

A hand in a dark sleeve points towards a glowing network diagram. The diagram consists of white lines connecting various nodes, some of which are bright white circles. A search icon, a magnifying glass with a colorful ring, is overlaid on the hand's finger. The background is dark with some bokeh light effects.

How to build cool & useful voice commerce applications (such as Alexa & Google Home)

Victoria Livschitz, Founder

February 2020

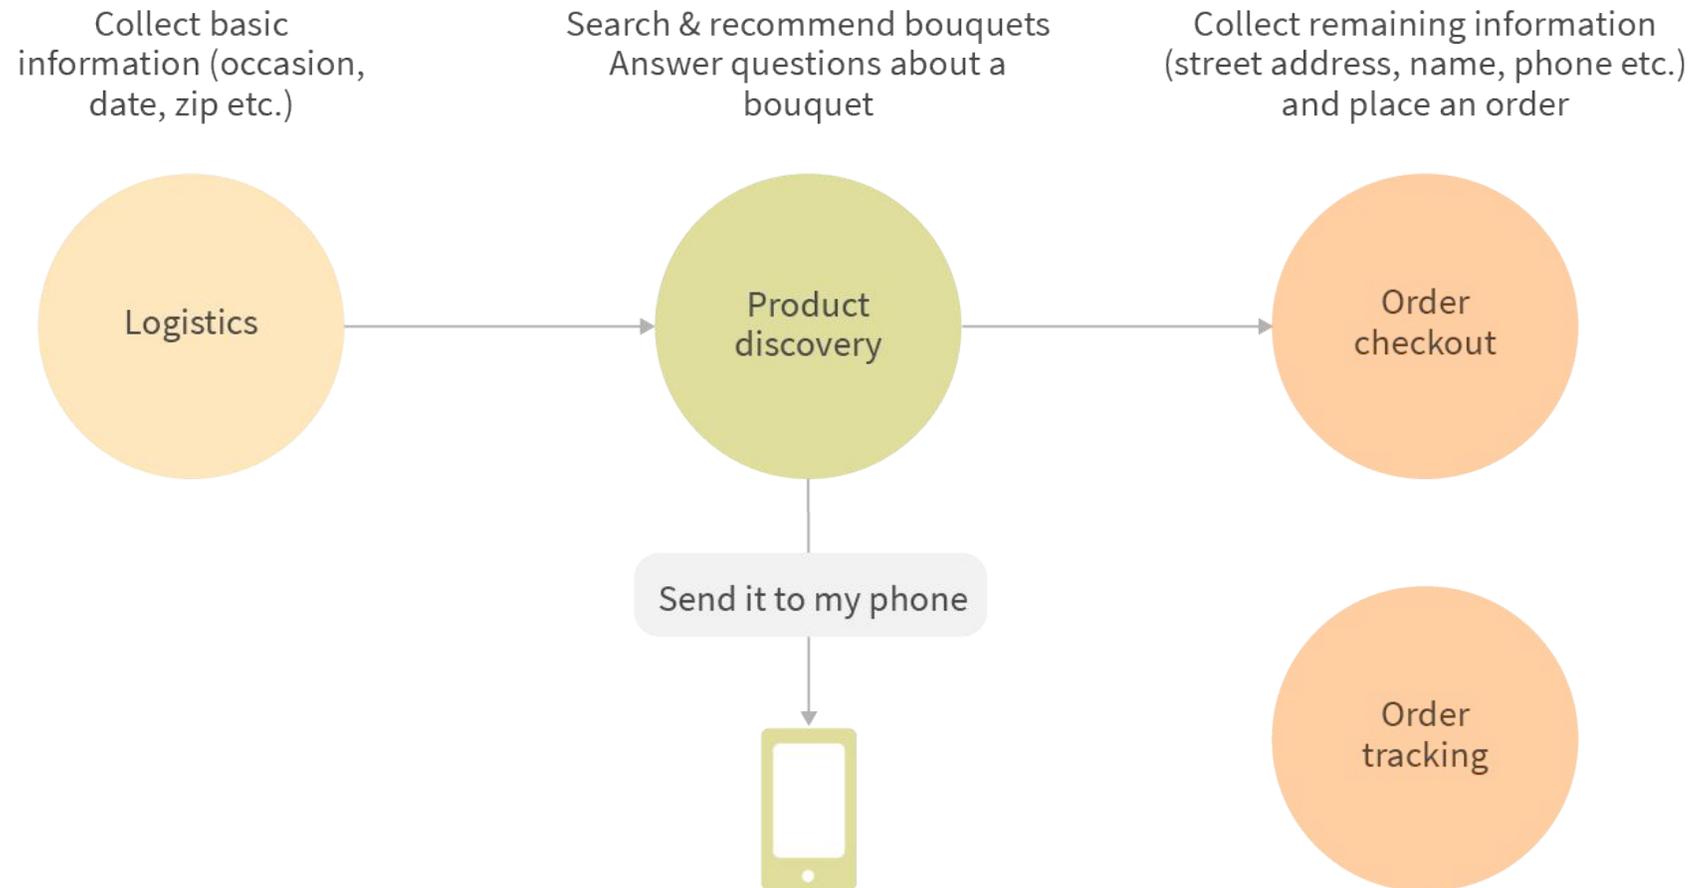
Today's talk

- At Grid Dynamics, we've been working on conversational systems since 2017
- Particular focus on voice commerce: selling products & services via Alexa/Google Home/etc.
- So, what does it take to write a conversational application?

Let's take a close look at the "Flower Genie" and "Data Genie": our "Petshops of Conversational AI"



So, you want to order some flowers for you mom?



Plausible actual conversation

I want to send my friend a bouquet for her birthday

Sure! I need to know basic delivery information to show you available bouquets. What is the recipient's ZIP code?

I don't know... it is in San Ramon, California

Ok. San Ramon, California. When is your friend's birthday?

Next Friday

Ok. Friday, June 15. Let's see what we can suggest...

Wait, her birthday is on June 16th.

Sure. Changed to Saturday, June 16. How about red roses and lilies in a glass vase? It costs \$59



She doesn't like red roses...

Ok, this one doesn't have red roses.



Do you have something with roses, tulips, and without a vase?

How about this beautiful collection of cream roses, white and violet tulips? It costs \$60.



Looks good... Do you have something bigger?

Sure. Here is a bouquet with white and cream roses, and mixed tulips. The price is \$69.



How many flowers in this bouquet?

11 roses and 15 tulips

And what about the previous? And remind me the price...

9 roses and 11 tulips, the price is \$60

Ok, I want to order it.

But it's not all roses. Typical gotchas and hiccups

Context switching



What is the recipient's address?

Wait, I forgot to ask, do the flowers arrive closed? Or are they already blooming?



The flowers will arrive in bud form and will need at least a couple of days to open up. Would you like to continue your order?

I changed my mind, let's go back to other bouquets.

Correct the previous answer



What date would you like your flowers to arrive?

Next Friday



Ok. Friday, October 15th. And what is the occasion for this gift?

Oh, no, I need them on October 16th



Sure, changed to October 16th. What is the occasion for this gift?

User cannot provide some information



What is the recipient's ZIP code?

I don't know... she lives in Chicago, Illinois.



Ok, Chicago, Illinois. And what is your phone number?

