MACQUARIE University

Chatbots

A HANDS-ON SESSION

The Plan

MACQUARIE University

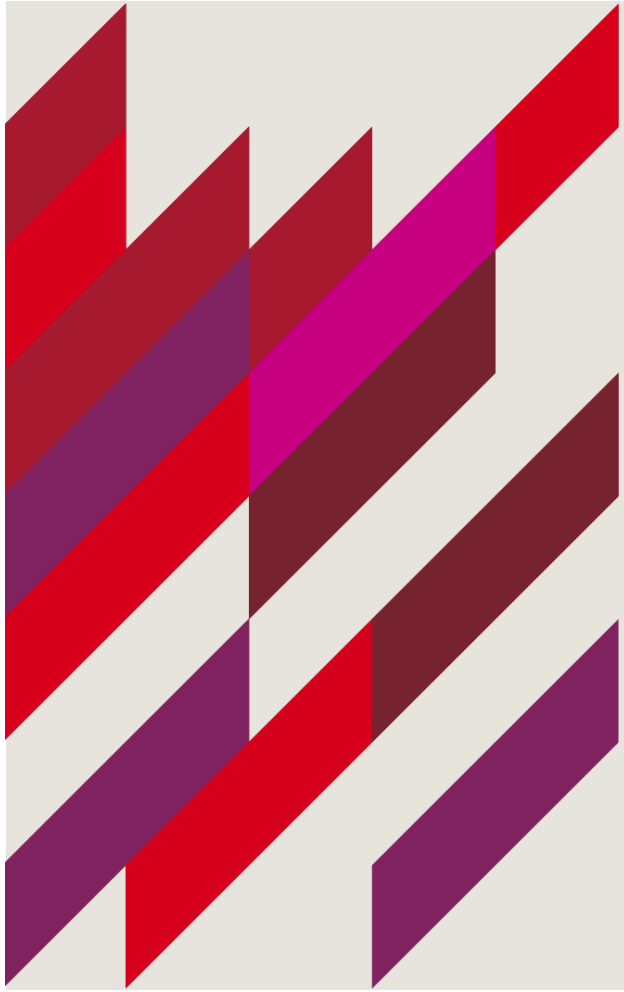

1. A Whirlwind History
2. Exploring Some Chatbots
3. Pandorabots
4. Building Your Own

The Idea of Artificially Intelligent Life …

FICTION

Greek myth:

Pygmalion and Galatea

Gothic novel:

Frankenstein and Frankenstein's monster

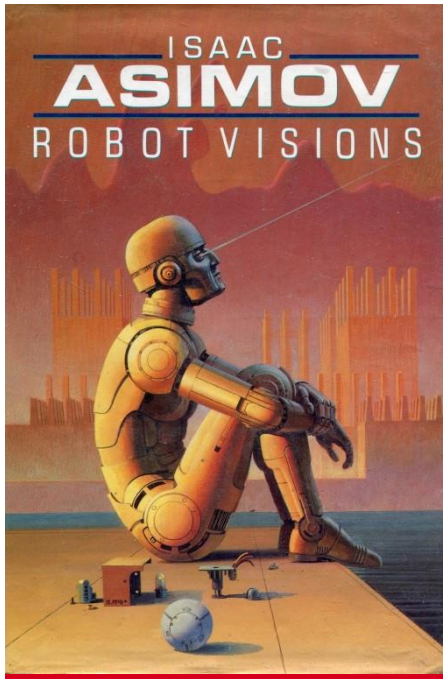

Scifi:

Asimov's robots

Human-like Automata

EARLY ATTEMPTS

- Ismail al-Jazari's drink-serving waitress
  - 12th century Turkey
  - a waitress that could serve water, tea or drinks.
    - "The drink was stored in a tank with a reservoir from where the drink drips into a bucket and, after seven minutes, into a cup, after which the waitress appears out of an automatic door serving the drink."
- Wolfgang von Kempelen's Mechanical Turk
  - 18th century Austria
  - an "automaton" that played chess
  - revealed to be a hoax: actually a human chess master hiding inside to operate the machine

