


Language Conditioned Imitation Learning Over Unstructured Data

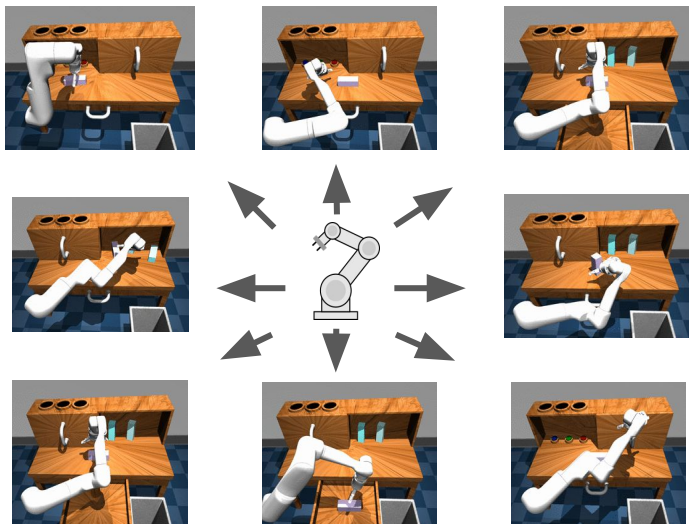
"open the door...now
press the green
button"



Corey Lynch, Pierre Sermanet


Setting

- One robot, **many tasks**.
- **Natural language** conditioning.
- Simple **imitation learning**.

