

**UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA
ALEXANDRIA DIVISION**

DIALECT, LLC,

Plaintiff,

v.

AMAZON.COM, INC., and
AMAZON WEB SERVICES, INC.,

Defendants.

No. 1:23-cv-581-TSE-JFA

JURY TRIAL DEMANDED

**AMENDED COMPLAINT FOR PATENT INFRINGEMENT AND DAMAGES
AND DEMAND FOR JURY TRIAL**

Plaintiff Dialect, LLC (“Dialect” or “Plaintiff”) files this Amended Complaint¹ for Patent Infringement and Damages against Amazon.com, Inc.; and Amazon Web Services, Inc. (“AWS”) (collectively, “Amazon” or “Defendants”) and alleges as follows:

INTRODUCTION

1. The novel inventions disclosed in U.S. Patent Nos. 7,693,720 (the “720 Patent”); 8,015,006 (the “006 Patent”); 8,140,327 (the “327 Patent”); 8,195,468 (the “468 Patent”); 9,031,845 (the “845 Patent”); 9,263,039 (the “039 Patent”); 9,495,957 (the “957 Patent”) (collectively, the “Asserted Patents”) in this matter were invented by VoiceBox Technologies (“VoiceBox”). VoiceBox was a key pioneer in the fields of voice recognition technology and natural language understanding (“NLU”). These technologies power a wide variety of consumer electronics and provide key functionality for smart phones, tablets, TVs, and Internet of Things

¹ Attached hereto as Appendix A is a redlined version of the complaint detailing the amendments made.

(“IoT”) devices. VoiceBox spent more than a decade developing and building key early NLU inventions producing one of the most valuable portfolios of technology according to the Institute of Electrical and Electronics Engineers (“IEEE”), including the Asserted Patents. In recognition of its many innovations, the U.S. Patent & Trademark Office awarded and issued the Asserted Patents to VoiceBox. The Asserted Patents in this case are the result of this substantial investment and research and were fundamental to the development of voice commerce technology.

2. Over the years, the inventions claimed in the Asserted Patents have been licensed to key companies in the industry.

3. Despite opportunities to do so and its knowledge of the Asserted Patents, Amazon has never licensed the Asserted Patents. VoiceBox’s opportunities to promote and build a business based on these patents were thwarted when Amazon introduced the infringing Echo and other Alexa Products² and used its enormous size and clout to poach dozens of VoiceBox Technologies’ engineers and scientists.

THE PARTIES

4. Plaintiff is the current owner and assignee of the Asserted Patents.

² “Alexa Products” collectively refers to Amazon’s Alexa virtual assistant and offerings that include Alexa, including the Echo product line (such as Echo 1st Gen., Echo 2nd Gen., Echo Dot 1st Gen., Echo Dot 2nd Gen., Echo Dot 3rd Gen., Echo Dot Kids Edition, Echo Show 1st Gen., Echo Show 2nd Gen., Echo Show 5, Echo Spot, Echo Plus 1st Gen., Echo Plus 2nd Gen., Echo Auto, and Echo Look), Amazon’s Alexa apps, Music apps, and Shopping apps on a smartphone or other mobile device, Amazon’s Alexa cloud, Alexa Voice Services, and Amazon.com website, and any other device, app, or instrumentality that includes, provides access to, or works with Alexa (such as Amazon Tap, Amazon Dash Wand, Echo Wall Clock, AmazonBasics Microwave, Amazon SmartPlug, Amazon Fire TV Sticks, Amazon Fire TVs, Amazon Fire TV Cubes, and Amazon Fire and Fire HD tablets) as well as software, hardware, and cloud infrastructure associated with any of the foregoing. Plaintiff reserves the right to expand upon or otherwise modify the above list during the discovery process in this case and is in no way limiting the scope of the accused products to what is currently listed.

5. Plaintiff is a Texas limited liability company with its principal place of business located at 133 E. Tyler St., Longview, TX 75601-7216.

6. Amazon.com, Inc. is a corporation incorporated in Delaware and has a principal place of business at 410 Terry Avenue North, Seattle, WA, 98109. Amazon.com, Inc. is the ultimate parent company of various subsidiaries, including AWS, which are responsible for making, using, selling, offering for sale and/or importing Alexa Products. Amazon maintains a regular and established place of business in this District through multiple permanent physical facilities, including in particular Amazon’s second headquarters that is located in this District. Amazon’s second headquarters is located in the National Landing neighborhood of Arlington, Virginia, which includes at least the leased properties at 241 18th St. S., 1770 Crystal Dr., 2100 Crystal Dr., 2345 Crystal Dr., and 1800 S. Bell St., Arlington, Virginia and included at least two additional office buildings that were developed for and have been delivered to Amazon for occupancy.³

7. AWS is a corporation incorporated in Delaware and has a principal place of business at 410 Terry Avenue North, Seattle, WA, 98109. AWS provides the cloud computing platforms that power one or more Alexa Products or are for use with one or more Alexa Products.⁴ AWS is registered to do business in Virginia.⁵

Entity ID	Entity Name	Name Type	Entity Type	Series LLC	Principal Office Address	RA Name	Status
T0395192	AMAZON WEB SERVICES LLC	Legal Name	Limited Liability Company	No	1200 12TH AVE STE 1200, SEATTLE, WA, 98144 - 0000, USA	CORPORATION SERVICE COMPANY	Inactive
E1918947	AMAZON WEB SERVICES, INC.	Legal Name	Stock Corporation		410 Terry Ave N, Seattle, WA, 98109 - 5210, USA	CORPORATION SERVICE COMPANY	Active

³ See Presentation by Joe Chapman, Director, Amazon Global Real Estate & Facilities, PenPlace, Amazon Arlington HQ, Site Plan Review Committee: SPRC #4, February 10, 2022 at 3, available at https://www.arlingtonva.us/files/sharedassets/public/projects/documents/site-planprojects/penplace/2022_02-07-sprc-4-deck.pdf.

⁴ <https://aws.amazon.com/blogs/aws/majority-of-alex-now-running-on-faster-more-cost-effective-amazon-ec2-infl-instances/>

⁵ <https://cis.scc.virginia.gov/EntitySearch/Index>

8. On information and belief, Amazon directly and/or indirectly develops, designs, manufactures, distributes, markets, offers to sell and/or sells infringing products and services in the United States, including in the Commonwealth of Virginia and within the Eastern District of Virginia, and otherwise direct infringing activities to this District in connection with their products and services as set forth in this Complaint.

JURISDICTION

9. This civil action arises under the Patent Laws of the United States, 35 U.S.C. § 1 *et seq.*, including without limitation 35 U.S.C. §§ 271, 281, 283, 284, and 285. Accordingly, this Court has subject matter jurisdiction under, *inter alia*, 28 U.S.C. §§ 1331 and 1338(a).

10. This District has general and specific personal jurisdiction over Amazon because Amazon has committed acts, directly or through intermediaries, in this District, giving rise to this action; is present in and transacts and conducts business in this District and the Commonwealth of Virginia; and transacts and conducts business with residents of this District and the Commonwealth of Virginia.

11. Plaintiff's causes of action arise, at least in part, from Amazon's contacts with and activities in this District and the Commonwealth of Virginia.

12. Amazon has infringed the Asserted Patents within this District and the Commonwealth of Virginia by making, using, selling, offering for sale, and/or importing in or into this District and elsewhere in the Commonwealth of Virginia, products that infringe the Asserted Patents, including, without limitation, products that practice the claimed methods of the Asserted Patents. Amazon, directly and through intermediaries, makes, uses, sells, offers for sale, imports, ships, distributes, advertises, promotes, and/or otherwise commercializes such infringing products in or into this District and the Commonwealth of Virginia. Amazon regularly conducts and solicits

business in, engages in other persistent courses of conduct in, and/or derives substantial revenue from goods and services provided to residents of this District and the Commonwealth of Virginia.

13. This Court has general and specific personal jurisdiction over Amazon. Amazon has substantial contacts with the forum as a consequence of actively working on and establishing its second headquarters in Virginia and in this District, and Amazon conducts substantial business in Virginia. Amazon states in its most recent 10-K filing to the SEC for fiscal year ended December 31, 2022, dated February 3, 2023, that “[w]e own and lease our corporate headquarters in Washington’s Puget Sound region and Arlington, Virginia.”

14. Amazon sells, makes, uses, and offers for sale its products and services, including products and services that infringe the Plaintiff’s Asserted Patents, within the Commonwealth of Virginia, including to customers in Virginia. Amazon further operates data centers in Ashburn, Virginia⁶ and has announced \$35 billion in additional investment in data centers in Virginia through 2040.⁷

15. A Court in this District has observed in April 2020:

It must be said that Amazon is nothing if not ubiquitous in the United States. Furthermore, after considering 238 cities, Amazon chose Arlington in the Eastern District of Virginia as the location for its HQ2 and will invest \$2.5 billion and 25,000 jobs in the undertaking. As such, Amazon cannot in good faith represent to the Court that E.D. Va. is an undesirable or inconvenient location to operate and do business. Litigating should not be an additional significant strain.

⁶ <https://www.datacenters.com/amazon-aws-ashburn>

⁷ <https://apnews.com/article/technology-data-management-and-storage-amazoncom-inc-virginia-business-c75df1f34069b09549fe15c99335b8fb>

Maghula, Ltd. v. Amazon.com, Inc., No. 1:19-cv-01570, ECF No. 52 at 32-33 (E.D. Va. Apr. 9, 2020).

16. Amazon has further stated that it is specifically hiring for its Alexa teams at HQ2 located in Alexandria, Virginia. At least as early as September 2019, Amazon listed 25 open positions for the Alexa team at its HQ2.⁸ Amazon's website stated that as of September 2021, Amazon was still growing its teams at HQ2, "including: software development engineers, technical sales representatives, program managers, and solutions architects on teams across" divisions including "Alexa."⁹ In 2022, Amazon publicly promoted the "Alexa Teams in HQ2," including multiple "Software Development Managers" that moved to HQ2 in Alexandria.¹⁰ Amazon's description of this video describes taking "a dive into the teams that make up Alexa at Amazon's new second headquarters in Arlington, Virginia. In this video, meet different managers leading the Alexa teams in HQ2 to learn more about what they're building and their favorite things about living in Northern Virginia."¹¹ Amazon further stated in the same description that "Alexa is hiring on all teams."¹² As of December 2022, Amazon reported that students from the University of the District of Columbia and the University of Maryland College Park "presented their Alexa Skills final projects" to "Alexa employees, and Amazon leaders at the company's headquarters in Virginia (HQ2)."¹³

⁸ <https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/amazon-hiring-for-cloud-services-alexa-products-at-hq2-in-arlington-va-53798578>

⁹ <https://www.aboutamazon.com/news/amazon-offices/the-next-chapter-for-hq2-sustainable-buildings-surrounded-by-nature>

¹⁰ <https://www.youtube.com/watch?v=vU2szgs2M7c>

¹¹ *Id.*

¹² *Id.*

¹³ <https://www.amazon.science/academic-engagements/amazonnext-program-hosts-final-project-presentations-at-virginia-hq2>

17. This Court also has personal jurisdiction over Amazon, in part, because Amazon does continuous and systematic business in this District, including by providing infringing products and services to the residents of this District that Amazon knew would be used within this District, and by soliciting business from the residents of this District.

18. Amazon has pled in legal filings that venue in this District is proper. In a 2020 action filed by Amazon in this District, Amazon asserted that venue was proper under 28 U.S.C. § 1391(b) because, among other things “it is a district in which Plaintiff [Amazon] maintains headquarters and/or substantial business operations”. *Amazon.com, Inc. v. WDC Holdings LLC*, No. 1:20-cv-484, ECF No. 1, ¶ 26 (E.D. Va. Apr. 27, 2020).

19. Venue is proper in this Court under 28 U.S.C. §§ 1391 and 1400(b).

BACKGROUND

20. In 2001, three brothers, Mike, Rich, and Bob Kennewick founded VoiceBox to bring NLU to a wide array of computer applications. They recognized that the typical computer speech-recognition systems forced human operators to adhere to a limited number of rigid speech prompts, typically through verbal menus of a so-called “Command and Control” system. These rigid prompts limited how systems were used and inhibited the widespread adoption of speech-recognition systems. The brothers believed that VoiceBox could become the first company to improve voice recognition systems to enable people to naturally and effectively interact with computer speech systems.

21. From its inception, VoiceBox engaged in intense research efforts to develop its NLU technology. As part of these efforts, VoiceBox Technologies achieved a significant milestone when it developed an early prototype called “Cybermind.” As demonstrated on Seattle-area

television news,¹⁴ Cybermind was a voice-controlled speaker that could provide weather, recipes, sports scores, calendar updates, or play a song.



22. On information and belief, consumer focus groups being introduced to VoiceBox conversational voice technology described it as “cool,” “unbelievable,” “so fast,” “it makes you feel like you’re in the future already,” and “I feel like I’m in the Jetsons.”¹⁵

23. Throughout its research and development efforts, VoiceBox realized that its technology could be deployed in a wide range of applications from connected home to mobile personal assistants.

24. VoiceBox’s groundbreaking work did not go unrecognized. By January 2012, VoiceBox was a leader in NLU and conversational voice technology. Leading companies throughout the world, including Toyota, Lexus, TomTom, Pioneer, Chrysler, Dodge, and Magellan used VoiceBox’s award-winning and patented contextual speech technology. VoiceBox had

¹⁴ <https://www.youtube.com/watch?v=DDcRyPnvWhw>

¹⁵ <https://www.youtube.com/watch?v=WCOGNnH-Bws>

software applications that ran on smart speakers, in-car systems, smartphones, smart TVs, computers, tablets, e-readers, and personal navigation devices.

25. On information and belief, in 2013 the Institute of Electrical and Electronics Engineers (“IEEE”) ranked VoiceBox number 13 in patent power for the computer software industry, ranking between SAP AG and Sony Computer Entertainment Inc.

IEEE Patent Power Rankings

Computer Software

Rank	Company / Organization	Country of Headquarters	2012 U.S. Patents	Pipeline Growth Index	Pipeline Impact	Self-Citations (%)	Adjusted Pipeline Impact	Pipeline Generality	Pipeline Originality	Pipeline Power
1	Microsoft Corp.	United States	2665	1.14	1.07	0.22	1.07	1.19	1.01	3909.67
2	VMware Inc.	United States	106	1.89	3.07	0.16	3.07	3.02	1.06	1966.63
3	Citrix Systems Inc.	United States	112	1.56	2.83	0.27	2.83	2.62	1.12	1441.86
4	Symantec Corp.	United States	379	1.34	1.54	0.16	1.54	1.63	1.01	1297.47
5	Digimarc Corp.	United States	94	0.9	5	0.88	2.08	4.8	1.12	944.96
6	Oracle Corp.	United States	913	0.92	1	0.12	1	1.13	0.98	930.36
7	CommVault Systems Inc.	United States	52	1.3	5	0.88	2.1	5	1.22	866.39
8	Cadence Design Systems Inc.	United States	158	1.15	2.45	0.15	2.45	1.76	0.89	699.87
9	Adobe Systems Inc.	United States	332	1.13	1.14	0.18	1.14	1.24	1	526.88
10	Rovi Corp.	United States	97	1.47	1.85	0.25	1.85	1.91	1.02	514.34
11	TeleCommunication Systems Inc.	United States	57	1.36	2.35	0.42	2.06	2.52	1.12	451.72
12	SAP AG	Germany	601	1.1	0.74	0.23	0.74	0.85	1.02	424.91
13	Voicebox Technologies Inc.	United States	11	1.83	5	0.65	3.26	5	1.29	423.56
14	Sony Computer Entertainment Inc.	Japan	220	1.33	1.1	0.36	1.03	1.26	1.08	409.7
15	Bally Technologies Inc.	United States	98	1.78	1.46	0.38	1.35	1.83	0.9	388.27
16	Smith Micro Software Inc.	United States	18	3	2.8	0.17	2.8	2.41	0.97	353.44
17	McAfee Inc.	United States	84	1.33	2.02	0.39	1.85	1.61	1.04	347.02
18	Nuance Communications Inc.	United States	160	1.15	1.19	0.3	1.19	1.56	1.02	345.99
19	Synopsys Inc.	United States	148	0.95	1.61	0.08	1.61	1.17	1.06	280.46
20	Infosys Ltd.	India	29	1.93	2.52	0.04	2.52	1.75	1.02	253.69

Source: IEEE Spectrum Patent Power 2013

26. After learning about VoiceBox Technologies’ technology, Toyota hired it to build a sophisticated NLU speech interface for its Lexus automobiles. VoiceBox Technologies built the voice and NLU capability for Toyota’s award-winning Entune multimedia system. As part of the development effort of an NLU interface for Lexus, VoiceBox demonstrated a personal assistant called “Alexus” that showcased the power of its Conversational Voice technology.

27. On information and belief, the VoiceBox “Alexus” concept was introduced to the public more than six months before Amazon announced “Alexa.”¹⁶

AMAZON’S KNOWLEDGE OF THE ASSERTED PATENTS

28. In 2011, VoiceBox contacted Amazon to explore a potential business relationship where VoiceBox would provide core NLU services to Amazon. Amazon’s corporate development department expressed interest and asked for “company and/or product overview slides” to facilitate an October 7, 2011 teleconference. In response, VoiceBox provided Amazon with a presentation that described its award-winning patented technology and explicitly referred to VoiceBox’s patented technology, including the following slides:



¹⁶ <https://www.youtube.com/watch?v=ymBJcuTSHJw> and <https://www.youtube.com/watch?v=w046TcaQu40>

**Pioneer and Leader in
Conversational Voice Applications**

Gartner
Cool Vendor
User Interface Innovation

“ Mindblowingly
Impressive ”

“ VoiceBox...makes
Applications and
devices understand ”

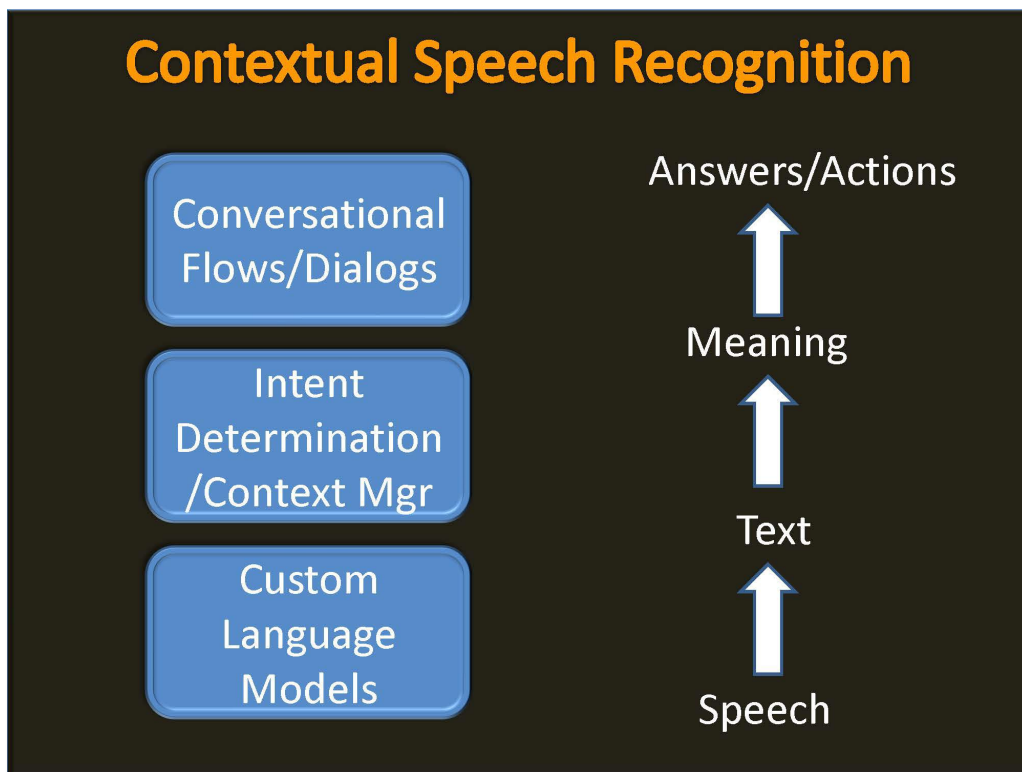


International
Academy of
Science

Outstanding Technology
of the Year, Nominee

THE WALL STREET JOURNAL.
Software Technology Innovator

“ Lets people use
everyday language ”



Key Customers



Embedded Applications



- Hands Free Dialing
- Music Search
- XM Radio
- XM Traffic
- XM Sports, Stocks, News
- Country Wide Address Entry
- POI Search
- Map Control
- Climate Control

Cloud Applications and Connected Services APIs



- Local Search
- SMS Messaging
- EMAIL
- Web Search
- Music Search
- Movies, Stocks, Sports,
- Traffic
- Hands Free Dialing
- Internet Radio
- On Demand Music
- Navigation/VDE

Connected TV

<http://www.youtube.com/watch?v=1j1hPdIQgoQ>



“What’s Ichiro’s batting average?”

“Show it on the TV”

“Search NetFlix for a Clint Eastwood Movie”

Tablets, Laptops and SmartPhones (CyberMind in the Cloud)



Personified User Experience

Personal Digital Assistant

Web Search with Intent
Determination

Internet Content (music, videos and
apps)

Skype Calling

Personal Reports

The slide deck is attached as Exhibit 8.

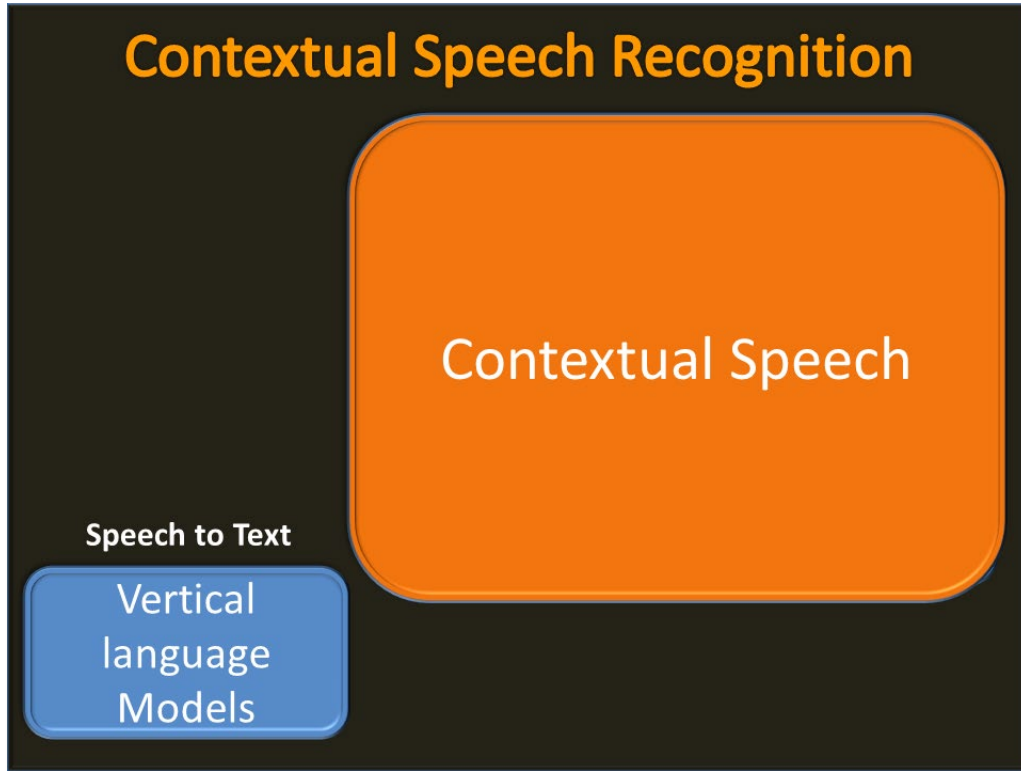
29. Amazon was so impressed by the technology VoiceBox presented that its representative emailed VoiceBox the next business day to invite VoiceBox to visit Amazon's offices on October 19, 2011. That meeting was with Douglas Booms, Amazon's Vice President of Worldwide Corporate Development, as well as engineers and product/business development members of Amazon's devices and digital teams.

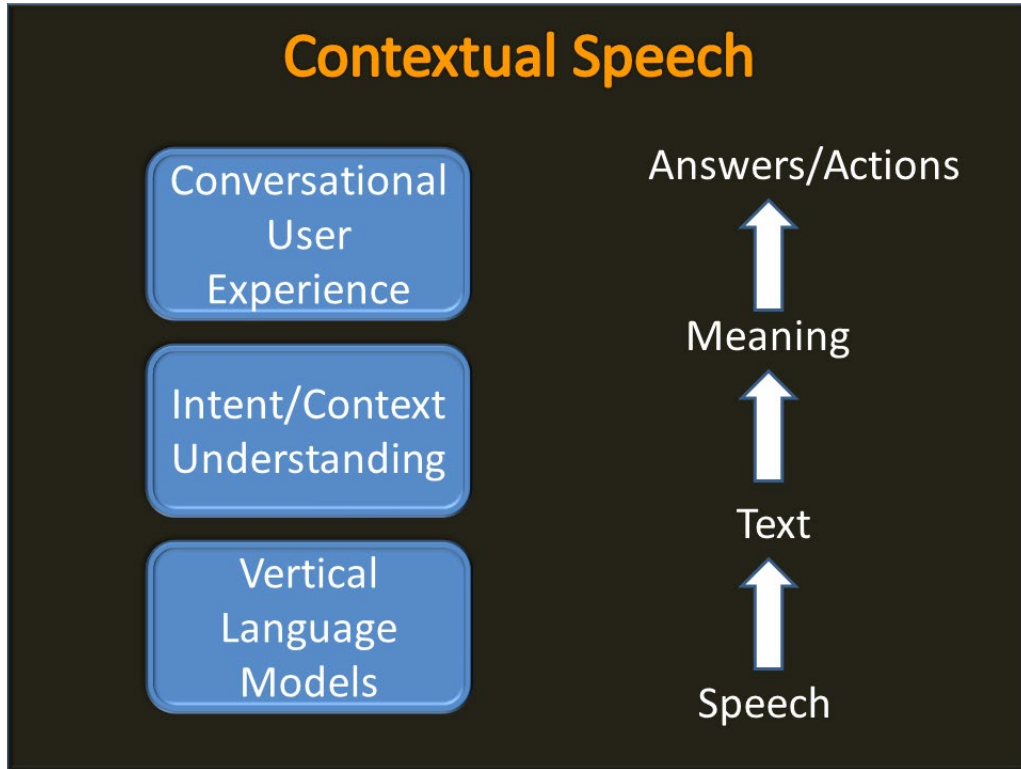
30. Two days after the meeting at Amazon, on Friday October 21, 2011, Amazon contacted VoiceBox and asked to visit the company's office in the following week. VoiceBox agreed to host Amazon's personnel for a meeting at VoiceBox office on October 26, 2011.

31. On information and belief, during the October 26, 2011 meeting, VoiceBox presented a deck of 42 slides to the visiting Amazon personnel. The slides provided even more detail about VoiceBox's patented technology and informed Amazon that VoiceBox had 12 patents

at the time with an additional 14 pending applications, including a number of the Asserted Patents.

The slides included the following:



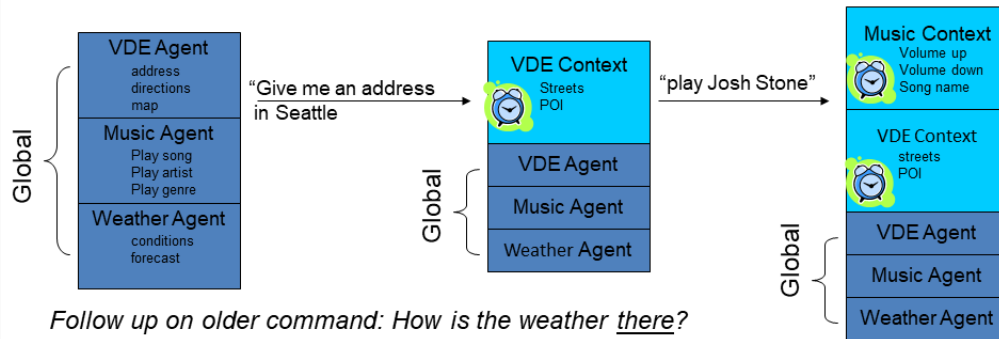


How It Works: Speech To Text

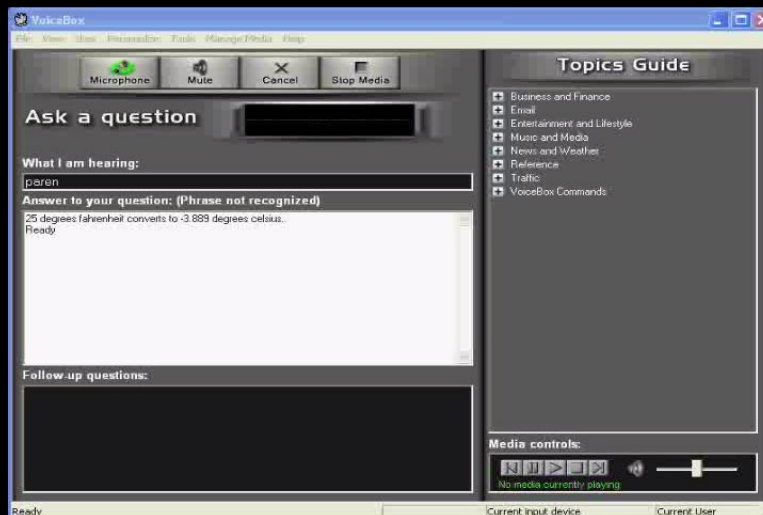



How It Works: Context Understanding

- Global Concepts → expected requests at any point in time
- Contextual Concepts → follow up requests with implicit information



Cybermind Digital Assistant Conversational User Experience





VoiceBox Patents

- VoiceBox has been awarded 12 patents for contextual and conversational speech technologies
- VoiceBox has an additional 14 pending applications
- Our most recent patent (at right) for processing natural language with context specific domain agents also incorporates personification

US80001506682

(12) **United States Patent**
Kennewick et al.

(10) **Patent No.:** US 8,015,006 B2
(45) **Date of Patent:** Sep. 6, 2011

(54) **SYSTEMS AND METHODS FOR PROCESSING NATURAL LANGUAGE SPEECH UTTERANCES WITH CONTEXT-SPECIFIC DOMAIN AGENTS**

(75) **Inventors:** Robert A. Kennewick, Seattle, WA (US); David Locker, Redmond, WA (US); Michael R. Kennewick, Sr., Bellevue, WA (US); Michael R. Kennewick, Jr., Bellevue, WA (US); Richard Kennewick, Woodinville, WA (US); Tom Freeman, Mercer Island, WA (US)

(73) **Assignee:** VoiceBox Technologies, Inc., Bellevue, WA (US)

(*) **Notice:** Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 617 days.

(21) **Appl. No.:** 12/130,397

(22) **Filed:** May 30, 2008

(65) **Prior Publication Data**
US 2008/0235023 A1 Sep. 25, 2008

Related U.S. Application Data

(63) Continuation of application No. 10/452,147, filed on Jun. 3, 2003, now Pat. No. 7,338,209.

(60) Provisional application No. 60/384,388, filed on Jun. 3, 2002.

(51) **Int. Cl.**
G10L 15/08 (2006.01)
G10L 15/22 (2006.01)
G10L 1/00 (2006.01)
G06F 17/27 (2006.01)
G06F 17/30 (2006.01)

(52) **U.S. Cl.** 704/236; 704/246; 704/257; 707/709; 707/723; 707/771

(58) **Field of Classification Search** 704/236; 704/246; 257; 270; 244; 707/709; 723; 728; 707/771
See application file for complete search history.

(56) **References Cited**
U.S. PATENT DOCUMENTS
4,430,669 A 2/1984 Chong 358/122 (Continued)


FOREIGN PATENT DOCUMENTS
EP 1 320 043 A2 6/2003 (Continued)


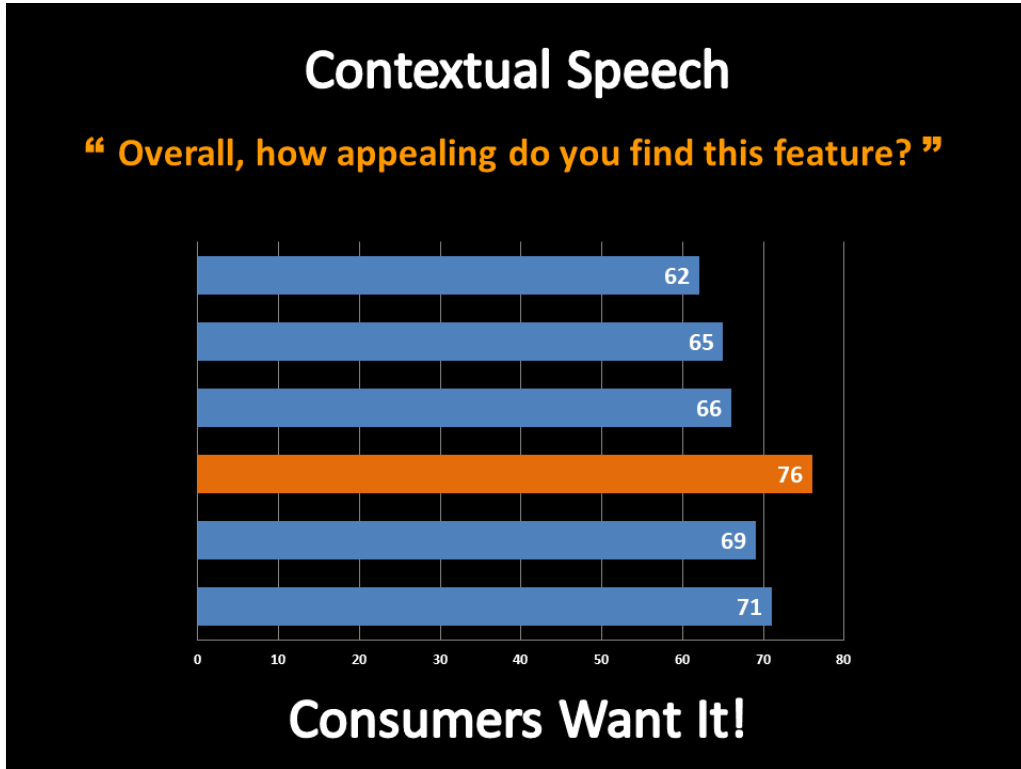
OTHER PUBLICATIONS
Reuters, "IBM to Enable Honda Drivers to Talk to Car", Charles Schwab & Co., Inc., Jul. 28, 2002, 1 page. (Continued)

Primary Examiner — Martin Lerner
[74] Attorney, Agent, or Firm — Pillsbury Winthrop Shaw Pittman LLP

(57) **ABSTRACT**
Systems and methods for receiving natural language queries and/or commands and executing the queries and/or commands. The systems and methods overcomes the deficiencies of prior art speech query and response systems through the application of a complete speech-based information query, retrieval, presentation and command environment. This environment makes significant use of context, prior information, domain knowledge, and user specific profile data to achieve a natural environment for one or more users making queries or commands in multiple domains. Through this integrated approach, a complete speech-based natural language query and response environment can be created. The systems and methods creates, stores and uses extensive personal profile information for each user, thereby improving the reliability of determining the content and presenting the expected results for a particular question or command.

