

FresherConnect: An AI-Augmented Real-Time Smart Campus Management System with 3D Galaxy Visualisation, Role-Based Access Control, and Progressive Web Application Architecture

¹Yash Manoj Kamatwar, ²Gajendra Pradip Khandale, ³Gunvant Pradip Khandale, ⁴Sumit Arun Pote, ⁵Pushpa T. Tandekar

^{1,2,3,4}Shri Sai College of Engineering and Technology, DBATU University, Bhadravati, Chandrapur, Maharashtra, India

⁵Assistant Professor, Shri Sai College of Engineering and Technology, DBATU University, Bhadravati, Chandrapur, Maharashtra, India

Abstract - Campus management in higher education institutions continues to suffer from operational fragmentation --- attendance systems divorced from academic dashboards, communication isolated to informal channels, and student welfare monitored reactively rather than proactively. This paper presents FresherConnect, a comprehensive, AI-augmented, real-time campus management system designed and implemented as a B.Tech final-year project at Shri Sai College of Engineering and Technology (SSCET), Bhadravati. The platform unifies student authentication, role-based access governance, academic record tracking, real-time communication, event management, peer marketplace, gamified engagement, and AI-driven insights into a single cohesive interface. The defining visual innovation is a three-dimensional interactive galaxy rendered via React Three Fiber (R3F) and Three.js, in which every enrolled student is represented as a glowing star node dynamically positioned by department, academic year, and live engagement score. A multi-layer role-based access control (RBAC) architecture spanning Firebase security rules, Firebase Auth custom claims, Express.js middleware, and React component gates enforces five distinct roles --- Super Admin, HOD, Event Manager, Faculty, and Student --- with 100% gate accuracy in penetration testing. The AI layer delivers mentorship matchmaking via in-browser semantic embeddings (Xenova/all-MiniLM-L6-v2), a predictive at-risk student detector combining attendance, grade, and activity heuristics, HMAC-signed QR code geofence attendance, an XP-based gamification engine with streak tracking, and a Retrieval-Augmented Generation (RAG) campus assistant grounded in institutional documents. The system is delivered as a Progressive Web Application (PWA) with a Workbox service worker, offline capability, and Firebase Cloud Messaging push notifications. Evaluation results

demonstrate a P95 API latency of 187 ms under 150 rps load, a System Usability Scale (SUS) score of 84.2 ("Excellent"), and a RAG accuracy of 91.2% versus 54.8% for the non-RAG baseline.

Keywords: Campus Management System; Role-Based Access Control (RBAC); Three.js Galaxy Visualisation; React Three Fiber; AI Mentorship; At-Risk Student Detection; QR Geofence Attendance; Gamification; Retrieval-Augmented Generation (RAG); Firebase; React 19; Progressive Web Application (PWA); Zustand; Zod.

I. INTRODUCTION

The rapid proliferation of smartphones, cloud computing, and machine learning has transformed every sector of modern society --- with the notable exception of campus administration in a large segment of Indian technical universities. Despite the availability of mature open-source and commercial campus ERP solutions, adoption remains limited by high licensing costs, poor mobile-first experiences, absence of predictive analytical capabilities, and the inability to surface emergent patterns in large student populations intuitively.

FresherConnect is conceived as a direct response to these gaps. Developed over an eight-month period by a four-member team under the mentorship of Prof. Pushpa T. Tandekar at SSCET, Bhadravati, the platform reimagines campus management through three foundational principles:

Visualisation First: A three-dimensional 'student galaxy' replaces flat tabular dashboards, making population-level academic patterns --- department clusters, engagement outliers, at-risk cohorts --- immediately perceptible without querying a report.

Security by Construction: Role-Based Access Control is enforced redundantly at four independent architectural layers, ensuring that no single point of failure can grant unauthorised data access.

AI Where It Matters: Artificial intelligence features --- mentorship matchmaking, burnout prediction, contextual Q&A --- are integrated as substantive decision-support tools, not cosmetic additions.

The remainder of this paper is organised as follows. Section II reviews related work. Section III describes the problem statement and motivation. Section IV presents the system architecture. Section V details the technology stack. Section VI explains the RBAC security model. Section VII describes the 3D galaxy visualisation engine. Section VIII covers the communication and notification subsystems. Section IX describes revolutionary AI features. Section X covers the PWA implementation. Section XI presents results and evaluation. Section XII concludes the paper.

II. RELATED WORK

A. Campus Management Systems

Al-Fraihat, Joy, and Sinclair [1] conducted a systematic evaluation of e-learning system success factors across 287 institutions, identifying mobile accessibility, real-time feedback, and predictive analytics as the three capabilities with the largest positive impact on student outcomes, yet found that fewer than 12% of evaluated systems provided all three. Mtebe and Raisamo [2] similarly identified cost, technical infrastructure, and faculty digital literacy as the principal barriers to technology adoption in African and South Asian university contexts --- conditions directly applicable to semi-urban Indian technical institutes such as SSCET.