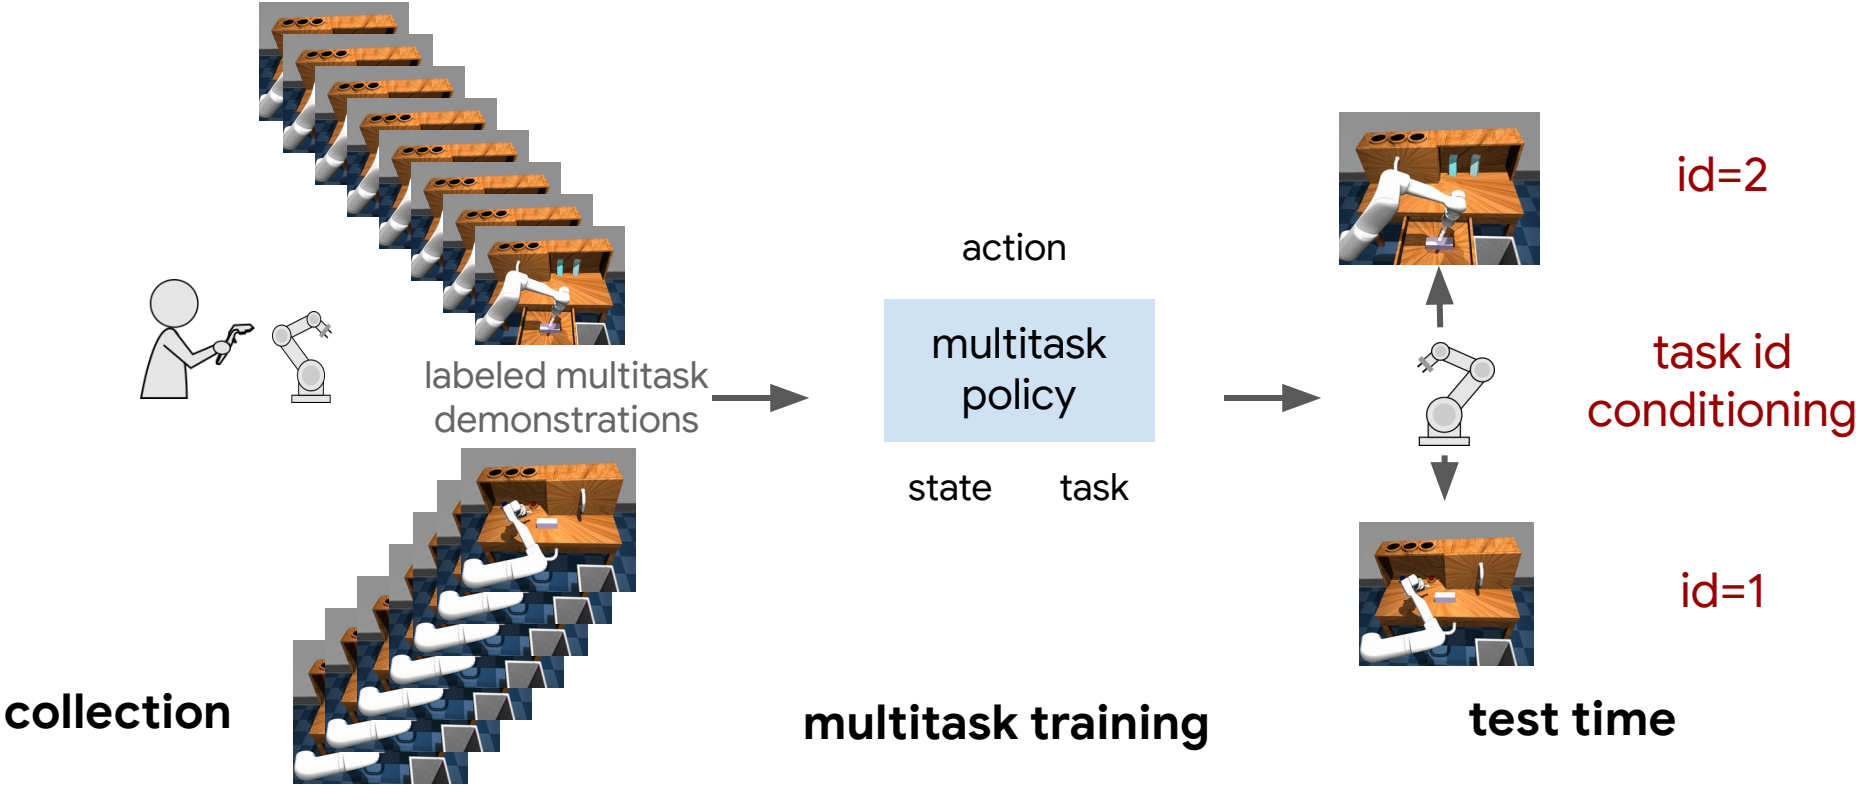


"open the door..."
"...now press the
green button"

**task specification:
natural language**



Prior work: multitask imitation learning




(Rahmatizadeh et al 2018)

How can we learn **many** useful behaviors **self-supervised** from **unstructured** data?


How can we condition these behaviors on free-form **natural language**?

How can we follow new **synonyms at test time**?

Human teleoperated “play” data for training



Learning Latent Plans
from Play
(Lynch et al 2019)



no upfront
task
definition

no resets

humans
cover state
space for the
robot

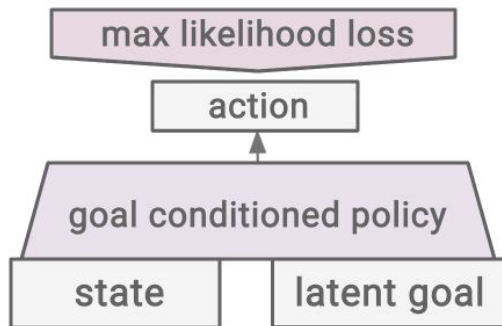
(2.5x)

Pair free-form **play** with free-form crowd-sourced **language**

Teleoperated play



Multicontext imitation learning: follow **image** or **language** goals






less than 1% of
windows
require
language labels



Follow **natural language** instructions at test time over **long horizon**

more videos at [language-play.github.io](https://github.com/language-play)