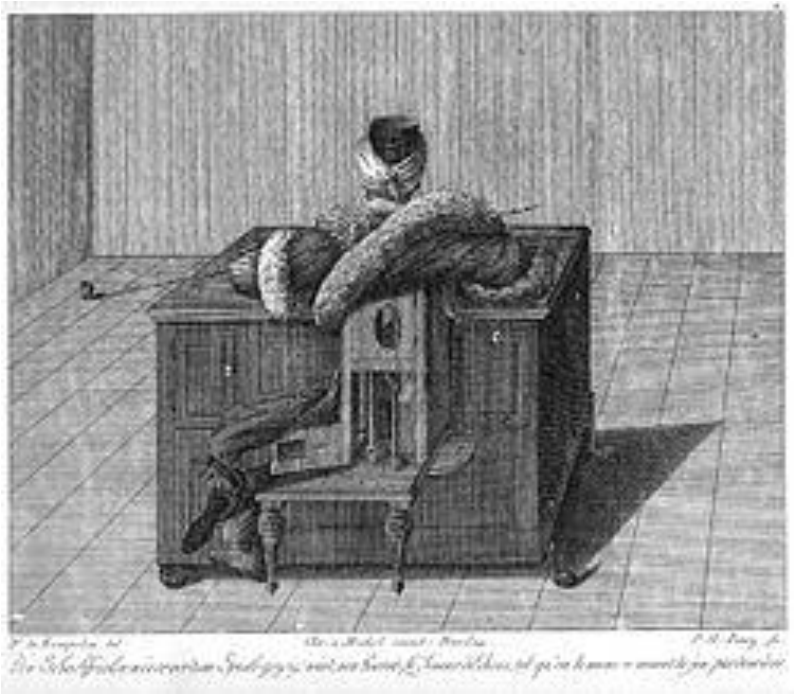

The Beginnings of AI

ALAN TURING

- Alan Turing: a key figure in the history of computing
- WWII: codebreaker at Bletchley Park
  - helped crack the German Enigma machine codes
- Turing Machine
  - the theoretical device that defines the limits of computation
  - e.g. Halting Problem
- Turing Test
  - proposed in 1950
  - asked the question about what it meant for a computational device to be intelligent

The Turing Test

THE IMITATION GAME

- How best to ask whether machines can think?
- Frame questions as: "Can machines do what we (as thinking entities) can do?“
- Inspired by a party game called The Imitation Game
  - a man and a woman go into separate rooms
  - guests try to tell them apart by writing a series of questions and reading the typewritten answers sent back
  - both the man and the woman aim to convince the guests that they are the other
- The computational version has a computer and a human instead of a man and a woman

