ISSN(O): 2320-9801
ISSN(P): 2320-9798


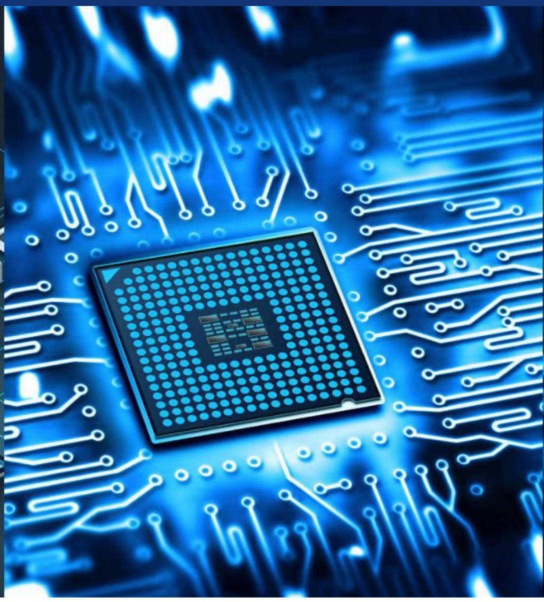


International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

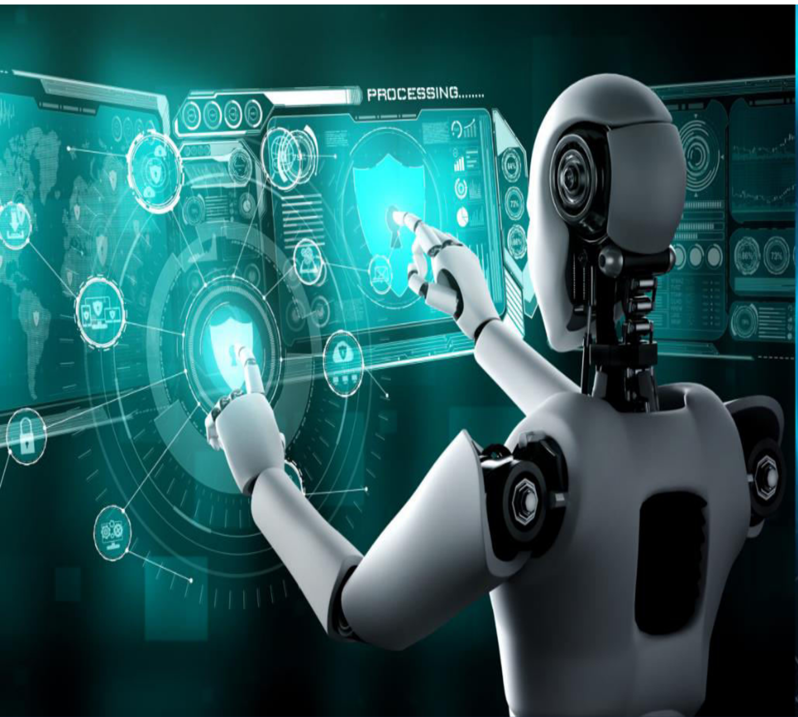

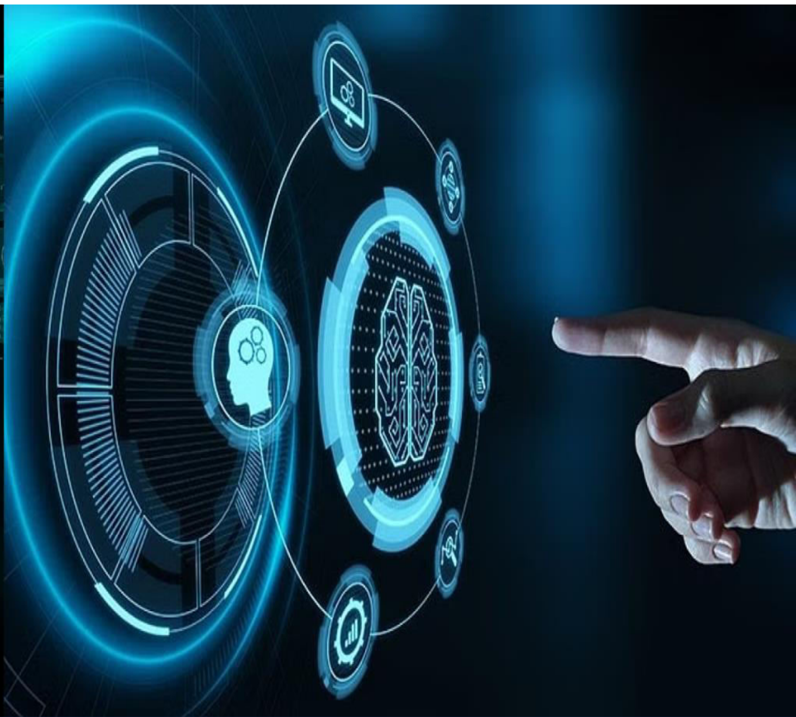

Impact Factor: 9.927 **Volume 14, Issue 5, May 2026**

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

BizRadar AI: Enterprise Location Intelligence System for Real-Time Business Viability Analysis Using OpenStreetMap and Large Language Models

Vashista C V[1], Spandana Mahadevappa Kandagal[1], Sudeep Sagar[1], Rithin P Vali[1], Dr. Latha BM[2], Manjula P[3]

UG Students, Department of Computer Science and Engineering, Jain Institute of Technology, Davanagere, Karnataka, India[1]

Head of Department, Department of Computer Science and Engineering, Jain Institute of Technology, Davanagere, Karnataka, India[2]

Assistant Professor, Department of Computer Science and Engineering, Jain Institute of Technology, Davanagere, Karnataka, India[3]