“open the door
... now pick up
the block
... now push the
red button”



now: **pull the drawer handle all the way**
next: put the block into the drawer

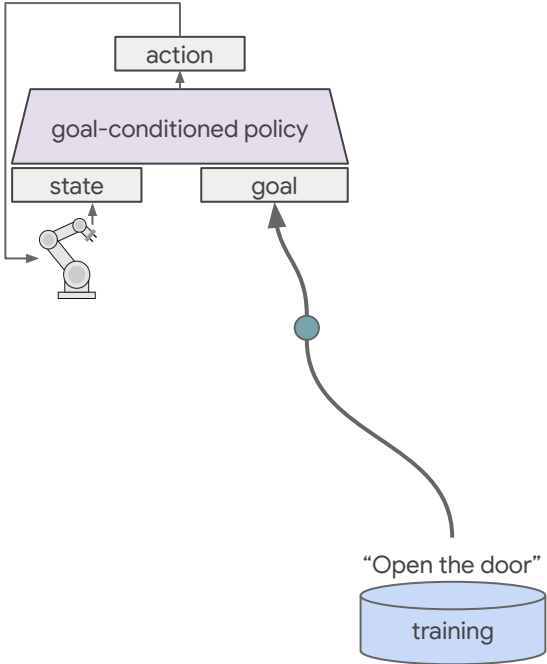
← **current instruction**

~~How can we learn many useful behaviors from~~
~~unstructured~~ data? Play + relabeled imitation.

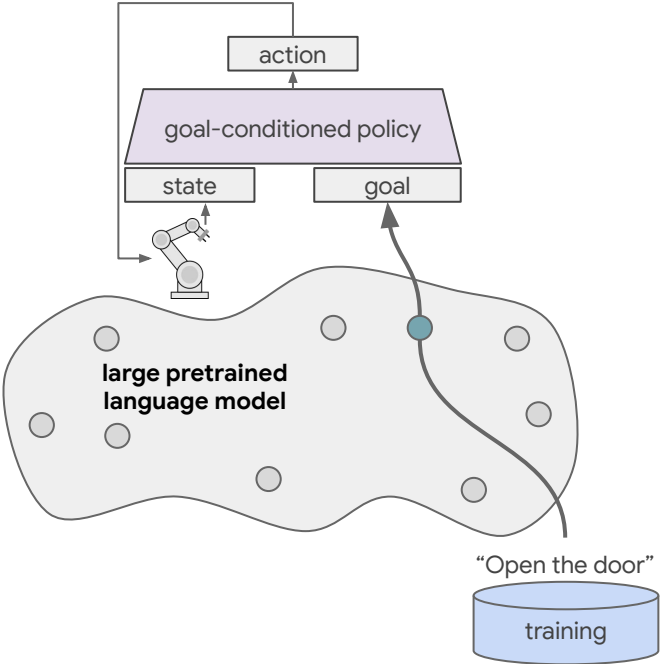
~~How can we condition these behaviors on free form~~
~~natural language~~? Pair random play windows with
language.

How can we be robust to **synonyms at test time**?

