

What are Chatbots? Beginner's Guide To Chatbots

Kashif Manzoor

July , 2017



Topics

- Chatbot 101
 - Chatbot Rudiments
 - Types of Chatbot
- Conversational Interface -Chatbots
- Natural Language Processing (NLP), Artificial Intelligence (AI), Machine Learning (ML)
- Rise of the Chatbot economy

Chatbot Use On The Rise

Gartner, the usage of chatbots (called virtual customer assistants by Gartner) will triple through 2019 as enterprises seek to increase customer satisfaction and reduce operating costs.

Source: "Seven Decision Points for Success with Virtual Customer Assistants," Gartner, July 26, 2016.



A man wearing a straw hat, glasses, a maroon sweater, and a yellow backpack is standing on a train platform, looking at his smartphone. The platform has tracks and a red and white railing in the background.

Can I book a room tonight?

Hi Kashif, we have a room
available tonight, room rate
\$379.



EDWARDIAN
HOTELS
LONDON

Chatbot Rudiments

Talkbot, Chatterbot, Bot, Chatterbox, Artificial Conversational Entity?

“Chatbot” refers to a broad range of technologies that allow consumers to use a conversational interface to accomplish tasks.

“A chatbot is a service, powered by rules and sometimes artificial intelligence, that we interact with through a chat interface.”

A chatbot (also known as a **talkbot, chatterbot, Bot, chatterbox, Artificial Conversational Entity**) is a computer program which conducts a conversation via auditory or textual methods. – Wikipedia

A Brief History of Chatbots

1950 • The Turing Test

Alan Turing theorized that a truly intelligent machine would be indistinguishable from a human during a text-only conversation. **Turing's ideas helped lay the groundwork for the chatbot revolution.**



1966 • ELIZA

The world's first chatbot, ELIZA could mimic the responses of a psychotherapist and -- for short spurts -- carry on convincingly human conversations.

1972 • PARRY

In the 1970s, ELIZA met its first (non-human) patient: PARRY, a chatbot that could imitate a person with paranoid schizophrenia.

1988 • Jabberwacky

Developed in the 1980s and released online in 1997, the Jabberwacky chatbot was designed to "simulate natural human chat in an interesting, entertaining and humorous manner."

1995 • ALICE

The "Artificial Linguistic Internet Computer Entity" chatbot served as the inspiration for the 2013 film *Her*, in which a human falls in love with a computer program.

2001 • SmarterChild

The first chatbot to achieve widespread adoption, SmarterChild joined the buddy lists of millions of AIM and MSN Messenger users in the early 2000s.


The SmarterChild approach -- building a chatbot for popular messaging platforms -- is the blueprint that most modern chatbot builders are now following.

Source: drift.com

The term "ChatterBot" was originally coined by [Michael Mauldin](#) (creator of the first [Verbot](#), Julia) in 1994 to describe these conversational programs – Wikipedia



Today, chatbots are part of [virtual assistants](#) such as [Google Assistant](#), and are accessed via many organizations' apps, websites, and on [instant messaging](#) platforms such as [Facebook Messenger](#) – Wikipedia





Examples of Chat Bots

Weather bot. Get the weather whenever you ask.



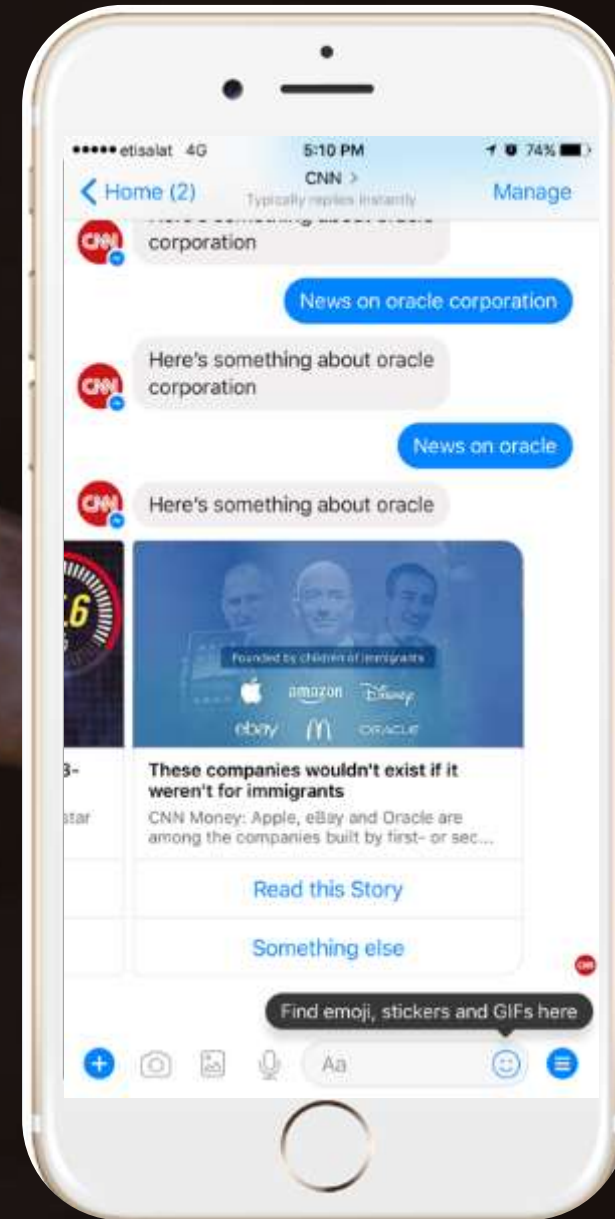
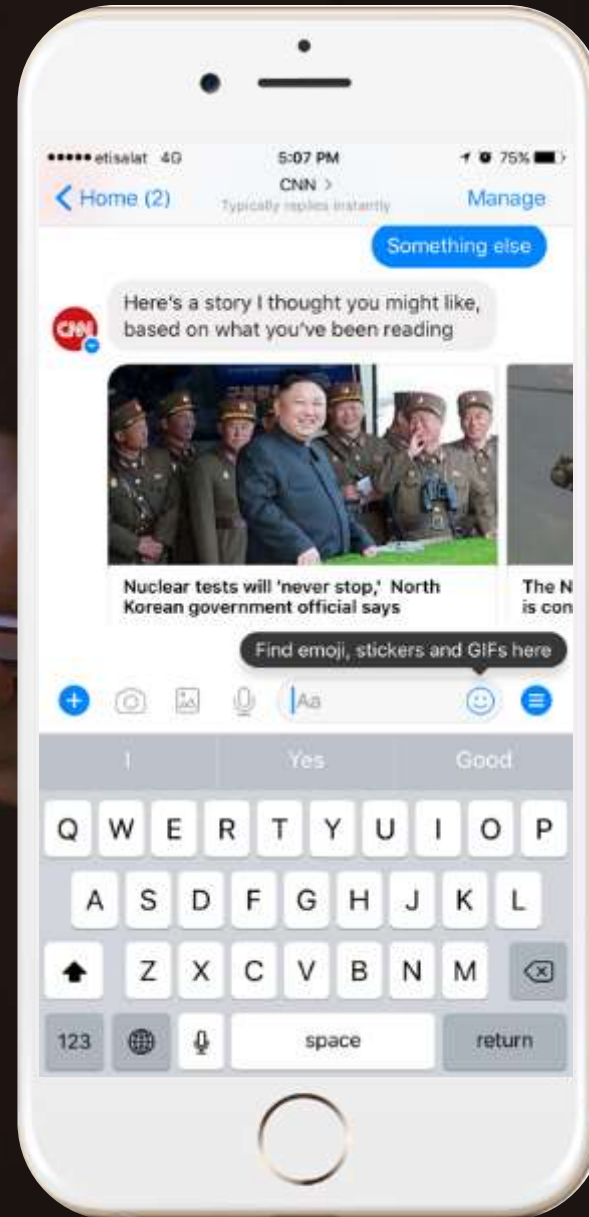
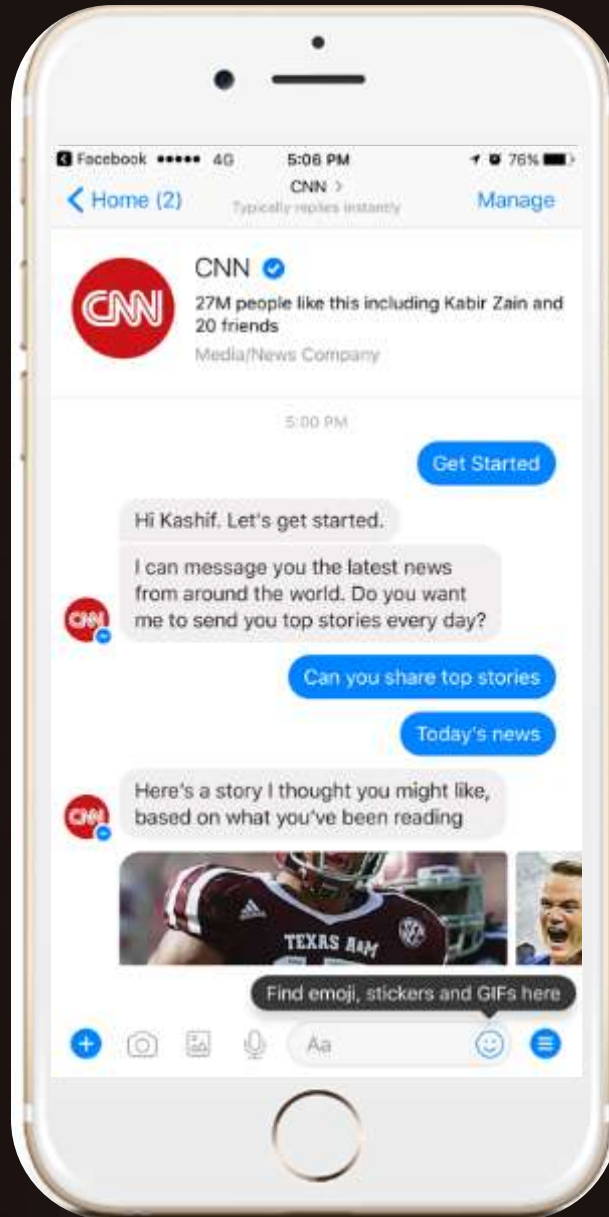
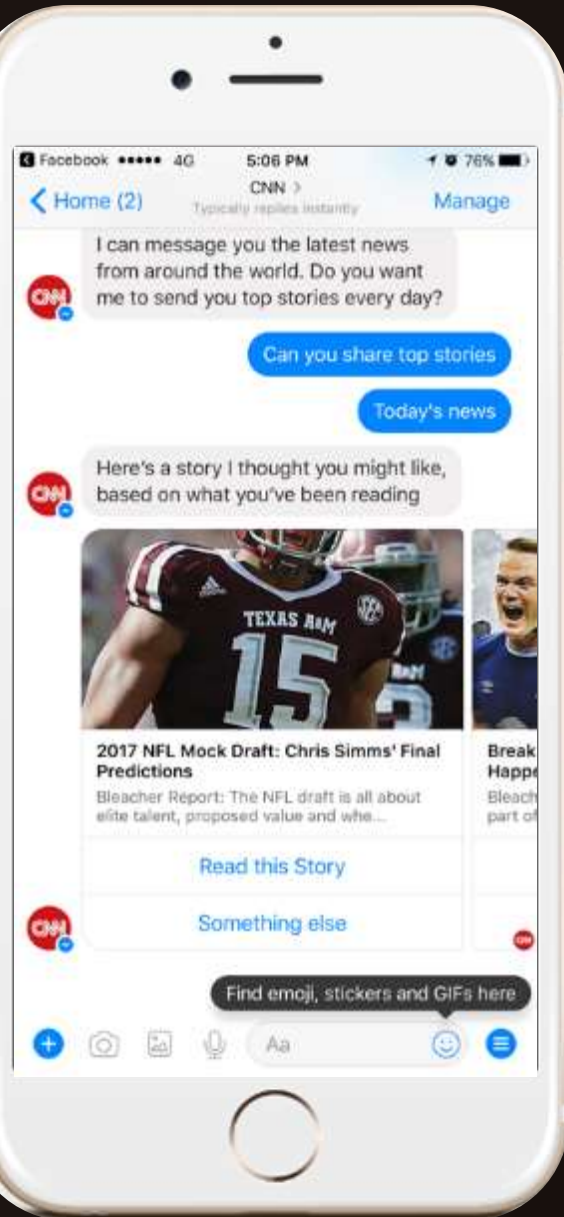
Grocery bot. Help me pick out and order groceries for the week.



