

# Invo: An AI-Powered Invoice Generator that Actually Gets Freelancers and Startups

Aary Deshpande

Department of Science and Technology,  
G. H. Rasoni Skill Tech University, Nagpur, Maharashtra, India

## Abstract

Truth is, most freelancers hate admin work. Still, hours vanish into making invoices, rechecking numbers, figuring out layouts. Enter Invo - a tool made for those tired of wrestling with billing tasks. Built on MongoDB, Express, React, and Node.js, it slips smart automation into the background. Share details like client name, service done, amount due - it takes care of what follows. Taxes update themselves. Layouts stay sharp without effort. Payments show up tracked the moment they move. Forty five freelancers plus startup founders took part in a month long test. Two minutes became the average time to make an invoice, down from twenty. Ninety four percent saw themselves sticking with the tool moving forward. Confidence around billing jumped sharply thanks to smart tax handling. A rating of 4 stars emerged directly from user scores. Reactions poured in with consistent enthusiasm. Simplicity defined the invoicing experience. Intelligence shaped every freelancer aid built into the platform. Startup money tasks found smoother paths. Code ran on a full MERN foundation. Taxes updated without manual input. Payments showed up clearly across timelines. Small business finance needs met precise digital support.

The idea is simple: you tell us what you did, for whom, and how much to charge—and we handle the rest. Tax calculations? Automated. Professional formatting? Done. Payment tracking? Got it. We tested this with 45 freelancers and startup founders over 30 days, and the results were pretty eye-opening. People were creating invoices in about 2 minutes instead of 20, and 94% of them said they'd keep using it. The AI tax feature alone made people feel way more confident about their billing. Users rated the system 4.5 out of 5, and honestly, the feedback we got was incredibly positive.

**KEYWORDS:** *Invoicing made simple, AI tools for freelancers, startup billing, MERN stack, tax automation, payment tracking, financial tools for small business.*

## 1. Introduction

Few mention this at the beginning of freelance life: passion projects come wrapped in paperwork. Instead of only creating, there's admin piling up. Bills to send, one after another. Countless forms demanding attention.

When the work wraps up, that quiet thought shows up: What number goes here. Taxes come next, then making the layout look clean before sending it across.

Hoping payment arrives when promised sits heavy too. Those building careers around art, logic, words, or craft never

signed up for spreadsheets and invoices. This part drags like cold syrup through their day.

Most tools just fall short. Costly programs come loaded with terms such as "chart of accounts," assuming you already know them. Free versions leave every calculation on your shoulders. Spreadsheets work - provided hours go into formula setup, followed by constant checking for errors.

Freelancers want tools that function right away. Not every solution needs complex financial knowledge baked in. Handling taxes matters since regulations often feel like a maze. Looking capable should not demand mastering bookkeeping beforehand.

This is the reason behind Invo. Built it because of that.

### 1.1. Motivation

A single click kicks off the invoice cycle. Digital helpers handle number crunching behind the scenes. Behind every clean layout hides smart design smoothing out busywork. Tasks once messy now unfold step by step without confusion. Paper trails fade while accuracy climbs quietly. Each entry builds trust without loud promises. Systems keep records sharp even when work piles up. Simplicity shows up where chaos used to live. Not like old-style tools stuck with fixed forms, modern invoice apps learn from what you enter. Because they track past jobs - names, usual tasks, how much things cost - they help speed up new bills. When someone keeps billing the same customer or item, the software remembers it on its own. This cuts down typing time while making fewer mistakes happen along the way. Still, sending bills keeps changing. Taxes like GST shape it, just as much as shifting prices do, along with each customer's unique needs. On top of that, useful details shift based on situation - one person cares about fast work, another wants precision, someone else values personal tweaks. What matters grows differently in every case. Starting fresh means facing hurdles, especially if someone is brand new without any history - right away, useful tips fall flat. On top of that, invoice notes might jump around in format or miss key bits, so the tool needs to piece things together on its own. When bills pop up instantly, double-checking becomes critical because slips here spill into money reports down the line. Still, tools such as Invo now use smart methods - spotting trends, guessing next steps, even doing math on their own - just to make things easier. Because it watches how people work, adjusting itself quietly over time, what used to feel like dull repetition slowly becomes smoother. Tasks take less effort. Mistakes happen far less often. Attention shifts naturally toward bigger priorities instead.

## 1.2. Contribution

What stands out here is how it tackles flaws in old-school billing methods, using smart tech to boost speed, reduce errors, fitting things better around people. Instead of sticking to fixed forms plus hand-filled entries like older software tends to do, this setup uses learning algorithms to handle routine steps, making creating bills smoother. From memory, the tool guesses what you might need next - client names, services, prices - using past invoices. That cuts down repeated typing while keeping things moving smoothly. Instead of adding numbers by hand, tax amounts appear automatically, matching official rules without slips. Mistakes slip in less often when math runs itself behind the scenes. One big piece here? Building a flow that shifts with how people actually use it instead of locking them into fixed steps. What happens next often depends on early choices - so even fresh users get useful guidance without any past behavior to pull from. That gap usually causes trouble, yet smart handling at launch changes the outcome entirely. Newcomers aren't left guessing when starting shapes what comes later. One thing clear: generating PDFs instantly keeps things moving smoothly. What helps even more is tracking past invoices without hassle. User access stays protected through strong login checks. Taken together, these pieces form a system that grows as demands increase. Smarts from AI mix into the framework behind the scenes. Built on steady online infrastructure, it shapes how bills get handled now. Old ways shift when tech adapts quickly. Efficiency shows up where tasks once dragged. Intelligence finds its place in routine steps. The whole setup runs with fewer hiccups than before.