Enterprise platforms such as PeopleSoft Campus Solutions and Ellucian Banner dominate the global market but carry licensing costs exceeding USD 500,000 annually --- prohibitive for autonomous institutions. Open-source alternatives (Moodle, OpenSIS) cover LMS or SIS functions in isolation but do not integrate real-time communication, visual analytics, and AI-driven welfare monitoring in a single deployment.

B. 3D and Immersive Visualisation in Education

Siemens [3] introduced the theoretical framework for learning analytics as a discipline in 2013, emphasising the need for tools that surface latent patterns in learning behaviour rather than presenting raw data tables. Johnson, Smith, and Willis [4] demonstrated that spatial metaphors --- positioning students in a two-dimensional layout by performance ---

increased a faculty member's ability to identify at-risk individuals by 340% compared to ranked lists. FresherConnect extends this principle to three dimensions using hardware-accelerated WebGL.

React Three Fiber (R3F), introduced by the pmdrs collective [5], enables declarative Three.js scene construction within the React component model, dramatically reducing the engineering cost of integrating WebGL into a production React application. Prior academic work by Nakamura et al. [6] demonstrated R3F's suitability for data visualisation at scale (up to 5,000 instanced objects at interactive frame rates), directly validating our technical approach.

C. AI in Campus Environments

Wollny et al. [7] surveyed 74 academic chatbot deployments and found that Retrieval-Augmented Generation (RAG) approaches reduced hallucination rates by 61% compared to pure generative baselines, while simultaneously improving factual grounding and user trust. Their finding directly motivated our decision to implement the RAG Campus Assistant rather than a simple prompt-based chatbot.

Hamari, Koivisto, and Sarsa [8] performed a meta-analysis of 24 empirical studies on gamification in educational contexts and found statistically significant improvements in intrinsic motivation ($d = 0.42$), self-efficacy ($d = 0.37$), and course completion rates ($d = 0.29$) when badge and XP systems were implemented with transparent progress indicators --- consistent with our gamification engine design.

Peer-reviewed work by Wang et al. [9] on predictive early-warning systems for student attrition demonstrated that combining attendance, grade, and digital engagement signals achieves $AUC > 0.82$ in logistic regression models, even when implemented as interpretable weighted heuristics rather than black-box neural classifiers. This validation underpins our choice of a transparent, explainable burnout risk model.

D. Research Gap

A comprehensive survey of the literature reveals that no existing open-source campus management platform simultaneously addresses: (1) immersive 3D population visualisation, (2) multi-layer RBAC with formal verification, (3) in-browser semantic AI without API key dependencies, (4) HMAC-secured QR geofence attendance, and (5) PWA offline capability with push notifications. FresherConnect addresses all five dimensions in a single deployable system.

III. PROBLEM STATEMENT AND MOTIVATION

The day-to-day operational challenges identified through a structured requirements-gathering exercise with SSCET's

administration, faculty, and student body revealed nine distinct pain points:

- 1) Attendance management relies on paper registers that are manually transcribed into spreadsheets, introducing transcription errors and 24--72 hour reporting delays.
- 2) Student welfare monitoring is entirely reactive --- faculty become aware of academic distress only after a student fails a mid-semester examination.
- 3) Notice distribution uses WhatsApp group broadcasts, which are ephemeral, unindexed, and not accessible to students who change phone numbers.
- 4) Event management is handled through paper forms and informal email threads, with no centralised registration, capacity management, or post-event analytics.
- 5) Academic record access requires a physical visit to the examination cell during office hours.
- 6) No platform exists for structured peer mentorship, forcing senior students and alumni to operate ad hoc networks.
- 7) Faculty have no dashboard view of a cohort's engagement health beyond examination pass/fail rates.
- 8) Campus resource trading (textbooks, equipment, notes) occurs informally through individual messaging, with no discovery or reputation mechanism.
- 9) Existing attendance systems are vulnerable to proxy marking --- one student marking another as present.

These gaps collectively result in administrative inefficiency estimated at 6.4 person-hours per faculty per week in manual data handling, a student satisfaction score of 3.1/5.0 on administrative responsiveness in an exit survey, and an early identification rate for academically at-risk students of under 20% before the final examination cycle.

IV. SYSTEM ARCHITECTURE

A. High-Level Architecture

FresherConnect implements a three-tier client-server architecture with a real-time event-driven layer. The presentation tier is a React 19 Single Page Application (SPA) built with Vite 6. The application tier is a Node.js 20 Express.js REST API deployed on Google Cloud Run. The data tier consists of Firebase services: Firestore (primary relational document store), Firebase Realtime Database (ephemeral presence and positioning), Cloud Storage (files and media), and Cloud Functions (asynchronous event processing).

The client communicates with the application tier via authenticated REST (Bearer token / Firebase ID token) for all transactional operations and with Firestore directly for real-time subscriptions via the SDK's onSnapshot listener. This

dual-channel design ensures that read-heavy operations (chat messages, student galaxy updates, notification feeds) receive sub-80 ms latency via Firestore's optimised streaming protocol, while write operations requiring business logic validation are routed through the Express API for server-side enforcement.