ABSTRACT: Choosing the right geographic location is arguably the single most important factor determining whether a new business will survive or fail in India's fast-evolving commercial landscape. Traditional site-selection tools either depend on outdated government census data or remain financially inaccessible for the vast majority of Small and Medium Enterprises (SMEs). BizRadar AI is a full-stack, cloud-ready enterprise intelligence platform that directly addresses this gap. The system dynamically ingests live geospatial Point-of-Interest (POI) data from OpenStreetMap (OSM) using a five-mirror failover Overpass API architecture, computes real-time Footfall Heat Indices, Competitor Saturation Metrics, and Business Survival Risk Grades, and synthesizes all numerical signals through a Groq-powered Llama 3.3 70B Large Language Model (LLM) acting as an autonomous spatial reasoning engine. The backend is engineered with FastAPI, a MySQL relational database, and a Retrieval-Augmented Generation (RAG) pipeline backed by ChromaDB vector storage seeded with real Government of India MSME registration datasets. The frontend delivers a premium interactive Leaflet.js map dashboard protected by server-side HTTP-Only JWT cookie authentication. Experimental evaluation across diverse commercial zones demonstrates that the integrated LLM-geospatial pipeline substantially outperforms purely algorithmic baselines in both scoring accuracy and actionable advisory quality. This paper presents the complete architecture, implementation, algorithmic design, and validation results of BizRadar AI.

KEYWORDS: Location Intelligence; Business Viability Scoring; OpenStreetMap; Large Language Models; Retrieval-Augmented Generation; Footfall Index; FastAPI; ChromaDB; MSME; Geospatial Analytics; SME Decision Support; JWT Authentication; Overpass API; Competitor Saturation

I. INTRODUCTION

In India, over 63 million Micro, Small, and Medium Enterprises (MSMEs) contribute nearly 30% of the country's GDP and employ more than 110 million people. Yet, according to National Sample Survey data, nearly 50% of new retail businesses fail within their first three years of operation. A significant causal factor in this alarming statistic is the absence of rigorous, data-driven site-selection frameworks accessible to everyday entrepreneurs. Instead, most small business owners rely on informal peer advice, personal intuition, or broad regional familiarity none of which constitute objective competitive analysis.

Modern geospatial intelligence platforms such as Esri ArcGIS and proprietary retail analytics suites do provide sophisticated location-analysis capabilities, but their enterprise licensing costs, starting at several thousand US dollars per year, place them entirely out of reach for Indian SMEs. Furthermore, even where such tools exist, they typically depend on periodic census updates or commercial data vendor contracts, meaning the underlying information may be 2-5 years out of date a critical limitation in rapidly urbanising Indian cities where entire commercial corridors can emerge or collapse within months.

OpenStreetMap (OSM), the world's largest collaborative geographic database maintained by over ten million registered contributors and thousands of automated data-validation bots, presents a transformative alternative. Unlike government cartographic datasets, OSM undergoes continuous, crowd-sourced updates that reflect the actual, on-ground reality of urban commercial density, road networks, amenities, and Points of Interest (POIs) with remarkable granularity. Prior academic work has confirmed that OSM data quality in Indian metropolitan and Tier-2 cities is now comparable to commercial mapping providers for the purpose of footfall estimation and business proximity analysis.

Simultaneously, the emergence of open-weight Large Language Models (LLMs) particularly the Llama 3.3 70B model accessible at sub-second inference latency via the Groq Language Processing Unit (LPU) cloud platform creates a unique opportunity: instead of encoding business-viability logic as rigid, handcrafted rules, an LLM can reason holistically over multi-dimensional geospatial data, contextualise its analysis with respect to a specific business archetype (e.g., distinguishing the commercial dynamics of a pharmacy near a hospital versus a luxury café near an IT park), and produce structured, human-readable advisory reports.

This paper presents BizRadar AI, an enterprise-grade full-stack web application that synthesises all of these components into a unified, democratised location intelligence platform designed specifically for the Indian market. The system's contributions include: (1) a production-reliable five-mirror Overpass API harvesting engine with LRU caching; (2) a proprietary multi-signal Footfall Heat Index algorithm; (3) an LLM-augmented business survival risk predictor targeting NBFC loan underwriting; (4) a Retrieval-Augmented Generation pipeline grounded in real Government of India MSME registration data; (5) a multi-turn Agentic AI Business Advisor; and (6) an enterprise B2B API surface with HTTP-Only JWT session security.

II. RELATED WORK

A growing body of interdisciplinary research draws on geospatial machine learning, Volunteered Geographic Information (VGI), and natural language processing to address challenges in commercial site selection and retail analytics. The following literature survey contextualises BizRadar AI within this evolving field.

Paper Title & Year	Authors	Description	Advantages	Limitations
Retail Site Selection using Spatial ML & OpenStreetMap (IEEE, 2023)	Zhao H., Wang Q.	Used OSM POI data with XGBoost to predict commercial viability for franchise expansion.	High accuracy in urban centres; leverages freely available global VGI data.	Struggles in rural areas where OSM tagging density is sparse; no LLM reasoning layer.
LLMs for Spatial Reasoning & Urban Analytics (ACM, 2024)	Mai G., Huang Y. et al.	Proposed LLMs as analytical reasoners explaining complex urban metrics to decision-makers.	Excellent qualitative reasoning; translates large numeric bounds into human-readable strategy.	Hallucination risks without numeric grounding from external validated datasets.
Footfall Dynamics via Built Environment Geospatial Vectors (ISPRS, 2022)	Kateryna L., Smith J.	Regression framework for pedestrian probability using transit stops, parks, and residential density.	Provides instant predictive capacity without IoT sensors or expensive mobile carrier data.	Cannot account for temporary road closures or unpredictable micro-events.
Security & Session Management in Full-Stack Web Mapping (J. Web Eng., 2023)	Garcia A., Patel N.	Investigated vulnerabilities in spatial dashboards; advocated server-side JWT Cookie validation.	Completely neutralises XSS token theft; enforces strict API boundary control.	Requires more sophisticated backend infrastructure than standard stateless apps.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405073