11 Claims, 7 Drawing Sheets







- a. Currently shipping embedded in all Lexus cars and Many Toyota brands since 2007
- b. Currently shipping connected-car cloud-services
- c. Will ship both embedded software and cloud-services with 1.5 million cars a year in North America from 2013 through 2016 (Revenue of over \$7M per year by 2013)
- d. Toyota is evaluating VoiceBox for Global Markets including Japan and China (additional 4 million cars)



- a. Shipping in Chrysler and Dodge Brands in North America**
- b. Chinese language for exports in 2012**
- c. Cloud Applications and Services in 2013/2014 (Revenue of over \$5M per year by 2013)**



Longstanding Relationship with Renault in Europe


- a. Will ship in Q1 2012 in 9 Languages**
- b. VBT is providing input to 2014/2015 specification**
- c. Key new feature is Cloud**
- d. 1.5 million cars a year in 2014/2015 for Hybrid Embedded and Cloud Services (Revenue of over \$7M a year by 2013)**




1. TomTom is shipping in all mid-range and high-end PNDs in North America to strong consumer reviews
2. European languages for PND will ship in Q1 2012 (Revenue of over \$5M per year by 2013)
3. TomTom Automotive will ship in Renault cars in June of 2012
4. TomTom Automotive will ship with Mazda cars in Q1 2012
5. TomTom evaluating VoiceBox Connected Services for 2012

Tier 1 Integrators







“Now available in Dealer Lots”



Pioneer has been shipping with VoiceBox since 2007

- a. Shipping in the aftermarket in-dash navigation system called AVICs
- b. Will ship Cloud Services product on iPhone, iPad and Android App Stores in Q4 (Ad Model)
- c. Will ship connected product on Scion in 2012
- d. Will ship with Audio Head Product with Toyota CY 13/14 Head Units



Connected TV



“What’s Ichiro’s batting average?”

“Show it on the TV”



Tablets and SmartPhones (SIRI Alternative)



Navigation

Music and Videos

Shopping

Voice Dialing (Phone and Skype) and Messaging

EMAIL

Calendar/Reminders

Social Networking

Web Search

Smart Bots/Wolfram

Voice Calculator

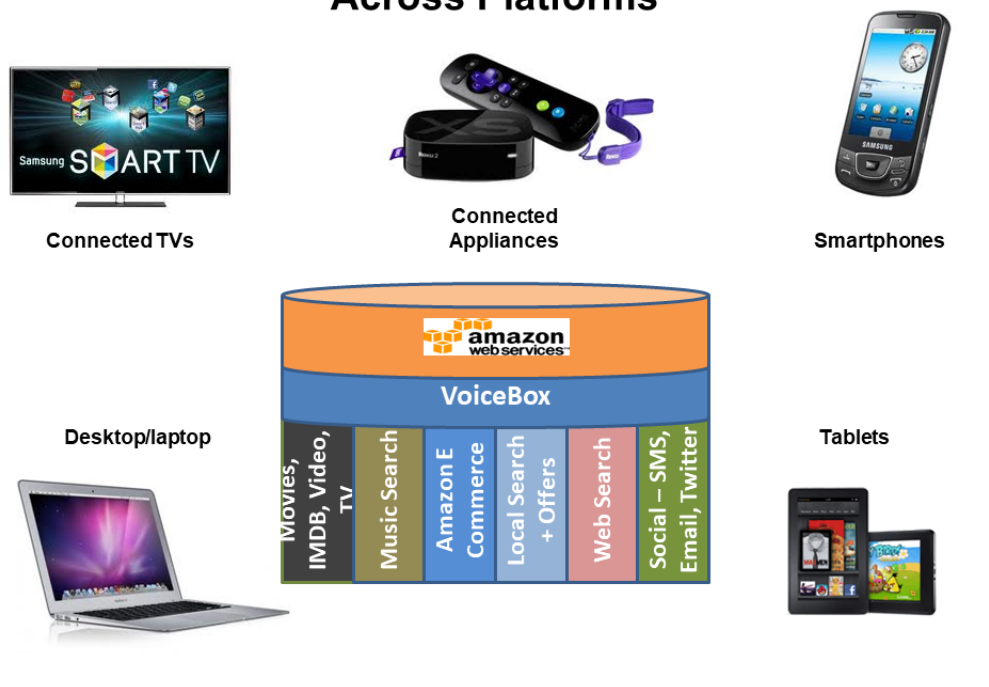
Dictionary/Wikipedia

Personal Reports

Example Voice Use Models for Amazon Kindle



Voice Services on Amazon Web Services Across Platforms



The slide deck is attached as Exhibit 9.

32. On information and belief, some of the Amazon personnel involved in the 2011 meetings became technical leaders for Amazon's Alexa Products while others became high-level executives with close working relationships with the senior leadership of Amazon.

33. Amazon did not pursue a business relationship with VoiceBox. Instead, on information and belief, Amazon decided to build its Alexa Products—without telling VoiceBox or asking permission to use VoiceBox's patented technology.

34. In 2014, Amazon announced the launch of Alexa, a virtual assistant, along with the first-generation Echo product, a smart speaker. Amazon's Alexa and first-generation Echo product were strikingly similar to the patented technology that VoiceBox showed Amazon in 2011.

35. In 2016, Amazon abruptly hired Philippe Di Cristo, who was VoiceBox's Chief Scientist. While at VoiceBox, Mr. Di Cristo gained knowledge of the company's voice technology and had full access to VoiceBox's intellectual property. As Mr. Di Cristo explains on his LinkedIn Page, he had worked on an "Amazon Echo- like system" while at VoiceBox.¹⁷

¹⁷ <https://www.linkedin.com/in/pdicristo/>



Speech and Voice Technology Leader

Amazon

Oct 2016 - Jan 2022 · 5 yrs 4 mos

Seattle, WA

I am building a team taking to take broad responsibilities in building the ultimate Personal Assistant experiences.

...

...see more



Chief Scientist, Vice-President Speech and NLU R&D

VoiceBox Technologies

Jul 2003 - Sep 2016 · 13 yrs 3 mos

Bellevue, WA

I joined the company in its infant stage as employee number 9. I cover broad responsibilities including strong technology lead, top-level management and architect/SE.

Some of my functions include leading an international team of over 100 engineers and scientists designing and implementing new Speech Recognition and Natural Language Understanding algorithms to enhance company's IP, researching and identifying technology trends to define mission critical strategic and tactical roadmap and patents, researching and evaluating worldwide potential technology partners, communicating company's technology IP to investors, executives partners, customers and stakeholders as well as driving competitor analysis and managing Speech partner relations.

Spearheaded and delivered major projects such as Toyota Connected Services for Lexus Enform/Entune, Samsung S-Voice (now on all S6 and S7 phones), as well as worldwide embedded car products (Chrysler, Renault, and more).

Designed and implemented in 2006, VoiceBox "Cybermind", the Amazon Echo-like system. Look for it on YouTube! (<http://bit.ly/cybermind>). In 2013, IEEE ranked VoiceBox Technologies #13 in Patent Power for the Computer Software industry.

36. On information and belief, Mr. Di Cristo played an active role in soliciting additional VoiceBox employees to join Amazon. For example, shortly after Mr. Di Cristo joined Amazon from VoiceBox, Amazon ramped-up its efforts to recruit VoiceBox employees.

37. On January 10, 2017, Amazon hosted an invite only networking event for VoiceBox employees, with the stated purpose of allowing VoiceBox employees to talk about opportunities at Amazon. For this event, Amazon rented out Seastar, the premier seafood restaurant near VoiceBox's office.

38. On January 17, 2017, after Amazon's rampant poaching of VoiceBox employees came to light, Mike Kennewick, CEO of VoiceBox, sent a letter to Jeff Bezos, CEO of Amazon.

THE ASSERTED PATENTS

39. The VoiceBox inventions contained in the Asserted Patents in this case relate to groundbreaking improvements to voice recognition and NLU and have particular application in consumer electronics such as smart phones, tablets, and IoT devices.

U.S. PATENT NO. 7,693,720

40. On April 6, 2010, the U.S. Patent Office duly and legally issued the '720 Patent, entitled "Mobile Systems And Methods For Responding To Natural Language Speech Utterance." A true and correct copy of the '720 Patent is attached hereto as **Exhibit 1**.

41. Dialect is the owner and assignee of all right, title, and interest in and to the '720 Patent, including the right to assert all causes of action arising under the '720 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

42. As described in the '720 Patent, "many natural language questions may be ambiguous or subjective." '720 Patent, at 2:8-10. While human beings are adept at interpreting these types of communication, they are not easily "transformed to machine processable form," resulting in a uniquely challenging technological problem to building natural-language speech-recognition interfaces for computers and vehicles. *Id.* at 2:6-12.

43. Thus, "creating a natural language speech interface that is suitable for use in the vehicular environment" was a difficult problem. *Id.* at 1:34-36. "Conventional computer interfaces" were "generally not suitable for a vehicular environment, owing to the speed of interaction and the inherent danger and distraction" involved. *Id.* at 1:26-31. The challenges for existing natural language speech interfaces included the need to "accommodate commands and queries from a wide range of domains and from many users with diverse preferences and needs." *Id.* at 1:36-38. Existing speech interfaces required queries and commands to be "highly structured"

in a way that is “not inherently natural to the human user,” causing natural language inputs to be “fundamentally incompatible” with existing solutions. *Id.* at 1:49-54.

44. As acknowledged in the ’720 Patent, despite improvements in the accuracy of existing speech recognition and the application of natural language processing in “parsing of speech queries,” no existing system overcame the “significant barriers to creating of a complete natural language speech-based query and response environment.” *Id.* at 1:56-67. The obstacles addressed by the inventions of the ’720 Patent included the fact that “most natural language queries and commands are incomplete in their definition” and “some questions can only be interpreted in the context of previous questions, knowledge of the domain, or the user’s history of interests and preferences.” *Id.* at 2:1-8.

45. To solve these problems, the inventors of the ’720 Patent conceived of novel software techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing computer or vehicle systems. For example, claim 1 recites a novel combination of a speech recognition engine that uses data that includes “a plurality of dictionary and phrase entries that are dynamically updated based on at least a history of a current dialog and prior dialogs associated with the user,” and a parser that interprets recognized words or phrases and uses data received from a plurality of domain agents, determines a context for a natural language speech utterance, selects a domain agents based on the context, and transforms the recognized words or phrases into a question or a command formulated in a grammar that the domain agent uses to process the formulated question or command. The recited combination also includes an agent architecture that communicatively couples domain agents to a system agent in a vehicle system:

1. A mobile system responsive to a user generated natural language speech utterance, comprising:

- a speech unit connected to a computer device on a vehicle, wherein the speech unit receives a natural language speech utterance from a user and converts the received natural language speech utterance into an electronic signal; and
- a natural language speech processing system connected to the computer device on the vehicle, wherein the natural language speech processing system receives, processes, and responds to the electronic signal using data received from a plurality of domain agents, wherein the natural language speech processing system includes:
 - a speech recognition engine that recognizes at least one of words or phrases from the electronic signal using at least the data received from the plurality of domain agents, wherein the data used by the speech recognition engine includes a plurality of dictionary and phrase entries that are dynamically updated based on at least a history of a current dialog and one or more prior dialogs associated with the user;
 - a parser that interprets the recognized words or phrases, wherein the parser uses at least the data received from the plurality of domain agents to interpret the recognized words or phrases, wherein the parser interprets the recognized words or phrases by:
 - determining a context for the natural language speech utterance;
 - selecting at least one of the plurality of domain agents based on the determined context; and
 - transforming the recognized words or phrases into at least one of a question or a command, wherein the at least one question or command is formulated in a grammar that the selected domain agent uses to process the formulated question or command; and
 - an agent architecture that communicatively couples services of each of an agent manager, a system agent, the plurality of domain agents, and an agent library that includes one or more utilities that can be used by the system agent and the plurality of domain agents, wherein the selected domain agent uses the communicatively coupled services to create a response to the formulated question or command and format the response for presentation to the user.

'720 Patent at Claim 1.

46. Embodiments of these claimed elements are shown and described in the specification. For example, Figure 5 shows an overall diagrammatic view of the interactive natural language speech processing system according to one embodiment:

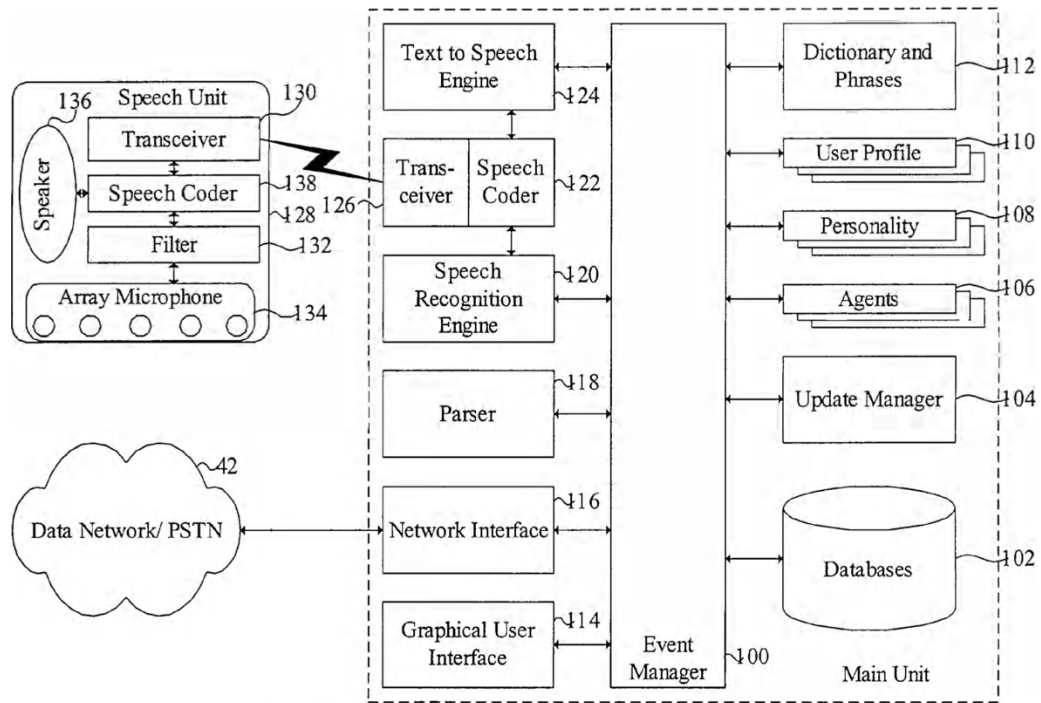


Figure 5. Speech Processing System Block Diagram

47. The specification of the '720 Patent describes how these claim elements help the overall system overcome the technical limitations of existing speech recognition systems. *See e.g., id.* at 21:61-22:57 (describing domain agents, system agents, and their interactions); 27:1-28:6 (describing the use of the speech recognition system and the dictionary and phrase entries); 28:25-29:50 (describing the interactions between system and domain agents in processing questions or commands).

48. In explaining the reasons or allowing the claims, the United States Patent and Trademark Office described how the closest existing prior art did not disclose or teach the claimed combination of inventive elements:

With respect to Claims 1 and 28, the prior art of record fails to explicitly teach or fairly suggest, either taken individually or in combination, a respective vehicle-based system and method as is set forth in these claims that features means/steps of receiving a natural

language speech input from a user, processing a natural language speech input using a speech recognition engine, and a specific parser, and a specific agent architecture in combination with the speech recognition engine including a plurality of dictionary and phrase entries that are dynamically updated based on at least combined histories from a present dialog and other prior dialogs and the agent architecture featuring a agent library that includes one or more utilities that can be used by the system agent and the plurality of domain agents.

'720 File History, Notice of Allowance (December 10, 2009) at 4 (attached as Exhibit 10).

U.S. PATENT NO. 8,015,006

49. On September 6, 2011, the U.S. Patent Office duly and legally issued the '006 Patent, entitled "Systems And Methods For Processing Natural Language Speech Utterances With Context-Specific Domain Agents." A true and correct copy of the '006 Patent is attached hereto as Exhibit 2.

50. Dialect is the owner and assignee of all right, title, and interest in and to the '006 Patent, including the right to assert all causes of action arising under the '006 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

51. As described in the '006 Patent, "[a] machine's ability to communicate with humans in a natural manner remains a difficult problem," in part because "machine-based queries (e.g., questions, commands, requests, and/or other types of communications) may be highly structured and are not inherently natural to the human user." '006 Patent at 1:38-41. Similarly, "[t]he fact that most natural language queries are incomplete in their definition is a significant barrier to natural human query-response interaction between humans and machines," and "many natural language questions are ambiguous or subjective," such that "the formation of a machine

processable query and returning of a natural language response may be difficult at best.” ‘006 Patent at 9:16-21.

52. Thus, while “speech recognition” (i.e., transcribing human speech into text) had “steadily improved in accuracy” and was “successfully used in a wide range of applications,” (id. at 1:46-48) simply translating uttered speech from a user into machine-readable text form, alone, did not and does not overcome the additional challenges of creating a natural language query and response system. Instead, existing systems were “generally unable to provide a complete environment for users to make natural language speech queries and receive natural-sounding responses” and “[t]here remain[ed] a number of significant barriers to creation of a complete natural language speech-based query and response environment.” Id. at 1:50-55.

53. To overcome these barriers, the inventors of the ‘006 Patent conceived of novel software techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing systems. The claimed invention “makes significant use of context, prior information, domain knowledge, and user specific profile data to achieve a natural environment for one or more users making queries or commands in multiple domains.” ‘006 Patent, Abstract. The inventions described and claimed in the ‘006 Patent overcome these challenges in various embodiments, for example by providing a system that uses domain agents to organize domain specific behavior and information. *Id.* at 2:53–3:7. The inventions in various embodiments further include a system that can “determine the user’s identity by voice and name for each utterance,” so that “[r]ecognized words and phrases may be tagged with this identity in all further processing” for security and other purposes. *Id.* at 16:60–17:4.

54. The novel features of the invention are recited in the claims. For example, Claim 5 of the ‘006 Patent recites a novel combination of parsing to determine a meaning and a context of

speech associated with a request involving a grammar by a domain agent, satisfying a predetermined confidence level, updating user specific vocabularies or dictionaries, and determining an identity of a user based on voice characteristics:

A method for processing natural language speech utterances with context-specific domain agents, comprising:

receiving, at a speech unit coupled to a processing device, a natural language speech utterance that contains a request;

recognizing, at a speech recognition engine coupled to the processing device, one or more words or phrases contained in the utterance using information in one or more dictionary and phrase tables;

parsing, at a parser coupled to the processing device, information relating to the utterance to determine a meaning associated with the utterance and a context associated with the request contained in the utterance, wherein the parsed information includes the one or more recognized words or phrases;

formulating, at the parser, the request contained in the utterance in accordance with a grammar used by a domain agent associated with the determined context, wherein formulating the request in accordance with the grammar used by the domain agent includes:

determining one or more required values and one or more optional values associated with formulating the request in the grammar used by the domain agent;

extracting one or more criteria and one or more parameters from one or more keywords contained in the one or more recognized words or phrases, wherein the parser extracts the one or more criteria and the one or more parameters using procedures sensitive to the determined context;

inferring one or more further criteria and one or more further parameters associated with the request using a dynamic set of prior probabilities or fuzzy possibilities; and

transforming the one or more extracted criteria, the one or more extracted parameters, the one or more inferred criteria, and the one or more inferred parameters into one or more tokens having a format compatible with the grammar used by the domain agent, wherein the one or more tokens include all the required values and one or more of the optional values associated with formulating the request in the grammar used by the domain agent;

processing the formulated request with the domain agent associated with the determined context to generate a response to the utterance; and

presenting the generated response to the utterance via the speech unit.

'006 Patent at Claim 5.

55. Embodiments of these claimed elements are shown and described in the specification. For example, Figure 1 shows an overall diagrammatic view of the interactive natural language speech processing system according to one embodiment:




FIG. 1

56. The specification of the '006 Patent describes how these claim elements help the overall system overcome the technical limitations of existing speech recognition systems. *See e.g., id.* at 10:56-12:18 (describing domain agents, system agents, and their interactions); 17:13-18:49 (describing the use of the speech recognition system and the dictionary and phrase entries, parser and domain agents to determine context and criteria); 18:50-21:25 (describing the interactions between system and domain agents in processing questions or commands).

57. In explaining the reasons for allowing the claims, the United States Patent and Trademark Office described how the closest existing prior art did not disclose or teach the claimed combination of inventive elements:

Independent Claim [5] is allowable because the prior art of record does not disclose or reasonably suggest a combination of parsing to determine a meaning and a context of speech associated with a request involving a grammar by a domain agent, satisfying a predetermined confidence level, updating user specific vocabularies or dictionaries, and determining an identity of a user based on voice characteristics. Sabourin (U.S. Patent No. 6,208,964) teaches updating user specific vocabularies or dictionaries, but not in combination with satisfying a predetermined confidence level and determining an identity of the user based on voice characteristics of the user. Although determining an identity of a user based on voice characteristics is known individually for a voice profile, the prior art of record does not disclose or reasonably suggest that feature in combination with updating a user specific vocabulary when a predetermined confidence level is not met.

'006 File History, Notice of Allowance (April 28, 2001) at 2 (attached as Exhibit 11).

U.S. PATENT NO. 8,140,327

58. On March 20, 2012, the U.S. Patent Office duly and legally issued the '327 Patent, entitled "System And Method For Filtering And Eliminating Noise From Natural Language Utterances To Improve Speech Recognition And Parsing." A true and correct copy of the '327 Patent is attached hereto as **Exhibit 3**.

59. Dialect is the owner and assignee of all right, title, and interest in and to the '327 Patent, including the right to assert all causes of action arising under the '327 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

60. As described in the '327 Patent, natural-language speech-recognition systems “may operate in many environments, including environments with background noise, point noise sources and people holding conversations.” '327 Patent, at 7:35-38.

61. To solve these problems, the inventors of the '327 Patent conceived of novel software and hardware techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing natural-language speech-recognition systems. For example, claim 14 recites a system comprising a novel combination of audio- and signal-processing elements that includes a microphone array configured to add one or more nulls to a beam pattern where the one or more nulls notch out point or limited area noise sources from the input speech signal; an adaptive filter that (1) compares environmental noise to the input signal to set one or more parameters associated with an adaptive filter, (2) uses band shaping and notch filtering remove narrow-band noise from the input speech signal, and (3) uses adaptive echo cancellation to suppress cross-talk and environmentally cause echoes; a speech coder that uses adaptive lossy audio compression to remove momentary gaps from the input speech signal; and a transceiver that communicates the digitized input speech signal to a speech recognition engine at a rate that depends on available bandwidth:

14. A system for filtering and eliminating noise from natural language speech utterances, comprising:

a microphone array configured to add one or more nulls to a beam pattern steered to point in a direction associated with a user speaking a natural language utterance to capture an input speech signal corresponding to the natural language utterance, wherein the one or more nulls notch out point or limited area noise sources from the input speech signal;

an adaptive filter coupled to the microphone array, wherein the adaptive filter is configured to:

receive the input speech signal corresponding to the natural language utterance from the microphone array and compare environmental, noise to the input speech signal to set one or more parameters associated with the adaptive filter;

use band shaping and notch filtering to remove narrow-band noise from the input speech signal received from the microphone array according to the one or more parameters; and


suppress cross-talk and environmentally caused echoes in the input speech signal received from the microphone array using adaptive echo cancellation;

a speech coder arranged between the adaptive filter and a speech recognition engine, wherein the speech coder is configured to receive the input speech signal passed through the adaptive filter and use adaptive lossy audio compression to remove momentary gaps from the input speech signal and variable rate sampling to compress and digitize the input speech signal, wherein the speech coder optimizes the adaptive lossy audio compression and the variable rate sampling to only preserve components in the input speech signal that will be input to the speech recognition engine; and

a transceiver configured to communicate the digitized input speech signal from a buffer in the speech coder to the speech recognition engine at a rate that depends on available bandwidth associated with a communication link that connects the transceiver and the speech recognition engine.

'327 Patent at claim 14.

62. Embodiments of these claimed elements are shown and described in the specification. For example, Figure 1 shows an overall diagrammatic view of the interactive natural language speech processing system according to one embodiment:



63. The specification of the '327 Patent describes how these claim elements help the overall system overcome the technical limitations of existing speech recognition systems. *See e.g., id.* at 15:42-60 (describing microphone arrays and beam forming); 15:61-16:10 (describing adaptive filters); 16:11-27 (describing speech encoders); 15:29-41 (describing adaptive rate limiting of data transmission); 16:11-27 (same).

64. In explaining the reasons or allowing the claims, the United States Patent and Trademark Office described how the closest existing prior art did not disclose or teach the claimed combination of inventive elements:

Regarding independent claims 1 and 14, the prior art of record does not disclose or reasonably suggest the combination of receiving a natural language utterance for speech recognition that includes all the details of a microphone array that adds one or more nulls to a beam pattern steered to a notch out point or limited area noise sources, comparing environmental noise to set one or more parameters associated with an adaptive filter, where the adaptive filter uses band shaping and notch filtering to remove narrow-band

noise, using adaptive echo cancellation to suppress cross-talk and environmentally caused echoes, sending the adaptively filtered speech to a speech coder that uses adaptive lossy audio compression to remove momentary gaps and variable rate sampling to compress the speech signal that is optimized to preserve components, and transmitting the digitized input speech from a buffer at a rate that depends upon available bandwidth between the speech coder and a speech recognition engine.

'327 File History, Notice of Allowance (November 21, 2011) at 2 (attached as Exhibit 12).

U.S. PATENT NO. 8,195,468

65. On June 5, 2012, the U.S. Patent Office duly and legally issued the '468 Patent, entitled "Mobile Systems And Methods Of Supporting Natural Language Human-Machine Interactions." A true and correct copy of the '468 Patent is attached hereto as **Exhibit 4**.

66. Dialect is the owner and assignee of all right, title, and interest in and to the '468 Patent, including the right to assert all causes of action arising under the '468 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

67. The '468 Patent explains that "Telematic systems are systems that bring human-computer interfaces to mobile environments. Conventional computer interfaces use some combination of keyboards, keypads, point and click techniques and touch screen displays. These conventional interface techniques are generally not suitable for a mobile environments, due at least in part to the speed of interaction and the inherent danger and distraction. Therefore, speech interfaces are being adopted in many telematic applications. However, creating a natural language speech interface that is suitable for use in the mobile environment has proved difficult." '468 Patent at 1:26-39 "The fact that most natural language requests and commands are incomplete in

their definition is a significant barrier to natural language query-response interaction. Further, some questions can only be interpreted in the context of previous questions, knowledge of the domain, or the user's history of interests and preferences. Thus, some natural language questions and commands may not be easily transformed to machine processable form. Compounding this problem, many natural language questions may be ambiguous or subjective. In these cases, the formation of a machine processable query and returning of a natural language response is difficult at best. *Id.* at 2:6-18.

68. As the '468 Patent explains, prior to the inventions disclosed therein, a machine's ability to communicate with humans in a natural manner was a difficult technical problem in need of a technical solution. As described in the specification, under the existing systems and devices "verbal communications and machine processing of requests that are extracted from the verbal communications may be fundamentally incompatible," because the existing systems and devices use requests that are "highly structured and may not be inherently natural to the human user." *Id.* at 1:53-58.

69. To solve these problems, the inventors of the '468 Patent conceived of novel software techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing mobile systems. For example, claim 19 recites a novel method of processing a combination of speech and non-speech inputs that receives multimodal natural language input from a user including a natural language utterance and a non-speech input, identifies the user, creates and merges transcripts of the inputs using a speech recognition engine and a semantic knowledge-based model that includes personalized and general models derived from prior interactions with the identified user and multiple users, and an environmental model derived from the identified user's environment. The method identifies entries in a context stack

matching information in the merged transcription, and determines the most likely context from the matched entries. It then identifies a domain agent associated with the most likely context, communicates a request to the domain agent and generates a response to the user from content provided by the domain agent:

A method for processing multi-modal natural language inputs, comprising:

receiving a multi-modal natural language input at a conversational voice user interface, the multi-modal input including a natural language utterance and a non-speech input provided by a user, wherein a transcription module coupled to the conversational voice user interface transcribes the non-speech input to create a non-speech-based transcription;

identifying the user that provided the multi-modal input;

creating a speech-based transcription of the natural language utterance using a speech recognition engine and a semantic knowledge-based model, wherein the semantic knowledge-based model includes a personalized cognitive model derived from one or more prior interactions between the identified user and the conversational voice user interface, a general cognitive model derived from one or more prior interactions between a plurality of users and the conversational voice user interface, and an environmental model derived from an environment of the identified user and the conversational voice user interface;

merging the speech-based transcription and the non-speech-based transcription to create a merged transcription;

identifying one or more entries in a context stack matching information contained in the merged transcription;

determining a most likely context for the multi-modal input based on the identified entries;

identifying a domain agent associated with the most likely context for the multi-modal input;

communicating a request to the identified domain agent; and

generating a response to the user from content provided by the identified domain agent as a result of processing the request.

'468 Patent at claim 19.

70. Embodiments of these claimed elements are shown and described in the specification. For example, Figure 8 illustrates one exemplary embodiment:

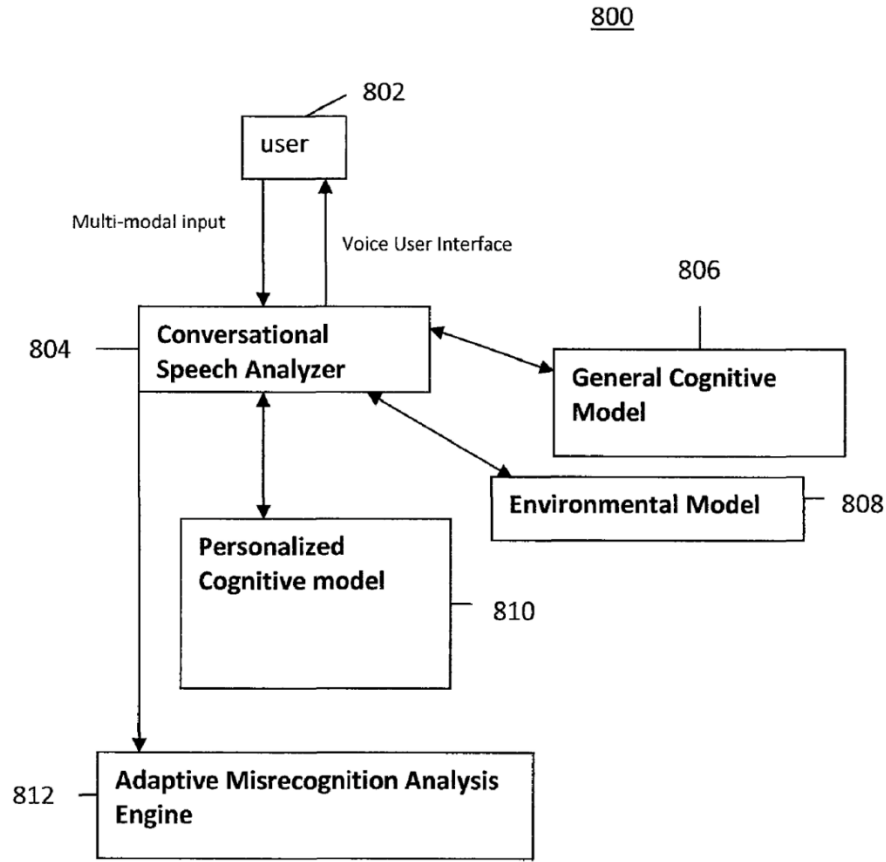


FIGURE 8

71. The specification of the '468 Patent describes how these claim elements help the overall system overcome the technical limitations of existing speech recognition systems. *See e.g.*, '468 Patent at 23:58-24:40.

72. In explaining the reasons or allowing the claims, the United States Patent and Trademark Office described how the closest existing prior art did not disclose or teach the claimed combination of inventive elements:

The prior art of record does not teach the combination of limitations in independent claims (1, 15, 19), including multi-modal natural language speech and non- speech input being transcribed and merged, identifying a user with a conversational speech analysis

engine which uses a semantic knowledge-based model including a personalized cognitive model derived from one or more prior interactions between the identified user and the mobile device, a general cognitive model derived from one or more prior interactions between a plurality of users and the mobile device, and an environmental model derived from an environment of the identified user and the mobile device, and a knowledge-enhanced speech recognition engine that identifies one or more entries in a context stack matching information contained in the merged transcription and determines a most likely context for the multi-modal natural language input based on the identified entries, and response generation by a domain agent associated with the most likely context identified by the system, where the domain agent receives a request.

'468 File History, Notice of Allowance (October 23, 2011) at 2 (attached as Exhibit 13).