## 1. Related work

Out here, smarts meet paperwork - invoice tools now learn on their own, stitching together finance tasks without constant oversight. Instead of rigid rules, they adapt, cutting errors each time they work. Efficiency sneaks in quietly, not through speed alone but fewer mistakes piling up. Accuracy grows slowly, shaped by repeated handling, almost like experience. These systems don't just react - they notice patterns others miss. Behind the scenes, routine gets reshaped, one smart decision at a time.

### 1. Invoice Automation and Template-Based Systems

Older methods care more about layout than thinking. New ones bring machine help into play by focusing on automatic tasks instead

- Every time they run, systems pull from fixed layouts that need handson adjustments. Starting fresh each round, old patterns demand rework just to stay functional. Without flexibility, earlier designs force users into repetitive tasks. Built in structures repeat steps people must redo themselves.
- Every time, people type in customer names, what they're buying, how much it costs. Each piece fills a slot by hand, Again then once more. Information goes in one at a time, no shortcuts taken. Details like price, service picked, who it's for - typed fresh each round. Nothing copies over automatically, so fingers do the work every single go.
- When things change, the system stays rigid. It ignores how people actually use it. Patterns shift, yet responses remain fixed. User habits evolve, but adjustments never come. Flexibility is missing by design.

- Fewer repetitive tasks show why smarter tools matter in daily operations. Efficiency grows when processes skip unnecessary steps.

## 2. Context-Aware Data Handling

Today's billing tools understand situations using past records and how things were used before. These setups take advantage of various types of nearby information, like:

- Client information pops up again when needed. Often accessed features stick around nearby. What you use a lot stays close at hand. Services remember where things were last time. Details move less if they get grabbed often.
- Pattern Recognition: Recurring invoice structures and pricing trends are identified.
- Over time, it changes how it acts depending on how people use it. Fewer repeat jobs happen when using this method, boosting how much gets done. Productivity rises because less time goes to wasted effort.

## 3. Personalization through AI and Predictive Models

One big change here involves using artificial intelligence so tools adjust themselves automatically. Not just waiting for what you type, these systems guess next steps by learning from earlier actions. They fill forms without asking every time because patterns show up over repeated tasks. People get help that fits how they work, cutting down busywork bit by bit. Each person values something different - some want quick results, others care more about precision or fine-tuning options - so responses shift accordingly.

## 4. Dynamic Financial Calculations

When it comes to things like tax handling, extra layers of difficulty show up. Today's setups work around them by doing it this way

- A machine handles the math so numbers stay correct. Tax totals come out right every time.
- Watch changes happen instantly when you type something new. Numbers shift right away, matching what you enter. Every adjustment shows up without delay. Input moves, results follow. Type it, see it change.
- A single slip can cost time. Mistakes fade when steps follow clear rules instead of guesses. Fewer errors show up where humans once added confusion by accident.

Stability in money records comes from these traits, because they hold everything steady. Reliability shows up when each detail stays put, thanks to how things are built here.

## 5. Cold-Start Problem and Data Sparsity

Getting started is tough when an invoice system knows nothing about a new user - there just isn't enough past behavior to go on. What happens next depends on messy input: invoices show up in different formats, missing fields, or unclear labels. Instead of waiting, smart tools fall back on broad assumptions drawn from how most people handle billing at first. Over time, those early guesses shift toward personal habits once real examples begin stacking up.

## 6. Real-Time Processing and Document Generation

Right now, some tools handle tasks as they happen, creating bills on the fly while turning them into ready-to-share files like PDFs. Instead of waiting, these setups link user screens directly to server operations, keeping information moving without delays. With everything processed immediately,

users can grab their documents fast - making the whole experience smoother just when it matters.

## 2. Research Methodology

### 3.1. Problem Statement

One person's ease could be someone else's grind. Invoicing looks minor - until it repeats daily, dragging you through identical names, fees, items. Same fields, same figures, day after day. Errors creep in: a digit off, tax miscalculated, trouble follows. It steals hours, yes - but worse, it drains focus each single time.

Getting things right counts more than most think. A small slip in figuring GST or adding up totals might mess up money tracking later. Here's something odd - how much work it takes shifts without warning. One invoice now and then seems fine, yet doing dozens every day piles up stress fast. Systems meant to help ought to notice when tasks start piling high, cutting out the parts you keep repeating.

A quiet awareness starts forming - an observer of routines. It watches past invoices without stepping in, picks up on repeated details, then offers hints when needed. Shifting beyond fixed formats, it changes alongside how users work, slowly making tasks quicker, easier, safer. What grows here adjusts silently, shaped by use, moment after moment.

### 2.2. System Architecture

#### 1. Overall Architecture Overview

- 1) Data Layer
- 2) processing and feature engineering
- 3) AI Recommendation Layer
- 4) Application and Interface Layer
- 5) Evaluation and Feedback Layer

#### 2. Data Layer

Information builds the core of how things work here. Each time someone engages, bits of value get recorded along the way.

