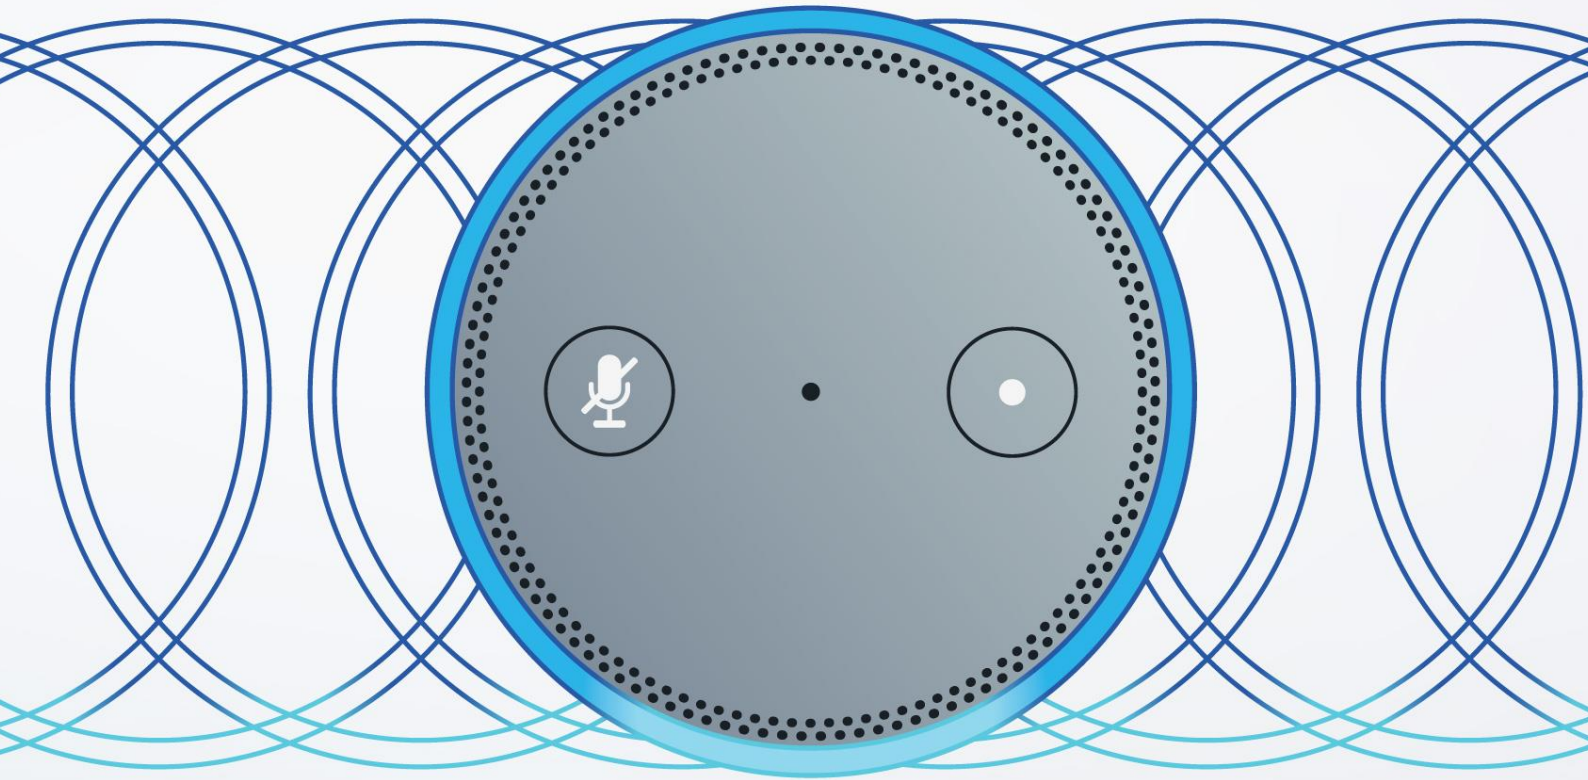


Internet of Things: A Research on the Integration of IoT Devices in Our Daily Life



Darmstadt University of Applied Sciences

Thesis by Ronald Yozafath Norori Pineda

Master's in Leadership in the Creative Industries

Supervised by

Prof. Dr. Frank Gabler

Prof. Sven Poguntke

ACKNOWLEDGMENTS

Thanks to my family, for backing me up during my studies.

Thanks to Julian Günther, José Pablo Umaña, Steven Weiss and Robert Veres for their support during the production of this document.

To my supervisors Prof. Dr. Frank Gabler and Prof. Sven Poguntke for their guidance.

DECLARATION OF AUTHORSHIP

I hereby certify that this thesis has been composed by me and is based on my own work, unless stated otherwise. No other person's work has been used without due acknowledgment in this thesis. All references and verbatim extracts have been quoted, and all sources of information, including graphs and data sets, have been specifically acknowledged.

Ronald Y. Noroni

05.03.2018

Signature

Date

ARCHIVING AND PUBLICATION

In submitting this thesis to the Darmstadt University of Applied Sciences. I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby.

Ronald Y. Noroni

05.03.2018

Signature

Date


Table of Contents

ABSTRACT	7
Introduction.....	9
History.....	11
Home	13
Industry 4.0	14
Smart Cities	15
Health	16
Other Applications.....	17
Smart Speakers	18
Amazon Echo Line.....	19
Google Home Line	24
Virtual Assistants	26
Speech Recognition	26
Other enablers of the modern virtual assistants	30
Cloud Computing	30
Wireless Technologies	32
ZigBee	32
Z-Wave.....	32
Bluetooth	33
Modern Virtual Assistants.....	34
Amazon Alexa	36
Google Assistant.....	37
Siri and HomePod	38
Adoption & Usage of Smart Assistants.....	39
Consumer Usage.....	42
Primary research on smart-speakers	47
Methodology.....	48
Opinion and Usage Survey	49
Distribution.....	49
Usage	50
Drawbacks	51
Outlook.....	51
Trials	52
Trials Survey.....	52
Usage	52
Drawbacks	53
Outlook.....	53
Trial Participants Assessment.....	54
Conclusions of the Primary Research.....	58
Third-party applications	58
Privacy	59

Closing notes	59
Conception of apps for Alexa.....	60
Amazon Echo Plus	61
Amazon Echo Dot 2 nd Generation	62
The development environment.....	63
Intents	64
Designing for voice	66
Developing for Amazon Alexa.....	67
Learning Skill: Darmstadt Tips	68
Amazon Developer Portal.....	69
Amazon Web Services Console.....	70
Learning Skill: Hessen Trivia	72
Custom Skill: Help Me Decide	74
Custom Skill: Walked the Dog	75
Prototypes	76
Conclusions.....	77
Personal Learnings	79
Appendix	81
Survey Results: Opinion and usage patterns of smart-speakers and smart-assistants	81
Survey Results: Trial Period Survey	86
Source Code: Help Me Decide (JavaScript).....	91
Source Code: Walked the Dog (Python)	93
References	96

DIGITAL VERSION

Scan this Near-Field Communication (NFC) tag with an enabled device to download the digital PDF version of this document.



thesis.norori.com/url/pdf

ABSTRACT

The Internet of Things has become a common topic in the technology industry, even though its roots can be tracked to decades ago, it didn't make an impact until the last few years with the dawn of the smart-speakers and smart-home appliances. This paper explores the origins of the industry and analyzes the current usage costumers give to the available devices.

Employing primary research tools like surveys, trials and interviews, this paper gives the reader an exhaustive overview of the current state of the smart-home market, and more precisely of the smart-speakers market and environment.

Focusing on the most popular family of smart-speakers, the Amazon Echo line, I analyze the patterns of interactions users have with the devices, their most used functions and desired missing features.

The study reveals that while the consumers are generally satisfied with the device, some challenges in the medium and long term exist, challenges developers and manufactures must not overlook. Recommendations and learnings from the study are shared in detail.

For the final part of the paper, the knowledge gathered during the research part is applied practically with the creation of prototypes of Alexa applications based on best-practices outlined from the behavior patterns observed during the trials.

The reader can expect to get a complete overview of the current state of the industry of smart-speakers, its challenges and outlook with this paper.



Introduction

During the past years, the relevance of the Internet of Things (IoT) has sky-rocketed to unexpected levels, reaching millions of users and thousands of devices ranging from simple sensors to dedicated smart-home systems with voice and video functions. Behind this powerful growth lays hefty innovation and dedication from tech companies and professionals who keep building the biggest and most connected network in the history of mankind.

Like with many technologies at the beginning of their development, it's difficult to predict the reach of impact that this massive connected network will give us, or if it will fade away with time.

The recent boost on the presence of IoT devices in homes, offices, industrial environments and public places all around the world is sparking the creativity and innovation of developers that find fresh approaches to solve problems and create applications for them.

It can't be denied that the IoT is already changing the world, it can be also assumed that in the next decades its impact will be multiplied affecting in positive and negative ways many aspects of our lives.

Cover Image: Designed with sources from FreePic

Motivated by a personal interest in being an early adopter, the broad range of possibilities at reach and significant lack of public information regarding the design and development of functional and practical applications, I embarked on this research about the history, current state and possible evolution of the IoT, while also creating prototypes, learning and designing guidelines to help others take advantage of its enormous potential.

I believe is now evident how the IoT will unquestionably affect our lives in more positive ways than negatives. Once the technology reaches and advanced maturity and its adoption is widespread, the possible appliances are endless.

This thesis is product of an extensive research based on primary and secondary sources.

Primary research is done throughout a series of surveys and trials that help me understand the way people use and interact with the devices of study.

The main objective of this paper, is to give anyone who reads it thoroughly an accurate and detailed understanding of the Internet of Things, more specifically of the smart-speakers and their adjacent markets and technologies while simultaneously inspiring and instructing them on the design and development of useful tools for the devices.

History

What is the Internet of Things? That was a common question I received during the writing of this paper, the term has floated around since the late 90's but it didn't become mainstream until a few years ago, nowadays is normal to see it online.

The term was coined by the British pioneer Kevin Ashton to describe “a system in which sensors could connect to the internet and other computers”. The inspiration for the term came from a Radio-frequency identification (RFID) system used to digitally track supplies for a mass-production factory without the intervention of human operators. “IoT” is also seen as an evolution of the term “M2M” or “Machine to Machine” used to describe similar interactions between machines since the 1950's¹.

One of the early modern examples of the IoT was an IP-enabled toaster that could be turned on and off over the internet², humans were still needed to put the bread in, but it was just a question of time until a IP-enabled mini-crane capable of picking up slices of bread and dropping them inside the toaster was added, giving the system a total autonomy. Less trivial appliances emerged with time: IP-enabled soda machines and coffee makers popped up in some US American universities.

Connecting new kinds of devices to the internet is not a new idea, for long time industrial appliances like ATMs, ticket dispensers, vending machines and traffic lights have been networked.

Notes

1. *The Silent Intelligence: The Internet of Thing* Book by Daniel Kellmerit and Daniel Obodovski
 2. *The Living Internet - "The internet Toaster"* - <http://thesis.norori.com/url/wkn49>
- Cover Image: Designed with sources from FreePic

These networks never caused as much attention as the current wave of devices getting online, and it's not difficult to understand why: This time the revolution is taking place in our own homes, affecting us in a direct and noticeable way.

Many factors have driven this notorious growth, and as natural in technological development. Moore's Law can be used to understand it better: devices tend to get consistently smaller, more efficient, powerful and considerably cheaper with time.

The reduction in size and the great power and efficiency of contemporary micro-chips and boards made possible for designers and engineers to integrate them into what before were unthinkable devices. The increased demand for these components also made them affordable.

Wireless technologies also contributed considerably to the emergence of the modern IoT. During the past decade Wi-Fi became an absolute basic need in consumer devices. Cellular networks are now covering great percentage of populated areas in the developed and emergent economies, while data costs have dropped continuously, being connected is no longer an unaffordable luxury.

As a direct consequence of the great number of connected users and devices, cloud computing raised as the final catalyst to the IoT formula: the cloud allowed service providers to collect, process and serve massive amounts to their user-base while being able to take advantage of all the analytics that could be extracted from their data.

The confluence of these factors and the fertile innovation era that the internet is currently going through sparked the hefty soar of The Internet of Things.

The Internet of Things is still in what can be considered an early stage of development, but despite this its reach is far and wide. Gartner estimated that in 2017, 8.4 billion devices varying from smart assistants to street lamps were connected to the internet¹. The number is expected to grow to almost 50 billion by just 2020.

Creativity and innovation push the boundaries of the IoT further every year, and today its present in all kinds of environments, they can be divided in:

Home

Smart Appliances have been driving the increased awareness on the IoT, since they are the major contact point the technology has with consumers.

In the Consumer Electronics Show 2018 it was clear the Smart Home will be a major focus point for manufactures in the next years. Samsung promised in 2016 that all of their products would be IoT-enabled by 2020, keeping their compromise, by 2018 already 90% of their products were ready to connect to the internet out of the box², their portfolio includes fridges, washing machines, thermostats, microwaves and entertainments systems, among dozens more.

Smart Home assistants like Amazon Alexa and Google Home are important part of the ecosystem, serving as main interaction point between the user and the devices connected to a home network. Their capabilities have been increasing exponentially in hand with their sales. CIRP reported³ that as of January 2018, in the United States there were:

- Amazon Echo (Alexa): 31 million units
- Google Home: 14 million units

A similar market-share is expected in the global markets, with Amazon ahead of Google in all reported countries. Apple joined the race in early 2018, with the HomePod, it was received with mixed reviews and no numbers on sales have been release yet.

Smart-speakers and other Smart Home appliances are driving the increasing growth of the IoT market in general, and since a few devices like the Amazon Echo Plus start to integrate smart-home hubs it also makes them an ideal “one-in-all” solution for simple smart-home setups.

Not many smart-speakers have integrated smart-home hubs, but this is a decisive function that will surely become a standard.


Notes

1. Gartner Press Release - February 7, 2017 – <http://thesis.norori.com/url/-tbki>

2. Fortune: Samsung is Betting Big on the Internet of Things – <http://thesis.norori.com/url/mbhpx>

3. Consumer Intelligence Research Partners (CIRP) Report – January 2018 – <http://thesis.norori.com/url/8qa8c>

The sales of Smart Assistants are expected to grow in great scale in the next years, in part triggered by new major manufacturers like Apple and Bose that join the trend releasing their own devices.



Source: Canalis estimates and forecasts, Smart Speaker Analysis, December 2017

Canalis expects the global market of smart-speakers to reach the staggering number of 56 million total shipping's by the end of 2018¹. This will provide manufactures and developers a fertile field for the development of creative and useful applications.

Industry 4.0

The integration of the Internet of Things into the mass production assembly lines is seen by many as the “fourth industrial revolution”, optimizing procedures and incrementing productivity while reducing costs on the long term are great benefits, but they carry the danger of the loss of jobs by the hands of automation.

Technologies that in principle enable Industry 4.0 have been around for a long time. Radio-Frequency identification, better known as RFID has been used in manufacturing industries for long time, mainly in inventory control and asset management.

RFID provides easy identification and exchange of data between passive tags and active readers. Tags can be as small as 0.05 mm × 0.05 mm and have been successfully glued to live ants for research of behaviour and location².

Notes

1. Canalis: January 2018 Report – <http://thesis.norori.com/url/xdwio>

2. BBC News: Ants' home search habit uncovered– <http://thesis.norori.com/url/hhqh->

The Big Data that results from massive and extended use of RFID provide manufactures with large amounts of information that help reveal patterns, trends, and associations. This becomes particularly handy when managing inventory and shipping in big scale, making each item easier to track and giving a full over-view that can help to make fast decisions.

Companies like Amazon use RFID tracking to organize their warehouses and keep track of inventory and shipping. Almost all fashion retailers already use RFID security solutions to prevent shop-lifting, have you seen the detectors at the entrance of any major retail shop? They're reading every tag that goes through the door and compare them to the database of products that are allowed out.

Smart Cities

Smart cities definitely broaden the definition of what can be considered part of the Internet of Things, with many devices that one would never expect or think would be connected to the internet. A great example of this are Dublin's Smart Trash Bins, that report the exact levels of usage of each one of their compartments and ask to be emptied once they're almost full. They give the city important insight in how to manage their trash disposal by letting them know what categories of trash are filling up quicker and which areas have more disposal.

Smart Traffic is a more traditional but effective approach cities have with the IoT. Gathering information from sources like cameras and in-asphalt sensors, computes can optimize the cycles of the traffic lights improving wait times and reducing pollution caused by standing cars. A study done by the Carnegie Mellon University in Pittsburgh, USA found that when "smart signal" technologies were used, the wait times were cut by 40% and the exhaust emissions were reduced by 21%³.