News bot. Ask it to tell you when ever something interesting happens.




CNN Chatbot on Facebook





Personal finance bot. It helps me manage money better.

Life advice bot. I'll tell it my problems and it helps me think of solutions.



A person with dark hair is seen from the back, writing on a white piece of paper. The word 'Meeting' is written in a large, cursive script. Below it, the time '14:30' is written in a similar style. The person is holding a black pen in their right hand and has their left hand on the paper. The background is a metallic surface with a bright light source creating a lens flare effect.

Meeting
14:30

Scheduling bot. Get me a meeting with someone on the Support team at

A bot that's your friend. In China there is a bot called Xiaoice, that over 20 million people talk to...



Bots are created with a purpose.

A retail store will likely want to create a bot that helps you purchase something, where someone like Services org might create a bot that can answer customer support questions.



How do they engage with customers?

You start to interact with a chatbot by sending it a message

Reactive: Customers can select chatbots from a menu or a button on a web page or in a mobile app. Other approaches include chatbots as “listeners” (for example, twitter, Facebook, or SMS) that react to inquiries as customers enter these channels.




- **Proactive:** Intelligent chatbots can operate in real time and predict customer intentions—offering specific help when they detect that a customer may need assistance. For example, a client has visited several mortgage pages and pauses on a specific page whereby the chatbot can proactively engage the client.

How Chatbots Work


Chatbot that functions based on rules:

- This bot is very very limited. It can only respond to very specific commands. If you say the wrong thing, it doesn't know what you mean.
- This bot is only as smart as it is programmed to be.



Chatbot that functions using machine learning:

- This bot has an artificial brain AKA artificial intelligence. You don't have to be ridiculously specific when you are talking to it. It understands language, not just commands.
- This bot continuously gets smarter as it learns from conversations it has with people.



How an Intelligent Chatbot Works...

1. Captures data in real time
2. Uses internal data
3. Combines data to predict customer intentions
4. The chatbot develops
5. Engages customers
6. Understands what is said
7. Formulates a response
8. Determines follow-up actions

What are the stages of Chatbot Maturity?



Informational

Understands natural language to provide answers to questions posed



Personalized

Presents user-specific responses by connecting to enterprise systems and clarifies user intent using menus or simple questions.