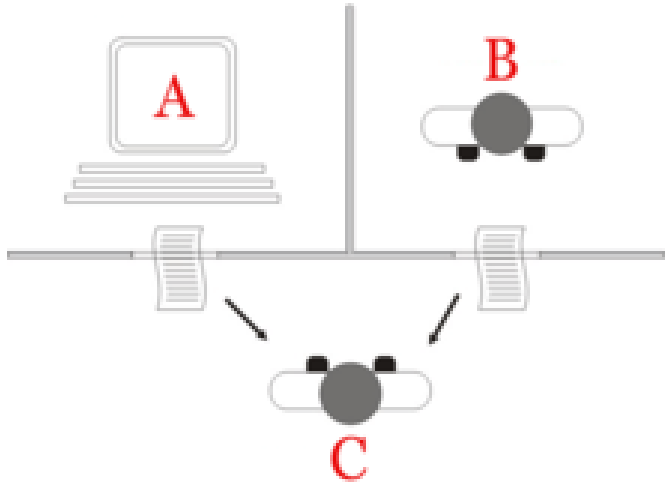

Philosophy of Mind: nature of intelligence in general still debated

One Strand: The Loebner Prize

MACQUARIE University

A CHATBOT CONTEST

- Inspired by the Turing Test
  - concrete programs with a panel of human judges
- Started in 1991, underwritten by Hugh Loebner
- Two one-time prizes
  - $100K for the first program to fool judges with textual, visual and auditory input
  - $25K for the first program to fool judges just based on text
- Annual competitions
  - typical prize ~$3K
  - best result in 2008 fooled 5/12 judges

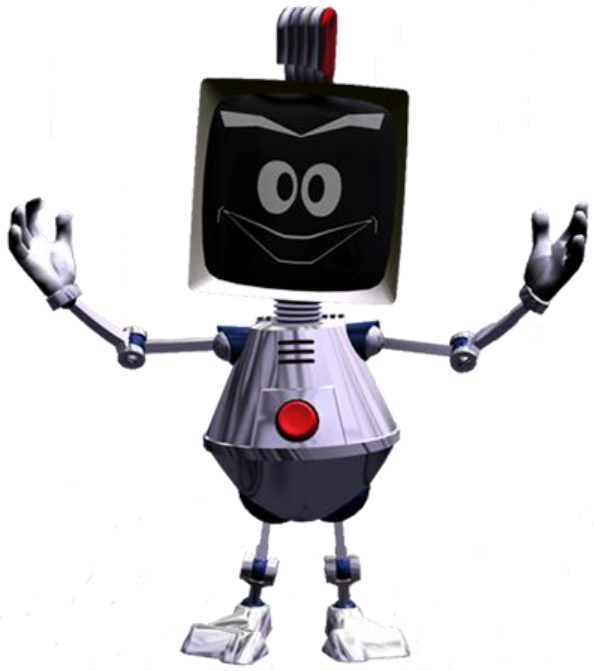

Academic view: seen as a bit of a gimmick

Another Strand: Natural Language Processing

A FIELD OF RESEARCH

- Started just after WWII, with idea of Machine Translation
- Lots of other problems now tackled:
  - question answering
  - summarisation
  - sentiment analysis
  - dialogue systems
  - …
- Lots of industry interest
  - Google, Microsoft, IBM, …
- Dialogue systems: research and deployment
  - current examples: Apple Siri (2010), Android Google Now (2013), Windows Cortana (2014), Amazon Echo home assistant (2014)
  - mostly for specific purposes, rather than for chatting

