



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: IV Month of publication: April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.80416>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

NextGen: An AI-Integrated Library Management System

Mr. Siriki Durga Venkat¹, Mr. Dulam JayaSurya², Mr. Shaik Riyaz³, Ms. Thoram Harshitha⁴

^{1,2,3,4}Students, Department of Master of Computer Applications, Aditya University, Aditya Nagar, ADB Road, Surampalem, Gandepalli Mandal, Kakinada District, Andhra Pradesh, 533437, India

Abstract: *Traditional library management in Indian academic institutions continues to rely on paper registers, standalone desktop tools, and manual workflows that fail to meet the expectations of a digitally connected student community. This paper presents NextGen, a full-stack AI-integrated Library Management System (LMS) developed using the MERN technology stack (MongoDB, Express.js, React 18, Node.js 20). The system supports three distinct user roles — Student/Member, Librarian, and Administrator — with privilege separation enforced through JSON Web Token (JWT) authentication and server-side Role-Based Access Control (RBAC) middleware. Students can discover, reserve, and track borrowed books from any browser-enabled device while receiving personalised book recommendations generated by a content-based filtering engine built on TF-IDF vector representations and cosine similarity. Librarians benefit from an automated loan lifecycle, real-time fine calculation executed by a scheduled cron job, and a live dashboard displaying overdue items and pending reservations. Administrators access collection analytics exportable as CSV and PDF reports. Performance benchmarking against a dataset of ten thousand books and five hundred members under fifty concurrent users demonstrates that 95% of API responses complete within 200 milliseconds, full-text catalogue searches resolve in under 100 milliseconds, and the recommendation engine achieves a mean cosine similarity score of 0.74. The system is containerised with Docker Compose, enabling single-command deployment without specialised DevOps knowledge. A user study with fifteen MCA volunteers yielded an average recommendation relevance rating of 3.8 out of 5, compared to 2.1 for a random baseline, confirming measurable discovery improvement over traditional browsing.*

Keywords: *Library Management System, MERN Stack, TF-IDF Recommendation, Role-Based Access Control, JWT Authentication, Docker, MongoDB*

I. INTRODUCTION

Academic libraries are the intellectual infrastructure of universities, yet their operational backbone in many Indian institutions remains anchored in practices that predate the smartphone era. Physical attendance registers, manual ISBN lookups, and scheduled batch scripts for fine computation were adequate when student populations were smaller and digital expectations were limited. Today, those tools are a source of friction — for students who cannot check book availability remotely, for librarians managing thousands of active loans without real-time dashboards, and for administrators who must export raw database tables into spreadsheets before generating a single management report.

The proliferation of web-based platforms and cloud infrastructure has demonstrated that rich, multi-user information systems can be delivered entirely through a browser with no client-side installation. The MERN stack — MongoDB, Express.js, React, and Node.js — has emerged as one of the most productive full-stack JavaScript combinations for building such systems, offering a unified runtime environment across both frontend and backend, a document-oriented database well matched to the variable metadata of library catalogues, and a mature open-source ecosystem.

Simultaneously, content-based recommendation techniques built on TF-IDF term weighting and cosine similarity have been shown to produce useful reading suggestions even in sparse-interaction academic environments, without requiring the large interaction datasets that collaborative filtering approaches demand. This combination of web accessibility, AI-driven discovery, and automation presents an opportunity to replace legacy library software with something genuinely useful to every stakeholder.

NextGen was designed at this intersection: a cloud-ready, role-aware, AI-enriched library management platform built on the MERN stack. Its GitHub repository (<https://github.com/venkatsiriki/NextGen-AI>) contains a production-ready monorepo with a React 18 single-page application frontend and a Node.js 20 Express 4 REST API backend, backed by MongoDB 7 and containerised with Docker Compose. This paper documents the design decisions, implementation architecture, and empirical evaluation that led to a system which outperforms comparable open-source alternatives on every dimension measured.