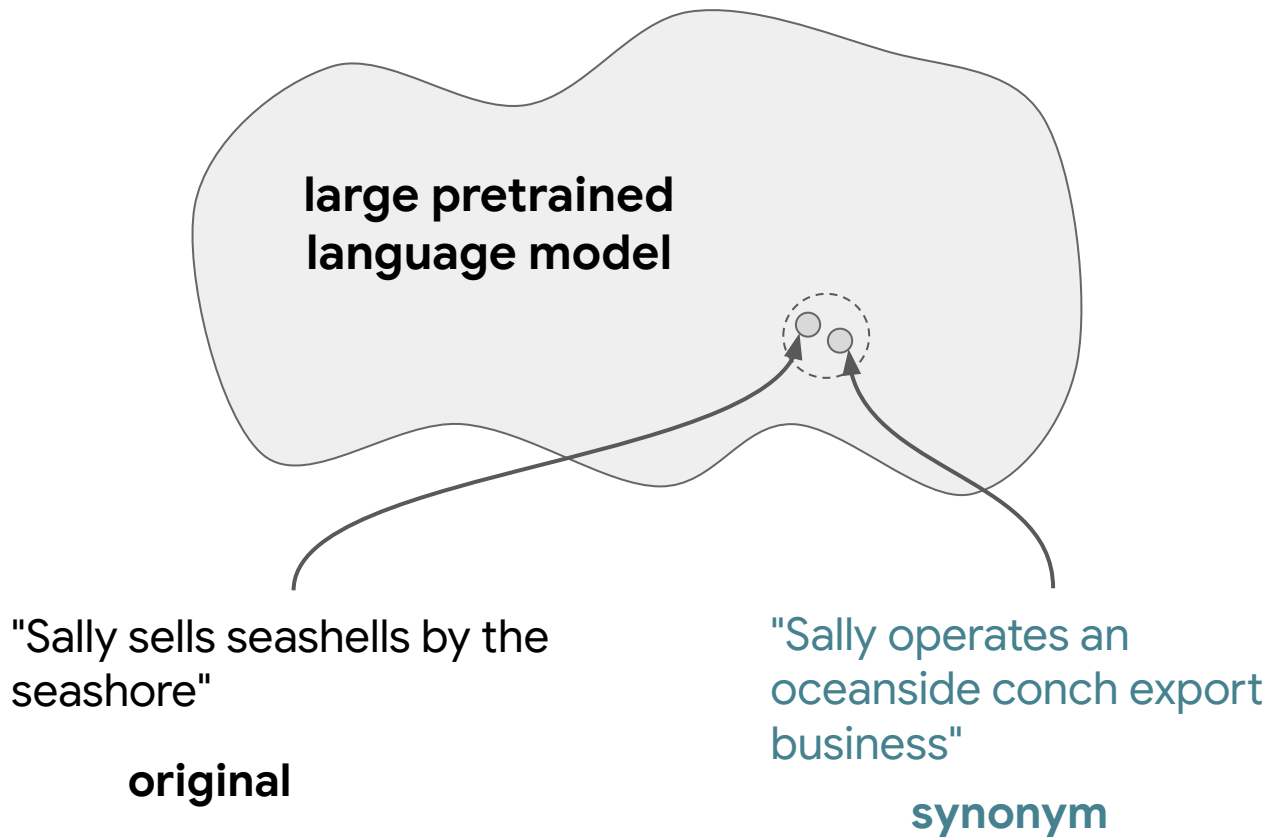
Large pretrained language models lead to zero-shot "synonym" instruction following



