

Your Customers Are Not the Average



Liam Bolling
CEO, Rehearsals

We build AI twins of individual people to
predict what your customers will do next.

Hi, I'm Liam.

- Microsoft and Tumblr.
- 7 years at Google, across Maps and Search.
- Recently Google Labs/Deepmind, shipped the first Gemini models.
- And for a while, I worked in this building.

A lot of memories in this building.

Let me tell you a story from
Google Maps.


Why do people love Maps?

Because it's the most accurate data on earth.

And a shocking amount of that accuracy comes from people fixing the map.

Move a pin. Correct an address. Millions of edits, from regular users.


But we kept seeing changes like this.



123 Example St, New York

the pin on the map

Suggested change



123 Example St, New York

same exact address but...

3,000 miles away in California

Not once. Constantly.

So these were bad users, right?

Maybe they were bad actors, trying to break the map?

Maybe they had no idea what they were doing?

Either way, we had a chart for it.

And the label on the chart was **bad users**

What were they doing?

Six months. A lot of charts. No answer.

It took dinner with a friend, (and complaining too much)

people keep moving addresses on google maps. i don't get it.

if they're changing an address, did they just move?

why would they move a public address when they're the ones who moved?

they probably thought that's the only way to change their home address.

They thought they were
changing their address.

But they were changing the map. For everyone.

They weren't bad users.

They were just using our product.

We built it for the average.

If we could have asked the thousand people who'd just moved, we'd have known before we shipped

How do you make decisions?

Every decision is a bet on what people will do next.

Do you ask what the **average** person would do?

Do you ask what an **average** persona would think?

You're betting on **averages**. Not on people.

EVERY COMPANY IN THIS ROOM HAS A VERSION OF THIS

If we raise premiums 8%, who leaves, and why?

Who opens the new premium card in the first 45 days, and why them?

Which patients stop refilling after 90 days, and why?

Two questions in every one: *what will happen, and why.*

We don't know. It takes too much time and money to find out.


HOW DO WE ANSWER THESE TODAY?

We ask ChatGPT.

Open a chat window. Type the question. Get a *“confident”* answer in four seconds.

If I raise Disney+ streaming service price by \$3, what happens and why?

AVERAGE ANSWER




"17% churn is significant.
You should not raise the price."

Wait.

What actually
happened?

WHAT ACTUALLY HAPPENED



Far more people just paid than the average would ever tell you.

This is real. This is what happened when Disney+ raised its price.

Your customers are not the average.

A language model is the mean. Your 10,000 customers aren't. They're individual people.

How do you predict what
your customers will do next?

and know why?

METHOD 1

Send an AI to interview them.

Accuracy	● Excellent. Real people.
Speed	● Month(s).
Cost	● A new invoice per question.
What-ifs	● Hard.

METHOD 2

Simulate a population.

Accuracy	● Low. Not your customers.
Speed	● Minutes.
Cost	● Cheap-ish.
What-ifs	● Confusing.

METHOD 3

Wait for AGI?

Accuracy	?
Speed	??
Cost	???
What-ifs	???

We need prediction
at the speed of simulation with the
why of every individual customer.


How do we solve this? Well...

“People who are really serious about software should make their own hardware.”

- Alan Kay

You need to own the
whole supply chain.


That's what we do at
Rehearsals



One person. One twin.
One AI model. One outcome.


Own all four. Then hold every link to a standard.

Person → Twin → AI model → Outcome



The AI twin has to say and do
what the individual human does.

Person → Twin → AI model → Outcome



**Models are built to be the mean average.
So we built our own, to get closer to a human.**

Engram, benchmarked with Google DeepMind. 2,000+ interviews, 22 categories. 83.5 overall vs. mid-50s for state of the art (SOTA) frontier models.

Person → Twin → AI model → Outcome

91%

Streaming price increase

98%

Flight demand after a fare
increase

97.6%

iPhone model mix vs.
published sales

72%

picked the winning ad
(3.5x ROAS)

**The model has to predict what people
actually do in real life.**

Predicted first, then observed. And any answer can be traced back to the real person it came from.

The decision isn't
in the average.

It's in the 10,000 people your dashboard calls bad.

3 things to take back

People are not a mean average. Stop asking the mean average what they'll do.

Strange behaviors aren't bad. They're an invitation to ask why.

Whoever you buy prediction from, make sure you can audit down to the real person.

Those weren't bad Google Maps users.

I just never asked them why. Ask your users why, and do it at Rehearsals.



Liam Bolling

CEO, Rehearsals

Liam@RunRehearsals.com

We build AI twins of individual people to predict what your customers will do next.