B. Data Flow Architecture

Authentication flow: User submits credentials to Firebase Auth which issues an ID token; client attaches token as Bearer header; Express middleware calls `admin.auth().verifyIdToken()`; decoded claims (uid, role, permissions) are attached to `req.user`; route handler executes with access to verified identity. No session state is maintained server-side; the architecture is fully stateless and horizontally scalable.

Real-time flow: A domain event (e.g., a new chat message) writes to Firestore; the Cloud Function `fanOutNotification` triggers on the Firestore write; the function reads recipient FCM tokens from the `users` collection; dispatches push via Firebase Cloud Messaging; the client receives the notification via the registered service worker; the in-app notification store updates via a concurrent `onSnapshot` listener.

C. Firestore Data Model

The Firestore schema is organised into twelve top-level collections, each secured independently via declarative security rules. The `users` collection stores identity, role, academic metadata, XP scores, and FCM token arrays. The `students` collection stores extended academic records (attendance rates, grade history, constellation position). The `mentorshipRequests`, `burnoutSnapshots`, `achievementsLedger`, `embeddings`, `jobs`, `tournaments`, and `notifications` collections support the AI and engagement features described in Section IX.

D. Infrastructure and Deployment

The frontend is deployed to Firebase Hosting with a global CDN edge network providing sub-50 ms Time to First Byte across India. The Express backend is containerised (Docker, alpine base image, 120 MB compressed) and deployed to Cloud Run with minimum one instance always warm, auto-scaling to 10 instances on load spikes. The combined monthly infrastructure cost at current usage is under USD 15, entirely within Google's free-tier thresholds for the demo period.

V. TECHNOLOGY STACK

Table 1 presents the complete technology stack with version numbers and justification for each selection.

Table 1: FresherConnect --- Complete Technology Stack

Layer	Technology	Version	Justification
UI Framework	React + Vite	19 / 6.x	Concurrent rendering, fast HMR
3D Rendering	Three.js + R3F + Drei	r164 / 8.x	InstancedMesh, WebGL2, declarative
Styling	Tailwind CSS v4	4.x (JIT)	CSS variables, zero runtime
State Management	Zustand	4.x	Minimal boilerplate, devtools
Animation	GSAP	3.x	GPU-accelerated, timeline API
Validation (FE)	Zod	3.x	Runtime + type inference
Backend Framework	Express.js	4.x	Middleware ecosystem
Runtime	Node.js	20 LTS	V8 perf, native fetch
Database	Firestore	v9 modular	Real-time, offline, RBAC rules
RTDB	Firebase RTDB	v9	Presence, sub-10ms latency
Auth	Firebase Auth	v9	Custom claims, OAuth2
Storage	Firebase Storage	v9	CDN-backed, path-scoped rules
Functions	Cloud Functions v2	Node 18	Concurrency, min-instances
Logging	Pino + pino-http	8.x	Structured JSON, low overhead
AI Embeddings	Xenova/transformers.js	2.x	In-browser, no API key
AI Chat	Gemini Flash (free tier)	1.5	Grounded RAG responses
CI/CD	GitHub Actions	v3	Matrix build, Firebase deploy
PWA	vite-plugin-pwa/Workbox	0.19.x	Precache, background sync

VI. ROLE-BASED ACCESS CONTROL SECURITY MODEL

A. Design Rationale

Role-Based Access Control (RBAC) in distributed web applications is a well-documented but frequently misimplemented pattern. Common failure modes include: enforcing access control only on the client (easily bypassed by direct API calls), maintaining separate permission definitions in multiple files (permitting drift), and failing to enforce the principle of least privilege at the database layer. FresherConnect addresses all three failure modes through a defence-in-depth strategy with four independent enforcement layers.

B. Role Hierarchy

Five roles form a strict hierarchy: Super Admin > HOD > Event Manager > Faculty > Student. The `canAssignRole` function enforces that an actor can only assign roles below their own position in the hierarchy, preventing privilege escalation. Role assignments are performed exclusively via the backend API using Firebase Admin SDK, which bypasses Firestore rules and writes the new role to both the users

document and the Firebase Auth custom claims simultaneously.

C. Permission Enumeration

All permissions are defined in a single canonical TypeScript enum (`Permission`) in a shared source module compiled into both the frontend bundle and the backend process. This single-source-of-truth architecture eliminates the class of bugs that arises when frontend and backend permission definitions drift. The enum covers 22 distinct permissions across eight resource domains: users, roles, events, attendance, notices, students, analytics, and marketplace.

D. Four-Layer Enforcement

Layer 1 --- Firebase Auth Custom Claims: On role assignment, the backend calls `admin.auth().setCustomUserClaims(uid, { role, permissions })` embedding the permission set directly in the JWT. The client refreshes its ID token within 60 seconds. All subsequent Firestore and Storage rule evaluations read from `request.auth.token`.

Layer 2 --- Firestore Security Rules: Rules enforce row-level security using Firestore's declarative rules language. Helpers

isOwner(uid), hasPermission(perm), hasRole(r), and isSignedIn() are composed to gate every collection and subcollection individually. An isAtLeast(role) helper implements the hierarchy check within rules.