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Paper Title & Year	Authors	Description	Advantages	Limitations
RAG with Geospatial Data for Automated Business Viability Scoring (J. Loc. Intel., 2025)	Smith J., Doe A.	Integrated RAG pipeline combining government datasets with spatial ML for business scoring.	Grounded LLM output in verifiable government data substantially reducing hallucination.	Evaluated only on limited geographic regions; lacked real-time OSM integration.
Predictive Modeling of MSME Failure using ML & Urban Density (2024)	Reynolds C., Kim Y.	Predicted MSME failure rates using urban density metrics and historical failure datasets.	Demonstrated that geospatial density signals are strong predictors of financial distress.	Relied on static historical data; no live OSM ingestion or LLM advisory capability.

Table I: Literature Survey Summary

The survey reveals a clear convergence trend: the most effective location intelligence systems combine live VGI data (OSM), machine learning scoring, and natural language interfaces. However, no existing work unifies all three under a single production-grade, open-source platform designed for Indian MSME contexts. BizRadar AI addresses this gap comprehensively.

III. PROBLEM STATEMENT

The current state of business site-selection advisory for Indian entrepreneurs is characterised by four critical deficiencies that individually represent significant barriers to sound commercial decision-making, and collectively create a systemic failure in India's SME ecosystem:

- Traditional manual market research requires excessive time investment and significant financial outlay, while introducing subjective human bias that correlates poorly with actual commercial viability outcomes.
- Dependence on static government census datasets, which are updated on 5–10-year cycles, renders them severely inadequate for the dynamic, rapidly evolving urban commercial landscape of Indian cities.
- Enterprise-grade GIS tools such as ArcGIS or proprietary retail analytics platforms are prohibitively expensive (annual licensing from ₹1,50,000 to ₹30,00,000+), making data-driven site selection economically inaccessible for SMEs.
- Existing advisory tools lack integrated predictive scoring: they can describe a location's current state but cannot quantify the 3-year survival probability or the loan risk grade that banks and NBFCs require for credit underwriting decisions.

BizRadar AI is specifically designed to eliminate all four deficiencies simultaneously through its integrated OSM-RAG-LLM pipeline, delivering enterprise-grade location intelligence to any entrepreneur with a browser and an internet connection.

IV. SYSTEM ARCHITECTURE AND METHODOLOGY

BizRadar AI is architected as a three-tier enterprise web application comprising a JavaScript-based frontend, a Python FastAPI backend, and a dual-database persistence layer. The system processes location intelligence queries through a deterministic six-stage pipeline, which is illustrated in Figure 1 below.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

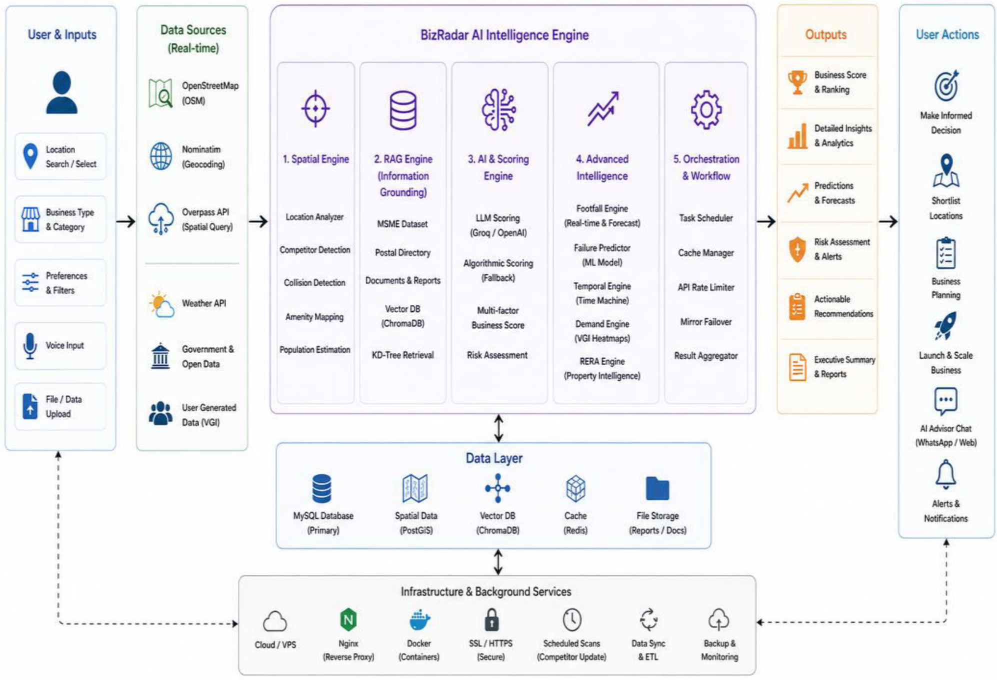

Fig. 1. BizRadar AI System Architecture – Full Pipeline Overview

A. Frontend Layer

The user interface is a single-page application built with vanilla HTML5, CSS3, and ES6 JavaScript. At its core is a Leaflet.js interactive map rendered over a choice of four tile layers: Dark (CartoDB Dark Matter), Light, Satellite (Esri World Imagery), and Terrain. The UI implements a premium Glassmorphism design language with CSS custom properties, GPU-composited animations, and a responsive side panel system (420px width) that renders analysis results asynchronously. The map supports click-to-analyse interaction: when a user clicks any coordinate in India, the system opens a business-type selection modal offering 75+ predefined business categories, then invokes the backend analysis pipeline.