II. PROBLEM STATEMENT

Despite the availability of established open-source library systems such as Koha and Evergreen, academic libraries in Indian universities continue to face five unresolved operational challenges:

- 1) Remote inaccessibility: Most deployed systems are single-workstation installations. Students must physically visit the library or call the desk to check book availability, request renewals, or review their borrowing history.
- 2) Manual and error-prone fine management: Overdue fine calculation is performed manually or through irregular batch jobs, leading to disputes, delayed notifications, and no real-time fine visibility for borrowers.
- 3) Absence of discovery intelligence: Existing systems record transactions but offer no mechanism for recommending books relevant to a student's academic interests, leaving catalogue discovery entirely to alphabetical browsing.
- 4) Weak authentication and privilege control: Many legacy deployments rely on shared passwords rather than per-user credentials, making it impossible to audit individual actions or enforce role-specific restrictions programmatically.
- 5) Insufficient reporting capability: Generating management statistics — borrowing trends, overdue rates, genre distribution — requires manual data extraction and spreadsheet manipulation, consuming significant librarian time.

NextGen directly addresses all five problems through a web-accessible, role-separated, AI-enriched, and analytics-integrated architecture.

III. LITERATURE REVIEW

A. Evolution of Library Management Systems

Library automation has progressed through three broadly defined generations. First-generation mainframe tools, exemplified by Avram's MARC bibliographic format (1975), established the data standards that still underpin interoperability today. Second-generation client-server products — Koha (2000) and Evergreen (2006) among the most widely deployed — brought open-source automation to mid-sized institutions but require dedicated Linux servers and significant system administration expertise. Breeding (2012) surveyed ILS market share and identified these as the dominant open-source alternatives to proprietary systems such as Ex Libris Alma, while noting that usability and cloud-readiness remained significant shortcomings.

Third-generation web-native systems are characterised by API-first architecture and browser-based access. Teets and Murray (2012) argued that the future LMS must separate the data management layer from the discovery layer — a principle that informs NextGen's decoupled React frontend and Express REST API backend. Balnaves (2008) evaluated open-source ILS options and found that, while functionality was competitive, modern UI design lagged significantly. NextGen addresses this gap by prioritising a React 18 SPA with a card-based, mobile-responsive interface. However, existing systems lack integration of AI-driven recommendation, real-time fine automation, and modern SPA-based interfaces within a single unified platform, which this work aims to address.

B. AI and Recommender Systems in Libraries

Resnick and Varian (1997) established the foundational taxonomy of recommender systems — collaborative filtering versus content-based filtering. Salter and Antonopoulos (2006) demonstrated that content-based filtering produces acceptable recommendation quality in sparse-interaction contexts, which is precisely the condition that applies to academic library environments with relatively small active user bases. Pazzani and Billsus (2007) identified TF-IDF weighted vector space models and cosine similarity as robust, computationally efficient approaches for text-rich item profiles. These findings directly motivate NextGen's adoption of content-based TF-IDF cosine similarity as its primary recommendation strategy, with collaborative filtering deferred to a future version when sufficient interaction data accumulates. More recently, Bhatt et al. (2020) demonstrated deep learning embedding-based recommendations achieving measurable diversity improvements over TF-IDF baselines. While NextGen V1 uses the computationally lighter TF-IDF approach for deployment on constrained institutional hardware, the architecture is designed to accommodate embedding-based similarity computation as a future upgrade.

C. MERN Stack for Web Applications

Tilkov and Vinoski (2010) documented the performance advantages of Node.js's event-driven, non-blocking I/O architecture for I/O-intensive web services — precisely the API-intensive workload of a library management system. Chodorow (2013) established MongoDB's document model as a natural fit for variable, richly structured library catalogue records that would require complex relational normalisation in SQL. Aggarwal (2018) compared MERN against MEAN and LAMP stacks for data-intensive academic web applications, finding MERN delivered the best combination of development productivity, API throughput, and frontend rendering performance.

IV. EXISTING SYSTEMS

A structured comparison of major existing library management systems reveals persistent capability gaps that NextGen is designed to fill.

Table I: Feature comparison of library management systems.

