

Research & MI All-Hands

April 2nd, 2015

© Google 2015 - Confidential and Proprietary

REDACTED FOR PUBLIC FILING & ABRIDGED

Redacted

Ex. No.
UPX0205

1:20-cv-03010-APM

GOOG-DOJ-12363161

Many Research teams Q4 2014

- Aimee (Anna)
- Google Brain (Jeff)
- Descartes (Ray)
- DeepMind (Demis)
- Machine Intelligence (JG)
- Strategic Technologies (Andrei)
- Google Research (Alfred)
- Replicant (James)



which door will she follow?

Priors
How many people followed each?

Google Confidential and Proprietary

A simple but powerful algorithm for “the choice problem” is the “Priors algorithm”.

The idea is to score the doors based on how many people took it.
In other words, you rank the choices based on how popular it is.

This is simple, yet very powerful: it is one of the strongest signals for much of Google's search and ads ranking!
If we know nothing about the user, this is probably the best thing we can do.



which door will she follow?

Priors

How many people followed each?

Personalization: refined priors

Our Twist

who went through each door?

*what **actions** describe them?*

Google Confidential and Proprietary

We bring two twists to the traditional heuristic.

Instead of attempting to describe -- through a noisy process -- what each door is about, we describe it based on the people who took it.

We can do this at Google, because at our scale, even the most obscure choice would have been exercised by thousands of people.

When a new user walks in, we measure their similarity to the people behind each door.

This brings us to the second twist, which is that while describing a user, we don't use demographics or other stereotypical attributes.

We simply use a user's past actions to describe them and match users based on their behavioral similarity.