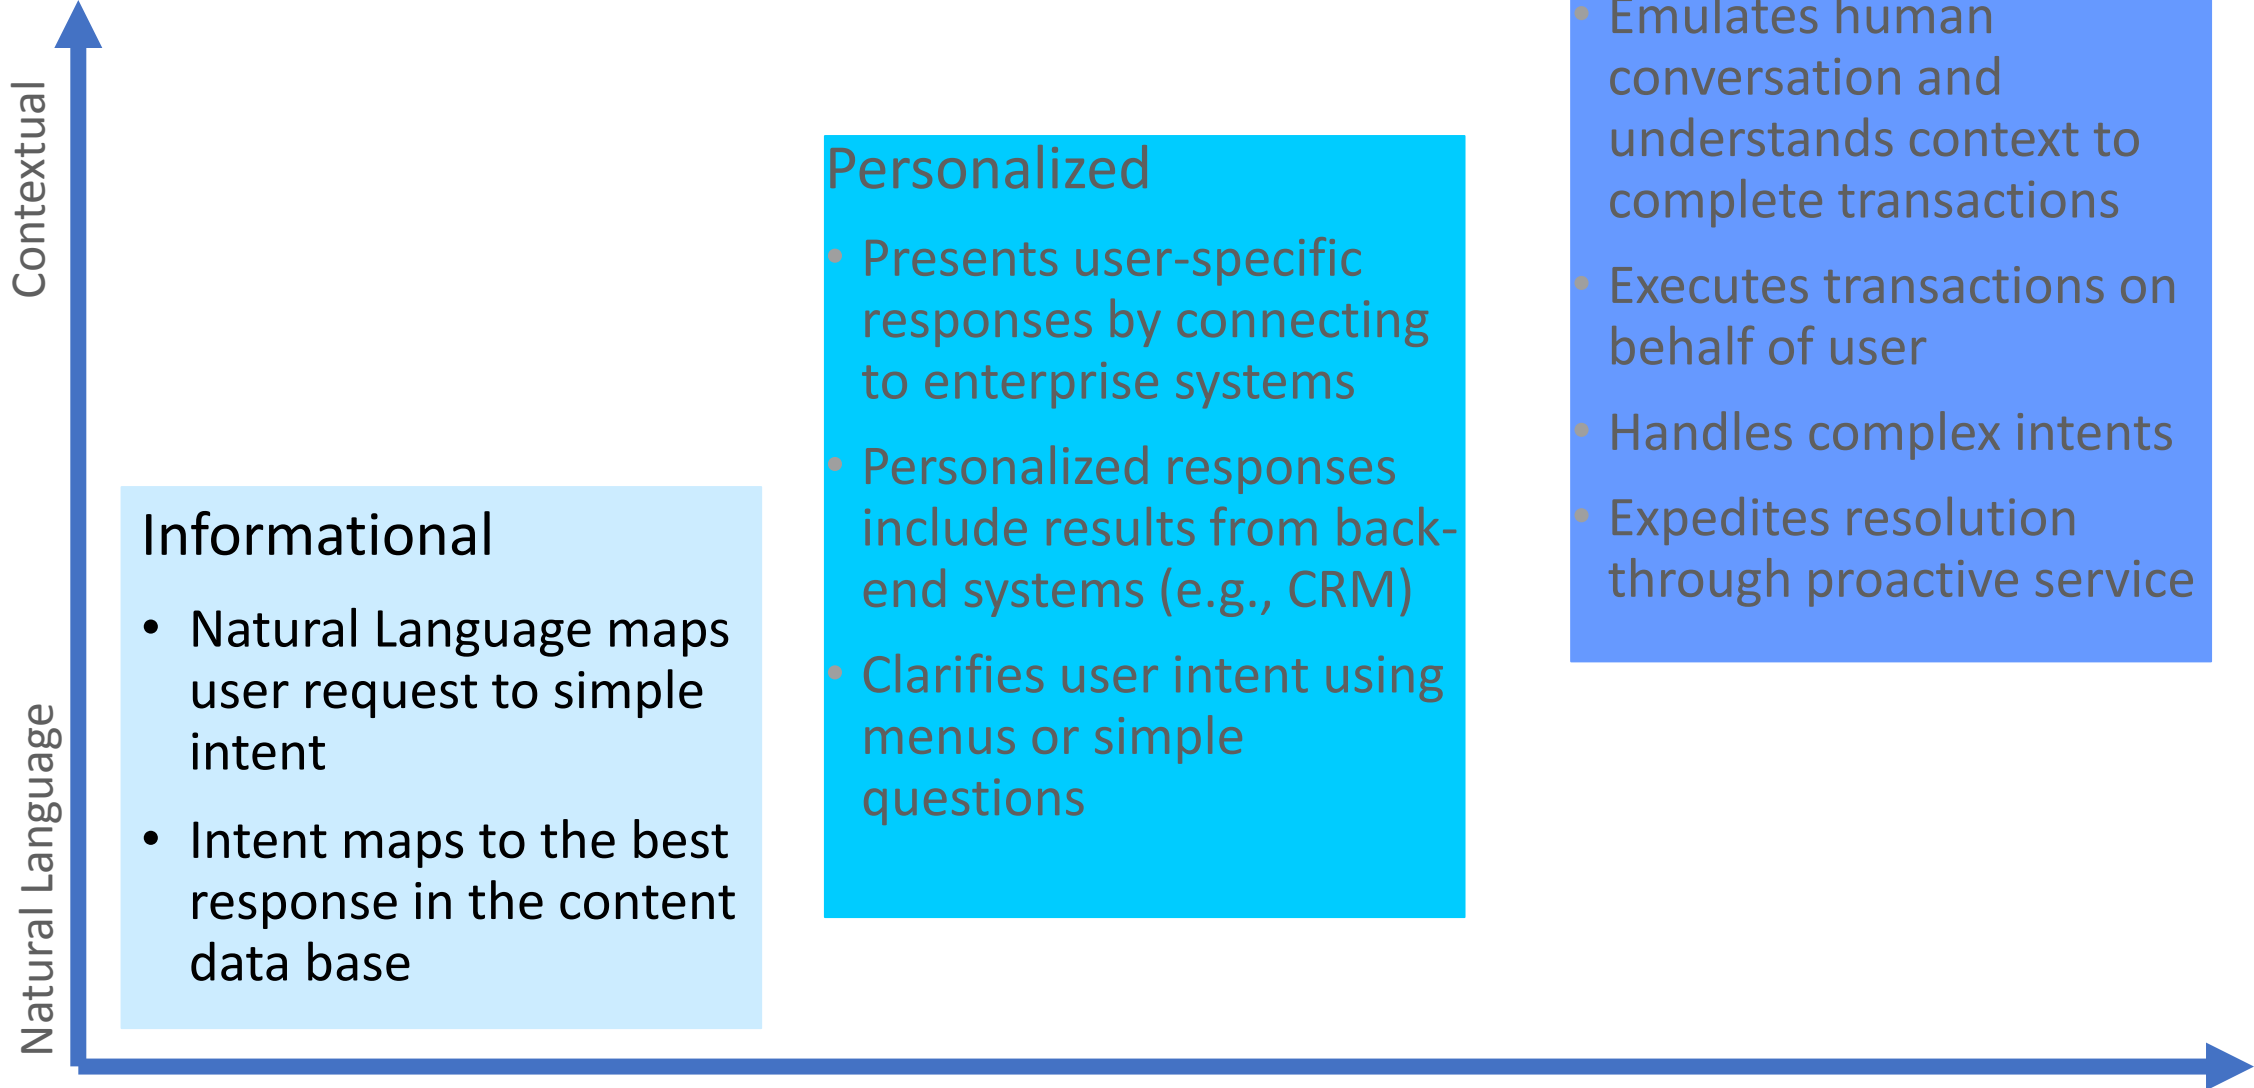
Transactional

Guides the user through a series of steps to complete a task (can also be conversational) and can integrate customer data