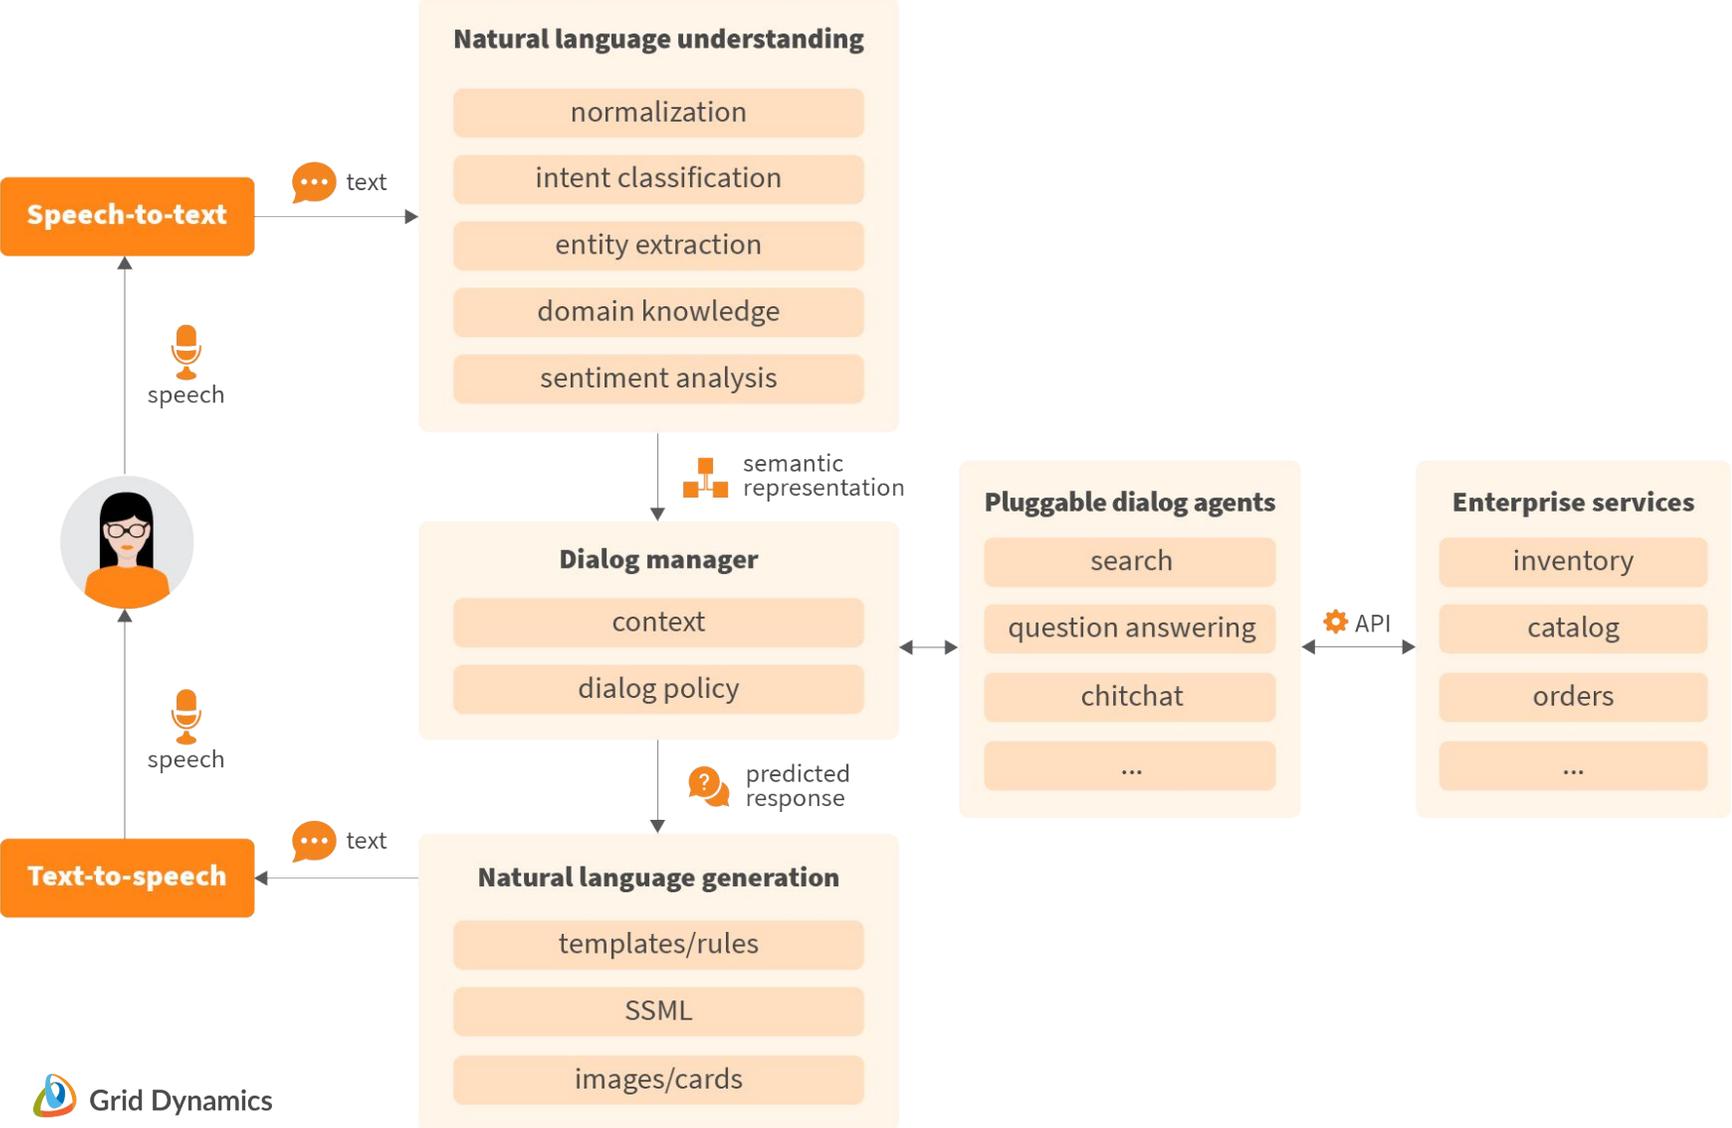
Why do you need to know that?