Notes

1. RFID Journal – Amazon RFID – <http://thesis.norori.com/url/ce3pi>
2. Smart Dublin Project - <http://thesis.norori.com/url/jjzjo>
3. "The Internet of Things: Changing the World" – Book by Michael Miller

While other areas of application are still not as developed as Smart Traffic¹, some experimenting has been done in other aspects of city regulation and management. In Asia, the government of Japan has been testing early-warning earthquake systems that notify the population about upcoming telluric movements before they start feeling them, sometimes even a few minutes before the expansion waves hit the furthest areas from the epicentre. US-American start-up “Zizmos”³, funded by the American National Science Foundation is also researching and developing similar systems of connected sensors that can be deployed in highly populated areas, helping people take shelter in the crucial seconds before the earthquake hit.

Health

According to the World Health Organization there are more than 1 billion overweight adults, 860 million chronic disease patients and over 600 million people over 60 years of age. Humans are also living longer than ever, with the average of age in developed nations already around 75 years of life expectancy.

The Integration of the IoT into health-care is another big revolution. Doctors and nurses can track the health of specific patients over time using sensors and researchers can make use of the big data generated by larger pools of patients to find patterns.

Blood pressure and heart rate monitors, blood glucose and composition meters and sleep trackers are some of the sensors already approved to be used by patients at home in some countries⁵, transmitting results instantly to clouds that can give early readings to the patients while alerting about unusual or dangerous signs to family and medical staff. This approach is especially helpful with patients living alone or having limited capacities due to dementia.

Notes

1. “Internet of Things. User-Centric IoT” – Book by Raffaele Gaffè, Radu-Laurentiu Vieriu, Edna Pasher, Gabriel Bendersky, Antonio J. Jara, Joel J.P.C. Rodrigues, Eliezer Dekel, Benny Mandler / LNICST

2- Huffington Post: IoT Provides Affordable Earthquake Early Warning to Communities by Robert Armitano - <http://thesis.norori.com/url/tbmkk>

3. Zizmos – <https://www.zizmos.com>

4. World Health Organization – <http://www.who.int/research/en/>

5. NXP Medical Applications Guide – <http://thesis.norori.com/url/gmior>

Development of IoT appliances on the health-care field are of great help, while there is plenty of space for further developments, their implementation speed is hindered by the complex set of rules and regulations that any new technology in the health-care industry has to go through. Despite this, many companies are investing important efforts in research and development of medical applications using the Internet of Things¹.

Fitness trackers are also an important part on this category, and one of the business models of the IoT with most revenue². Wearables like Smart Watches with GPS and Wireless Internet functions are getting more advanced, cheaper and smaller every year. Statista calculates that 141 million smart-watches will be sold in 2018.

Many fitness trackers act like health trackers as well, being equipped with heart rate sensors and oxygen meters they provide users important real-time insight on their health. Some cases of users detecting early signs of heart arrest or pre-conditions have been reported to the press in the past⁴, proving that even the most superficial IoT devices can save lives.

Other Applications

The Internet of Things have been proved to be flexible and full of potential in almost every field.

Nowadays, even Agriculture and Livestock industries take advantage of the benefits of the IoT, tracking the levels of water in the soil to establish irrigation patterns or locating livestock in large farms while also providing insight on their health.

The IoT creates great opportunities for the development of Environmental, Industrial, Commercial, Health, Urban Planning, Transportation and Home applications. We have yet to discover the boundaries of this exciting technological revolution.

Notes

1. NXP Medical Applications Guide - <http://thesis.norori.com/url/gmior>

2. Fitbit Revenue Report - <http://thesis.norori.com/url/gjuar>

3. Statista - Smart Watches - <http://thesis.norori.com/url/miomn>

4. World Economic Forum: A smartwatch just saved a man from having a heart attack - <http://thesis.norori.com/url/hyjbz>



Smart Speakers

For the purpose of this thesis paper, out of the whole spectre of the Internet of Things I will focus in the ecosystem of Smart-Speakers developed by Amazon and in certain degree, the ones created by Google.

Since smart-speakers are arguably the first contact point most costumers are having with the realm of the IoT, they represent a remarkable research opportunity, to study the way every-day people, non-techies, adapt and behave with the technology, to survey their use-patterns and opinions and to develop design guidelines that help developers create applications centered on the user experience.

This chapter will introduce you to the specific details of the product lines of the Amazon Echo and Google Home, their respective Artificial Intelligences and technologies, giving you a full over-view of their capabilities and limitations.

Certain emphasis will be given to some standard-complying peripherals that can be controlled with either brand, but the focus will stay within the speakers themselves.

Cover Image: Amazon Echo Press Images

Amazon Echo Line

The Amazon Echo was designed and developed by Lab126, Amazon's computer hardware research and development company based in California, United States. The first patents related to the project were registered as early as 2010.



Originally code named “Project Doppler”, it was thought as a way to expand Amazon’s device portfolio beyond the Kindle book readers and tablets. Jeff Bezos, Amazon’s CEO, had the intention to create a speaker with additional functionality based on patents that Amazon had filed in 2010 related to “augmented reality” — and, specifically, voice control.

The development of the first generation of Echo devices took place in Lab126 between 2011 and late 2014, with great difficulty for engineers and speech recognition specialists that encounter several problems training the speech recognition algorithms, the artificial intelligence (AI) and refining technical specifications.

Finding a name for the AI wasn’t an easy task either, Lab126 had to find a word or name that wasn’t widely used in every-day life, but that was easy to pronounce and to be recognized. After surveying diverse options, the team found the word “Alexa” as a suitable candidate, because of its mix of soft vowels and an “x”, making it sound unique. David Limp, the vice president in charge of Amazon Devices, also mentioned that “the name was a little reminiscent of the library of Alexander, which was at one time the keeper of all knowledge. The idea was you could ask Alexa anything, and it would know and answer.”

The decision was nevertheless, not easy, Amazon’s most prominent employee: Jeff Bezos, was favoring the word “Amazon” to be the name of the AI and wake-word of the Echo.

Notes

1. Business Insider - “Amazon engineers had one good reason and one geeky reason for choosing the name Alexa” - <http://thesis.norori.com/url/xcw3j>

2. pymnts: “How Alexa Found Her Voice” - <http://thesis.norori.com/url/ae7ig>

The idea was seen as quite problematic by many team members of Lab126, considering how successful the company is and how often its name is mentioned on the media or by costumers in ordinary conversations. The main concern was that the Echo would hear its name wrongly often and accidentally place online orders on behalf of its owner.

The wake-word and name of the AI assistant was finally decided to be “Alexa”, with the option to change it to “Echo,” “Amazon,” and “Computer” in case that a member of the family shares the same name. If you ask Alexa now why she was named like that, she would answer “I’m named in honor of the library of Alexandria, that stored the knowledge of the ancient world.”

Before the launch of the first-generation Echo, there was a lot of expectation and suspense in Lab126, since the latest project of the company, the Fire Phone, was not commercially successful and represented millions of dollars in losses for Amazon¹.

The first Amazon Echo was released to the public on November 6, 2014 for a retail price of 180 USD. At first, limited to Amazon Prime members or by invitation only, access to the product didn’t open for everyone until June 23, 2015.

The first wave of reactions about the product were negative, many technology media outlets called the speaker “a gimmick that would gain no traction among consumers”. They were wrong, the Amazon Echo proved to be a success. Consumer Intelligence Research Partners (CIRP) reported in 2015 that 3 million Echo were sold in U.S.

The huge success of the product is attributed by CIRP to how “Amazon very carefully but aggressively promoted Amazon Echo during the year. It dedicated Amazon.com homepage real estate to Echo, promoted it on Amazon Prime Day in July, and also advertised it heavily during the holiday shopping season”, said CIRP partner Josh Lowitz³.

Notes

1. *Fortune: Amazon is killing off the Fire Phone* - <http://thisis.norori.com/url/sqa5x>

2. *Consumer Intelligence Research Partners (CIRP) Report - January 2018* - <http://thisis.norori.com/url/8qa8c>

3. *Twice: “Amazon Has Sold 3 Million Echo’s in U.S.”* - <http://thisis.norori.com/url/d3ayo>

Amazon and Lab126 subsequently expanded and updated the Echo Line. In March 2016 the Echo Dot, a smaller and cheaper version of the Echo was released, bringing the smart-speakers to the mass market with a very attractive price-tag, a compact form and the same capabilities as the bigger Echo.

The Echo Dot was an aggressive attempt of Amazon to get a bigger market-share advantage before Google had the chance to get into the game with their own smart-speaker, the Google Home, which was rumored to be released in just a few months. It also worked as a strategy to lure customers to buy more than one Echo per home, one for each room. The strategy was successful, the sales of the devices were up 9x compared to the previous year. “Despite our best efforts and ramped-up production, we still had trouble keeping them in stock,” said Jeff Wilke, Chief Executive Officer of Worldwide Consumer at Amazon after 2016’s holiday season was over.

The next year, the company decided to expand their portfolio even more, with up-dated versions of the classic Echo and Echo Dot and the addition of new members of the family: The Echo Show, Echo Spot, Echo Look and the Echo Tap. Most of the new devices included video-cameras and free-video calls to cell-phones and other Echo devices, in a clear attempt to include video capabilities into the Amazon Smart-Home.

The Echo Show was a far departure from the traditional design of the family, with an 18 centimeters touch-screen the focus now wasn’t only on the spoken interaction. It is equipped with far- range voice recognition and the same Alexa assistant of the other members of the family, but it took advantage of its screen to interact with smart-home video devices like surveillance cameras or baby monitors. It also has an integrates a camera that allows video-calls. The Echo Show got positive review for critics, but its high price has kept it from the top-selling spots.

Notes

1. *Forbes: Amazon Says Echo Was the Best-Selling Product This Holiday - Sells Millions - <http://thesis.norori.com/url/6vfg->*

The Echo Spot is an interesting cross-over between an Echo Show and an Echo Dot: It has a circular 6.2 centimetres screen, a small integrated camera, far-field microphones and a speaker. It can interact with the same with smart-home video devices as the Echo Show while maintaining the small shape-factor and capabilities of the Echo Spot. Despite being advertised as a smart table-clock it's capabilities has been exploited well by the community of developers.

The Echo Look was the most experimental device released on that year, with a clear focus on the niche of video content-creators like "YouTubers" or "social media influencers" that posted new content constantly. Named by Amazon as a "Hands-Free Camera and Style Assistant" it includes cloud-based machine-learning algorithms that let it use "Style Check to get a second opinion on your outfit". As the most experimental device of the line, it's available exclusively by invitation and was not released to the public at the moment of publishing of this paper. The Echo Look features built-in LED lighting, depth-sensing cameras (multiple cameras that can discern the background), far field microphones and a small speaker.

Apart from the new experimental devices, the already established Echo and Echo Dot received well-deserved updates in hardware and design-wise. The second generation of the Echo Dot features improved speakers, an auxiliary audio exit and a better design, its software included the Echo Spatial Perception that enables Multi-Room music and let multiple devices work together when responding or listening to the user. Its price was also slashed in half in comparison to the first generation.

The second generation of the Echo had major design changes, it comes in a much smaller size and with different design styles, three with cloth covers and three with sleek finishes, to make it fit well into every room. It also improved its audio quality with the inclusion of Dolby speaker technology, all this while reducing the price by 50\$/50€.

The last addition to the Echo family was a big bet of Amazon on becoming the centre of the smart-home: The Echo Plus, a device that kept the same dimensions and design of the original Echo but received important hardware updates under the hood. Apart from the Dolby speaker technology that the other devices also received, the Echo Plus includes a built-in “smart-home hub” with ZigBee signal capabilities, allowing it to directly control thousands of “smart” devices without the need of buying an additional hub.


Unsurprisingly, the sales of the Echo devices were an immense success for Amazon, securing the company a higher market-share and an outstanding revenue.

CIRP survey data suggests that the total smart speaker sales in the U.S. reached nearly 45 million at the end of 2017, with 18 Million sold in just the last quarter of the year¹.

The future of the smart-home and speakers is promising. Forrester, an American market research company that “provides advice on existing and potential impact of technology, to its clients and the public” forecasts that by 2022, the smart speaker consumer will have a reach of 172.4 million US consumers, increasing steadily every year.

FORRESTER®

Households with smart home devices (Millions)



F = Forecast

Source: Forrester Data: Smart Home Devices Forecast, 2017 To 2022 (US)

Notes

1. Forbes: Amazon Says Echo Was the Best-Selling Product This Holiday - Sells Millions - <http://thisis.norori.com/url/qmgz6>
2. Forrester Data: Smart Home Devices Forecast, 2017 To 2022 - <http://thisis.norori.com/url/b-919>



Google Home Line

The first generation of Google Home was released in the United States in November 2016, with subsequent releases globally throughout 2017, reaching Germany officially on August 2017. As of February 2018, the Google Home devices are sold in limited markets, you can only buy one if you have an address in Australia, Canada, France, Germany, Japan, United Kingdom or the United States.

The Google Home ecosystem is currently composed by the Google Home, Google Home Mini and Google Home Max.

The original Google Home is a small rounded speaker with an angled top featuring LED-lights and capacitive touch controls to start and stop music, and adjust volume, the lower part houses a speaker and far-range microphones. The lower-half is also interchangeable to personalize it. It was originally priced at 140 USD.

Upon its launch, Google Home received mixed reviews¹, with many critics being disappointed by the lack of functionality the device had and the poor integration with Google's native services like Contacts, Maps, Gmail or Now. The lag behind Amazon's Echo was also a reason for critics to feel let down, but with no much surprise, since the Echo had a longer time out in the market that allowed developers to build much more apps and experiment with the device better.

The first year for the Google Home was not easy, it got smarter and more apps were added to its catalog every day, but it was still far behind the Amazon Echo in terms of sales. Statista estimated in May 2017 that only 22% of the interactions with smart-assistants in the US would be with a Google Home, specific sale figures were not released by the Mountain View gigantic, but nobody expected substantial numbers.

Notes

1. Business Insider: I was so excited, and so let down, by Google Home - <http://thesis.norori.com/url/qczdb>

2. Business Insider: Amazon's Echo isn't going to give up its lead anytime soon - <http://thesis.norori.com/url/vcdx4>

The Home family got its first expansion in late-2017, with the addition of the Google Home Mini, a device directed specifically to compete with the Amazon Echo Dot and the Google Home Max, a massive device with house-filling speakers.

The Google Home Mini is a small speaker with a simple but elegant design, equipped with far-field microphones and Wi-Fi connectivity. Its price-tag was just 50 USD, a much affordable option when compared with its bigger brother.

Google Home Max is a massive speaker: 336.6 mm wide, 190.0 mm tall, 154.4 mm deep and with 5kg of weight. Equipped with Dual 4.5" woofers, premium acoustic transparent fabric and an equally big price-tag: 399 USD.

Both new additions were received with better reviews than the original Google Home, with critics praising their design and sound-quality while complimenting the great leaps the Google Assistant made during the last months, with more compatible apps, smarter answers and a more seemingly experience.

Overall, Google is still behind Amazon in terms of development, experience and sales on the smart-speaker market, but the company is taking good steps to close the breach. As of today, Google has sold 14 million units of Google Home devices¹ and is becoming a major player on the industry, benefiting consumers with a more competitive environment.

Notes

1. *Consumer Intelligence Research Partners (CIRP) Report – January 2018 –*
<http://thesis.norori.com/url/8qa8c>



Virtual Assistants

-Alexa, what are you?

-I'm Alexa, and this is an Amazon Echo. I'm here, but my head is in the cloud. I'm here to help you.

Speech Recognition

Decades ago, talking to computers was a mere ambition of science fiction movies or futuristic books, but since the introduction of the first mass-use modern virtual assistant, Siri, who lived on the smartphone iPhone 4S and was introduced in October 2011.

Siri was the first smartphone assistant with millions of users that could recognize a wide range of words, but the modern history of the technology behind it starts much earlier: in 1961.

The IBM Shoebox¹, a computer with the shape and size of a normal American shoebox equipped with a microphone. It could recognize 16 spoken words and the numbers from zero to nine.

Notes

1. IBM: The Shoebox - <http://thesis.norori.com/url/hr87g>

Cover Image: Designed with sources from FreePic.

When users ordered it to solve basic arithmetic problems it immediately did and printed the results. It worked by converting the audio input into electrical impulses and then classifying the signal using a measuring circuit that matched the input with the saved patterns and triggered the actions. It was not the most advanced machine, but it was “a forerunner of today's voice recognition systems.”

No major advances were done until more than a decade later, when in 1971 the Carnegie Mellon University (CMU) in Pittsburgh, Pennsylvania (US), IBM, and the Stanford Research Institute in close collaboration with the United States Department of Defense and its DARPA agency, started the development of the “Harpy”¹.

In contrast with its predecessors, the Harpy could understand full sentences and recognize 1,011 different words, it had approximately the same vocabulary of an average three-year-old human.