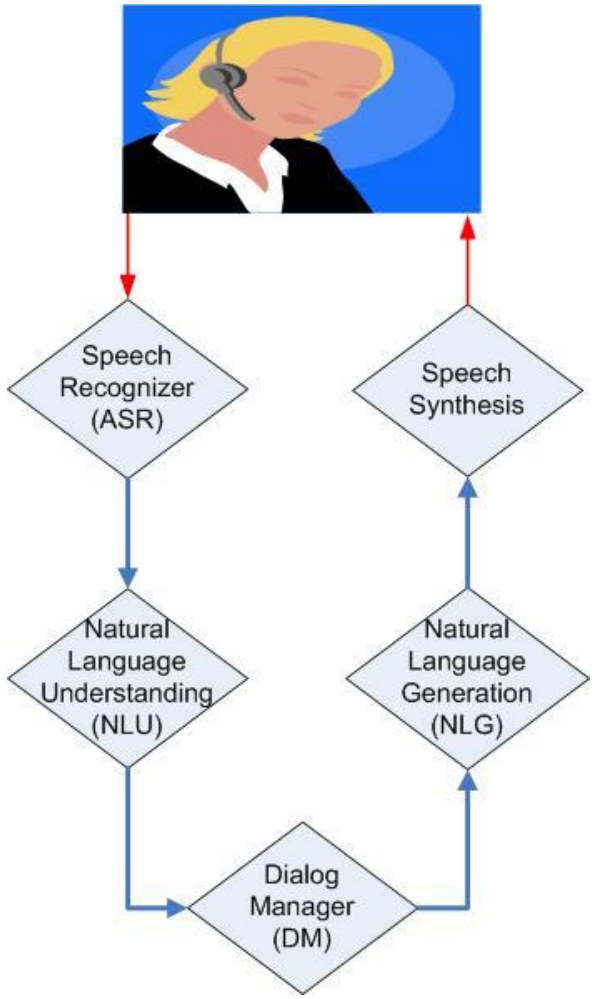

Two Approaches

BOTH FOR CHATBOTS AND NLP

- Rule-based
  - can be quite complex
  - narrow coverage
  - fairly brittle
  - can produce decent quality output
- Statistical / machine learning-based
  - broader coverage
  - robust
  - output quality may be variable

