


From seeing risk to acting on it

STUDENT PERFORMANCE TRACKER

The Command Center showed who was at risk. The Tracker decides who to contact today, why, and leaves the draft ready – early intervention turned into an automatic routine.



STATUS	STACK	MODEL	BASE
Production	Next.js · Firebase · Slack	6 triggers	reuses Command Center

01

THE PROBLEM

VISIBILITY WASN'T ACTION

The Command Center [Ch. 01] solved the *seeing*: in a single view, the risk of the whole cohort ordered by urgency. But between seeing the red and doing something there was still a manual chasm. Each intervention required the mentor to open the student's profile, read the history, decide on the message and write it by hand – and repeat it, student by student, every morning.

Risk detection was a static snapshot: it said *who* was struggling today, not *what pattern* had led them there nor whether anyone had already intervened before.

No memory of the intervention remained. A mentor could reach out twice about the same thing, or forget to follow up on a case left open the week before.

Prioritizing under time pressure was done by eye: with dozens of students in yellow, which five do you write to before the first session?

And outreach—the step that actually changes the student's trajectory—was precisely the one with the most friction. The bottleneck was no longer information; it was the manual work of turning it into a sent message.

02

THE OBJECTIVE

AUTOMATE DETECTION AND OUTREACH

PRIMARY OBJECTIVE

Move from passive visibility to automatic early intervention: have the system detect risk patterns, prioritize the students who most need contact today, and leave outreach one click away.

As secondary objectives: break risk down into **nameable patterns**—not an opaque score, but six concrete triggers a mentor recognizes—; keep an *intervention history* per student so that no action is duplicated or lost; reduce the morning routine to a single artifact readable in minutes; and draft the first version of the outreach message automatically, leaving the mentor the final judgment.

Success was defined up front: each morning the mentor opens a single view, sees the five priority students with the explicit why, and can trigger contact without writing from scratch again.

03

THE ARTIFACT

THE "MORNING COFFEE" DASHBOARD

The product was designed around a single moment of use: morning coffee. The mentor opens **Morning Coffee**, sees the **five priority students** of the day with the *trigger* that fired each case, a **Quality Score** that measures the student's recent health, and a **Slack draft** already written and ready to send. What used to be dozens of profiles opened by hand is now resolved on a single screen.

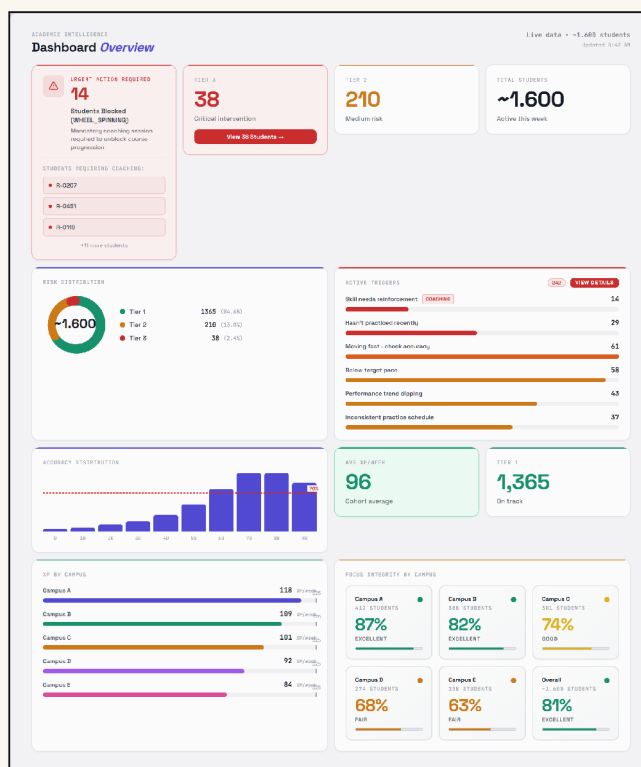


FIG. 1 – "MORNING COFFEE" DASHBOARD · TOP-5 PRIORITY + TRIGGER + QUALITY SCORE + SLACK DRAFT · THE MORNING ROUTINE IN A SINGLE VIEW