We need to know the phone number to call if we have additional questions or issues with your delivery.

Live demo!

General system's architecture



How about an AI recommendation system for a camera?

Help me choose the right camera. How hard is it?

- Hint: a lot harder than flowers!
- 1,000s of products; wide range on categories, from cheap point-n-shoot to professional
- Technical product with wildly different features
- Some customers know exactly what they want (hobbists, pros); others - nothing at all

What else do we want from the dialog?

- Know the difference between advisory vs. order-taking. Adopt the dialog accordingly
- Determine customer's knowledge level. Adopt the dialog accordingly
- Graceful switch between “leading the witness” & provide useful information.
- Deal with “I dunno”

Live demo!

Scenario 1: “I want a camera for hiking”, “I want a camera for travels”

- Step 1: quickly determine that I don't know much about cameras
- Step 2: gently find out what I need camera for, then lead me through selection
- Step 3: ask and answer relevant (to the purpose) questions
- Step 4: close the deal

Scenario 2: Expert that knows exactly what he wants

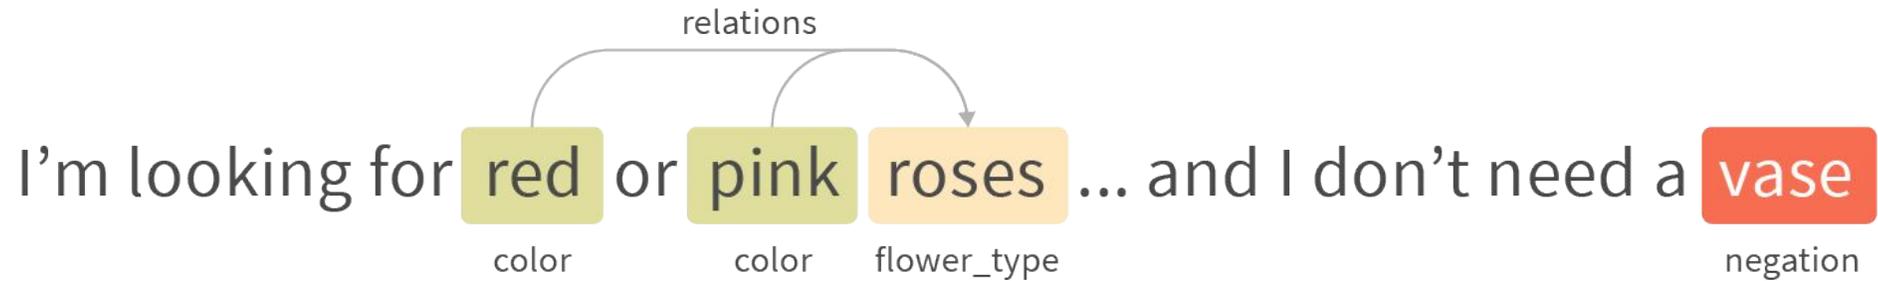
- Step 1: quickly determine that I already know a lot about what I want.
- Step 2: provide direct, complete, factual information. Follow, not lead.
- Step 3: close the deal

