



When shall the meeting
be scheduled?

Henrik Grundmann and Nicolas Bär, Doodle AG

Why ask the machine when you have polls?


- Doodle polls big step forward from email ping-pong, but organizer not omniscient -> back and forth often still happening
- Doodle transitioning into subscription business -> professional scheduling tools for business customers
- Introduce HELPER to increase poll success - not an automated system

What makes a good time suggestion?

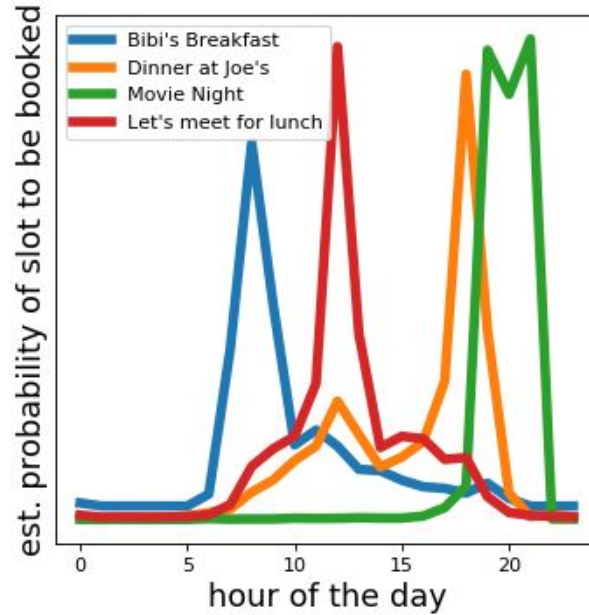
Accepted by the organizer and often voted by participants

- Very varied user base
- Complex to cater to everyone, first focus on premium customers
- Subscription business still young, almost no history on customers
- Use simple models based on aggregated data and strong features

Feature: Event times by timezone



Feature: Titles and descriptions



Creating suggestions from features

- simple models on top of pre-built features
- Features mostly build as probabilities -> naive multiplication
- Random Forest etc. to estimate $p(\text{time is good}|\text{features})$
- Gather proxy training data from past polls by calculating features on creation -> if user adds slot to poll, it is probably not a bad one

Limits of proxy data and the need to go live

Several issues with backtesting data


- Not clear how good proxy data is (no actual user interaction)
- Availability seemed irrelevant (deeper analysis -> blockers)
- Can not measure actual poll success as no voting

Need to go live asap

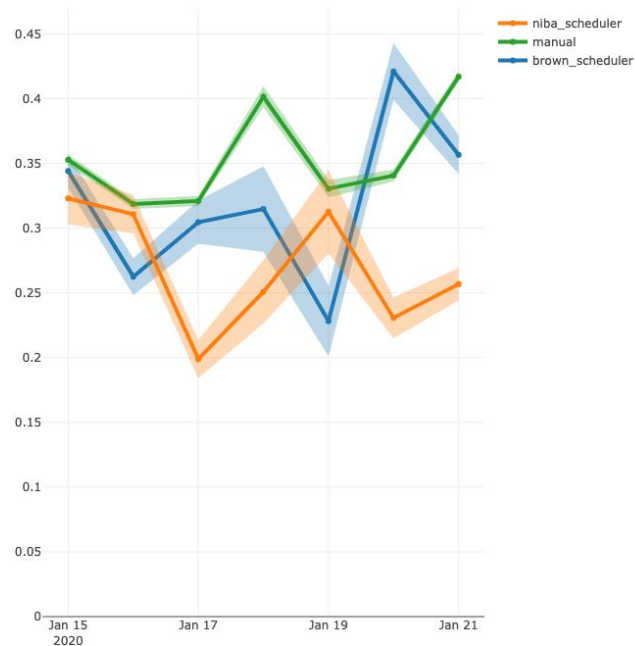
- first prototype not successful in user testing → additional filtering on top

