

Appendix I - U.S. Patent No. 8,620,659

Cortana

Element	Analysis for Cortana
[42.0] A method of processing natural language utterances, the method being implemented by a computer system that includes one or more processors executing one or more computer program instructions which, when executed, perform the method, the method comprising:	Cortana meets this element.  https://www.youtube.com/watch?app=desktop&v=xuiKfUgEvfo What is the Cortana search box in Windows? The Cortana search box in Windows is a central place to enter queries and access information quickly. You can use it to search for files, apps, settings, and perform web searches without opening a browser. Just type or speak your query in the search box, and Cortana will provide relevant results. https://www.lenovo.com/us/en/glossary/cortana/?orgRef=https%253A%252F%252Fwww.google.com%252F

Element	Analysis for Cortana
	Cortana typically serves as a dictionary, thesaurus, or encyclopedia to its users. It comes in handy when you intend to save time on redundant tasks that machines can accomplish on their own. You can just type-in or ask your questions and get immediate responses without the need for any intervention. What's the synonym of analyze?, What is the length of the Great Wall of China?, What's the latitude and longitude of San Francisco?, and so on, are examples of questions that can be asked. https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-cortana/ Cortana in Windows 10 The Cortana application in the Windows 10 version 2004 and later is a chat application with your personal productivity assistant. This application helps users to quickly get information across the Microsoft 365 platform by just using simply typed or spoken Keywords in queries. https://www.infiflex.com/making-microsoft-365-smarter-with-cortana
[42.1] receiving a first input of a user that comprises a natural language utterance;	Cortana meets this element. See element [42.0].
[42.2] generating an interpretation of the natural language utterance based on one or more recognized words of the natural language utterance;	Cortana meets this element. For example, Cortana uses automated speech recognition (ASR) technology and natural language processing (NLP) technology to generate an interpretation of natural language command, taking into account contextual information and information relating to a user.

Element	Analysis for Cortana
	Cortana, like other Windows voice-enabled functionality, relies on automated speech recognition (ASR) systems that run both on the device and in the cloud. These systems include enhanced abilities to filter out background noise, recognize accented speech, and otherwise tell the difference between words and phrases that should be acted upon and words and phrases that should be ignored. However, to access the cloud ASR systems, Windows must have user consent and a connection to the internet. These may not be available during the OOBE flow. In that case, Windows uses a more limited, offline-only ASR system that does not have the level of robustness that Cortana usually offers. https://learn.microsoft.com/en-us/windows-hardware/customize/desktop/cortana-voice-support How Cortana learns In the initial stages, when you set up Cortana for your Microsoft account, it gathers your personal information such as name, email, age, and other facts that you provided for that account. Logging into your Microsoft account is needed to use Cortana. Secondly, the more you use Cortana, the better it performs. As stated earlier, Cortana learns and fine tunes itself when you tend to ask more and more questions. Also, allowing it to access your location, search history, photos, documents, music, emails, etc., gives it a chance to optimize its learning. https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-cortana/ Reddy says he is now working on helping make user conversations with Cortana more personal and contextual. Sometimes, people ask Cortana follow-up questions, and the team is working on making sure she handles those as skillfully as the initial query. https://news.microsoft.com/life/why-cortanas-awesome-and-she-knows-it/ (emphasis added)

Element	Analysis for Cortana
	Furthermore, if given permission to access phone data, Cortana can read calendars, contacts, and email to improve its knowledge of context and connections. Heck, who plays classical trumpet in a local orchestra, might receive a calendar update about a change in rehearsal time. Cortana would let him know about the change and alert him if the new time conflicts with another appointment. https://www.microsoft.com/en-us/research/blog/anticipating-more-from-cortana/ Our PDA LU is built to handle multi-domain, multi-turn, contextual query understanding [14, 9, 15, 16], subject to the constraints of the back-end data sources and experiences. The LU module semantically parses and analyses the queries according to a domain-specific semantic schema. Given a requirement to minimize latency and computational costs (dimensions 7,8) deep semantic analysis is not performed. Instead, LU delivers a flat analysis matching the experience requirements while allowing for free form natural language expressions that represent different user intents and slots.


Element	Analysis for Cortana
	Follow-up queries in isolation are ambiguous, <i>e.g.</i>, • Turn 1: “how is the weather in New York” (weather) • Turn 2: “what about the weekend” (weather) or • Turn 1: “how is my schedule” (calendar) • Turn 2: “what about the weekend” (calendar) LU modeling thus must be performed in a contextual manner. Here, conversational session information from the history greatly reduces the ambiguity of the current turn. We built a set of models addressing contextual domain, intent and slot modeling [14, 15, 16]. Our current approach involves modeling contextual intent prediction as a sequential tagging problem with previous utterances as features that condition this turn’s prediction. Contextual information from beyond LU, such as dialog state (where this is tracked by an experience), and PDA responses are also exploited as effective external features. Contextual modeling of queries reduces the likelihood of abrupt intent or domains switching leading to more coherent interaction during a multi-turn session.