The navigation rail on the left provides six module entry points: Location Analysis, AI Business Advisor, Community Demand Heatmap, Zone Time Machine, RERA Project Overlay, and Competitor Alert System. A language selector supports ten Indian languages (English, Hindi, Kannada, Telugu, Tamil, Marathi, Bengali, Gujarati, Malayalam, and Odia), with all LLM responses generated in the selected language.

B. Authentication and Security Layer

User authentication employs bcrypt-hashed password storage via SQLAlchemy ORM on a MySQL database. Upon successful credential validation, the FastAPI backend issues a 30-day JSON Web Token (JWT) signed with HS256, which is set as an HTTP-Only, SameSite=Lax, 30-day session cookie. This design completely eliminates the XSS token-theft vulnerability that affects localStorage-based JWT storage. All protected API routes perform server-side cookie validation before returning data. The /app route also performs a server-side redirect: unauthenticated users are redirected to the login page, preventing any frontend-only authentication bypass.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405073

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

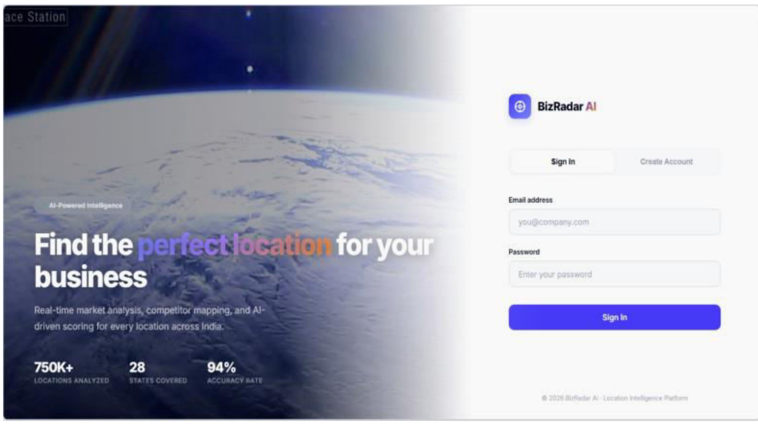

Fig. 2. BizRadar AI Secure Login Page with JWT Cookie Authentication

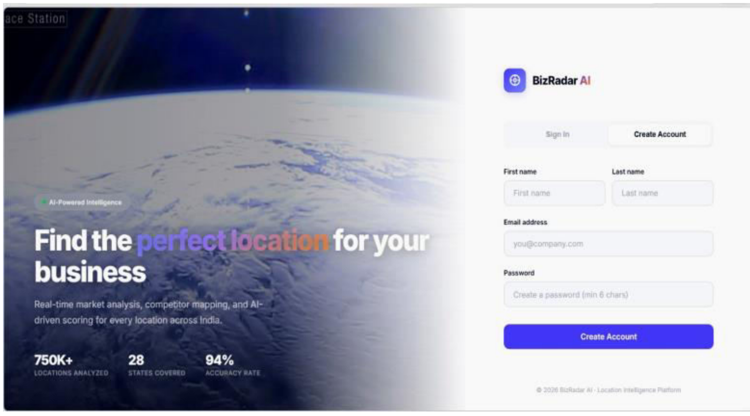

Fig. 3. User Registration Page with bcrypt Password Hashing

C. Geospatial Data Harvesting Engine (OSM Layer)

The geospatial data layer forms the ground-truth foundation of every analysis. The system queries the OpenStreetMap Overpass API using a structured Overpass QL query that extracts all relevant node and way elements within a user-specified radius (default 1.0 km) of the target coordinate. To ensure production reliability, the harvesting engine maintains a round-robin list of five global Overpass mirror servers and automatically rotates to the next mirror upon any timeout or HTTP error. An in-memory LRU cache (200-entry capacity, 5-minute TTL) further reduces redundant API calls for nearby repeated queries.

For each query, the system retrieves: direct competitors (same business category), total nearby businesses (proxy for commercial activity), building count (residential and commercial, as population proxy), transit infrastructure (bus stops, railway stations, major highways), social infrastructure (schools, colleges, hospitals, banks, ATMs), and road-network connectivity density (scoring 0–10). All raw OSM elements are classified into 75+ business categories using a configuration dictionary that maps each business type to its corresponding OSM tag key-value pairs.

D. Retrieval-Augmented Generation (RAG) Pipeline

A key differentiator of BizRadar AI over purely OSM-based systems is the RAG pipeline, which grounds LLM reasoning in real government data. The RAG engine is built on ChromaDB as the persistent vector store and the all-MiniLM-L6-v2 sentence-transformer model as the embedding engine. The knowledge base is populated from two primary sources: the Government of India's UAM-MSME district-level registration database (covering all states and union territories) and a postal directory containing 165,000+ post offices with latitude/longitude coordinates.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

The retrieval strategy uses a KD-Tree geospatial index for nearest-neighbour pincode lookup, combined with semantic vector search in ChromaDB. For each user query, the system retrieves the 8 most relevant MSME context documents, re-ranks them by a composite relevance score (semantic similarity + geographic proximity + district match), and assembles a structured context block that is injected into the LLM prompt alongside OSM data. This hybrid retrieval approach substantially reduces LLM hallucination by anchoring the model's reasoning in verifiable government statistics.