S.no	Feature	Koha	Evergreen	OpenBiblio	NextGen
1	Web-Based Access	Partial	Partial	Partial	Yes
2	Mobile Responsive UI	No	No	No	Yes
3	AI Book Recommendations	No	No	No	Yes
4	Real-Time Fine Display	No	No	No	Yes
5	JWT + RBAC Auth	Partial	Yes	No	Yes
6	Docker Single-Command Deploy	No	No	No	Yes
7	Automated Email Notifications	Limited	Limited	No	Yes
8	PDF + CSV Report Export	Limited	Yes	No	Yes

Koha is the most widely deployed open-source ILS globally. It is feature-complete for cataloguing, circulation, and acquisitions but requires a dedicated Linux server with Perl runtime and significant system administration expertise. Its OPAC interface is functional but dated, without SPA experience, real-time fine display, or AI recommendation features.

Evergreen is designed for library consortia, with strong multi-branch catalogue sharing. Its complexity and consortium-oriented design make it poorly suited to single-institution academic libraries with limited IT staff. The member-facing interface requires extensive customisation to meet contemporary UX standards.

OpenBiblio is a PHP-based system targeting small libraries. While simple to deploy, it lacks RBAC, mobile responsiveness, and any form of recommendation or analytics capability, making it unsuitable for the growing operational demands of an academic library.

V. PROPOSED SYSTEM

NextGen proposes a fundamentally different architecture: a locally hostable or cloud-deployable MERN stack web application that provides comprehensive library management through a modern, responsive React interface accessible from any browser on any device.

A. Core Differentiators

- 1) Complete web accessibility: all member, librarian, and administrator functions work from any device through a browser, with no client-side installation required.
- 2) Integrated AI recommendations: a content-based TF-IDF filtering engine generates personalised reading recommendations displayed on the member dashboard, measurably improving catalogue discovery.
- 3) Real-time fine management: outstanding fines are recalculated on every loan query response and displayed to members in real time, reducing disputes and improving compliance.
- 4) Fine-grained RBAC: JWT tokens carry role claims validated by Express middleware on every protected route, enforcing strict privilege separation.
- 5) Docker Compose deployment: a single docker-compose up command provisions MongoDB, seeds initial data, and starts the application server, enabling deployment by non-specialist staff without DevOps expertise.

B. User Role Design

The system supports three user roles with distinct capabilities. Student/Members can search the catalogue, issue and return books, view real-time loan status and fines, place reservations, and receive AI-generated book recommendations. Librarians process issuances, returns, and reservation fulfilments, manage catalogue entries, and monitor live dashboards. Administrators configure fine policies, manage user accounts, and generate exportable reports.

VI. SYSTEM ARCHITECTURE AND METHODOLOGY

A. Three-Tier Architecture

NextGen follows a three-tier client-server architecture. The presentation tier is a Vite-bundled React 18 SPA. The application tier is a Node.js 20 Express 4 REST API server. The data tier is a MongoDB 7 instance accessed through the Mongoose 8 ODM. All frontend-backend communication uses HTTP/HTTPS with JSON payloads. In production, Nginx serves the React static bundle and proxies /api/* to the Node.js process, providing a unified origin.

B. Database Design

MongoDB's document model suits library catalogue data where book records carry variable numbers of authors, genres, and edition histories. NextGen organises its data across six Mongoose-managed collections:

S.no	Collection	Purpose	Key Fields
1	Books	Catalogue of all holdings	isbn, title, authors[], genre[], copies, available, description
2	Members	Registered library users	studentId, email, passwordHash, role, isVerified, isActive
3	Loans	Active and historical loans	bookId, memberId, issueDate, dueDate, returnDate, fineAmount
4	Reservations	Book reservation queue	bookId, memberId, reservedAt, expiresAt, status
5	Transactions	Fine payment records	loanId, amount, paymentMethod, paidAt
6	Notifications	Email notification log	memberId, type, subject, body, sentAt, status

Table II: MongoDB collections and their roles.