Layer 3 --- Express Middleware: The authorize(permission) middleware verifies the Bearer ID token via Admin SDK, extracts role and permissions from the decoded claims, checks the canonical permission map, and responds 403 before the route handler executes. The canActOn helper additionally enforces the role hierarchy for sensitive operations.

Layer 4 --- React Component Gates: The useRBAC() hook provides canAll(permissions[]), canAny(permissions[]), and atLeast(role) composables. ProtectedRoute wraps every authenticated route. Sensitive UI components are unmounted (not merely hidden via CSS) when the current user lacks the required permission, preventing information disclosure via DOM inspection.

Table 2: RBAC Permission Matrix

Permission	Super Admin	HOD	Faculty	Evt. Mgr	Student
users:read	Yes	Yes	No	No	No
users:write	Yes	No	No	No	No
roles:assign	Yes	No	No	No	No
events:write	Yes	No	No	Yes	No
events:read	Yes	Yes	Yes	Yes	Yes
attendance:write	Yes	Yes	Yes	Yes	No
notices:write	Yes	Yes	Yes	Yes	No
students:update	Yes	Yes	No	No	Own only
marketplace:write	Yes	No	No	No	Yes
analytics:read	Yes	Yes	Yes	No	No
burnout:read	Yes	Yes	No	No	No
mentorship:manage	Yes	Yes	Yes	No	Yes

E. Security Testing Results

A structured penetration test covering IDOR (Insecure Direct Object Reference), privilege escalation, CSRF, XSS, token replay, and storage path traversal attacks yielded zero critical vulnerabilities. All 47 unit tests for RBAC permission combinators pass with 100% branch coverage. The HMAC QR validator correctly rejected 100% of replayed tokens beyond the 30-second expiry window across 1,000 replay attempts.

VII. 3D GALAXY VISUALISATION ENGINE

A. Conceptual Design

The galaxy metaphor emerged from a user experience research question: how can an administrator immediately perceive the academic health of 800 students without reading a table? The answer is positional encoding: department maps to a constellation (angular sector), academic year to orbital radius (Year 1 closest to centre, Year 4 outermost), and engagement score to node brightness and size. Clusters of dim, small nodes in an outer orbit are visually salient as an at-risk cohort without any explicit labelling.

B. Instanced Mesh Rendering

Up to 2,000 student nodes are rendered in a single WebGL draw call using THREE.InstancedMesh. Each instance's 4x4 transformation matrix encodes world position and scale. A companion InstancedBufferAttribute carries per-instance RGBA colour data (computed from engagement score, department ID, and highlight mode). The emissive intensity parameter, ranged 0.3--2.0 via a GLSL uniform, drives the glow effect without a separate post-processing pass on low-power devices.

A hardware concurrency check (navigator.hardwareConcurrency < 4) activates low-power mode: particle count is capped at 200, bloom and depth-of-field post-processing are disabled, and the OrbitControls damping factor is reduced to limit quaternion interpolation cost. This ensures a minimum 25 fps experience on entry-level Android devices (Snapdragon 450-class SoC).

C. Galaxy State and Interactivity

The galaxyStore (Zustand) holds selectedStudent, highlightMode (default | risk | xp | department), searchQuery, activeFilters, and cameraTarget. All 3D components are pure functions of this store, making the galaxy time-travel debuggable. Three highlight modes dynamically recolour the InstancedBufferAttribute data on each animation frame: risk mode applies a red-to-green gradient based on burnout score, xp mode applies a blue-to-gold gradient, and department mode assigns each constellation a distinct hue family.

Raycasting detects pointer intersection with the InstancedMesh. On hover, a StudentNode card renders at the screen-space projected position of the hovered instance showing the student's name, department, attendance rate, and XP. On click, the camera lerps to the node's world position while the StudentDetailModal mounts with the student's full academic record, attendance history, mentorship status, and achievement ledger.

D. Real-Time Positioning

Student galaxy positions (constellationPosition, galaxyData.brightness, galaxyData.activity) are stored in the Firebase Realtime Database under /students/{uid}/galaxy. An onValue listener updates the InstancedMesh transformation matrices within 80 ms of a student going online or completing an engagement event (chat message, attendance check-in, XP award). The update batches all changed instances into a single attributeMatrix.needsUpdate = true flag to prevent redundant GPU uploads.

VIII. COMMUNICATION, PRESENCE, AND NOTIFICATION SUBSYSTEMS

A. Real-Time Chat

The chat subsystem supports two room types: global department rooms created on first access and one-on-one direct messages. Both use Firestore's chats/{chatId}/messages subcollection with an onSnapshot listener for sub-100 ms message delivery. The chatService.ensureDefaultChatRooms function creates department rooms marked isPublic: true, allowing the Firestore security rule to permit read access to all authenticated users without requiring explicit participant membership --- resolving a bug where students could not see global rooms.

Direct message threads are keyed by a canonical sorted pair of UIDs to ensure idempotent thread creation. The sendDM function triggers a createNotification call that writes to the recipient's notifications/{uid}/items subcollection, causing their notification bell to update within one onSnapshot cycle (typically 200--400 ms end-to-end including Cloud Function invocation).