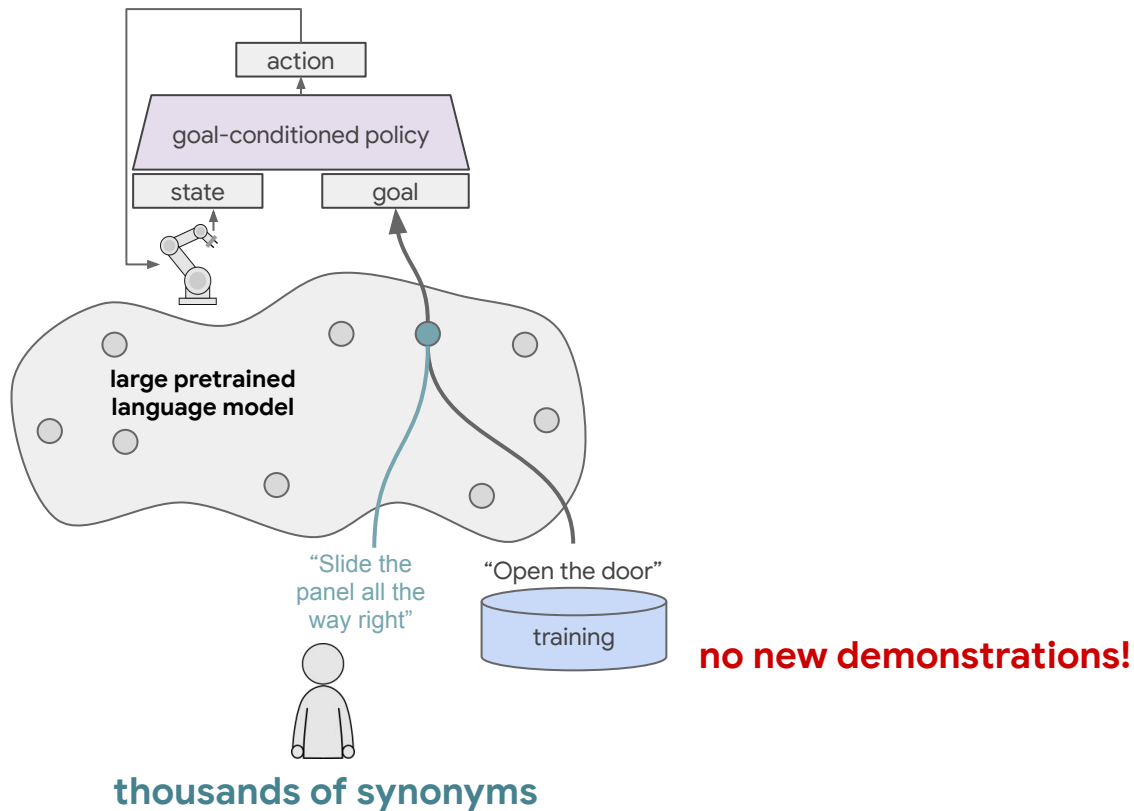
Large pretrained language models lead to zero-shot "synonym" instruction following



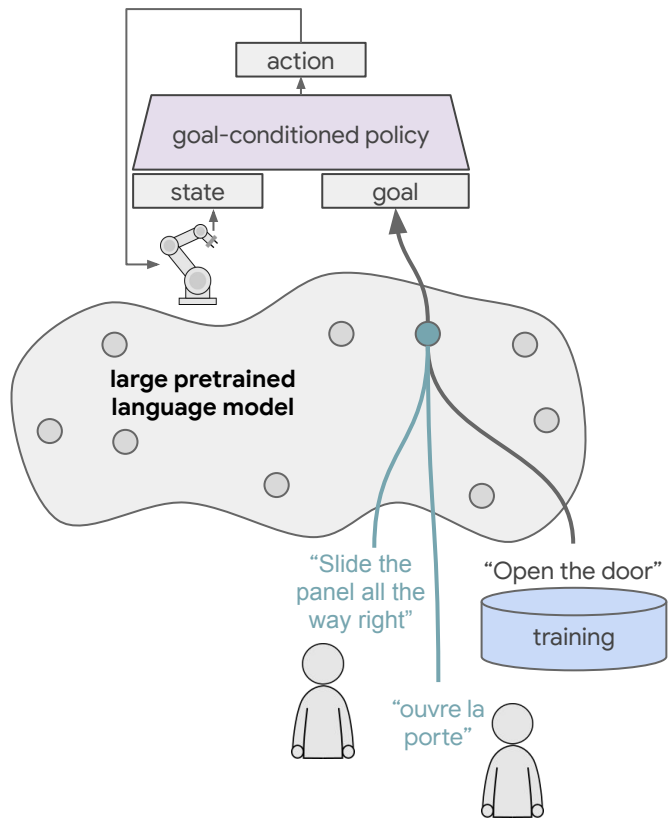
Large pretrained language models lead to zero-shot "synonym" instruction following