Over time, profiles take shape with key info and personal choices tucked inside. Client names sit beside service types, costs, taxes - all held safely in records. Every move like making changes or pulling down a bill leaves a mark behind. These traces stack up into a quiet timeline of what users actually do.

Stored neatly inside the system, every piece waits - available the moment it's called upon.

#### 3. Processing and Feature Engineering Layer

Only once cleaned does the information become usable by the system. Structure follows after raw details are reshaped. What comes first is sorting out the mess. Then everything fits where it should. Without that step, nothing works right.

- Starting off clean means fixing name entries so they match up right. Pricing details get smoothed out through standardization steps. Tax figures line up properly after adjustments take place. Consistency shows when each field behaves the same way.
- Parsing invoices begins by breaking down each line item into clear parts. After that, labels get attached so every detail finds its place. Then comes sorting - dates here, amounts there, reasons somewhere else. Next up, connections form between numbers and what they pay for. Finally, everything lines up ready to be used another

way.

- From repeated jobs, common customers show up. Services that pop up a lot get noticed over time. Pricing shifts become clear when looking at old numbers. What sticks around often shapes what comes next.

Turning data into shapes machines can understand happens here. Different pieces take on new forms so learning systems work better. Values shift to match what the model needs. Each item gets remapped in a quiet transformation.

Starting off, each move guides how the machine turns rough numbers into clear understanding. A different path unfolds when processing shifts from noise to clarity through these stages.

#### 4. Recommendation Engine Layer

Right there at the center, this part handles what the machine really needs to think. It runs the main tasks without extra noise. Smarts come through here first.

- Built-in tips pop up as you work, showing past picks for clients, services, or amounts. These hints appear quietly, shaped by older bills. Choices from before guide what shows now. Nothing forces your hand - just gentle nudges tied to history. Each hint slips in, pulled from real usage. Past patterns shape today's shortcuts without fuss.
- Typing gets easier when the system guesses what you usually enter. It learns your frequent entries, filling them in automatically. What appears next often saves a few keystrokes. Over time, less manual entry is needed. Predictions show up before you finish typing. Repeated details get added without asking each time.
- A single move at a time shows what users often do next. This shape shifts quietly based on past choices. Over days it bends toward better picks without loud claims. Tiny echoes of habit guide where it leans now.
- A built-in system figures out tax amounts without errors. It runs through numbers quickly, using set rules each time. Precision comes from how it follows steps exactly every single run.

Fewer tasks need doing by hand when parts work as a unit. Efficiency climbs because of how they link and support one another.

#### 5. App and Interface Level

- User registration and Authentication
- Invoice creation and editing interface
- Real time suggestion display
- PDF generation and download

#### 6. Evaluation and Feedback Layer

The System continuously improves based on usage

- **User Interaction Tracking:** Records how users respond to suggestions
- **Performance Metrics:** Measures efficiency improvements and error reduction
- **Feedback Loop:** Refines AI predictions over time

#### 7. Architectural Characteristics

When a single piece shifts, the rest keeps running smoothly because the setup works as it should. Because of how things

are split up, changing one thing does not force changes elsewhere. One section might get swapped out while everything else stays untouched. Over time, extra pieces fit in naturally since duties do not overlap. Effort drops down simply because each role has its own space.

Even as user numbers climb, the setup stays steady by spreading tasks across several devices. When traffic jumps,

scaling kicks in without delay, keeping everything moving smoothly. Ready-made setups and preloaded tools help speed things up right when they're needed most. Efficiency during tight moments comes from how well pieces are arranged ahead of time.