Reference: *Your Best Agent Is a Chatbot* by www.247-inc.com

Chatbot Maturity Framework

UNDERSTANDING



Information Provision

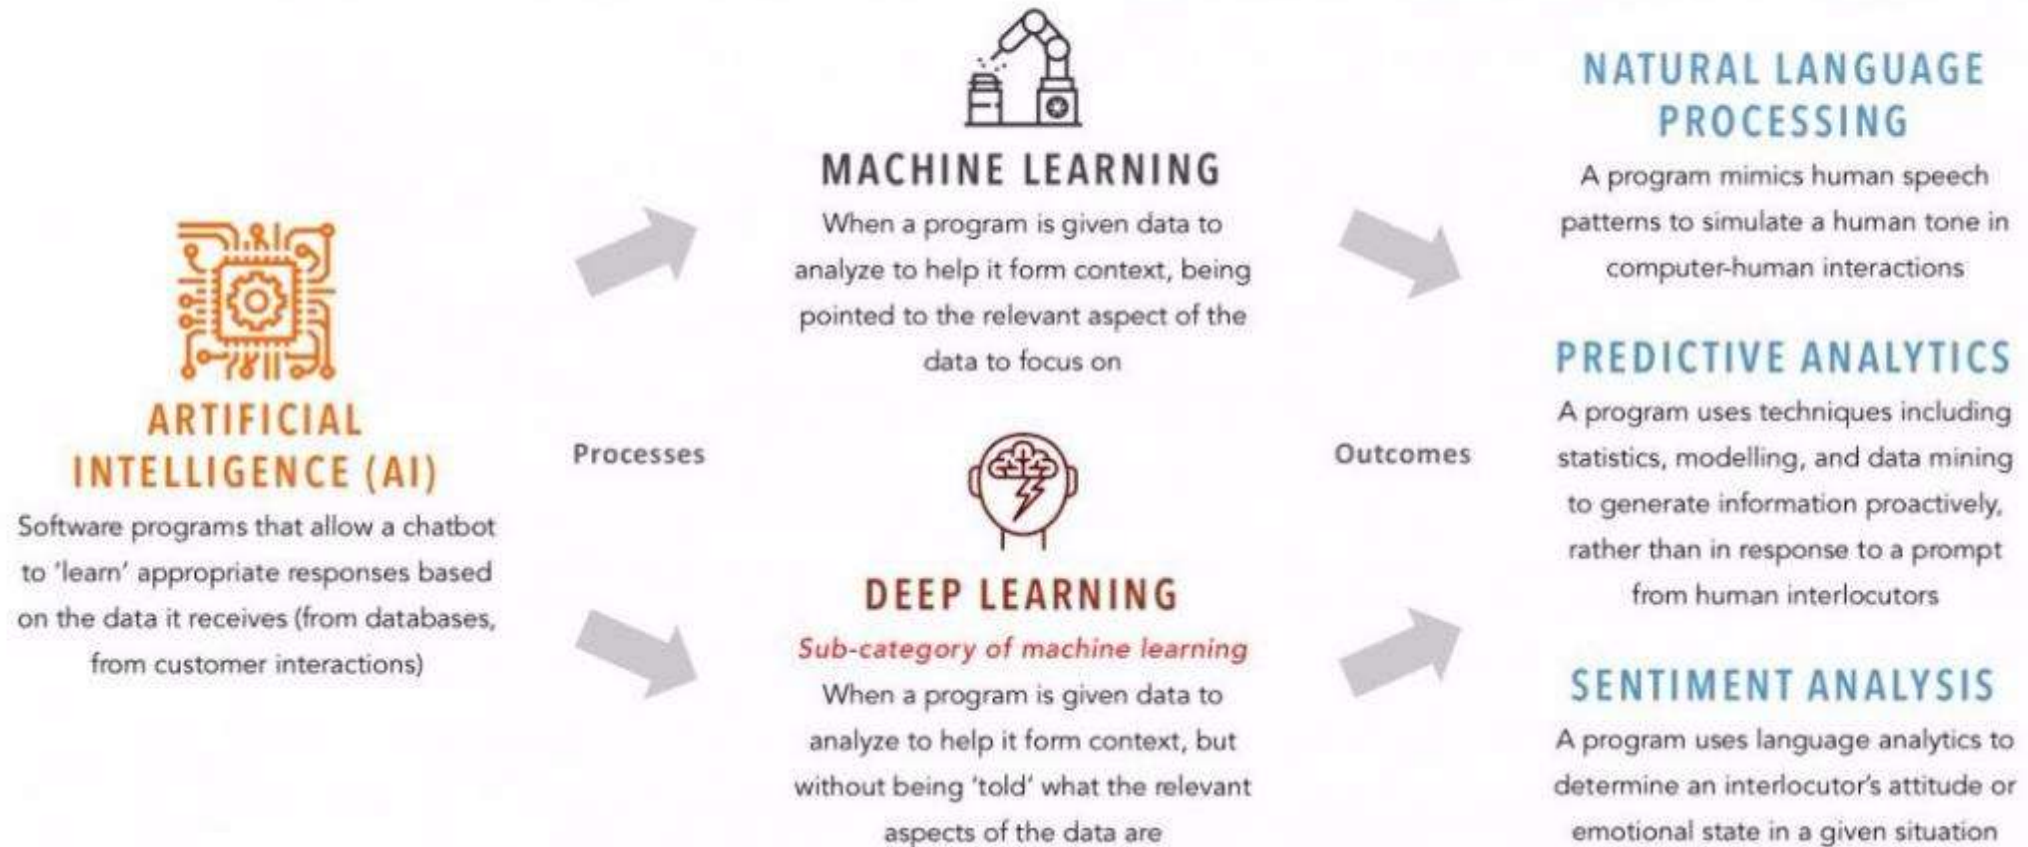
ENGAGEMENT

Transactional

Drivers for deploying Enterprise Chatbots?

1. Changing customer expectations
2. Decreasing customer satisfaction
3. Losing sales
4. Increasing volume
5. Reducing operating costs
6. Increasing visibility

CURRENT UNDERLYING CHATBOT TECHNOLOGY



Source: Business Insider

Pattern matchers

- Early chatbots used pattern matching to classify text and produce a response. This is often referred to as “brute force” as the author of the system needs to describe every pattern for which there is a response.
- A standard structure for these patterns is “AIML” (artificial intelligence markup language). Its use of the term “artificial intelligence” is quite an embellishment, but [that’s another story](#).


Algorithms

- The brute-force mechanism is daunting: for each unique input a pattern must be available to specify a response. This creates a hierarchical structure of patterns, the inspiration for the idiom “rats nest”.
- To reduce the classifier to a more manageable machine, we can approach the work *algorithmically*, that is to say: we can build an equation for it. This is what computer scientists call a “reductionist” approach: the problem is *reduced* so that the solution is simplified.
- A classic text classification algorithm is called “Multinomial Naive Bayes”, [taught in courses at Stanford](#) and elsewhere. Here is the equation:

$$\hat{P}(t|c) = \frac{T_{ct} + 1}{\sum_{t' \in V} (T_{ct'} + 1)} = \frac{T_{ct} + 1}{(\sum_{t' \in V} T_{ct'}) + B'}$$

Neural Networks

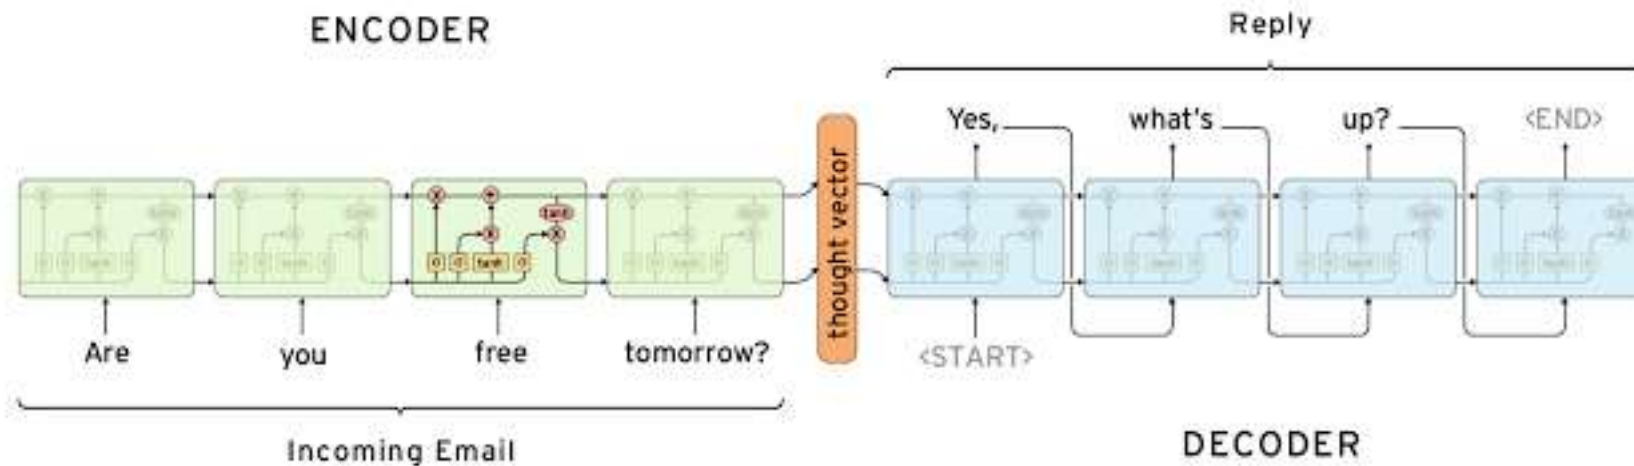
- Artificial neural networks, invented in the 1940's, are a way of calculating an output from an input (a classification) using weighted connections ("synapses") that are calculated from repeated iterations through training data. Each pass through the training data alters the weights such that the neural network produces the output with greater "accuracy" (lower error rate).



Retrieval-based vs. Generative models

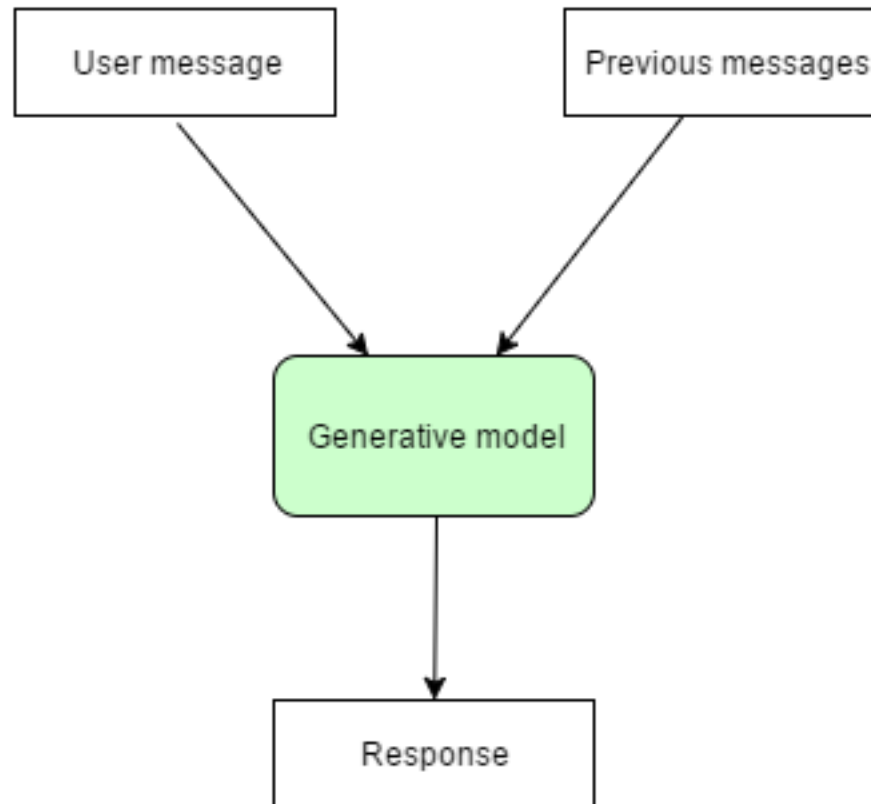
Retrieval-based models (easier) use a repository of predefined responses and some kind of heuristic to pick an appropriate response based on the input and context. The heuristic could be as simple as a rule-based expression match, or as complex as an ensemble of Machine Learning classifiers. These systems don't generate any new text, they just pick a response from a fixed set.