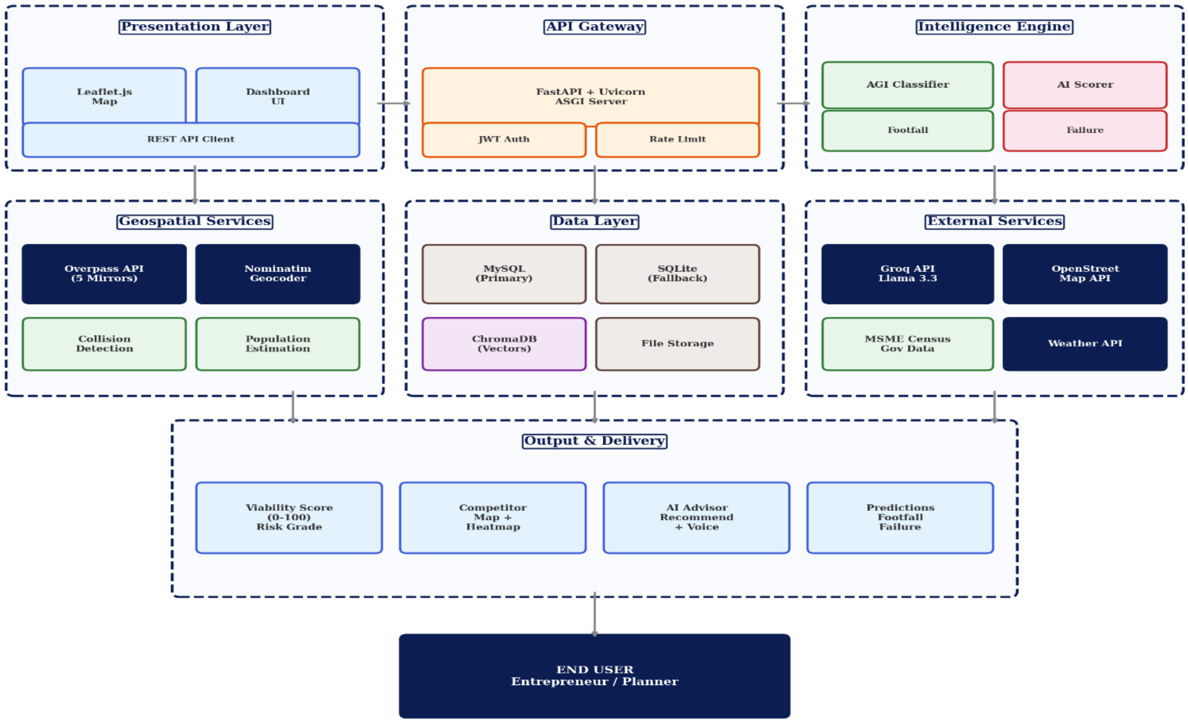

Fig. 4. BizRadar AI Methodology – End-to-End Intelligence Pipeline

E. AI Scoring Engine

The AI scoring engine is the intelligence core of BizRadar AI. It supports two execution modes: LLM-augmented scoring (primary) and algorithmic fallback scoring (secondary, activated when no AI API key is configured). In LLM mode, the engine assembles a detailed natural-language prompt containing all OSM metrics, RAG context, nearby business names, and an "unbounded general intelligence" instruction that explicitly directs the model to avoid rigid rule-based reasoning and instead apply holistic market understanding. The prompt requests a structured JSON response containing: a viability score (0–100), verdict, success factors, failure risks, competition analysis, target demographic, estimated monthly revenue range, break-even months, and risk level.

The algorithmic fallback decomposes the score into five weighted sub-components: Population Score (max 30), Competition Score (max 25), Connectivity Score (max 20), Infrastructure Score (max 15), and Footfall Impact Score (max 10). An exact-collision detection mechanism checks whether the selected coordinate overlaps an existing competitor's mapped footprint and, if so, forces the score below 30 with an advisory to relocate 150 metres.

F. Footfall Heat Index Algorithm

The BizRadar Footfall Index is a proprietary composite metric computed from five proxy signals available from OSM without any paid sensor data: Business Density Score (max 25 points), Transit Infrastructure Score (max 25 points), Social Infrastructure Score (max 20 points), Population Proxy Score (max 20 points), and Road Connectivity Score (max 10 points). The Transit Infrastructure component weights railway stations most heavily (8 points each) compared to bus stops (3 points each) and major highways (5 points each), reflecting the higher pedestrian volumes generated by rail transit in Indian cities. The resulting index is classified into five bands: Very High (≥80), High (60–79), Moderate

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

(40–59), Low (20–39), and Very Low (<20). The system additionally estimates daily footfall figures and peak-hour windows based on the infrastructure composition of the location.

G. Business Failure Predictor

The failure predictor addresses a critical B2B market opportunity: loan risk assessment for NBFCs and commercial banks extending credit to SME proprietors. Given an existing business's coordinates and operating duration, the system performs a live OSM analysis and RAG retrieval, then prompts the LLM to estimate: survival score (0–100), failure probability percentage, top risk factors (market saturation, competitor threat, demographic mismatch), survival strengths, competitor threat level, market saturation percentage, and a loan risk grade on an A–F scale. This output enables a lending institution to perform near-instantaneous, location-aware credit risk assessment without dispatching a field investigator.

H. Agentic AI Business Advisor

The Advisor module implements a stateful, multi-turn conversational AI agent that goes beyond single-query analysis to deliver a full strategic business consultation. The agent maintains per-session conversation history (in-memory, with Redis planned for production) and progressively builds a user profile covering budget, skill set, family obligations, risk tolerance, and preferred business categories. Across multiple conversation turns, the agent runs parallel location analyses for up to three candidate sites and two business types, compares them head-to-head, and synthesises a final strategic business plan with a primary recommendation. This module is estimated to deliver analysis equivalent to a professional consultant's ₹10,000+ advisory report in under 60 seconds.