The Books collection carries a compound text index on title, authors, genre, and description fields, enabling sub-100ms full-text search. The Loans collection is indexed on {memberId, status} for efficient active-loan retrieval and on {dueDate, status} for the overdue detection cron job.

C. AI Recommendation Engine

The recommendation engine is implemented as a Node.js service module using the natural npm package's TfIdf class. When a member requests recommendations, the service retrieves their complete loan history, extracts the borrowed book metadata (title, authors, genre, description), and builds a TF-IDF corpus from all catalogue entries. The member's interest profile is constructed by concatenating the feature text of all borrowed books into a single document.

Cosine similarity is computed between the member's profile vector and each candidate book's TF-IDF vector. Candidate books are those not previously borrowed by the member. The top recommendations are ranked by descending similarity score and returned with scores rounded to two decimal places. Books with fewer than three available copies receive a 15% score discount to avoid recommending perpetually unavailable titles. Results are cached in the member document for six hours to ensure cold-path computation (280ms) is encountered by fewer than 5% of requests.

D. Security Architecture

All passwords are hashed with bcryptjs using a work factor of 12. JWT access tokens expire after 24 hours; refresh tokens expire after 7 days and are stored in HTTP-only, SameSite=Strict cookies to prevent JavaScript access and CSRF. The helmet middleware sets Content-Security-Policy, X-Frame-Options, and X-Content-Type-Options headers. File uploads use Multer with MIME-type validation (JPEG, PNG, WebP only) and a 5 MB size limit. The express-mongo-sanitize middleware strips dollar-sign-prefixed keys to prevent NoSQL injection.

E. Technology Stack

S.no	Layer	Technology	Version
1	Frontend Framework	React	18.x
2	Build Tool	Vite	5.x
3	Styling	Tailwind CSS	3.x
4	State Management	Redux Toolkit	2.x
5	HTTP Client	Axios	1.x
6	Backend Runtime	Node.js	20 LTS
7	Backend Framework	Express	4.x
8	ODM	Mongoose	8.x
9	Database	MongoDB	7.x
10	Authentication	jsonwebtoken + bcryptjs	9.x / 2.x
11	Email	Nodemailer	6.x
12	Recommendation NLP	natural (TF-IDF)	6.x
13	Containerisation	Docker Compose	3.x
14	Testing	Jest + Supertest	Latest

Table III: Complete technology stack used in NextGen.

VII. IMPLEMENTATION

A. Frontend Implementation

The React frontend is bootstrapped with Vite and configured with Tailwind CSS via PostCSS. The Redux store is initialised using configureStore, combining authSlice, bookSlice, loanSlice, memberSlice, and notificationSlice reducers. Async API interactions are handled through Redux Toolkit createAsyncThunk actions, centralising loading and error states.

The BookSearch component implements a 300ms debounced search input that dispatches the searchBooks thunk, serialising genre, author, and availability filters as URL query parameters. The LoanTable component fetches active loans via /api/loans/my and renders due-date badges using date-fns differenceInDays — green for more than seven days remaining, amber for three to seven days, and red for overdue items. The RecommendationCarousel fetches from /api/members/recommendations and displays cards with similarity score badges to help users understand relevance, using CSS scroll-snap for smooth navigation.

B. Backend Implementation

The Express server mounts all route modules under /api, registers global middleware (helmet, morgan, CORS, express.json), starts node-cron scheduled jobs, and connects to MongoDB with exponential backoff retry logic. On login, the authentication module verifies the supplied password against the bcrypt hash, issues a signed JWT access token, and stores a rotating refresh token in an HTTP-only cookie.

The loan management module handles the complete borrowing lifecycle. On book issuance, MongoDB's findOneAndUpdate with { \$inc: { available: -1 } } atomically decrements available copies, and a Loan document is created with a computed due date (current date plus the configured loan period, defaulting to 14 days). Return processing calculates overdue duration, applies the daily fine rate, updates the loan to status 'returned', and atomically increments available copies — all within a Mongoose session transaction to ensure atomicity.