Generative models (harder) don't rely on pre-defined responses. They generate new responses from scratch. Generative models are typically based on Machine Translation techniques, but instead of translating from one language to another, we "translate" from an input



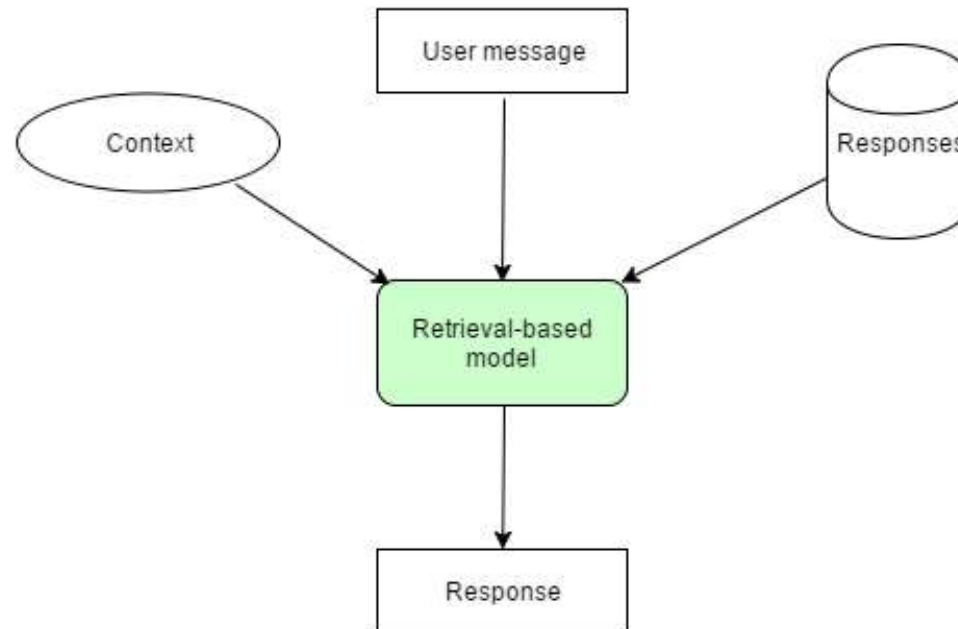
Generative models

- Generative models are the future of chatbots, they make bots smarter. This approach is not widely used by chatbot developers, it is mostly in the labs now.

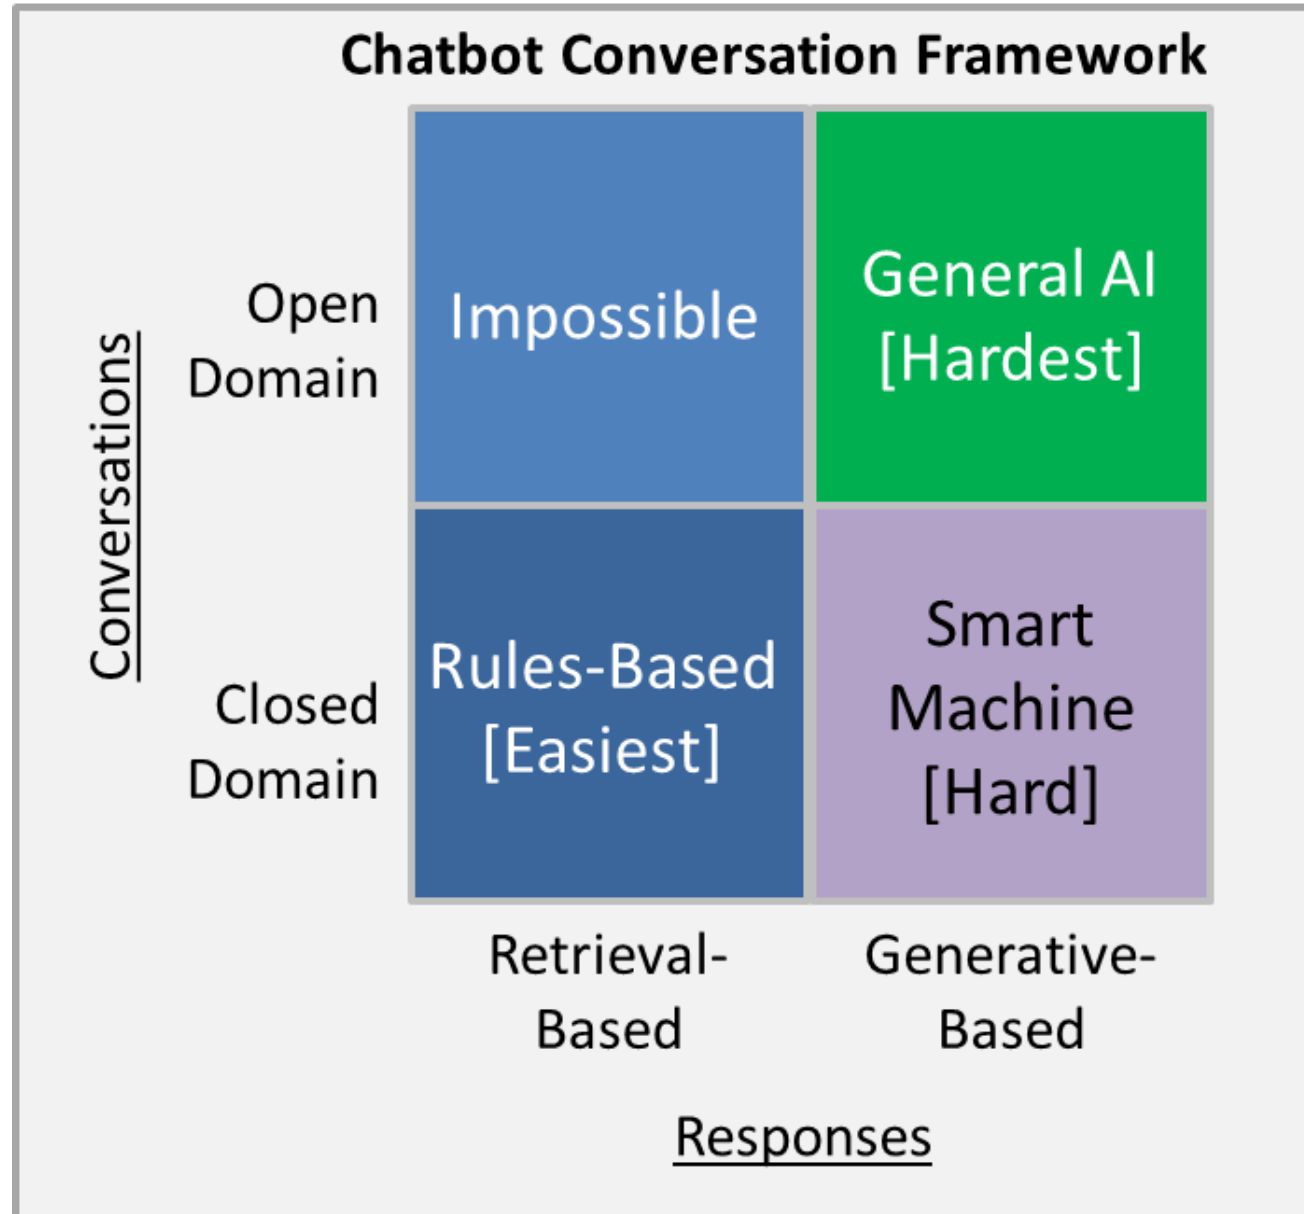


Retrieval-based models

- Retrieval-based models are much easier to build. They also provide more predictable results. You probably won't get 100% accuracy of responses, but at least you know all possible responses and can make sure that there are no inappropriate or grammatically incorrect responses.
- Retrieval-based models are more practical at the moment, many algorithms and APIs are readily available for developers.