“We don’t want to look things up in dictionaries – so I wanted to build a machine to translate speech, so that when you speak in one language, it would convert what you say into text and then do machine translation to synthesize the text, all in one.

With sentences, you get words flowing into each other, you get a lot of confusion and don’t know where the words end and where they begin. So, you have things like ‘euthanasia’, which could be ‘youth in Asia’,” says Waibel. “Or if you say, ‘Give me a new display’ it could be understood as ‘give me a nudist play”” says Alexander Waibel, a computer science professor at Carnegie Mellon University who worked on the development of the Harpy¹.

Despite the difficulties faced by the researchers and developers, the Harpy became, in the words of Jaime Carbonell, director of the Language Technologies Institute at CMU “*the first system that successfully used a language model to determine which sequences of words made sense together, and thus reduce speech recognition errors.*”¹

Notes

1. BBC - The machines that learned to listen- <http://thesis.norori.com/url/-uc6y>

The next advancement on the field came in the shape of a voice-activated typewriter in the 1981¹, the “Tangora”. The device, developed by IBM, used an IBM Personal Computer and was based on a Hidden Markov Model, a kind of machine-learning technique still widely used today. The computer created statistic models on digital signal processing techniques and “learned” to recognize patterns².

Tangora could recognize words and print them on paper, the system was very sensitive to different voices and thus the researches had to train the system for each speaker. The speakers had to individually train the system to recognize their voices, making short pauses between words. Despite the cumbersome procedure, the team managed to get the Tangora to expand its vocabulary to 20,000 words and demonstrated the effectivity of the Hidden Markov Model and machine-learning methods on the field of speech recognition.

By the 1990’s speech recognition reached a phase where commercial appliances were possible. Dragon Systems, one of IBM’s biggest competitor on the field released DragonDictate 30K in 1990, for the huge sum of 9,000 USD. The software enabled users to control computers using only their voices, which was a very convenient way of interaction for people with limited mobility, it was reported that even the actor Christopher Reeve used the software after becoming quadriplegic³.

Further technological advances made computers faster and cheaper, enabling big improvements on speech recognition. Dragon Systems released Dragon NaturallySpeaking in 1997, becoming a huge sales hit and the center of a considerable media attention. It was featured on big television chains like CNN and BBC while taking home all the COMDEX (the predecessor of CES) computer expo trade show awards of that year³.

Notes

1. *Issues in Practical Large Vocabulary Isolated Word Recognition: The IBM Tangora System* -

<http://thesis.norori.com/url/wl4f4>

2. *IBM 100: Pioneering Speech Recognition*- <http://thesis.norori.com/url/buaer>

3. *MIT: Enter the Dragon* - <http://thesis.norori.com/url/btadv>

IBM counter by hurrying the development of their own software, the IBM ViaVoice, and releasing it to the public just a few months after Dragon NaturallySpeaking.

Speech Recognition reached a plateau in terms of Gartner's hype-cycle in the early 2000s, when the average accuracy was well over 80%, staying in stagnation.

Windows Vista¹ and Mac OS X 10.2² featured accessibility speech recognition tools that allowed users to control the cursor or write long texts using their voice only, but they were not as accurate as the enterprise solutions offered by IBM and Dragon Systems³.

Better speech recognition for the masses wasn't a reality until 2010, when Google integrated "personalized recognition" to Voice Search. Melanie Pinola, editor of PC World believes that this was a turning point on speech recognition: *"The impact of Google's app is significant for two reasons. First, cell phones and other mobile devices are ideal vehicles for speech recognition, as the desire to replace their tiny on-screen keyboards serves as an incentive to develop better, alternative input methods. Second, Google had the ability to offload the processing for its app to its cloud data centers, harnessing all that computing power to perform the large-scale data analysis necessary to make matches between the user's words and the enormous number of human-speech examples it gathered."*

The speech recognition technology made such considerable progress leaps that it wasn't long until it was used on the first modern virtual assistants.

Notes

1. How to configure and use Text-to-Speech in Windows XP and in Windows Vista - <http://thesis.norori.com/url/x4qq1>
2. A Decade of Apple Accessibility - <http://thesis.norori.com/url/5s2wx>
3. Vista's Voice Recognition Stammers - <http://thesis.norori.com/url/dm340>

Other enablers of the modern virtual assistants

Apart from speech-recognition, a number of other technologies contributed to the creation of the right conditions for the flourishing of Google Now, Siri or Amazon Alexa.

As mentioned before, the IT industry is governed by Moore's Law: computing gets faster and cheaper with time. The same principle applies to other resources like data, power and functionality.

The always shrinking size of the smart-phones created the need to move processing outside of the device, while mobile-data got faster with the implementation of 3G and 4G standards: What followed was the booming of cloud computing.

Cloud Computing

Microsoft¹, IBM² and Amazon³ all agree on one definition of what cloud computing means:

“Cloud computing, often referred to as simply “the cloud,” is the delivery of on-demand computing resources — everything from servers, storage, databases, networking, software, analytics and more— over the internet on a pay-for-use basis.”

Such arrangements were necessary to allow the thriving of certain developments in IT, like web-search, remote access to data, online banking and more. Recently, cloud computing started to be used to power resource-draining mobile applications: AI-based photo filters like Prism or FaceApp, music identifiers like Shazam and Soundhound, mobile versions of the Adobe Suite and especially cloud-based storage solutions like Dropbox, OneDrive or Google Drive.

The nature of the cloud fits perfectly well with the needs of the internet of things, where connectivity and processing boards need to share tiny spaces with other important components of the devices they occupy.

Notes

1. Microsoft Azure: What is cloud computing? A beginner's guide - <http://thesis.norori.com/url/0-d17>
2. IBM: What is cloud computing? - <http://thesis.norori.com/url/ynfj>
3. Amazon Web Services: What is Cloud Computing? - <http://thesis.norori.com/url/roykq>

All modern smart-assistants process their request in the cloud, which allow them to keep getting smarter and adding new features day by day without any changes to hardware used by the users.

Locating the AI in the cloud also provide other advantages that surely outweigh the power of any modern compact micro-processor. With the extremely vast amount of complex data amassed in the cloud by the constant use of the virtual assistants by millions of users, data stops being just “data” and it becomes “Big Data”, large data sets that can reveal patterns of use. Big Data becomes the seed for something even more complex and still not fully understood: Machine Learning (ML).

SAS gives one of the most accurate definitions of the term, while providing great examples that are easy to understand: “Machine learning uses sophisticated algorithms to “learn” from massive volumes of Big Data. The more data the algorithms can access, the more they can learn. Real-world machine learning examples are everywhere. Think of personalized product recommendations on Amazon, facial recognition on Facebook, or fastest route suggestions in Google Maps”. While the term might be a little misleading and make us believe that computers are truly becoming self-aware, they’re just only learning to recognize patterns and predict when they’re most likely to be repeated while acting accordingly.

Smart-assistants learn to understand the way different users pronounce the same words and improve their accuracy, or to understand the different ways one may ask for the same thing. While some people may say “what’s the weather?” others would ask “Is it rainy today?”, while the question is different, both users just want to know the weather. Neither Amazon or Google have revealed what amount of the total “knowledge” of their assistants is mined from machine learning, but one might suspect is not that little.

We can expect more exciting developments on the big data and machine learning fields, as for now, they’re still in a very early stage of development, but they will surely become important technologies in the next decades.

Notes

1. SAS: What Is Big Data? - <http://thesis.norori.com/url/mxy18>
2. SAP: What is Machine Learning? <http://thesis.norori.com/url/uczuo>

Wireless Technologies

Hardware-wise, smart-speakers contain no extraordinary technology, despite this some of their capabilities are quite inherent to the internet of things niche, especially the ones regardless connectivity.

It might sound surprising, but Wi-Fi is not the default connectivity standard of the IoT industry. Although all the smart-hubs that we use to control IoT-enabled devices have Wi-Fi, most of the time they don't "talk" with the other IoT devices using Wi-Fi frequencies, they do so by communicating using ZigBee, Z-Wave or Bluetooth Low-Energy.

ZigBee

ZigBee is a wireless technology developed and ruled by the ZigBee Alliance, a global alliance of companies that got together to create a wireless standard.

ZigBee uses the 2.4 GHz radio frequency with IEEE 802.15.4-based specification to communicate with receptors located up to 100 meters away from the emitter. It's important to note, that ZigBee is also capable of transmitting information using a "mesh network", where receptors repeat the signal until it reaches a destination, increasing the original range of 100 meters even further. A ZigBee module can be as small as 10mmx10mm, they are low-cost and low-power¹.

The standard was a pretty unknown standard until Amazon started using it on the Amazon Echo Plus, their first smart-speaker with a smart-home included. Amazon didn't include ZigBee's competitor standard: The Z-Wave, and considering how significant impact Amazon is having on the IoT industry, it was a definitive push for ZigBee and a great loss for Z-Wave.

Z-Wave

Z-Wave is a proprietary standard developed in 2001 by Zensys, a Danish company. In 2008 the company was bought by the US-based Sigma Designs, who is now the owner of the standard².

Notes

1. ZigBee Specification FAQ <http://thesis.norori.com/url/odv2r>

2. Z-Wave: About <http://thesis.norori.com/url/n35tp>

3. Electronics360: ZigBee vs Z-Wave for the IoT- <http://thesis.norori.com/url/enppa>

The standard has the same low-energy features and mesh network capabilities of ZigBee.

Being a proprietary design has restrained Z-Wave to reaching a bigger market share, as of Q1 2017 there were only 1,700 compatible devices. It would not come as a surprise if the standard ends dying up soon.

Bluetooth Low Energy

Also known as Bluetooth 4.0, Bluetooth Low Energy (BLE) is an evolution of the same old Bluetooth we all know, with the key difference of BLE being its low power consumption¹.

BLE is a popular standard in the wearables industry, being the default method of connectivity for all kinds of wearable devices, like smart-watches, fitness trackers or even smart-glasses like Snapchat's Spectacles².

Despite being a well-known standard in other industries, BLE is being largely ignored for IoT appliances because of its shorter range and lack of mesh network capabilities, two vital needs.

All smart-speakers still integrate BLE or even the newest standard, Bluetooth 5.0, but it's mainly used to connect with smartphones for music streaming, one of the most common uses of Bluetooth.

Notes

1. Bluetooth Low Energy Specifications - <http://thesis.norori.com/url/se89h>

2. Snapchat - Spectacle Features <http://thesis.norori.com/url/t6b3d>

Modern Virtual Assistants

Cloud-based processing, lower cost of data and the ubiquity of the smart-phones created the perfect conditions for the first Virtual-Assistants, companies like Apple and Google were the pioneers on this field, and were its major contributors for a long time.

Apple introduced “Siri” on October 2011 as one of the major updates of its new phone, the iPhone 4S. Their key-note presentation highlighted how easy it was to complete simple tasks like asking about the weather, set-up alarms, directions or check the status of the stock markets by just using your voice¹.

Siri used cloud-based speech-recognition technology provided by Nuance Communications, a company that bought Dragon Systems, the leaders of the voice-recognition market back in the 1990s². The voice of Siri was based on the voice of Susan Bennett, an American voice-over artist born in the state of New York³.

Critics received Siri with excitement⁴, viewing the voice-interaction as the next big revolution of smart-phones, just after the touch-screens, but many were concerned that Siri was confined to the iPhone OS only with no way to integrate it with third-party applications. This was the case until 2017, with the release of iOS 11, that allowed third-party apps make use of the smart-assistant, this integration made possible to call taxis, send messages over WhatsApp and compose tweets with your voice only.

After the initial release of Siri, it was not long until other competitors joined the Virtual Assistants race, Google launched Google Now with the update of Android 4.1, "Jelly Bean". It lived inside the Google Search screen and took skillful use of the ecosystem of Google Apps: It could analyze e-mails, text messages, flight confirmations, calendars and notes and organized them in a chronological screen. It also learned to be ahead of its user: For example, with the help of the GPS it would learn where the users spent their nights and where they were during working-hours to define the location of your home and of your office and provide you with suggestions of better traffic routes.

Notes

1. Apple iPhone 4S Keynote - <http://thesis.norori.com/url/y5w9t>

2. Nuance confirms its voice technology is behind Apple's Siri - <http://thesis.norori.com/url/gzo6z>

3. Susan Bennett -The official website of Susan Bennett, the voice of Siri - <http://thesis.norori.com/url/omwuf>

4. Review: With Siri, iPhone finds its voice - CNN <http://thesis.norori.com/url/ydjw3>

Google Now was more focused on providing valuable information when you needed it and less interested on the voice-interaction, like Apple did. Google Now and Apple Siri evolved year by year, with the help of the substantial base of users they could improve their accuracy considerably while adding additional features continuously.

Despite this, the popularity of phone-based virtual assistants decreased sharply, in 2012 a 62% of iPhone 4S users used Siri at least several times a week, including 35% of respondents who reported daily use¹. By 2016, a report by Verto Analytics revealed that “between May 2016 and May 2017 Siri’s active user base declined with 7.3 million monthly users (nearly 15 percent of its total) – at the same time Android mobile devices have gained popularity.”

Following the data release from Verto Analytics, the Fortune magazine reported that while the use of smart-phone based virtual assistants was in decline, the use of voice controlled IoT-based devices was soaring³.

The smartphone is not a natural environment for voice-centred interaction, but with the dawn of the home-based internet of things devices, the machines were finally finding the place where their voices belonged.

Notes

1. Park Associates: *The Rest of the Siri Story* <http://thesis.norori.com/url/nouif>

2. Verto - *New Research Report: Sluggish Adoption of AI-Driven Personal Apps*
<http://thesis.norori.com/url/nkav6>

3. Fortune: *Apple's Siri Virtual Assistant Has Lost Millions of Users*
<http://thesis.norori.com/url/podo7>

Amazon Alexa

“Smart speaker uptake has grown faster than any other consumer technology we’ve recently encountered, such as AR, VR or even wearables. While 2017 has been a banner year for smart-speakers in terms of hardware sales, especially for Google and Amazon, smart-speakers in 2018 will move beyond hardware, with strategic attempts to monetize the growing installed base in the US and beyond. The possibilities to do this are endless, be it discreet advertising, content subscription bundles, premium services or enterprise solutions. The technology is still in transition, and increased investments from multiple players of the ecosystem will fuel growth.” – Lucio Chen, Research Analyst at Canalys.

Smart-speakers are here to stay, homes are getting “smarter” all around the world day by day, enabling endless opportunities for developers to create useful applications for all kinds of home-appliances.

Alexa, the artificial intelligence (AI) behind the Amazon Echo line is one of the most famous virtual voices nowadays, powered by robust speech-recognition technology and complex cloud-based processes, it provides users with unrivalled voice-interaction and is capable of controlling music, making lists, playing audiobooks, giving weather and traffic forecasts and reading the news. Thanks to its internet-of-things enabled technology is also able to control all kinds of home-devices, like lights, thermostats, fans, TVs, fridges, blinds and pretty much everything with Wi-Fi or ZigBee connectivity.

Alexa is currently limited to two languages: English (with sub-variants for the UK, Ireland, Canada, the US, Australia and India), German and Japanese.

Developers can create applications for Alexa, the apps are called “skills” by Amazon and are distributed in a market-place inside the smart-phone based Alexa App.

Notes

1. Voicebot: 56 Million Smart Speaker Sales in 2018 Says Canalys <http://thesis.norori.com/url/tamco>

2- Quartz: Amazon says there are now 5,000 people working on Alexa — <http://thesis.norori.com/url/gywqm>

Google Assistant

Google Assistant is the evolution of Google Now, the interface that lived inside the Google Search mobile-app and provided predictive cards about upcoming events or tasks, speech recognition for search inquiries and daily updates.

Google Assistant was launched to the public for first time in May 2016, as part of Google's messaging app "Allo" and consequently on the Google Home and the Pixel and Pixel XL smart-phones. It was not possible to install it on other smart-phones until February 2017, when a stand-alone app was released on the Google Play Store¹, an iOS version was released in May of the same year.

Google Assistant is capable of engaging in two-ways conversations and completing tasks standard for a smart-assistant of its kind, like setting up alarms, reminders and timers, checking the weather, news or the results of sport matches and doing web-search, it's also able to control IoT-enabled devices when connected to a smart-hub.

Although Google kept the assistant locked into their own hardware for a time, it changed its approach sharply, in the CES 2018 Google announced that it was integrating the Google Assistant in devices from other hardware-manufactures, in a clear attempt to stop the domination of the smart-speaker market that Amazon has. The list of manufactures includes: Anker, JBL, LG, Lenovo, Sony and many others. Google also announced that Assistant was soon going to jump into smart-screens, but it gave not much more details on that.

Google is one of the most competitive players in the fields of smart-assistants, speech recognition and smart-speakers, driving the development of the industry faster and providing consumers with high-quality hard-ware and soft-ware.