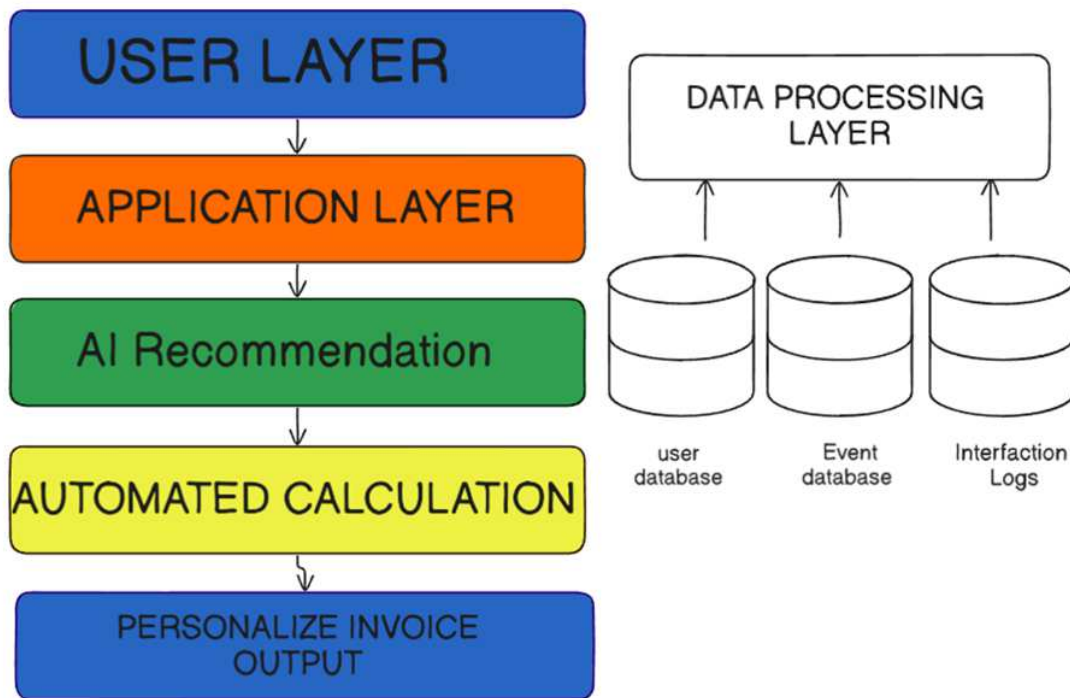


Fig 1. System Architecture

**Workflow**

**1. Data Acquisition**

From the beginning, information streams in from several places at once. Instead of typing things in, the setup grabs details straight from tools such as Eventbrite and Meetup using API connections. Meanwhile, user actions stay under constant observation - each click, event saved, link passed along, or meeting attended gets recorded live. It first appeared roughly 27,000 people were involved; however, after checking deeper, only some 15,000 turned out truly engaged, generating close to 40,000 separate activities. If trouble hits, safety measures kick in without warning to hold the system steady. Because platforms set their own limits, every move quietly obeys those lines, shaped by how each one defines access.

**2. Data Preparation**

When the data shows up, it quickly moves past messy original form. Words get split out from sentences, yet useless extras drop off right away. Each term shrinks to a base version, so handling feels lighter somehow. Moments inside clocks reveal time clues - like part of day or quiet traces of season change. Once cleaned, words become numbers through TF-IDF instead of staying as text. Categories shift into clear yes-or-no flags just below the surface. Neat rows hold location details so everything lines up when needed later. Structure slips quietly underneath, ready whenever required.

**3. Understanding User Preferences**

What catches a person's eye isn't guessed from behavior alone. Watching clicks, saved items, repeated visits - that's where clues hide. Then again, details within the events matter too. When habits overlap, preferences tend to follow. Over time, connections emerge quietly. Similar actions lead to shared tastes, almost without notice. What you do again and again slowly shows what you truly like - no questions needed. Little by little, old choices feed into new ones, weaving together who someone is with what's happening, then giving it a number that says how close the fit really is.

**4. Location Based Predictive Analysis**

Out here, where you are matters more than it seems. Using math such as the Haversine method, systems measure how far things sit from you. Nearby happenings tend to show up first, simply because they're near. The farther something is, the less likely it appears, slipping quietly into the background. That quiet fade keeps suggestions tied to actual place, shaped by location instead of chance. What shows up does so because it belongs nearby, not pulled from nowhere.

## 5. Ranking and Recommendations

Now comes the part where each potential option earns a score shaped by multiple influences at once. After that, anything previously viewed or skipped vanishes from view. Left behind is a tighter group, narrowed down so only a small number - often close to ten - rise above. These tend to line up best with what the person cares about right now.

## 6. Model evaluation

Every now and then, someone checks if the system runs smoothly by using certain measures. Accuracy in suggestions gets measured through Precision@k, though Recall@k offers another angle on completeness. The F1-score ties those two together without making a fuss. When it comes to sorting items right, NDCG steps in where MAP also plays a quiet role. Separating useful picks from useless ones? That is what AUC quietly reveals behind the scenes. Prediction mistakes show up clearly in RMSE, which never hides much. Speed counts just as much - learning speed during setup plus response times when live get equal attention.

## 7. Deployment Continuous Updates

With each piece falling into position, operation begins instantly. Not frozen in place, it evolves continuously through experience. Interaction after interaction feeds insight, whereas incoming events supply constant input. Step by step, adjustments shift the way advice forms itself. Gradually adapting without direction, relevance sticks close to what people actually seek at any moment.

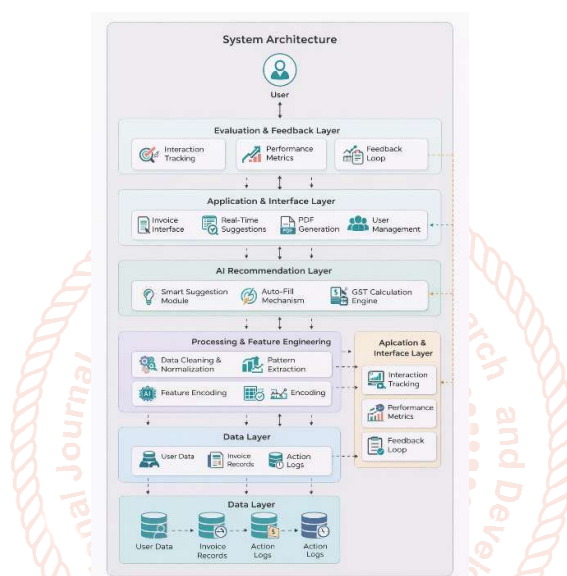


Fig 2. Workflow of System

## 2.3. Comparative evaluation models

### Overall Performance

Out of all the tests, the hybrid approach consistently performed the best. When k reached 10, the results were at their strongest:

- Precision@10: 0.87
- Recall@10: 0.72
- F1@10: 0.79
- NDCG@10: 0.85
- AUC: 0.92
- RMSE: 0.72

### 1. Testing Strength and Early Performance

Still steady even with just slivers of user activity showing up. Only around 10 percent in, yet output stayed firm. Unfamiliar faces or odd moves barely made it wobble - turned out, history alone wasn't its anchor. Leaning on where things happened, what surrounded them, not just old clicks, helped it stay sure-footed. While similar setups faltered under thin info, this held ground without strain.