Conclusion

- Availability of high-quality AI models is spreading =>
- The price & complexity of conversational applications is rapidly coming down =>
- Best practices in design, testing and certification of conversational apps are emerging =>

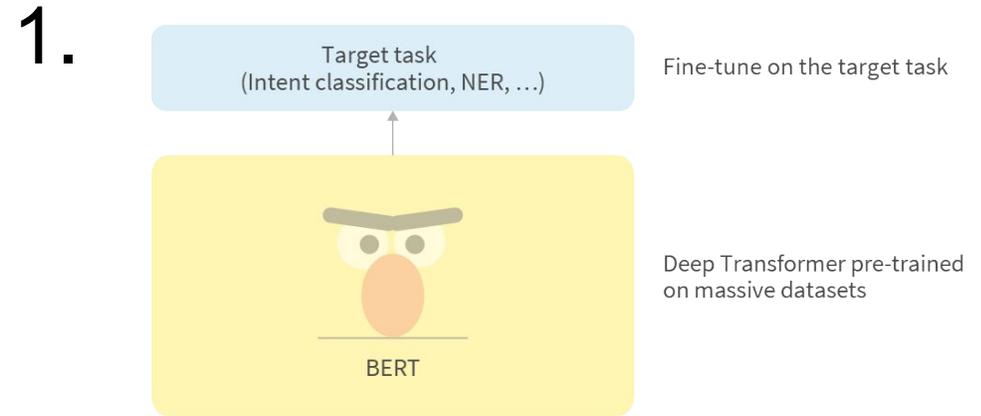
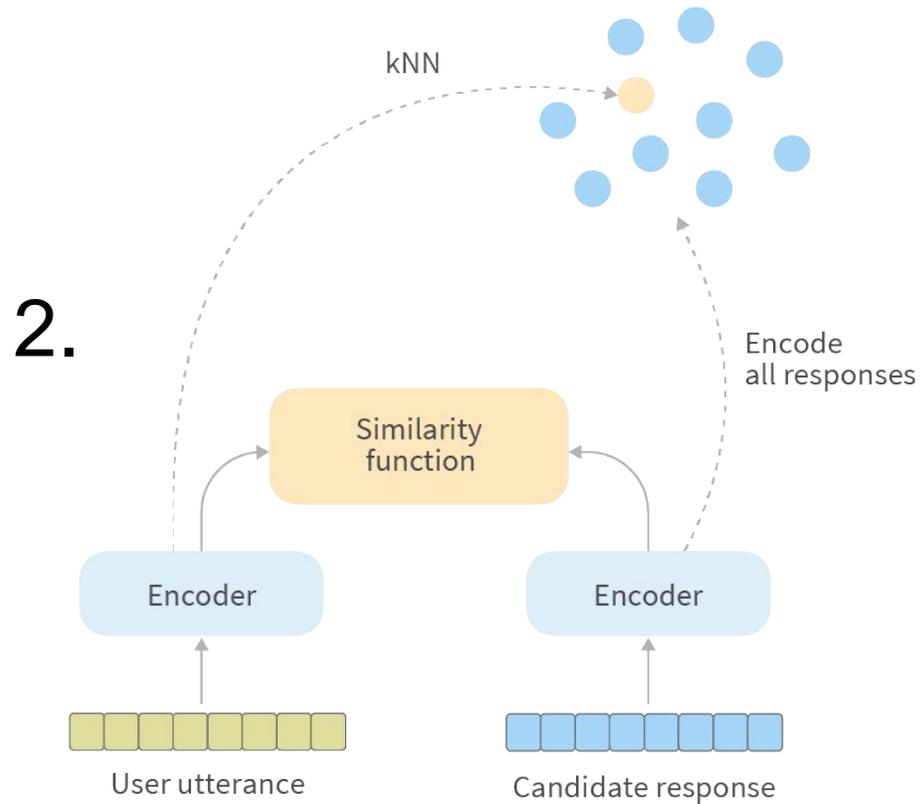
- This is already practical to create a wide range of useful eCommerce voice / visual apps today