04

THE STUDENT CARD MEMORY OF THE INTERVENTION

If Morning Coffee answers "who do I attend to now?", the student card answers "what has happened to them and what have we done?". Each student —identified only by their **R-/Y-/G-** code, with no personal data— has an **intervention timeline**: which trigger fired, which outreach was sent, what happened afterward. The memory that used to live in the mentor's head is now recorded and queryable.

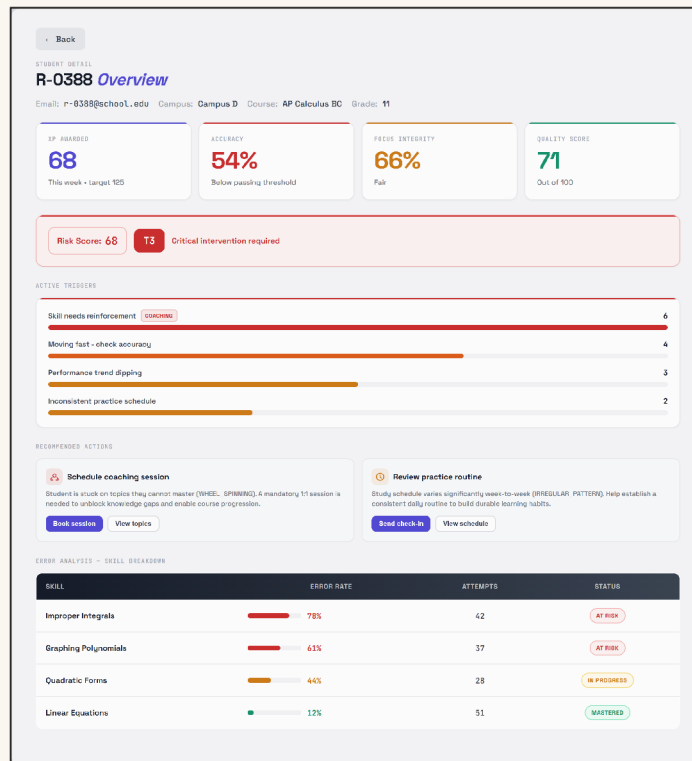


FIG. 2 – STUDENT CARD · INTERVENTION HISTORY IN A TIMELINE · NO CONTACT IS DUPLICATED OR LOST

05

THE MODEL

SIX TRIGGERS, ONE RISK SCORE, ONE OUTREACH TIER

The heart of the Tracker is a pattern-detection model. Six independent *triggers* watch distinct signals in each student's activity; when one or more fire, they feed a **Risk Score** that classifies the student into a **tier**, and the tier decides the form of the **outreach** — from a light touch to an immediate-contact alert.

