through weak ties, the numerous lower-intensity connections that bridge otherwise separate clusters [3]. A platform intended to support development must therefore support both participation and weak-tie information flow, not merely provide a place to converse.

Proximity and boundedness. A student sits at the center of a field of communities at varying relational distance, and influence falls off with distance for structural reasons: spatial proximity makes help reachable in person; contextual proximity (shared curriculum, professors, calendar) lets advice transfer without translation; and social proximity (mutual acquaintance) shortens the path to the trust that collaboration requires. The supporting literature converges on the campus as the right unit of analysis. Tinto's interactional theory holds that a student's academic and social integration into the institution is a primary determinant of persistence and development [6]. Walther's work on computer-mediated communication shows that bounded environments, where participants share an institutional frame, exhibit participation dynamics measurably different from open networks — reduced distance and reduced social risk draw sustained participation from students who would self-exclude on open platforms [4]. A campus-bounded platform is thus not a narrower version of a global one but

a structurally different, and for this purpose superior, participation environment.

Why existing platforms do not close the gap. Each widely used platform addresses a fragment — professional networking, community chat, event ticketing, external opportunity aggregation — and none addresses campus engagement as a whole. The reason is structural: all of them dissolve the proximity that gives campus engagement its value, treating distance as a feature when, for this problem, distance is the defect. Beneath this lies an incentive structure that forbids the fix: the value of a for-profit social or professional network rises with the size and diversity of its global graph, so a platform that capped itself at one campus would forfeit the growth its model depends on. Even where campus boundaries are accommodated, they are an opt-in metadata facet over an architecture pulling the other way. An institutionally funded platform inverts this calculus — there is no external acquisition target and no engagement-scaled advertising inventory — so proximity preservation can become the operating principle rather than a constraint to work around.

Enabling technologies. Closing the visibility gap is made practical by recent technical work. Retrieval-augmented generation, formalized by Lewis et al., couples a retrieval component with a generative model so that reasoning is decoupled from corpus size [7]; this underpins natural-language discovery of entities and people. Structured team formation has been shown to be computationally hard in general [10] and is best approached through criteria-based filtering with human selection [11]; this underpins the collaboration channel. Governance research indicates that calibrated, lightweight governance can increase rather than restrict participation [12], with automated, rule-conditioned moderation [13][14] as a mature direction for later scale. Learning-analytics work establishes principles for decision-supporting dashboards [8] and for converting structured analytics into narrative summaries for non-technical audiences [9]. QR-code-based verification with signed payloads provides hardware-free, forgery-resistant attendance capture [15]. No single source describes the integrated system; the contribution here is synthesis under one principle.



Fig. 1. The proximity-first thesis: from student growth to a campus-bounded platform that makes an existing community fabric visible and accessible.

3. METHODOLOGY: SYSTEM DESIGN AND ARCHITECTURE

As a systems contribution, the methodology of this work is its design. We describe the design principle, the model that operationalizes it, the architecture that realizes it, and the genuine algorithms it contains.

3.1 Design principle and the three-entity model

The governing principle is proximity-first: the platform is bounded to a single institution by architecture, not by optional filter, so that every entity and interaction preserves the spatial, contextual, and social proximity established in Section 2 as the source of developmental value. This principle is operationalized through a three-entity model that mirrors the three relational scales at which campus communities form: clubs (officially sanctioned communities of practice), communities (self-organized interest groups, which the platform also lets students propose and create), and peer connections (one-to-one relationships). Each entity carries both channels: clubs surface a knowledge channel (discussion) and an opportunity channel (events and recruitment); communities surface discussion forums and collaboration requests; peer connections open direct messaging and collaboration once established. Visibility across all three is provided by natural-language discovery. The design target is therefore a complete three-entities-by-two-channels matrix, with discovery closing the visibility axis over the whole.

3.2 Layered architecture

The system follows a conventional three-layer separation (Fig. 2): a responsive single-page client; an application/API layer organized into functional modules (authentication and identity, posts and feed, clubs and recruitment, event lifecycle, communities and discussion, peer connection and messaging, search, analytics, administration); and a single document-oriented datastore. A hosted large-language-model service and the cloud hosting platforms are external dependencies outside the system boundary. Client-server interactions use HTTPS request/response; direct messaging between connected students uses a real-time WebSocket channel; the community general channel uses short-interval client polling — an honest simplification of the prototype rather than a uniform real-time fabric.

A deliberate identity-design decision underpins the model. The persisted user record stores a single role of administrator or student; the richer roles the interface depends on — fresher, club member, club president, community manager — are never stored but computed at request time from organizational state. This keeps the data model small and prevents the role-drift and denormalization bugs that storing such roles would invite.




Fig. 2. Layered system architecture: client, application/API, and data layers, with the hosted LLM service and cloud hosting as external dependencies.