Notes

1. Google Assistant - Android Apps on Google Play - <http://thesis.norori.com/url/jxj2e>

2. The Verge: Google is introducing a new Smart Display platform - <http://thesis.norori.com/url/-gu6b>

Siri and HomePod

Apple is not what it used to be, what was once one of the biggest innovators in the IT industry has lagged behind the competition in the last years. Failing to foresee upcoming trends and unable to keep up with the technological advances of its competition Apple is now in a difficult position in terms of innovation. Even though their revenue is still increasing, especially because of their huge sales of iPhone smart-phones, they haven't produce a ground-breaking device since the iPad in 2010.

They surely joined to the bandwagon of major trends like wearables, with the Apple Watch series, music and video streaming with Apple Music and large phones with the iPhone X. All of them adopted quite late when compared to competitors. The smart-speakers trend was no exception the way Apple does things now.

Apple is not new with smart-assistants, Siri was a pioneer and one of a kind when it was released in 2011, so it comes as a great surprise that Apple took so long to release a smart-speaker of its own.

The HomePod, Apple's response to the Amazon Echo and Google Home was announced on June 2017¹. Its release was delayed several times until it finally went on sale in February 2018².

Critics were not impressed³: The HomePod is one of the most expensive smart-speakers available in the market, priced at 349.99 USD \$, it delivers high-quality sound but its smart-assistant lags behind the competition by far. At the time of its release, music could only be streamed using an Apple Music subscription, that goes for 10 USD \$ per month. It also lacks Bluetooth streaming or audio inputs.

Further failures plagued the release of the HomePod, with the so called "Ring Gate" being the most serious: Oils from the bottom of the speaker leaked and damaged furniture leaving circular-shaped stains⁴.

The HomePod seems more like a prototype than a fully finished premium device.

Notes

1. Apple: HomePod reinvents music in the home - <http://thesis.norori.com/url/q6l4j>

2. Apple: HomePod arrives February 9, available to order this Friday - <http://thesis.norori.com/url/5pwla>

3. Engadget - Apple HomePod review: A great speaker that's not so smart <http://thesis.norori.com/url/sqjle>

4. Metro: Apple HomePod is leaving a nasty stain on owners' tables - <http://thesis.norori.com/url/yhk7o>



Adoption & Usage of Smart Assistants

For a technology to truly take-off in the consumer market a series of conditions and events must take place, one of them is widespread adoption.

In the past, we have seen many new products that seemed to create an irreversible revolution: In late 2010, the first iPad caused excitement for tablets, soon enough dozens of manufactures were flooding the market with their own versions, from 2010 to 2013, 660 million tablets were sold worldwide and by 2015 one billion devices were sold¹. For a moment, it appeared that computers were a thing of the past and from now on we would all use our tablets to work. It wasn't the case, as the buzz cooled down and the sales got slower².

It can be said that the debut of the tablets was an exaggerated event, but it's undeniable that they did make an important impact, creating a considerable market and influencing other devices in many ways, the size of the smart-phones got considerably larger on the years following the appearance of the first tablets, and the same can be said about the increasing number of laptops with touch-screens, an inherent capability of tablets.


Notes

1. Statista: Number of tablet users worldwide from 2013 to 2020 - <http://thesis.norori.com/url/ilvg>

2. Time: Tablet Sales Are Down: 5 Reasons Why <http://thesis.norori.com/url/t6m6s>

Cover Image: Amazon Echo Press Images

It's difficult to know if a certain trend -whether it is tablets, wearables or smart-speakers- is going to make a lasting impact based only on its initial hype. All successful new technologies follow the cycle of the *Gartner Hype Cycle for Emerging Technologies*, as it is a fundamental trait of prosperous new advances to get an incredible amount of attention from the press and public followed by a phase of disillusionment.



Gartner Hype Cycle as of August 2017

The Gartner Hype Cycle for Emerging Technologies is a graphic tool developed by Gartner and used to represent the maturity, adoption and commercial and social application of a particular technology according to its life cycle and the way in which it is potentially relevant for solving real business problems and to exploit new opportunities. This methodology gives a vision of how a technology or application can evolve over time, thus providing an overview of its maturity, helping to manage its implementation in the specific context of each business and to take better advantage of the opportunities that come from it. As seen on the graphic above, it is divided into 5 phases of adoption and exposure over time.

Notes

1. Gartner: Top Trends in the Gartner Hype Cycle for Emerging Technologies, 2017 - <http://thesis.norori.com/url/vsily>

As of late 2017, Gartner situated Connected Home, Virtual Assistants and the IoT platform all around the summit of the “Peak of Inflated expectations”, this comes as no surprise, considering how consumers and press have been highly aware of all of these technologies and sales of products related to them have been skyrocketing. The three technologies are similar and are intricately connected to each other, but are considered separate entities by the firm.

Not all technologies around the Peak of Inflated expectations meet commercial success, older technologies like Virtual Reality¹ and Augmented Reality², which are already around the Slope of Enlightenment, just sold around 1 million devices world-wide until today, a much lower number compared to the 45 million³ smart-speakers sold in the US alone.

Sales are not everything, but they surely help a certain new technology to set foot inside the market and to accelerate its development by increasing the interest and investment of manufacturers, the attention and coverage of the media and the interest and support from the consumers.

The Connected Home and Virtual Assistants are expected to reach the plateau of productivity within 5 to 10 years. A technology reaches Gartner’s “plateau of productivity” when its benefits are widely demonstrated, developed and accepted. In this phase, the criteria for determining commercial viability begin to be clarified and the adoption of the technology is either massive or stable. It’s also when they start to be truly profitable.

While the Connected Home and Virtual Assistants might be profitable for their manufacturers, there is still no concrete idea of how to generate revenue for the third-party developers that create applications and maintain the sense of novelty on the devices. Possibly the reason why Gartner considers they won’t reach the plateau in a medium-term.

On the other hand, the IoT platform is situated in a much closer position from the plateau of productivity, a certainly surprising fact considering how volatile and ambiguous the IoT is.

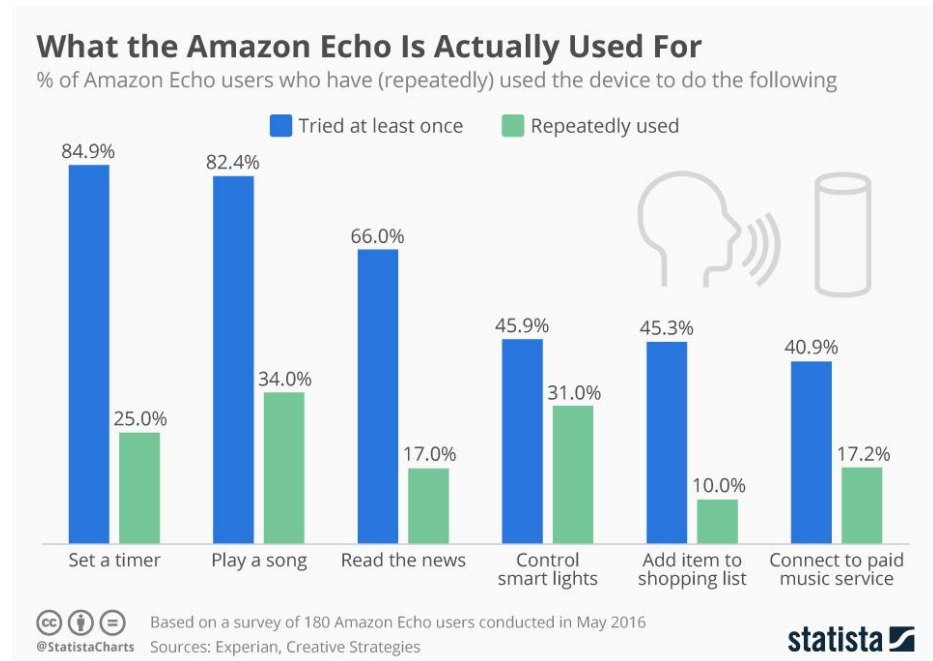
Notes

1. *Forbes: Virtual Reality Headset Sales Hit 1 Million* - <http://thesis.norori.com/url/drdsz>
2. *Statista: Global virtual & augmented reality headset sales revenue 2016-2020* - *Statistic* <http://thesis.norori.com/url/7yfqh>
3. *Consumer Intelligence Research Partners (CIRP) Report – January 2018* – <http://thesis.norori.com/url/8qa8c>

Consumer Usage

With 45 million smart-speakers in the US alone¹, one might wonder: What are consumers doing with their smart-speakers, virtual assistants and connected IoT home devices?

Statista conducted a survey with early-adopters of the first generation of the Amazon Echo and asked which of the basic functions of the device they used the most.



Statista Survey on Amazon Echo Usage²

Setting timers and playing songs led the list, albeit other functions like reading news, shopping and connecting to paid music services were used at least once by many of the respondents, when asked if they used the same function repeatedly their usage slashed in half.

It's reasonable that many early-adopters try all the functions of a device at the beginning, but forget about them once the novelty of the device is gone, a behavior often observed with other devices like smart-phones, TVs and computers.


Notes

1. Consumer Intelligence Research Partners (CIRP) Report – January 2018 – <http://thesis.norori.com/url/8qa8c>
2. Statista: What the Amazon Echo Is Actually Used For - <http://thesis.norori.com/url/ixht9>

This pattern of use behavior is currently also observed with the third-party applications developed for the Amazon Echo ecosystem, which is unquestionably troubling for the future development of the smart-speakers.

While the offer of apps, called “skills” by Amazon is increasing steadily in all markets covered by the company, up to 70 percent of the available Alexa “Skills” have no or few customer reviews, indicating a low usage.

Total Number of Alexa Skills as of January 2018 - US, UK & Germany



VoiceBot report on total number of Alexa Skills as of early 2018¹

The offer of Amazon Skills in the Alexa Store reached over 25,000 apps in the US in early 2018, while the UK and Germany stayed behind with only 9,000 and 3,100 apps respectively. Many big brands have been joining the voice-apps craze, with Alexa Skills providing content from media outlets like ESPN, NBC, Spiegel Online and Reuters, to mention a few, and skills based on television or movies, like Dunkirk, Deadpool or Spider-Man.

However, according to VoiceLabs and Alpine AI², when users do try a new app, there is only a 3% chance they will be using it for more than two weeks. A truly bleak outlook, especially when compared to the retention rate of Android and iOS apps, that can go over 6% for many categories of applications³.

Notes

- 1. Voicebot: Amazon Closes Year With 266% Alexa Skill Growth in U.S. - <http://thesis.norori.com/url/n2z2o>
- 2. VoiceLabs and AlpineAI: The 2017 Voice Report - <http://thesis.norori.com/url/c8q6p>
- 3. Statista: Mobile app user retention & churn 2017 - <http://thesis.norori.com/url/7hwt9>

The low usage and retention rates of the Alexa Skills is not the only or even the biggest challenge the voice-centred ecosystem is facing in the long-term. The lack of a monetization platform for developers of Alexa Skills is a clear impending threat that will get acute as soon as the smart-speakers move past Gartner's peak of expectations into the trough of disillusionment.

Both Amazon and Google don't provide a platform that helps all developers to monetize their creations. They do reward the apps that get more engagement, but few apply to these benefits¹.

By early 2017, more than a year after the introduction of the first generation of Echo devices, Alpine AI reported² that not a single Alexa Skill was monetized successfully.

Amazon started to offer new ways to monetize Amazon Skills in November 2017³, enabling in-skill purchases and paid skills, but it also introduced stricter guidelines for advertising inside the skills, killing the monetization model of Sponsored Messages used by major developers⁴.

The newly introduced monetization models were immediately adopted by several developers, skills like "Jeopardy!" started to offer add-ons subscriptions to "Double Jeopardy!", and games like Heads Up and History Trivial offered additional content packages for a fee. Amazon also announced that in 2018, charities would be able to collect donations by voice-orders.

The pay-in skill feature was used in TGI Friday's app to allow customers to place delivery orders and pay using their Amazon account and Atom Tickets allowed users to buy movie theater tickets in the same way. Both applications were strongly restricted to certain regions of the US.

All the new monetization models were mostly harnessed by large companies or agencies, since a huge promotion effort was necessary to start reporting any revenue, leaving indie developers out of the game.

Notes

1. Amazon: Earn Money for Developing Alexa Skills that Customers Love <http://thesis.norori.com/url/3sjqk>

2. VoiceLabs and AlpineAI: The 2017 Voice Report - <http://thesis.norori.com/url/c8q6p>

3. Amazon: Amazon Announces New Ways to Earn Money with Your Alexa Skills <http://thesis.norori.com/url/5-5rp>

4. VoiceLabs is 'hibernating' its Sponsored Messages Amazon Alexa Advertising Network - <http://thesis.norori.com/url/x5ids>

The Alexa Developer Rewards program is still the best chance for indie developers to collect revenue from their contributions to the Alexa Skills store.

"Every skill makes Alexa smarter or more useful. We can't do that by ourselves and we want to enable indie developers to innovate and extend Alexa capabilities at a rapid pace. If our developer community succeeds, we succeed." said Rob Pulciani, Director of Amazon Alexa, in a statement to CNET¹.

Indie developers that join the Alexa Developer Rewards program can expect a difficult outlook, with little chances of monetary success due to the now over-crowded offer and lack of proper ways to promote skills to costumers.

Overwhelming success cases have occurred¹, Joel Wilson released two simple trivia apps called "Question of the Day" and "Three Questions" in February 2017 "just for the fun of it", a few months later the popularity of his apps sky-rocketed and Amazon was sending him monthly checks that ranged from 2,000 USD to 9,000 USD.

He joined the developers' community at an early stage, when for his own luck few trivia apps were available, now this market is extremely over-crowded but it's still one of the most dynamic categories of Alexa Skills.

Nowadays, the pay-outs have been dramatically reduced and developers can rarely expect more rewards than free server storage and processing from the Amazon platform.

In the same CNET article published in late 2017¹, Joseph Jaquinta, an American, IBM senior developer gave insightful feedback on the current situation for developers of the smart-speakers ecosystem: He developed a dozen Alexa-skills, mostly games intended to generate revenue, and while he got over 5,000 USD the first months, this was reduced with time. "For there to be quality skills on Alexa, quality developers are going to need to spend quality time. If you can't make money off it, no one's going to seriously engage." said Jaquinta.

Notes

1. CNET: What Amazon's Alexa economy pays the people building its skills <http://thesis.norori.com/url/py3m>

The current overview of the adoption of smart-speakers is good, with an ongoing steadily growing user-base and predictions of this to continue for a while, the compromise of major manufactures to develop new and better devices and the interest of the public, the smart-speakers market will continue to develop at fast phase.

Although massive monetization challenges plague the voice-apps ecosystem, few other obstacles seem to get on the way of their expansion. As Rob Pulciani, Director of Amazon Alexa said¹: “If our developer community succeeds, we succeed.”

Consumer market adoption of smart-speakers will continue to flourish, as more people are getting comfortable interacting with voice-first devices and developers get rewarded more thoroughly, one can expect a few exciting next years for the market of the smart-speakers.

Costumers are getting more comfortable interacting with voice-first devices, and while some discomfort over possible privacy breaches exist and some people find odd requesting things to machines by voice.

Notes

1. CNET: What Amazon's Alexa economy pays the people building its skills <http://thesis.norori.com/url/pyi3m>



Primary research on smart-speakers

The position smart-speakers will take in our homes is still unclear, their current role is ambiguous for many users, manufacturers and developers. We still don't know precisely what will be their impact or how users will accommodate to their presence in the long term.

Voice-first devices like the smart-speakers Amazon Alexa and the Google Home are still a big novelty in many households, where users tend to play with their devices more due to its freshness, but what would happen once this wears out is what is more intriguing, will users forget about their existence and save them in the cellar to collect dust, or will they interact with them less, but in a more natural and practical way? The latter scenario is where all the bets are placed.

As, mentioned in the book *The Silent Intelligence*¹, “the most profound technologies are those that disappear from the view. They weave themselves into the fabric of everyday life until they are indistinguishable from it.” Will the smart-speakers integrate completely in the conventional routine of the millions of households they are inhabiting? It's perhaps still too soon to have a concrete answer.

Notes

1. Daniel Obodovski, Daniel Kellmerit. *The Silent Intelligence: The Internet of Things*
Cover Image: Designed with sources from FreePic

Since smart-speakers are still in a quite early phase of introduction, not much information about the way users interact with the devices or best practices that developers can follow when programming voice-first applications exist. As a result of this situation, primary research becomes the best tool to get an understanding of how smart-speakers are being used.


In this chapter, the findings and methodology of a primary research on user-interaction with voice-first devices will be exposed.

Methodology

For the purpose of the primary research, I set up a website¹ that served as a central point for all the phases of research: Collection of surveys and questionnaires, registration for trials with the studied devices and the consequential feedback acquisition, publication of some insights and the release of applications developed based on the learnings of the study.