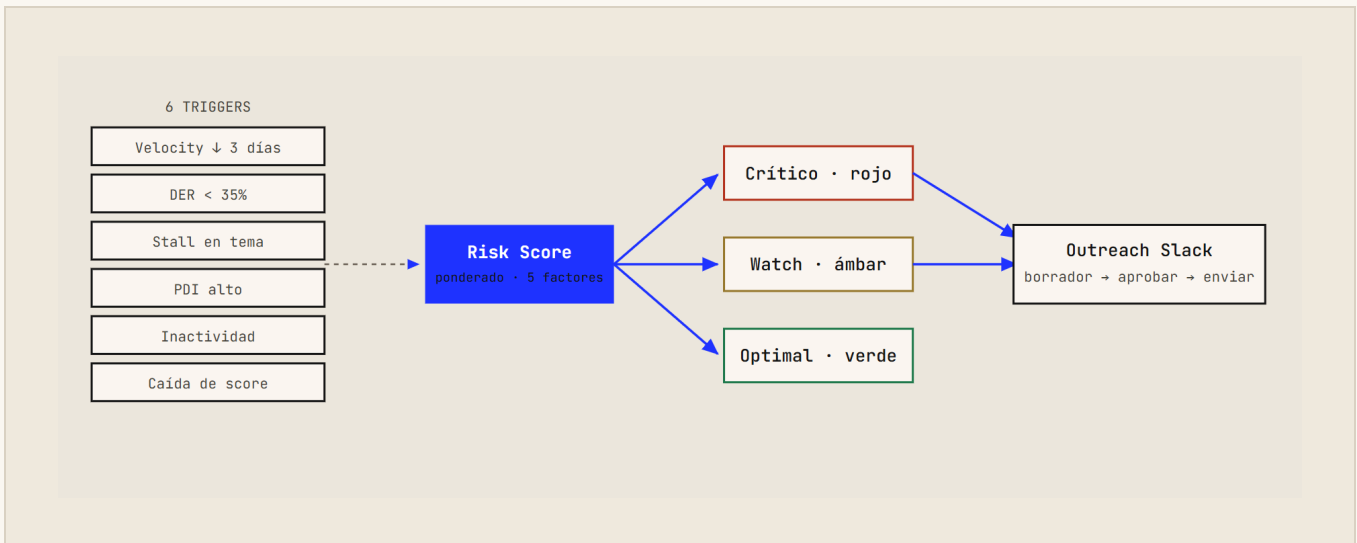


FIG. 3 — MODEL FLOW: SIX TRIGGERS → RISK SCORE → TIER → OUTREACH

TRIGGER	WHAT PATTERN IT DETECTS
Stall	The student stops advancing — no measurable progress for several days in a row.
Velocity drop	The rate of advancement falls below the standard expected for their cohort.
Precision decay	Accuracy collapses within the session — a signal of fatigue or guessing.
Guessing	A pattern of fast, erratic answers that betrays random responses.
Knowledge debt	Accumulation of unmastered foundational topics the student carries.
Disengagement	A sustained drop in activity — less time and fewer tasks than their baseline.

It's not an opaque number: each top-5 case arrives with the trigger that fired it, so the mentor understands the why before writing the first word.

The six triggers are the model's fixed structure; the exact thresholds of each are calibrated against the cohort's observed behavior and live centralized so they can be recalibrated without touching the detection logic.

06

BUILD

AI-FIRST PROCESS ON AN EXISTING BASE

The Tracker didn't start from scratch: it was built **on top of** the Command Center's visibility layer, reusing its sync pipeline and its Firestore.

Development ran with AI as the primary co-developer —trigger-model architecture, implementation, refactors—, while the decision of **what signals to measure and how to translate them into outreach** was always my own. Carried through to production on that base, with the Slack integration to push the message drafts directly to the mentor's channel.

The draft generator was the piece that closed the loop: taking the trigger and the student's context and producing an editable first message resolved the friction the Command Center had left open. The mentor keeps the final judgment —edits and sends—, but no longer starts from a blank page.

07

THE RESULT

STATUS AND LEGACY

In production, connected to the real cohort, emitting its priority top-5 each morning with drafts ready. The Tracker completed the arc the Command Center had begun: from *seeing* risk to *acting* on it without repeated manual work.

Detection model	6 triggers → Risk Score → tier
Build	AI-first on an existing base · reuse of sync + Firestore
Integrations	Practice API · Firebase · Slack · Vercel
Artifacts	Morning Coffee · Student card · Outreach draft
Status	In production

LEARNINGS

- **Product.** Visibility isn't value until it becomes action. The leap from the Command Center to the Tracker was closing the last meter: from visible risk to a sent message.
- **Design.** A nameable pattern —six triggers the mentor recognizes— builds more trust than a perfect but opaque score. The why matters as much as the who.
- **Process.** Building on a previous portfolio product —reusing its sync and its Firestore— made it possible to reach production by reusing infrastructure instead of starting over.