U.S. PATENT NO. 9,031,845

73. On May 12, 2015, the U.S. Patent Office duly and legally issued the '845 Patent, entitled "Mobile Systems And Methods For Responding To Natural Language Speech Utterance." A true and correct copy of the '845 Patent is attached hereto as **Exhibit 5**.

74. Dialect is the owner and assignee of all right, title, and interest in and to the '845 Patent, including the right to assert all causes of action arising under the '845 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

75. As described in the '845 Patent, "many natural language questions may be ambiguous or subjective." '845 Patent, at 2:13-15. While human beings are adept at interpreting these types of communication, they are not easily "transformed to machine processable form," resulting in a uniquely challenging technological problem to building natural-language speech-recognition interfaces for computers and vehicles. *Id.* at 2:10-17. Furthermore, "retrieval of both

local and network hosted online information and processing of commands in a natural manner” was “a difficult problem in any environment, especially onboard vehicles.” *Id.* at 1:47-50.

76. Thus, “creating a natural language speech interface that is suitable for use in the vehicular environment” was a difficult problem. *Id.* at 1:38-40. “Conventional computer interfaces” were “generally not suitable for a vehicular environment, owing to the speed of interaction and the inherent danger and distraction” involved. *Id.* at 1:30-35. The challenges for existing natural language speech interfaces included the need to “accommodate commands and queries from a wide range of domains and from many users with diverse preferences and needs.” *Id.* at 1:40-42. Existing speech interfaces required queries and commands to be “highly structured” in a way that is “not inherently natural to the human user,” causing natural language inputs to be “fundamentally incompatible” with existing solutions. *Id.* at 1:53-58.

77. Compounding these problems, a natural language speech interface in a vehicular environment has its own unique challenges. Vehicular speech interfaces using “wireless communications” increase “the chances that queries will not complete or return useful results.” *Id.* at 2:29-31. Additionally, the system may include “controls of a critical nature or with safety implications,” which may need to be executed on-board the vehicle to “verify that a command will not create a hazardous condition before it is executed.” *Id.* at 9:48-52; *see also id.* at 16:53-17:4.

78. As acknowledged in the ’845 Patent, despite improvements in the accuracy of existing speech recognition and the application of natural language processing in “parsing of speech queries,” no existing system overcame the “significant barriers to creating of a complete natural language speech-based query and response environment.” *Id.* at 1:62-2:4. The obstacles addressed by the inventions of the ’845 Patent included the fact that “most natural language queries and commands are incomplete in their definition” and “some questions can only be interpreted in

the context of previous questions, knowledge of the domain, or the user's history of interests and preferences." *Id.* at 2:5-11.

79. To solve these problems, the inventors of the '845 Patent conceived of novel software techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing computer or vehicle systems. For example, claim 1 recites a novel configuration of software structures that processes a natural-language input in order to determine the domain and the context of the user's input, transforms the input into a command or query that can be processed by the system based on the domain and the context, and uses the transformed input to selectively execute the command or query on-board the vehicle or to invoke a wide area network wireless device to execute the command or query off-board the vehicle:

1. A mobile system for processing natural language utterances, comprising:
 - one or more physical processors at a vehicle that are programmed to execute one or more computer program instructions which, when executed, cause the one or more physical processors to:
 - receive a natural language utterance associated with a user;
 - perform speech recognition on the natural language utterance;
 - parse and interpret the speech recognized natural language utterance;
 - determine a domain and a context that are associated with the parsed and interpreted natural language utterance;
 - formulate a command or query based on the domain and the context;
 - determine whether the command or query is to be executed on-board or off-board the vehicle;
 - execute the command or query at the vehicle in response to a determination that the command or query is to be executed on-board the vehicle; and
 - invoke a device that communicates wirelessly over a wide area network to process the command or query such that the command or query is executed off-board the vehicle in response to a determination that the command or query is to be executed off-board the vehicle.

'845 Patent at Claim 1.

80. Embodiments of these claimed elements are shown and described in the specification. For example, Figure 2 shows an overall block diagram of the system according to

an embodiment in which “the main speech processing unit 98 and speech unit 128 are external to the [Telematics Control Unit(s)]”:

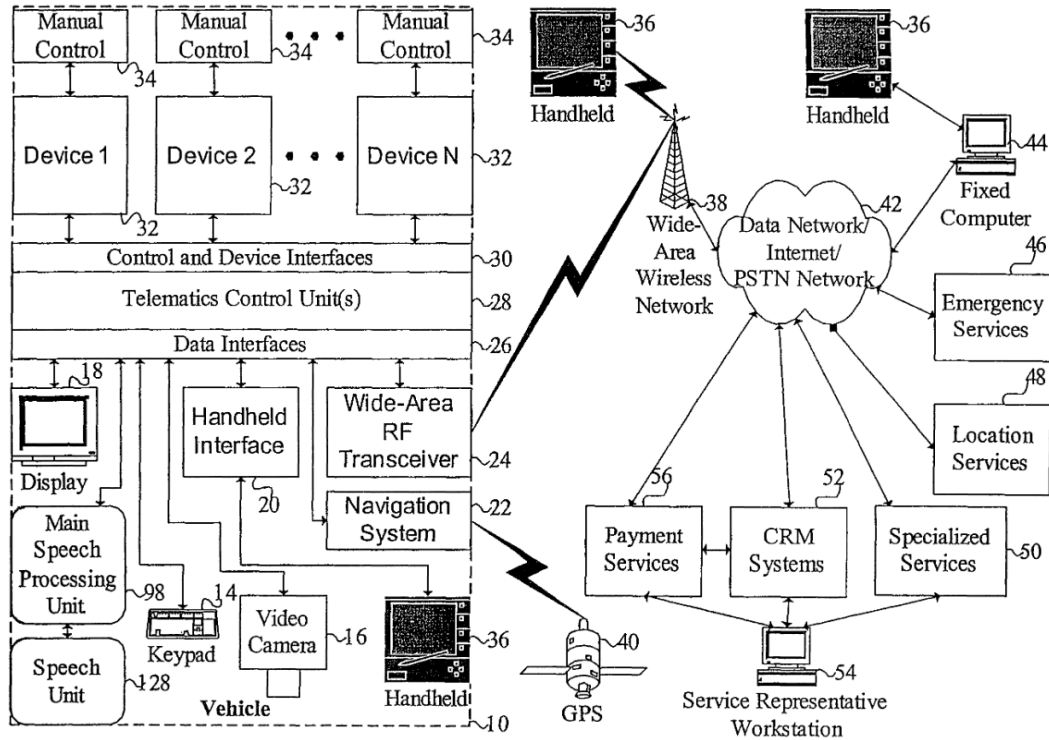


Figure 2. Second Embodiment System Block Diagram

See also, *id.* at 18:48-52. Additionally, Figure 5 shows an overall diagrammatic view of the interactive natural language speech processing system according to one embodiment:

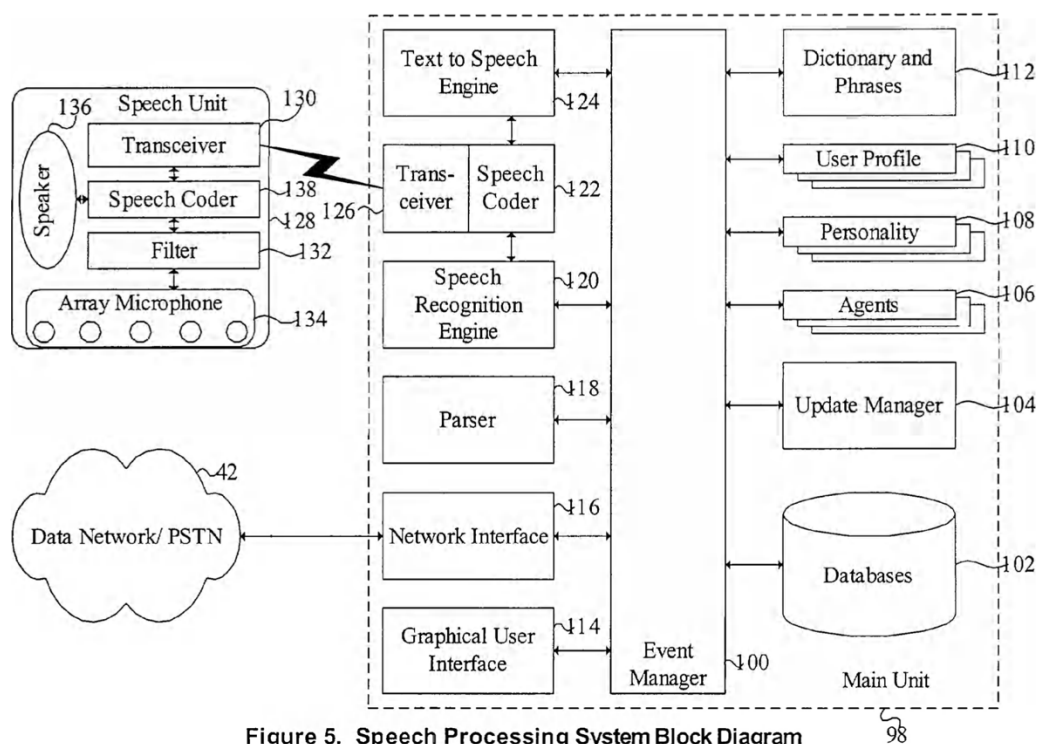


Figure 5. Speech Processing System Block Diagram

98

81. The specification of the '845 Patent describes how these claim elements help the overall system overcome the technical limitations of existing speech recognition systems. *See e.g., id.* at 3:52-65 (describing wireless communications with vehicle-related systems); 4:10-5:2 (describing the use of domain specific behavior and data); 7:24-63 (describe the use of domain and context information in formulating commands and queries); 16:36-18:39 (describing the integration and interaction between vehicle systems and wide area networks); 24:37-61 (describing the use of user-specific context in formulating commands and queries).

82. In explaining the reasons for allowing the claims, the United States Patent and Trademark Office described how the closest existing prior art did not disclose or teach the claimed combination of inventive elements:

However, in Funk, the decision of whether to process the Request on-board or off-board is based on keyword recognition of the input speech. Funk does not teach that it has an

on-board capability for natural language processing beyond keyword spotting. If a keyword corresponding to the car sensors is not spotted in the input speech, Funk sends the speech processing off-board such that the natural language processing for determining such a query is done by the Voice Portal which is not onboard. (Funk, Figures 3 or 4 and Col. 7, line 50 to Col. 8, line 9.)

Junqua (U.S. 6,598,018) teaches natural language processing and extraction of command or query entirely on-board a vehicle but not sending the query off-board.

The references provided by the IDS of 4/24/2014, that were prior to the filing date of 7/15/2003 of the parent application or the priority date of 7/15/2002 of the provisional application to which the parent claims priority, including Quilici (U.S. 7,974,875) which pertains to the use of telephone, were considered. These references do not combine the multimodal and natural language input in a car/vehicle/automobile and determining a command or query from the natural language input based on the context of the input and then making the decision to perform a command or query that is obtained from speech recognizing the natural language on-board or sending it off-board.

'845 File History, Notice of Allowance (December 30, 2014) at 10-11 (attached as Exhibit 14).

U.S. PATENT NO. 9,263,039

83. On February 16, 2016, the U.S. Patent Office duly and legally issued the '039 Patent, entitled "Systems And Methods For Responding To Natural Language Speech Utterance." A true and correct copy of the '039 Patent is attached hereto as **Exhibit 6**.

84. Dialect is the owner and assignee of all right, title, and interest in and to the '039 Patent, including the right to assert all causes of action arising under the '039 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

85. As described in the '039 Patent, a “machine’s ability to communicate with humans in a natural manner” was “a difficult problem.” '039 Patent, at 1:39-40. While human beings are adept at communicating with an understanding of the context and domain knowledge of the target they are attempting to communicate with, “machine-based queries may be highly structured and may not be inherently natural to the human user” resulting in a uniquely challenging technological problem to building natural-language speech-recognition interfaces. *Id.* at 1:40-49.

86. As acknowledged in the '039 Patent, despite improvements in the accuracy of existing speech recognition and the application of natural language processing in “parsing of speech queries,” no existing system overcame the “significant barriers to creating of a complete speech-based and/or non-speech-based natural language query and response environment.” *Id.* at 1:52-61. Among the challenges in developing natural-language speech-recognition interfaces was the fact that “the question[s] asked or the response[s] received” from users “are incomplete, ambiguous or subjective.” *Id.* at 1:30-35. The inventive elements of the '039 Patent are directed to a system that can process communications with “incomplete thoughts, incomplete sentences, incomplete phrases, slang terminology, repeated words, word variations, synonyms, or other imperfect information.” *Id.* at 2:1-7.

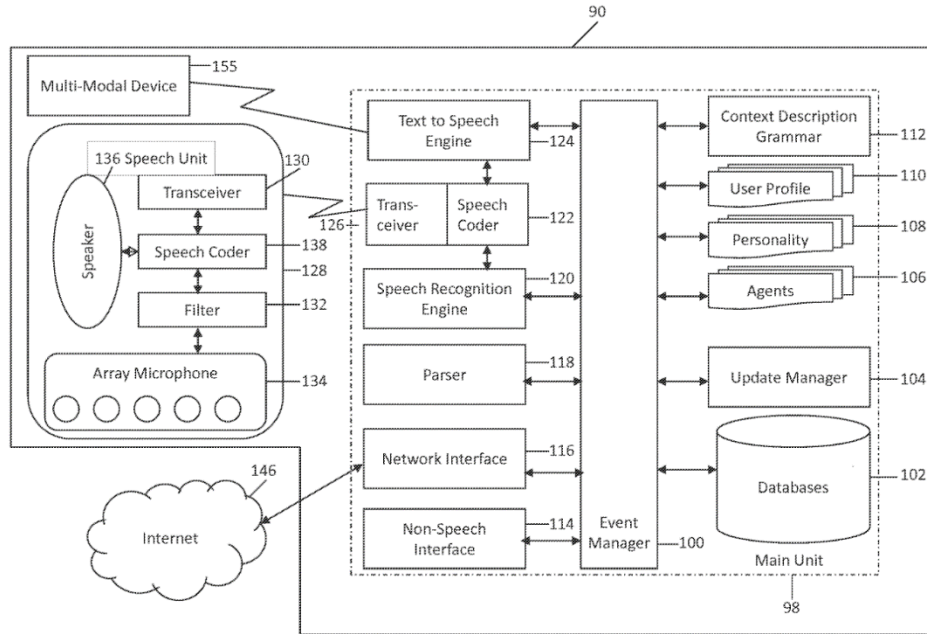
87. To solve these problems, the inventors of the '039 Patent conceived of novel software techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing speech-recognition computer control systems. For example, claim 13 recites a novel configuration of software structures that transcribes the speech and non-speech communications to create speech-based and non-speech-based textual messages, merges the speech-based and non-speech based textual messages; searches the merged query for text combinations, compares the text combinations to context description grammar, generates a

relevance score based on that comparison, selects a domain agent based on the relevance score, and organizing content based on the results from the relevance score to generate a response:

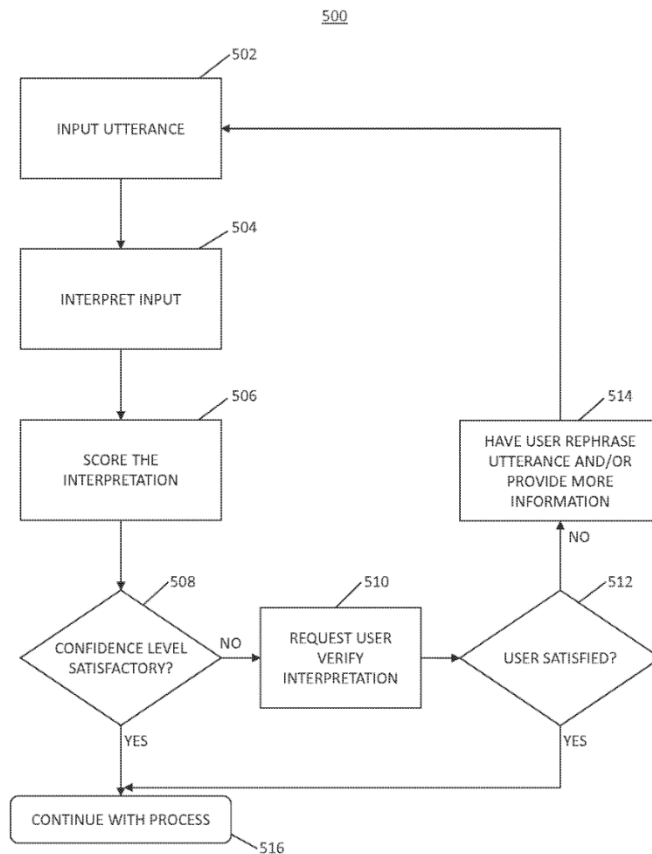
13. A method of processing speech and non-speech communications, comprising:
 - receiving the speech and non-speech communications;
 - transcribing the speech and non-speech communications to create a speech-based textual message and a non-speech-based textual message;
 - merging the speech-based textual message and the non-speech-based textual message to generate a query;
 - searching the query for text combinations;
 - comparing the text combinations to entries in a context description grammar;
 - accessing a plurality of domain agents that are associated with the context description grammar;
 - generating a relevance score based on results from comparing the text combinations to entries in the context description grammar;
 - selecting one or more domain agents based on results from the relevance score;
 - obtaining content that is gathered by the selected domain agents; and
 - generating a response from the content, wherein the content is arranged in a selected order based on results from the relevance score.

'039 Patent at Claim 13.

88. Embodiments of these claimed elements are shown and described in the specification. For example, Figure 1 shows an overall diagrammatic view of the interactive natural language speech processing system according to one embodiment:



Additionally, Figure 5 shows a process for correctly interpreting a user's utterance according to one embodiment:



89. The specification of the '039 Patent describes how these claim elements help the overall system overcome the technical limitations of existing speech recognition systems. *See e.g., id.* at 13:61-14:37 (describing comparison to context description grammar and relevance scoring); 20:20-58 (describing improved word recognition accuracy using data from context description grammar); 21:28-36 (describing a scoring system); 23:19-29 (describing a scoring system); 28:4-31 (describing the process of Figure 5); 28:56-29:8 (describing selection of agents).

90. In explaining the reasons or allowing the claims, the United States Patent and Trademark Office described how the closest existing prior art did not disclose or teach the claimed combination of inventive elements:

The prior art of record does not disclose or suggest the combination of a comparison module that compares text combinations to entries in a context description grammar, a scoring module that provides relevance scores based on the results from the comparison module, a domain agent selector that selects domain agents based on results from the scoring module, and a response generating module that generates a response from the content, wherein the content is arranged in a selected order based on results from the scoring module, as required by independent claim 1. Independent claims 13 and 19 recite similar limitations, and are allowed for similar reasons as claim 1.

'039 File History, Notice of Allowance (October 15, 2015) at 3 (attached as Exhibit 15).

U.S. PATENT NO. 9,495,957

91. On November 15, 2016, the U.S. Patent Office duly and legally issued the '957 Patent, entitled "Mobile Systems And Methods Of Supporting Natural Language Human-Machine Interactions." A true and correct copy of the '957 Patent is attached hereto as **Exhibit 7**.

92. Dialect is the owner and assignee of all right, title, and interest in and to the '957 Patent, including the right to assert all causes of action arising under the '957 Patent and the right to sue and obtain any remedies for past, present, or future infringement.

93. The '957 Patent explains that “Telematic systems are systems that bring human-computer interfaces to mobile environments. Conventional computer interfaces use some combination of keyboards, keypads, point and click techniques and touch screen displays. These conventional interface techniques are generally not suitable for a mobile environments, due at least in part to the speed of interaction and the inherent danger and distraction. Therefore, speech interfaces are being adopted in many telematic applications. However, creating a natural language speech interface that is suitable for use in the mobile environment has proved difficult.” '957 Patent at 1:35-46. “The fact that most natural language requests and commands are incomplete in their definition is a significant barrier to natural language query-response interaction. Further, some questions can only be interpreted in the context of previous questions, knowledge of the domain, or the user’s history of interests and preferences. Thus, some natural language questions and commands may not be easily transformed to machine processable form. Compounding this problem, many natural language questions may be ambiguous or subjective. In these cases, the formation of a machine processable query and returning of a natural language response is difficult at best. *Id.* at 2:13-24.

94. As the '957 Patent explains, prior to the inventions disclosed therein, a machine’s ability to communicate with humans in a natural manner was a difficult technical problem in need of a technical solution. As described in the specification, under the existing systems and devices “verbal communications and machine processing of requests that are extracted from the verbal communications may be fundamentally incompatible,” because the existing systems and devices

use requests that are “highly structured and may not be inherently natural to the human user.” *Id.* at 1:63-68.

95. To solve these problems, the inventors of the '957 Patent conceived of novel software techniques and structures (and novel combinations and ordering of techniques and structures) not found in existing mobile systems. For example, claim 1 recites a novel computer system for processing a natural language utterance that generates a context stack of context information that corresponds to prior utterances, performs speech recognition to determine words in the utterance, identifies context entries in the context stack that correspond to the words by comparing entries from the context stack to the words and generates rank scores for the context entries based on the comparison. Based on the words and the context information, the system determines a command or request associated with the utterance:

1. A system for processing a natural language utterance, the system including one or more processors executing one or more computer program modules which, when executed, cause the one or more processors to:
 - generate a context stack comprising context information that corresponds to a plurality of prior utterances, wherein the context stack includes a plurality of context entries;
 - receive the natural language utterance, wherein the natural language utterance is associated with a command or is associated with a request;
 - determine one or more words of the natural language utterance by performing speech recognition on the natural language utterance;
 - identify, from among the plurality of context entries, one or more context entries that correspond to the one or more words, wherein the context information includes the one or more context entries, wherein identifying the one or more context entries comprises:
 - comparing the plurality of context entries to the one or more words;
 - generating, based on the comparison, one or more rank scores for individual context entries of the plurality of context entries; and
 - identifying, based on the one or more rank scores, the one or more context entries from among the plurality of context entries; and

determine, based on the determined one or more words and the context information, the command or the request associated with the natural language utterance.

'957 Patent at claim 1.

96. Embodiments of these claimed elements are shown and described in the specification. For example, the specification explains:

According to yet another alternative embodiment of the invention, context information may be determined from a command or request that is presented as a text message and/or a command or request that is presented as a verbal utterance and processed using a multi-pass automatic speech recognition module that transcribes the verbal utterance to a text message....

According to one embodiment of the invention, the text message may be searched for a particular character, group of characters, word, group of words, and other text combinations. The text combination may be compared against entries in a context description grammar that is associated with each agent 106. If a match is identified between an active grammar in the context description grammar and the command and/or request, then the match may be scored. The agents 106 may be ranked based on the determined score. In generating an aggregate response from the one or more responses received from the agents 106, the ordering of the responses from the individual agents may be determined based on the rank of agents 106. The aggregate response may be generated by a response generator module. Agents 106 may update a context stack, that includes an ordered list of command contexts, to enable follow-up requests.

'957 Patent at 20:16-46.

AMAZON’S USE OF THE PATENTED TECHNOLOGY

97. Amazon is a digital and retail conglomerate that provides *inter alia* internet cloud services, software products, digital products, and internet-enabled hardware in the United States and worldwide. In its 10-K filing for fiscal year 2022, Amazon reported consolidated total net sales of over \$513 billion.

98. Among the products and services that Amazon makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States, are: the “Alexa Products”—Amazon’s Alexa virtual assistant and offerings that include Alexa, including the Echo product line (such as Echo 1st Gen., Echo 2nd Gen., Echo Dot 1st Gen., Echo Dot 2nd Gen., Echo Dot 3rd Gen., Echo Dot Kids Edition, Echo Show 1st Gen., Echo Show 2nd Gen., Echo Show 5, Echo Spot, Echo Plus 1st Gen., Echo Plus 2nd Gen., Echo Auto, and Echo Look), Amazon’s Alexa apps, Music apps, and Shopping apps on a smartphone or other mobile device, Amazon’s Alexa cloud, Alexa Voice Services, and Amazon.com website, and any other device, app, or instrumentality that includes, provides access to, or works with Alexa (such as Amazon Tap, Amazon Dash Wand, Echo Wall Clock—servers, network infrastructure, tablets, and internet of things (“IoT”) devices (collectively, the “Accused Products”), which infringe the Asserted Patents as described in the counts below.

99. On information and belief, Amazon also provides third-party developers of automobiles, electronic hardware, and software with interfaces to the Alexa Products in order to encourage those developers to design, make, use, import into the United States, offer to sell, and sell products and services capable of being voice-controlled by the Accused Products.

FIRST COUNT
(Infringement of U.S Patent No. 7,693,720)

100. Dialect incorporates by reference the allegations set forth in Paragraphs 1–99 as though fully set forth herein.

101. The claims of the '720 Patent are valid and enforceable.


102. The claims of the '720 Patent are directed to patentable subject matter. Specifically, Claim 1 of the '720 Patent is directed to a novel vehicular speech-recognition system including the claimed novel speech recognition engine, parser, and agent architecture. The inventive, tangible claimed structures of the '720 Patent improve on existing natural language processing systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

103. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and services that embody the invention disclosed and claimed in the '720 Patent, including at least the Alexa Products, operating via tablets, smartphones, automobile infotainment systems, or other devices supporting in-vehicle Alexa functionality such as the Echo Auto, Alexa Auto SDK and Alexa Mobile Accessory Kit (collectively, the “Accused Automotive Products and Services”).

104. Each of the Accused Automotive Products and Services contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 1 of the '720 Patent.

105. Each of the Accused Automotive Products and Services comprises a mobile system responsive to a user generated natural language speech utterance.


106. For example, Amazon describes how the Alexa Auto SDK and Alexa Mobile Accessory Kit provide mobile solutions for hands-free interactions.¹⁸



Head Unit with Alexa Built-in

By building Alexa into your vehicle with the Alexa Auto SDK, your customers can interact with Alexa, hands-free, using the in-cabin microphone and speakers. You can also display digital content on the infotainment screen, like album artwork, audiobook covers, search results, and more.

[Alexa Auto SDK »](#)



Aftermarket Device with Alexa Built-in

Bring Alexa to the millions of vehicles already on the road by building Alexa into your automotive aftermarket device. The Alexa Mobile Accessory (AMA) Kit enables device makers to easily connect Bluetooth-enabled audio devices with Alexa using the Amazon Alexa App on Android and iOS devices.

[Learn more about AMA »](#)

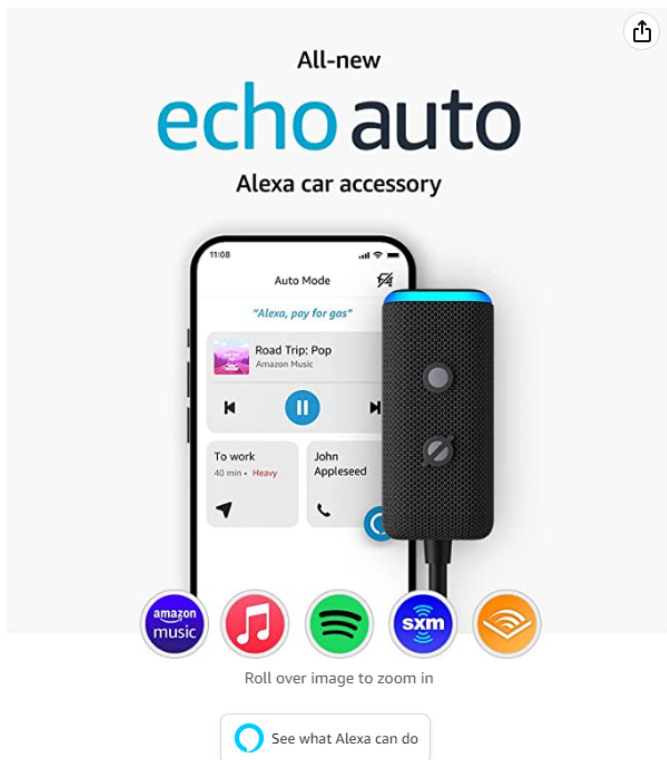
Amazon further describes the Alexa Auto SDK as enabling speech recognition, Automatic Speech Recognition and Natural Language Understanding.¹⁹

¹⁸ <https://developer.amazon.com/en-US/alexa/alexa-auto/vehicles#experiences>

¹⁹ <https://developer.amazon.com/en-US/blogs/alexa/alexa-auto/2019/09/enabling-offline-access-to-alexa-in-vehicles-with-local-voice-control-extension-to-alexa-auto-sdk-v2-0>

The Alexa Auto SDK enables core Alexa functionality, such as speech recognition and text-to-speech, and other capabilities such as streaming media, controlling smart home devices, notifications, weather reports, and tens of thousands of custom skills. It can be updated over-the-air and now includes the Local Voice Control extension, which ensures that customers get a reliable response from Alexa regardless of connectivity. On a normal commute, drivers lose connectivity or experience intermittent or transitional connectivity while driving in dense urban areas, through tunnels, in remote areas, or in parking garages. With the Local Voice Control extension to the Alexa Auto SDK, automakers embed a scaled-down version of Alexa's cloud-based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle's infotainment system that increases resiliency of the cloud-based capabilities. This ensures that a customer can always ask, "Alexa, turn on the A/C," "Alexa, tune to 90.5 FM," or "Alexa, call Jackie," and get the appropriate response and action from Alexa.

Amazon further describes Echo Auto as enabling hands-free features.²⁰



All-new
echo auto
Alexa car accessory

Amazon Music, Apple Music, Spotify, and more, or listen live to radio stations. Use Follow Me Music to resume media playback.

- **CALL AND MESSAGE WITH YOUR VOICE** - Use your voice to make calls, reply to text messages, drop in on Alexa-enabled devices in your home or broadcast announcements.
- **GET YOUR FAVORITE ENTERTAINMENT HANDS-FREE** - Binge a hot new podcast, catch up on the news, or listen to best-selling Audible books.
- **ALEXA CAN HELP WITH AMAZON ROADSIDE ASSISTANCE** - Just say, "Alexa, call Roadside Assistance" and you'll be connected to an agent who can assess the situation and request pay-as-you-go services as needed.
- **CONTROL YOUR SMART HOME FROM THE ROAD** - Ask Alexa to set the thermostat, turn off the lights, check if your front door is locked, and more while you're away from home.
- **DESIGNED TO PROTECT YOUR PRIVACY** - Built with multiple layers of privacy controls, including a mic-off button and light indicator.

We want you to know

Echo Auto is designed to enable Alexa hands-free features, such as play music, make calls, answer text messages and more for cars that don't have a voice assistant built-in. Echo Auto connects to the Alexa app on your phone using your existing data plan and plays through your car's speakers via Bluetooth connection or Aux input. Carrier charges may apply.

Car and Phone compatibility - Check if your car is compatible and check if your phone is compatible. Note: Echo Auto is compatible with most cars with Bluetooth or Aux input.

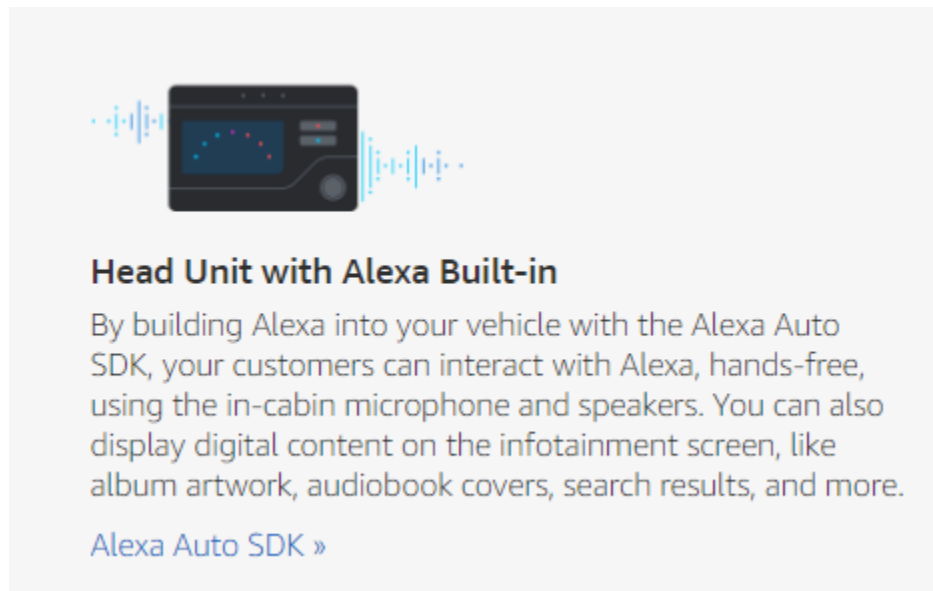
Fast car charger - This product is Quick Charge 3.0 certified for fast charging.

Mounting your Echo Auto - The included adhesive car mount should not be placed on soft plastic, leather, or fabric surfaces. Choose a hard, durable surface away from the airflow.

107. Each of the Accused Automotive Products and Services comprises a speech unit connected to a computer device on a vehicle, wherein the speech unit receives a natural language speech utterance from a user and converts the received natural language speech utterance into an electronic signal.

²⁰ <https://developer.android.com/training/cars/apps/poi>

108. For example, Amazon describes how Alexa can be built into a vehicle that comprises an in-cabin microphone and speakers connected to the automotive head unit.²¹



Amazon further describes how Alexa Mobile Accessory Kit manufacturers must support encoding for the received natural language utterance to convert the utterance into an electronic signal.²²

Codec requirements for AMA

AMA supports the following codecs:

- Opus
- mSBC

The following table shows the recommended parameter settings for the supported codecs.

Note: An AMA enabled device must support at least one codec with the corresponding parameter settings.