The website was promoted via social media, among groups of expats in Germany, English-speaking Germans, friends, colleagues and family.

People who owned or have interacted with smart-speakers of any kind had the chance to fill-up opinion and usage surveys, while people interested on having a first-hand smart-speaker experience were able to register for trials.



Website for the Internet of Things Research Thesis

A total of 39 opinion and usage surveys were recorded, while the trials accounted to a total of 7, 5 of which were of short duration (1 week) and two medium-term trials (28 days).

Notes

1. Ronald Norori's Master's Thesis Portal - <http://thesis.norori.com/>

Opinion and Usage Survey

A survey about the opinion and usage patterns of smart-speakers was distributed among groups of expats in Hamburg via social-media and to employees of start-ups and technology companies.

While the targeted samples were in general more tech-savvy than the average European, many people I approached had not interacted with smart-speakers in the past, reducing the size of the sample considerably. The final sample size was of 39 participants.

The Survey was titled “Opinion and usage patterns of smart-speakers and smart-assistants”, the questions centered on defining which brands of speakers were used most frequently and what were the patterns of use of the users.

Distribution

The only prerequisite to enter the survey was interacting with smart-speakers prior to filling up the questionnaire, all the participants owned a smart-speaker of their own. The most popular line of smart-speakers among the participants was, unsurprisingly, the Amazon Echo Line, with over 85% reporting using one in the past, the Amazon Echo Dot 2nd Generation was the most popular device over-all, with a share of 74%.

The Google Home line positioned in second place, with 69% of the participants using one of their devices. The traditional Google Home was the most used device of the line, with 64%.

No participants interacted with an Apple HomePod at the time of the closing of the survey, this was an anticipated result since the HomePod was on the US-market for just over two weeks then and is not being released in Europe until the late spring of 2018.

Despite Apple being a non-factor in the smart-speaker territory, its smart-assistant, Siri, positioned first of its category

with an 87% reported usage, this can be attributed to its longer life-time and larger distribution, Siri has been around since 2011 and included in each iPhone release since then.

Amazon Alexa was second with 74% and Google Assistant third with over 66%, the lead of Alexa demonstrates that many Android-users are not aware or not interested in using Google Assistant with, which is partially built-into the Google Search app and available as a separated smartphone-app.

In contrast, participants were generally satisfied after interacting with smart-assistants using a smart-speaker, two-thirds of the sample said that they had a “positive or very positive” first experience, while just 23% reported not being impressed or underwhelmed with their first attempts using the devices.

Usage

The survey revealed interesting patterns of usage. While naturally, the most popular reported activity was playing music over Wi-Fi or Bluetooth with 100% reporting doing so. Built-in features like weather forecasts (97%), timers (89%) and reminders (87%) also positioned on the higher tier.

Most users seemed satisfied with the basic functions of the smart-speakers, which leaves third-party app developers with a much-smaller user-share.

Just 18% said they were using third-party apps frequently 23% sometimes, and outstanding 21% reported never using third-party applications at all, a majority of 38% rarely used one. For a growing ecosystem of over 25,000 apps in the case of Alexa, this pattern of use is discouraging and represents a major challenge that the manufactures need to over-come in order to maintain a prosperous developer community and new interesting applications coming to their devices.

43% stated using their smart-speakers to control other smart-home devices like lights, thermostats and plugs. Although there is no comparison data, one can foresee this number increasing during the next months, mostly driven by the newer editions of smart-speakers, like the Echo Plus, now including integrated smart-hubs, slashing the upfront costs while at the same time the price of smart-home devices is decreasing steadily.

Out of the owners of smart-home devices, most (47%) had a single smart-home device of any kind in their residence, others reported more with two at 35%, three with 17%, and four with just 6%. All the respondents reported having at least one smart-lamp installed, with thermostats being the only other type of smart-home device reported.

Drawbacks

While satisfaction levels were high, negative perception or experiences with smart-speakers were also expressed. Worries about privacy issues and other security vulnerabilities led the list, with about half of the participants expressing discomfort over the security and privacy factors.

The quality of the sound was another reported downside with 38%, it's important to point out that most of the participants had just the lower-end version of each line, with either the Echo Dot or Google Home Mini, both speakers known of small size and little sound-power.

Limited features were also reported as downsides of the experience, with about 28% being not completely satisfied with the spectre of features available at that moment.

The design and price of the smart-speakers and accessories like smart-home enabled devices were listed as other small downside, with around one-quarter of the survey not satisfied with the current price tags.

Outlook

Despite noticeable drawbacks, the participants expect good a prospect for the industry of smart-speakers, with more than half saying they think smart-speakers will be a product one will "definitely or probably need", this opinion was most common among the younger the participants.

Trials

A series of trials of short and medium-term usage of smart-speakers were conducted. Participants were called to join the trials via the thesis's portal and approached in person as well.

Participants of the trial had the chance to use the Amazon Echo Dot 2nd Generation for periods of 7 or 28-days, followed by a simple survey and quick interview to learn from their experience and analyse their point of view.

A total of 7 participants took part on the trial, 5 in the short-term trial of 7-days and 2 in the medium-term trial of 28-days.

Trials Survey

Almost half of the participants were younger than 24, two between 25 and 34, and one each between 35 and 44 and 55 and 64. Most participants were living in shared-flats known as "WG", while others lived with their families, none of them lived alone.

Most used a smart-speaker for first time during the trial, just 2 participants had use a similar device (in this case, the Google Home) in the past.

In contrast, usage of smart-assistants was almost total, with 6 out of the 7 participants having experience with virtual assistants, Siri and Google Assistants were the most popularly reported, respectively.

Usage

Almost all participants reported having a positive first impression of the device, with just one saying he had "a neutral opinion".

The usage of the device was comprehensively high, with all trial participants talking to the device several times a day. A quarter used the device "more than 5 times a day", about half did it "up to 5 times a day" while the remaining 29% used the device "1 to 3 times a day".

All played music, set-up timers and asked for the weather during their trial, almost used reminders as well. The interest of the participant on other specific functionalities was moderate, with 70% trying-out other functionalities like conversation apps, news readers, games and other utilities.

Drawbacks

Similar drawbacks to the ones reported in the Opinion Survey were raised between participants of the trial. Privacy concerns were again cited as the top downside. Followed by the lack of interesting features outside the basic offer. Poor sound quality was also mentioned, but this can be overlooked since the test-device was the smallest smart-speaker available, with a very limited sound range.

Outlook

All participants were excited after the trial, with 2 saying that they think of smart-speakers as a “must have” and half saying they will “probably get one within the next years”, just one participant had no intention of getting one of their own after the trial.


Trial Participants Assessment

Direct feedback was also collected in a series of conversations and interviews with the participants after completing the trial. Several valuable findings were identified.

“It was useful for many daily tasks” – Long Trial #1

Older couple – Household of 2 people – device located in kitchen

During this trial, the most-used features were the timers and music reproduction. A detailed breakdown of each day’s activity can be seen below.



Almost all the used applications are native functions of the Amazon Echo. Third-party apps were used just a few times over the whole trial, signalling that the participants got comfortable with the basic functions and had little interest in exploring its capabilities.

The device was placed in the kitchen, which is one of the most common locations for smart-speakers, a survey by Experian revealed that 51% of owners of an Amazon Echo have placed the device in their kitchens. Several basic functions of smart-speakers, like timers, conversions and basic search are very handy when cooking.

Notes

1. Recode: 51 percent of Amazon Echo owners have it in the kitchen - <http://thesis.norori.com/url/8up36>

The participants of this trial expressed their satisfaction with how easy it was to communicate with the device, even though they were not native speakers of English. Difficulties only raised when asking to play music from specific artist with names hard to pronounce.

Timers were used only when cooking or baking. Weather forecasts were requested only when rain was possible, but not as a daily habit.

The lack of usage of third-party apps demonstrates that users want simple and straightforward apps that require little interaction but offer great practicality.

The participants of this trial had a quite casual approach to the device, which is a useful perspective that represents a great portion of the user base smart-speakers now have: people not used to adopt new technologies right away but that were interested on it because of its hype, smooth and practical use and low price.

Privacy concerns were also in the mind of the participants of this trial, the “mute” function of the speaker was used with regularity when guests were over or private conversations were held. This behaviour continued even after it was explained to them that the device will not record any audio unless it hears the wake word “Alexa”.

Amazon, Google and the developers of apps for their smart-speaker ecosystem need to be aware that the long-term usage of their devices will be more casual and less engaging than the one expected in other platforms, apps need to be more passive and universal user interaction guidelines need to be established and strictly enforced.


Gaining the trust of their user-base is another essential priority, all providers of smart-speakers have so far secured their systems well enough and no data-breaches have been reported so far, still, more effort communicating the lack of risks needs to be done.

“It’s fun to use” – Long Trial #2

Shared flat of 3 young people – device located in living room

The second long-trial took place in a shared flat, the household was composed of young people under the age of 25.

The larger number of people using the device and the fact it was located in the living-room, where they spent much time contributed to a high interaction-rate.



The usage peaked during the first days, when many functions were tested for fun and the participants got used to interacting with the device and its commands.

After a few weeks of use, the most-used functions were music playing, weather forecasts and simple web-searches. Less frequently used were games, timers and new apps.

Participants saw the device as a good media hub thanks to its voice-controlled music experience, with which they could play songs from different sources: radio channels, Bluetooth and streaming services like Spotify and Amazon Music.

Utilities like weather forecasts and timers were also found handy. Participants expressed their desire for better reminders, with more specific functions that helped them keep track of their shared-tasks, like watering the plants, doing the dishes or cleaning.

No concerns about privacy and security were hinted, but other aspects of the device were cited as in need of being improved, most notably the quality of the audio. This, again, is due to the use of the smallest version of the Echo line, the Echo Dot 2nd Generation, that has decent but not great speakers. The work-around played by this group was connecting the device via auxiliary audio cable to a bigger and better speaker, the Echo device was still used as the “brain” and “ears” of the system, but the audio came out through external speakers.

Third-party applications were used in seldom, the users were more than happy using the basic capabilities of the device. When third-party apps were used, it was for trivia games and news-flash from local news providers.

The group praised the usability of the device and the price of the entry-level version, after the trial they wanted to get one smart-speaker of their own, but agreed on saving for an upper-level version like the Echo Plus that features Dolby Audio.

Providers of smart-speakers have the terrific opportunity to capitalize on the market of young-adults living in shared accommodation, so far Amazon and Google have been promoting their smart-speakers as a family oriented device, but should consider including other groups of people into their marketing strategy.

The challenges for their app ecosystem remained visible during this test, third-party apps are not being used by the users in the long term.

Conclusions of the Primary Research

Despite the small sample size of this primary research, its findings are aligned with the results of other studies that employed larger samples.

Third-party applications

The global management consulting firm Accenture, after surveying 21,000 consumers in 19 countries released a report in early 2018 that stated that while the adoption of smart-speakers was sky-rocketing, the real challenge was keeping users engaged in the long term.

The voice-industry start-up Alpine detected the same problem one year earlier, when smart-speakers were still a breakthrough: *“While there are now more than 7,000 Skills to choose from on the Alexa platform, only 31% have more than one consumer review. This indicates that many of Alexa voice applications are ‘Zombie Skills’: they are accessible but are not heavily used or appreciated. While this issue is not unique to the Alexa app store (present in iOS and Android ecosystems), it is an issue that must be addressed.”*

During the trials, all users got back to using the same applications after a while: the basic functions of the device like music, reminders and timers. Although this is not necessarily a terribly bad behaviour, it undermines the capabilities of the device and jeopardizes the developer’s community that have to deal with not only a lack of revenue but low usage. In the long term, the growth of the app stores of Amazon Echo, Google Home and other providers of smart-speakers might experience a dramatic deceleration.

Supporting the developers is a decision that lays completely on the providers, maintaining a third-party ecosystem represents a major responsibility that consumes sources and can deliver irregular experiences to the users.

Amazon is very committed to its ecosystem of developers and has already addressed several issues pointed out by different sources.

Notes

1. Accenture: Digital Consumer Survey - <http://thesis.norori.com/url/c3mb9>
2. Alpine: The 2017 Voice Report - <http://thesis.norori.com/url/e2dxe>

The Alexa Developer Rewards rolled-out a special payment program¹ that rewarded “*top performing custom skills based on customer engagement, that are correctly categorized in any of eight eligible skill categories.*” The program recently expanded to “Skills for Kids” in the US under special conditions².

Amazon also shares guidelines, tutorials and source-code material to its community in several languages throughout the Alexa Developer portal, a web that has a high engagement and impact.

Privacy

Despite the lack of security breaches of smart-speaker data until now and the fact that we all already carry devices with microphones with us, people expressed certain discomfort with the possibility of being spied over smart-speakers.

Google Home and Amazon Echo never record users unless they hear the “wake word” and all their recording are easily manageable with the companion apps, this is communicated by both providers³, that also provide a physical mute-button on their devices that stops all listening from happening, even if the wake word is said.

Closing notes

Consumers are excited with the dawn of the market of smart-speakers and smart-assistants, while the adoption of other devices from the real of the Internet of Things like Smart Home Appliances (light-bulbs, thermostats and so on) is slow, all the conditions for an acceleration of the market exist.

Amazon, Google and Apple can expect a steady growth of their user base of smart-speakers, while general challenges exist and the industry is still trying to define its own identity and place, one can expect exciting developments in its near future.

Detailed graphics of the data collected during the research can be found on the appendix of this paper.

Notes

1. Amazon Announces New Ways to Earn Money with Your Alexa Skills - <http://thesis.norori.com/url/5-5rp>

2. Alexa Developer Rewards Program Expands to Skills for Kids - <http://thesis.norori.com/url/rxhsn>

3. Amazon.com Help: Alexa and Alexa Device FAQs <http://thesis.norori.com/url/cgp3h>



Conception of apps for Alexa

For the purpose of this paper, I will focus only on the development of voice-controlled applications for Alexa, the smart-assistant of the Amazon Echo line.

Other assistants were not considered for this document, as their market share is significantly smaller than the one of the Echo and their development ecosystem is not as advanced as the one offered by Amazon.

Applications for the Amazon Alexa are officially called “Alexa Skills” and are distributed to the device using the Alexa App on Android or iOS or by commanding a device with voice to install certain app.

The purpose of this chapter is to set guidelines for the development of voice-interaction centred applications, from their conception to the final deployment on the desired device and its distribution throughout the Amazon Skills Store.

We are used to interacting with digital applications using touch-screens, keyboards or a computer mouse. Voice is a totally novel approach to this rule, but it still feels somehow natural to talk to a computer, and this is no surprise, since our voice is the main interaction tool, we use it every day to communicate with others. Despite this familiarity, designers and developers are still figuring out what are the best practices to follow when interacting with machines using our voice.

One of my major goals with this paper is to contribute on the research of this best practices, with first-hand research and development on this field. The specific devices used for this research are the Amazon Echo Plus and the Amazon Echo Dot 2nd Generation. General technical details of each one will be specified further on.

Cover Image: Designed with sources from FreePic

Amazon Echo Plus

The Amazon Echo Plus was released in the US and Germany on the 31st of October 2017.

It is a direct evolution of the first Amazon Echo, with a similar design with improved technical specifications that include speakers with Dolby Sound, 7 second-generation far field microphones and an integrated smart-hub with ZigBee capabilities.

Its price-tag on release was 149.99 USD \$ and 149.99 EUR € depending on the market. It includes a free Philips Hue Light Bulb to let users experience the smart-home capabilities of the device out-of-the-box.

A full breakdown of the specifications of the device can be seen below¹:



Amazon Echo Plus Specifications

<i>Width</i>	8.4 cm
<i>Depth</i>	8.4 cm
<i>Weight</i>	954 grams
<i>Height</i>	23.5 cm
<i>Colours</i>	White, Black, Silver
<i>Dimensions</i>	23.5 cm (H): 8.4 cm (W): 8.4 cm (D)
<i>Sound</i>	Dolby Sound (2.5" woofer + 0.8" tweeter)
<i>Microphones</i>	7-far-field-microphone array
<i>Connectivity</i>	Wi-Fi (Dual-band, dual-antenna 802.11 a/b/g/n/ac @ 2.4GHz/5GHz) – Bluetooth 5.0 – ZigBee Pro
<i>Outputs</i>	3.5mm audio-jack
<i>Release Date</i>	October 2017

Notes

1. Amazon Product Page: Amazon Echo Plus - <http://thesis.norori.com/url/c9pts>



Amazon Echo Dot 2nd Generation

The second generation of the entry-level Echo Dot was released in October 2016.

Its shares almost the same design with the original Echo Dot except for the removal of the volume control dial. The most significant changes are under-the-hood, like the improved voice recognition, the Echo Spatial Perception (ESP) technology that allows multiple Echo devices to work together detecting and answering voice-requests and a dedicated Speech Processor.