Open domain vs. Closed domain



MACHINE INTELLIGENCE 3.0

ENTERPRISE INTELLIGENCE

VISUAL

Orbital Insight, Clarifai, DeepVision, Cortica, Iqonic, Space_Know, Captricity, Netra, Deepomatic

AUDIO

Gridspace, TalkIQ, Nexidia, Twilio, CAPIO, Expect Labs, Clover, Mobvoi, Curious.AI, Pop2P, Archive

SENSOR

PREDIX, G3 IoT, MAANA, Sentenai, Planet OS, Uptake, Imubit, Kinetica, Thingworx, KONIX, Alluvium

INTERNAL DATA

PRIMER, IBM Watson, Dycorp, Palantir, ARIMO, Alation, Sapho, Outlier, Digital Reasoning

MARKET

Mattermark, Quid, Datafex, Premise, Bottlenose, Motiva, Enigma, CB Insights, Tracxn, Predata

ENTERPRISE FUNCTIONS

CUSTOMER SUPPORT

DigitalGenius, Kasisto, Eloquent, Wise.io, ActionIQ, Zendesk, Proact, CLARABRIDGE

SALES

collective[i], Gsense, Fuse|machines, AVISO, Salesforce, Inside Sales, Clari, Zensight, Sales.com

MARKETING

MINTIGO, Lattice, Radius, LiftIgniter, Persado, Brightfuel, Retention, Cognicor, AIRPR, MSGQ

SECURITY

CYLANCE, Darktrace, Zimperium, Dexpinstruct, Sentinel, Demisto, Graphistry, Drawbridge, SignalSense, AppZen

RECRUITING

textio, Entelo, Wade & Wendy, Hi, Unifive, SpringRole, GIGSTER, HireVue

AUTONOMOUS SYSTEMS

GROUND NAVIGATION

drive.ai, AdasWorks, Zoox, Mobileye, Uber, Google, Tesla, Autonomy, Auto Robotics

AERIAL

SKYDIO, SHIELD AI, Airware, DJI, Lily, DroneDeploy, Pilot.ai, SKYCATCH

INDUSTRIAL

JAYBRIDGE, OSARO, CLEARPATH, Fetch, Kinross, Rethink Robotics, HARVEST

PERSONAL

Amazon Alexa, Cortana, Allo, Facebook, Siri, Replika

AGENTS

PROFESSIONAL

butter.ai, Pogo, Skipflag, Clara, X.ai, Slack, Talla, Zoom, Sudo

INDUSTRIES

AGRICULTURE

Blue@River, Mavrz, Tule, TRACE, Pivot Bio, TerraNova, Agri-Data, Descartes Labs, Udi, Spunforce

EDUCATION

Knewton, Volley, Gradescope, CTI, Coursero, Udacity, All School

INVESTMENT

Bloomberg, Sentient, iSentium, Kenshc, AlphaSense, Dotaminr, Cerbellum Capital, Quandl

LEGAL

BlueJ, Beagle, Everlaw, Ravel, Seal, Ross, Legal Robot

LOGISTICS

NAUTO, Acerta, Preteckt, Clearmetal, Routific, Marble, Pitstop

INDUSTRIES CONT'D

MATERIALS

Zymergen, Citrine, Eigen Innovations, Sight Machine, Ginkgo Bioworks, Nanotronics, Calculario

RETAIL FINANCE

TALA, Best Finance, Lendo, Earnest, Affirm, Mirador, Wealthfront, Betterment

PATIENT

PULSE, CareSkore, Zephyr Health, Watson Health, Oncoda, Sentriam, Atomwise, Numerate

IMAGE

Butterfly, 3SCAN, Arterys, Enlitic, Baylabs, Imagia, Google DeepMind

BIOLOGICAL

CarbonX, Color, Grail, Deep Genomics, Recursion, Luminst, Numerate, Atomwise, Verily, Whoish

TECHNOLOGY STACK

AGENT ENABLERS

OCTANE.AI, howdy, Maluuba, KITT.AI, OpenAI Gym, Kasisto, AUTOMAT, semanticmachines

DATA SCIENCE

DOMINO, SPARKBEYOND, rapidminer, kaggle, DataRobot, yhat, AYASDI, data iku, seldon, yseop, bigml

MACHINE LEARNING

CognitiveScale, GoogleML, context, relevant, Dycorp, HyperScience, nora, logics, minds.ai, H2O.ai, SCALED INFERENCE, sparkcognition, loop, GEOMETRIC INTELLIGENCE, deepsense.io, reactive, skymind, bonsai

NATURAL LANGUAGE

agolo, PYLIEN, LEXALYTICS, Narrative Science, spaCy, LUMINOSO, cortical.io, MonkeyLearn

DEVELOPMENT

SIGOPT, HyperOpt, fuzzyio, pkite, rainforest, lobe, Anodot, Signifai, LAYER 6, bonsai

DATA CAPTURE

CrowdFlower, diffbot, CrowdAI, import, Paxata, DATASIFT, amazon, mechanical, turk, enigma, WorkFusion, DATALOGUE, TRIFACTA, parsehub

OPEN SOURCE LIBRARIES

Keras, Chainer, CNTK, TensorFlow, Caffe, H2O, DEEPLARNING4J, theano, torch, DSSTNE, Scikit-learn, AzureML, neon, MXNet, DMTK, Spark, PaddlePaddle, WEKA

HARDWARE


KNUPATH, TENSTORRENT, Cirrascale, NVIDIA, intel, nervana, Movidius, tensilica, GoogleTPU, 10th Labs, Qualcomm, Cerebras, Isosemi

RESEARCH

OpenAI, sense, ELEMENT, vicarious, KNOGGIN, Numeria, Kimera Systems, Cogitai

Example: Airport Wifi Bot

Airport Wifi Bot
Active on Messenger

 Airport Wifi Bot
4 people like this
Product/Service

10:56AM

Hey Kashif, I'm Airport Wifi Bot (I'm changing my name to Porter soon to avoid a lawsuit)!

I'm here to help you get wifi at any airport you're at. All you need to do is type in the airport code (for example: DEN) and I will give you the information I have on it.

Feel free to submit info below and it will be added to my list of information!

— To restart me press the menu in the bottom left or type restart —

[Submit Wifi Password](#)

Some information may be incorrect seeing that it has not been updated recently.

Some passwords with inclusion of the a year may also be incorrect. If that is the case, try the current year in the password!

Go ahead! Type the airport code or type show me examples!

Get

DXB

You need to find one of the following!

The Emirates Lounge

The password is:
ekloungedxb

The British Airways Lounge

The password is:
vancouver

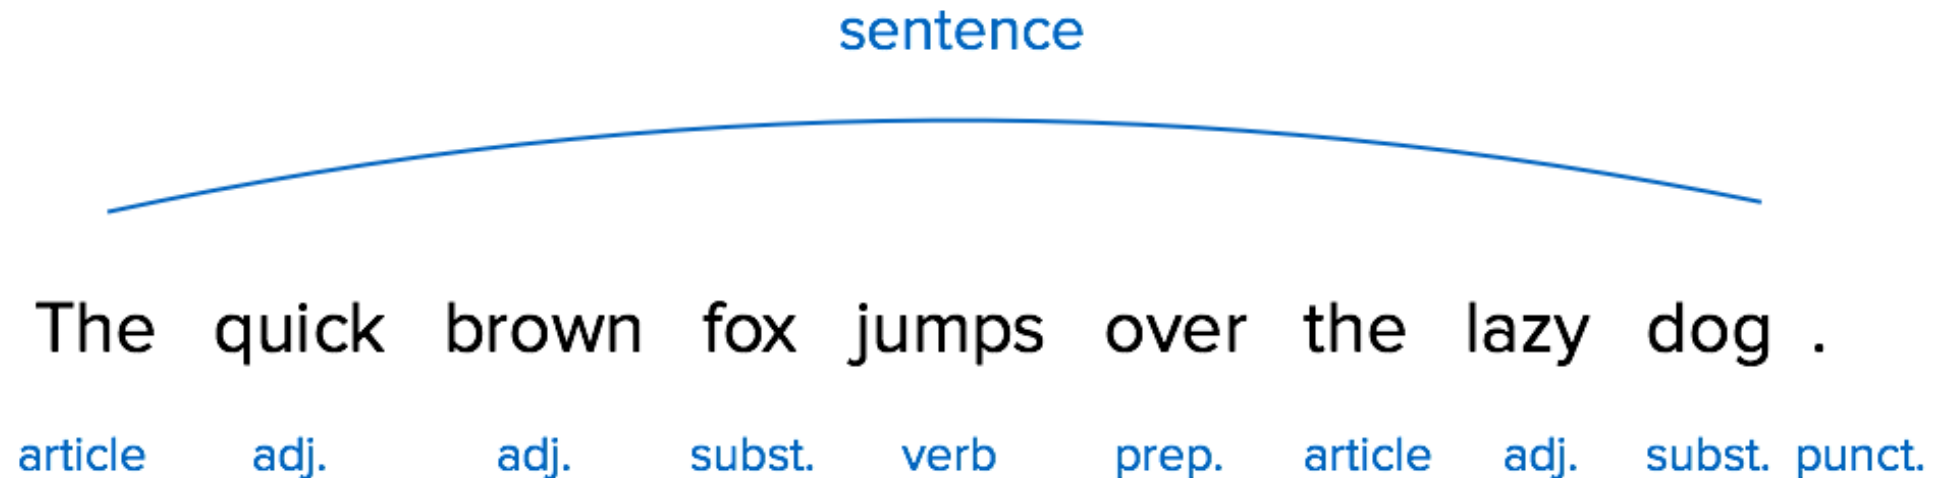


Natural language processing (NLP)

Natural language processing (NLP) is a field of [computer science](#), [artificial intelligence](#) and [computational linguistics](#) concerned with the interactions between [computers](#) and [human \(natural\) languages](#), and, in particular, concerned with programming computers to fruitfully process large [natural language corpora](#). Wikipedia

The process of NLP consists of roughly 5 steps.