²¹ <https://developer.amazon.com/en-US/alexa/alexa-auto/vehicles#experiences>

²² <https://developer.amazon.com/en-US/docs/alexa/ama-kit/ama-kit-hw-security-reqs.html>

AMA software requirements

Your Bluetooth-enabled accessory must support AMA transport and control protocols, including the following requirements:

- The accessory must encode the user's utterances with the encoding technologies described in the [Codec Requirements for AMA](#). Then, the accessory sends the encoded utterances to the Alexa app on a phone by using the [AMA Kit Protocol Specification](#).
- The accessory must implement functionality to receive Alexa responses from a mobile phone by using the [AMA Kit Protocol Specification](#).
- The accessory must implement Over-the-Air (OTA) updates through the device companion app, and not through the Alexa app.
- The accessory must provide device information, such as serial number, name, supported transport, and device type to the AMA protocol stack.

Amazon further requires that the Alexa Mobile Accessory Kits include Bluetooth connections that can connect to a computing device on board the vehicle.²³

Bluetooth and Bluetooth Low Energy (BLE) requirements for AMA

The Bluetooth implementation for a device must meet the following requirements to support AMA:

- Bluetooth v4.2 dual mode + LE compliance (with DLE support)
- Standard pairing, authentication, link key, and encryption operation
- BT Protocols support for SPP/A2DP/HFP /SDP/RFCOMM/SCO
- Accessories that implement mSBC must be MFi Certified

109. Each of the Accused Automotive Products and Services comprises a natural language speech processing system connected to the computer device on the vehicle, wherein the natural language speech processing system receives, processes, and responds to the electronic signal using data received from a plurality of domain agents.

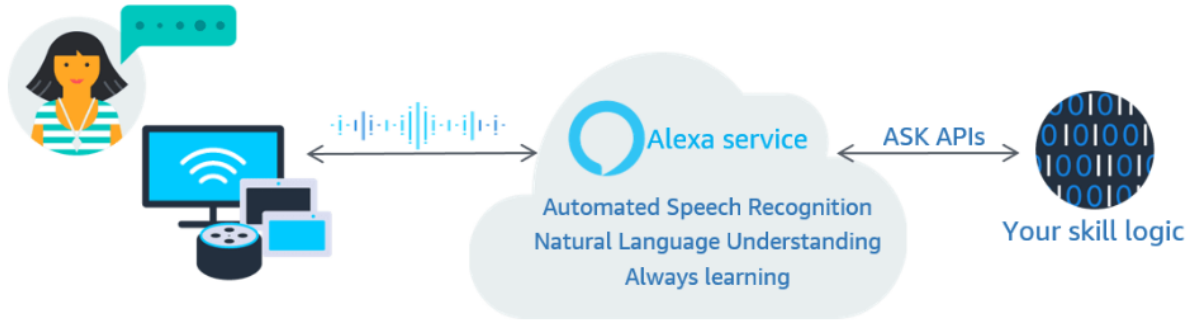
110. For example, Amazon describes Alexa as recognizing speech and determining what the user wants before sending a request to invoke a skill that can fulfill the request.²⁴

²³ *Id.*

²⁴ <https://developer.amazon.com/en-US/docs/alexa/ask-overviews/what-is-the-alexa-skills-kit.html>

A user accesses content in a skill by asking Alexa to *invoke* the skill. Alexa is always ready to invoke new skills. When a user says the wake word, "Alexa," and speaks to an Alexa-enabled device, the device streams the speech to the *Alexa service* in the cloud. Alexa recognizes the speech, determines what the user wants, and then sends a request to invoke the skill that can fulfill the request. The Alexa service handles the speech recognition and natural language processing. Your skill runs as a service on a cloud platform. Alexa communicates with your skill by using a request-response mechanism over the HTTPS interface. When a user invokes an Alexa skill, your skill receives a POST request containing a JSON body. The request body contains the parameters necessary for your skill to understand the request, perform its logic, and then generate a response.

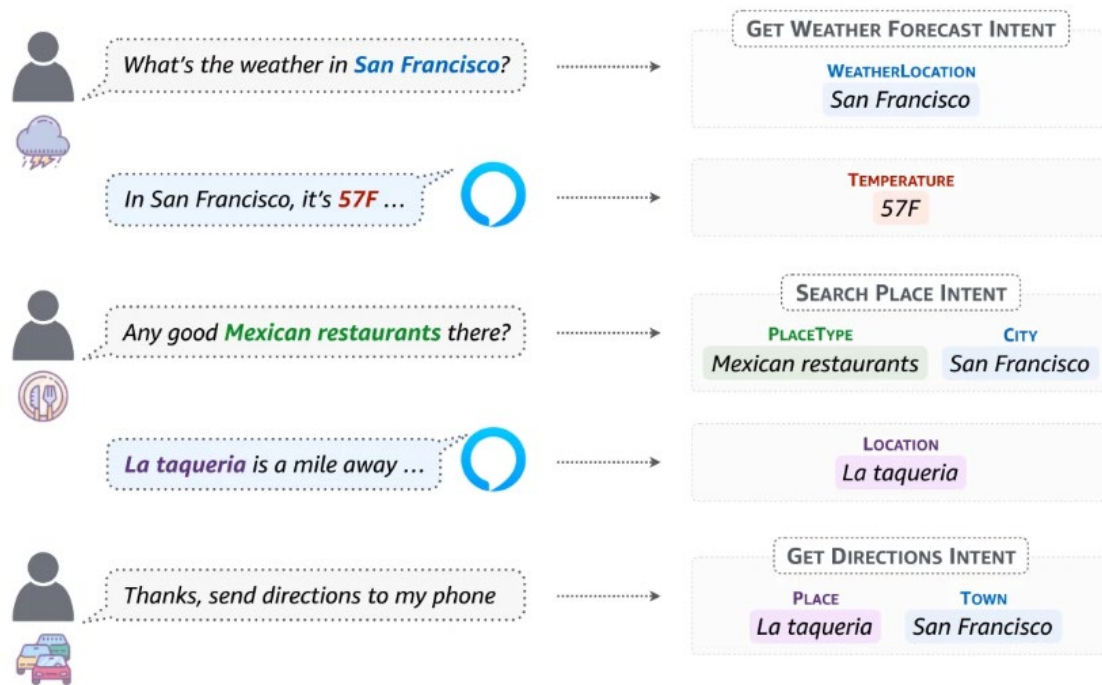
The following diagram shows the voice-activated processing flow to invoke a skill with the Alexa service.



111. Each of the Accused Automotive Products and Services comprises a speech recognition engine that recognizes at least one of words or phrases from the electronic signal using at least the data received from the plurality of domain agents, wherein the data used by the speech recognition engine includes a plurality of dictionary and phrase entries that are dynamically updated based on at least a history of a current dialog and one or more prior dialogs associated with the user.

112. For example, Amazon describes how Alexa recognizes words or phrases using information from different Alexa intents and based on at least a history of the current dialog.²⁵

²⁵ <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

Amazon further describes Alexa tracking “previously provided information.”²⁶

Alexa Conversations is especially useful for use cases where the conversation can take a number of unanticipated paths as the user naturally talks to the skill, such as when a user chooses a movie, orders food, or makes a reservation. For example, when ordering a pizza, a user might do the following:

- Answer more than one question at once ("Medium, two toppings.")
- Ask questions and expect Alexa to track previously provided information ("How many people does that feed?")
- List values ("Pepperoni and green pepper.")
- Make a correction ("Make that a large.")

Further, Amazon describes how Alexa uses phrase slot types.²⁷

²⁶ <https://developer.amazon.com/en-US/docs/alexa/conversations/about-alexa-conversations.html>

²⁷ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/slot-type-reference.html>

Phrases

Phrase slot types allow for input from a user with fewer constraints on format and content. Make sure that your skill uses no more than one phrase slot per intent.

The following phrase slot type is supported.

Slot Type	Short Description	Examples	Supported Languages
AMAZON.SearchQuery	A search query like you might enter into a search engine, for example for searching a database.	Why is the sky blue What is the capital of Maryland Restaurants near me	<ul style="list-style-type: none"> • Arabic (SA) • English (AU) • English (CA) • English (IN) • English (UK) • English (US) • French (CA) • French (FR) • German (DE) • Hindi (IN) • Italian (IT) • Japanese (JP) • Portuguese (BR) • Spanish (ES) • Spanish (MX) • Spanish (US)

As another example, Amazon describes best practices for maintaining lists of key words and phrases.²⁸

The most important rule is to design intents with a list of sample utterances that captures key words and phrases that can invoke the skill in a name-free manner. Keep your utterances specific to your skills use-cases and avoid targeting generic utterances. The Alexa ML/AI models are designed in part to route traffic to skills with a high correlation between skill functionality and target utterances.

²⁸ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/understand-name-free-interaction-for-custom-skills.html>

Alexa Products further record previous interactions and maintain a history “to learn more about you as they listen.”²⁹

Your interactions with Alexa are used to help make Alexa more accurate and useful to you over time and to build new Alexa experiences. But what does that actually mean?

Our customers want Alexa to get smarter and more helpful to them every day. To do that, we use your Alexa requests to train our speech recognition and natural language understanding systems using machine learning. In more simple terms, we analyze data to teach Alexa to more accurately make decisions or predictions. For example, when you ask if it's raining outside, Alexa decides that you're asking about the weather—and when you say, “Alexa, play music,” Alexa predicts what music you might like based on your previous requests.

113. Each of the Accused Automotive Products and Services comprises a parser that interprets the recognized words or phrases, wherein the parser uses at least the data received from the plurality of domain agents to interpret the recognized words or phrases, wherein the parser interprets the recognized words or phrases by determining a context for the natural language speech utterance; selecting at least one of the plurality of domain agents based on the determined context; and transforming the recognized words or phrases into at least one of a question or a command, wherein the at least one question or command is formulated in a grammar that the selected domain agent uses to process the formulated question or command.

114. For example, Amazon describes the Alexa Auto SDK as enabling speech recognition.³⁰

²⁹ <https://www.amazon.com/b?ie=UTF8&node=23608617011>

³⁰ <https://developer.amazon.com/en-US/blogs/alexa/alexa-auto/2019/09/enabling-offline-access-to-alexa-in-vehicles-with-local-voice-control-extension-to-alexa-auto-sdk-v2-0>

The Alexa Auto SDK enables core Alexa functionality, such as speech recognition and text-to-speech, and other capabilities such as streaming media, controlling smart home devices, notifications, weather reports, and tens of thousands of custom skills. It can be updated over-the-air and now includes the Local Voice Control extension, which ensures that customers get a reliable response from Alexa regardless of connectivity. On a normal commute, drivers lose connectivity or experience intermittent or transitional connectivity while driving in dense urban areas, through tunnels, in remote areas, or in parking garages. With the Local Voice Control extension to the Alexa Auto SDK, automakers embed a scaled-down version of Alexa's cloud-based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle's infotainment system that increases resiliency of the cloud-based capabilities. This ensures that a customer can always ask, "Alexa, turn on the A/C," "Alexa, tune to 90.5 FM," or "Alexa, call Jackie," and get the appropriate response and action from Alexa.

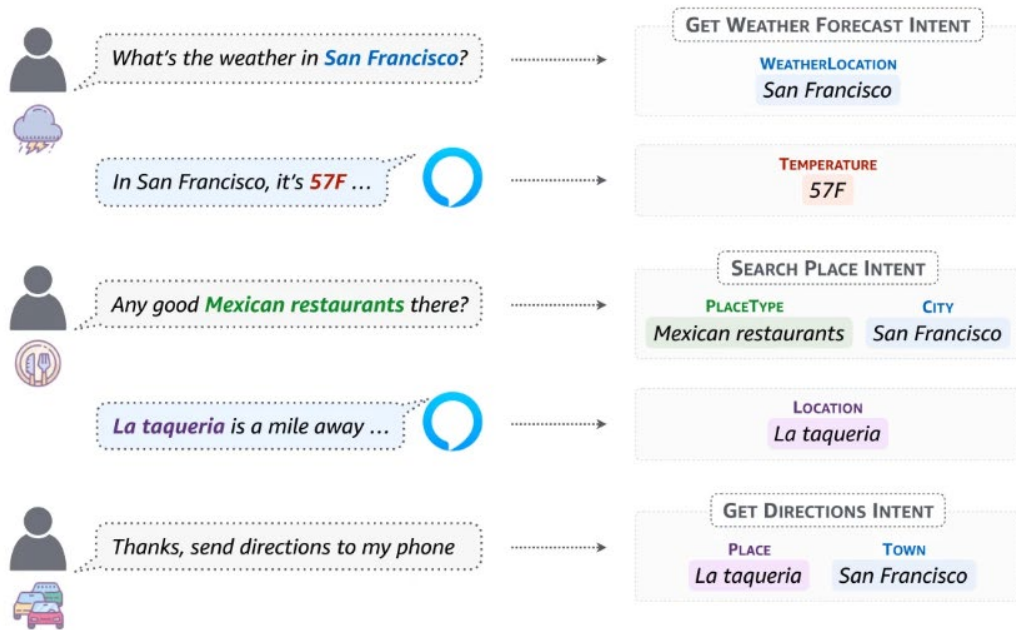
Amazon further describes how Alexa finds the most relevant skill for a given utterance..³¹

Alexa uses a two-step, scalable, and efficient neural shortlisting-reranking approach to find the most relevant skill for a given utterance. This post describes the first of those two steps, which relies on a neural model we call Shortlister. (I'll describe the second step in a follow-up post tomorrow.) Shortlister is a scalable and efficient architecture with a shared encoder, a personalized skill attention mechanism, and skill-specific classification networks. We outline this architecture in our [paper](#) "Efficient Large-Scale Neural Domain Classification with Personalized Attention", which we will present at the 56th Annual Meeting of the Association for Computational Linguistics ([ACL 2018](#)) in July.

Amazon further describes how Alexa formulates a question or command based on the grammar of different Alexa intents..³²

³¹ <https://www.amazon.science/blog/the-scalable-neural-architecture-behind-alexas-ability-to-select-skills>

³² <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.


Amazon further describes how grammar has “been a tool in Alexa’s NLU toolkit since well before the first Echo device shipped.”³³

³³ <https://www.amazon.science/blog/tools-for-generating-synthetic-data-helped-bootstrap-alexa-s-new-language-releases>

Rules of Grammar

Grammars have been a tool in Alexa’s NLU toolkit since well before the first Echo device shipped. A grammar is a set of rewrite rules for varying basic template sentences through word insertions, deletions, and substitutions.

Below is a very simple grammar, which models requests to play either pop or rock music, with or without the modifiers “more” and “some”. Below the rules of the grammar is a diagram of a computational system (a finite-state transducer, or FST) that implements them.



A toy grammar, which can model requests to play pop or rock music, with or without the modifiers “some” or “more”, and a diagram of the resulting finite-state transducer. The question mark indicates that the some_more variable is optional.

115. Each of the Accused Automotive Products and Services comprises an agent architecture that communicatively couples services of each of an agent manager, a system agent, the plurality of domain agents, and an agent library that includes one or more utilities that can be used by the system agent and the plurality of domain agents, wherein the selected domain agent uses the communicatively coupled services to create a response to the formulated question or command and format the response for presentation to the user.

116. For example, Amazon describes how the Alexa platform provides the infrastructure to support a variety of skills.³⁴

The following diagram shows the voice-activated processing flow to invoke a skill with the Alexa service.



In addition to voice interaction, skills might include complementary visuals and touch interactions.

How does a user interact with a skill?

Every Alexa skill has a *voice interaction model* that defines the words and phrases users can say to make the skill do what they want. This model determines how users communicate with and control your skill. A voice user interface is similar to a graphical user interface in a traditional app. Instead of clicking buttons and selecting options from dialog boxes, users make their requests and respond to questions by voice. Often, the voice interaction is of a much shorter duration than interaction with an app. When a user asks questions and makes requests, Alexa employs the interaction model to interpret and translate the words into a specific request to the identified skill.

The following table compares a voice user interface skill with a graphical user interface app for making airline reservations.

Action	Voice user interface	Typical graphical user interface
Make a request	User says, "Alexa, I want to fly to Denver from Seattle."	User clicks on the app, and then selects the origin and destination airports. User scrolls through the list of airports to find Seattle , and then scrolls to find Denver .
Collect more information from the user	Alexa replies, "When would you like to travel?" and then waits for a reply.	App displays a calendar, and then waits for the user to select a date.
Provide needed information	User replies, "February first." The skill makes the reservation and waits for confirmation.	User opens the calendar, selects February 1 , and then chooses OK . User clicks a button to complete the request, and then waits for confirmation.
Request complete	Alexa replies, "Your reservation from Seattle to Denver on Monday, February first is all set."	App displays the result of the request. The user closes the app.

³⁴ <https://developer.amazon.com/en-US/docs/alexa/ask-overviews/what-is-the-alexa-skills-kit.html#:~:text=Skills%20are%20like%20apps%20for,to%20control%20cloud%2Dconnected%20devices>

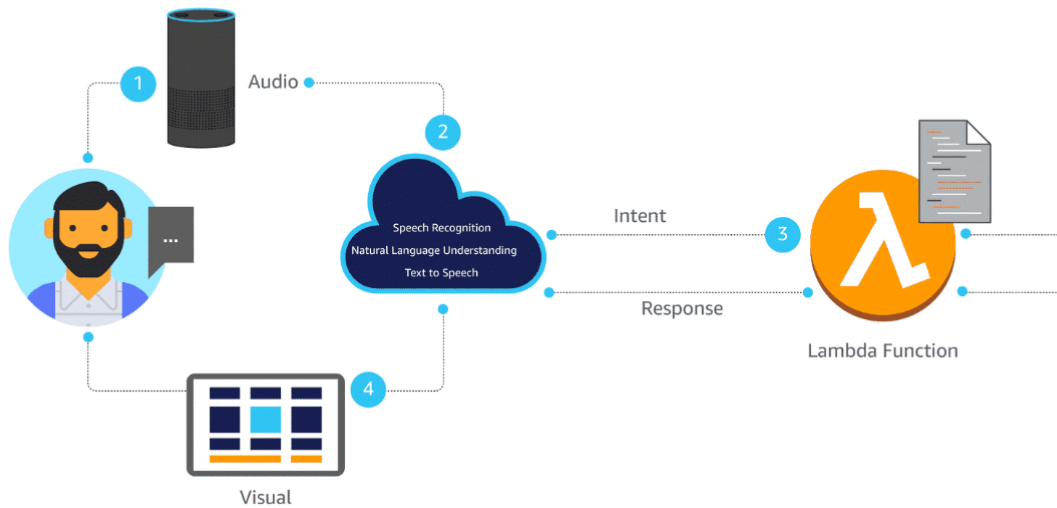
Further, Amazon describes how the Alexa infrastructure supports a user’s interaction with an Alexa skill.³⁵

How an Alexa skill works

The following simple workflow demonstrates how an Alexa skill works. In this example, the user invokes a simple Alexa skill called "Hello World."

1. To launch the skill, the user says, "Alexa, open Hello World." A complete statement or request to Alexa by the user is called an utterance.
2. The Alexa-enabled device sends the utterance to the Alexa service in the cloud. The Alexa service processes the utterance through automatic speech recognition to convert the utterance to text, and then through natural language understanding to recognize the intent of the text.
3. Alexa sends a JavaScript Object Notation (JSON) request to handle the intent to the resource compute service that hosts your skill (usually an AWS Lambda function). The resource compute service acts as the back end and executes your skill code to handle the intent. In this case, the server returns, "Welcome to the Hello World skill."

The following diagram demonstrates what happens when a user interacts with an Alexa skill. It assumes that you are using AWS Lambda's serverless compute service to host your skill code.



117. Amazon has long known about the '720 Patent.

118. Amazon has been aware of the '720 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.

119. Amazon knew or should have known that Amazon’s actions infringe one or more of the claims of the '720 Patent because Amazon has the technical expertise to understand the

³⁵ <https://developer.amazon.com/it-IT/docs/alexa/workshops/build-an-engaging-skill/why-build/index.html>

scope and content of the '720 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Automotive Products and Services, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '720 Patent at least as of the filing of this Complaint.

120. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 1 of the '720 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

121. Users of the Accused Automotive Products and Services directly infringe at least Claim 1 of the '720 Patent when they use the Accused Automotive Products and Services in the ordinary, customary, and intended way.

122. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Automotive Products and Services within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Automotive Products and Services to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Automotive Products and Services in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Automotive Products and Services, which Amazon knew infringes at least Claim 1 of the '720 Patent, or, alternatively, was willfully blind to the infringement.

123. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the

Accused Automotive Products and Services within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Automotive Products and Services in the United States by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Automotive Products and Services, which Amazon knew infringes at least Claim 1 of the '720 Patent, or, alternatively, was willfully blind to the infringement.

124. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 1 of the '720 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '720 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

125. On information and belief, in violation of 35 U.S.C. § 271(f)(1), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 1 of the '720 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

126. On information and belief, in violation of 35 U.S.C. § 271(f)(2), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 1 of the '720 Patent that are

especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

127. Amazon is not licensed or otherwise authorized to practice the claims of the '720 Patent.

128. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '720 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 1.

129. As a result of Amazon's infringement of the '720 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

130. On information and belief, in addition to Amazon's knowledge of the '720 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '720 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '720 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

131. Accordingly, Amazon's infringement of the '720 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

132. Amazon's infringement of Dialect's rights under the '720 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SECOND COUNT
(Infringement of U.S Patent No. 8,015,006)

133. Dialect incorporates by reference the allegations set forth in Paragraphs 1–132 as though fully set forth herein.

134. The claims of the '006 Patent are valid and enforceable.

135. The claims of the '006 Patent are directed to patentable subject matter. Specifically, Claim 5 of the '006 Patent is directed to a novel natural-language speech-recognition system including the claimed novel combination of parsing to determine a meaning and a context of speech associated with a request involving a grammar by a domain agent, satisfying a predetermined confidence level, updating user specific vocabularies or dictionaries, and determining an identity of a user based on voice characteristics. The inventive, tangible claimed structures of the '006 Patent improve on existing natural language processing systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

136. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and

services that embody the invention disclosed and claimed in the '006 Patent, including at least the Accused Products.

137. Each of the Accused Products contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 5 of the '006 Patent.

138. Each of the Accused Products comprises a method for processing natural language speech utterances with context-specific domain agents.

139. For example, Amazon describes Alexa as processing the speech of a customer.³⁶

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



140. Each of the Accused Products comprises receiving, at a speech unit coupled to a processing device, a natural language speech utterance that contains a request.

141. For example, Amazon describes Alexa as receiving a user request at an Alexa device such as an Echo Dot or Echo Show.³⁷

³⁶ <https://developer.amazon.com/en-US/alexa/alexa-skills-kit/get-deeper/custom-skills>

³⁷ *Id.*

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



142. Each of the Accused Products comprises recognizing, at a speech recognition engine coupled to the processing device, one or more words or phrases contained in the utterance using information in one or more dictionary and phrase tables.

143. For example, Amazon describes Alexa as processing the user request with speech recognition and natural language understanding.³⁸

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



Amazon further describes Alexa as using a “wide range of sentences, phrases, and words that users are likely to say.”³⁹

³⁸ *Id.*

³⁹ <https://developer.amazon.com/en-US/alexa/alexa-haus/voice-fundamentals>

Match a variety of utterances

Don't assume that users say the exact phrase that you anticipate for an intent. While the user might say, "Plan a trip," they also might say "Plan a vacation to Hawaii." To make sure your skill can respond to a variety of user utterances, provide a wide range of sentences, phrases, and words that users are likely to say.

144. Each of the Accused Products comprises parsing, at a parser coupled to the processing device, information relating to the utterance to determine a meaning associated with the utterance and a context associated with the request contained in the utterance, wherein the parsed information includes the one or more recognized words or phrases.

145. For example, Amazon describes Alexa as using intents, slots, and sample utterances for a skill, where an intent represents an action that fulfills a request.⁴⁰

Create Intents, Utterances, and Slots

Note: Sign in to the [developer console](#) to build or publish your skill.

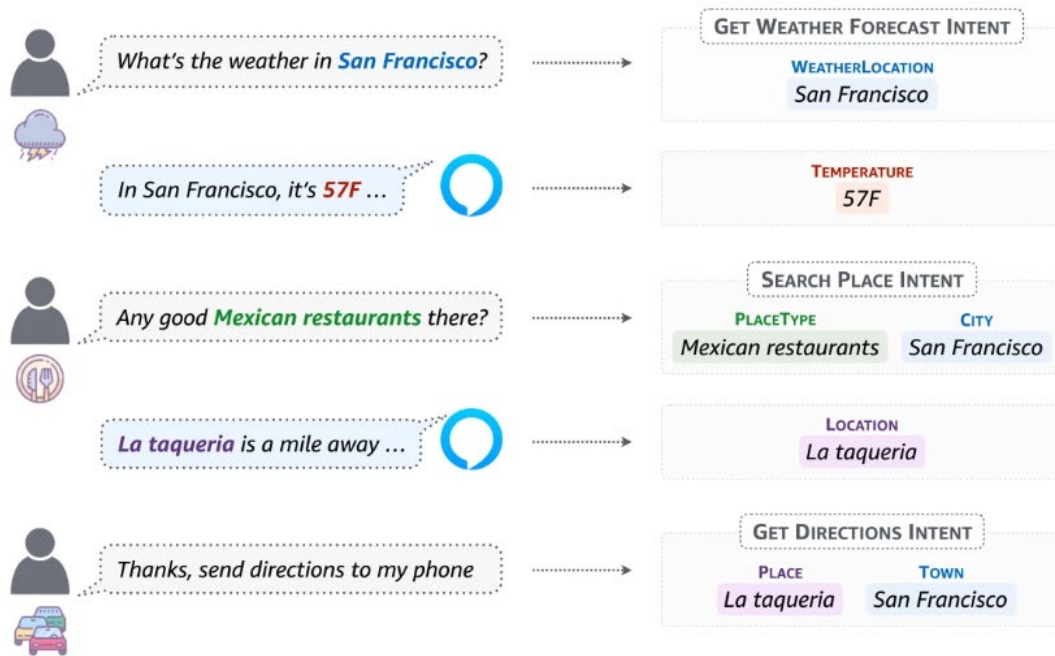
To create your custom interaction model, you create the *intents*, *slots*, and *sample utterances* for your skill. An *intent* represents an action that fulfills a user's spoken request. Intents can optionally have arguments called *slots*. The *sample utterances* are set of likely spoken phrases mapped to the intents.

For details and code examples that show how you access intents and slots sent in user requests during skill runtime, see [Handle Requests Sent by Alexa](#).

Amazon further describes Alexa intents assigning identified data to different slots based on the intent.⁴¹

⁴⁰ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/create-intents-utterances-and-slots.html>

⁴¹ <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

Amazon's Alexa Machine Learning team further describes a system for referring back to slots in the context during a conversation.⁴²

Abstract

In the slot-filling paradigm, where a user can refer back to slots in the context during a conversation, the goal of the contextual understanding system is to resolve the referring expressions to the appropriate slots in the context. In large-scale multi-domain systems, this presents two challenges - scaling to a very large and potentially unbounded set of slot values, and dealing with diverse schemas. We present a neural network architecture that addresses the slot value scalability challenge by reformulating the contextual interpretation as a decision to *carryover* a slot from a set of possible candidates. To deal with heterogeneous schemas, we introduce a simple data-driven method for transforming the candidate slots. Our experiments show that our approach can scale to multiple domains and provides competitive results over a strong baseline.

Index Terms: spoken dialog, state tracking, anaphora resolution, contextual understanding.

⁴² "Contextual Slot Carryover for Disparate Schemas," Naik (Amazon) et al. 2018.

146. Each of the Accused Products comprises formulating, at the parser, the request contained in the utterance in accordance with a grammar used by a domain agent associated with the determined context.

147. For example, Amazon describes Alexa as formulating requests contained in the user utterance in accordance with Alexa intents.⁴³



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

Amazon further describes how grammar has “been a tool in Alexa’s NLU toolkit since well before the first Echo device shipped.”⁴⁴

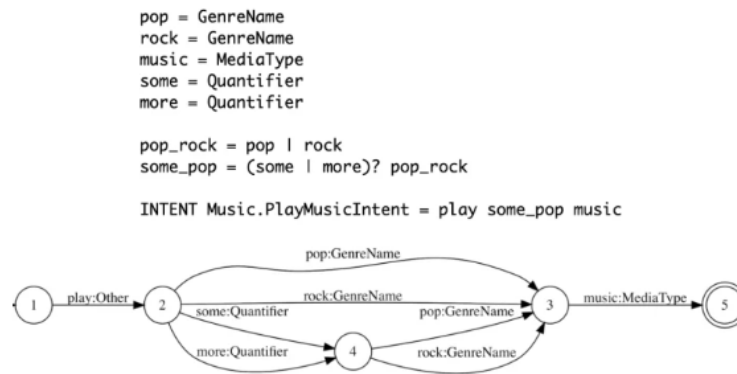
⁴³ <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>

⁴⁴ <https://www.amazon.science/blog/tools-for-generating-synthetic-data-helped-bootstrap-alexa-s-new-language-releases>

Rules of Grammar

Grammars have been a tool in Alexa’s NLU toolkit since well before the first Echo device shipped. A grammar is a set of rewrite rules for varying basic template sentences through word insertions, deletions, and substitutions.

Below is a very simple grammar, which models requests to play either pop or rock music, with or without the modifiers “more” and “some”. Below the rules of the grammar is a diagram of a computational system (a finite-state transducer, or FST) that implements them.



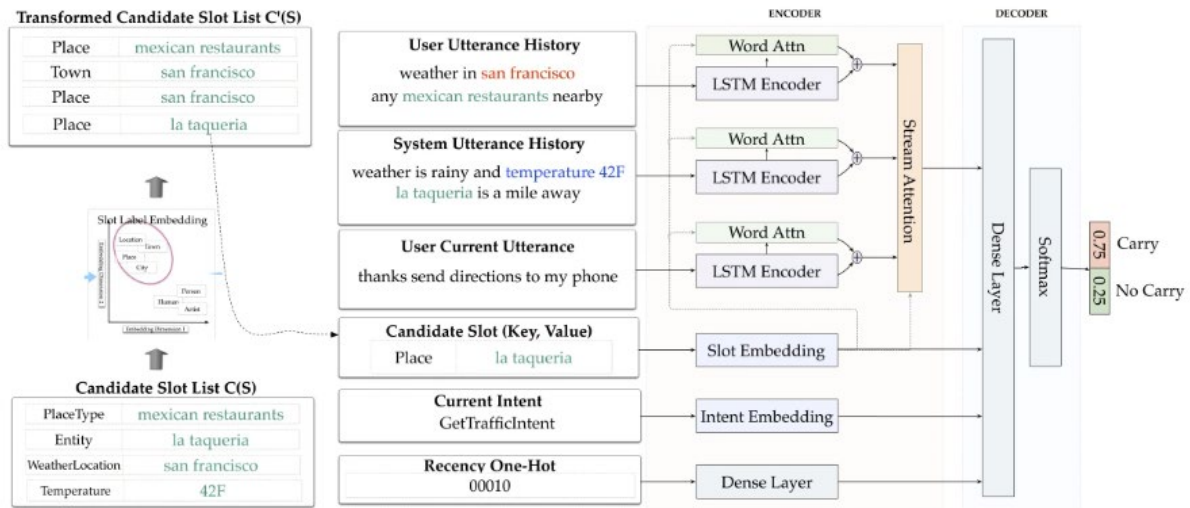
A toy grammar, which can model requests to play pop or rock music, with or without the modifiers “some” or “more”, and a diagram of the resulting finite-state transducer. The question mark indicates that the some_more variable is optional.

148. Each of the Accused Products comprises determining one or more required values and one or more optional values associated with formulating the request in the grammar used by the domain agent.

149. For example, Amazon describes Alexa as supporting slots that are variables that a function acts upon.⁴⁵

⁴⁵ <https://www.amazon.science/blog/how-alexas-learning-to-converse-more-naturally>

Today, Alexa analyzes the semantic content of utterances according to the categories domain, intent, and slot. "Domain" describes the type of application — or "skill" — that the utterance should invoke; for instance, *mapping* skills should answer questions about geographic distance. "Intent" describes the particular function the skill should execute, such as measuring driving distance. And "slots" are variables that the function acts upon, such as point of origin and destination.

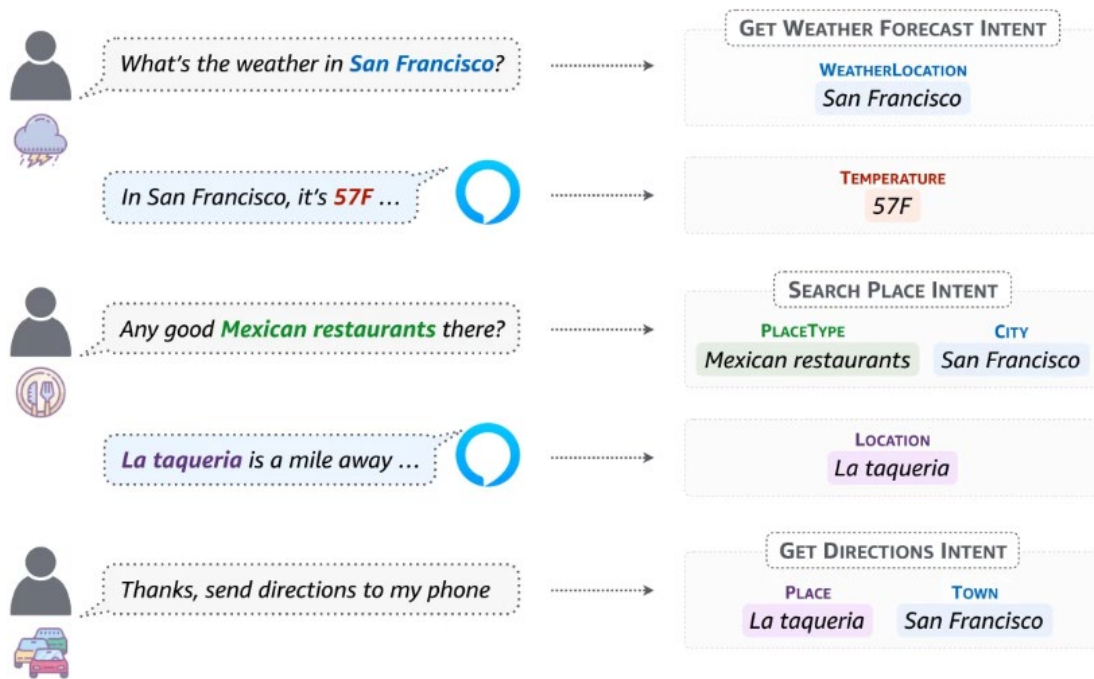


Our system architecture

150. Each of the Accused Products comprises extracting one or more criteria and one or more parameters from one or more keywords contained in the one or more recognized words or phrases, wherein the parser extracts the one or more criteria and the one or more parameters using procedures sensitive to the determined context.