Its price was almost cut in half when compared to the first generation as well, at 49.99 USD \$ / 59.99 EUR €. Amazon applied considerable discount during the holiday season, decreasing the price 60% more during Christmas 2017.

A full breakdown of the specifications of the device can be seen below¹:

Amazon Echo Dot Second Generation Specifications

<i>Width</i>	8 cm
<i>Depth</i>	8.5 cm
<i>Weight</i>	163 grams
<i>Height</i>	3.5 cm
<i>Colours</i>	White, Black
<i>Sound</i>	0.6" single speaker
<i>Microphones</i>	7-far-field-microphone array
<i>Connectivity</i>	Wi-Fi (Dual-band, dual-antenna 802.11 a/b/g/n/ac @ 2.4GHz/5GHz), Bluetooth 5.0
<i>Outputs</i>	3.5mm audio-jack
<i>Release Date</i>	October 2016

Notes

1. Amazon Product Page: Echo Dot (2nd Generation) - <http://thesis.norori.com/url/4fixd>

The development environment

Before the conception process starts, one must understand the development platform, with its specifications, limitations and guidelines in mind.

Amazon offers the Alexa Developer Portal¹, that serves as a perfect starting point for newcomers to the platform.

Skills for Amazon Alexa can be built using C#, Java, JavaScript, Go or Python, giving a wide choice of options to start getting used to the specifics of the Alexa platform.

Flexibility from Amazon doesn't stop with the wide choice of programming languages for the platform, but it goes as far as offering hosting and cloud-processing for the apps. It's also possible to self-host one's application, but it's not as practical.

Skills developed for Alexa can be hosted on Amazon Web Services (AWS), a cloud-computing environment that offers developers all the necessary tools for the creation of complex applications throughout the AWS Lambda server-less computing platform.

A simple Alexa application would follow the next interaction model, in which the user's input is recognized by the Amazon Web Services platform and send to the appropriate AWS Lambda, which process the request based on its own code and generates an answer that travels back to the user



Notes

1. Amazon Alexa Developers <http://thesis.norori.com/url/u-rby>

More complex interactions can of course, occur, and since the AWS platform supports several protocols that vary from databases to IoT Gateways and an open API the opportunity of developing complex and useful apps exists.

Intents

Alexa has a strict syntax for invoking third-party apps, users must always mention the name of the application they want to use and what they want from it. Alexa recognizes these phrases as “intents”, for example, if a user wants to open an app, they have to say:

“Alexa, open Darmstadt Tips”

“Alexa” is the wake word that puts the device in listening mode, the word “open” is the intent or desired action followed by the name of the app to which the intent is directed. Alexa interprets this as an open intent for the application “Darmstadt Tips”.

Users can also trigger specific actions from Alexa Skills that support them, for example:

“Alexa, ask Help Me Decide Apple or Banana?”

In this case, Alexa triggers the intent “ask” for the app “Help Me Decide” which gets a request containing the sentence “Apple or Banana”. Amazon refers to the whole sentence as an “utterance”.

Since the app is programmed to recognize two different words or phrases separated by the word “or” and create “slots” it results in the following string of data:

SlotChoiceA: Apple
SlotChoiceB: Banana

The app then proceeds to shuffle them and pick one at random for the user, creating an answer sentence that reads as:

“I think you should go for SlotRandomChoice”

Were “SlotRandomChoice” refers to the slot picked up at random by the computer. All this process occurs in just a few seconds, giving the user a prompt reply that makes the interaction feel smooth and natural.

Intents and Slots are the most important interaction tool the developers have, since it's the main input source for instructions coming from the user.

Developers are able to define which words or phrases trigger which intents using a tool called "Interaction Model", which is a simple JavaScript Object Notation (JSON) integrated on Amazon's Developer platform.

A typical interaction model contains extensions of the basic intents Amazon Alexa has, like Stop, Cancel or Help, the name of the skill and as many custom intents as the developer requires. A simple interaction model looks like this:

```
{
  "LanguageModel": {
    "intents": [
      {
        "name": "AMAZON.CancelIntent",
        "samples": [
          "Stop"
        ]
      },
      {
        "name": "AMAZON.HelpIntent",
        "samples": [
          "Help me",
          "What can I do with you",
          "how can you help me",
          "what do you do"
        ]
      },
      {
        "name": "AMAZON.StopIntent",
        "samples": [
          "Stop",
          "Bye",
          "That's all"
        ]
      }
    ],
    "invocationName": "darmstadt tips"
  }
}
```

Every time a request is directed to a skill, Alexa will search for matches among the intents and will trigger the closest match. In case no match is found an error will be returned.

Developers must have a clear idea of what their skill can make and how it will communicate this to the user. Keeping interaction brief, simple and direct is the best decisions.

Designing for voice

What can users accomplish with my skill? How do I think they will normally interact with it? And how can I make them go from start to finish faster? Interaction must be simple, direct and natural.

Developers must adopt the perspective of the user's and build from there, a clever idea for this is constructing scripts and making the conversation linear and uncluttered. One single sentence must be enough for everything.

The name of the skill is another crucial element that will impact the usage of our application considerably. The name of the skill must answer a simple question: What's the need my skill is covering? For example: Help Me Decide? Did Someone Walked the Dog? Did Someone Watered the Plants?

Once the name and purpose of the skill are defined, the scripting phase must commence. Scripts help you identify the most likely way users will communicate with your application. For example, in case of a skill that keeps track of when the dog walked, users might ask:


- Did the dog walk already?
- Was the dog out today?
- Did someone take the dog out today?
- When was the dog out?

Plan Alexa to react accordingly to each one of the sentences below:

- The dog was walked last time 6 hours ago.
- Yes, the dog walked today already.
- No, the dog was not out today.
- The dog was out just over an hour ago.

While writing your skill, keep in mind that the best is to have few intents at first, so you can cover them as good as you can. Other intents can be added in further updates.

Brainstorm for as many ideas of interaction as you can and write them all down, see it as if you were training Alexa hear and talk.



Developing for Amazon Alexa

The best way to learn the dynamics of the creation of Amazon Skills is doing one by your own, for that reason, and following all the knowledge gathered during the writing of this paper I created a series of prototypes for the Amazon Echo.

The learnings from the primary research and suggestions from its participants were taken in consideration for the conception and testing of the skills as well.

In this chapter, you will be presented with a set of skills divided into two categories: Learning Skills and Custom Skills.

The Learning Skills were developed following all the guidelines and using substantial source-code provided by Amazon. They serve as elemental learning tools that I used to introduce all the concepts related to developing, training, configuring and publishing the skills.

The Custom Skills are skills developed from scratch based on ideas collected from the primary research. They are simple but effective applications that cover specific needs.

Cover Image: Designed with sources from FreePic and Amazon Echo Press Images

Learning Skill: Darmstadt Tips

User: Alexa, ask Darmstadt Tips for a tip.

*Alexa: Here's an interesting tip about Darmstadt:
Darmstadt has one of the best Mexican restaurants in
Hessen: "La Hacienda" in Holzhofallee 2.*

Facts and Tips skills are one of the most popular types of skills in the Alexa Skills Store. They are simple, direct and quite useful when the topic interests the user.

Amazon divides the developing process of Alexa Skills in two different portals: The Amazon Developer Portal and the Amazon Web Services (AWS) Console.

The Amazon Developer Portal is where the skill is initially created and its basic properties like name, type and icon defined. Here is also where the Interaction Model is built, the Testing and Debugging takes place and where the submission for final publishing is requested.

The AWS Console is where the code behind the skill is stored and executed. All other dependencies like databases, gateways and APIs are also executed here.

One straightforward way to visualize the dynamic between the two portals is to see the Developer Portal as the body and senses of the skill and the AWS Console as the brain that process and reacts to the information it receives.

The app in the Amazon Developer Portal (ADP) and the code in the AWS Console connect to each other using identification numbers called Skill ID in the DP and Amazon Resource Name (ARN) in the case of the AWS Console.

An account in both services is required to start developing.

Amazon Developer Portal

The Amazon Developer Portal is where the developing process begins. Developers define parameters using a quite straightforward interface. All the basic information of the skill is defined in just a few minutes, including the invocation name, supported languages and platform.

The screenshot shows the Amazon Developer Console interface for configuring a skill. The top navigation bar includes 'amazon // DEVELOPER CONSOLE' and the user name 'RONALD NORORI'. Below the navigation bar, there are tabs for 'DASHBOARD', 'APPS & SERVICES', 'ALEXA', 'REPORTING', 'SUPPORT', 'DOCS', and 'SETTINGS'. The 'ALEXA' tab is selected. The main content area shows the skill configuration for 'Darmstadt Tips'. On the left, there is a sidebar with a checklist of categories: 'Skill Information', 'Interaction Model', 'Configuration', 'Publishing Information', and 'Privacy & Compliance', each with a green checkmark. The 'Interaction Model' is set to 'Builder'. The main configuration area includes: 'Skill Type' set to 'Custom'; 'Language' set to 'English (U.S.)'; 'Application Id' set to 'amzn1.ask.skill.9c3399dc-e175-458d-9af6-1cd62d1a41c8'; 'Name' set to 'Darmstadt Tips'; and 'Invocation Name' set to 'darmstadt tips'. A blue information box at the bottom states: 'For successful Alexa Skills Certification, please review and follow our Invocation Name Guidelines as well as our Certification Requirements.'

The most complex process done in the ADP is the creation of the Interaction Model, that dictates the array of different orders the skill can understand and organize and defines how this information is processed before being sent to the AWS Console.

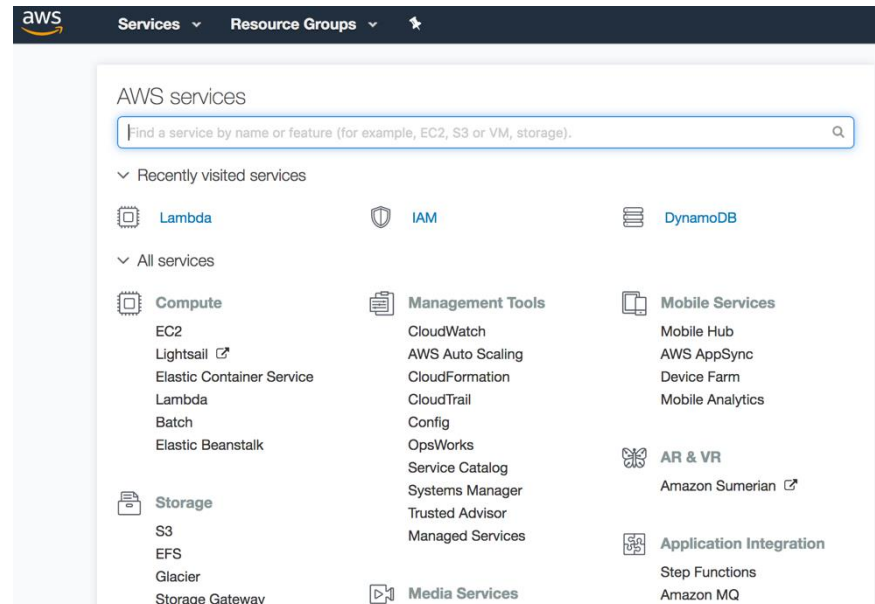
Developers can also request special permits that allow extended functionality for their skill, like getting the exact location of the device, modifying shopping lists or connecting external accounts.

Ultimately, once the development of the skill is finalized, the process of submission for approval and publishing is initiated in the ADP interface.

Category, icon, description, global availability and search keywords are defined on this last step. Most importantly, the instructions for the testing Amazon's staff will make are also established. Once a skill is approved by Amazon it becomes available for all the stores selected by the developer.

Amazon Web Services Console

The AWS Console is a far more complex interface, because it offers a broad range of services to all kinds of platforms and not only to the Amazon Alexa.



For the development of Alexa Skills, the only function needed is the Lambda service. Amazon defines Lambda¹ as “a compute service that lets you run code without provisioning or managing servers. AWS Lambda executes your code only when needed and scales automatically, from a few requests per day to thousands per second.”

Lambda supports several programming languages and is capable of connecting to other internal services provided by AWS like Dynamo Databases for storage or Cloud Watch for logging and debugging. A series of API's and gateways allow it to connect with external servers as well.

The platform provides a complex but yet easy to use interface that enables developers to create intricate applications.


During the months this paper was written, the AWS Console and the ADP undergone several re-designs that changed its appearance and functionality drastically.

Notes

1. Amazon: What Is AWS Lambda? - <http://thesis.norori.com/url/wpbcn>

When a user invokes a skill, the ADP interface converts the audio into a text order that is then send to the correspondent AWS Lambda application.

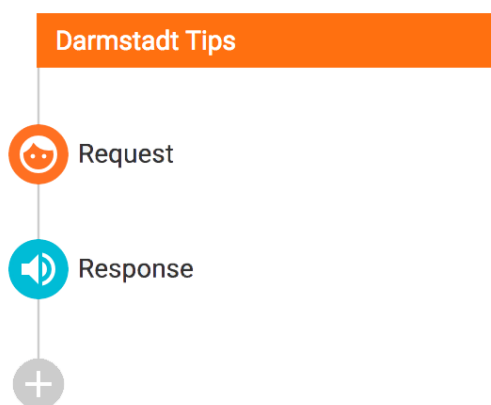
In the case of the skill Darmstadt Tips and as with almost all other Alexa Skills, this connection is done via the Alexa Skill Kit, a simple “trigger” interface that connects the AWS Lambda and ADP using simple identification numbers and a shared API.



After receiving an order from the Alexa Skills Kit, Lambda proceeds to follow the instructions given on the Function Code, in this case called “Darmstadt-Tips”. The function code defines the different ways the skill will react to input. A log of all interactions is temporary saved on the Cloud Watch logs.

The function code of Darmstadt Tips consists a simple JavaScript that selects an item, in this case a “tip”, from an array of options at random.

In the case of Darmstadt Tips, the communication is a one-way street, where the user invokes the skill and Alexa responds back. No dialog is necessary since no additional information is required.



More complex communication is advised only when necessary, as is the case with the second prototype.

Learning Skill: Hessen Trivia

User: Alexa, launch Hessen Trivia.

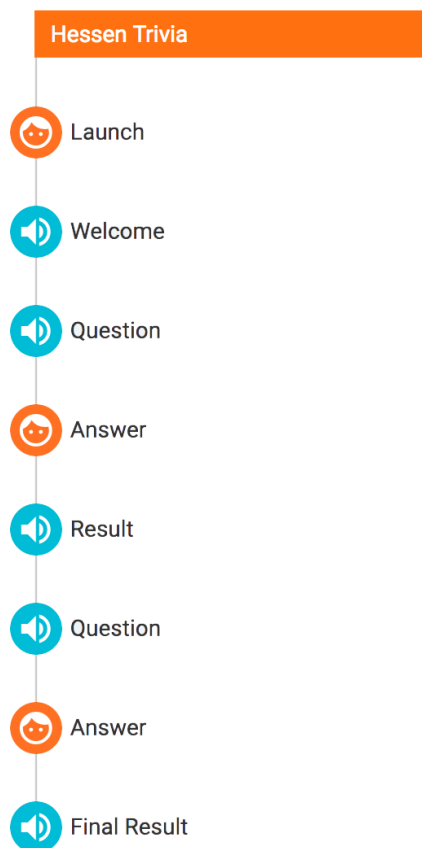
Alexa: Welcome to Hessen Trivia. I will ask you three questions, try to get as many right as you can. Just say the number of the answer. Let's begin.

Alexa: Which of the following drinks originated in Hessen? Glühwein, Vita Cola or Club Mate?

User: One.

Alexa: That's correct! Your score is one. Next question.

Skills like games require a more complex Interaction Model, in the case of the Hessen Trivia, the users are asked for their feedback three times before getting a final score.



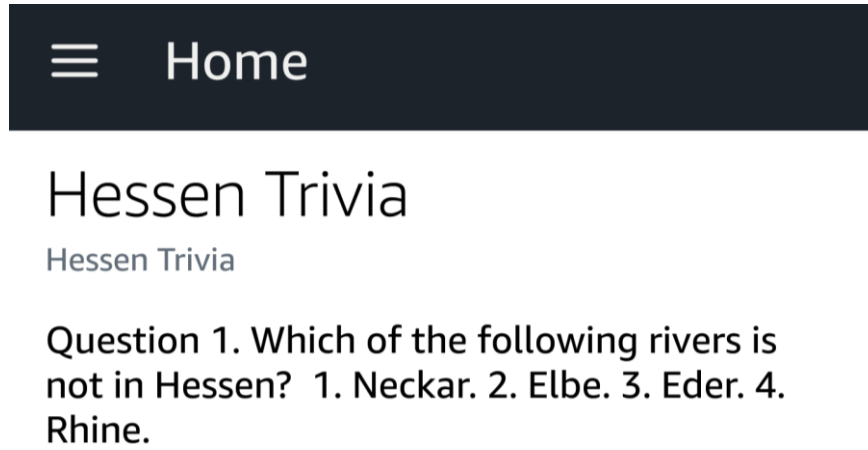
For skills like this to work well, the developer must take in consideration the creation of information slots and several sample utterances on the ADP and the programming of the required intent responses on the AWS Lambda Function Code.