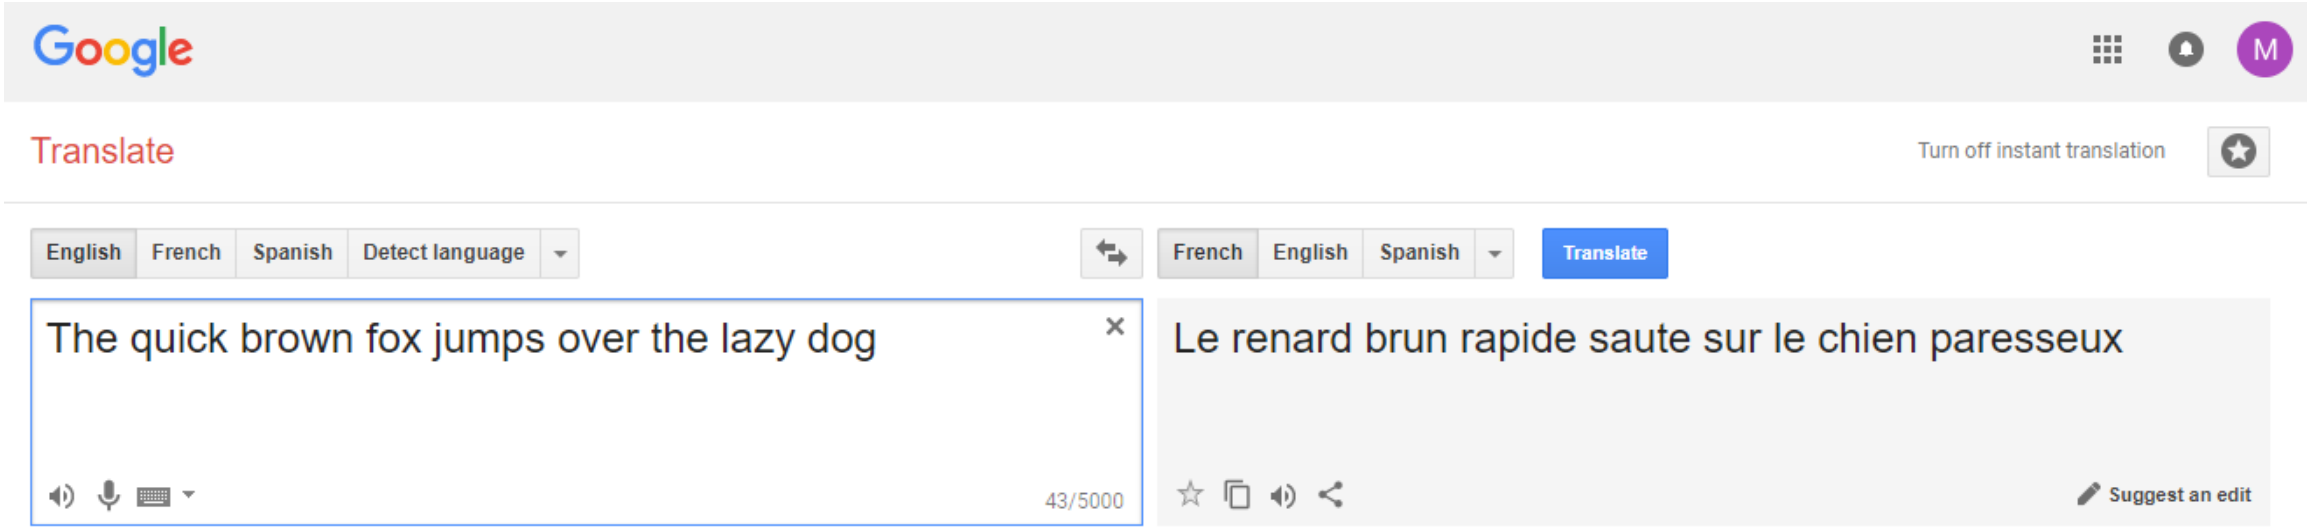

Two Approaches

- Deep learning
  - the latest machine learning approach
  - dialogue systems treated like Machine Translation
  - still early days