3.3 The applied-AI discovery layer

Natural-language discovery is the platform's primary instrument of visibility. The system uses a hosted large-language-model service as a black-box inference layer —

not a model trained or fine-tuned by this work — a choice justified by limited labelled data and prototype scale, and made provider-agnostic so the model can be substituted without architectural change. The current pipeline gathers candidate clubs, communities, and students from the datastore, issues a single request to the service under a strict JSON contract, parses the ranked response defensively, and returns matches with one-line reasons. This realizes a simplified instance of the retrieval-augmented pattern [7], appropriate to a catalogue of tens to low-hundreds of entities, with a documented evolution toward two-stage embedding retrieval at larger scale (Section 6).

3.4 Algorithm details

Three components are genuine algorithms rather than routine data handling; the remainder of the system is deliberately not dressed up as such.

Algorithm 1 — Computed role evaluation. Given a student and current organizational state, the role is derived rather than read: the student is a president if any club names them president; otherwise a community manager if they manage a community; otherwise a member if they appear in a club's member set; otherwise a fresher. Presidency takes precedence, and the structural limits (at most one club membership, at most one managed community) are enforced at the write paths rather than by the reader.

Algorithm 2 — Idempotent event attendance via signed tickets. On registration to a free event, a ticket bearing an unguessable identifier is issued immediately; for a paid event no ticket is issued until an organizer verifies a submitted payment receipt, at which point the identifier is minted. At the venue the organizer scans the ticket: the first scan marks the registrant present with a timestamp; any re-scan reports the prior attendance without recording a second admission. Idempotency is the correctness property — re-scanning never double-counts — and a uniqueness constraint prevents duplicate registration.

Algorithm 3 — LLM-backed search pipeline. The query and a gathered candidate set are sent to the external service in a single call; the response passes through a three-layer extraction (direct parse, then fenced-code stripping, then brace extraction) with a graceful empty result on failure, yielding ranked matches with reasons. The algorithm's robustness lies in treating the external service as fallible and degrading cleanly rather than in any ranking computation of our own.

4. RESULTS: FEASIBILITY DEMONSTRATION

The result of this work is a complete, working, deployed system that demonstrates the feasibility of the proximity-first design. We report what the artifact is and that it functions; we do not report measured engagement outcomes, which would require the controlled study identified in Section 6.

Realized functionality. The deployed prototype implements the full three-entity model with both channels in each. Students authenticate against institution-supplied identities with a forced first-login password change and an administrative restriction path. The post-and-feed subsystem presents recruitment, event, general, and collaboration posts through a single model with deadline-aware ordering. Clubs support a complete recruitment workflow — drive, application, review by department and year, and accept/reject with automatic withdrawal of an accepted applicant's competing applications. The event subsystem realizes the full lifecycle of Algorithm 2, including the paid-event verification branch and QR attendance. Communities support proposal-and-approval, public and private joining, discussion threads, and collaboration requests. Peer connection gates direct messaging behind an accepted connection, after which a real-time conversation opens. Natural-language search returns ranked clubs, communities, and students with reasons. Institutional analytics serve both club leadership (membership, recruitment, events) and the campus administrator, whose view spans all communities and positions them to create opportunity where engagement is thin.

Verification. Because the prototype ships without an automated test suite, verification was manual and phase-gate-driven within an iterative development cycle: endpoints exercised in isolation, multi-step workflows verified end-to-end, and full user journeys smoke-tested on the deployed build against the specified requirements. This process surfaced and resolved several substantive defects whose fixes shaped the design — most notably a structural flaw in the original email scheme, which embedded the mutable study year and would therefore drift at each academic-year rollover; it was re-anchored to the immutable joining year, yielding a stable lifelong identifier and a dual-mode onboarding import. Cross-tab session isolation, authoritative attribute resolution, and resilient database connectivity were similarly hardened in response to testing findings. The artifact is deployed on free-tier cloud infrastructure and is operable end-to-end by the four actor classes.

Scope of the result. The demonstrated result is feasibility: the proximity-first principle and the three-entity model can be operationalized into a coherent, functioning system that closes the visibility-and-access gap for an institution. No claim is made regarding measured improvement in engagement, retention, or outcomes; the seeded prototype operated at the scale of a single cohort.

5. DISCUSSION

The contribution. The contribution of this work is threefold: a design principle (proximity-first engagement, in which institutional boundedness is an architectural commitment rather than a filter); its operationalization (the three-entity model with two channels each, plus computed roles and an applied-AI discovery layer); and a working artifact that demonstrates the principle is buildable and coherent. The synthesis is the novel element — the constituent ideas exist separately in the literature, but no prior source integrates engagement theory, communities-of-practice, proximity, and modern retrieval into a single campus-bounded platform whose central commitment is the very property the literature recommends and existing platforms erase.