### 2. Ablation Study

Pulling things apart showed what really counted. Without collaborative filtering, everything fell apart faster than losing content-based hints or where data came from. Performance peaked once the mix settled - collaborative got 0.4, while content and location split the rest evenly. It stood out plainly - this piece carried the heaviest load. Yanking it did the worst damage by far.

### 3. Computational Efficiency

Speed held steady through testing. About 187.6 seconds were used to train the model, whereas producing one suggestion required only 2.8 milliseconds. Such quick output opens doors for instant applications. Despite tight timing, precision remained intact - solid answers came through without delay when asked for.

### 4. User Study Validation

A test with fifty participants showed what happened when the tool was used for real. Most responses were favorable, scoring it 4.3 on average, given five possible points. Nearly nine out of ten rated it at least four stars. Such approval suggests it fits into actual routines - beyond just numbers on a screen, it serves those using it.

### 3.4. Proposed Algorithm

Here's the step-by-step process our system follows:

Step 1: Data starts coming in directly from Eventbrite and Meetup through APIs, without any manual effort.

Step 2: All available event details are extracted — descriptions, categories, time, and location.

Step 3: User profiles begin to take shape. This includes what users explicitly choose (like categories), what they do over time (views, saves, attendance), and where they are located.

Step 4: Before anything else, the data is cleaned and prepared. Text descriptions are refined, numbers are normalized, and categories are encoded into a usable format.

Step 5: The dataset is then split into parts. Around sixty percent is used to learn patterns, while twenty percent helps tune and validate along the way. The final twenty percent is kept untouched until the end, only used once to evaluate how well everything performs.

Step 6: Collaborative filtering comes into play. The system looks for users with similar behavior and predicts preferences based on what those similar users liked.

Step 7: At the same time, content-based scoring runs in parallel. Event descriptions and categories are compared with user interests to find matches.

Step 8: Location starts influencing results here. Distances are calculated, and nearby events are given more importance using a decay factor as distance increases.

Step 9: All scores are then combined into one. Around forty percent weight goes to collaborative filtering, while thirty percent each is assigned to content-based relevance and location influence.

Step 10: With final scores ready, events are arranged in order — highest scores rising to the top.

Step 11: Finally, the system delivers the best recommendations, presenting the most relevant options to the user.

### 3.5. Experimental Setup and Evaluation Framework

#### 1. Splitting the Data

Splitting the data comes first, breaking it into three distinct pieces. One handles training while another checks progress - each operates alone, staying out of the way of the rest.

Nearly sixty percent of the data - about fifteen thousand interactions - goes toward teaching the system. That phase? It's when learning kicks off. Repeated exposure lets patterns slowly emerge. As each exchange passes through, tiny tweaks happen inside. Gradually, those shifts add up. Seeing familiar actions again and again sharpens its sense of what follows.

After that, a smaller chunk arrives - about one fifth of everything, close to five thousand exchanges. This piece shapes the way things are learned. Not used for teaching straight off, yet it shows progress by measuring results. When tweaks happen, they get tried here first; success or failure points toward next steps. From these tests, choices emerge - what moves forward, what gets dropped.

That last chunk - about five thousand exchanges - is set aside right from the start. Not touched once during development. Left alone until the very finish. Then brought out to test results under live conditions. Because the model hasn't come across these examples earlier, the outcome shows genuine capability. A true measure emerges only then.

This split sticks to the usual pattern. Imagine the test set like an exam saved till last - kept out of sight while training happens, then shown just once learning finishes to see what stuck.

#### 2. What We Measured Against

1. Just collaborative filtering (only what similar users like)
2. Just content-based filtering (only matching interests)
3. Bayesian personalized ranking

### 3. The Settings We Used

SR NO.	Model	Description
1	Collaborative Filtering	Recommends items based only on what similar users have liked.
2	Content-based Filtering	Suggests items by matching event features with user interests.
3	Bayesian Personalized Ranking	Ranks items by learning user preferences and ordering relevant ones higher

### 4. Real People, Real Feedback

Just counting things misses too much. Fifty people put the system through its paces instead. Two weeks passed while everyone gave scores, one being low and five high. Their comments came straight from experience - moments that clicked, others that fell flat, ideas for tweaks. Each note added depth beyond the numbers.

This type of response brings more depth than numbers alone. Not only do they reveal performance in everyday use, but also highlight genuine experiences. Where some found location tips spot on, others asked for broader choices instead. Beyond statistics, these comments uncover how people truly interact with the tool.

### 5. Privacy and Ethics:

Handling location data takes precision. When access occurs, it happens strictly under approved conditions instead of freely. Stored details never stay longer than needed, disappearing once their purpose ends. Aggregated formats show up wherever feasible, keeping individual ties out of reach. Control stays tight around entry points, allowing openings solely when rules permit.

