

TEAM 12

MEMBERS: Sam Benson Devine · Jack Halverson · Tobias Knight · Qiqi Li · Yehan Wang

Inside the Recommender

How TheaterIQ turns ratings, metadata, and timing into a Match Score for every film × audience × slot.

STAGE 1
ALS Affinity
Who likes what

STAGE 2
XGBoost Re-rank
In what context

OUTPUT
Match Score 0–100
For every slot

ACTION
Greedy Scheduler
Screens × slots

THE PROBLEM

A single weekend is a thousand small bets

An independent operator with four to eight screens and a dozen bookable films has to make hundreds of small decisions each week: which film on which screen, at what time, for which audience. Get those right and the auditorium fills. Get them wrong and you pay studio minimums on empty seats.

Existing ticketing tools — Veezi, CineSys, Vista — report on what happened last week. They do not tell a programmer what to try next week. That is the gap TheaterIQ closes.

STAGE 1 — WHO LIKES WHAT

ALS builds a taste map

The first stage uses **Alternating Least Squares** on roughly 100,000 MovieLens ratings. ALS learns a 32-number taste profile for every viewer and a matching flavor profile for every film. When a viewer's profile aligns with a film's profile, the model predicts they will rate it highly.

Viewers are then grouped into eight demographic segments (age band × gender). Averaging profiles within a segment gives one representative taste vector per audience — the *latent affinity signal* that feeds Stage 2.

STAGE 2 — IN WHAT CONTEXT

XGBoost adds the real world

Affinity alone is not enough. A family film can match a segment beautifully and still flop at 10pm on a Tuesday. Stage 2 is an **XGBoost** ranker that takes the segment taste vector, the film's taste vector, its genre flags, TMDb metadata (popularity, runtime, budget tier), and the target time slot — then predicts the probability of a positive outcome. That probability, scaled to 0–100, is the **Match Score**.

THE MATCH SCORE, EXPLAINED

A single 0–100 number for every combination of film × audience segment × time slot.

75 + High confidence. Book it, no promotion needed.

50 – 74 Playable with promotional support.

Under 50 Weak match — consider dropping or swapping.

Scores are calibrated on a held-out slice of ratings, so a 70 at Riverview means the same thing as a 70 at Trylon.

WHAT IT OUTPUTS

A schedule, not a dashboard

The scored grid feeds a **greedy scheduler** that walks every (screen × slot) cell and picks the best available film. A theater archetype multiplier — suburban multiplex versus urban arthouse — nudges segment weights so the same model specialises to each venue without retraining.

Each slot comes out of the scheduler labeled **high-confidence**, **promotional-support**, or **consider-dropping**. A programmer sees not just what to book, but which slots need a targeted email or a matinee discount to fill.

HOW IT IS EVALUATED

Honest metrics, known limits

Predictions are tested on a **time-based holdout** — the newest 20% of ratings are withheld from training, so the model is graded on predicting a future it never saw. Stage 2 currently reports AUC (how well it separates positive from negative outcomes) against a rating ≥ 4 target.

Rating timestamps stand in for showtimes in this first pass. The synthetic performance layer attaches plausible occupancy and revenue to real scraped showtimes, allowing end-to-end validation without requiring proprietary ticketing data.

The Little Guys in the Twin Cities

How Riverview, Lagoon, and Trylon actually book films — and where TheaterIQ earns its keep.

HOW INDEPENDENTS GET FILMS

The booking call

An independent like Trylon, Riverview, or Lagoon does not pick freely from a catalog. Each film is governed by a **licensing agreement** with the studio or distributor's booking rep, negotiated title by title.

The core lever is the **film rental rate** — the share of ticket revenue paid back to the studio. Wide releases in opening week run **70–80%** to the studio. Independent and arthouse titles run closer to **35–50%**, which is why arthouses can survive on smaller crowds. Studios also set **minimum run lengths** (often 2–4 weeks for blockbusters) and may require the largest auditorium for opening weekend.

For Trylon (single screen, 100-odd seats, nonprofit), booking is almost entirely repertory and arthouse — no wide-release pressure. Riverview and Lagoon both mix first-run indies with selective studio holdovers; their edge is playing a new A24 release at week three while a chain has moved on.

HOW SHOWTIMES GET BUILT

Scheduling by hand, every week

Once a film is booked, showtimes are laid out against running time, **15–20 minute turnover windows**, and peak-demand slots. Friday and Saturday evenings belong to the strongest titles; matinees draw families and retirees; late-night is reserved for genre, cult, and repertory.

At Trylon, Riverview, and Lagoon this is almost always done **manually in a spreadsheet**, informed by decades of neighborhood knowledge, patron email chatter, and the owner's gut. That instinct is real and valuable — but it does not scale, does not generate a forecast, and cannot tell you whether moving Auditorium 2 from a mid-run indie to a prestige holdover on a Tuesday would lift the weekly gross.

This is the exact moment a **Match Score recommendation** earns its seat at the table: not replacing instinct, augmenting it.

THE DATA GAP

What the big chains have that the little guys don't

AMC · Regal · Cinemark	Riverview · Lagoon · Trylon
Loyalty program tracking individual viewing & spend	Door counts, POS receipts, email-list intuition
ZIP-code demographic overlays around every site	Neighborhood knowledge built over decades
Nielsen, Movio, and paid data-provider contracts	Box Office Mojo and The Numbers — public data only
Revenue-per-seat dashboards for dynamic scheduling	A weekly spreadsheet and a programmer's memory
A/B tests on pricing, promos, and messaging	One shot per week, no control group

THE BUSINESS CASE — IN ONE LINE

Independents do not need a data team. They need one trusted recommendation per screen, per week, that explains itself. TheaterIQ is that recommendation.

Rip-and-replace? No. TheaterIQ sits on top of Veezi, CineSys, or a spreadsheet.

Who approves? Always the human programmer. Every recommendation is reviewed before it goes live.

What does it cost to try? Far less than a single weekend of empty auditoriums at studio minimums.

WHERE IT PAYS OFF FIRST

The soft January window. Match Scores surface repertory, event cinema, and foreign titles likely to draw loyal audiences during the industry's weakest month.

The mid-run holdover decision. Should Lagoon hold the A24 release a third week, or give the screen to the incoming release? The scheduler scores both.

The mid-week underperformer. When observed occupancy drops 20+ points below the Match Score, the system flags the gap and identifies which segment to target with a promotional brief.

Instinct still matters. TheaterIQ makes it auditable.