151. For example, Amazon describes Alexa as identifying and assigning keywords into different slots based on the intent.⁴⁶

⁴⁶ <https://www.amazon.science/blog/how-alex-a-is-learning-to-converse-more-naturally>

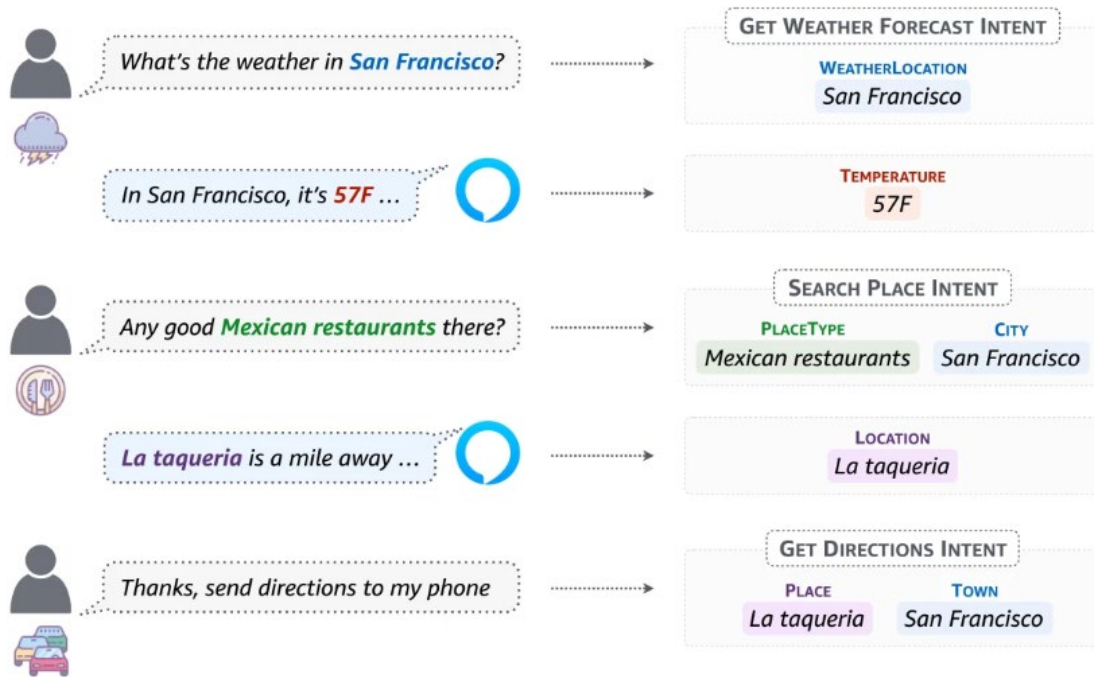


Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

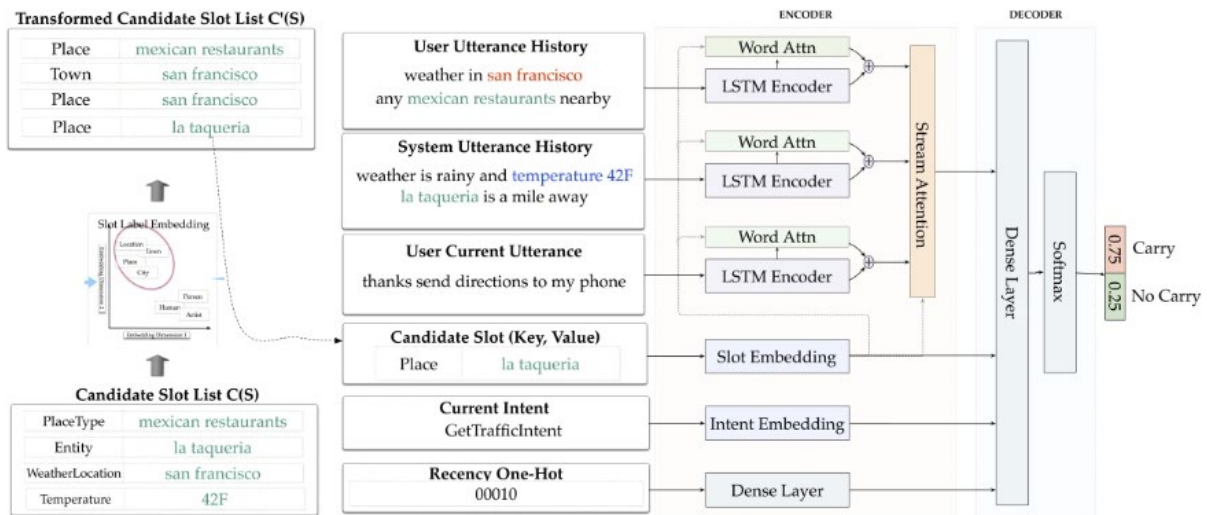
152. Each of the Accused Products comprises inferring one or more further criteria and one or more further parameters associated with the request using a dynamic set of prior probabilities or fuzzy possibilities.

153. For example, Amazon describes Alexa as inferring parameters for a request using the history of the current interaction.⁴⁷

⁴⁷ <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.



Our system architecture

Amazon further states that their approach makes decisions “about slot values mentioned in context” and “the probability that any given carryover decision is the correct one.”⁴⁸

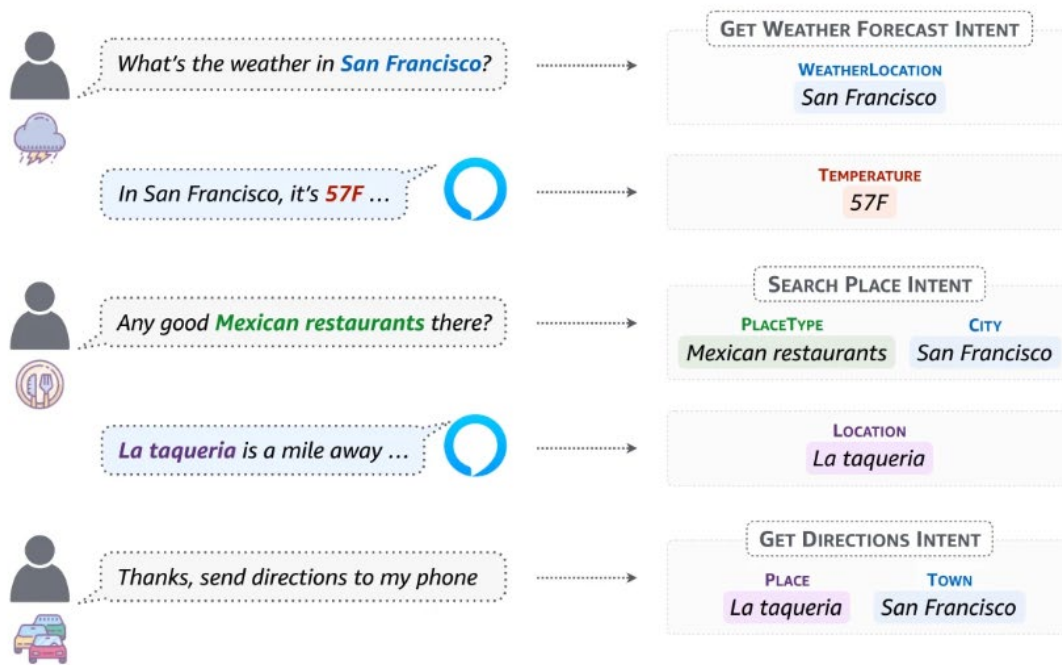
⁴⁸ *Id.*

We make an important distinction between our approach and conventional dialogue state tracking, which maintains a probability distribution across all possible values that a given target slot can take on. Our system, by contrast, (i) cares only about slot values mentioned in context and (ii) makes independent decisions about the probability that any given carryover decision is the correct one. This helps us scale to the large number of slots across all Alexa skills.

154. Each of the Accused Products comprises transforming the one or more extracted criteria, the one or more extracted parameters, the one or more inferred criteria, and the one or more inferred parameters into one or more tokens having a format compatible with the grammar used by the domain agent, wherein the one or more tokens include all the required values and one or more of the optional values associated with formulating the request in the grammar used by the domain agent.

155. For example, Amazon describes Alexa as formulating requests based on the combination of data assigned to slots based on the intent.⁴⁹

⁴⁹ <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

156. Each of the Accused Products comprises processing the formulated request with the domain agent associated with the determined context to generate a response to the utterance.

157. For example, Amazon describes Alexa as processing a skill interaction model and skill application logic to produce a response.⁵⁰

⁵⁰ <https://developer.amazon.com/en-US/alexa/alexa-skills-kit/get-deeper/custom-skills>

How Custom Skills Work

When you build a custom skill, you define your own interaction model, or voice use interface, rather than relying on one provided by Alexa. Your custom interaction model specifies how customers' spoken input maps to the requests, or intents, that your skill can handle. When a customer speaks to your skill, Alexa uses your interaction model to determine the underlying customer intent, and sends it to your skill application logic.



158. Each of the Accused Products comprises presenting the generated response to the utterance via the speech unit.

159. For example, Amazon describes Alexa as providing text to speech services to generate the response to the user’s request.⁵¹

How Custom Skills Work

When you build a custom skill, you define your own interaction model, or voice use interface, rather than relying on one provided by Alexa. Your custom interaction model specifies how customers' spoken input maps to the requests, or intents, that your skill can handle. When a customer speaks to your skill, Alexa uses your interaction model to determine the underlying customer intent, and sends it to your skill application logic.



160. Amazon has long known about the '006 Patent.

161. Amazon has been aware of the '006 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.

⁵¹ *Id.*

162. Amazon knew or should have known that Amazon's actions infringe one or more of the claims of the '006 Patent because Amazon has the technical expertise to understand the scope and content of the '006 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Products, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '006 Patent at least as of the filing of this Complaint.

163. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 5 of the '006 Patent in violation of at least 35 U.S.C. § 271(b) and (c).

164. Users of the Accused Products directly infringe at least Claim 5 of the '006 Patent when they use the Accused Products in the ordinary, customary, and intended way.

165. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Products in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Products, which Amazon knew infringes at least Claim 5 of the '006 Patent, or, alternatively, was willfully blind to the infringement.

166. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the

Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Products in the United States by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Products, which Amazon knew infringes at least Claim 5 of the '006 Patent, or, alternatively, was willfully blind to the infringement.

167. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 5 of the '006 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '006 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

168. Amazon is not licensed or otherwise authorized to practice the claims of the '006 Patent.

169. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '006 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 5.

170. As a result of Amazon's infringement of the '006 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

171. On information and belief, in addition to Amazon's knowledge of the '006 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and

continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '006 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '006 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

172. Accordingly, Amazon's infringement of the '006 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

173. Amazon's infringement of Dialect's rights under the '006 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

THIRD COUNT
(Infringement of U.S Patent No. 8,140,327)

174. Dialect incorporates by reference the allegations set forth in Paragraphs 1–173 of the Complaint as though fully set forth herein.

175. The claims of the '327 Patent are valid and enforceable.

176. The claims of the '327 Patent are directed to patentable subject matter. Specifically, Claim 14 of the '327 Patent is directed to a novel system for eliminating noise from natural language speech utterances including the claimed novel microphone array, adaptive filter, speech coder, and transceiver. The inventive, tangible claimed structures of the '327 Patent improve on

existing natural language processing input systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

177. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and services that embody the invention disclosed and claimed in the '327 Patent, including at least the Accused Products.

178. Each of the Accused Products contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 14 of the '327 Patent.

179. Each of the '327 Patent Accused Products comprises a system for filtering and eliminating noise from natural language speech utterances.

180. For example, Amazon describes Alexa devices as using the same technology as the Amazon Alexa Premium Voice Far-Field Development Kit.⁵²

⁵² <https://developer.amazon.com/blogs/alexa/post/80facfd2-1176-4c4f-94ac-4c5c781011ca/amazon-alexa-premium-far-field-voice-development-kit#:~:text=The%20Amazon%20Alexa%20Premium%20Far%2DField%20Voice%20Development%20Kit%20uses,playing%20music%20at%20loud%20volumes>

PRODUCT BRIEF

Amazon Alexa Premium Voice Far-Field Development Kit



Engineered for Performance

The Amazon Alexa Premium Voice Far-Field Development Kit is designed to help commercial device manufacturers enable high quality far-field voice experiences in their products. This development kit gives manufacturers access to the audio technology that powers the Amazon Echo family of devices. It supports both tabletop 360° far-field voice activation applications, as well as applications that require voice-activation from one direction, such as devices with a display or placed near a wall. Both solutions incorporate Amazon's proprietary software and algorithm technology for "Alexa" wake word recognition, beam forming, noise reduction, and acoustic echo cancellation, and provides accurate far-field voice recognition in noisy environments and from long distances.

Turnkey Kit for Evaluation, Prototyping, and Engineering Development

This development kit simplifies the task of designing high-quality audio front ends for voice processing systems, allowing OEMs to focus on core product functionality. Included in the kit are two microphone array boards, a digital signal processor board which contains the audio processing hardware to optimize the quality of the voice signal, and a Raspberry Pi 3[®] single-board computer with the AVS Device SDK to handle interactions with the AVS cloud. Simply register the device, add a set of powered speakers, and download and install the software.

Features and Benefits

Premium Far-Field Performance


- 7- and 8- microphone arrays optimized for premium far-field audio performance in devices with 360° and 180° use cases
- Amazon audio processing and wake word trigger algorithms provide excellent voice recognition performance, even in noisy environmental conditions

Easy to Integrate

- Modular "far-field in a chip" architecture incorporates the audio processing algorithms and wake word trigger engine on the digital signal processor, to simplify integration with the device processor and software system stack
- Raspberry Pi 3[®] single-board computer can be replaced with production ready SoC for system prototyping and product development
- AVS Device SDK simplifies development of production-ready AVS client software

Amazon further describes this technology as allowing devices to pick up voice requests in noisy environments and while playing music at loud volumes.⁵³

⁵³ <https://developer.amazon.com/en-US/alexa/solution-providers/dev-kits/amazon-premium-voice>



The image shows the Amazon Alexa Premium Far-Field Voice Development Kit. It is a transparent plastic enclosure housing a printed circuit board (PCB) with various components. A brown leather-like strap is attached to the top. To the left of the device is a logo consisting of a grid of blue circles with a white 'X' in the center, and the text 'Far-Field' below it.

Amazon Alexa Premium Far-Field Voice Dev Kit

With the Amazon Alexa Premium Far-Field Voice Development Kit, Amazon is helping commercial device manufacturers more efficiently build high-quality Alexa voice experiences. Developers can integrate Alexa into products with the dev kit's simplified architecture of "far-field in a chip": premium-level voice technology on a single digital signal processor (DSP). The kit is based on the same premium-level voice recognition technology found in the latest Amazon Echo family of devices. Products built with the dev kit can pick up voice requests in noisy environments and while playing music at loud volumes.

- Far-Field Voice Recognition
- Amazon wake word engine (WWE) and audio processing algorithms running on Intel's dual DSP with inference engine
- 7-mic circular array and 8-mic rectangular array configurations from the Echo and Echo Show, respectively

Please speak with your Amazon AVS business development contact, or fill out the [request form](#) below, if interested in requesting the Amazon Alexa Premium Far-Field Voice Development Kit.

[See All Solutions »](#)

181. Each of the '327 Patent Accused Products comprises a microphone array configured to add one or more nulls to a beam pattern steered to point in a direction associated with a user speaking a natural language utterance to capture an input speech signal corresponding to the natural language utterance, wherein the one or more nulls notch out point or limited area noise sources from the input speech signal.

182. Amazon describes the technology of Alexa devices as using a microphone array that supports beam forming.⁵⁴

⁵⁴ <https://developer.amazon.com/blogs/alexa/post/80facfd2-1176-4c4f-94ac-4c5c781011ca/amazon-alexa-premium-far-field-voice-development-kit#:~:text=The%20Amazon%20Alexa%20Premium%20Far%2DField%20Voice%20Development%20Kit%20uses,playing%20music%20at%20loud%20volumes>

PRODUCT BRIEF

Amazon Alexa Premium Voice Far-Field Development Kit



Engineered for Performance

The Amazon Alexa Premium Voice Far-Field Development Kit is designed to help commercial device manufacturers enable high quality far-field voice experiences in their products. This development kit gives manufacturers access to the audio technology that powers the Amazon Echo family of devices. It supports both tabletop 360° far-field voice activation applications, as well as applications that require voice-activation from one direction, such as devices with a display or placed near a wall. Both solutions incorporate Amazon's proprietary software and algorithm technology for "Alexa" wake word recognition, beam forming, noise reduction, and acoustic echo cancellation, and provides accurate far-field voice recognition in noisy environments and from long distances.

Turnkey Kit for Evaluation, Prototyping, and Engineering Development

This development kit simplifies the task of designing high-quality audio front ends for voice processing systems, allowing OEMs to focus on core product functionality. Included in the kit are two microphone array boards, a digital signal processor board which contains the audio processing hardware to optimize the quality of the voice signal, and a Raspberry Pi 3[®] single-board computer with the AVS Device SDK to handle interactions with the AVS cloud. Simply register the device, add a set of powered speakers, and download and install the software.

Features and Benefits

Premium Far-Field Performance

- 7- and 8- microphone arrays optimized for premium far-field audio performance in devices with 360° and 180° use cases
- Amazon audio processing and wake word trigger algorithms provide excellent voice recognition performance, even in noisy environmental conditions

Easy to Integrate

- Modular "far-field in a chip" architecture incorporates the audio processing algorithms and wake word trigger engine on the digital signal processor, to simplify integration with the device processor and software system stack
- Raspberry Pi 3[®] single-board computer can be replaced with production ready SoC for system prototyping and product development
- AVS Device SDK simplifies development of production-ready AVS client software

Amazon's applied scientist in the Alexa Speech group specifically discussed the beam former for dealing with the noise robustness of the system.⁵⁵

Your group has three papers at ICASSP this year. What are they about?

Two are about multichannel acoustic modeling. Right now, we are using seven microphones, and we beamform the signal to do acoustic modeling. We optimize the beam former based on speech quality, but we optimize the acoustic model based on word error rate.

With our new work, we are incorporating beamforming into the acoustic model, so we are jointly optimizing both the beam former and the acoustic model, and both of them are optimized toward reducing word error rate.

The other paper is about improving the conventional beamforming system. The challenge we are dealing with is the noise robustness of the system. Under clean conditions, ASR is already performing very well, but we still have problems dealing with noisy conditions. So the main idea is to train a teacher model on the clean data, then artificially add noise — music noise, audio playback noise, household noise.

Then we use the noisy data to train the student model, paired with the teacher's predictions from the clean data. We are hoping to create a student that performs similarly to the teacher but in the noisy domain.

Another example is in beamforming. You can do traditional beamforming where you combine the microphones with filters that affect the amplitude and phase to focus on sound from a specific direction. But if that's all you do, then you still have to figure out what direction to point it in. So we train it on the sound of interest. This is the speech that we want to listen to, so when it combines the microphones together, it can combine them in a way that optimizes for the specific direction of the speech and the overall recognition end to end.

Amazon further describes in its documentation of Alexa “Hardware Configurations” how beamforming results “in an increase in SNR and a reduction in reverberation in the audio signal.”⁵⁶

Beamforming

Beamforming is a signal processing technique for multi-microphone arrays that emphasizes user speech from a desired direction when suppressing audio interference from other directions. These algorithms result in an increase in SNR and a reduction in reverberation in the audio signal from the desired direction that improves the accuracy of speech recognition systems, especially for far-field. For ASR, always use linear processing based versions of beamforming.

Amazon has further published papers about the SIR Beam Selector for Amazon Echo Devices.⁵⁷

SIR Beam Selector for Amazon Echo Devices Audio Front-End

Xianxian Zhang, Trausti Kristjansson, Philip Hilmes

Amazon Inc., Sunnyvale, California, USA
{xianxi, traustik, philmes}@amazon.com

Amazon further describes how “Amazon scientists recommend car manufacturers use a signal-processing technique called beamforming that steers microphone arrays toward a voice signal.”⁵⁸

In design meetings, Amazon scientists collaborate with car manufacturers on techniques that allow the car’s built-in system to focus on the person speaking to Alexa and cancel out acoustic echo and background noise.

For example, Amazon scientists recommend car manufacturers use a signal-processing technique called [beamforming](#) that steers microphone arrays toward a voice signal coming from a particular source, such as the driver, while suppressing audio interference from all other directions.

183. Each of the ’327 Patent Accused Products comprises an adaptive filter coupled to the microphone array, wherein the adaptive filter is configured to: receive the input speech signal corresponding to the natural language utterance from the microphone array and compare

environmental noise to the input speech signal to set one or more parameters associated with the adaptive filter.

184. Amazon’s applied scientist in the Alexa Speech group specifically discussed using an adaptive filter for echo cancellation changes based on what it’s hearing.⁵⁹

How do you combine signal processing and machine learning?

Acoustic echoes are handled really well by traditional signal-processing techniques, but we may use machine learning to contrive a control mechanism to figure out how to adapt that signal processing to a specific situation. So if you have an Echo device that’s playing music, we use acoustic-echo cancellation filters that have learned what is playing, and then, listening to the microphones, it uses that reference to directly cancel the audio. But its performance is very dependent on how it’s adapting the filter to cancel the audio.

So we use machine learning algorithms to understand what’s going on in the environment and then adapt our canceling the sound. The adaptive filter for echo cancellation changes based on what it’s hearing, but you can tell it to change or you can tell it, no, don’t change. So we listen to understand, for instance, somebody’s trying to talk to Alexa. Let’s freeze the filter so that it doesn’t change and try to cancel what the person is saying.

Amazon further describes the AEC algorithm that “estimates the acoustic echo path and the acoustic echo between the loudspeaker and microphone components,” which “is then subtracted from the microphone signal.”⁶⁰

In speech recognition systems, the term "acoustic echo" refers to the signal that a loudspeaker plays and that a microphone captures from the loudspeaker. The acoustic echo is a source of interference for the ASR engine because the microphone simultaneously captures the echo and the user utterance. The goal of AEC is to remove the acoustic echo component from the microphone signal so that the ASR engine accurately understands the user utterance. The AEC algorithm adaptively estimates the acoustic echo path and the acoustic echo between the loudspeaker and microphone components. The estimated acoustic echo is then subtracted from the microphone signal to obtain a near echo-free microphone signal. An AEC-processed microphone signal should be free from acoustic echo.

185. Each of the ’327 Patent Accused Products comprises an adaptive filter configured to use band shaping and notch filtering to remove narrow-band noise from the input speech signal received from the microphone array according to the one or more parameters.

⁵⁶ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/audio-hardware-configurations.html>

⁵⁷ <https://assets.amazon.science/da/c2/71f5f9fa49f585a4616e49d52749/sir-beam-selector-for-amazon-echo-devices-audio-front-end.pdf>

⁵⁸ <https://www.amazon.science/news-and-features/the-science-behind-alexa-in-vehicles>

⁵⁹ https://www.amazon.jobs/fr/landing_pages/icassp

⁶⁰ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/audio-hardware-configurations.html>

186. Amazon's applied scientist in the Alexa Speech group specifically discussed using an adaptive filter for echo cancellation changes.⁶¹

How do you combine signal processing and machine learning?

Acoustic echoes are handled really well by traditional signal-processing techniques, but we may use machine learning to contrive a control mechanism to figure out how to adapt that signal processing to a specific situation. So if you have an Echo device that's playing music, we use acoustic-echo cancellation filters that have learned what is playing, and then, listening to the microphones, it uses that reference to directly cancel the audio. But its performance is very dependent on how it's adapting the filter to cancel the audio.

So we use machine learning algorithms to understand what's going on in the environment and then adapt our canceling the sound. The adaptive filter for echo cancellation changes based on what it's hearing, but you can tell it to change or you can tell it, no, don't change. So we listen to understand, for instance, somebody's trying to talk to Alexa. Let's freeze the filter so that it doesn't change and try to cancel what the person is saying.

187. Each of the '327 Patent Accused Products comprises an adaptive filter configured to suppress cross-talk and environmentally caused echoes in the input speech signal received from the microphone array using adaptive echo cancellation.

188. Amazon's applied scientist in the Alexa Speech group specifically discussed using an adaptive filter for echo cancellation changes.⁶²

How do you combine signal processing and machine learning?

Acoustic echoes are handled really well by traditional signal-processing techniques, but we may use machine learning to contrive a control mechanism to figure out how to adapt that signal processing to a specific situation. So if you have an Echo device that's playing music, we use acoustic-echo cancellation filters that have learned what is playing, and then, listening to the microphones, it uses that reference to directly cancel the audio. But its performance is very dependent on how it's adapting the filter to cancel the audio.

So we use machine learning algorithms to understand what's going on in the environment and then adapt our canceling the sound. The adaptive filter for echo cancellation changes based on what it's hearing, but you can tell it to change or you can tell it, no, don't change. So we listen to understand, for instance, somebody's trying to talk to Alexa. Let's freeze the filter so that it doesn't change and try to cancel what the person is saying.

⁶¹ https://www.amazon.jobs/fr/landing_pages/icassp

⁶² https://www.amazon.jobs/fr/landing_pages/icassp

Amazon further describes in its documentation of Alexa “Hardware Configurations” how “Acoustic Echo Cancellation” subtracts acoustic echo “from the microphone signal to obtain a near echo-free microphone signal.”⁶³

Acoustic Echo Cancellation (AEC)

In speech recognition systems, the term "acoustic echo" refers to the signal that a loudspeaker plays and that a microphone captures from the loudspeaker. The acoustic echo is a source of interference for the ASR engine because the microphone simultaneously captures the echo and the user utterance. The goal of AEC is to remove the acoustic echo component from the microphone signal so that the ASR engine accurately understands the user utterance. The AEC algorithm adaptively estimates the acoustic echo path and the acoustic echo between the loudspeaker and microphone components. The estimated acoustic echo is then subtracted from the microphone signal to obtain a near echo-free microphone signal. An AEC-processed microphone signal should be free from acoustic echo.

Amazon further describes how it recommends a “technique called acoustic echo cancellation” to car manufacturers for integration of Alexa.⁶⁴

Another recommended technique called acoustic echo cancellation blocks out the noise from the car’s speakers so that music played by the speaker does not interfere with the customer requests. Manufacturers can also deploy various noise reduction algorithms, which learn to separate a speaker’s voice from other background noise.

Inside the vehicle lab, Amazon scientists use software and acoustic tools to simulate road noise, music, and other acoustic elements as they test and validate Alexa integration with the car manufacturer’s built-in systems.

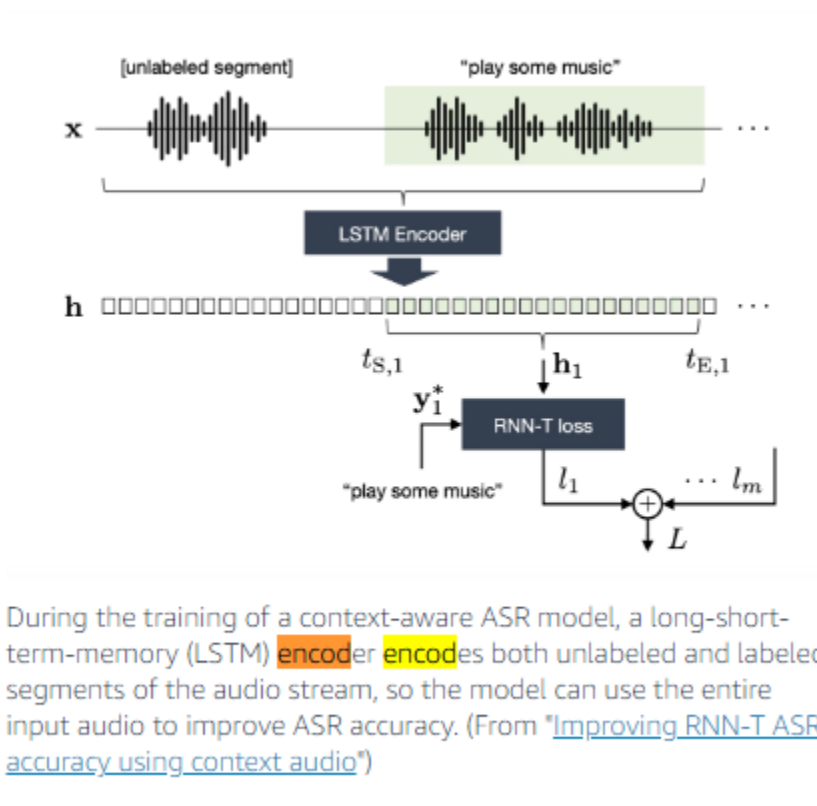
189. Each of ’327 Patent Accused Products comprises a speech coder arranged between the adaptive filter and a speech recognition engine, wherein the speech coder is configured to receive the input speech signal passed through the adaptive filter and use adaptive lossy audio compression to remove momentary gaps from the input speech signal and variable rate sampling

⁶³ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/audio-hardware-configurations.html>

⁶⁴ <https://www.amazon.science/news-and-features/the-science-behind-alexa-in-vehicles>

to compress and digitize the input speech signal, wherein the speech coder optimizes the adaptive lossy audio compression and the variable rate sampling to only preserve components in the input speech signal that will be input to the speech recognition engine.

190. For example, Amazon describes Alexa encoding segments of the audio stream.⁶⁵



Furthermore, Amazon describes the compression techniques used by Alexa.⁶⁶

On top of these innovations, however, we still had to develop some new compression techniques to get the RNN-T to run efficiently on-device. A neural network consists of simple processing nodes each of which is connected to several others. The connections between nodes have associated *weights*, which determine how much one node's output contributes to the computation performed by the next node.

⁶⁵ <https://www.amazon.science/blog/on-device-speech-processing-makes-alexa-faster-lower-bandwidth>

⁶⁶ *Id.*

191. Each of '327 Patent Accused Products comprises a transceiver configured to communicate the digitized input speech signal from a buffer in the speech coder to the speech recognition engine at a rate that depends on available bandwidth associated with a communication link that connects the transceiver and the speech recognition engine.

192. For example, Amazon describes Alexa sending audio streams to the cloud in frames for off-device processing.⁶⁷

With cloud-based ASR, encrypted audio streams to the cloud in small snippets called "frames". With on-device ASR, only the lattice is sent to the cloud, where a large and powerful neural language model reranks the hypotheses. The lattice can't be sent until the customer has finished speaking, as words later in a sequence can dramatically change the overall probability of a hypothesis.

Amazon further describes how Alexa uses Opus to send captured audio to the Alexa Voice Service (AVS).⁶⁸

Codec specifications

Send all captured audio to AVS in either PCM or Opus, and adhere to the following specifications:

Specification	PCM	Opus
Number of channels	Single channel (mono)	Single channel (mono)
Sample size	16-bit linear PCM (LPCM)	16 bit
Sample rate	16 kHz	16 kHz
Bitrate	256 Kpbs	32 Kpbs or 64 Kpbs, hard constant bitrate
Byte order	Little endian	Little endian
Frame size	10 ms	20 ms
DATA frame size	320 bytes	80 bytes (32 Kpbs) or 160 bytes (64 Kpbs)
Complexity	N/A	= 4

⁶⁷ <https://www.amazon.science/blog/on-device-speech-processing-makes-alexa-faster-lower-bandwidth>

⁶⁸ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/speechrecognizer.html>

Furthermore, the Opus codec is “highly tuned to allow it to quickly optimize between optimal fidelity for the available bandwidth.”⁶⁹

Opus is highly tuned to allow it to quickly optimize between optimal fidelity for the available bandwidth. It also has a variable bitrate (VBR) mode that can range from 6 kbit/s to 510 kbit/s to help minimize bandwidth while maintaining consistent quality for applications that require that.

Furthermore, the Internet Engineering Task Force (IETF) RFC 6716—Definition of the Opus Audio Codec—specifies that an Opus encoder “should select which coding mode to use at run-time depending on the conditions.”⁷⁰

For a normal encoder where both the SILK and the CELT modules are included, an optimal encoder should select which coding mode to use at run-time depending on the conditions. In the reference implementation, the frame size is selected by the application, but the other configuration parameters (number of channels, bandwidth, mode) are automatically selected (unless explicitly overridden by the application) depending on the following:

- o Requested bitrate
- o Input sampling rate
- o Type of signal (speech vs. music)
- o Frame size in use

193. Amazon has long known about the '327 Patent.

194. Amazon has been aware of the '327 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.

195. Amazon knew or should have known that Amazon’s actions infringe one or more of the claims of the '327 Patent because Amazon has the technical expertise to understand the

⁶⁹ <https://www.wowza.com/blog/opus-codec-the-audio-format-explained>

⁷⁰ <https://datatracker.ietf.org/doc/html/rfc6716>

scope and content of the '327 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Products, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '327 Patent at least as of the filing of this Complaint.

196. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 14 of the '327 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

197. Users of the Accused Products directly infringe at least Claim 14 of the '327 Patent when they use the Accused Products in the ordinary, customary, and intended way.

198. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Products in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Products, which Amazon knew infringes at least Claim 14 of the '327 Patent, or, alternatively, was willfully blind to the infringement.

199. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of

infringement with respect to the Accused Products in the United States by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Products, which Amazon knew infringes at least Claim 14 of the '327 Patent, or, alternatively, was willfully blind to the infringement.

200. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 14 of the '327 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '327 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

201. On information and belief, in violation of 35 U.S.C. § 271(f)(1), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 14 of the '327 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

202. On information and belief, in violation of 35 U.S.C. § 271(f)(2), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 14 of the '327 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and

intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

203. Amazon is not licensed or otherwise authorized to practice the claims of the '327 Patent.

204. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '327 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 14.