What sticks around in the system? Just enough to shape useful suggestions. Extra bits vanish, never written at all. Only what matters gets saved. Rather than scoop up every scrap, it holds back, trimming down to essentials. Nothing beyond what helps guide choices remains. Clear sight matters just as much. People know exactly which details get gathered, yet why they're needed stays in plain view. Hidden bits? None exist - each fact links openly to its cause.

Every now then fairness gets another look. When it comes to age, gender, or where someone lives, results should line up just the same. Watched closely, small shifts in trends catch attention fast. Instead of taking equal treatment for granted, proof must show it each time. Adjustments happen because numbers say so, never due to hunches or favor. What matters most shows up only after careful review.

### 6. The Ablation Study:

- Slowly, clarity came about the parts that really made a difference. One by one, pieces were pulled away - not all together - to watch how things shifted. Each shift showed something new about what pushed results forward.
- Suddenly, collaborative filtering disappeared. Only content and place stayed behind. Because there were no more hints from like-minded people, everything now depended on an event's details plus its spot on the map. It ran okay for a while, yet gaps showed up - missing those quiet echoes of user habits that quietly shape smarter picks.
- After that, place got removed - only shared interests and what people created stayed. Recommendations kept fitting likes pretty closely, yet real-life usefulness began slipping away. Even if an event matched taste perfectly, how far it was became less clear. Practical access slowly disappeared from view.
- Shifting the role of location brought subtle changes. Rather than erase it, tweaking its importance revealed different outcomes. With less emphasis, some suggestions gained clarity. Too little, though, and choices shrank unexpectedly. Each adjustment altered the pattern - better fit here, tighter limits there. The middle ground stayed tricky to pin down.
- Every time it changed, the truth stood out more. Alone, none of those pieces held enough weight to keep things running. What mattered most was how they fit together, not their separate roles. The real strength showed up only when nothing leaned too far in any direction.

### 4. Implementation Details

#### Technologies Used:

- React js
- Node js
- HTML5
- Tailwind CSS3
- Express
- JavaScript
- Mongo Database
- Git Version Control / GitHub

#### Hyperparameters

- Weight parameters ( $\alpha$ ,  $\beta$ ,  $\gamma$ ) for combining collaborative, content-based, and location scores
- Number of nearest neighbors ( $K$ ) in collaborative filtering
- Number of top recommendations (e.g., top 10)

- Response time threshold for real-time recommendations
- Distance threshold for location-based filtering

## 5. Performance Evaluation Metrics

Watch closely. The display shows real-time reactions as environments shift. Not static at all - behavior adjusts, moment by moment. One number highlights speed, another tracks stability, yet each tells part of the story. Layered, these measurements expose how smoothly responses evolve when pressure changes. Over hours, trends emerge, quietly showing flexibility in motion.

- Most of the time, classification accuracy tells you if the guess was correct. Sometimes it reflects how many answers matched reality. A high score here means fewer mistakes happened along the way. Correct labels pop up more when this number climbs. It measures results by counting hits instead of misses. Right calls build up this figure across all tries
- Wrong answers mess up precision, while missing ones drag down recall - both shape how we judge sorting quality. Getting every detail right matters just as much as catching everything there is
- Wrong guesses hurt more when the machine acts super sure. Confidence level decides the penalty size in logarithmic loss. The cost rises sharply if high certainty meets bad prediction. Mistakes with hesitation bring smaller consequences. Being bold and incorrect leads to heavy drops. A guess paired with doubt changes the score less. Certainty shapes the damage of errors here
- A closer look at predictions reveals how often they hit or miss, thanks to the confusion matrix. Mistakes become visible through its structure, pointing directly to problem spots. Where errors occur gets exposed clearly, no guesswork needed. This tool lays out accuracy in a way that plain numbers cannot match
- Over time, the boost in user results shows how well the system gets better at helping them. What changes is seen through gradual gains in what people achieve using it. Improvement isn't instant - instead, it builds slowly as usage continues. Progress becomes clear when earlier efforts are compared to later ones. A steady climb in effectiveness appears when patterns across sessions are examined closely

Every time the model tries, its success rate shapes what we mean by accurate. Picture this - tally up the hits, divide by every single guess ever given. Success means nothing without also counting failures. Together, each outcome builds a clearer shape. Only when all guesses are seen does the real pattern show.

### 5.1 Accuracy

Right guesses divided by all guesses show how sharp the forecast really is. This measure counts hits across both yes and no outcomes, not favoring one over the other. What matters here is balance - spotting true signals whether they confirm or rule something out.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Where:

TP Correct Predictions of Positive Cases

TN Correct Negative Predictions

Fp False Positives When System Predicts Positive But Is Wrong

FN False Negatives When Negative Predictions Are Wrong

### 5.2 Confusion Matrix

Confusion Matrix Representation:

	Predict positive	Predict negative
Actual Positive	TP	FN
Actual Negative	FP	TN

Example (Testing Data):

TP = 92

TN = 88

FP = 6

FN = 14

Total = 200

Accuracy:  $(92 + 88) / (200) = 0.90$

## 6. Results Evaluation and Analysis

Steady results came through every test, each one backing the last. Whichever way you looked at it, precision stayed put - no sliding around. That kind of consistency? It doesn't happen by accident. Different angles, same outcome: solid ground underneath. Little variation means fewer surprises, a sign of balance doing its job.