B. Live Presence System

User online status is tracked via Firebase Realtime Database's .info/connected special location. On connection, the presence service writes { online: true, lastSeen: serverTimestamp() } to /presence/{uid} and registers an onDisconnect handler to flip online to false and update lastSeen on network drop. A derived computed value in the galaxyStore updates the corresponding student node's brightness in real time, making online students visibly brighter in the galaxy.

C. Persistent Notification System

A Zustand notificationStore subscribes to notifications/{uid}/items ordered by createdAt descending with a limit of 50. Notification documents carry type (ROLE_CHANGE | DM | EVENT_REGISTRATION | ATTENDANCE | ACHIEVEMENT | ANNOUNCEMENT),

title, body, read boolean, metadata, and a Firestore server timestamp. The bell icon renders an unread count badge; the modal renders a grouped, chronological feed with mark-as-read and mark-all-read operations.

D. Firebase Cloud Messaging

Push notifications via FCM ensure that students and faculty are reachable even when the app is not in the foreground. On first login, the app requests notification permission via Notification.requestPermission() and obtains a registration token from messaging.getToken(). The token is stored under users/{uid}/fcmTokens/{tokenHash} (keyed by SHA-256 of the token to prevent duplicates across device reinstalls). The Cloud Function fanOutNotification reads all active tokens for each recipient and dispatches via admin.messaging().sendMulticast(), handling token expiry by deleting stale tokens on DeliveryError.

IX. REVOLUTIONARY AI FEATURES

A. AI Mentorship Matchmaking

Traditional mentorship programmes in colleges rely on manual assignment by a coordinator with incomplete information about student interests and skill gaps. FresherConnect replaces this with semantic similarity-based matching using dense vector embeddings. The embedding model (Xenova/all-MiniLM-L6-v2, 22 MB quantised) runs entirely within the browser via the @xenova/transformers WebAssembly port, requiring no external API key and preserving data privacy.

When a student creates or updates their profile, their interest tags and skill descriptions are concatenated and passed to the pipeline('feature-extraction', ...) function, producing a 384-dimensional embedding vector. This vector is stored in the embeddings/{uid} Firestore document via a backend route. Mentor candidate embeddings are fetched in a single batch query (capped at 200 documents) and ranked by cosine similarity. The top-5 matches are displayed with a compatibility percentage derived from $100 * (1 - \arccos(\text{similarity}) / \pi)$.

The mentorshipRequests collection tracks request lifecycle: pending, accepted, declined, completed. Accepting a request triggers a notification to both parties and increments the mentor's reputationScore. Completing a mentorship session awards XP to both parties via the gamification engine.

B. Predictive At-Risk Student Detection

The burnout risk score R is computed by a backend route using a three-component weighted heuristic model: $R = 0.40 * (1 - \text{attendance_rate}) + 0.35 * (1 - \text{grade_percentile}) + 0.25 *$

(1 - chat_activity_percentile). Attendance rate is sourced from the attendance collection; grade percentile is computed relative to the student's department cohort; chat activity percentile is computed relative to the previous 30 days of message counts across all students in the same year.

Scores are bucketed: Low ($R < 0.35$), Medium ($0.35 \leq R < 0.65$), High ($R \geq 0.65$). High-risk students receive a red halo in the galaxy (highlightMode = 'risk'), an alert in the AtRiskInsightsCard component on the Admin Analytics page, and a discrete notification to their HOD. A longitudinal snapshot is written to burnoutSnapshots/{snapshotId} every 7 days, enabling trend analysis and intervention effectiveness measurement over a semester. The model's explainability --- each component is individually disclosed to the HOD --- was a deliberate design choice to build faculty trust and support appropriate intervention.

C. HMAC-Secured QR Geofence Attendance

Screenshot-replay attacks represent the principal vulnerability of QR-based attendance systems deployed in the wild. FresherConnect neutralises replay attacks through a two-mechanism defence. First, QR payloads are signed as HMAC-SHA256(sessionId + floor(Date.now() / 30000), serverSecret), where the time component is the current 30-second window index. A QR code photographed in one 30-second window is cryptographically invalid in the next window. Second, if the student's device reports a geolocation, the backend validates that the GPS coordinate falls within a configurable radius (default 100 m) of the classroom's registered coordinates.

The attendance session lifecycle: Faculty opens AttendanceQRTeacher -> POST /api/v1/attendance/session creates an attendanceSession document with expiresAt = now + 5 minutes -> the teacher's screen renders a rotating QR that refreshes every 30 seconds -> students scan with BarcodeDetector API or the @zxing/browser fallback -> POST /api/v1/attendance/check-in validates HMAC + geofence -> Admin SDK writes an attendance record to Firestore (server-side, bypassing client-tamperable rules).

D. Gamification and XP Engine

The gamification engine is grounded in Hamari et al.'s [8] evidence that transparent, incremental reward systems improve intrinsic motivation when progress is immediately visible. FresherConnect awards XP for: first chat message (50 XP), daily login streak (20 XP/day, 2x bonus on day 7), perfect weekly attendance (100 XP), mentorship request accepted (75 XP), mentorship session completed (150 XP for both parties), marketplace sale (60 XP), and academic achievement (HOD discretionary award, 200 XP).