205. As a result of Amazon's infringement of the '327 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

206. On information and belief, in addition to Amazon's knowledge of the '327 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '327 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '327 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

207. Accordingly, Amazon's infringement of the '327 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case

warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

208. Amazon's infringement of Dialect's rights under the '327 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

FOURTH COUNT
(Infringement of U.S Patent No. 8,195,468)

209. Dialect incorporates by reference the allegations set forth in Paragraphs 1–208 of the Complaint as though fully set forth herein.

210. The claims of the '468 Patent are valid and enforceable.

211. The claims of the '468 Patent are directed to patentable subject matter. Specifically, Claim 19 of the '468 Patent is directed to a novel method for processing multimodal natural language inputs using the claimed speech recognition engine and semantic knowledge-based model, context stack and domain agent. The inventive, tangible claimed structures of the '468 Patent improve on existing natural language processing systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

212. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and services that embody the invention disclosed and claimed in the '468 Patent, including at least the Alexa Products, operating via tablets, smartphones, or other devices supporting non-speech inputs such as the Echo Show (collectively, the "Accused Multi-modal Products and Services")..

213. Each of the Accused Multi-modal Products and Services contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 19 of the '468 Patent.

214. Each of the Accused Multi-modal Products and Services comprises a method for processing multi-modal natural language inputs.

215. For example, Amazon describes Alexa as processing natural language requests from a customer via multi-modal input systems such as the Echo Show.⁷¹

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



Amazon further describes the Alexa Multimodal Response Builder.⁷²

⁷¹ <https://developer.amazon.com/en-US/alexa/alexa-skills-kit/start>

⁷² <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-authoring-tool.html#use-mmrbr>

Use the Multimodal Response Builder

The Multimodal Response Builder provides a guided experience to create a visual response in three steps:

1. Select from a set of templates designed to look good across a broad range of devices.
2. Customize the response by specifying the content to display within the template
3. Preview the response both in the developer console and on a device.

After you finish customizing the response, the Multimodal Response Builder generates a code example you can copy into a request handler in your skill to display the response. Rebuild the interaction model for your skill and then test your skill on a device or with the developer console simulator.

You can create a new document in the Multimodal Response Builder to get started, and then edit the document in the [full authoring tool](#) if you want to do more complex customizations.

To open the Multimodal Response Builder

1. In the developer console, open the skill for which you want to create this document.
2. In the left-hand navigation, click **Multimodal Responses**.
3. Click **Create with Response Builder**.
4. Step through the pages to complete the three steps
 - Select a template. The set of templates includes [responsive templates](#) as well as other visual designs.
 - Customize
 - Preview and test

Note: Populating the primary text or title fields of a template with long strings might create a poor visual experience. Make sure you test your response across different viewport and make sure all of your content displays as you expect.

Amazon further describes the Alexa Multi-Modal API for supporting multi-modal inputs.⁷³

⁷³ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/avs-apl-overview.html>

Alexa Multi-Modal API Overview

Note: Learn how to improve your skills with APL with [Build visually rich experiences using APL](#) at the Alexa Learning Lab.

The Alexa Voice Service (AVS) provides a collection of APIs for Alexa Built-in devices with visual displays, such as smart TVs, which allow those devices to receive and present Alexa visual responses. Some of these APIs help device makers to publish the visual characteristics and functionality of their visual displays to enable Alexa Skills and services to respond with appropriate visual content for a device display.

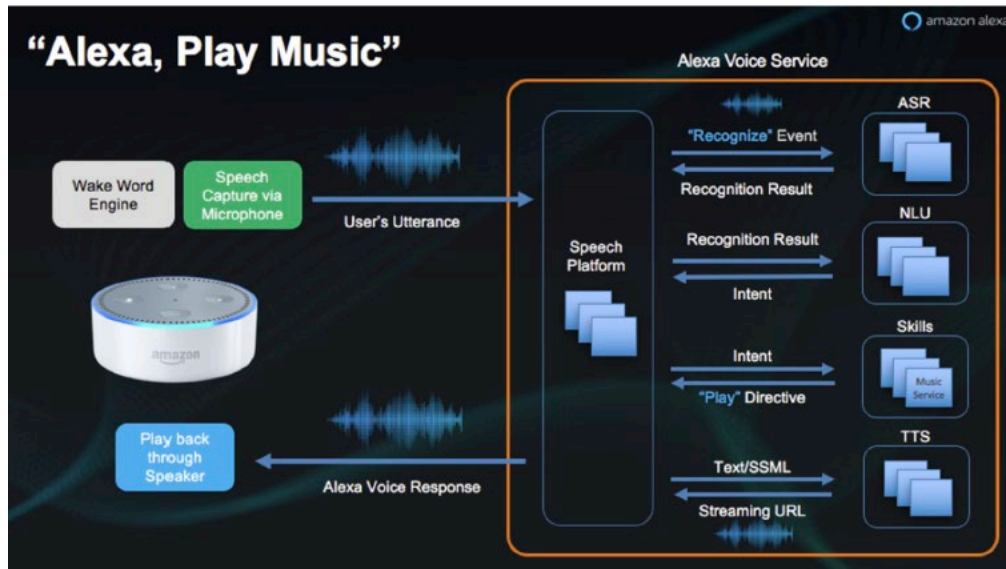
To add visual Alexa experiences to your device, you have several options for different levels of implementation:

- **Full APL integration** – To add a visual Alexa experience to your device, integrate the APL Core Library and implement an APL View Host in your code. See the [APL Core Library repository and documentation](#) to learn how to implement these two requirements.
- **Receive and render APL on a device** – To receive and render APL documents on your device screen, implement the [Alexa.Presentation.APL](#) namespace.
- **Publish device characteristics to Alexa** – To publish information about the visual characteristics of your device and its display for Alexa Skill developers, you can implement [Alexa.Display](#), [Alexa.Display.Window](#), and [Alexa.InteractionMode](#).
- **AVS Device SDK** – Simplify your APL implementation by integrating the AVS Device SDK, which includes a sample app for voice-only devices and an IPC Server sample app for smart screen devices: [AVS Device SDK repository](#).

216. Each of the Accused Multi-modal Products and Services comprises receiving a multi-modal natural language input at a conversational voice user interface, the multi-modal input including a natural language utterance and a non-speech input provided by a user, wherein a transcription module coupled to the conversational voice user interface transcribes the non-speech input to create a non-speech-based transcription.

217. For example, Amazon describes how the Alexa Voice Service processes a natural language request.⁷⁴

⁷⁴ <https://developer.amazon.com/fr/blogs/alexa/post/b015c5bd-2b01-4f7f-923e-dd34c998f3d0/avs-tech-series-the-basics-of-amazon-alexa-developer-tools-and-services>



Amazon further discloses that Alexa converts the sound to text and that text is processed to understand what the user means.⁷⁵

How Alexa Works

When I'm in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I'm saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as **Automatic Speech Recognition (ASR)**. Once the text is known, the second phase is to understand what I mean. This is **Natural Language Understanding (NLU)**. The output of NLU is an Intent (what does the customer want) and associated parameters. In this example ("Alexa, what's the weather today?"), the intent might be "`GetWeatherForecast`" and the parameter can be my postcode, inferred from my profile.

This whole process uses Artificial Intelligence heavily to transform the sound of my voice to **phonemes**, phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

Amazon further describes how Alexa Presentation Language processes multi-modal inputs including speech and non-speech elements.⁷⁶

⁷⁵ <https://aws.amazon.com/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost-effective-amazon-ec2-inf1-instances/>


⁷⁶ <https://developer.amazon.com/en-US/blogs/alexa/alexa-skills-kit/2020/07/new-alexa-presentation-language-1-4>

Build Interactive, Multimodal Experiences with Alexa Presentation Language (APL) 1.4

Arunjeet Singh Jul 22, 2020

Share: [f](#) [in](#) [t](#)

Multimodal Intermediate Alexa Live



We are excited to announce the next version of the Alexa Presentation Language (APL) that enables you to easily build interactive visual experiences with new capabilities and improved tooling. APL 1.4 lets you add editable text boxes, drag and drop UI controls, and back navigation so customers can return to previous screens. You can utilize new Alexa responsive components and templates to quickly add visuals to your skills to enhance the voice experience, and live preview your APL documents in the authoring tool. Also, we added a major update to the Alexa Skills Kit (ASK) toolkit for Visual Studio Code (VS Code) that adds APL rendering and local debugging. Learn more about APL 1.4 in our [technical documentation](#).

New APL Features

APL 1.4 supports user gestures and new components you can incorporate into your multimodal skills.

- **New Gestures:** Take advantage of new drag and drop UI controls such as single finger drag, swipe to delete, and long press gestures.
- **Editable Text Boxes:** Add new editable text boxes that allow customers to input a text response via touch or TV remote.
- **Grid Layout Support:** Use the GridSequence component to add a list of text and images to a fixed grid layout that scrolls in a single direction, vertically or horizontally. You can take advantage of the pre-built AlexaGridList responsive template that uses GridSequence.
- **Back Navigation:** Add back navigation so customers can easily return to previous screens in your skill with touch. Back navigation requires requesting the backstack extension from the APL extension framework.

Amazon further describes how Alexa captures and responds to key presses by the user.⁷⁷

Keyboard handlers

APL provides support for keyboard events. You can define keyboard handlers to capture and respond to key presses. For details, see [Keyboard Events and Handlers](#).

Amazon provides sample code of keyboard event handlers.⁷⁸

handleKeyDown

An array of [keyboard event handlers](#) to run when the user presses a key on the keyboard or when a key auto-repeats. The `keyDown` event is generated whenever possible, not just for text entry. For example, pressing the "shift" key should generate a `keyDown` event.

The event generated has the form:

```
"event": {
  "source": {
    "type": "COMPONENT_TYPE", // The type of the component (e.g., "Pager", "TouchWrapper")
    "handler": "KeyDown",
    ... // Component source properties
  },
  "keyboard": {
    "altKey": Boolean
    "code": String,
    "ctrlKey": Boolean,
    "key": String,
    "metaKey": Boolean,
    "repeat": Boolean,
    "shiftKey": Boolean
  }
}
```

The `event.source.type` property contains the name of the component, such as `"TouchWrapper"` or `"ScrollView"`. Refer to [Event source](#) for a description of `event.source` properties.

For more details on the `keyboard` property, see [Keyboard Event Handlers](#).

The `handleKeyDown` event handler runs in normal mode.

218. Each of the Accused Multi-modal Products and Services comprises identifying the user that provided the multi-modal input.

219. For example, Amazon describes Alexa Voice ID, which recognizes users when they speak.⁷⁹

⁷⁷ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-for-screen-devices.html>

⁷⁸ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-actionable-component.html>

⁷⁹ <https://www.amazon.com/gp/help/customer/display.html?nodeId=GYCXKY2AB2QWZT2X>

[Digital Services and Device Support](#) › [Alexa Settings](#) › [Alexa Profile Settings](#) ›

What Is Alexa Voice ID?

Alexa voice ID helps Alexa recognize you when you speak and provide a personalized experience.

You can create an Alexa voice ID for a personalized experience when Alexa recognizes your voice. With Alexa voice ID, Alexa can call you by name and provide enhanced personalization. To set up Alexa voice ID, go to [Create an Alexa Voice ID](#).

Amazon further describes how Voice ID is enabled for Alexa skills.⁸⁰

Request Recognized Speaker Contact Information for Use in Your Skill

Note: Sign in to the [developer console](#) to build or publish your skill.

If your skill implements personalization for recognized speakers, and the Alexa user is a recognized speaker with a voice ID and **Personalize skills** enabled in the Alexa app, your skill can request the user's permission to access the contact information associated with the user's voice ID. The contact information includes the name — either the full name or given name (first name) — and mobile number.

After receiving consent from the user, your skill can use the contact information to enable helpful experiences like the following:

- Personalize responses to the user, such as *"Welcome back, <first name>!"*
- Make a booking or appointment with your service on behalf of the user
- Send detailed or follow-up information to the user by text

You access this contact information by using the Person Profile API.

Note

For sample code, see the [Person Profile Demo](#) in GitHub.

For details about personalization, see [Add Personalization to Your Skill](#).

If you want to enable the user to verbally consent to share their contact information with your skill, see [Use Voice-Forward Consent in Your Alexa Skill](#).

⁸⁰ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/request-recognized-speaker-contact-information.html>

220. Each of the Accused Multi-modal Products and Services comprises creating a speech-based transcription of the natural language utterance using a speech recognition engine and a semantic knowledge-based model, wherein the semantic knowledge-based model includes a personalized cognitive model derived from one or more prior interactions between the identified user and the conversational voice user interface, a general cognitive model derived from one or more prior interactions between a plurality of users and the conversational voice user interface, and an environmental model derived from an environment of the identified user and the conversational voice user interface.

221. For example, Amazon describes how Alexa converts the sound into text.⁸¹

How Alexa Works

When I'm in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I'm saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as **Automatic Speech Recognition (ASR)**. Once the text is known, the second phase is to understand what I mean. This is **Natural Language Understanding (NLU)**. The output of NLU is an Intent (what does the customer want) and associated parameters. In this example ("Alexa, what's the weather today?"), the intent might be "**GetWeatherForecast**" and the parameter can be my postcode, inferred from my profile.

This whole process uses Artificial Intelligence heavily to transform the sound of my voice to **phonemes**, phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

Amazon further describes how Voice ID allows Alexa to personalize responses based on prior interactions with the identified user.⁸²

⁸¹ <https://aws.amazon.com/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost-effective-amazon-ec2-inf1-instances/>

⁸² <https://www.amazon.com/gp/help/customer/display.html?nodeId=GYCXKY2AB2QWZT2X>

Digital Services and Device Support › Alexa Settings › Alexa Profile Settings ›

What Is Alexa Voice ID?

Alexa voice ID helps Alexa recognize you when you speak and provide a personalized experience.

You can create an Alexa voice ID for a personalized experience when Alexa recognizes your voice. With Alexa voice ID, Alexa can call you by name and provide enhanced personalization. To set up Alexa voice ID, go to [Create an Alexa Voice ID](#).

Amazon further describes Alexa’s model as including training from a large generalized corpus.⁸³

Our overarching strategy for Alexa’s AI has been to combine machine learning (ML) — in particular, deep learning — with the large-scale data and computational resources available through AWS. But these performance improvements are the result of research on a variety of specific topics that extend deep learning, including

- semi-supervised learning, or using a [combination](#) of unlabeled and labeled data to [improve](#) the ML system;
- active learning, or the [learning strategy](#) where the ML system selects more-informative samples to receive manual labels;
- large-scale distributed training, or [parallelizing](#) ML-based model training for efficient learning on a large corpus; and
- context-aware modeling, or using a wide variety of information — including the [type of device](#) where a request originates, [skills](#) the customer uses or has enabled, and [past requests](#) — to improve accuracy.

Amazon further describes how the Intent Request History REST API can “get the aggregated and anonymized transcriptions of user speech data” from “customer interactions over the past 30 days.”⁸⁴

⁸³ <https://www.amazon.science/blog/alexa-at-five-looking-back-looking-forward>

⁸⁴ <https://developer.amazon.com/en-US/docs/alexa/smapi/intent-request-history.html>

Use the Intent Request History REST API to get the aggregated and anonymized transcriptions of user speech data and intent request details for a skill. The API returns data from customer interactions over the past 30 days, updated each day. You can use the data to learn how users interact with your skill to identify improvements to your interaction model.

This API supports text search and sort on the results. Data is available for skills that have at least 10 unique users per locale in a day. Alexa returns data for those locales that meet the threshold only.

Amazon further describes Alexa Hunches using information from the user's environment in supporting skill.⁸⁵

4. More context-aware and proactive

Today, through an optional feature called [Hunches](#), Alexa can learn how you interact with your smart home and suggest actions when she senses that devices such as lights, locks, switches, and plugs are not in the states that you prefer. We are currently expanding the notion of Hunches to include another Alexa feature called [Routines](#). If you set your alarm for 6:00 a.m. every day, for example, and on waking, you immediately ask for the weather, Alexa will suggest creating a Routine that sets the weekday alarm to 6:00 and plays the weather report as soon as the alarm goes off.

Earlier this year, we launched Alexa Guard, a feature that you can activate when you leave the house. If your Echo device [detects](#) the sound of a smoke alarm, a carbon monoxide alarm, or glass breaking, Alexa Guard sends you an alert. Guard's acoustic-event-detection model uses multitask learning, which reduces the amount of labeled data needed for training and makes the model [more compact](#).

This fall, we will begin previewing an extended version of Alexa Guard that recognizes additional sounds associated with activity, such as footsteps, talking, coughing, or doors closing. Customers can also create Routines that include Guard — activating Guard automatically during work hours, for instance.

222. On information and belief, each of the Accused Multi-modal Products and Services comprises merging the speech-based transcription and the non-speech-based transcription to create a merged transcription.

⁸⁵ <https://www.amazon.science/blog/alex-at-five-looking-back-looking-forward>

223. Each of the Accused Multi-modal Products and Services comprises identifying one or more entries in a context stack matching information contained in the merged transcription.

224. For example, Amazon describes how Alexa compares text combinations to a list of sample utterances that can invoke a skill in a name-free manner.⁸⁶

⁸⁶ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/understand-name-free-interaction-for-custom-skills.html>

The most important rule is to design intents with a list of sample utterances that captures key words and phrases that can invoke the skill in a name-free manner. Keep your utterances specific to your skills use-cases and avoid targeting generic utterances. The Alexa ML/AI models are designed in part to route traffic to skills with a high correlation between skill functionality and target utterances.

Guidelines

To help ensure that Alexa directs to your skill only those customer requests, questions, or statements that can be supported by the skill, build sample utterances that capture your skill's functionality. Provide sample utterances that you predict customers can use to interact with the skill naturally. For a better user experience, provide a large variety of sample utterances written in different forms. For example, name-free utterances for a horoscope skill might include:

- *"What is the horoscope for gemini?"*
- *"Today's horoscope"*
- *"I want to know my horoscope"*
- *"What does my star sign say?"*

Amazon further describes how HypRank in Alexa performs “intent-slot semantic analysis for a skill” such that user inputs are compared to “a list of hypotheses.”⁸⁷

⁸⁷ <https://www.amazon.science/blog/hyprank-how-alexa-determines-what-skill-can-best-meet-a-customers-need>

HypRank comprises two components:

1. Hypothesis representation for each skill; and
2. A bidirectional long short-term memory (bi-LSTM) model for re-ranking a list of hypotheses.

For each skill in the k-best list, we form a hypothesis based on additional semantic and contextual signals. For example, we perform intent-slot semantic analysis for a skill. If a user says “play Michael Jackson,” the Pop Music skill might infer the intent PlayMusic, while the Classic Music skill might infer the intent PlayTune. But the confidence scores that the skills assign their inferences could be useful for skill re-ranking. The hypothesis generator is constantly being updated, re-weighting signals and accommodating new functionality and changes in usage patterns.

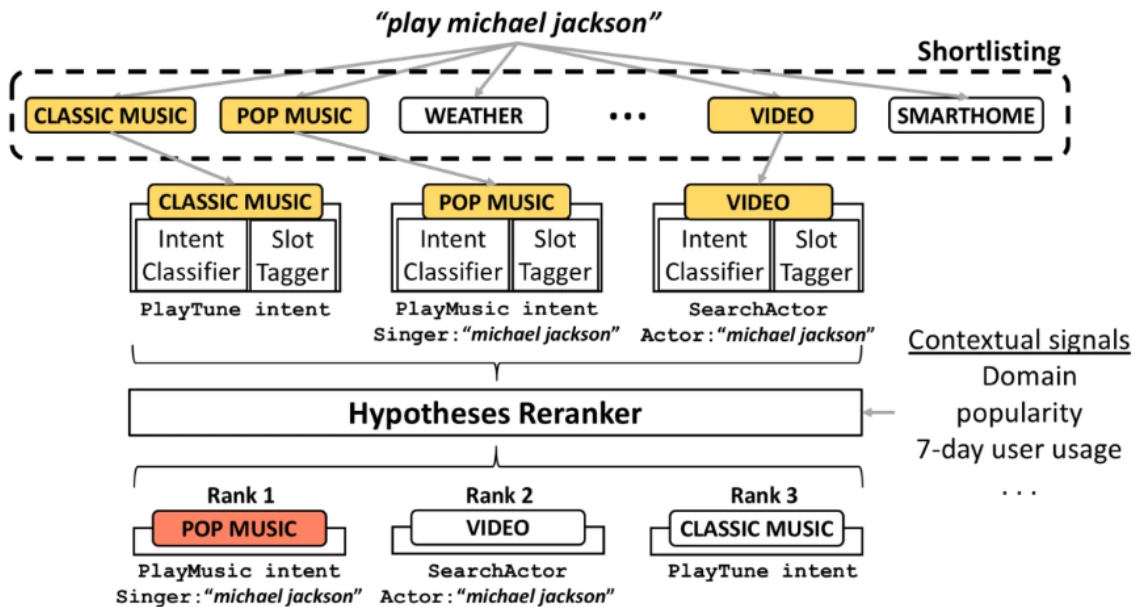
HypRank is unique because of its list-wise ranking approach using a bi-LSTM layer. LSTM models are common in natural-language processing because they factor in the order in which data are received: if you’re trying to understand the sixth word in an utterance, it helps to know what the previous five were. Bidirectional LSTM models consider data sequences both forward and backward.

By leveraging the bi-LSTM layer, HypRank can evaluate an entire list of skill hypotheses before providing a re-ranking score for each hypothesis. This is distinct from point-wise approaches that look at each hypothesis in isolation or pair-wise approaches that look at pairs of hypotheses in a series of tournament-like competitions.

225. Each of the Accused Multi-modal Products and Services comprises determining a most likely context for the multi-modal input based on the identified entries.

226. For example, Amazon describes how the Hypotheses Reranker in Alexa uses “rich contextual signals” in order “to select the most pertinent skills.”⁸⁸

⁸⁸ <https://www.amazon.science/blog/hyprank-how-alexa-determines-what-skill-can-best-meet-a-customers-need>



A high-level flow of the two-step shortlisting-reranking approach

The Challenge

The problem here is essentially a domain classification problem over the k-best candidate skills returned by the shortlisting system, which we call Shortlister. The goal of Shortlister is to achieve high recall — to identify as many pertinent skills as possible — with maximum efficiency. On the other hand, the goal of the reranking network, HypRank, is to use rich contextual signals to achieve high precision — to select the most pertinent skills. Designing HypRank comes with its own challenges:

227. Each of the Accused Multi-modal Products and Services comprises identifying a domain agent associated with the most likely context for the multi-modal input.

228. For example, Amazon describes Alexa as using a two-step approach to find the most relevant skill for a given utterance.⁸⁹

⁸⁹ <https://www.amazon.science/blog/the-scalable-neural-architecture-behind-alexa-s-ability-to-select-skills>

Alexa uses a two-step, scalable, and efficient neural shortlisting-reranking approach to find the most relevant skill for a given utterance. This post describes the first of those two steps, which relies on a neural model we call Shortlister. (I'll describe the second step in a follow-up post tomorrow.) Shortlister is a scalable and efficient architecture with a shared encoder, a personalized skill attention mechanism, and skill-specific classification networks. We outline this architecture in our [paper](#) "Efficient Large-Scale Neural Domain Classification with Personalized Attention", which we will present at the 56th Annual Meeting of the Association for Computational Linguistics ([ACL 2018](#)) in July.


The shared encoder network is hierarchical: Its lower layers are character-based and orthography sensitive and learn to represent each word in terms of character structure or shape; its middle layers are word-based, and with the outputs from the lower layers, they learn to represent an entire utterance. The skill attention mechanism is a separate network that is personalized per user. It computes a summary vector that describes which skills are enabled in a given user's profile and how relevant they are to the utterance representation. Both the utterance representation vector and the personalized skill-summary vector feed into a battery of skill-specific classification networks, one network for each skill.

Amazon further describes how Alexa surfaces the most relevant skills.⁹⁰

⁹⁰ <https://developer.amazon.com/fr/blogs/alexa/post/0fecdb38-97c9-48ac-953b-23814a469cfc/skill-discovery>


Surfacing the Most Relevant Skills via Shortlister and HypRank

It's important that customers are able to easily reengage with their skills – even if they describe what they want instead of calling the skill by name. Let's look again at the phrase "Alexa, remove an ink stain" where we have an utterance, "remove a {stain type} stain," and a slot value, stain type of "ink," that can be mapped to a skill's interaction model. To make this work, two techniques are employed to deliver the right experience. The first is a neural model architecture we call Shortlister, which solves a domain classification problem to find the most statistically significant (k -best) matches of candidate skills (details in a recent [blog](#) and [research paper](#)). The goal of Shortlister is to efficiently identify as many relevant skills as possible – it's high volume. The second is a hypothesis reranking network, we call HypRank. It uses contextual signals to select the most relevant skills – it's high precision (details in a recent [blog](#) and [research paper](#)). When combined, these two help classify which skill can provide the most likely correct response from an invocation name-free utterance. Beyond helping people reengage with the skills they've manually or automatically enabled, the Shortlister and HypRank surface specific and relevant skills to people that have never used them.



Amazon further provides examples of how the hypothesis reranker operates.⁹¹

⁹¹ <https://www.amazon.science/blog/hyprank-how-alexa-determines-what-skill-can-best-meet-a-customers-need>



The Challenge

The problem here is essentially a domain classification problem over the k-best candidate skills returned by the shortlisting system, which we call Shortlister. The goal of Shortlister is to achieve high recall — to identify as many pertinent skills as possible — with maximum efficiency. On the other hand, the goal of the reranking network, HypRank, is to use rich contextual signals to achieve high precision — to select the most pertinent skills. Designing HypRank comes with its own challenges:

229. Each of the Accused Multi-modal Products and Services comprises communicating a request to the identified domain agent.

230. For example, Amazon describes how Alexa communicates a request to an Alexa skill.⁹²

The following diagram shows the voice-activated processing flow to invoke a skill with the Alexa service.



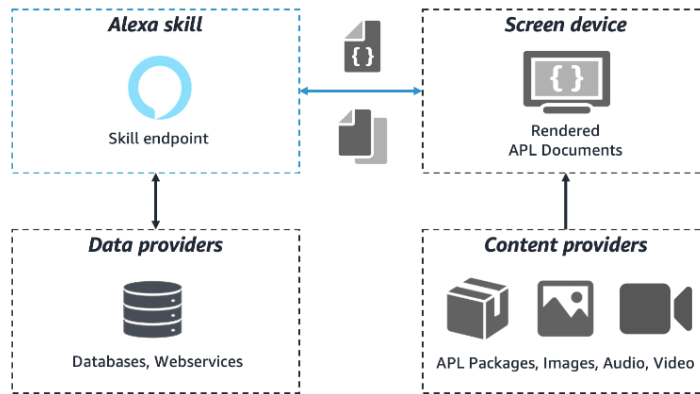
In addition to voice interaction, skills might include complementary visuals and touch interactions.

⁹² <https://developer.amazon.com/en-US/docs/alexa/ask-overviews/what-is-the-alexa-skills-kit.html>

231. Each of the Accused Multi-modal Products and Services comprises generating a response to the user from content provided by the identified domain agent as a result of processing the request.

232. For example, Amazon describes how an Alexa skill processes a multi-modal response on a screen device.⁹³

To plan your APL experiences, you need to understand the four technical *actors* that play a role in rendering APL *documents* on Alexa-enabled devices with screens. The following diagram shows a typical APL architecture at a high level.



APL architecture with four technical actors.

- **Alexa skill** — A skill endpoint that initiates an APL experience by sending out an APL *document* template in response to a regular skill request such as an `IntentRequest`.
- **Data providers** — APL templates often come with data that should be rendered with it so that the skill often attaches a JSON data source to the template that is sent to a screen device. This data typically doesn't reside within a skill, but resides in external data stores, such as a database that the skill reads from and writes to.
- **Screen device** — When the skill has sent an APL template with an attached data source, an Alexa-enabled device with a screen, such as an Echo Show or Fire TV, is responsible for inflating and rendering the template on screen.
- **Content providers** — APL templates that are sent by a skill often contain URL references to external content and media files; such as APL *packages*, images, audio and video files; that are stored publicly on the internet. An Alexa-enabled device with a screen resolves these references by downloading or streaming source files from external content providers, such as *content delivery networks* (CDN).

233. Amazon has long known about the '468 Patent.

234. Amazon has been aware of the '468 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.

⁹³ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-bp-understand-apl-architecture.html>

235. Amazon knew or should have known that Amazon's actions infringe one or more of the claims of the '468 Patent because Amazon has the technical expertise to understand the scope and content of the '468 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Multi-modal Products and Services, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '468 Patent at least as of the filing of this Complaint.

236. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 19 of the '468 Patent in violation of at least 35 U.S.C. § 271(b) and (c).

237. Users of the Accused Multi-modal Products and Services directly infringe at least Claim 19 of the '468 Patent when they use the Accused Multi-modal Products and Services in the ordinary, customary, and intended way.

238. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Multi-modal Products and Services within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Multi-modal Products and Services to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Multi-modal Products and Services in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Multi-modal Products and Services, which Amazon knew infringes at least Claim 19 of the '468 Patent, or, alternatively, was willfully blind to the infringement.

239. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the Accused Multi-modal Products and Services within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Multi-modal Products and Services in the United States by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Multi-modal Products and Services, which Amazon knew infringes at least Claim 19 of the '468 Patent, or, alternatively, was willfully blind to the infringement.

240. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 19 of the '468 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '468 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

241. Amazon is not licensed or otherwise authorized to practice the claims of the '468 Patent.

242. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '468 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 19.

243. As a result of Amazon's infringement of the '468 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

244. On information and belief, in addition to Amazon's knowledge of the '468 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '468 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '468 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

245. Accordingly, Amazon's infringement of the '468 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

246. Amazon's infringement of Dialect's rights under the '468 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

FIFTH COUNT
(Infringement of U.S Patent No. 9,031,845)

247. Dialect incorporates by reference the allegations set forth in Paragraphs 1-246 as though fully set forth herein.

248. The claims of the '845 Patent are valid and enforceable.

249. The claims of the '845 Patent are directed to patentable subject matter. Specifically, Claim 1 of the '845 Patent is directed to a novel vehicular speech-recognition system including the claimed novel software structures and combination which formulate a command or query based on the domain and the context and selectively invoke wide area networking wireless communication devices based on a determination of whether the command or query is to be executed on-board or off-board the vehicle. The inventive, tangible claimed structures of the '845 Patent improve on existing natural language processing systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

250. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and services that embody the invention disclosed and claimed in the '845 Patent, including at least the Accused Automotive Products and Services.

251. Each of the Accused Automotive Products and Services contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 1 of the '845 Patent.

252. Each of the Accused Automotive Products and Services comprises a mobile system for processing natural language utterances.

253. For example, Amazon describes how vehicles integrating the Alexa Auto SDK and Alexa Mobile Accessory Kit provide mobile solutions for hands-free interactions.⁹⁴


⁹⁴ <https://developer.amazon.com/en-US/alexa/alexa-auto/vehicles#experiences>



Head Unit with Alexa Built-in

By building Alexa into your vehicle with the Alexa Auto SDK, your customers can interact with Alexa, hands-free, using the in-cabin microphone and speakers. You can also display digital content on the infotainment screen, like album artwork, audiobook covers, search results, and more.

[Alexa Auto SDK »](#)



Aftermarket Device with Alexa Built-in

Bring Alexa to the millions of vehicles already on the road by building Alexa into your automotive aftermarket device. The Alexa Mobile Accessory (AMA) Kit enables device makers to easily connect Bluetooth-enabled audio devices with Alexa using the Amazon Alexa App on Android and iOS devices.

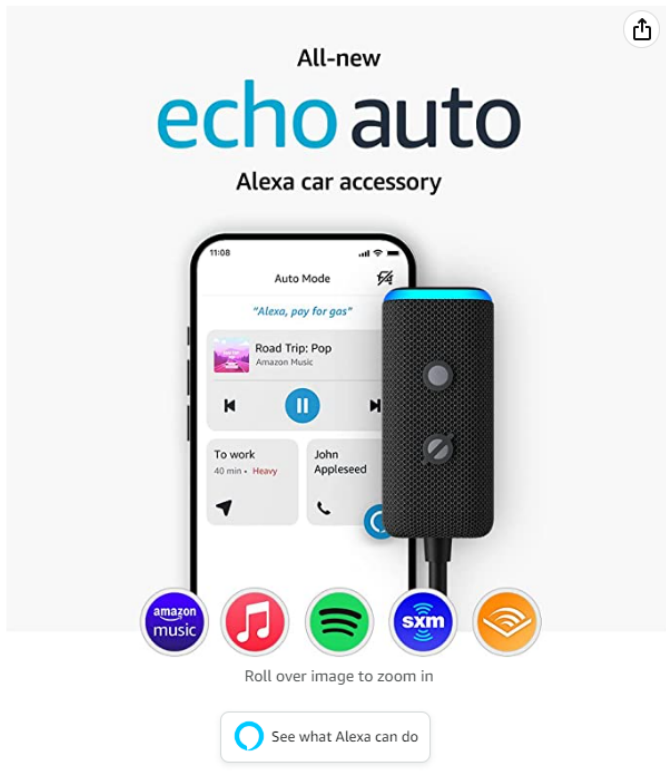
[Learn more about AMA »](#)

Amazon further describes the Alexa Auto SDK as enabling speech recognition, Automatic Speech Recognition and Natural Language Understanding.⁹⁵

⁹⁵ <https://developer.amazon.com/en-US/blogs/alexa/alexa-auto/2019/09/enabling-offline-access-to-alexa-in-vehicles-with-local-voice-control-extension-to-alexa-auto-sdk-v2-0>

The Alexa Auto SDK enables core Alexa functionality, such as speech recognition and text-to-speech, and other capabilities such as streaming media, controlling smart home devices, notifications, weather reports, and tens of thousands of custom skills. It can be updated over-the-air and now includes the Local Voice Control extension, which ensures that customers get a reliable response from Alexa regardless of connectivity. On a normal commute, drivers lose connectivity or experience intermittent or transitional connectivity while driving in dense urban areas, through tunnels, in remote areas, or in parking garages. With the Local Voice Control extension to the Alexa Auto SDK, automakers embed a scaled-down version of Alexa's cloud-based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle's infotainment system that increases resiliency of the cloud-based capabilities. This ensures that a customer can always ask, "Alexa, turn on the A/C," "Alexa, tune to 90.5 FM," or "Alexa, call Jackie," and get the appropriate response and action from Alexa.