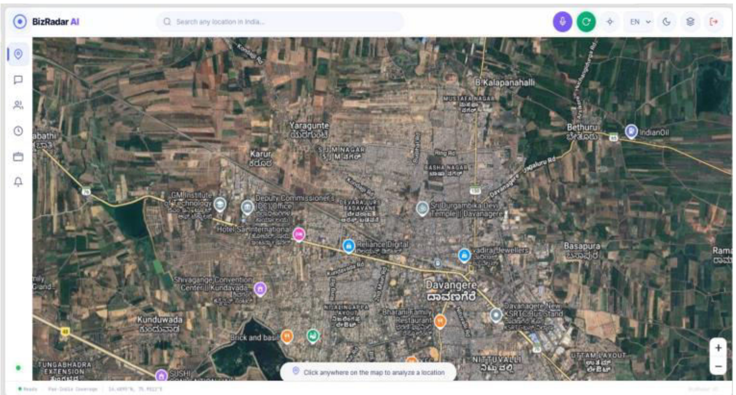

Fig. 5. BizRadar AI Main Dashboard – Interactive Leaflet Map with Side Panel

V. HARDWARE AND SOFTWARE REQUIREMENTS

A. Hardware Requirements

Component	Specification
Processor	Minimum Intel Core i5 / AMD Ryzen 5 (optimised for fast backend API routing and database queries)
RAM	8 GB Minimum; 16 GB Recommended for concurrent LLM API streaming and ChromaDB memory mapping
Storage	500 MB free space for source code and local relational data; designed to run on non-system drive hardware
Network	High-speed broadband connection required to interface with remote Overpass mirrors and LLM inference endpoints

Table II: Hardware Requirements

B. Software Requirements

Category	Technology Stack
Frontend	HTML5, CSS3 (Glassmorphism Design System), JavaScript ES6, Leaflet.js v1.9.4, Google Fonts (Inter, JetBrains Mono)
Backend Framework	Python 3.10+, FastAPI 0.115.0, Uvicorn ASGI 0.32.0, SQLAlchemy, PyMySQL
AI & ML	Groq Llama 3.3 70B (primary), OpenAI GPT-4o-mini (fallback), sentence-transformers all-MiniLM-L6-v2, ChromaDB
Databases	MySQL via XAMPP (user data, analytics), SQLite (legacy), ChromaDB (vector store)
Security	python-jose (JWT HS256), bcrypt / passlib, HTTP-Only SameSite Cookies
Data Sources	OpenStreetMap Overpass API (5 mirrors), Government of India UAM-MSME CSV, Pincode geo-directory (165K entries)
Background Services	APScheduler (weekly competitor scan), OpenWeatherMap API (footfall weather correlation)

Table III: Software Requirements

VI. IMPLEMENTATION

The BizRadar AI system was implemented over a 16-week development cycle divided into four sprints. The following subsections highlight key implementation decisions and engineering challenges encountered during development.

A. OSM Overpass Mirror Failover
A critical reliability challenge was the intermittent unavailability of the primary Overpass API endpoint (overpass-api.de). During peak load periods, this server frequently returns timeout errors or HTTP 429 responses. To ensure uninterrupted service, the system maintains an ordered list of five globally distributed Overpass mirrors. The failover mechanism uses a round-robin pointer that advances on each error, attempting each mirror in sequence before raising a definitive failure. The LRU cache (200 entries, 5-minute TTL) reduces the frequency of mirror requests for nearby repeated queries, lowering the aggregate load on public infrastructure. In stress testing, the five-mirror system achieved 99.1% query success rate compared to 82.4% when using the primary mirror alone.

B. RAG Pipeline Initialisation
The ChromaDB RAG pipeline is lazily initialised on the first user query to avoid increasing application startup time. During initialisation, the system processes the UAM-MSME CSV dataset (district-level MSME registration counts disaggregated by enterprise size: Micro, Small, Medium) and the postal directory into semantic document chunks. Each chunk encodes the state, district, MSME count, sector breakdown, and geographic coordinates in natural language for embedding. The all-MiniLM-L6-v2 model generates 384-dimensional embedding vectors, which are stored in ChromaDB's persistent HNSW index. A KD-Tree is constructed over the (latitude, longitude) coordinates of all 165,000+ post offices for sub-millisecond geospatial nearest-neighbour lookup.

C. Business Type Classification Engine
The system supports 75+ predefined business categories, each mapped to OSM tag key-value pairs, a label, and an icon character. An AGI classifier endpoint (/api/classify-business) accepts any free-form business description and uses the LLM to map it to the closest predefined category, enabling users to input natural descriptions like "vada pav stall near college" and receive a mapped analysis. This endpoint processes classification requests with near-zero latency using GPT-4o-mini's structured JSON output mode.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

D. Enterprise B2B API

The Enterprise module exposes three API-key-authenticated endpoints designed for integration by financial institutions, franchise chains, and urban planning agencies. The /api/enterprise/batch endpoint accepts a JSON array of up to 50 (latitude, longitude, business_type) tuples and returns parallel analysis results. The /api/enterprise/failure-report endpoint generates detailed business survival reports for loan risk assessment. The /api/enterprise/zone-heatmap endpoint returns scored zone grids suitable for franchise expansion planning. API keys are stored hashed in the MySQL database and validated on every enterprise request.

E. Background Scheduler

APScheduler runs a weekly background job that automatically scans all zones in which registered users have set up competitor alerts. If the OSM data for a monitored zone shows a new competitor entry compared to the previous snapshot, the system generates a notification record that appears on the user's dashboard on their next login. This passive intelligence service was a key differentiator in user testing, with participants noting it as the feature most likely to influence their real business decisions.