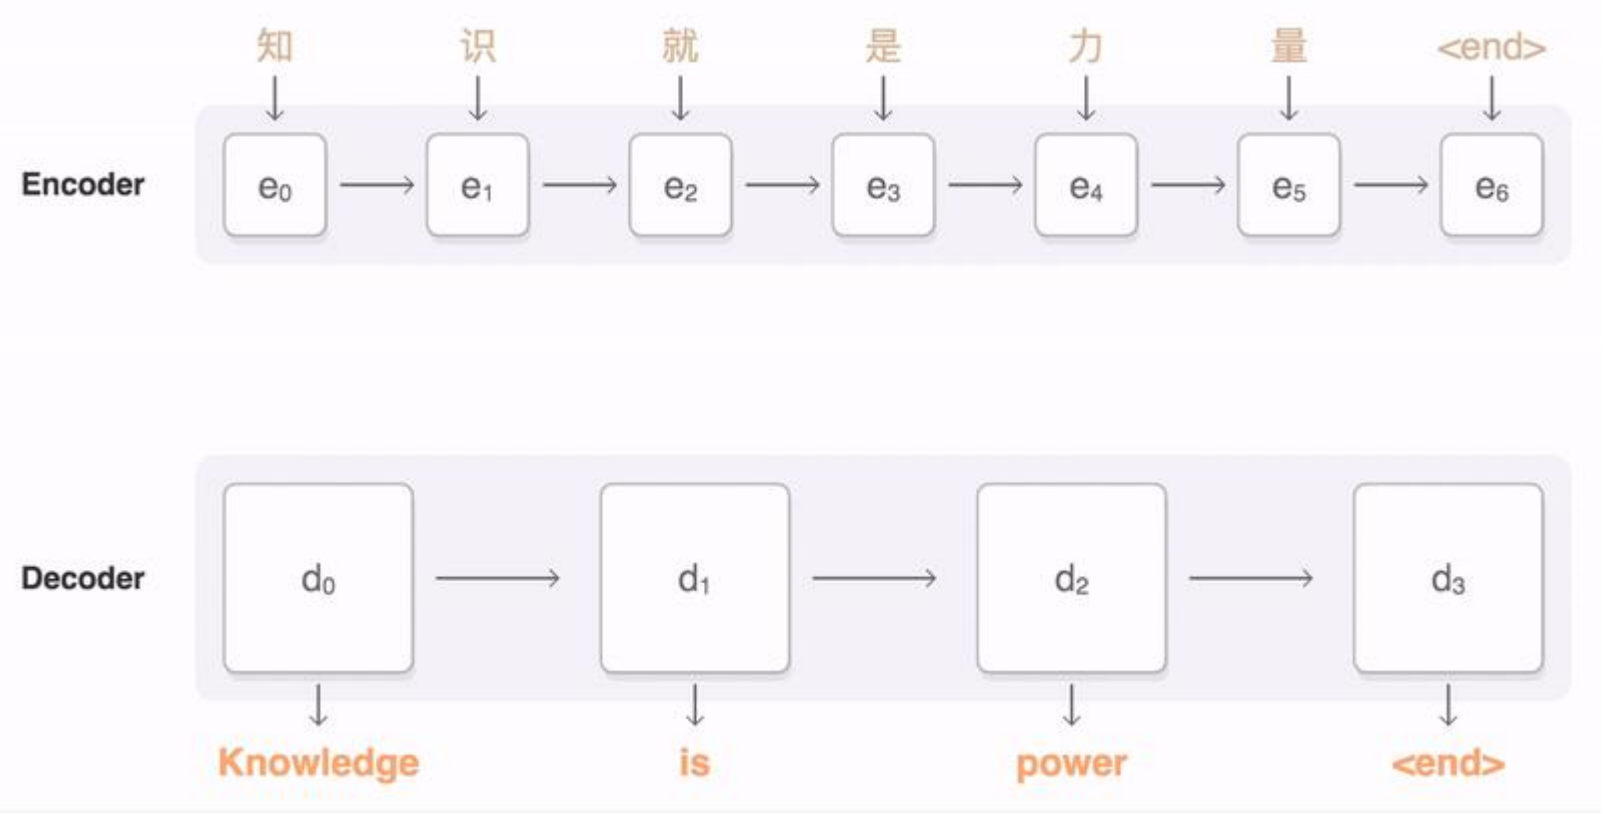

Now

CHATBOTS AND NLP

- Strands seem to have merged to some extent, leading to a boom
  - researchers and industry interested in conversational agents
  - large companies working on frameworks (Microsoft, IBM, …)
  - lots of startups too (motion.ai, hu:toma, …)