Why boundedness is the right inversion. The discussion in Section 2 implies a design lesson that the prototype confirms in practice: the features that a for-profit platform cannot offer without harming its growth model — strict institutional bounding, no cross-campus pull, proximity preservation — are exactly the features an institutionally funded platform can adopt as first principles. The economic constraint that prevents commercial platforms from closing the gap is absent here, which is why an institutional platform is the natural home for the solution.

Honest limitations. Several limitations are intrinsic to the prototype stage and are reported plainly. First and most important, no controlled user study has been conducted; the paper demonstrates feasibility, not impact, and any claim of engagement improvement awaits that study. Second, the system runs at single-cohort scale on free-tier infrastructure, with a cold-start latency on first request that is a property of the tier rather than a defect. Third, several engineering choices are prototype-appropriate: discussion upvotes are not yet de-duplicated server-side, images are stored inline in the datastore, the community general channel uses polling rather than real-time push, and the session token is held in browser session storage with a documented migration path to a more secure cookie. Fourth, governance is entirely human-administered; automated, rule-conditioned moderation

[13][14] is deferred. None of these affects the feasibility result, and each has a defined evolution path.

Implications. For institutions, the result suggests that the engagement gap is closable with modest infrastructure and without commercial dependencies, and that the closing is principally a matter of visibility — surfacing a fabric that already exists — rather than of creating community from nothing. For system designers, it illustrates that a clearly articulated design principle (here, proximity preservation) can be carried consistently from theory through data model to deployed behavior, and that an external language model can serve as a discovery primitive without being the object of the research.

6. CONCLUSION

This paper presented CampusNet, a campus-bounded social platform that makes an institution's existing community fabric visible and accessible while preserving the proximity that gives that fabric its developmental value. The work contributes a design principle (proximity-first engagement), its operationalization into a three-entity, two-channel model with computed roles and an applied-AI discovery layer, and a working, deployed prototype that demonstrates the feasibility of the approach. The system is honest about its stage: it shows that the principle is buildable and coherent, not that it measurably improves outcomes.

The principal extension is a controlled evaluation — a deployment with a real student cohort and pre/post measures of engagement, discovery success, and belonging — which would convert the present feasibility claim into an empirical one and is the natural dissertation continuation. Technically, the discovery layer is designed to evolve into a two-stage architecture in which a local embedding model performs first-stage vector retrieval and the hosted language model operates only over the retrieved shortlist, decoupling discovery cost from catalogue size at institutional scale. Further documented paths include an automated regression suite with continuous integration, server-side upvote de-duplication, object storage for media, consolidation of the community channel onto the real-time transport, rule-conditioned assisted moderation, and a multi-institution deployment in which each campus remains its own bounded community while the platform hosts many — growth that multiplies bounded communities without widening any single boundary. Through every such path, the platform's central commitment is preserved: proximity, the principle the literature recommends and

commercial platforms erase, remains the operating commitment of the system at any scale.

ACKNOWLEDGEMENT

The authors thank their esteemed guide Dr. P.N. Railkar (professor at department of Computer Engineering, Smt. Kashibai Navale College of Engineering) for her guidance throughout this project.

REFERENCES

1. Nassar, S. et al.: Extracurricular Participation and Student Academic and Non-Academic Outcomes: A Multi-Institution Study. (2024)
2. University of California Institutional Research: Student Organization Participation and Academic Outcomes. (2023)
3. Granovetter, M.: The Strength of Weak Ties: A Network Theory Revisited. *Sociological Theory* 1 (1983) 201–233
4. Walther, J.B.: Computer-Mediated Communication: Impersonal, Interpersonal, and Hyperpersonal Interaction. *Communication Research* 23 (1996) 3–43
5. Lave, J., Wenger, E.: *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press (1991)
6. Tinto, V.: *Leaving College: Rethinking the Causes and Cures of Student Attrition*. 2nd edn. University of Chicago Press (1993)
7. Lewis, P. et al.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems* 33 (2020). arXiv:2005.11401
8. Pei, B. et al.: Design Principles for Learning Analytics Dashboards in Educational Contexts. (2024)
9. Kim, J. et al.: Large-Language-Model-Driven Learning Analytics and Narrative Reporting. (2024)
10. Anagnostopoulos, A. et al.: Online Team Formation in Social Networks. *Proc. WWW* (2012) 839–848
11. Awal, G.K., Bharadwaj, K.K.: Team Formation in Social Networks Based on Collective Intelligence — An Evolutionary Approach. *Applied Intelligence* 41 (2014) 627–648
12. Dwork, C. et al.: Content Moderation and the Formation of Online Communities. (2024)
13. Ribeiro, M.H. et al.: Automated Content Moderation Increases Adherence to Community Guidelines. *Proc. WWW* (2023)
14. Xin, Q. et al.: Rule-Conditioned Community Moderation. (2024)

15. QR-Code-Based Event Ticketing and Attendance Verification with Signed Payloads. (2024)