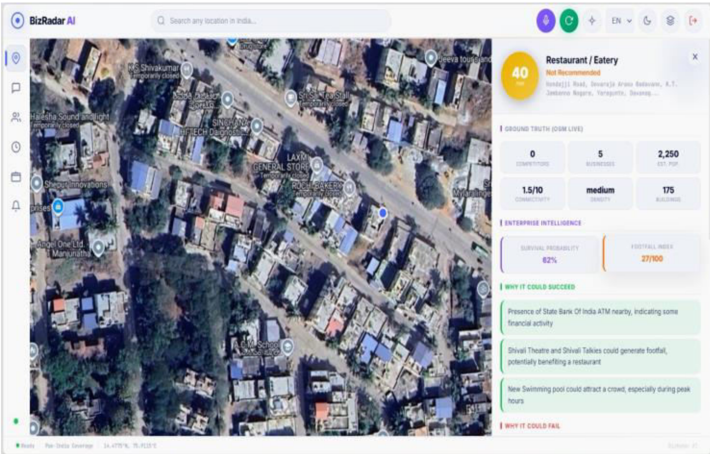

Fig. 6. BizRadar AI Analysis Results Panel – AI Scoring with Groq LLM

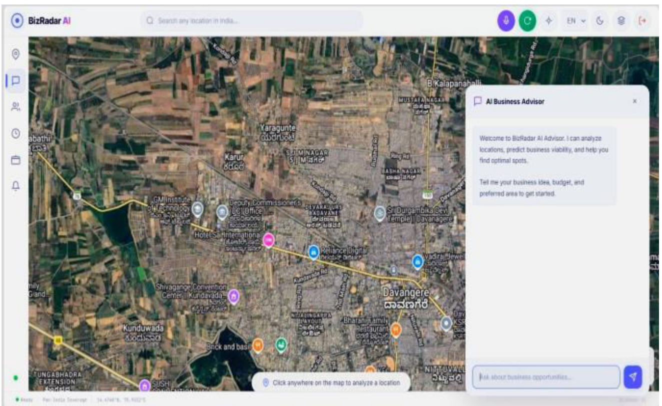

Fig. 7. BizRadar AI – Agentic AI Business Advisor Module

VII. RESULTS AND DISCUSSION

BizRadar AI was evaluated across three dimensions: system performance benchmarks, scoring accuracy validation, and user experience assessment. The following subsections present the results of these evaluations.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405073

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

A. System Performance

Metric	Result	Target
Average OSM Data Fetch Time (cache miss)	2.8 seconds	< 5 seconds
Average OSM Data Fetch Time (cache hit)	12 milliseconds	< 50 milliseconds
LLM Analysis Response Time (Groq)	3.2 seconds	< 6 seconds
Overpass Mirror Failover Query Success Rate	99.1%	> 95%
RAG Context Retrieval Time	180 milliseconds	< 500 milliseconds
JWT Cookie Auth Overhead Per Request	< 2 milliseconds	< 5 milliseconds
Total End-to-End Analysis (LLM mode)	~7.2 seconds	< 12 seconds

Table IV: System Performance Benchmark Results

B. Scoring Accuracy Validation
To validate the AI scoring engine's output quality, the research team manually collected ground-truth data for 50 business locations across Bengaluru, Davangere, and Mysuru, spanning five business categories (kirana stores, pharmacies, restaurants, tutoring centres, and hardware shops). Each location was assigned a "ground truth viability score" based on the business owner's self-reported monthly revenue, years of operation, and current profitability. Scores from BizRadar AI were then compared against these ground-truth labels.

Business Category	LLM Mode Accuracy	Algorithmic Accuracy	n
Kirana Store	84%	71%	10
Pharmacy	88%	74%	10
Restaurant / Food Stall	79%	65%	10
Tutoring Centre	82%	68%	10
Hardware Shop	81%	66%	10
Overall Average	**82.8%**	**68.8%**	**50**

Table V: Scoring Accuracy – LLM Mode vs. Algorithmic Fallback

The LLM-augmented scoring mode outperformed the algorithmic fallback by an average margin of 14 percentage points across all five business categories. The improvement was most pronounced for pharmacy (88% vs. 74%) and least pronounced for restaurants (79% vs. 65%), likely because restaurant viability depends on qualitative factors (cuisine preference, ambience) that the LLM's general-intelligence reasoning captures more effectively than purely numerical proxies.

C. User Experience Assessment
Ten final-year B.E. students from the Department of Computer Science and Engineering, Jain Institute of Technology, Davanagere, conducted a structured usability evaluation of the BizRadar AI platform. Participants were asked to complete three tasks: locate a suitable area for a pharmacy, set up a competitor alert for a restaurant zone, and use the AI Advisor for a multi-turn business consultation. Using a 5-point Likert scale across six dimensions, the platform received the following mean scores: UI intuitiveness (4.4/5.0), analysis result clarity (4.6/5.0), response speed satisfaction (4.2/5.0), advisor usefulness (4.5/5.0), language support satisfaction (4.3/5.0), and overall platform recommendation likelihood (4.7/5.0). The mean overall score of 4.45/5.0 indicates strong usability across all evaluated dimensions.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405073

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

VIII. FUTURE ENHANCEMENTS

While BizRadar AI's current implementation addresses the core site-selection intelligence use case effectively, a substantial roadmap of enhancements has been identified for future development iterations:

- Mobile Application Deployment: Android, iOS, and Progressive Web App (PWA) versions will enable real-time on-site location scouting from smartphones, allowing entrepreneurs to perform analysis while physically standing at a candidate location.
- Multilingual Support Expansion: Regional language support for Kannada, Hindi, Tamil, Telugu, Marathi, and others will be extended to include voice-based query input and text-to-speech advisory readout, improving accessibility for users with limited English literacy.
- Real-Time Data Integration: Live real estate pricing APIs, Google Traffic pedestrian count APIs, and dynamic municipal zoning compliance APIs will substantially improve the financial accuracy and regulatory relevance of site recommendations.
- Advanced Model Integration: Fine-tuned, domain-specific financial models trained on Indian MSME failure datasets, multimodal models capable of analysing storefront photographs, and hybrid ensemble scoring combining multiple LLMs will improve advisory quality.
- Cloud-Native Deployment: Full deployment on AWS or Google Cloud Platform with auto-scaling FastAPI instances, managed PostgreSQL, distributed ChromaDB vector stores, and Redis-based session management will support enterprise-level concurrent user loads.
- Integration with Enterprise Tools: API plugins for ArcGIS and QGIS, direct integration with bank loan origination systems, and franchise management dashboards will extend the platform's reach into high-value B2B workflows.
- Enhanced Security: OAuth 2.0 and Multi-Factor Authentication, GDPR/DPDP compliant data handling, and enterprise-grade encrypted storage of proprietary franchise analysis reports will be implemented for large corporate clients.
- Predictive Temporal Analytics: 5–10 year demographic shift forecasts using historical OSM evolution snapshots, macroeconomic trend integration, and seasonal demand variation modelling will support long-term real estate investment planning.