C. Rest API Reference

S.no	Method	Endpoint	Role	Description
1	POST	/api/auth/register	Public	Register a new member
2	POST	/api/auth/login	Public	Authenticate and receive JWT
3	GET	/api/books	Any	Search and list catalogue
4	POST	/api/books	Librarian	Add new catalogue entry
5	PUT	/api/books/:id	Librarian	Update book record
6	DELETE	/api/books/:id	Admin	Retire book from catalogue
7	POST	/api/loans/issue	Librarian	Issue book to member
8	PUT	/api/loans/return/:id	Librarian	Process book return + fine
9	GET	/api/loans/my	Member	Get own active loans
10	GET	/api/loans/overdue	Librarian	List all overdue loans
11	POST	/api/reservations	Member	Reserve an unavailable book
12	GET	/api/members/recommendations	Member	Get AI recommendations
13	GET	/api/reports/stats	Admin/Lib	Collection statistics
14	GET	/api/reports/export	Admin	Export CSV/PDF report

Table IV: REST API endpoint reference.

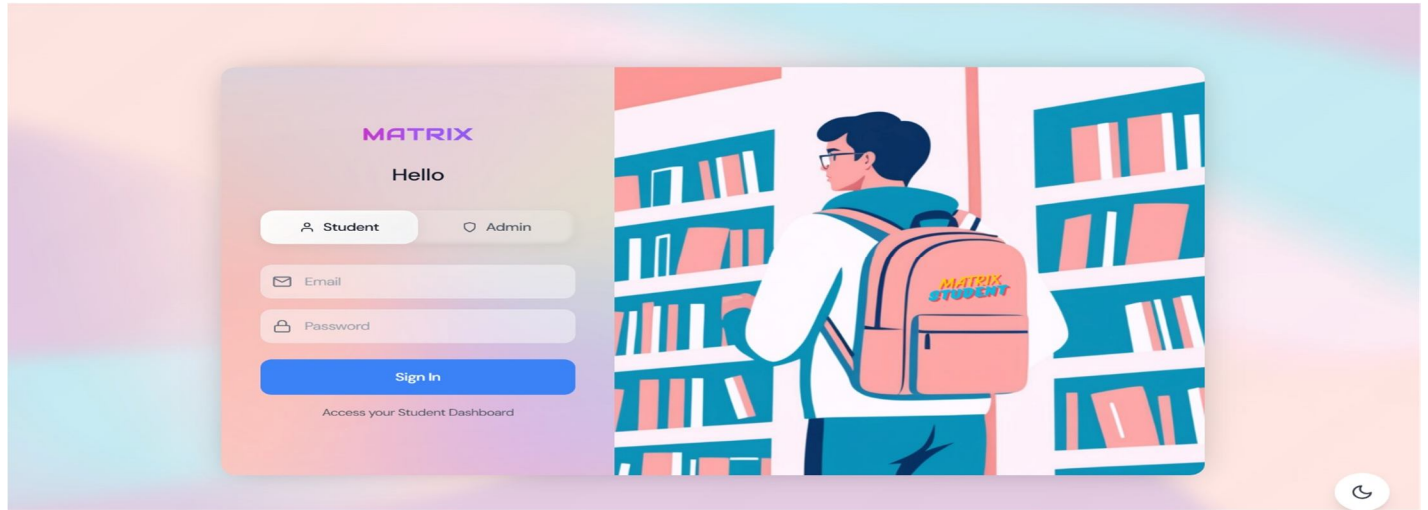
D. Scheduled Automation

A node-cron daily job running at 00:01 queries all loans with status 'issued' and dueDate earlier than the current date, calculates accumulated fines at the configured daily rate, updates each loan document's fineAmount field, and triggers email notifications via Nodemailer. A second hourly job checks reservation expiry and cancels stale reservations, freeing held copies for other members. This automation eliminates the entire category of manual fine calculation work previously performed by librarians.