First successful results from silent launch

fraction of polls that accepted suggestions



fraction of upvotes in polls that added suggestions



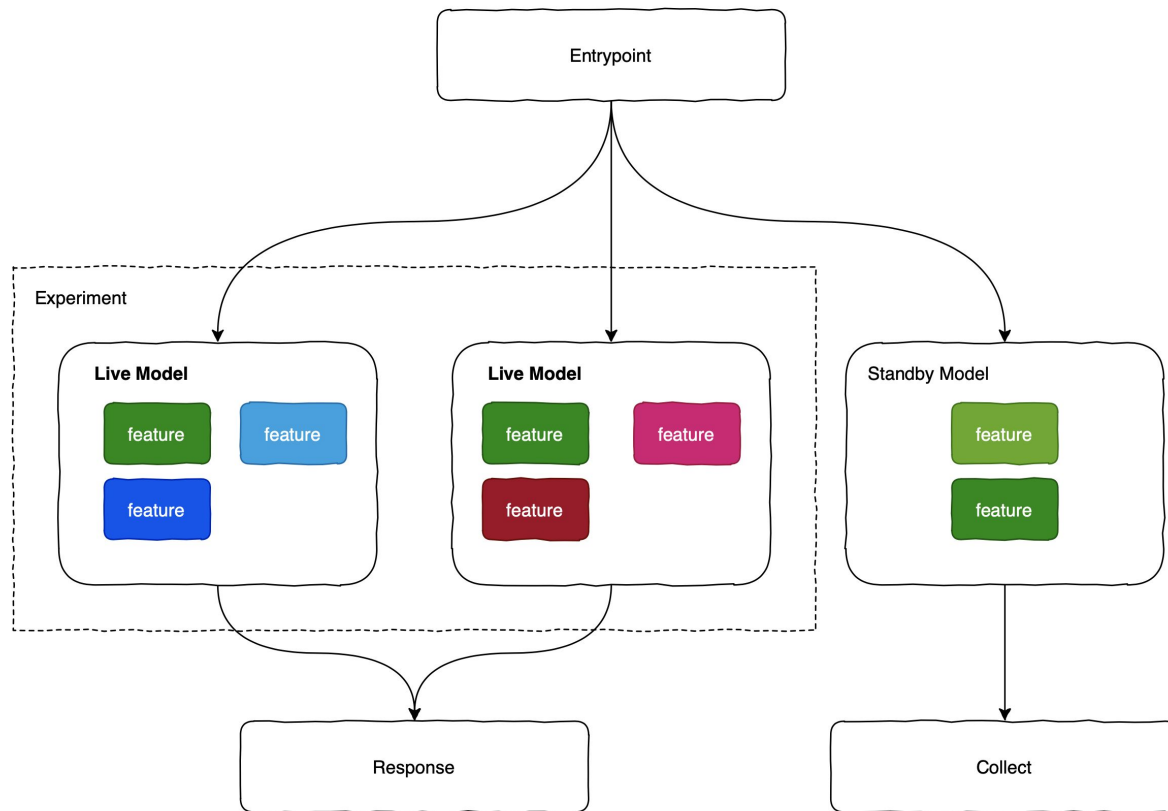
How to improve performance?

- add more personalized features
 - typical usage of users instead of typical usage by country
 - typical meeting times of user-pairs
- currently testing several neural net approaches for end-to-end models that avoid pre-training of features