Amazon further describes Echo Auto as enabling hands-free features.⁹⁶



All-new
echo auto
Alexa car accessory

Amazon Music, Apple Music, Spotify, and more, or listen live to radio stations. Use Follow Me Music to resume media playback.

- **CALL AND MESSAGE WITH YOUR VOICE** - Use your voice to make calls, reply to text messages, drop in on Alexa-enabled devices in your home or broadcast announcements.
- **GET YOUR FAVORITE ENTERTAINMENT HANDS-FREE** - Binge a hot new podcast, catch up on the news, or listen to best-selling Audible books.
- **ALEXA CAN HELP WITH AMAZON ROADSIDE ASSISTANCE** - Just say, "Alexa, call Roadside Assistance" and you'll be connected to an agent who can assess the situation and request pay-as-you-go services as needed.
- **CONTROL YOUR SMART HOME FROM THE ROAD** - Ask Alexa to set the thermostat, turn off the lights, check if your front door is locked, and more while you're away from home.
- **DESIGNED TO PROTECT YOUR PRIVACY** - Built with multiple layers of privacy controls, including a mic-off button and light indicator.

We want you to know

Echo Auto is designed to enable Alexa hands-free features, such as play music, make calls, answer text messages and more for cars that don't have a voice assistant built-in. Echo Auto connects to the Alexa app on your phone using your existing data plan and plays through your car's speakers via Bluetooth connection or Aux input. Carrier charges may apply.

Car and Phone compatibility - Check if your car is compatible and check if your phone is compatible. Note: Echo Auto is compatible with most cars with Bluetooth or Aux input.

Fast car charger - This product is Quick Charge 3.0 certified for fast charging.

Mounting your Echo Auto - The included adhesive car mount should not be placed on soft plastic, leather, or fabric surfaces. Choose a hard, durable surface away from the airflow.

254. Each of the Accused Automotive Products and Services comprises one or more physical processors at a vehicle that are programmed to execute one or more computer program instructions which, when executed, cause the one or more physical processors to: receive a natural language utterance associated with a user.

⁹⁶ <https://developer.android.com/training/cars/apps/poi>

255. For example, Amazon describes the use of the Alexa Auto SDK to provide a “scaled-down version of Alexa’s cloud based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle’s infotainment system.”⁹⁷

The Alexa Auto SDK enables core Alexa functionality, such as speech recognition and text-to-speech, and other capabilities such as streaming media, controlling smart home devices, notifications, weather reports, and tens of thousands of custom skills. It can be updated over-the-air and now includes the Local Voice Control extension, which ensures that customers get a reliable response from Alexa regardless of connectivity. On a normal commute, drivers lose connectivity or experience intermittent or transitional connectivity while driving in dense urban areas, through tunnels, in remote areas, or in parking garages. With the Local Voice Control extension to the Alexa Auto SDK, automakers embed a scaled-down version of Alexa’s cloud-based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle’s infotainment system that increases resiliency of the cloud-based capabilities. This ensures that a customer can always ask, “Alexa, turn on the A/C,” “Alexa, tune to 90.5 FM,” or “Alexa, call Jackie,” and get the appropriate response and action from Alexa.

256. Each of the Accused Automotive Products and Services comprises performing speech recognition on the natural language utterance.

257. For example, Amazon describes Alexa Auto SDK as supporting Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU).⁹⁸

258. Each of the Accused Automotive Products and Services comprises parsing and interpreting the speech recognized natural language utterance.

259. For example, Amazon describes how in “vehicles with Alexa” that the system can interpret natural language questions.⁹⁹

“Alexa, where’s the nearest coffee shop?”

In [vehicles with Alexa](#), drivers can pose questions like that — while keeping their eyes on the road and hands on the wheel. Amazon’s cloud-based voice assistant technology works with the car’s navigation system to figure out where the nearest coffee shop is and guide drivers to it.

⁹⁷ <https://developer.amazon.com/en-US/blogs/alexa/alexa-auto/2019/09/enabling-offline-access-to-alexa-in-vehicles-with-local-voice-control-extension-to-alexa-auto-sdk-v2-0t>

⁹⁸ *Id.*

⁹⁹ <https://www.amazon.science/news-and-features/the-science-behind-alexa-in-vehiclest>

Amazon states that its intention is to take the home Alexa experience “and extend that experience to cars.”¹⁰⁰ Amazon further describes Alexa as matching the transcribed text of the user’s speech to the user’s intent.¹⁰¹

How Alexa Works

When I’m in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I’m saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as **Automatic Speech Recognition (ASR)**. Once the text is known, the second phase is to understand what I mean. This is **Natural Language Understanding (NLU)**. The output of NLU is an Intent (what does the customer want) and associated parameters. In this example (“Alexa, what’s the weather today?”), the intent might be “**GetWeatherForecast**” and the parameter can be my postcode, inferred from my profile.

This whole process uses Artificial Intelligence heavily to transform the sound of my voice to **phonemes**, phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

260. Each of the Accused Automotive Products and Services comprises determining a domain and a context that are associated with the parsed and interpreted natural language utterance.

261. For example, Amazon describes Alexa as using a two-step approach to find the most relevant skill for a given utterance.¹⁰²

¹⁰⁰ *Id.*

¹⁰¹ <https://aws.amazon.com/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost-effective-amazon-ec2-inf1-instances/>

¹⁰² <https://www.amazon.science/blog/the-scalable-neural-architecture-behind-alexa-s-ability-to-select-skills>

Alexa uses a two-step, scalable, and efficient neural shortlisting-reranking approach to find the most relevant skill for a given utterance. This post describes the first of those two steps, which relies on a neural model we call Shortlister. (I'll describe the second step in a follow-up post tomorrow.) Shortlister is a scalable and efficient architecture with a shared encoder, a personalized skill attention mechanism, and skill-specific classification networks. We outline this architecture in our [paper](#) "Efficient Large-Scale Neural Domain Classification with Personalized Attention", which we will present at the 56th Annual Meeting of the Association for Computational Linguistics ([ACL 2018](#)) in July.

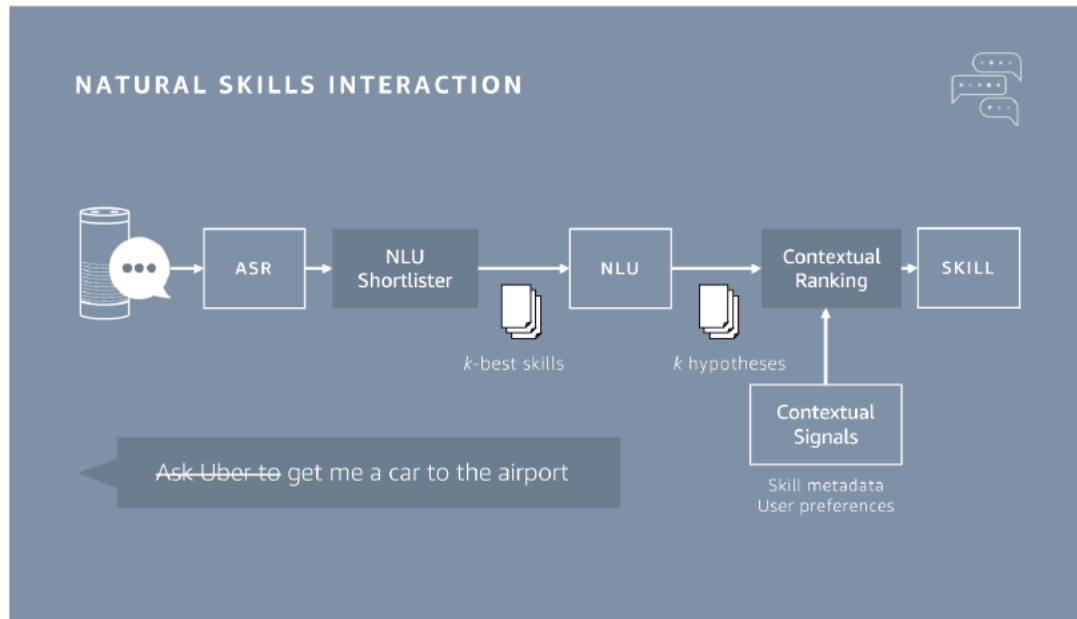
The shared encoder network is hierarchical: Its lower layers are character-based and orthography sensitive and learn to represent each word in terms of character structure or shape; its middle layers are word-based, and with the outputs from the lower layers, they learn to represent an entire utterance. The skill attention mechanism is a separate network that is personalized per user. It computes a summary vector that describes which skills are enabled in a given user's profile and how relevant they are to the utterance representation. Both the utterance representation vector and the personalized skill-summary vector feed into a battery of skill-specific classification networks, one network for each skill.

Amazon further describes how Alexa surfaces the most relevant skills.¹⁰³

¹⁰³ <https://developer.amazon.com/fr/blogs/alexa/post/0fecdb38-97c9-48ac-953b-23814a469cfc/skill-discovery>

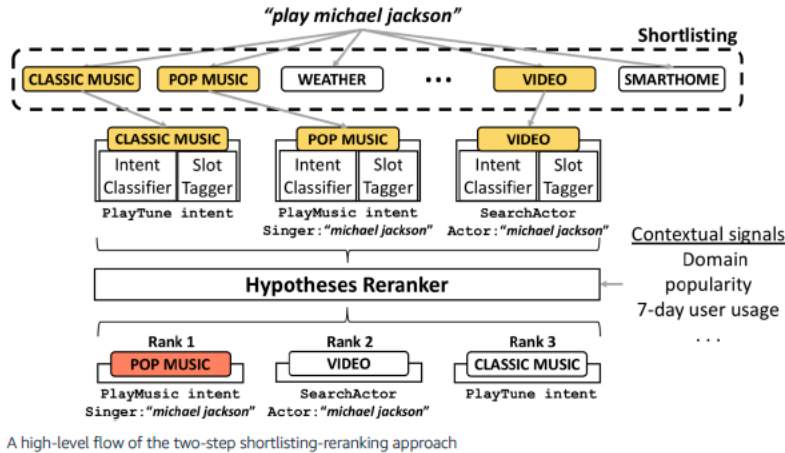
Surfacing the Most Relevant Skills via Shortlister and HypRank

It's important that customers are able to easily reengage with their skills – even if they describe what they want instead of calling the skill by name. Let's look again at the phrase "Alexa, remove an ink stain" where we have an utterance, "remove a {stain type} stain," and a slot value, stain type of "ink," that can be mapped to a skill's interaction model. To make this work, two techniques are employed to deliver the right experience. The first is a neural model architecture we call Shortlister, which solves a domain classification problem to find the most statistically significant (*k*-best) matches of candidate skills (details in a recent [blog](#) and [research paper](#)). The goal of Shortlister is to efficiently identify as many relevant skills as possible – it's high volume. The second is a hypothesis reranking network, we call HypRank. It uses contextual signals to select the most relevant skills – it's high precision (details in a recent [blog](#) and [research paper](#)). When combined, these two help classify which skill can provide the most likely correct response from an invocation name-free utterance. Beyond helping people reengage with the skills they've manually or automatically enabled, the Shortlister and HypRank surface specific and relevant skills to people that have never used them.



Amazon further provides examples of how the hypothesis reranker operates.¹⁰⁴

¹⁰⁴ <https://www.amazon.science/blog/hyprank-how-alex-a-determines-what-skill-can-best-meet-a-customers-need>



The Challenge

The problem here is essentially a domain classification problem over the k-best candidate skills returned by the shortlisting system, which we call Shortlister. The goal of Shortlister is to achieve high recall — to identify as many pertinent skills as possible — with maximum efficiency. On the other hand, the goal of the reranking network, HypRank, is to use rich contextual signals to achieve high precision — to select the most pertinent skills. Designing HypRank comes with its own challenges:

262. Each of the Accused Automotive Products and Services comprises formulating a command or query based on the domain and the context.

263. For example, Amazon describes Alexa as processing the user’s command to tune to 90.5 FM.¹⁰⁵



¹⁰⁵ <https://developer.amazon.com/en-US/blogs/alexa/alexa-auto/2019/09/enabling-offline-access-to-alexa-in-vehicles-with-local-voice-control-extension-to-alexa-auto-sdk-v2-0t>

264. Each of the Accused Automotive Products and Services comprises determining whether the command or query is to be executed on-board or off-board the vehicle.

265. For example, Alexa describes the Local Voice Control extension to the Alexa Auto SDK as providing for commands or queries to be executed on-board the vehicle instead of on the Alexa cloud hardware.¹⁰⁶

The Alexa Auto SDK enables core Alexa functionality, such as speech recognition and text-to-speech, and other capabilities such as streaming media, controlling smart home devices, notifications, weather reports, and tens of thousands of custom skills. It can be updated over-the-air and now includes the Local Voice Control extension, which ensures that customers get a reliable response from Alexa regardless of connectivity. On a normal commute, drivers lose connectivity or experience intermittent or transitional connectivity while driving in dense urban areas, through tunnels, in remote areas, or in parking garages. With the Local Voice Control extension to the Alexa Auto SDK, automakers embed a scaled-down version of Alexa's cloud-based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle's infotainment system that increases resiliency of the cloud-based capabilities. This ensures that a customer can always ask, "Alexa, turn on the A/C," "Alexa, tune to 90.5 FM," or "Alexa, call Jackie," and get the appropriate response and action from Alexa.

"The vehicle environment is unique and requires online and offline voice capabilities to be truly successful," says Matt Arcaro, Research Manager, Next-Generation Automotive Research at IDC. "This dual-mode support not only reduces the impact of variable network conditions, but will give vehicle users the confidence to increasingly think of voice as an effective and efficient means for vehicle interaction and customization."

The following features are now available through the Alexa Auto SDK:

- **Music and Radio:** control media playback, stream music, audiobooks, and podcasts from leading providers, tune radio, change local media sources, turn up the volume, and adjust equalizer.
- **Calling and Messaging:** dial by name and number, control call in progress, make Alexa announcements, use Alexa Drop In, join meetings, and send and receive Alexa messages.
- **Search and Navigation:** navigate to favorite destinations, search for places by name, brand, category, address, or intersection, get details such as hours of operation, phone number, and ratings, determine route info like traffic and ETA, and cancel navigation in progress.
- **Car Control:** turn on air conditioning, set desired cabin temperature, control temperature by zone, set fan mode and speed, defrost front and rear windshield, and control interior cabin lighting.

266. Each of the Accused Automotive Products and Services comprises executing the command or query at the vehicle in response to a determination that the command or query is to be executed on-board the vehicle.

267. For example, Amazon describes commands that can be executed at the vehicle when the command is to be executed on-board.¹⁰⁷

¹⁰⁶ *Id.*

¹⁰⁷ <https://developer.android.com/training/cars/media>

The Alexa Auto SDK enables core Alexa functionality, such as speech recognition and text-to-speech, and other capabilities such as streaming media, controlling smart home devices, notifications, weather reports, and tens of thousands of custom skills. It can be updated over-the-air and now includes the Local Voice Control extension, which ensures that customers get a reliable response from Alexa regardless of connectivity. On a normal commute, drivers lose connectivity or experience intermittent or transitional connectivity while driving in dense urban areas, through tunnels, in remote areas, or in parking garages. With the Local Voice Control extension to the Alexa Auto SDK, automakers embed a scaled-down version of Alexa's cloud-based Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) into the vehicle's infotainment system that increases resiliency of the cloud-based capabilities. This ensures that a customer can always ask, "Alexa, turn on the A/C," "Alexa, tune to 90.5 FM," or "Alexa, call Jackie," and get the appropriate response and action from Alexa.

"The vehicle environment is unique and requires online and offline voice capabilities to be truly successful," says Matt Arcaro, Research Manager, Next-Generation Automotive Research at IDC. "This dual-mode support not only reduces the impact of variable network conditions, but will give vehicle users the confidence to increasingly think of voice as an effective and efficient means for vehicle interaction and customization."

The following features are now available through the Alexa Auto SDK:

- **Music and Radio:** control media playback, stream music, audiobooks, and podcasts from leading providers, tune radio, change local media sources, turn up the volume, and adjust equalizer.
- **Calling and Messaging:** dial by name and number, control call in progress, make Alexa announcements, use Alexa Drop In, join meetings, and send and receive Alexa messages.
- **Search and Navigation:** navigate to favorite destinations, search for places by name, brand, category, address, or intersection, get details such as hours of operation, phone number, and ratings, determine route info like traffic and ETA, and cancel navigation in progress.
- **Car Control:** turn on air conditioning, set desired cabin temperature, control temperature by zone, set fan mode and speed, defrost front and rear windshield, and control interior cabin lighting.

268. Each of the Accused Automotive Products and Services comprises invoking a device that communicates wirelessly over a wide area network to process the command or query such that the command or query is executed off-board the vehicle in response to a determination that the command or query is to be executed off-board the vehicle.

269. For example, Amazon describes the use of online functionality for accessing Alexa in the car.¹⁰⁸

Today, we announced version 2.0 of the Alexa Auto SDK, which now offers a full-suite of offline functionality, in addition to online functionality, for accessing Alexa in the car, including: music and radio, search and navigation, calling and messaging, and car control. Automakers and their suppliers can use the Alexa Auto SDK to integrate Alexa into their vehicles, helping keep customers entertained and productive—even when connectivity is intermittent or not available.

270. Amazon has long known about the '845 Patent.

¹⁰⁸ *Id.*

271. Amazon has been aware of the '845 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.

272. Amazon knew or should have known that Amazon's actions infringe one or more of the claims of the '845 Patent because Amazon has the technical expertise to understand the scope and content of the '845 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Automotive Products and Services, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '845 Patent at least as of the filing of this Complaint.

273. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 1 of the '845 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

274. Users of the Accused Automotive Products and Services directly infringe at least Claim 1 of the '845 Patent when they use the Accused Automotive Products and Services in the ordinary, customary, and intended way.

275. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Automotive Products and Services within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Automotive Products and Services to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Automotive Products and Services in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and

distribution of the Accused Automotive Products and Services, which Amazon knew infringes at least Claim 1 of the '845 Patent, or, alternatively, was willfully blind to the infringement.

276. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the Accused Automotive Products and Services within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Automotive Products and Services in the United States by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Automotive Products and Services, which Amazon knew infringes at least Claim 1 of the '845 Patent, or, alternatively, was willfully blind to the infringement.

277. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 1 of the '845 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '845 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

278. On information and belief, in violation of 35 U.S.C. § 271(f)(1), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 1 of the '845 Patent, where such components are uncombined in whole or in part, in such

manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

279. On information and belief, in violation of 35 U.S.C. § 271(f)(2), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 1 of the '845 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

280. Amazon is not licensed or otherwise authorized to practice the claims of the '845 Patent.

281. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '845 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 1.

282. As a result of Amazon's infringement of the '845 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

283. On information and belief, in addition to Amazon's knowledge of the '845 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '845 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that

infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '845 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

284. Accordingly, Amazon's infringement of the '845 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

285. Amazon's infringement of Dialect's rights under the '845 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SIXTH COUNT
(Infringement of U.S Patent No. 9,263,039)

286. Dialect incorporates by reference the allegations set forth in Paragraphs 1-285 as though fully set forth herein.

287. The claims of the '039 Patent are valid and enforceable.

288. The claims of the '039 Patent are directed to patentable subject matter. Specifically, Claim 13 of the '039 Patent is directed to a novel natural-language speech-recognition system including the claimed novel software structures and combination which compare speech and non-speech text inputs to a context description grammar, calculates a relevance score, selects a domain agent based on a relevance score, and arranges the content of the response based on the results of the scoring module. The inventive, tangible claimed structures of the '039 Patent improve on

existing natural language processing systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

289. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and services that embody the invention disclosed and claimed in the '039 Patent, including at least the Alexa Products, operating via tablets, smartphones, or other devices supporting non-speech inputs such as the Echo Show (collectively, the “Accused Multi-modal Products and Services”).

290. Each of the Accused Multi-modal Products and Services contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 13 of the '039 Patent.

291. Each of the Accused Multi-modal Products and Services comprises a method of processing speech and non-speech communications.

292. For example, Amazon describes Alexa as performing speech recognition on user inputs.¹⁰⁹

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



¹⁰⁹ <https://developer.amazon.com/en-US/alexa/alexa-skills-kit/get-deeper/custom-skills>

Amazon further describes Alexa as processing non-speech inputs.¹¹⁰

Alexa Multi-Modal API Overview

Note: Learn how to improve your skills with APL with [Build visually rich experiences using APL](#) at the Alexa Learning Lab.

The Alexa Voice Service (AVS) provides a collection of APIs for Alexa Built-in devices with visual displays, such as smart TVs, which allow those devices to receive and present Alexa visual responses. Some of these APIs help device makers to publish the visual characteristics and functionality of their visual displays to enable Alexa Skills and services to respond with appropriate visual content for a device display.

To add visual Alexa experiences to your device, you have several options for different levels of implementation:

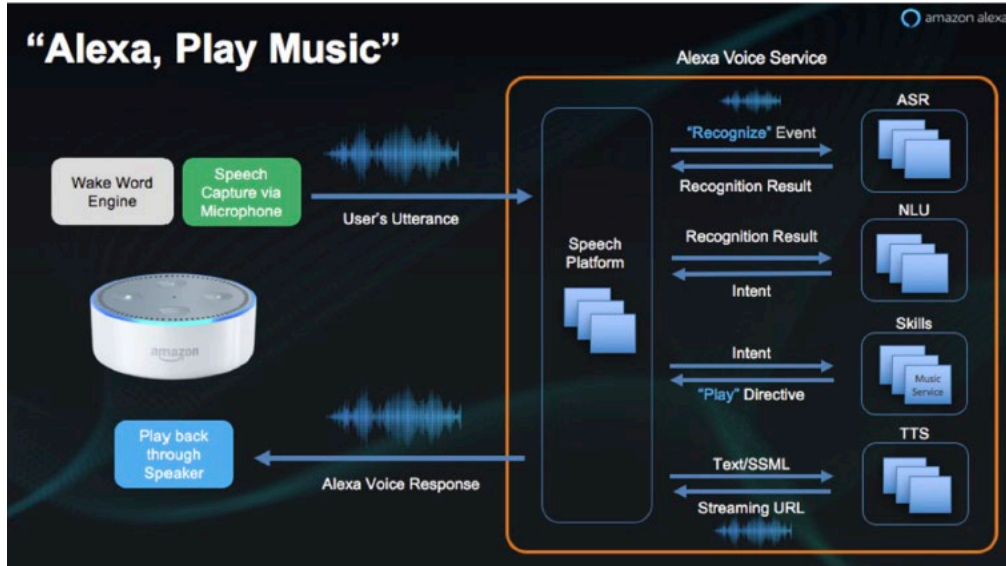
- **Full APL integration** – To add a visual Alexa experience to your device, integrate the APL Core Library and implement an APL View Host in your code. See the [APL Core Library repository and documentation](#) to learn how to implement these two requirements.
- **Receive and render APL on a device** – To receive and render APL documents on your device screen, implement the [Alexa.Presentation.APL](#) namespace.
- **Publish device characteristics to Alexa** – To publish information about the visual characteristics of your device and its display for Alexa Skill developers, you can implement [Alexa.Display](#), [Alexa.Display.Window](#), and [Alexa.InteractionMode](#).
- **AVS Device SDK** – Simplify your APL implementation by integrating the AVS Device SDK, which includes a sample app for voice-only devices and an IPC Server sample app for smart screen devices: [AVS Device SDK repository](#).

293. Each of the Accused Multi-modal Products and Services comprises receiving the speech and non-speech communications.

294. For example, Amazon discloses that Alexa receives speech inputs.¹¹¹

¹¹⁰ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/avs-apl-overview.html>

¹¹¹ <https://developer.amazon.com/fr/blogs/alexa/post/b015c5bd-2b01-4f7f-923e-dd34c998f3d0/avs-tech-series-the-basics-of-amazon-alexa-developer-tools-and-services>



Amazon further discloses Alexa processing the sound input into text and then parsing it to understand what the user means.¹¹²

How Alexa Works

When I'm in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I'm saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as **Automatic Speech Recognition (ASR)**. Once the text is known, the second phase is to understand what I mean. This is **Natural Language Understanding (NLU)**. The output of NLU is an Intent (what does the customer want) and associated parameters. In this example ("Alexa, what's the weather today?"), the intent might be "**GetWeatherForecast**" and the parameter can be my postcode, inferred from my profile.

This whole process uses Artificial Intelligence heavily to transform the sound of my voice to **phonemes**, phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

Amazon further describes Alexa's ability to receive non-speech inputs.¹¹³

¹¹² <https://aws.amazon.com/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost-effective-amazon-ec2-inf1-instances/>

¹¹³ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-authoring-tool.html#use-mmrb>

Use the Multimodal Response Builder

The Multimodal Response Builder provides a guided experience to create a visual response in three steps:

1. Select from a set of templates designed to look good across a broad range of devices.
2. Customize the response by specifying the content to display within the template
3. Preview the response both in the developer console and on a device.

After you finish customizing the response, the Multimodal Response Builder generates a code example you can copy into a request handler in your skill to display the response. Rebuild the interaction model for your skill and then test your skill on a device or with the developer console simulator.

You can create a new document in the Multimodal Response Builder to get started, and then edit the document in the [full authoring tool](#) if you want to do more complex customizations.

To open the Multimodal Response Builder

1. In the developer console, open the skill for which you want to create this document.
2. In the left-hand navigation, click **Multimodal Responses**.
3. Click **Create with Response Builder**.
4. Step through the pages to complete the three steps
 - Select a template. The set of templates includes [responsive templates](#) as well as other visual designs.
 - Customize
 - Preview and test

Note: Populating the primary text or title fields of a template with long strings might create a poor visual experience. Make sure you test your response across different viewport and make sure all of your content displays as you expect.

Amazon further describes the Alexa Multi-Modal API.¹¹⁴

¹¹⁴ <https://developer.amazon.com/en-US/docs/alexa/alexa-voice-service/avs-apl-overview.html>

Alexa Multi-Modal API Overview

Note: Learn how to improve your skills with APL with [Build visually rich experiences using APL](#) at the Alexa Learning Lab.

The Alexa Voice Service (AVS) provides a collection of APIs for Alexa Built-in devices with visual displays, such as smart TVs, which allow those devices to receive and present Alexa visual responses. Some of these APIs help device makers to publish the visual characteristics and functionality of their visual displays to enable Alexa Skills and services to respond with appropriate visual content for a device display.

To add visual Alexa experiences to your device, you have several options for different levels of implementation:

- **Full APL integration** – To add a visual Alexa experience to your device, integrate the APL Core Library and implement an APL View Host in your code. See the [APL Core Library repository and documentation](#) to learn how to implement these two requirements.
- **Receive and render APL on a device** – To receive and render APL documents on your device screen, implement the [Alexa.Presentation.APL](#) namespace.
- **Publish device characteristics to Alexa** – To publish information about the visual characteristics of your device and its display for Alexa Skill developers, you can implement [Alexa.Display](#), [Alexa.Display.Window](#), and [Alexa.InteractionMode](#).
- **AVS Device SDK** – Simplify your APL implementation by integrating the AVS Device SDK, which includes a sample app for voice-only devices and an IPC Server sample app for smart screen devices: [AVS Device SDK repository](#).

295. Each of the Accused Multi-modal Products and Services comprises transcribing the speech and non-speech communications to create a speech-based textual message and a non-speech-based textual message.

296. For example, Amazon describes Alexa as transcribes the speech.¹¹⁵

How Alexa Works

When I'm in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I'm saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as [Automatic Speech Recognition \(ASR\)](#). Once the text is known, the second phase is to understand what I mean. This is [Natural Language Understanding \(NLU\)](#). The output of NLU is an Intent (what does the customer want) and associated parameters. In this example ("Alexa, what's the weather today?"), the intent might be "

`GetWeatherForecast` " and the parameter can be my postcode, inferred from my profile.

This whole process uses Artificial Intelligence heavily to transform the sound of my voice to [phonemes](#), phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

Amazon further describes how to capture and respond to key presses by the user.¹¹⁶

¹¹⁵ <https://aws.amazon.com/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost-effective-amazon-ec2-inf1-instances/>

¹¹⁶ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-for-screen-devices.html>

Keyboard handlers

APL provides support for keyboard events. You can define keyboard handlers to capture and respond to key presses. For details, see [Keyboard Events and Handlers](#).

297. On information and belief, each of the Accused Multi-modal Products and Services comprises merging the speech-based textual message and the non-speech-based textual message to generate a query.

298. For example, Amazon describes Alexa processing the transcript to determine what the user means.¹¹⁷

How Alexa Works

When I'm in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I'm saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as **Automatic Speech Recognition (ASR)**. Once the text is known, the second phase is to understand what I mean. This is **Natural Language Understanding (NLU)**. The output of NLU is an Intent (what does the customer want) and associated parameters. In this example ("Alexa, what's the weather today?"), the intent might be "**GetWeatherForecast**" and the parameter can be my postcode, inferred from my profile.

This whole process uses Artificial Intelligence heavily to transform the sound of my voice to **phonemes**, phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

299. Each of the Accused Multi-modal Products and Services comprises searching the query for text combinations.

300. For example, Amazon describes Alexa as listening for recognized text combinations to select an appropriate skill.¹¹⁸

¹¹⁷ <https://aws.amazon.com/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost-effective-amazon-ec2-inf1-instances/>

¹¹⁸ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/understand-name-free-interaction-for-custom-skills.html>

When Alexa receives a request from a customer without a skill name, such as "Alexa, play relaxing sounds with crickets," Alexa looks for skills that might fulfill the request. Alexa determines the best choice among eligible skills and hands the request to the skill. To provide a signal for Alexa to consider when routing name free requests and enable customers to launch your skill and intents without knowing or having to remember the skill's name, consider adding Name Free Interactions (NFI) container to the skill.

About Skill and Intent Launch phrases

Skill launch Phrases are an optional new way to teach Alexa how an end-customer might invoke your skill as a modal launch, along with the standard invocation pattern of "Alexa, open ". For example, "Alexa, Open Tasty Recipes Skill" you might have something like "can you give me a tasty recipe".

Intents and Intent Launch Phrases: The nameFreeInteraction container allows you to identify which intents are natural ingress points (or invocation points) for your skill. Where end users might currently need to say "Alexa, open Tasty Recipes Skill and give me a recipe for meatballs", you can indicate to Alexa that your getRecipe intent is a natural ingress point, and provide common natural phrases that would be used to deep-link into this intent such as simply, "give me a tasty recipe for meatballs". The phrases provided in the nfi container give you the mechanism to indicate which sample utterances/use cases are of the highest importance so that Alexa can weight these utterances among the rest. Note: slots are not supported for skill or intent launch phrases within the nameFreeInteraction container, however the sample utterances you provide in your interaction model that contain slots will still be used for training data.

The most important rule is to design intents with a list of sample utterances that captures key words and phrases that can invoke the skill in a name- free manner. Keep your utterances specific to your skills use-cases and avoid targeting generic utterances. The Alexa ML/AI models are designed in part to route traffic to skills with a high correlation between skill functionality and target utterances.

Guidelines

To help ensure that Alexa directs to your skill only those customer requests, questions, or statements that can be supported by the skill, build sample utterances that capture your skill's functionality. Provide sample utterances that you predict customers can use to interact with the skill naturally. For a better user experience, provide a large variety of sample utterances written in different forms. For example, name-free utterances for a horoscope skill might include:

- *"What is the horoscope for gemini?"*
- *"Today's horoscope"*
- *"I want to know my horoscope"*
- *"What does my star sign say?"*

301. Each of the Accused Multi-modal Products and Services comprises comparing the text combinations to entries in a context description grammar.

302. For example, Amazon describes Alexa as comparing text combinations to a list of sample utterances that can invoke a skill in a name-free manner.¹¹⁹

The most important rule is to design intents with a list of sample utterances that captures key words and phrases that can invoke the skill in a name-free manner. Keep your utterances specific to your skills use-cases and avoid targeting generic utterances. The Alexa ML/AI models are designed in part to route traffic to skills with a high correlation between skill functionality and target utterances.

Guidelines

To help ensure that Alexa directs to your skill only those customer requests, questions, or statements that can be supported by the skill, build sample utterances that capture your skill's functionality. Provide sample utterances that you predict customers can use to interact with the skill naturally. For a better user experience, provide a large variety of sample utterances written in different forms. For example, name-free utterances for a horoscope skill might include:

- *"What is the horoscope for gemini?"*
- *"Today's horoscope"*
- *"I want to know my horoscope"*
- *"What does my star sign say?"*

Amazon further describes how grammar has “been a tool in Alexa’s NLU toolkit since well before the first Echo device shipped.”¹²⁰


¹¹⁹ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/understand-name-free-interaction-for-custom-skills.html>

¹²⁰ <https://www.amazon.science/blog/tools-for-generating-synthetic-data-helped-bootstrap-alexa-s-new-language-releases>

Rules of Grammar

Grammars have been a tool in Alexa’s NLU toolkit since well before the first Echo device shipped. A grammar is a set of rewrite rules for varying basic template sentences through word insertions, deletions, and substitutions.

Below is a very simple grammar, which models requests to play either pop or rock music, with or without the modifiers “more” and “some”. Below the rules of the grammar is a diagram of a computational system (a finite-state transducer, or FST) that implements them.



A toy grammar, which can model requests to play pop or rock music, with or without the modifiers “some” or “more”, and a diagram of the resulting finite-state transducer. The question mark indicates that the some_more variable is optional.

303. Each of the Accused Multi-modal Products and Services comprises accessing a plurality of domain agents that are associated with the context description grammar.