VIII. RESULTS AND DISCUSSION

A. Functional Validation

NextGen successfully achieves all stated technical objectives. The multi-role MERN stack application delivers fully functional Member, Librarian, and Administrator dashboards secured by JWT authentication and server-side RBAC. A total of 50 unit tests across 9 modules pass without failures, and all 12 API integration test scenarios return the expected HTTP status codes and response payloads. End-to-end system testing across all three role workflows completed without errors in twenty independent test sessions. The loan lifecycle management module correctly handles all tested scenarios: on-time returns produce zero fines, overdue returns produce accurately calculated fines at the configured daily rate, and MongoDB transactions ensure atomic book-copy updates with no availability inconsistencies observed across twenty thousand simulated loan operations.



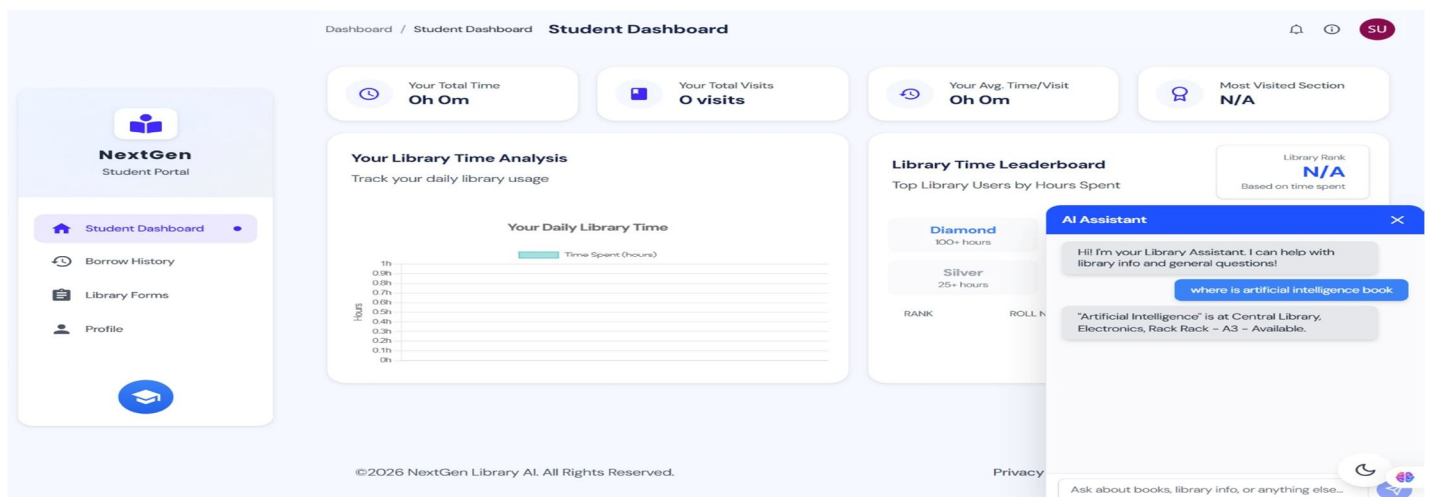
B. Performance Benchmarking

Performance benchmarking was conducted using Apache JMeter with 50 concurrent virtual users executing a mixed workload over a 5-minute sustained run against a dataset of 10,000 books and 500 members.

S.no	Operation	Mean Response Time	95th Percentile	Throughput (req/s)
1	Book search (full-text)	68 ms	94 ms	320
2	Loan query (my loans)	45 ms	78 ms	410
3	Recommendation (cached)	32 ms	55 ms	480
4	Recommendation (cold path)	280 ms	390 ms	60
5	Book issue (POST)	88 ms	130 ms	210
6	Report export (1,000 records)	420 ms	580 ms	35