1 / The first step is lexical analysis. The lexicon of a language is, simply put, a collection of words and phrases in a language. As a first step, the computer will thus analyse the text and divide it into paragraphs, sentences and words.



Example of lexical analysis

The process of NLP consists of roughly 5 steps.

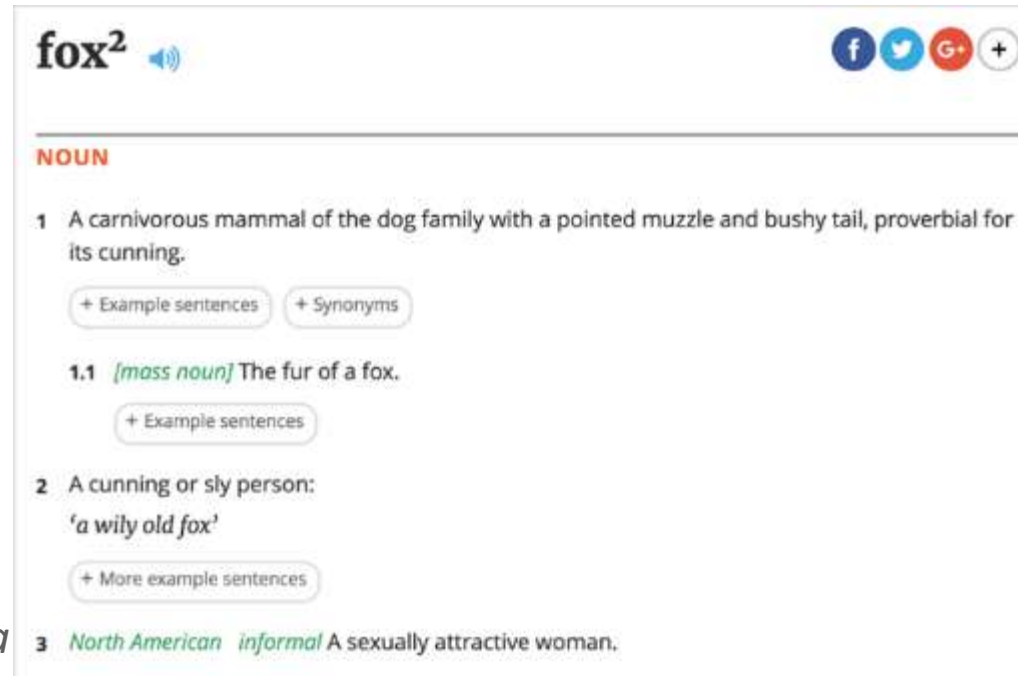
2 / The second step is the syntactic analysis: the computer analyses the grammatical role of each word in a sentence and identifies the relationship between each word. This is something you probably learned in school: what is the subject of the sentence? Is there a predicate?



Example of syntactic analysis

The process of NLP consists of roughly 5 steps.

3 / In the third step, the semantic analysis, the computer checks the intrinsic meaning of the words, so that means looking up the meaning of the words as stated in the dictionary. A word can have several meanings, so the computer also needs to map this with the syntactic structures analysed in the previous step to derive the correct meaning.



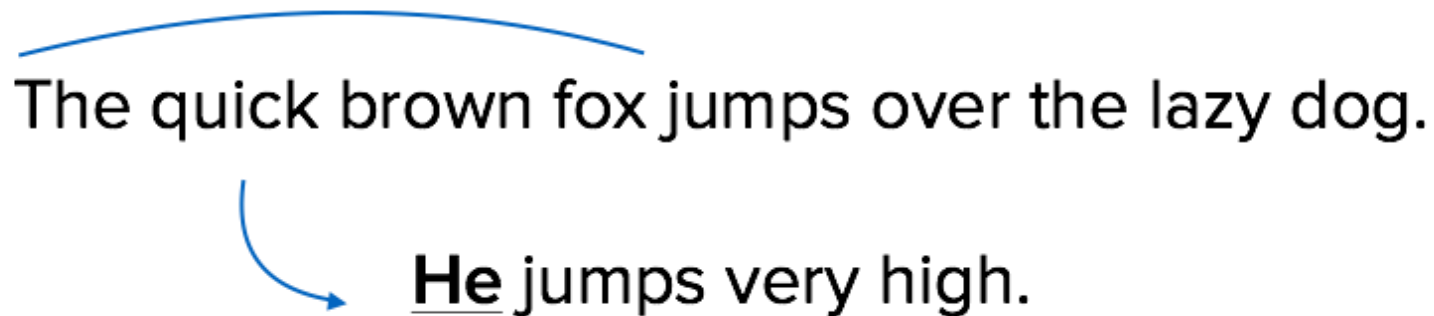
The screenshot shows a dictionary entry for the word "fox". At the top, it says "fox²" with a speaker icon and social media icons for Facebook, Twitter, Google+, and a plus sign. Below this, the word is categorized as a "NOUN". There are three numbered definitions:

- 1 A carnivorous mammal of the dog family with a pointed muzzle and bushy tail, proverbial for its cunning.
+ Example sentences + Synonyms
- 1.1 [mass noun] The fur of a fox.
+ Example sentences
- 2 A cunning or sly person:
'a wily old fox'
+ More example sentences
- 3 North American informal A sexually attractive woman.

Example of different meanings of a

The process of NLP consists of roughly 5 steps.

4 / The fourth step is discourse integration, which means looking at the meaning of a sentence compared to the sentence that comes before it. We can assume that there is cohesion between the different sentences in a text, so NLP must also take this into account.



The quick brown fox jumps over the lazy dog.

He jumps very high.

The diagram shows two sentences. The first sentence is "The quick brown fox jumps over the lazy dog." A blue curved arrow starts from the end of this sentence and points to the second sentence, "He jumps very high." The word "He" in the second sentence is underlined, indicating a coreference with the "fox" in the first sentence.

Example of discourse integration

The process of NLP consists of roughly 5 steps.

5 / Finally, there is the pragmatic analysis, which is also the most difficult step for a computer. The pragmatic analysis involves re-interpreting what is said as what was actually meant. This involves taking knowledge from the real world into account because as humans, what we say is not always what we mean. Take for example the sentence: “There’s beer in the fridge”. If you say this to a guest entering your house, you are not simply describing the contents of your fridge, you are actually offering them a drink. This ambiguity is hard for a computer to handle.

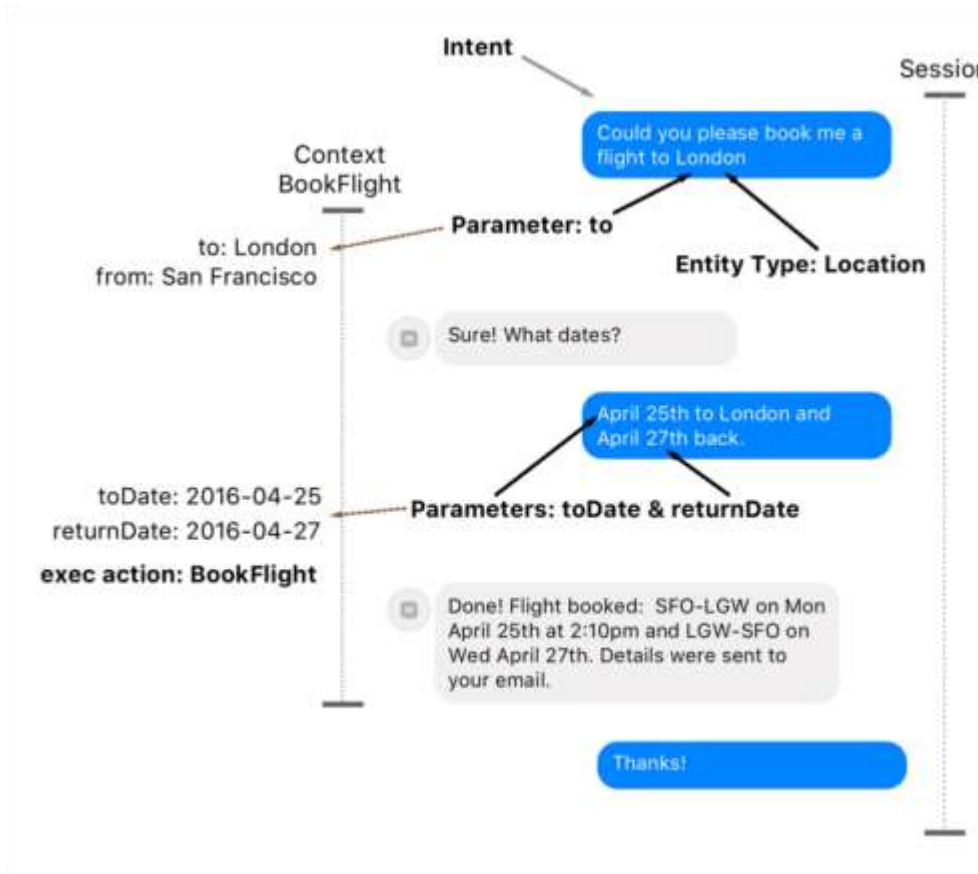
Example of discourse integration

<http://www.icapps.com/the-linguistics-behind-chatbots/>

Machine Learning

Machine learning makes bots smarter

Machine learning (ML) can be defined as an algorithm of making systems learn, by using observations or past experience. Instead of hand-coding large sets of rules, NLP can rely on ML to automatically learn these rules by analysing a corpus. A corpus can be a book, news articles, reports or even conversations. If a bot contains algorithms for machine learning, it becomes smarter the more people talk to it.