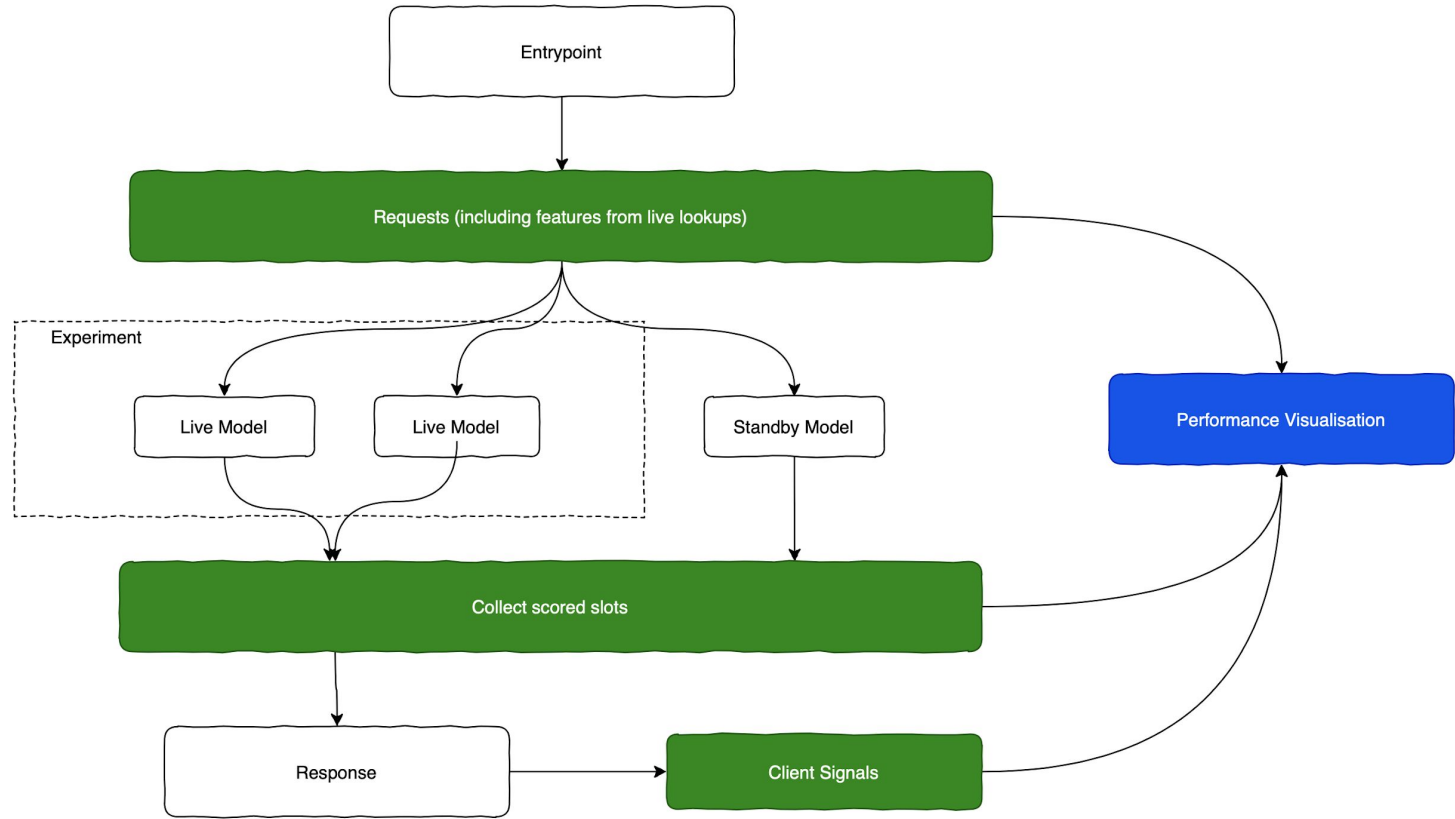
Technical Implementation

Continuous learning (also for engineers)