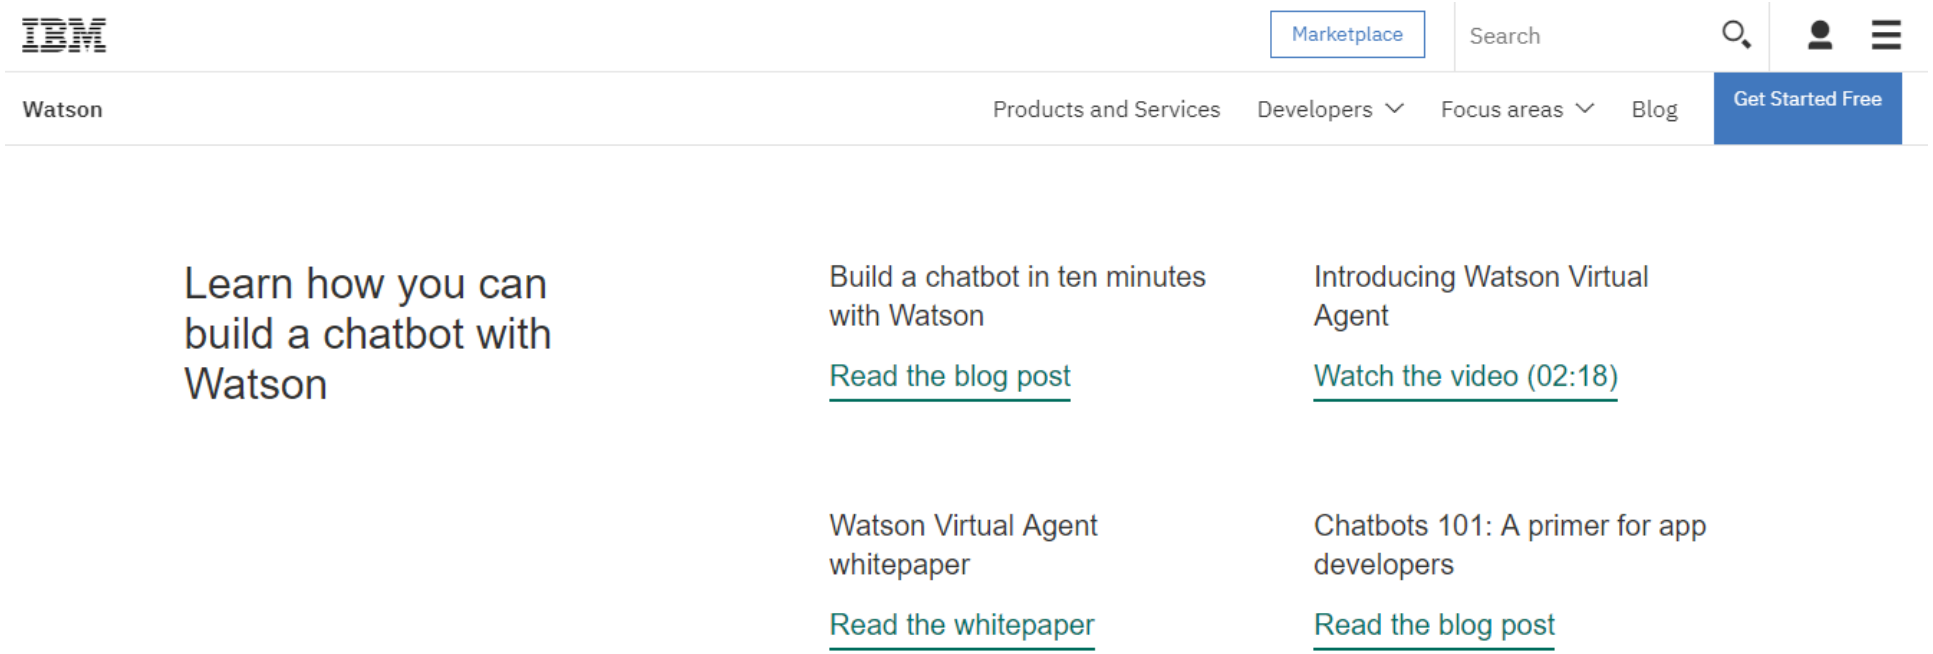

Some Examples

CHATBOTS AND NLP

- Rule-based (“retrieval-based”)
  - Mitsuku
  - Rose
  - Rosie
  - Chomskybot
  - Cleverbot [video]
  - Elbot
  - Woebot
  - Insomnobot
  - lots listed at chatbots.org
- Machine learning-based (“generative”)
  - Neuralconvo
  - IBM Watson Conversation
  - Microsoft’s Tay [Wikipedia]
  - Nadia

Pandorabots

WRITING RULES: AIML MARKUP

- An environment for building rule-based bots
  - Evolved from an early Loebner prize-winning bot (ALICE)
  - Now Rosie is available
- Rules are described using a special mark-up language, AIML
  - XML-compliant
- Consists of three key parts
  - `<category>`: the tag that marks a "unit of knowledge" in a bot's knowledge base
  - `<pattern>`: used to contain a simple pattern that matches what a user may say or type to a bot
  - `<template>`: contains the response to a user input