In the case of Hessen Trivia, the ADP and AWS Lambda Function code are trained to recognize the following intents:

- **Cancel, Stop and Help Intents:** All of them are part of the mandatory intents every skill must support. As their names, they act as ways to exit the skill or get help.
- **Answer Intent:** Designed to recognize the number the user says every time they're prompted for an answer.
- **I don't know Intent:** Handles the reaction of the app every time the user says, "I don't know the answer".
- **Start Over Intent:** Re-starts the skill every time the user says, "I want to start again", "New game" or similar requests.

Developers must include intents on the ADP so Alexa can recognize the input as a separated action and direct it to AWS Lambda.

The Function Code of the Hessen Trivia skill is capable of handling all the intents described above in a simple way, giving voice answers accordingly and pushing "information cards" to the Alexa app for users to review the interaction later on their smartphones.



The information cards allow users to visualize the interaction in case they want to review what they heard. It's also helpful to direct users to external sources of information, for example: If the user requested for an address it can launch a navigation to get to the intended place.

Custom Skill: Help Me Decide

User: Alexa, ask Help Me Decide, Pancakes or Cereal?

Alexa: So, Pancakes or Cereal. I think you should go for Pancakes!

Based on the feedback received during the Primary Research, I maintained my prototypes as simple as possible while keeping the interaction natural.

Help Me Decide is an archetype of this design philosophy: Interaction is short and direct, while the answer from Alexa is concise and brief.

The skill prompts the user for the input of two options separated by the word “or”, that helps the skill differentiate between them.

Once the information is inserted by the user, the skill responds back with a confirmation sentence, mentioning the two options inserted by the user and finalizing the sentence with the decision the computer made.

The decision of the skill is done using a simple random selector, the skill is built entirely on JavaScript.

Another important detail of the skill is its name, after all this is the way users will invoke its action. “Help Me Decide” is one of the most natural ways to ask Alexa for her decision.

Because no native function for selecting between an array of options exist, the third-party skill is triggered instead. Since its release, the skill has registered an average usage of 100 unique sessions per week.

Help Me Decide is a great example of how to design skills for Alexa:

- Their functionality must be easy to understand and must be related to its name.
- Interaction must be brief and direct.
- Engagement should take just a few seconds from beginning to end.

Custom Skill: Walked the Dog

User: Alexa, ask Walked the Dog if the dog was out already?


Alexa: Last time I heard the dog was out was over 6 hours ago.

User: Alexa, tell Walked the Dog that I just went for a walk with Bubbles.

Alexa: All right! I will remember you just walked the dog.

The mechanic behind the “Walked the Dog” skill is another example of simple and direct communication. Although the interaction on the user’s side is simple, the actual functionality of the skill is far more complex than the other prototypes presented until now.

The “Walked the Dog” skill makes use of databases to keep track of the information inserted by the user. This enables the skill to write and read information from any other Echo device connected to the same account.



The database is managed using *DynamoDB*, a service integrated into the AWS platform that allows the creation and management of simple data tables.

Every time the user invokes the skill, they are prompted with a Welcome Message, if after a few seconds no answer is received a Help Message with examples of how to use the skill is broadcast to the Echo device and the Alexa App.

When users want to report that they walked the dog, they can just say “Alexa, tell Walked the Dog I was just out with the dog” and the skill will register the exact time this action was done.

When users want to know when was the last time someone walked the dog they can just ask “Alexa, ask Walked the Dog”.

The inspiration for the development of this prototype came during the primary research for this paper, when users living in the same flat couldn’t keep track of daily tasks like walking or feeding the dog, taking care of the plants or cleaning.


Walked the Dog’s model can be applied to other purposes where users need to keep track of cyclical tasks. To prove this point, “Watered the Plants” was created using the same source code.

The same mechanic would work for skills that help users remember, for instance, if the dog or cat were fed, the apartment cleaned or the mail picked up.

Prototypes

All prototypes described on this chapter were approved by Amazon and published on the Alexa’s Skill Store, to have a first-hand look at them install them from the following page:

thesis.norori.com/prototypes





Conclusions

According to Gartner, 5.5 million new things get connected every day. By 2020, they estimate that there will be up to 21 billion connected devices¹.

It starts getting clear that the Internet of Things is here to stay, affecting industries of all kinds with its applications for Home, Health, Industry, Transportation and more.

Smart-speakers will remain the main touch-point most consumers will have with the industry at least for the next few years. For now, they are primarily used for simple tasks like music streaming and weather forecasts, but smart-speakers are a critical first step in the impending shift toward the “connected home”.

Most homes are not ready to support the full-on integration with the Internet of Things that the smart-speakers enable, but once users start using voice-enabled devices they are more likely to take smaller steps making the shift to connected devices.

This shift towards smart-homes will generate big revenue in the upcoming years, with Amazon and Google leading the way and other big companies like Apple joining the race later.

While consumers get more comfortable with the presence of always-listening devices and equip their homes with connected appliances manufacturers of IoT devices gear up to get their share of the market.

Notes

1. Gartner Says 6.4 Billion Connected “Things” Will Be in Use in 2016, Up 30 Percent From 2015 “2020” <http://thisis.norori.com/url/c2x45>
Cover Image: Designed with sources from FreePic

Germany is already seeing the first innovators on this field, with companies like Innogy, Bosch and Homematic-IP designing all kind of products for the smart-home and distributing them throughout traditional technology-stores like Saturn² or Media Markt³.

In a global scale, consumer spending on smart-home devices is expected to increase an average of 30% by 2022¹, with the North American, Asia Pacific and Western European markets leading the way. IoT-hardware is selling and there are no signs of this to decelerate in the near future.

Smart-speakers will follow the same trends as the IoT, with unstoppable sales. But while the bottom line looks good, specific challenges for smart-speakers start to appear.

How will manufacturers of smart-speakers keep their consumers engaged? Is there more to the devices than voice-controlled music playback and weather reports?

For smart-speakers to triumph in the long term, manufacturers need to nurture their communities of third-party developers in the same way that smartphone manufacturers did when the mobile smartphone app-stores started to appear.

Manufacturers are already identifying the challenges they face and taking steps to countermeasure them. Only with clever and premeditated planning will they make of the smart-speakers a new, permanent and wide-spread segment of user-interaction.

The constant innovation from developers, the growing availability of smart-home appliances and the steady sales of smart-speakers in combination almost guarantee that living in a smart-home and talking with our smart-assistants will be the standard of the future.

Notes

1. *Visual Capitalist: The Battle for Smart Speaker Market Share* <http://thesis.norori.com/url/pnlip>

2. *Smart Home Saturn* - <http://thesis.norori.com/url/glnoi>

3. *Smart Home Media Markt* - <http://thesis.norori.com/url/wk8hz>

Personal Learnings

During the writing of this paper I had the chance to get in-deep insight into an exciting and thriving new industry.

Thanks to the learnings gathered during the primary research I could understand better how users interact with smart-speakers and other smart-home appliances and got to get a better notion of where things are headed.

I personally believe smart-speakers, smart-assistants and the smart-home will all become entangled with our daily lives. They will get more advanced, but will maintain an unprecedented natural simplicity.

Seeing how the hype for this topic didn't vanish but instead got more intense during the last months validated the fact that the Internet of Things is the "next big thing".

Participants of my trials and surveys also assured me this with their great excitement and optimism for the industry.



The prototyping phase helped me train my coding and conception skills presenting me the challenge to develop for a platform with a totally new user-interaction environment.

Developing Alexa Skills that were published on the Amazon Store, and are now being used by hundreds of costumers per week, make me feel proud that the products of this last months of work have a practical usability.

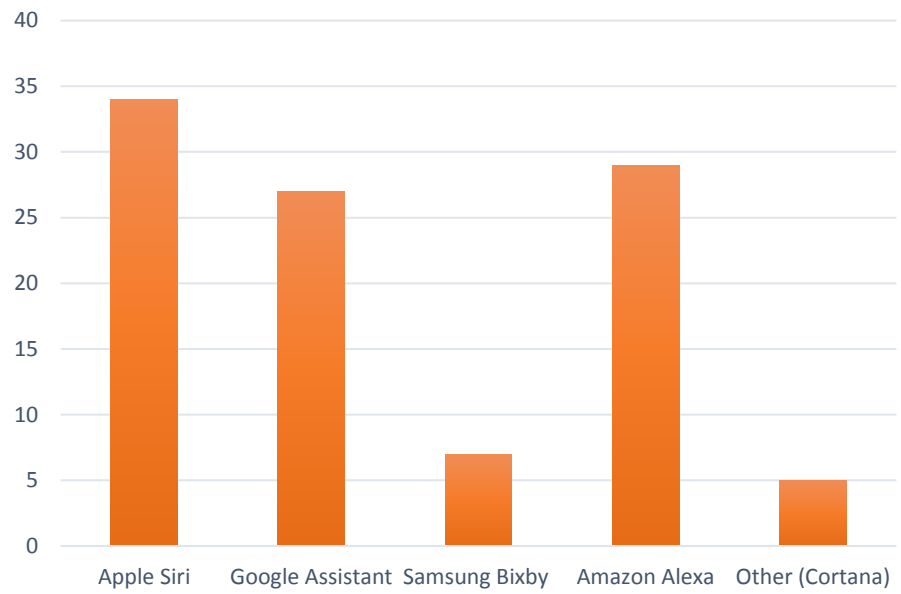
The knowledge and skills I gathered during the production of this paper helped me become a more thorough professional and makes me feel ready to face future challenges with more confidence.

Appendix


Survey Results: Opinion and usage patterns of smart-speakers and smart-assistants




Which of the following smart-assistants have you use?




What was your first impression after interacting with a smart-assistant using a smart-speaker?



Which of the following functions of smart-speakers have you use?




Do you think of smart-speakers as something you need or don't need in the next 5 years?




How many smart-devices do you own?



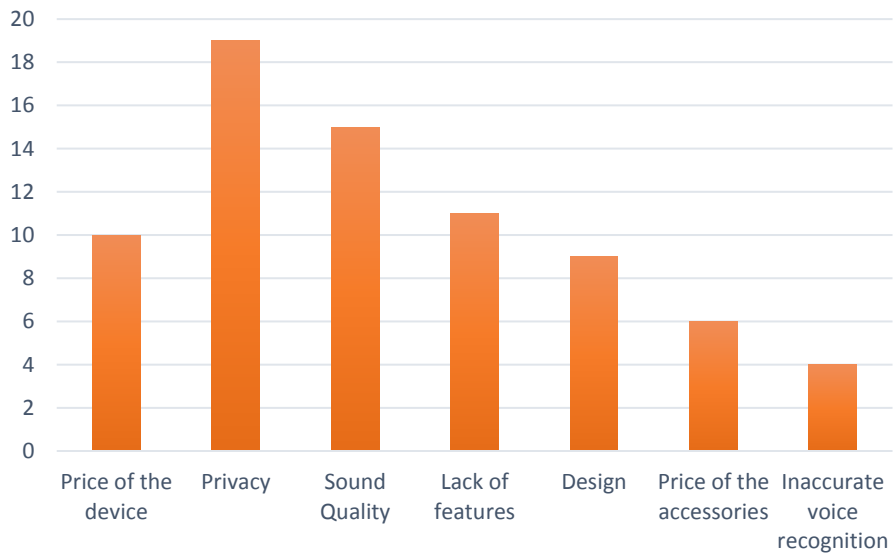
How often do you control smart-devices with smart-speakers?



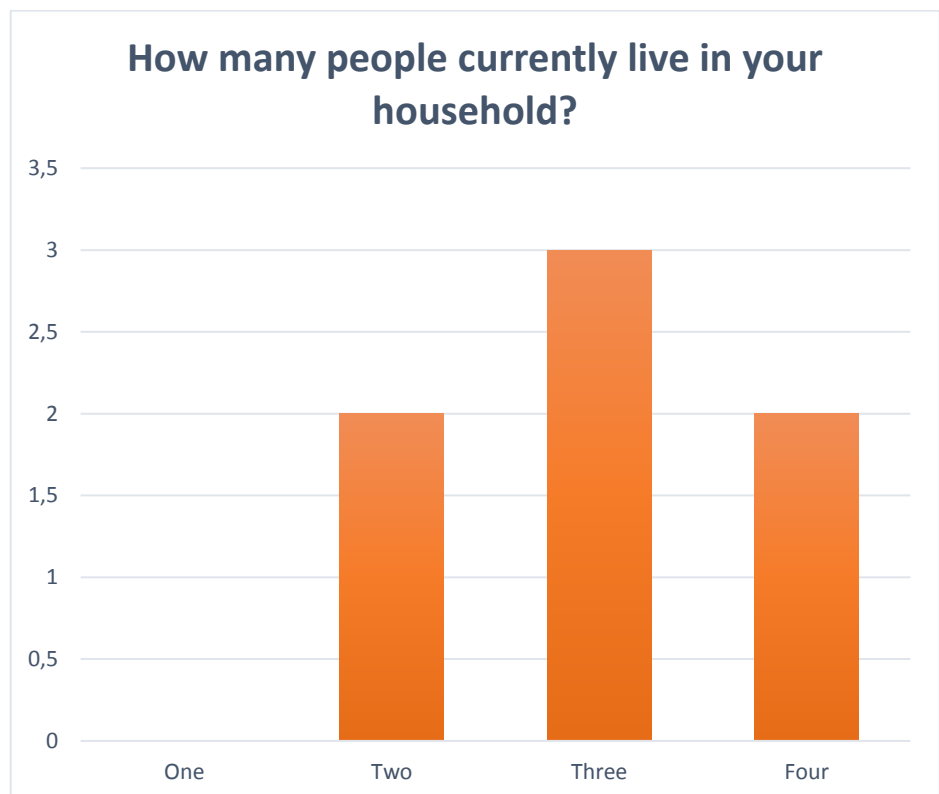

How often do you use third-party apps with your smart-speaker?




Which of the following negative features of smart-speakers you consider more discouraging?




Survey Results: Trial Period Survey



What is your first reaction to Amazon Echo Dot 2nd Generation?




When you think about Amazon Echo, do you think of it as something you need or don't need?




How long was your trial?


7 days 28 days




How often did you use the Amazon Echo during the trial period?



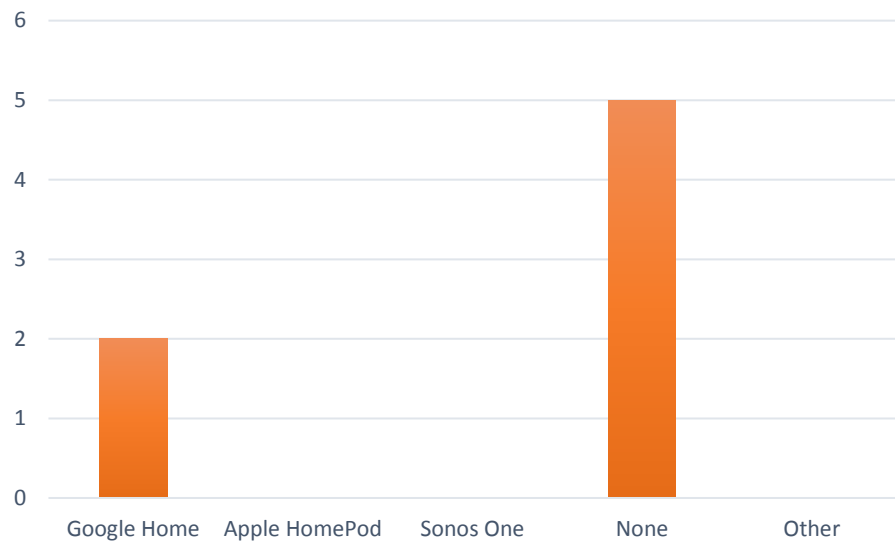
Which of the following functions did you use the most?