Rise of the Chatbot Economy



38%

prefer to use instant messaging to communicate with brands, compared to 19% of Generation X.



69%

of them say they feel good about the company when they solve a problem without talking to customer service. They're experience-loyal than brand-loyal.

\$700

billion in annual buying power by 2017



MILLENNIALS

Is your business ready to serve the "me me me generation"?

2020



Is the year that they will account for 40% of all consumers.

Millenials to brands:

"I won't come to you, but I might meet you half way"



THEY WANT BRANDS TO BE AWARE OF THEM AS INDIVIDUALS RATHER THAN A NAME OR NUMBER ON A LIST.



Siri gets over 1 billion requests a week

Do better than meeting your customers half way.

Give them their own personal intelligent customer service agent who knows and addresses them personally and available 24/7 in their pocket.



Visit <http://botego.com/GiveThem> for a free demo and more details.



Millennials don't use email

THEN

Email / Semi-inflexible
messaging tools




NOW

Slack, FB Messenger, WhatsApp,
Instagram, Snapchat, WeChat



Gartner Hype Cycle

- Machine Learning
- Natural-Language question answering



Conversational Systems: The Post-App Digital Mesh Experience

Chatbots and Personal Assistants as
Conversational Intermediaries
With Intelligent Cloud Services



Xiaoi



Siri



Cortana



Amazon



... Just the Beginning

Conversational +

#TopTechTrends

26 CONFIDENTIAL AND PROPRIETARY | © 2016 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner and ITxpo are registered trademarks of Gartner, Inc. or its affiliates.

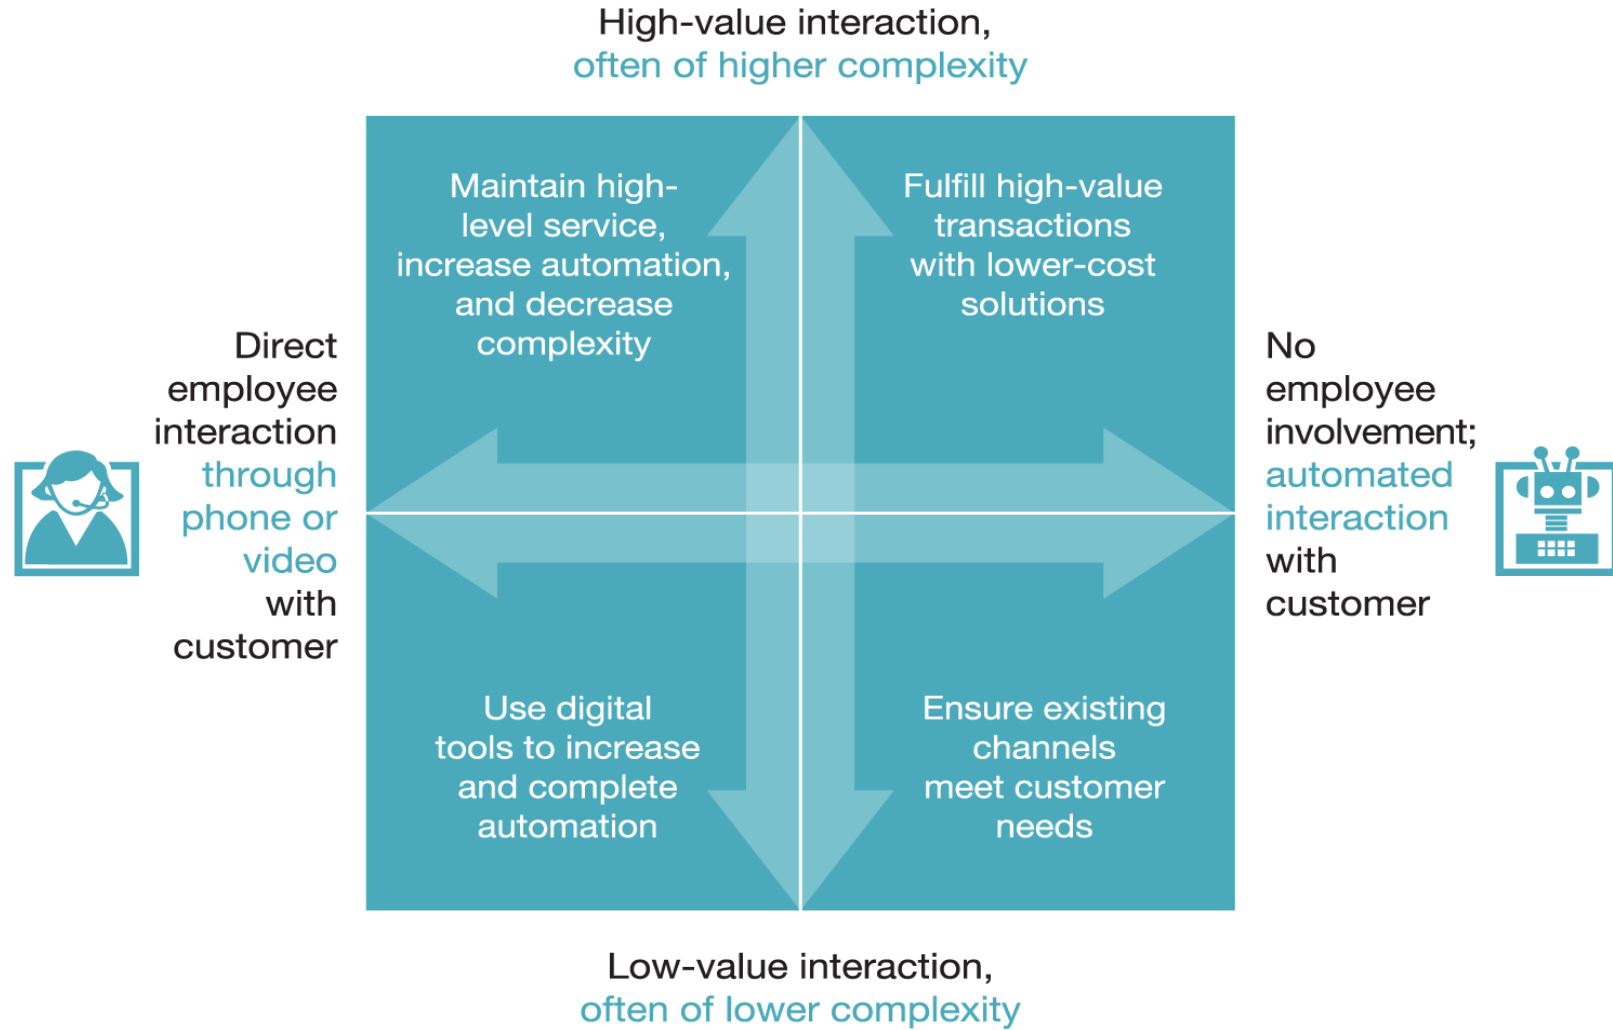
Gartner

Gartner predicts that artificial intelligence will amount for 85% of customer relationships by 2020

A [McKinsey survey](#) from 2015 estimated that digital-care channels (e.g. web chat, social media, and email) accounted for 30 percent of customer-care interactions and that by 2020 it is expected to grow to 48 percent.

In the future, customer-care organizations will balance personal interactions with automation.

Service description by segment




Chatbots – the virtual customer service assistant

2016 witnessed many retailers experimenting with their own virtual chatbots and in total, [34,000 different bots](#) have emerged. Chatbots, powered by a combination of machine learning, natural language processing, and live operators, can provide customer service, sales support and make suggestions for what to buy at a much greater level of detail than ever before.

Messaging Apps Have Surpassed Social Networks

Monthly active users for top 4 social networks and messaging apps

— Big 4 Messaging Apps
— Big 4 Social Networking Apps



BI INTELLIGENCE

Source: Companies, BI Intelligence

SO YOU WANT TO BUILD YOUR
OWN ?