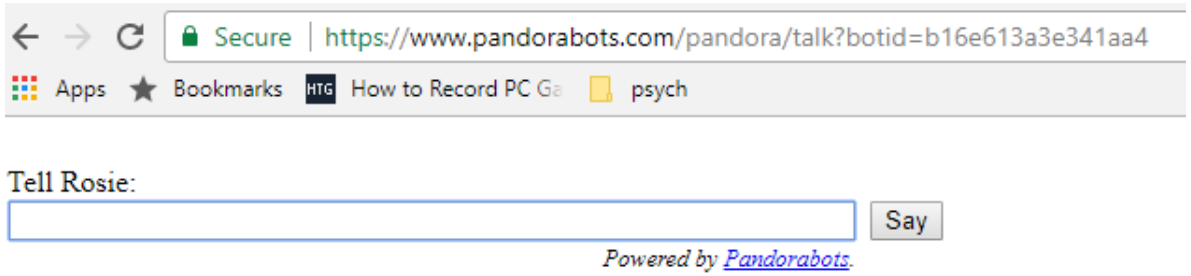


```
<category>
    <pattern>WHAT ARE YOU</pattern>
    <template>
        <think><set name="topic">Me</set></think>
        I am the latest result in artificial intelligence,
        which can reproduce the capabilities of the human brain
        with greater speed and accuracy.
    </template>
</category>
```

Pandorabots

THINGS TO DO

START

- Register with pandorabots
  - You'll need a Facebook, Twitter, Google+ or Yahoo account
- Go through Quick Start
- Try adding the ability to tell a knock knock joke (at end of article)

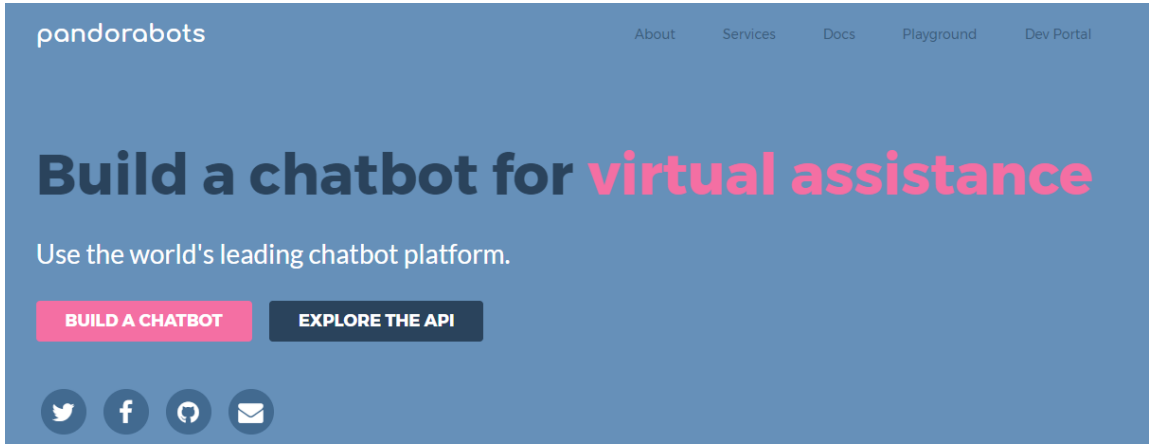

INSTALL ROSIE IN PANDORABOTS

- Go to github repository
- Click on 

; choose "Download ZIP"
- Unzip the zipfile
- Create a new bot and go to the `Files` tab
- Upload the files in the zipfile's `lib/` directory to pandorabots

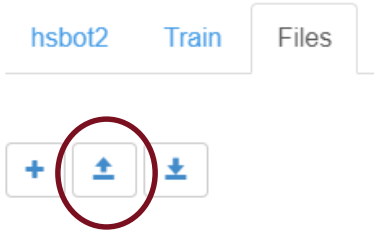

Pandorabots

THINGS TO DO

MODIFY ROSIE

- Try out the bot in a conversational environment
  - click `Publish Bot`
  - go to the Playground Clubhouse
- Change the bot's name and other properties
  - in file **`SYSTEM`** `> properties`
  - search for property `name`
- Glance through the tutorial for details on the various tags (very long, so don't read all of it)
  - useful ones: `<random>`, `<star/>`, `<set>`, `<srai>`, `<that>`, `<topic>`
- Add your own chat topics
- **Challenge**: try adding the knock knock joke to this chatbot
- Publish your final bot so you can try out each other's.