304. For example Amazon describes Alexa invoking custom skills.¹²¹

¹²¹ <https://developer.amazon.com/en-US/docs/alexa/custom-skills/understanding-how-users-invoke-custom-skills.html>

About invoking custom skills

Users can invoke Alexa custom skills in two ways:

- Use a phrase supported by the Alexa service in combination with the invocation name for a custom skill to request information, ask a question, or tell Alexa to do something ("Alexa, *ask* Daily Horoscopes *for* the horoscope for Gemini"). See the following sections:
 - [Invoke a skill with a specific request \(intent\)](#)
 - [Invoke a skill with no specific request \(no intent\)](#)
 - [Make sure that the sample utterances support the invocation phrases](#)
 - [Other words and phrases users can include in their utterances](#)
- Indirectly invoke an Alexa custom skill through a name-free interaction. In this case, the user asks Alexa to perform a task, without naming the skill that should fulfill the request. ("Alexa, what is my horoscope?"). See the following section: [Indirectly invoke a skill with name-free interaction](#).

Amazon describes a variety of methods for invoking skills.¹²²

invocation

The act of beginning an interaction with a particular Alexa ability. For example, if a customer wants to wake Alexa to use the Horoscope skill, "Alexa, ask Horoscope for today's reading."

Alexa then follows up after the invocation and asks, "What horoscope sign would you like?"

Types of invocations include: [full intent invocation](#), [partial intent invocation](#), and [no intent invocation](#).

full intent invocation

A customer's request that contains all information Alexa needs to make the request actionable. For example, "Alexa, ask History Buff what happened on June 3rd."

partial intent invocation

A customer's request that contains the customer's intent, but is missing a required slot. For example, "Alexa, ask Horoscope for today's reading." Here, the required zodiac sign is missing, and Alexa needs to obtain that information from the customer.

no intent invocation

A customer's request with no intent or slot information. For example, "Alexa, open History Buff."

305. Each of the Accused Multi-modal Products and Services comprises generating a relevance score based on results from comparing the text combinations to entries in the context description grammar.

306. For example, Amazon describes Alexa's two-step neural shortlisting-reranking approach to find the most relevant skill for a given utterance.¹²³

¹²² <https://developer.amazon.com/en-US/docs/alexa/ask-overviews/alexa-skills-kit-glossary.html>

¹²³ <https://www.amazon.science/blog/the-scalable-neural-architecture-behind-alexa-ability-to-select-skills>

Alexa uses a two-step, scalable, and efficient neural shortlisting-reranking approach to find the most relevant skill for a given utterance. This post describes the first of those two steps, which relies on a neural model we call Shortlister. (I'll describe the second step in a follow-up post tomorrow.) Shortlister is a scalable and efficient architecture with a shared encoder, a personalized skill attention mechanism, and skill-specific classification networks. We outline this architecture in our [paper](#) "Efficient Large-Scale Neural Domain Classification with Personalized Attention", which we will present at the 56th Annual Meeting of the Association for Computational Linguistics ([ACL 2018](#)) in July.

The shared encoder network is hierarchical: Its lower layers are character-based and orthography sensitive and learn to represent each word in terms of character structure or shape; its middle layers are word-based, and with the outputs from the lower layers, they learn to represent an entire utterance. The skill attention mechanism is a separate network that is personalized per user. It computes a summary vector that describes which skills are enabled in a given user's profile and how relevant they are to the utterance representation. Both the utterance representation vector and the personalized skill-summary vector feed into a battery of skill-specific classification networks, one network for each skill.


307. Each of the Accused Multi-modal Products and Services comprises selecting one or more domain agents based on results from the relevance score.

308. For example, Amazon describes how Alexa uses Shortlister and HypRank to select the most relevant skills.¹²⁴

¹²⁴ <https://developer.amazon.com/fr/blogs/alexa/post/0fecdb38-97c9-48ac-953b-23814a469cfc/skill-discovery>

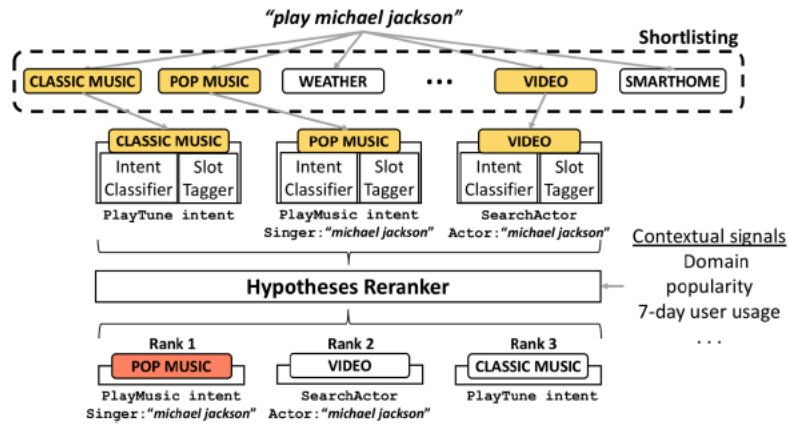
Surfacing the Most Relevant Skills via Shortlister and HypRank

It's important that customers are able to easily reengage with their skills – even if they describe what they want instead of calling the skill by name. Let's look again at the phrase "Alexa, remove an ink stain" where we have an utterance, "remove a {stain type} stain," and a slot value, stain type of "ink," that can be mapped to a skill's interaction model. To make this work, two techniques are employed to deliver the right experience. The first is a neural model architecture we call Shortlister, which solves a domain classification problem to find the most statistically significant (k -best) matches of candidate skills (details in a recent [blog](#) and [research paper](#)). The goal of Shortlister is to efficiently identify as many relevant skills as possible – it's high volume. The second is a hypothesis reranking network, we call HypRank. It uses contextual signals to select the most relevant skills – it's high precision (details in a recent [blog](#) and [research paper](#)). When combined, these two help classify which skill can provide the most likely correct response from an invocation name-free utterance. Beyond helping people reengage with the skills they've manually or automatically enabled, the Shortlister and HypRank surface specific and relevant skills to people that have never used them.



Amazon has further provided a specific example of the operation of the Hypotheses Reranker.¹²⁵

¹²⁵ <https://www.amazon.science/blog/hyprank-how-alex-a-determines-what-skill-can-best-meet-a-customers-need>



A high-level flow of the two-step shortlisting-reranking approach

The Challenge


The problem here is essentially a domain classification problem over the k -best candidate skills returned by the shortlisting system, which we call Shortlister. The goal of Shortlister is to achieve high recall — to identify as many pertinent skills as possible — with maximum efficiency. On the other hand, the goal of the reranking network, HypRank, is to use rich contextual signals to achieve high precision — to select the most pertinent skills. Designing HypRank comes with its own challenges:

309. Each of the Accused Multi-modal Products and Services comprises obtaining content that is gathered by the selected domain agents.

310. For example, Amazon describes how Alexa collects information from content providers by an Alexa skill.¹²⁶

¹²⁶ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-bp-understand-apl-architecture.html>

To plan your APL experiences, you need to understand the four technical *actors* that play a role in rendering APL *documents* on Alexa-enabled devices with screens. The following diagram shows a typical APL architecture at a high level.



APL architecture with four technical actors.

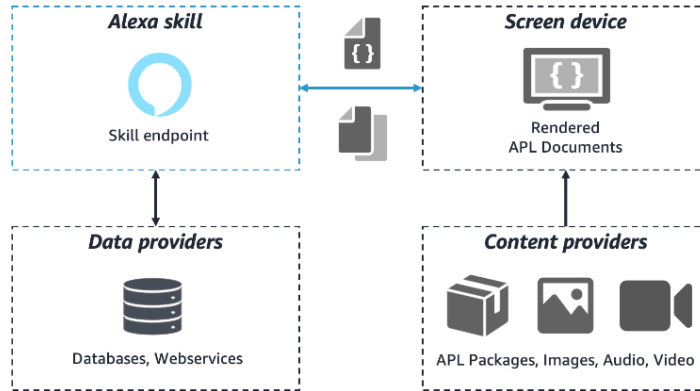
- **Alexa skill** — A skill endpoint that initiates an APL experience by sending out an APL *document* template in response to a regular skill request such as an `IntentRequest`.
- **Data providers** — APL templates often come with data that should be rendered with it so that the skill often attaches a JSON data source to the template that is sent to a screen device. This data typically doesn't reside within a skill, but resides in external data stores, such as a database that the skill reads from and writes to.
- **Screen device** — When the skill has sent an APL template with an attached data source, an Alexa-enabled device with a screen, such as an Echo Show or Fire TV, is responsible for inflating and rendering the template on screen.
- **Content providers** — APL templates that are sent by a skill often contain URL references to external content and media files; such as APL *packages*, images, audio and video files; that are stored publicly on the internet. An Alexa-enabled device with a screen resolves these references by downloading or streaming source files from external content providers, such as *content delivery networks* (CDN).

311. Each of the Accused Multi-modal Products and Services comprises generating a response from the content, wherein the content is arranged in a selected order based on results from the relevance score.

312. For example, Amazon describes how Alexa arranges the content from the content providers based on the APL template to be displayed on the screen device.¹²⁷

¹²⁷ <https://developer.amazon.com/en-US/docs/alexa/alexa-presentation-language/apl-bp-understand-apl-architecture.html>

To plan your APL experiences, you need to understand the four technical *actors* that play a role in rendering APL *documents* on Alexa-enabled devices with screens. The following diagram shows a typical APL architecture at a high level.



APL architecture with four technical actors.

- **Alexa skill** — A skill endpoint that initiates an APL experience by sending out an APL *document* template in response to a regular skill request such as an `IntentRequest`.
- **Data providers** — APL templates often come with data that should be rendered with it so that the skill often attaches a JSON data source to the template that is sent to a screen device. This data typically doesn't reside within a skill, but resides in external data stores, such as a database that the skill reads from and writes to.
- **Screen device** — When the skill has sent an APL template with an attached data source, an Alexa-enabled device with a screen, such as an Echo Show or Fire TV, is responsible for inflating and rendering the template on screen.
- **Content providers** — APL templates that are sent by a skill often contain URL references to external content and media files; such as APL *packages*, images, audio and video files; that are stored publicly on the internet. An Alexa-enabled device with a screen resolves these references by downloading or streaming source files from external content providers, such as *content delivery networks* (CDN).

Amazon further describes how Alexa Shopping Actions API can add a shopping experience to a skill and how the API can further provide recommended products to a user in response to a request.¹²⁸

¹²⁸ <https://developer.amazon.com/en-IN/docs/alexa/alexa-shopping/alexa-shopping-actions-for-alexa-skills-api-reference.html#recommend-action>

Request

To enable the customer to hear recommendations for specific products, add a `Connections.StartConnection` directive to your response.

Request example

```
{
  "type": "Connections.StartConnection",
  "uri": "connection://AMAZON.RecommendShoppingProducts/1",
  "input": {
    "products": [{
      "asin": "ASIN1"
    },
    {
      "asin": "ASIN2",
    },
    {
      "asin": "ASIN3"
    }
  ],
  "attribution": {
    "associateId": "some.associate.id.1",
    "trackingId": "some.tracking.id.1"
  }
},
"token": "RecommendProducts",
"onCompletion": "SEND_ERRORS_ONLY"
}
```

313. Amazon has long known about the '039 Patent.

314. Amazon has been aware of the '039 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.

315. Amazon knew or should have known that Amazon's actions infringe one or more of the claims of the '039 Patent because Amazon has the technical expertise to understand the scope and content of the '039 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Multi-modal Products and Services, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '039 Patent at least as of the filing of this Complaint.

316. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 13 of the '039 Patent in violation of at least 35 U.S.C. § 271(b), and (c).

317. Users of the Accused Multi-modal Products and Services directly infringe at least Claim 13 of the '039 Patent when they use the Accused Multi-modal Products and Services in the ordinary, customary, and intended way.

318. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Multi-modal Products and Services within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Multi-modal Products and Services to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Multi-modal Products and Services in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Multi-modal Products and Services, which Amazon knew infringes at least Claim 13 of the '039 Patent, or, alternatively, was willfully blind to the infringement.

319. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the Accused Multi-modal Products and Services within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Multi-modal Products and Services in the United States by activities related to selling, marketing, advertising, promotion,

support, and distribution of the Accused Multi-modal Products and Services, which Amazon knew infringes at least Claim 13 of the '039 Patent, or, alternatively, was willfully blind to the infringement.

320. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 13 of the '039 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '039 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

321. Amazon is not licensed or otherwise authorized to practice the claims of the '039 Patent.

322. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '039 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 13.

323. As a result of Amazon's infringement of the '039 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

324. On information and belief, in addition to Amazon's knowledge of the '039 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '039 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that

infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '039 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

325. Accordingly, Amazon's infringement of the '039 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

326. Amazon's infringement of Dialect's rights under the '039 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SEVENTH COUNT
(Infringement of U.S Patent No. 9,495,957)

327. Dialect incorporates by reference the allegations set forth in Paragraphs 1-326 of the Complaint as though fully set forth herein.

328. The claims of the '957 Patent are valid and enforceable.

329. The claims of the '957 Patent are directed to patentable subject matter. Specifically, Claim 1 of the '957 Patent is directed to a novel and tangible computerized system for processing natural language utterances generating ranked scores for context entries from a context stack by comparing context entries from the context stack with words determined by performing speech recognition on the utterance, identifying a context entry based on the ranked scores, and determining a command or request associated based on the words and the context information. The inventive, tangible claimed structures of the '957 Patent improve on existing natural language

processing systems as described above. The claimed inventions provide specific concrete solutions not known or used in the prior art.

330. In violation of 35 U.S.C. § 271(a) and without authority from Plaintiff, Amazon has directly infringed by making, using, offering for sale, selling, and/or importing into the Commonwealth of Virginia, this judicial district, and elsewhere in the United States products and services that embody the invention disclosed and claimed in the '957 Patent, including at least the Accused Products.

331. Each of the Accused Products contains elements that are identical or equivalent to each claimed element of the patented invention pointed out by at least Claim 1 of the '957 Patent.

332. Each of the Accused Products comprises a system for processing a natural language utterance, the system including one or more processors executing one or more computer program modules.

333. For example, Amazon's Alexa devices comprise one or more processors executing one or more computer program modules. Furthermore, Amazon describes Alexa as processing natural language utterances.¹²⁹

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



¹²⁹ <https://developer.amazon.com/en-US/alexa/alexa-skills-kit/get-deeper/custom-skills>

334. Each of the Accused Products comprises generating a context stack comprising context information that corresponds to a plurality of prior utterances, wherein the context stack includes a plurality of context entries.

335. For example, Amazon describes the use of a shortlisting-reranking approach used by Alexa.¹³⁰

Alexa uses a two-step, scalable, and efficient neural shortlisting-reranking approach to find the most relevant skill for a given utterance. This post describes the first of those two steps, which relies on a neural model we call Shortlister. (I'll describe the second step in a follow-up post tomorrow.) Shortlister is a scalable and efficient architecture with a shared encoder, a personalized skill attention mechanism, and skill-specific classification networks. We outline this architecture in our [paper](#) "Efficient Large-Scale Neural Domain Classification with Personalized Attention", which we will present at the 56th Annual Meeting of the Association for Computational Linguistics ([ACL 2018](#)) in July.

The shared encoder network is hierarchical: Its lower layers are character-based and orthography sensitive and learn to represent each word in terms of character structure or shape; its middle layers are word-based, and with the outputs from the lower layers, they learn to represent an entire utterance. The skill attention mechanism is a separate network that is personalized per user. It computes a summary vector that describes which skills are enabled in a given user's profile and how relevant they are to the utterance representation. Both the utterance representation vector and the personalized skill-summary vector feed into a battery of skill-specific classification networks, one network for each skill.

Amazon further provides examples of how Alexa uses records of past interactions to interpret a natural language request.¹³¹

¹³⁰ <https://www.amazon.science/blog/the-scalable-neural-architecture-behind-alexas-ability-to-select-skills>

¹³¹ <https://www.amazon.science/blog/how-alexa-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

336. Each of the Accused Products comprises receiving the natural language utterance, wherein the natural language utterance is associated with a command or is associated with a request.

337. For example, Amazon describes Alexa as receiving the natural language utterance, which includes a request.¹³²

How an Alexa Skill Works

An Alexa skill has both an interaction model—or voice user interface—and application logic. When a customer speaks, Alexa processes the speech in the context of your interaction model to determine the customer request. Alexa then sends the request to your skill application logic, which acts on it. You provide your application logic as a back-end cloud service hosted by Alexa, AWS, or another server.



338. Each of the Accused Products comprises determining one or more words of the natural language utterance by performing speech recognition on the natural language utterance.

¹³² <https://developer.amazon.com/en-US/alexa/alexa-skills-kit/get-deeper/custom-skills>

339. For example, Amazon discloses that Alexa transcribes the speech to text before processing the text to understand what the user means.¹³³

How Alexa Works

When I'm in my living room and ask Alexa about the weather, I trigger a complex system. First, the on-device chip detects the wake word (Alexa). Once detected, the microphones record what I'm saying and stream the sound for analysis in the cloud. At a high level, there are two phases to analyze the sound of my voice. First, Alexa converts the sound to text. This is known as **Automatic Speech Recognition (ASR)**. Once the text is known, the second phase is to understand what I mean. This is **Natural Language Understanding (NLU)**. The output of NLU is an Intent (what does the customer want) and associated parameters. In this example ("Alexa, what's the weather today?"), the intent might be "**GetWeatherForecast**" and the parameter can be my postcode, inferred from my profile.

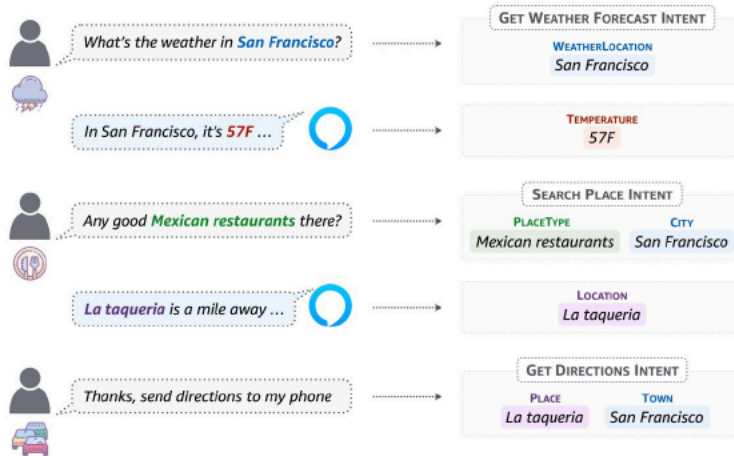
This whole process uses Artificial Intelligence heavily to transform the sound of my voice to **phonemes**, phonemes to words, words to phrases, phrases to intents. Based on the NLU output, Alexa routes the intent to a service to fulfill it. The service might be internal to Alexa or external, like one of the skills activated on my Alexa account. The fulfillment service processes the intent and returns a response as a JSON document. The document contains the text of the response Alexa must say.

340. Each of the Accused Products comprises identifying, from among the plurality of context entries, one or more context entries that correspond to the one or more words, wherein the context information includes the one or more context entries, wherein identifying the one or more context entries comprises: comparing the plurality of context entries to the one or more words.

341. For example, Amazon describes Alexa identifying one or more words based on the Alexa intent to assign the data to different slots.¹³⁴

¹³³ <https://aws.amazon.com/blogs/aws/majority-of-alex-a-now-running-on-faster-more-cost-effective-amazon-ec2-infl-instances/>

¹³⁴ <https://www.amazon.science/blog/how-alex-a-is-learning-to-converse-more-naturally>



Different Alexa intents (weather, search, directions) might assign the same data (San Francisco) to different slots (WeatherLocation, City, Town). Tracking references by carrying slots over from one intent to the next is crucial for maintaining a coherent conversation.

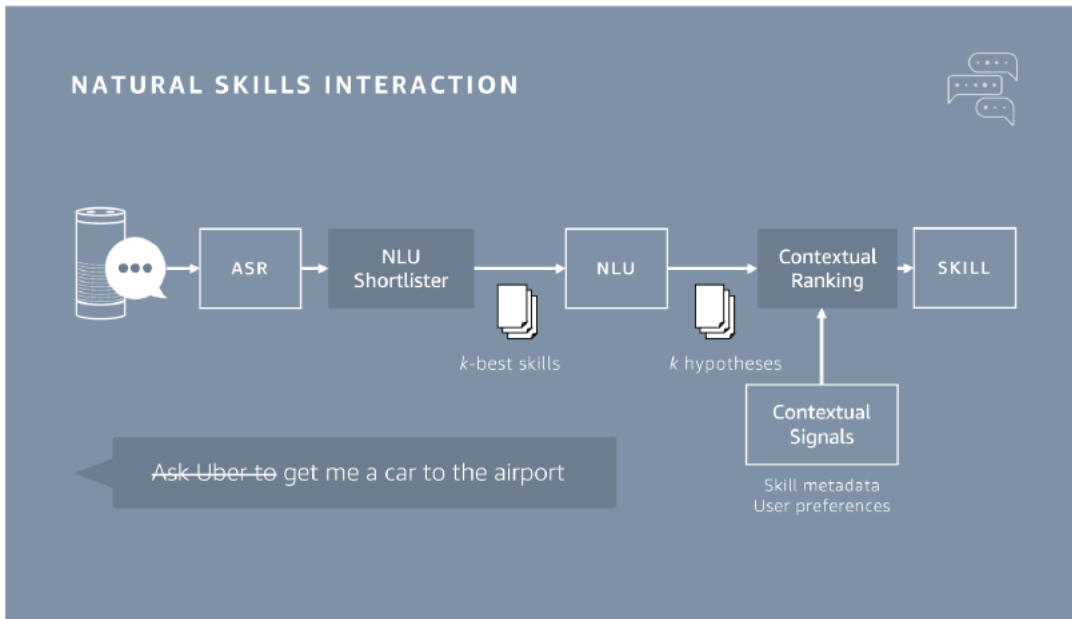
342. Each of Accused Products comprises generating, based on the comparison, one or more rank scores for individual context entries of the plurality of context entries.

343. For example, Amazon describes how Alexa uses Shortlister and HypRank to surface the most relevant skills.¹³⁵

¹³⁵ <https://developer.amazon.com/fr/blogs/alexa/post/0fecdb38-97c9-48ac-953b-23814a469cfc/skill-discovery>

Surfacing the Most Relevant Skills via Shortlister and HypRank

It's important that customers are able to easily reengage with their skills – even if they describe what they want instead of calling the skill by name. Let's look again at the phrase “Alexa, remove an ink stain” where we have an utterance, “remove a {stain type} stain,” and a slot value, stain type of “ink,” that can be mapped to a skill's interaction model. To make this work, two techniques are employed to deliver the right experience. The first is a neural model architecture we call Shortlister, which solves a domain classification problem to find the most statistically significant (*k*-best) matches of candidate skills (details in a recent [blog](#) and [research paper](#)). The goal of Shortlister is to efficiently identify as many relevant skills as possible – it's high volume. The second is a hypothesis reranking network, we call HypRank. It uses contextual signals to select the most relevant skills – it's high precision (details in a recent [blog](#) and [research paper](#)). When combined, these two help classify which skill can provide the most likely correct response from an invocation name-free utterance. Beyond helping people reengage with the skills they've manually or automatically enabled, the Shortlister and HypRank surface specific and relevant skills to people that have never used them.



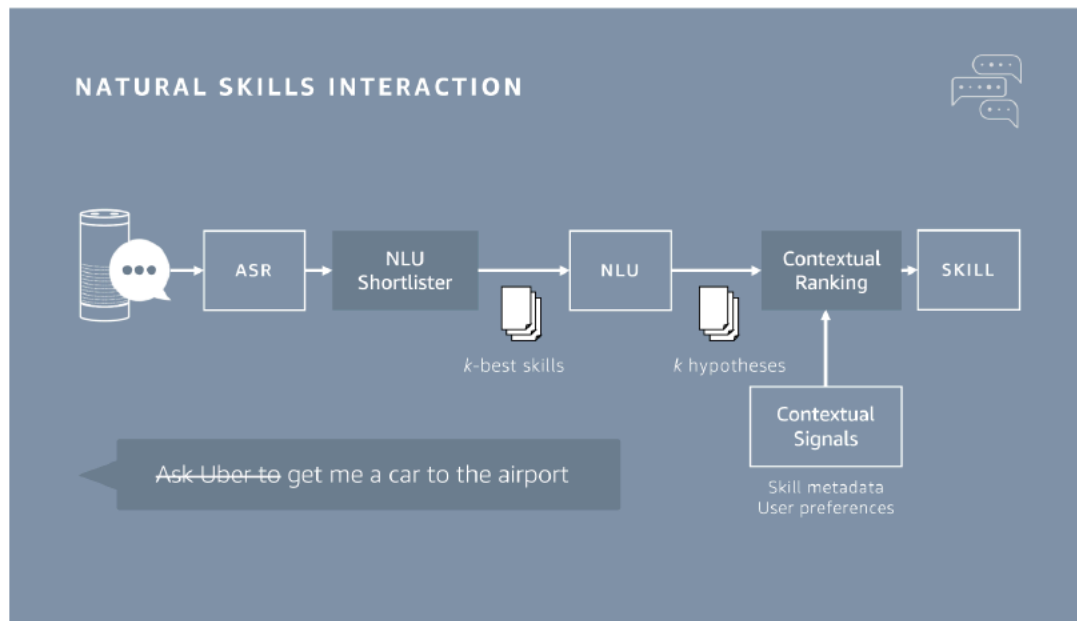
344. Each of the '957 Patent Accused Products comprises identifying, based on the one or more rank scores, the one or more context entries from among the plurality of context entries.

345. For example, Amazon describes how Alexa uses Shortlister and HypRank to identify the most relevant skills.¹³⁶

¹³⁶ *Id.*

Surfacing the Most Relevant Skills via Shortlister and HypRank

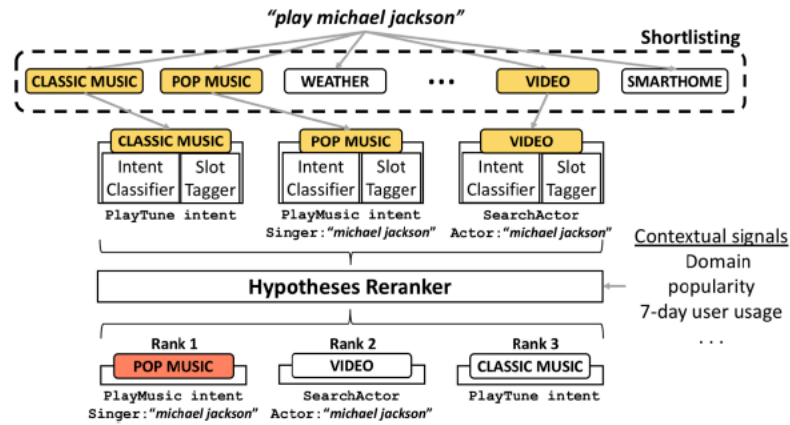
It's important that customers are able to easily reengage with their skills – even if they describe what they want instead of calling the skill by name. Let's look again at the phrase "Alexa, remove an ink stain" where we have an utterance, "remove a {stain type} stain," and a slot value, stain type of "ink," that can be mapped to a skill's interaction model. To make this work, two techniques are employed to deliver the right experience. The first is a neural model architecture we call Shortlister, which solves a domain classification problem to find the most statistically significant (k -best) matches of candidate skills (details in a recent [blog](#) and [research paper](#)). The goal of Shortlister is to efficiently identify as many relevant skills as possible – it's high volume. The second is a hypothesis reranking network, we call HypRank. It uses contextual signals to select the most relevant skills – it's high precision (details in a recent [blog](#) and [research paper](#)). When combined, these two help classify which skill can provide the most likely correct response from an invocation name-free utterance. Beyond helping people reengage with the skills they've manually or automatically enabled, the Shortlister and HypRank surface specific and relevant skills to people that have never used them.



346. Each of the Accused Products comprises determining, based on the determined one or more words and the context information, the command or the request associated with the natural language utterance.

347. For example, Amazon describes how Alexa uses the Hypotheses Reranker to determine the command in a specific example.¹³⁷

¹³⁷ <https://www.amazon.science/blog/hyprank-how-alexa-determines-what-skill-can-best-meet-a-customers-need>



A high-level flow of the two-step shortlisting-reranking approach

The Challenge

The problem here is essentially a domain classification problem over the k-best candidate skills returned by the shortlisting system, which we call Shortlister. The goal of Shortlister is to achieve high recall — to identify as many pertinent skills as possible — with maximum efficiency. On the other hand, the goal of the reranking network, HypRank, is to use rich contextual signals to achieve high precision — to select the most pertinent skills. Designing HypRank comes with its own challenges:

348. Amazon has long known about the '957 Patent.
349. Amazon has been aware of the '957 Patent from multiple communications, presentations, and licensing negotiations between VoiceBox Technologies and Amazon between at least 2011 and 2017 and from VoiceBox Technologies personnel hired by Amazon.
350. Amazon knew or should have known that Amazon's actions infringe one or more of the claims of the '957 Patent because Amazon has the technical expertise to understand the scope and content of the '957 Patent, because Amazon is a major provider of voice recognition products and services, and because Amazon knows the design, function, and operation of the Accused Products, as well as the nature and extent of their use by others. At a minimum, Amazon has knowledge of the '957 Patent at least as of the filing of this Complaint.

351. Further, on information and belief, Amazon has actively induced and/or contributed to infringement of at least Claim 1 of the '957 Patent in violation of at least 35 U.S.C. § 271(b), (c), and (f).

352. Users of the Accused Products and Services directly infringe at least Claim 1 of the '957 Patent when they use the Accused Products in the ordinary, customary, and intended way.

353. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing businesses and consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to businesses and consumers within the United States, and instructing and encouraging such businesses and consumers to use the Accused Products in the ordinary, customary, and intended way by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Products, which Amazon knew infringes at least Claim 1 of the '957 Patent, or, alternatively, was willfully blind to the infringement.

354. On information and belief, Amazon's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing Amazon's customers to commit acts of infringement with respect to the Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Products in the United States by activities related to selling, marketing, advertising, promotion, support, and distribution of the Accused Products, which Amazon knew infringes at least Claim 1 of the '957 Patent, or, alternatively, was willfully blind to the infringement.

355. On information and belief, in violation of 35 U.S.C. § 271(c), Amazon's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least Claim 1 of the '957 Patent, constituting a material part of the invention. On information and belief, Amazon knows and has known the same to be especially made or especially adapted for use in an infringement of the '957 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use.

356. On information and belief, in violation of 35 U.S.C. § 271(f)(1), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the patented invention of at least Claim 1 of the '957 Patent, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

357. On information and belief, in violation of 35 U.S.C. § 271(f)(2), Amazon's infringement further includes without authority supplying or causing to be supplied in or from the United States components of the patented invention of at least Claim 1 of the '957 Patent that are especially made or especially adapted for use in the invention and not staple articles or commodities of commerce suitable for substantial noninfringing use, where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components will be combined outside of the United States in a manner that would infringe the patent if such combination occurred within the United States.

358. Amazon is not licensed or otherwise authorized to practice the claims of the '957 Patent.

359. Thus, by its acts, Amazon has injured Dialect and is liable to Dialect for directly and/or indirectly infringing one or more claims of the '957 Patent, whether literally or under the doctrine of equivalents, including without limitation Claim 1.

360. As a result of Amazon's infringement of the '957 Patent, Dialect has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty with interest and costs.

361. On information and belief, in addition to Amazon's knowledge of the '957 Patent as set forth above both prior to and as a result of the filing of this Complaint, Amazon has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '957 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made Amazon aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '957 Patent. On information and belief, discovery will reveal additional facts and circumstances from which Amazon's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

362. Accordingly, Amazon's infringement of the '957 Patent has also been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

363. Amazon's infringement of Dialect's rights under the '957 Patent will continue to damage Dialect, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

NOTICE

364. Plaintiff has complied with the notice requirement of 35 U.S.C. § 287 and has not and does not currently distribute, sell, offer for sale, or make products embodying the Asserted Patents.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment and seeks relief from Defendants as follows:

- a. For judgment that Amazon has infringed and continues to infringe the claims of the '720, '006, '327, '468, '845, '039, and '957 Patents;
- b. For a permanent injunction against Amazon and its respective officers, directors, agents, servants, affiliates, employees, divisions, branches, subsidiaries, parents, and all other acting in active concert therewith from infringement of the '720, '006, '327, '468, '845, '039, and '957 Patents;
- c. For an accounting of all damages sustained by Plaintiff as a result of Amazon's acts of infringement;
- d. In the event Amazon is not permanently enjoined, for a mandatory future royalty payable on each and every future sale by Amazon of a product or service that is found to infringe one or more of the Asserted Patents and on all future products and services which are not colorably different from products and services found to infringe;
- e. For a judgment and order finding that Amazon's infringement is willful and/or egregious and awarding to Plaintiff enhanced damages pursuant to 35 U.S.C. § 284;

- f. For a judgment and order requiring Amazon to pay Plaintiff's damages, costs, expenses, and pre- and post-judgment interest for its infringement of the '720, '006, '327, '468, '845, '039, and '957 Patents as provided under 35 U.S.C. § 284;
- g. For a judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees; and
- h. For such other and further relief in law and in equity as the Court may deem just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiff demands a trial by jury in this action for all issues triable by a jury.

Dated: July 31, 2023

Respectfully Submitted,
/s/ Walter D. Kelley, Jr. _____
Walter D. Kelley, Jr. (VA Bar No. 21622)
Tara Zurawski (VA Bar No. 73602)
HAUSFELD LLP
888 16th Street N.W., Suite 300
Washington, DC 20006
Tel: (202) 540-7200
Fax: (202) 540-7201
wkelley@hausfeld.com
tzurawski@hausfeld.com

Garland Stephens
garland@bluepeak.law
Robert Magee
robert@bluepeak.law
Richard Koehl
richard@bluepeak.law
Jeff Risher
jeff@bluepeak.law
BLUE PEAK LAW GROUP LLP
Telephone: 281-972-3036