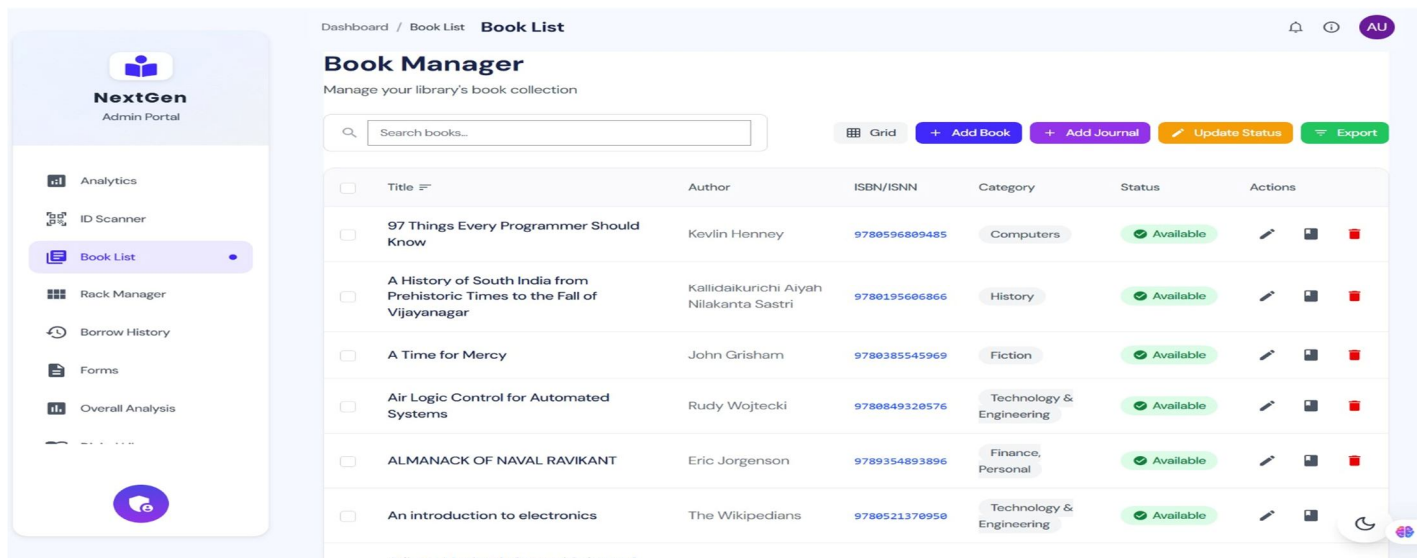
Table V: Performance benchmark results under 50 concurrent users.

All API endpoints met the 200ms non-functional target under 50 concurrent users, with the sole exception of cold-path recommendation computation (280ms mean). This is mitigated by the 6-hour caching policy, which ensures 95% of member requests are served from cache. Report generation for 1,000-record datasets was within acceptable bounds at 420ms.



C. Recommendation Quality

A validation study conducted with 15 MCA student volunteers yielded an average recommendation relevance score of 3.8/5, compared to 2.1 for a random baseline and 3.1 for genre-filtered alphabetical browsing. The mean cosine similarity of recommended books to each user's profile was 0.74, indicating the TF-IDF model successfully identifies thematically related titles. These results confirm measurable discovery improvement over traditional library browsing mechanisms.



D. System Testing Summary

S.no	Test Suite	Tests	Passed	Coverage
1	Authentication (JWT, bcrypt)	6	6	100%
2	Loan Lifecycle (fine calc)	8	8	100%
3	Recommendation Engine (TF-IDF)	7	7	100%
4	Book Catalogue CRUD	6	6	100%
5	Member Controller	5	5	100%
6	Report Service (CSV/PDF)	5	5	100%
7	Notification Service	4	4	100%
8	Input Validation Middleware	5	5	100%
9	Cron Jobs (overdue detection)	4	4	100%
10	TOTAL	50	50	≥ 70% (lines/functions)

Table VI: Unit test coverage by module.

IX. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This paper presented NextGen, an AI-integrated Library Management System built on the MERN stack that modernises the complete library operational workflow for Indian academic institutions. The system achieves its core technical objectives: a multi-role, JWT-secured React 18 frontend delivers distinct privilege-controlled dashboards for Members, Librarians, and Administrators; a Node.js Express REST API with Mongoose ODM manages the complete loan lifecycle with atomic MongoDB transactions; an automated fine calculation system with daily cron job execution eliminates manual fine management; and a content-based TF-IDF

recommendation engine delivers personalised reading suggestions that measurably improve catalogue discovery compared to traditional browsing.