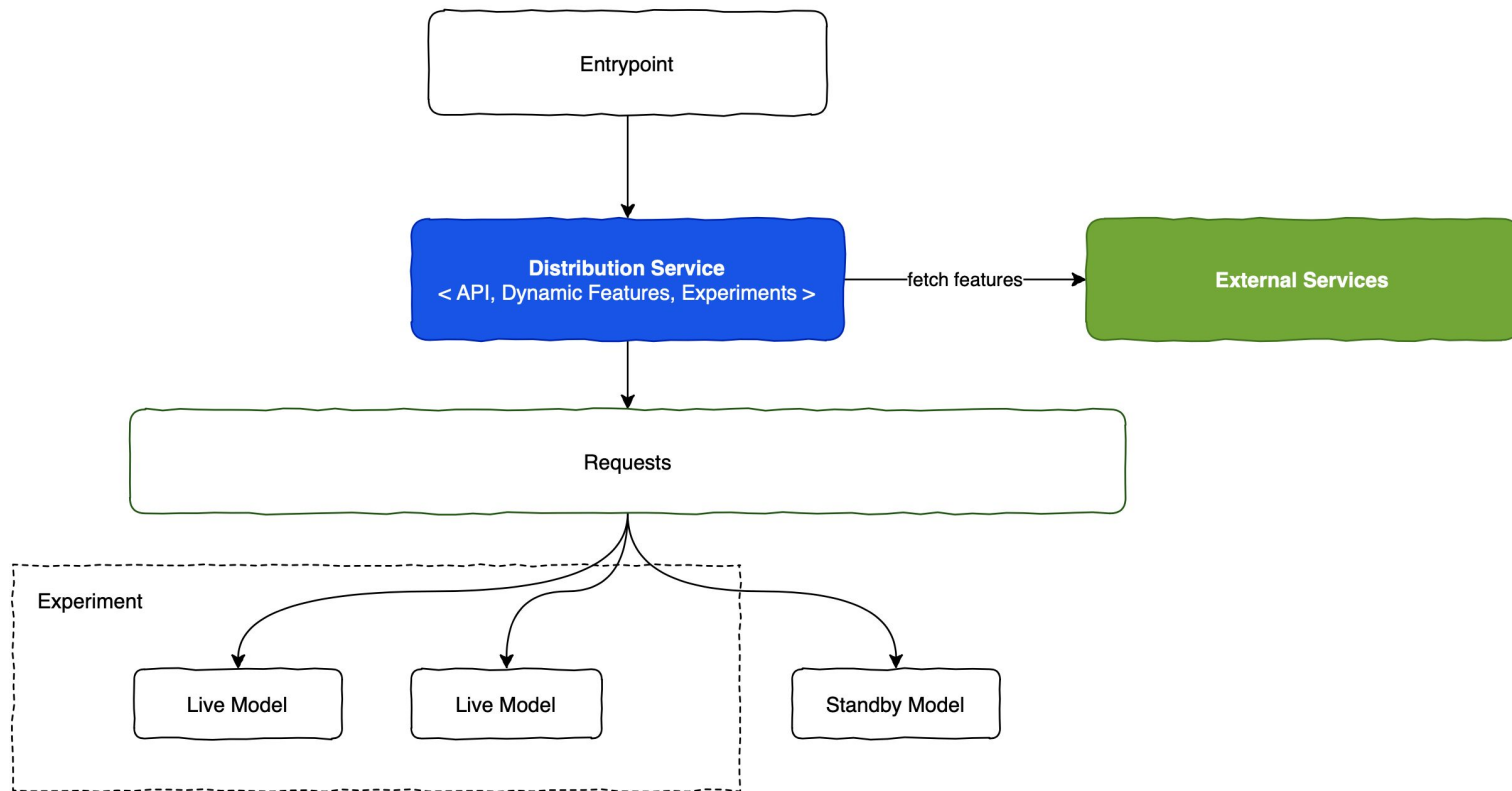
Run multiple models in production at the same time



Record request, response and interaction for model training & performance



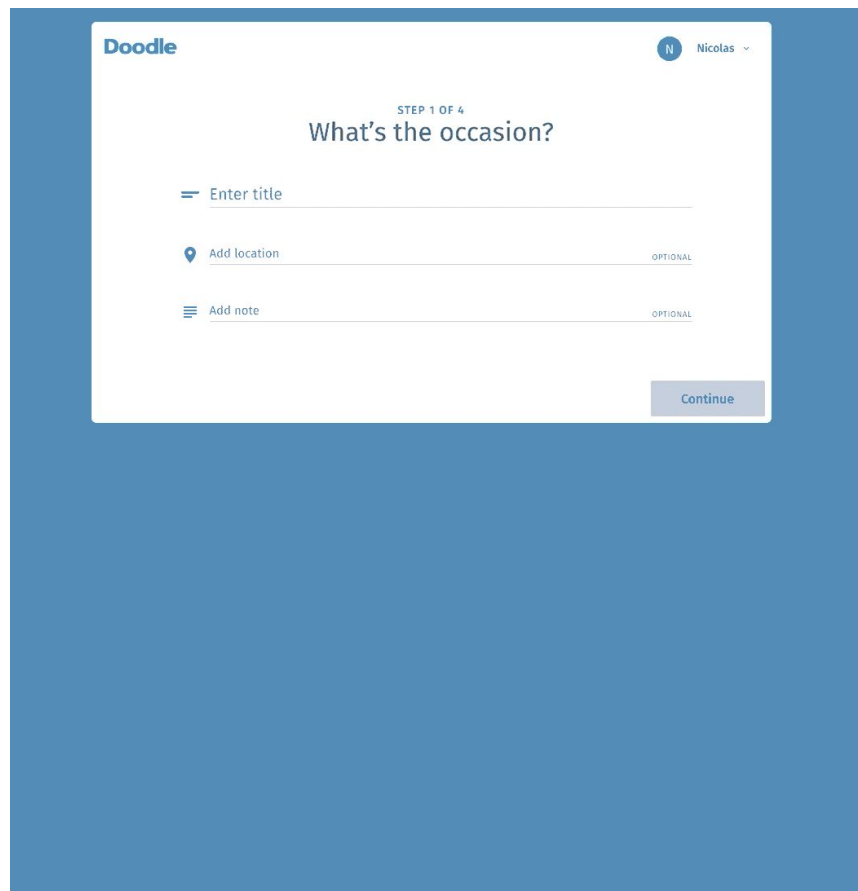
Suggest in real-time and flexibility to integrate into different products