Understanding customer queries using Deep Learning

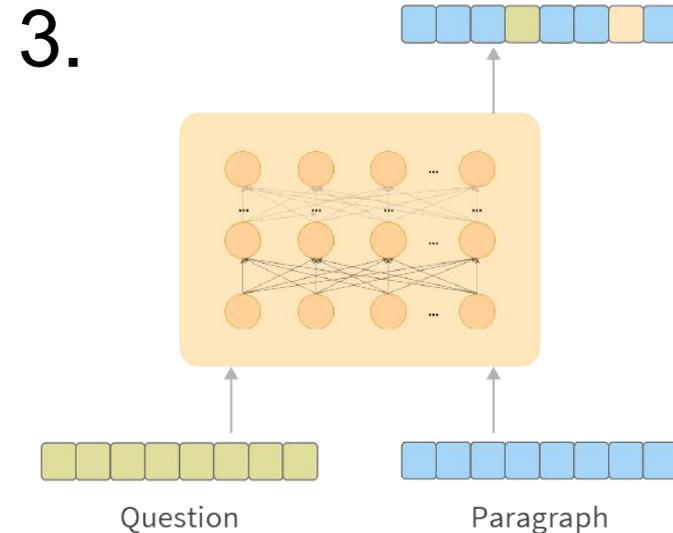


Intent: SearchFlowers

ML bag of tricks

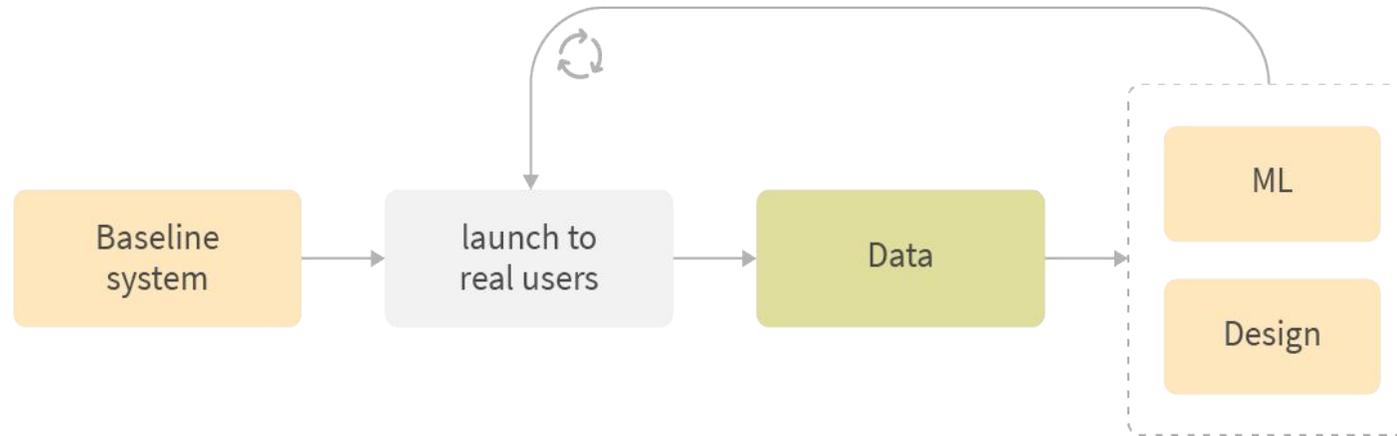


Transfer learning



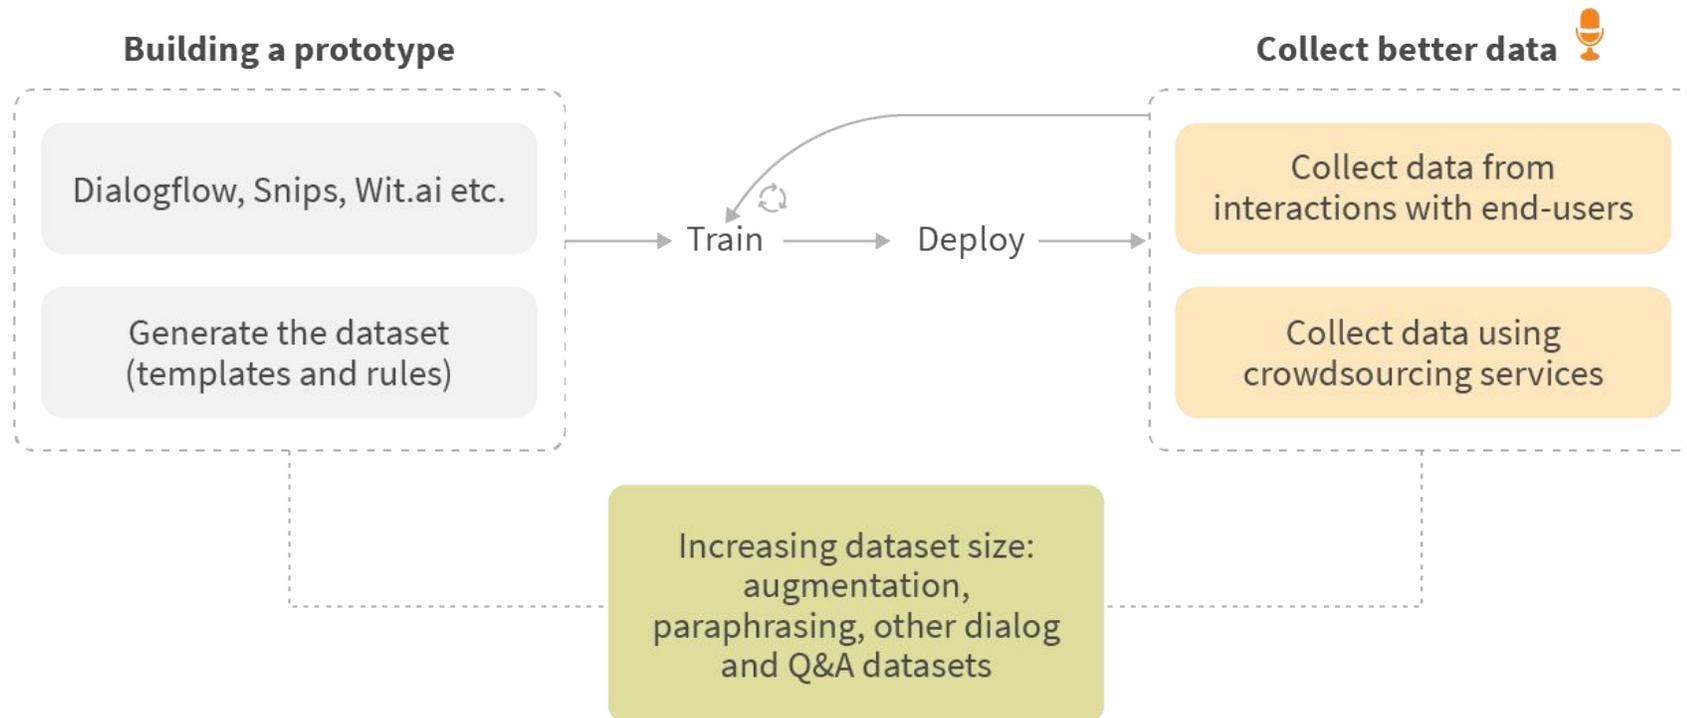
Reading comprehension

Design cycle: yep, it's a cycle. Particularly in AI systems.



If you don't understand your customer's behavior, your customers will not understand your AI

Where do you get the training data?



Finally, testing and certification. There is a lot to it.



To learn more

- Detailed blog post that spills all the beans:

<https://blog.griddynamics.com/how-we-built-a-conversational-ai-for-ordering-flowers/>

- More about conversational AI development services:

<https://www.griddynamics.com/technologies/ai/voice-application-development-services>

- Write to us!

vlivschitz@griddynamics.com

info@griddynamics.com