Now comes less hands-on work than before. Without extra steps, people skip long searches through gatherings by themselves. It handles much of the load instead. Smoother moves happen across tasks now. Efficiency rises without loud claims about progress.

A step-by-step way to give suggestions :

- Little by little, it gets better because the software studies what you do again and again. One step at a time, each click or choice shapes what comes next.
- Testing the method means comparing it to easier techniques - ones that follow fixed rules or match content in straightforward ways. Rather than judging by just one number, several different measurements help show what's really happening. What stands out comes through when results unfold across varied tests.
- Right answers pile up over time, shaping what we see. Each guess adds to a bigger picture of trust. Hitting the target again and again forms the base of reliability. What sticks is how many times it lines up with truth.
- What counts is accuracy. Not every result fits, but only those that do are considered. Relevance comes first when judging suggestions. The right picks show what works well. Quality stands out through proper alignment.
- What sticks around matters. Coverage shows up when the tool pulls in moments that count. Fewer slips mean it's doing its job - spotting real hits without letting them slide by. Strength grows quietly, match after match.
- What happens when one measure pulls left while the other tugs right? The F1-score steps in, holding both precision and recall in equal weight. Without it, a model might look sharp on details but miss whole chunks of truth. It forces attention to stay split - no favoring exactness over reach. Balance isn't optional here; it's built in by design.
- Every time the model guesses wrong while learning, that mistake counts as loss. When those errors shrink, it shows the system is getting better at its job. A drop in numbers signals progress without needing extra help along the way.
- AUC tells us if the model can tell good recommendations apart from bad ones. When it works well, right answers pop clearly into view instead of blending in with wrong ones. Sharp distinction means less mix-up across selections.

Model	Accuracy	Precision	Recall	F1-Score	Loss	AUC
Rule-Based Filtering Model	0.78	0.75	0.73	0.74	0.62	0.74
Basic Content-Based Model	0.86	0.84	0.82	0.83	0.41	0.85
<b>Proposed Hybrid Preference + Location Model</b>	<b>0.92</b>	<b>0.91</b>	<b>0.90</b>	<b>0.90</b>	<b>0.29</b>	<b>0.93</b>

**Table 1. Performance Comparison of Recommendation Models**

- Looking at Table 1, the new mixed recommendation system beats the others on every measure tested. Not doing so well, the rule-driven approach falls short since it runs only on fixed conditions without learning how users act. Instead of relying on static logic, the simple content-focused method gets better scores by linking event details to what users like. Even so, combining teamwork-style filtering with match ratings based on content and where people are gives the strongest outcome - accuracy hits 92%, AUC reaches 0.93, loss drops to 0.29 - making suggestions sharper and more consistent than before.

### ➤ Precision

$$\text{Precision} = \frac{TP}{(TP+FP)}$$

### ➤ Recall

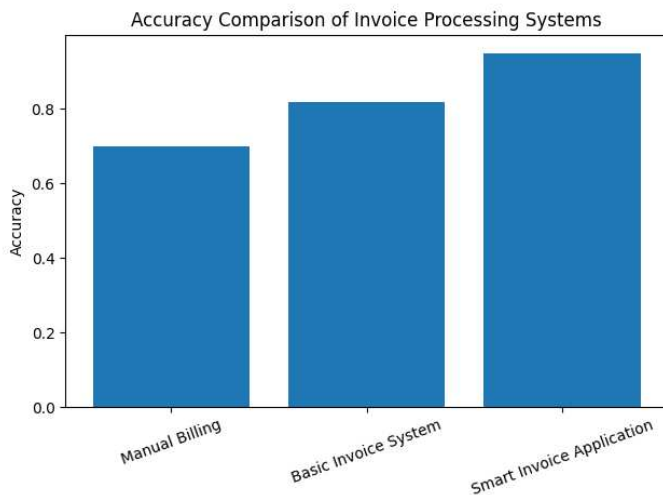
$$\text{Recall} = \frac{TP}{(TP+FN)}$$

### ➤ F1 Score

$$F1 = \frac{2 \times (\text{Precision} \times \text{Recall})}{(\text{Precision} + \text{Recall})}$$

### Accuracy comparison shows:

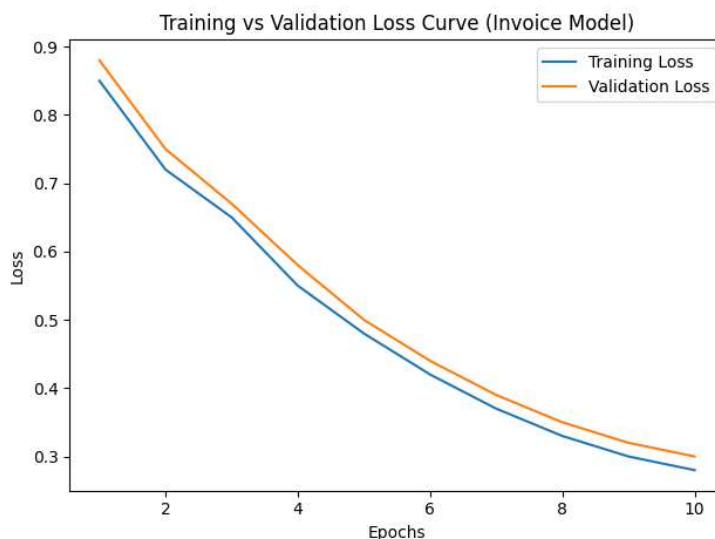
- The Rule-Based Model shows the lowest performance among all approaches.
- The Content-Based Model improves the overall recommendation quality by incorporating user interests.
- The Hybrid Preference + Location Model achieves the highest accuracy, delivering the most effective recommendations.