Integrations and Outlook

if you have a hammer everything looks like a nail

When is the best time?

The image shows a screenshot of the Doodle poll creation interface. The interface is set against a solid blue background. At the top left, the 'Doodle' logo is visible. At the top right, there is a user profile icon with the letter 'N' and the name 'Nicolas'. The main heading in the center is 'What's the occasion?'. Above this heading, it says 'STEP 1 OF 4'. Below the heading, there are three input fields: 'Enter title', 'Add location', and 'Add note'. Each field has a small icon to its left (a horizontal line for title, a location pin for location, and a list icon for note). To the right of the 'Add location' and 'Add note' fields, the word 'OPTIONAL' is written. At the bottom right of the form, there is a grey button labeled 'Continue'.

When you know the best time to meet, you can just book it

aml-d-prep

You created this private channel today. This is the very beginning of the **aml-d-prep** channel.

 Set a description  Add an app  Add people to this private channel








Today



Nicolas Bär 2:01 PM

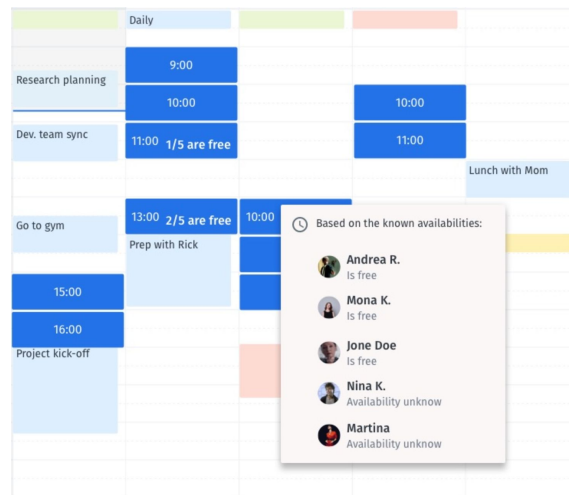
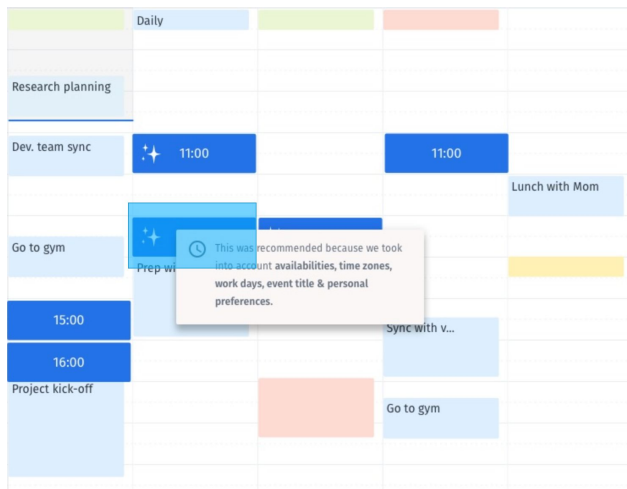
joined aml-d-prep.

/doodle_staging-meet meet @hg about Suggestion Service Architecture next week for 30min

 **B** *I*      

What's next?

- Better explain suggestions to the users
- Foundation for new products and workflows that speed up the scheduling process.



Doodle

Made with ❤️ in Belgrade, Berlin, Tel Aviv, and Zürich
and help from:



kubernetes



Features: Time of poll creation