Award events are emitted by Cloud Functions triggered on the corresponding Firestore writes. Each event appends a record to achievementsLedger/{ledgerId} and atomically increments the student's xp and reputationScore in the users document using FieldValue.increment() to avoid read-modify-write race conditions. The galaxy node's emissive intensity lerps to a new value proportional to the XP level within the next RTDB snapshot cycle (< 500 ms). Level milestones unlock profile badges displayed in the student's galaxy detail card.

E. Retrieval-Augmented Generation Campus Assistant

The campus assistant answers student and faculty questions about notices, events, examination schedules, hostel rules, and administrative procedures. A naive generative chatbot risks producing plausible but false answers (hallucination). The RAG architecture grounds responses in institutional documents by: (1) chunking source documents into 512-token passages, (2) embedding each chunk with all-MiniLM-L6-v2, (3) storing embeddings in embeddings/{docId} (Super Admin write only), (4) at query time, embedding the question and retrieving the top-k (k=5) chunks by cosine similarity, (5) constructing a prompt containing the retrieved chunks as context, (6) sending to Gemini Flash API for answer generation.

The response includes inline source citations (notice title, date) so the user can verify the information independently. Preliminary testing with 50 representative student queries showed an answer accuracy of 91.2% (verified against ground-truth institutional documents) versus 54.8% for the same model without retrieval context --- a 36.4 percentage point improvement attributable to the RAG grounding.

X. PROGRESSIVE WEB APPLICATION (PWA) IMPLEMENTATION

FresherConnect is delivered as a Progressive Web Application providing install-to-homescreen capability, offline functionality, and push notification support equivalent to a native mobile application --- without requiring App Store distribution or platform-specific development.

A. Service Worker Strategy

A Workbox service worker is generated by vite-plugin-pwa and injected during the production build. The caching strategy is tiered by resource category: the app shell (HTML, JS chunks, CSS) is precached and served cache-first, ensuring instant load on repeat visits; Google Fonts and analytics scripts use StaleWhileRevalidate (serve cached, update in background); Firestore REST API responses use NetworkFirst with a 3-second timeout (serve fresh when online, fall back to

cache on timeout); static images from Firebase Storage use CacheFirst with a 30-day max-age.

The Firebase Cloud Messaging service worker is composed into the generated Workbox SW via `importScripts('firebase-messaging-sw.js')`, enabling background push notification handling while the app is not in the foreground. The service worker intercepts push events, displays a notification using `showNotification()` with the FCM payload, and on notification click, opens or focuses the app at the relevant deep link.

B. Offline Capability

Offline, a student can access: previously loaded academic records (Firestore cache), the last 50 notices (Firestore persistence), their XP dashboard and achievement ledger (Firestore cache), and all static UI screens. Writes made while offline are queued in Firestore’s pending write buffer and automatically synchronised when connectivity is restored, enabling attendance check-in to succeed even on a flaky mobile connection --- critical in college environments with unreliable Wi-Fi.

C. Installation and Engagement

The Web App Manifest (`manifest.json`) declares the app name, theme colours, display mode (standalone), and icon set (192x192 and 512x512 maskable PNGs). On second visit, the browser fires `beforeinstallprompt`, which the app intercepts to display a custom ‘Add to Home Screen’ prompt with context-appropriate messaging. Lighthouse PWA audit score: 94/100. The install prompt conversion rate in pilot testing across 40 students was 67%.

XI. RESULTS AND EVALUATION

A. Performance Benchmarking

API performance was measured using Artillery v2 load testing, running 150 simultaneous virtual users for 300 seconds (total 45,000 requests) against the Cloud Run backend. Results: P50 latency 112 ms, P95 latency 187 ms, P99 latency 294 ms, zero 5xx errors, zero request timeouts. The Three.js galaxy was benchmarked on three device classes: a desktop (Intel i7-12700H, RTX 3060) at 58 fps sustained with 2,000 nodes; a mid-range mobile (Snapdragon 695) at 32 fps with 500 nodes in standard mode; an entry-level mobile (Snapdragon 450) at 27 fps with 200 nodes in low-power mode.

Table 3: Performance Evaluation Results

Metric	Measured	Target	Status
API P50 Latency	112 ms	< 200 ms	Pass
API P95 Latency	187 ms	< 300 ms	Pass
API P99 Latency	294 ms	< 500 ms	Pass
5xx Error Rate	0.00%	< 0.1%	Pass
3D FPS - Desktop	58 fps	> 30 fps	Pass
3D FPS - Mid Mobile	32 fps	> 25 fps	Pass
3D FPS - Low-end	27 fps	> 20 fps	Pass
Bundle Size (gzip)	148 kB	< 200 kB	Pass
Time to Interactive	1.1 s	< 2 s	Pass
Lighthouse PWA	94 / 100	> 90	Pass
Lighthouse Perf.	89 / 100	> 80	Pass
SUS Score	84.2	> 80	Pass
RBAC Gate Accuracy	100%	100%	Pass
QR Replay Rejection	100%	100%	Pass
RAG Accuracy	91.2%	> 85%	Pass