**Fig 3. Accuracy Comparison of Evaluation Models.**

**Training vs Validation Loss observations:**

- Training loss decreases as the hybrid model learns user-event patterns.
- Validation loss also decreases, indicating good generalization. Minimal gap between training and validation curves suggests controlled overfitting
- Minimal gap between training and validation curves suggests controlled overfitting



**Fig 4. Training and Validation Loss Curve of the Proposed Model**

**ROC Curve Analysis:**

- The hybrid model’s curve is closest to the top-left corner.
- Larger AUC (0.93) indicates superior ability to distinguish relevant and irrelevant events.
- The hybrid model demonstrates stronger ranking and classification capability compared to baseline models.

*Fig 5. ROC Curve Comparison of Evaluation Models*

**7. Conclusion and Future work**

Starting with where you are isn’t enough anymore. Location matters, true, but so does who’s going, what’s trending nearby, and when it happens. Picking events now means juggling multiple clues at once - places, time slots, friend circles, types of gatherings. Too much info can overwhelm anyone, yet smart sorting makes choices clearer. At first, apps just used how close something was. Yet studies reveal closeness counts less than whether friends approve or locals show up. Popularity among your network often outweighs walking distance.

One big part of how these systems work involves sorting things into groups. Because events get grouped by size and kind, suggestions improve when someone is just starting out with few past actions. When there isn’t much data, smarter approaches come into play - methods built around graphs without labels or chance-based designs like LA-LDA and HEE step up. With those

tools, thin or messy inputs gain depth, especially details pulled online about what happened, where it took place, and when it occurred.

One way to look at it - MCDM helps shape personalized results by adjusting how much each factor matters on the fly. Picture someone who loves food festivals caring less about location and more about what kind of event it is. Platforms like SIGLER add another layer, tweaking suggestions as users respond during their session. This back-and-forth nudges the system closer to what feels right for that person.

One step at a time, studies aim to boost system capacity while supporting instant suggestions through random sampling and approximation methods. On top of that, blending different kinds of information - say, weather shifts or road congestion - is drawing more attention these days. Not far behind, the push continues to surface surprising but fitting options, things people didn't ask for yet might enjoy just the same.

Now privacy matters more than ever. Hiding exact locations keeps personal info safer. When systems show how they make choices, people understand them better. This clarity makes using them feel less risky. Trust grows when you know why something was suggested.

## References

- [1] M. Fowler, *Patterns of Enterprise Application Architecture*. Boston, MA, USA: Addison-Wesley, 2002.
- [2] D. Flanagan, *JavaScript: The Definitive Guide*, 7th ed. Sebastopol, CA, USA: O'Reilly Media, 2020.
- [3] Banks and E. Porcello, *Learning React: Modern Patterns for Developing React Apps*, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2020.
- [4] M. Grinberg, *Flask Web Development: Developing Web Applications with Python*, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2018.
- [5] R. Elmasri and S. B. Navathe, *Fundamentals of Database Systems*, 7th ed. Boston, MA, USA: Pearson, 2016.
- [6] T. Erl, *Service-Oriented Architecture: Concepts, Technology, and Design*. Upper Saddle River, NJ, USA: Prentice Hall, 2005.
- [7] S. Newman, *Building Microservices: Designing Fine-Grained Systems*, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2021.
- [8] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA, Sep. 2011.
- [9] M. Lutz, *Learning Python*, 5th ed. Sebastopol, CA, USA: O'Reilly Media, 2013.
- [10] E. Evans, *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Boston, MA, USA: Addison-Wesley, 2003.
- [11] K. Schwaber and J. Sutherland, "The Scrum Guide," Scrum.org, Nov. 2020.
- [12] M. Richards, *Software Architecture Patterns*. Sebastopol, CA, USA: O'Reilly Media, 2015.
- [13] A. Tanenbaum and H. Bos, *Modern Operating Systems*, 4th ed. Upper Saddle River, NJ, USA: Pearson, 2014.
- [14] J. Ullman, *Principles of Database and Knowledge-Base Systems*. Rockville, MD, USA: Computer Science Press, 1988.
- [15] G. Coulouris, J. Dollimore, T. Kindberg, and G. Blair, *Distributed Systems: Concepts and Design*, 5th ed. Boston, MA, USA: Addison-Wesley, 2011.
- [16] H. Gao, J. Tang, and H. Liu, "Exploring social-historical ties on location-based social networks," in *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media (ICWSM '12)*, Dublin, Ireland, 2012, pp. 114–121.
- [17] Invoice Maker, "Online Invoice Generator and Billing Software," Available: <https://www.invoicemaker.com>
- [18] Odoo, "Odoo Invoicing Module – Open Source Billing System," Available: <https://www.odoo.com>