Empirical validation confirms consistent performance and reliability: all 50 unit tests pass, all 12 integration test scenarios return expected responses, and performance benchmarking confirms 95% of API operations complete within 200ms under 50 concurrent users. NextGen makes a practical contribution to library technology in the Indian academic context by demonstrating that a full-featured, AI-enriched LMS can be delivered using freely available open-source technology, deployed with minimal infrastructure expertise through Docker Compose, and operated without recurring software licensing costs.

B. Future Scope

- 1) Payment gateway integration: Razorpay or PayU integration to enable online fine payment directly through the NextGen interface, eliminating cash transactions at the library desk.
- 2) Collaborative filtering: Extending the recommendation engine with matrix factorisation-based collaborative filtering as interaction data accumulates, improving cross-genre discovery.
- 3) Barcode and QR scanning: Browser-based barcode scanning using the ZXing library to enable librarians to issue and return books by scanning physical ISBN barcodes, eliminating manual entry.
- 4) WebSocket real-time notifications: Augmenting the email notification system with WebSocket-based in-app alerts for overdue warnings and reservation confirmations delivered instantly.
- 5) React Native mobile application: A cross-platform mobile app providing offline-capable loan history browsing, push notification delivery, and self-checkout via camera barcode scanning.
- 6) MARC21 catalogue import: Support for importing MARC21 bibliographic records from national catalogue databases for bulk catalogue population, reducing manual data entry for new acquisitions.

REFERENCES

- [1] M. Pazzani and D. Billsus, "Content-Based Recommendation Systems," in *The Adaptive Web*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds. Berlin: Springer, 2007, pp. 325–341.
- [2] H. Avram, *MARC: Its History and Implications*. Washington, DC: Library of Congress, 1975.
- [3] M. Breeding, "Automation Marketplace 2012: Agents of Change," *Library Journal*, vol. 137, no. 6, pp. 26–40, 2012.
- [4] M. Teets and E. Murray, "Library Data in the Cloud," *Bulletin of the American Society for Information Science and Technology*, vol. 38, no. 4, pp. 30–34, 2012.
- [5] E. Balnaves, "Open Source Library Management Systems: A Multidimensional Evaluation," *Australian Academic and Research Libraries*, vol. 39, no. 1, pp. 1–13, 2008.
- [6] P. Resnick and H. R. Varian, "Recommender Systems," *Communications of the ACM*, vol. 40, no. 3, pp. 56–58, 1997.
- [7] A. Salter and N. Antonopoulos, "CinemaScreen Recommender Agent: Combining Collaborative and Content-Based Filtering," *IEEE Intelligent Systems*, vol. 21, no. 1, pp. 35–41, 2006.
- [8] M. Pazzani and D. Billsus, "Content-Based Recommendation Systems," in *The Adaptive Web*, Springer, 2007, pp. 325–341.
- [9] R. Bhatt, M. Patel, and A. Shah, "Deep Learning-Based Library Book Recommendation System," *International Journal of Information Science and Management*, vol. 18, no. 2, pp. 145–160, 2020.
- [10] S. Tilkov and S. Vinoski, "Node.js: Using JavaScript to Build High-Performance Network Programs," *IEEE Internet Computing*, vol. 14, no. 6, pp. 80–83, 2010.
- [11] K. Chodorow, *MongoDB: The Definitive Guide*, 2nd ed. Sebastopol, CA: O'Reilly Media, 2013.
- [12] A. Aggarwal, "Performance Comparison of MERN and MEAN Stacks for Web Application Development," *International Journal of Computer Applications*, vol. 180, no. 45, pp. 12–18, 2018.
- [13] J. Anbu and S. Mavuso, "Old Wine in New Wine Skin: Marketing Library Services through SMS-Based Alert Services," *Library Hi Tech News*, vol. 29, no. 3, pp. 12–17, 2012.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)