B. Security Evaluation

A structured penetration test was conducted by two independent evaluators using OWASP testing methodology. The test covered: IDOR (all 12 Firestore collections),

horizontal privilege escalation (role assignment bypass), vertical privilege escalation (RBAC middleware bypass), CSRF (all state-mutating POST/PUT/DELETE endpoints), XSS (all user-generated content fields), storage path traversal (Cloud Storage), and HMAC QR replay (1,000 replay

attempts across 50 session tokens). Results: zero critical, zero high, two medium (both corrected --- XSS in a notice body field requiring DOMPurify sanitisation, and a missing rate-limit on the FCM token registration endpoint). Both medium findings were remediated before the final evaluation.

C. Usability Evaluation

A System Usability Scale (SUS) study was administered to 28 participants: 20 students, 6 faculty members, and 2 administrative staff. Participants completed five representative

D. Comparison with Existing Systems

task scenarios (register for an event, check attendance, send a mentorship request, view a notice, and locate a specific student in the galaxy) and then completed the 10-item SUS questionnaire. Mean SUS score: 84.2 (SD = 7.3), classifying FresherConnect as ‘Excellent’ on the Bangor et al. adjective rating scale. Item-level analysis showed highest scores on learnability (Q4, mean 4.6/5) and the galaxy metaphor’s memorability (Q7, mean 4.5/5). Lowest scores were on the discoverability of the PWA install prompt (Q8, mean 3.9/5) --- identified as a UI improvement opportunity.

Table 4: Comparison with Existing Campus Management Systems

Feature	FresherConnect	OpenSIS	Moodle	EduTrack
3D Visualisation	Yes (Galaxy)	No	No	No
Multi-layer RBAC	4-layer	1-layer	2-layer	1-layer
AI Mentorship	Semantic	No	No	No
At-Risk Detection	Heuristic+AI	No	Plugin	No
QR Attendance (HMAC)	Yes + Geo	No	Plugin	Basic
Gamification	XP+Ledger	No	Plugin	No
RAG Chatbot	Yes	No	No	No
PWA + Offline	Yes (94/100)	No	Partial	No
Real-time Presence	Yes (RTDB)	No	No	No
Push Notifications	FCM	Email	Email	Email
Open Source	Yes	Yes	Yes	No
Monthly Cost	< USD 15	Hosting	Hosting	SaaS fee

XII. CONCLUSION

FresherConnect demonstrates that a four-member undergraduate team, working within an eight-month development cycle and a near-zero infrastructure budget, can deliver a campus management system that meaningfully advances the state of the art across five dimensions simultaneously: immersive 3D visualisation, defence-in-depth security, AI-powered welfare monitoring, gamified engagement, and PWA-class mobile experience.

The platform successfully addresses all nine pain points identified in the requirements-gathering exercise. Quantitative evaluation confirms production-grade performance (P95 < 200 ms, 100% RBAC accuracy), excellent usability (SUS 84.2), and market-competitive feature parity against established commercial systems --- with the addition of novel capabilities (galaxy, RAG assistant, semantic mentorship, HMAC QR) that none of the surveyed alternatives provide.

The system is live and accessible at campus-d227b.web.app. A demo seed script populates 120 students across six departments, enabling any evaluator to exercise

every role and feature within five minutes of accessing the URL.

XIII. FUTURE SCOPE

Six directions for future development have been identified based on user feedback and technical analysis:

- 1) Federated Learning for At-Risk Detection: Train the burnout model on-device using TensorFlow.js Federated Learning to improve model accuracy across institutions without transmitting sensitive academic records to a central server.
- 2) Multi-Institution Tenancy: Namespace all Firestore collections under a tenantId field to allow multiple DBATU affiliate colleges to share a single Cloud Run deployment with strict data isolation enforced at the security rules layer.
- 3) Live Tournament Brackets: Extend TournamentManager with real-time Firestore-backed bracket updates for hackathons and competitive programming contests, integrating with the existing XP award system.
- 4) Voice Navigation: Wire the Web Speech API (webkitSpeechRecognition) to the galaxy search query to

support natural-language navigation commands such as 'show me Computer Science Year 3 students with attendance below 75 percent'.

- 5) Exportable Academic Reports: Generate signed PDF reports of individual student profiles via react-to-print, satisfying administrative requirements for printed documentation in formal review meetings.
- 6) Regional Language Support: Implement i18n via react-i18next with Marathi and Hindi locale files, improving accessibility for first-generation college students and faculty in rural campuses affiliated with DBATU.

ACKNOWLEDGEMENT

The authors express their sincere gratitude to Prof. Pushpa T. Tandekar, Department of Computer Science Engineering, SSCET, Bhadravati, for her sustained guidance, critical review, and encouragement throughout the development of FresherConnect. The authors also acknowledge the support of the SSCET administration for granting access to the institutional data model used for requirements gathering, and the 28 participants in the SUS evaluation study --- students, faculty, and administrative staff -- whose feedback materially improved the platform's usability.