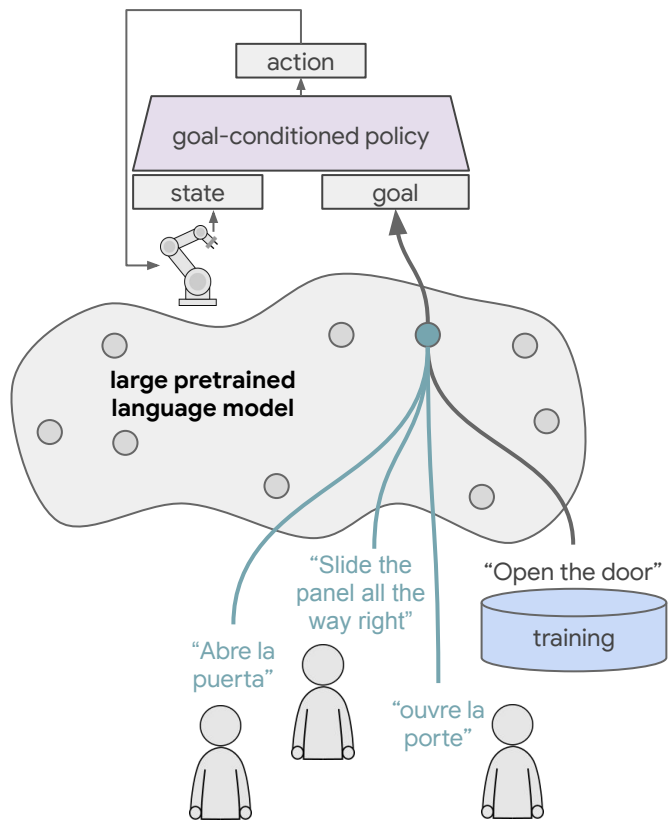
Large pretrained language models lead to zero-shot "synonym" instruction following



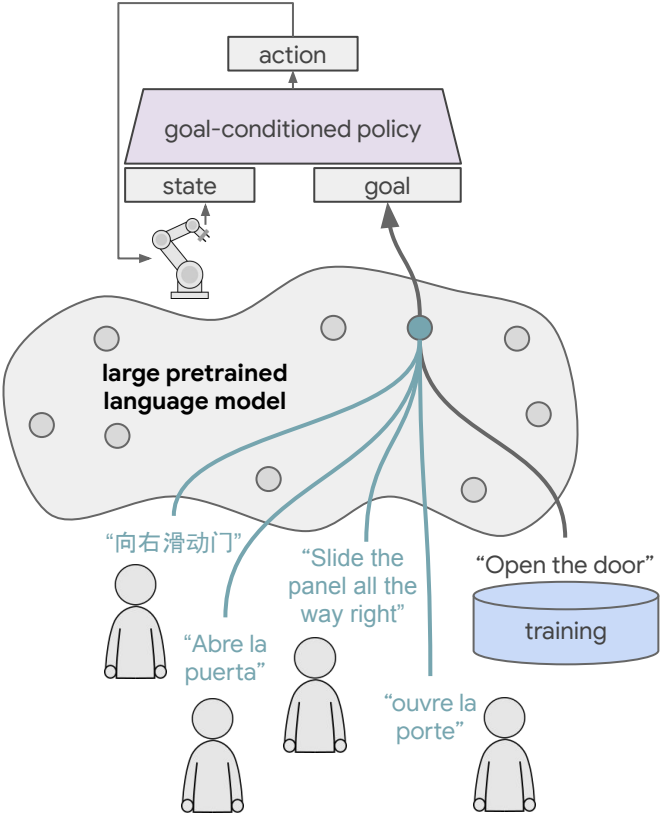
Large pretrained language models lead to zero-shot "synonym" instruction following