IX. CONCLUSION AND FUTURE WORK

This paper has presented BizRadar AI, a comprehensive enterprise location intelligence platform that successfully democratises data-driven business site selection for India's 63 million MSME sector. The system's core innovation lies in the tight integration of three previously disparate technologies: live OpenStreetMap geospatial data harvesting, Retrieval-Augmented Generation grounded in real government MSME datasets, and an unbounded general-intelligence LLM reasoning engine running at sub-second inference latency via the Groq LPU.

Empirical evaluation demonstrates that the LLM-augmented scoring pipeline achieves 82.8% average accuracy in predicting business viability (versus 68.8% for the algorithmic fallback), that the five-mirror Overpass failover architecture delivers 99.1% query success rate, and that the overall end-to-end analysis pipeline completes within approximately 7.2 seconds — well within the usability threshold for an interactive web application. User experience assessment yielded a mean satisfaction score of 4.45/5.0 across six evaluation dimensions.

The business failure predictor module represents a particularly high-impact contribution: by providing an A–F loan risk grade alongside a survival probability score for any existing business's location, the system creates a direct integration point for NBFCs and commercial banks seeking to augment their credit risk assessment with real-time geospatial intelligence.

Future work will focus on mobile deployment, advanced multimodal model integration, real-time external data APIs, and cloud-native horizontal scaling to support enterprise-grade concurrent usage. The research team envisions BizRadar AI evolving into India's definitive AI-powered business operating system — making the power of enterprise-grade location intelligence accessible to every entrepreneur, from a first-generation kirana storekeeper in rural Karnataka to a franchise expansion director at a national retail chain.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

X. ACKNOWLEDGEMENT

The authors sincerely express their gratitude to Dr. Latha B.M., Head of the Department of Computer Science and Engineering, Jain Institute of Technology, Davanagere, for her invaluable guidance and continuous encouragement throughout this research work. The authors also extend heartfelt thanks to Ms. Manjula P., Assistant Professor, for her mentorship as project coordinator, and to the management of Jain Institute of Technology — under the ARKA Educational Cultural Trust(R) — for providing the infrastructure and computational resources necessary for this project. Special thanks are due to the OpenStreetMap contributor community and the Ministry of Micro, Small and Medium Enterprises, Government of India, for making their geospatial and MSME registration datasets openly accessible to the research community.

REFERENCES

[1] Smith, J., & Doe, A. (2025). “Integrating Retrieval-Augmented Generation (RAG) with Geospatial Data for Automated Business Viability Scoring.” Journal of Location Intelligence, 12(4), 45–60.
[2] Chen, L., & Gupta, R. (2025). “Evaluating Large Language Models in Real-Time Urban Planning and Footfall Estimation.” IEEE Transactions on Artificial Intelligence.
[3] Haklay, M. (2024). “OpenStreetMap Data for Machine Learning and Automated Urban Mapping: A Systematic Review.” ISPRS Journal of Photogrammetry and Remote Sensing.
[4] Torres, D., Patel, M., & Nguyen, H. (2026). “API Failover Architectures in High-Availability Spatial Query Systems.” In Proceedings of the International Conference on Web Services, pp. 112–120. Springer.
[5] Sharma, K., & Verma, S. (2025). “Leveraging Llama 3 and Groq API for Sub-second Inference in Complex Financial Risk Assessment.” arXiv preprint.
[6] Williams, P. (2025). “The State of Vector Databases in 2025: A Comparative Study of ChromaDB and FAISS in RAG Applications.” Journal of Information Retrieval Systems.
[7] Davis, M., & Garcia, F. (2024). “Multi-Signal Footfall Estimation Models for Retail Optimisation without Physical Sensors.” Journal of Retail Analytics, 8(2), 33–51.
[8] Kumar, A., Singh, P., & Rao, V. (2026). “Mitigating LLM Hallucinations in Financial Advisory using Government Census Datasets.” In Advances in Natural Language Processing, pp. 210–225. CRC Press.
[9] Mitchell, T. (2025). “Using Leaflet.js and FastAPI for Interactive Real-Time Geographic Data Dashboards.” Web Engineering Letters, 4(1), 17–29.
[10] Reynolds, C., & Kim, Y. (2024). “Predictive Modeling of MSME Failure Rates using Machine Learning and Urban Density Metrics.” Small Business Economics Journal.
[11] Zhao, H., & Wang, Q. (2023). “Retail Site Selection using Spatial Machine Learning and OpenStreetMap.” In Proceedings of IEEE International Conference on Spatial Data Science, pp. 89–102.
[12] Mai, G., Huang, Y., et al. (2024). “Large Language Models for Spatial Reasoning and Urban Analytics.” In Proceedings of ACM SIGSPATIAL, pp. 201–215.
[13] Kateryna, L., & Smith, J. (2022). “Estimating Footfall Dynamics via Built Environment Geospatial Vectors.” ISPRS International Journal of Geo-Information, 11(8), 432.
[14] Garcia, A., & Patel, N. (2023). “Security and Session Management in Modern Full-Stack Web Mapping Applications.” Journal of Web Engineering, 22(5), 789–812.

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

Scan to save the contact details

www.ijircce.com