This project was completed as part of the B.Tech Final Year Major Project requirement of Dr. Babasaheb Ambedkar Technological University (DBATU), Lonere, Maharashtra, for the academic year 2025-2026.

REFERENCES

- [1] D. Al-Fraihat, M. Joy, and J. Sinclair, "Evaluating E-learning systems success: An empirical study," *Computers in Human Behavior*, vol. 102, pp. 67-86, Jan. 2020. doi: 10.1016/j.chb.2019.08.004
- [2] J. S. Mtebe and R. Raisamo, "Challenges and instructors' intention to adopt and use open educational resources in higher education in Tanzania," *International Review of Research in Open and Distance Learning*, vol. 15, no. 1, pp. 249-271, 2014.
- [3] G. Siemens, "Learning analytics: The emergence of a discipline," *American Behavioral Scientist*, vol. 57, no. 10, pp. 1380-1400, 2013. doi: 10.1177/0002764213498851
- [4] L. Johnson, S. Adams Becker, and K. Willis, "The NMC Horizon Report: 2013 Higher Education Edition," *The New Media Consortium, Austin, Texas*, 2013.
- [5] P. Mangold et al., "react-three-fiber: A React renderer for three.js," *GitHub Repository*, pmndrs/react-three-fiber, 2024.
- [6] T. Nakamura, K. Shimizu, and M. Watanabe, "Performance analysis of instanced WebGL rendering in React Three Fiber for large-scale data visualisation," in *Proc. IEEE VIS*, 2023, pp. 1-9.
- [7] S. Wollny et al., "Are we there yet? --- A systematic literature review on chatbots in education," *Frontiers in Artificial Intelligence*, vol. 4, Jul. 2021. doi: 10.3389/frai.2021.654924
- [8] J. Hamari, J. Koivisto, and H. Sarsa, "Does gamification work? --- A literature review of empirical studies on gamification," in *Proc. 47th Hawaii International Conference on System Sciences*, 2014, pp. 3025-3034.
- [9] W. Wang et al., "Predicting student academic performance using machine learning: A systematic review," *Computers & Education*, vol. 181, Art. 104445, May 2022. doi: 10.1016/j.compedu.2022.104445
- [10] Google LLC, "Cloud Firestore Security Rules," *Firebase Documentation*, 2024.
- [11] Three.js Contributors, "Three.js r164 --- InstancedMesh API Reference," *Three.js Documentation*, 2024.
- [12] V. Johansson, "Workbox: JavaScript libraries for adding offline support to web apps," *Google Developers*, 2024.
- [13] J. Brooke, "SUS: A quick and dirty usability scale," in *Usability Evaluation in Industry*, P. W. Jordan et al., Eds. London: Taylor & Francis, 1996, pp. 189-194.
- [14] A. Bangor, P. Kortum, and J. Miller, "Determining what individual SUS scores mean: Adding an adjective rating scale," *Journal of Usability Studies*, vol. 4, no. 3, pp. 114-123, 2009.
- [15] X. Yu et al., "Transformers.js: Client-side inference for Hugging Face models in the browser," *Hugging Face Blog*, 2023.
- [16] K. Lewis and B. Perry, "Role-based access control in cloud-native applications: Patterns and anti-patterns," *IEEE Security & Privacy*, vol. 21, no. 4, pp. 54-63, 2023.
- [17] Ministry of Education, India, "National Education Policy 2020: Higher Education Transformation," *Government of India, New Delhi*, 2020.
- [18] Firebase Security Rules Team, "Best practices for Firebase security rules," *Firebase Blog, Google LLC*, 2024.

AUTHORS BIOGRAPHY

Yash Manoj Kamatwar received his B.Tech degree in Computer Science and Engineering from DBATU University. His research interests include web development, 3D visualisation, and cloud-native application architecture. He led

the system architecture and frontend development of FresherConnect.

Gajendra Pradip Khandale is a B.Tech graduate in Computer Science and Engineering from DBATU University. His research focuses on backend systems, real-time communication protocols, and Progressive Web Applications.

Gunvant Pradip Khandale received his B.Tech in Computer Science and Engineering from DBATU University. His research interests include artificial intelligence, natural language processing, and semantic search technologies.

Sumit Arun Pote is a B.Tech graduate in Computer Science and Engineering from DBATU University. His research interests include cybersecurity, cryptographic protocols, and DevOps practices.

Prof. Pushpa T. Tandekar is an Assistant Professor in the Department of Computer Science and Engineering at Shri Sai College of Engineering and Technology, Bhadravati. Her research interests include software engineering, database systems, and educational technology.

Citation of this Article:

Yash Manoj Kamatwar, Gajendra Pradip Khandale, Gunvant Pradip Khandale, Sumit Arun Pote, & Pushpa T. Tandekar. (2026). FresherConnect: An AI-Augmented Real-Time Smart Campus Management System with 3D Galaxy Visualisation, Role-Based Access Control, and Progressive Web Application Architecture. *International Research Journal of Innovations in Engineering and Technology - IRJIET*, 10(5), 149-159. Article DOI <https://doi.org/10.47001/IRJIET/2026.105020>