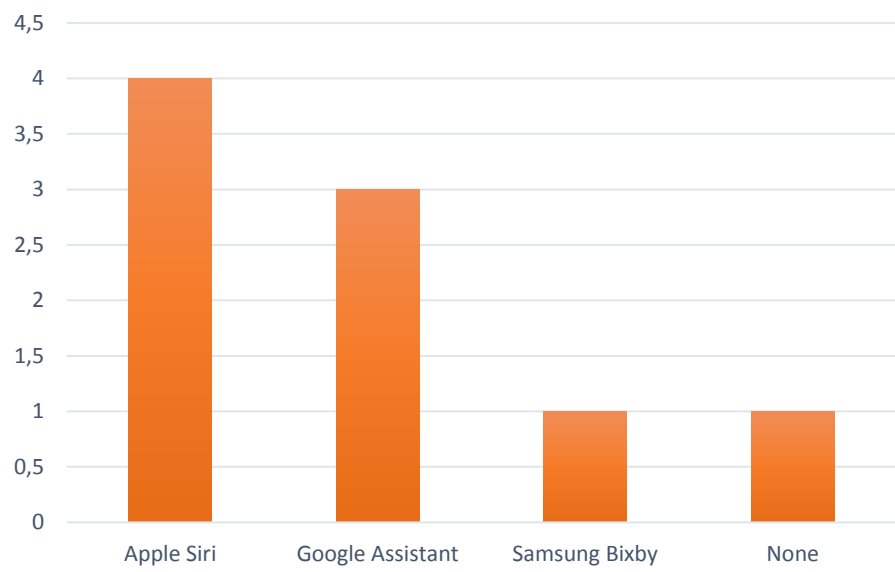
Which of the following negative features of smart-speakers you consider more discouraging?



Have you use other smart-speakers apart from the ones developed by Amazon?



Have you use other smart-assistants apart from Alexa?



Source Code: Help Me Decide (JavaScript)

```
var Alexa = require('alexa-sdk');
const skillName = "Help Me Decide";
var speechOutput = "";
var handlers = {
    "DidNotRecognised": function () {
        var speechOutput = "Sorry, I did not hear you. If
you need help, say help.";
        this.emit(':ask', speechOutput, speechOutput);
    },
    "RandomWordIntent": function () {
        var sentence =
this.event.request.intent.slots.FirstWord.value;
        var firstWord, secondWord;
        var sentenceLength = sentence.length;
        if (sentence.search(/\bor\b/i) !== -1) {
            var indexOfOrWord =
sentence.search(/\bor\b/i);
            firstWord = sentence.substr(0,
indexOfOrWord);
            secondWord =
sentence.substr(indexOfOrWord + 2, ((sentence.length) - indexOfOrWord + 2));
            setData(firstWord, secondWord);
        } else if (sentence.search(/\band\b/i) !== -1) {
            var indexOfAndWord =
sentence.search(/\band\b/i);
            console.log("indexOfAndWord :-
"+indexOfAndWord);
            firstWord = sentence.substr(0,
indexOfAndWord);
            console.log("firstWord :- " +
firstWord);
            secondWord =
sentence.substr(indexOfAndWord + 3, ((sentence.length) - indexOfAndWord +
3));
            console.log("secondWord :- " +
secondWord);
            setData(firstWord, secondWord);
        } else {
            speechOutput = "Sorry, I did not
hear you. If you need help, say help.";
            this.emit(':ask', speechOutput,
speechOutput);
        }
        this.emit(':tell', speechOutput, speechOutput);
    },
    "AMAZON.HelpIntent": function () {
        var speechOutput = "";
        speechOutput += "I can help you decide between
two different options, for example say Foxes or Pandas, Study or Watch TV,
Pizza or Sushi, and I will pick one up for you";
        this.emit(':ask', speechOutput, speechOutput);
    },
    "AMAZON.StopIntent": function () {
        var speechOutput = "Bye bye!";
        this.emit(':tell', speechOutput);
    },
    "AMAZON.CancelIntent": function () {
        var speechOutput = "Goodbye!";
        this.emit(':tell', speechOutput);
    },
    "LaunchRequest": function () {
        var speechText = "";
        speechText += "Welcome to Help Me Decide, tell me
two things and I will pick one up for you.";
        var repromptText = "For instructions you can say,
help.";
        this.emit(':ask', speechText, repromptText);
    }
};
function setData(a, b) {
    var result;
    var hrTime = process.hrtime();
    var time = hrTime[0] % 2;
    if (time == 0) {
        result = a;
    } else {
```

```
        result = b;
    }
    speechOutput = "So " + a + " or " + b + "...I think you should
go for " + result + ".";
}
exports.handler = function (event, context) {
    var alexa = Alexa.handler(event, context);
    alexa.APP_ID = "amzn1.ask.skill.c27391a7-425a-4d0c-8160-
f1cde2e99d36";
    alexa.registerHandlers(handlers);
    alexa.execute();
};
```

Source Code: Walked the Dog (Python)

```
from __future__ import print_function
import boto3
import time
from datetime import datetime

def lambda_handler(event, context):
    // Determine Intent
    if event['session']['new']:
        on_session_started({'requestId': event['request']['requestId']},
                           event['session'])

    if event['request']['type'] == "LaunchRequest":
        return on_launch(event['request'], event['session'])
    elif event['request']['type'] == "IntentRequest":
        return on_intent(event['request'], event['session'])
    elif event['request']['type'] == "SessionEndedRequest":
        return on_session_ended(event['request'], event['session'])
    // Action of each intent
def on_session_started(session_started_request, session):
    """ Called when the session starts """

    print("on_session_started requestId=" +
          session_started_request['requestId']
          + ", sessionId=" + session['sessionId'])

def on_launch(launch_request, session):
    return get_welcome_response()

def on_intent(intent_request, session):
    //
    print("on_intent requestId=" + intent_request['requestId']
          + ", sessionId=" + session['sessionId'])

    intent = intent_request['intent']
    intent_name = intent_request['intent']['name']

    if intent_name == "AMAZON.HelpIntent" or intent_name == "TheDogWhatWas":
        return get_help_response()
    elif intent_name == "TheDogReadLogIntent":
        return handle_read_log(intent_request, session)
    elif intent_name == "TheDogWriteLogIntent":
        return handle_log_timestamp(intent_request, session)
    elif intent_name == "AMAZON.CancelIntent" or intent_name ==
"AMAZON.StopIntent":
        return handle_session_end_request()
    else:
        return get_help_response()

def on_session_ended(session_ended_request, session):
    //
    print("on_session_ended requestId=" + session_ended_request['requestId']
          +
          ", sessionId=" + session['sessionId'])
    //

def get_welcome_response():

    session_attributes = {}
    card_title = "Welcome"
    speech_output = "Welcome to the Walked The Dog app!"
    reprompt_text = speech_output
    should_end_session = False
    return build_response(session_attributes,
                           build_speechlet_response(card_title, speech_output,
                                                      reprompt_text, should_end_session))

def handle_read_log(read_request, session):

    dt_timestamp = fetch_log(session['user']['userId'])

    session_attributes = {}
    card_title = "When did the dog go out?"
    speech_output = "The last time I heard someone walked the dog was " +
convert_to_speech(dt_timestamp)
    reprompt_text = speech_output
    should_end_session = True
    return build_response(session_attributes,
                           build_speechlet_response(card_title, speech_output,
                                                      reprompt_text, should_end_session))
```

```

def handle_log_timestamp(log_request, session):
    // Timestamp
    timestamp = log_request['timestamp']

    save_timestamp(session['user']['userId'])

    session_attributes = {}
    card_title = "Logged"
    speech_output = "All right! I will remember you walked the dog at this
time"
    reprompt_text = speech_output
    should_end_session = True
    return build_response(session_attributes,
        build_speechlet_response(card_title, speech_output,
    reprompt_text, should_end_session))

def get_help_response():
    session_attributes = {}
    card_title = "Help"
    speech_output = "You can tell me if you walked the dog, and I will
remember the exact time you reported this to me when someone asks me. For
example, if you say, «I just walked the dog», and hours later someone else
asks, «When was the dog out?», I will tell them how long ago I heard the dog
was out."
    reprompt_text = speech_output
    should_end_session = False
    return build_response(session_attributes,
        build_speechlet_response(card_title, speech_output,
    reprompt_text, should_end_session))

def handle_session_end_request():
    session_attributes = {}
    card_title = "Session Ended"
    speech_output = "Bye bye!"
    reprompt_text = speech_output
    should_end_session = True
    return build_response(session_attributes,
        build_speechlet_response(card_title, speech_output,
    reprompt_text, should_end_session))

// Database

def save_timestamp(userId):

    sdb = boto3.client('sdb')
    response = sdb.create_domain(DomainName='TheDog')
    response = sdb.put_attributes(DomainName='TheDog',
        ItemName=userId,
        Attributes=[
            {
                'Name' : 'lastTime',
                'Value' : str(time.time()),
                'Replace' : True
            },
        ],
    )

def fetch_log(userId):
//Store

    sdb = boto3.client('sdb')
    response = sdb.create_domain(DomainName='TheDog')
    response = sdb.get_attributes(DomainName='TheDog',
        ItemName=userId)
    lastFedTime = float(response['Attributes'][0]['Value'])

    simulated_timeshift = 3*7200
    return datetime.fromtimestamp(lastFedTime)

def convert_to_speech(dt_timestamp):
// Relative Time Conversion

    one_hour = 3600
    one_day = 24 * one_hour

    current_time = datetime.fromtimestamp(time.time())

    delta_time = int((current_time - dt_timestamp).total_seconds())
    print(delta_time)

    if (delta_time < one_hour):
        return "less than 1 hour ago."

```

```

elif (delta_time < one_day):
    hours = str(int(delta_time / one_hour))
    return "just over " + hours + " hours ago."
else:
    return "more than 1 day ago."
// Cards
def build_speechlet_response(title, output, reprompt_text,
should_end_session):
    return {
        'outputSpeech': {
            'type': 'PlainText',
            'text': output
        },
        'card': {
            'type': 'Simple',
            'title': 'I Walked The Dog - ' + title,
            'content': 'I Walked The Dog - ' + output
        },
        'reprompt': {
            'outputSpeech': {
                'type': 'PlainText',
                'text': reprompt_text
            }
        },
        'shouldEndSession': should_end_session
    }

def build_response(session_attributes, speechlet_response):
    return {
        'version': '1.0',
        'sessionAttributes': session_attributes,
        'response': speechlet_response
    }

```

The source-code of the learning examples can be found at:
thesis.norori.com

References

All references are managed throughout a self-hosted URL shortening service that makes it easier to get to the information.

To visit the source, replace the “code” with the corresponding short code.

<http://thesis.norori.com/url/code>

For example, to visit the first URL, change “code” for “wk8h2”.

Short code	Title
/wk8h2/	Smart Home MediaMarkt
/9ln01/	Smart Home Saturn
/pnl1p/	Amazon vs. Google: The Battle for Smart Speaker Market Share
/c2x45/	Gartner Says 6.4 Billion Connected "Things" Will Be in Use in 2016, Up 30 Percent From 2015
/wpbcn/	What Is AWS Lambda? - AWS Lambda
/u-rby/	Amazon Alexa
/rxhsn/	Alexa Developer Rewards Program Expands to Skills for kids: Alexa Blogs
/cgp3h/	Amazon.com Help: Alexa and Alexa Device FAQs
/c3mb9/	Digital Consumer Survey - CES 2018 - Accenture
/e2dxe/	The 2017 Voice Report by Alpine (fka VoiceLabs) - Alpine.AI
/8up36/	You're not alone: 51 percent of Amazon Echo owners have it in the kitchen - Recode
/pyi3m/	What Amazon's Alexa economy pays the people building its skills
/x5id5/	VoiceLabs is 'Hibernating' its Sponsored Messages Amazon Alexa Advertising Network - VoiceLabs
/5-5rp/	Amazon Announces New Ways to Earn Money with Alexa Skills: Alexa Blogs
/7lwt9/	Mobile app user retention & churn 2017 - Statistic
/tc613/	iOS app user retention rate by category 2016 - Statistic
/3sjqk/	Earn Money for Developing Alexa Skills that Customers Love
/n2r20/	Amazon Closes Year With 266% Alexa Skill Growth in U.S. - Voicebot
/c8q6p/	The 2017 Voice Report by Alpine (fka VoiceLabs) - Alpine.AI
/ixht9/	Chart: What the Amazon Echo Is Actually Used For - Statista
/7yfqh/	Global virtual & augmented reality headset sales revenue 2016-2020 - Statistic
/drd5z/	Virtual Reality Headset Sales Hit 1 Million
/vsily/	Top Trends in the Gartner Hype Cycle for Emerging Technologies, 2017 - Smarter with Gartner
/t6m6s/	Tablet Sales Are Down: 5 Reasons Why - Time
/ilvvg/	Tablet users worldwide 2013-2020 - Statistic
/4fixd/	Echo Dot (2nd Generation) - Alexa-enabled Bluetooth Speaker - Black

/cgpts/	Amazon Echo Plus — The simple way to start your smart home
/yhk7o/	Apple HomePod is leaving a nasty stain on owners' tables - Metro News
/sqjle/	Apple HomePod review: A great speaker that's not so smart
/5pwla/	HomePod arrives February 9, available to order this Friday - Apple
/q6l4j/	HomePod reinvents music in the home - Apple
/zbhs3/	Apple Reports First Quarter Results - Apple
/-gu6b/	Google is introducing a new Smart Display platform - The Verge
/jxjze/	Google Assistant - Android Apps on Google Play
/omuvf/	Susan Bennett - Voices & Vocals
/t6b3d/	Spectacles - Features
/se89h/	Bluetooth Vs. Bluetooth Low Energy: What's the Difference?
/enppa/	ZigBee vs Z-Wave for the IoT - Electronics360
/n35tp/	Z-Wave – Official Website
/odv2r/	ZigBee Specifications and FAQ
/uc2uo/	What is Machine Learning? - SAP
/mxy18/	What Is Big Data? - SAS US
/roykq/	What is Cloud Computing? - Amazon Web Services
/ynfjl/	What is cloud computing? - IBM Cloud
/o-d17/	What is cloud computing? A beginner's guide - Microsoft Azure
/9ywqm/	Amazon says there are now 5,000 people working on Alexa — Quartz
/tamco/	56 Million Smart Speaker Sales in 2018 Says Canals - Voicebot
/podo7/	Apple's Siri Virtual Assistant Has Lost Millions of Users - Fortune
/nkav6/	New Research Report: Sluggish Adoption of AI-Driven Personal Apps
/nouif/	And Now for the Rest of the Siri Story
/ydwj3/	Review: With Siri, iPhone finds its voice - CNN
/gzo6z/	Nuance confirms its voice technology is behind Apple's Siri
/y5w9t/	Apple Special Event 2011 - iPhone 4S Introduction - YouTube
/dm34o/	Vista's Voice Recognition Stammers - PCWorld
/x4qq1/	How to configure and use Text-to-Speech in Windows XP and in Windows Vista
/5s2wx/	Blind Faith: A Decade of Apple Accessibility – Maccessibility
/btadv/	Enter the Dragon - MIT Technology Review
/wl4f4/	Issues in Practical Large Vocabulary Isolated Word Recognition: The IBM Tangora System - SpringerLink
/buaer/	IBM100 - Pioneering Speech Recognition
/-uc6y/	BBC - Future - The machines that learned to listen
/hr87g/	IBM Archives: IBM Shoebox
/vcdx4/	Amazon Echo vs. Google Home sales estimates: CHART - Business Insider Deutschland
/qczdb/	Google Home review: Don't buy it - Business Insider Deutschland
/b-919/	Forrester Data: Smart Home Devices Forecast, 2017 To 2022 (US)

/qmgz6/	CIRP Says 18 Million Smart-speakers Sold in Q4 2017 Bringing Installed Base to 45 Million - Voicebot
/6vf9-/	Amazon Says Echo Was the Best-Selling Product This Holiday - Sells Millions
/d3ayo/	Amazon Has Sold 3 Million Echos In U.S. - Twice
/s9a5x/	Amazon Is Killing Off the Fire Phone - Fortune
/xcw3j/	Why Amazon called its smart assistant Alexa - Business Insider Deutschland
/hyjzb/	A smartwatch just saved a man from a heart attack - World Economic Forum
/m1omn/	Global smartwatch unit sales 2014-2018 - Statistic
/gmior/	NXP Healthcare
/tbnkk/	IoT Provides Affordable Earthquake Early Warning to Communities - HuffPost
/jjzjo/	Smart Bins in Dublin - Smart Dublin
/ce3pi/	Amazon Announces Program with Auburn RFID Lab
/hhqh-/	BBC NEWS - UK - England - Bristol - Ants' home search habit uncovered
/xdwio/	Smart-speakers are the fastest-growing consumer tech; shipments to surpass 50 million in 2018 - Canalys
/8qa8c/	Analyst firm: Google Home gains ground on Amazon Echo, now 44M total devices sold
/mbhpx/	CES 2018: What Does Samsung's Big Bet on IOT Mean for You? - Fortune
/-tbk1/	Gartner Says 8.4 Billion Connected "Things" Will Be in Use in 2017, Up 31 Percent From 2016
/wkn49/	Internet Toaster, John Romkey, Simon Hackett

*All references were re-validated in February 2018.