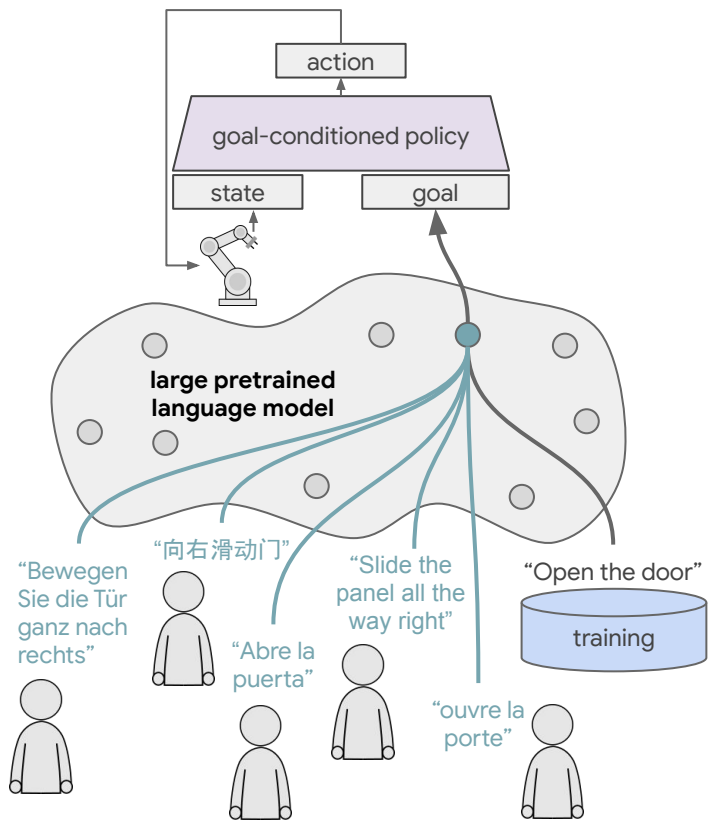
Large pretrained language models lead to zero-shot "synonym" instruction following



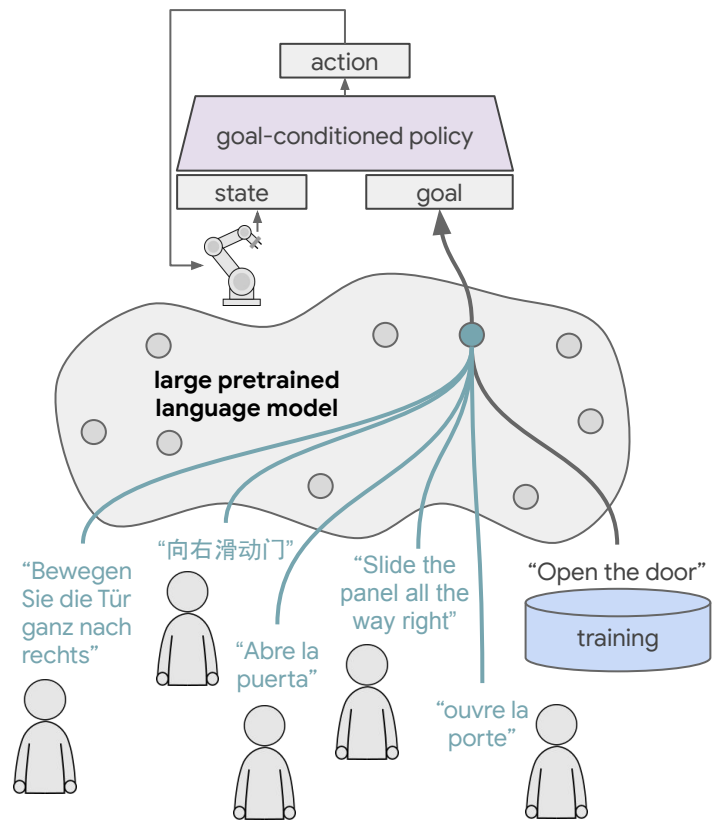
Large pretrained language models lead to zero-shot "synonym" instruction following



Large pretrained language models lead to zero-shot "synonym" instruction following



Large pretrained language models lead to zero-shot "synonym" instruction following



Outcome: follow thousands of new instructions (in 16 languages) without collecting new demonstrations

Thank
you!

Videos + paper:


language-play.github.io



Corey Lynch



Pierre Sermanet



now: **do not do anything**
next:



Robotics at Google
<http://g.co/robotics>