Element	Analysis for Cortana
	4.1. Domain, Intent, Slot Modeling Prior work has shown a variety of approaches can be used for language understanding [17, 8, 20, 21, 15]. Some early systems used rule based parsing, hand-authored by an expert for a specific domain [17]. Recent state-of-the-art systems use machine learned models [21, 20, 15]. Semantic meaning of a query is represented as a semantic frame (SF), which is a tuple of (DOMAIN, INTENT, SLOTS) in the semantic space. The slots are a list of key-value pairs: (SLOTNAME, SLOTVALUE). In our system, domain classification is done using machine learned (ML) classifiers such as SVMs [25], which generate a score for each domain. We use a separate binary one-against all classifier for each domain. For each domain, we also train a domain specific intent and slot model. Intents and slots can be shared across different domains. Intent detection also uses SVM classification but is framed as a multi-class classification problem. Slot tagging is considered as a sequence classification problem. Conditional random fields (CRFs) [26] and more recently deep learning techniques [9, 31, 16] are used for slot tagging. All of the LU models use weighted entity lexicons [18, 19] and are additionally trained with artificial out-of-vocabulary (OOV) words to improve generalization to unseen entities. The slot values can be further resolved into entities (e.g. a strongly typed object in some back-end data source) or canonicalized into a standard form (e.g. time/date values). To deal with the open domain nature of queries, each LU domain uses training data from all other domains as negative examples, including representative data for the web domain AN OVERVIEW OF END-TO-END LANGUAGE UNDERSTANDING AND DIALOG MANAGEMENT FOR PERSONAL DIGITAL ASSISTANTS, Microsoft Corporation, 2016, available at https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/CortanaLUDialog-FromSLTproceedings.pdf
[11.3] generating a request based on the interpretation;	Cortana meets this element. For example, Cortana generates a request to its skills based on the interpretation.


Element	Analysis for Cortana
	• Turn 1: “find french restaurants in seattle” • State 1: cuisine=“french”, place_type=“restaurants”, absolute_location=“seattle” • Turn 2: “how about chinese” • State 2: cuisine=“chinese”, place_type=“restaurants”, absolute_location=“seattle” The diagram illustrates the architecture of a Personal Digital Assistant (Cortana) across two turns. It is organized into three main columns: Input, Update state, and Policy. First Turn: Input: ASR / typed input feeds into a stack of boxes: Natural, Language, and Understanding (per domain). Update state: The Understanding box feeds into Slot/Entity Carry Over, which then feeds into Flexible Item Selection. Policy: Flexible Item Selection feeds into a stack of boxes: Service Providers, Specialist dialogs, Q&A, and Web Search. This stack feeds into Per Task Policy, which then feeds into HR + HS. HR + HS also receives input from a 'Previous turn context' block. Output: HR + HS feeds into a final box labeled 'Ranking & cross task policy'. Second (and subsequent) turns: - The process repeats, with the 'Previous turn context' block from the first turn feeding into the 'Understanding' box of the second turn. - The flow continues through Slot/Entity Carry Over, Flexible Item Selection, and the Policy stack (Service Providers, Specialist dialogs, Q&A, Web Search) to Per Task Policy and HR + HS. Fig. 1. Personal Digital Assistant Reactive System Architecture. AN OVERVIEW OF END-TO-END LANGUAGE UNDERSTANDING AND DIALOG MANAGEMENT FOR PERSONAL DIGITAL ASSISTANTS, Microsoft Corporation, 2016, available at https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/CortanaLUDialog-FromSLTproceedings.pdf
[11.4] transmitting the request to a domain agent for processing;	Cortana meets this element. On information and belief, Cortana uses multiple domain agents to handle user requests.

Element	Analysis for Cortana
	Can Cortana control my Windows settings? Yes, Cortana can control various Windows settings, such as adjusting volume, screen brightness, and power settings, all through voice commands. How can Cortana help with productivity? Cortana can enhance productivity by setting reminders, managing calendars, sending emails, providing weather updates, and answering quick questions, saving you time and effort. Does Cortana have location-based features? Yes, Cortana can use location data to offer location-based reminders, weather updates, traffic information, and restaurant recommendations, making it helpful for travel and navigation. Can Cortana perform math calculations? Yes, Cortana can perform basic to complex math calculations, including addition, subtraction, multiplication, division, and more, all by voice commands. https://www.lenovo.com/us/en/glossary/cortana/?orgRef=https%253A%252F%252Fgemini.google.com%252F

Element	Analysis for Cortana
	Top 5 Features of Cortana Cortana's ability to perform several tasks for its users, right from handling user queries, calendar events to managing smart home devices, makes it an appropriate multifunctional tool that anyone can use to ease their lives. It gives tough competition to Alexa and Siri when it comes to handling such a variety of functions. Let's take a look at the top five key features of Cortana: 1. Create reminders Cortana is a handy digital assistant that helps in creating reminders without much intervention. For example, you can address Cortana by saying, 'Hey Cortana, remind me to attend Tom's conference call at 6 p.m.'. You can also enable location settings and give a command to Cortana, such as, 'Hey Cortana, remind me to buy groceries when I am returning home from the office in the evening'. Another way of creating reminders is by doing so manually. You can follow these steps: Open up Cortana > click on the Notebook icon > Tap Reminder on the left menu > create a reminder. 2. Control smart home devices Cortana allows you to control several home devices such as Philips Hue Lights just by giving voice commands. For example, once you connect the Philips lighting system to Cortana, you can manage it with your voice. You can give instructions such as, 'Hey Cortana, please dim the lights after 15 minutes' or 'Hey Cortana, please switch on the lights in the living room'. 3. Look up words while browsing Cortana provides features to look up words or phrases while browsing the web. The assistant is already enabled on the Microsoft Edge browser, and makes it easy for users to search on the web. This feature can be availed by highlighting the words or phrases in contention, right-clicking on them, and then tapping on the 'Ask Cortana' option. This lets Cortana use the Bing search engine to identify the relevant information on the web and then show the results on the browser via a tiny pop up or window. 4. Get weather forecasts One of the easiest ways of checking weather forecasts on Cortana-enabled devices is by asking Cortana to report the weather conditions. You can give a voice command such as, 'Hey Cortana, what's the current humidity level in Los Angeles' or 'Hey Cortana, what will the outside temperature be in my city at 2 p.m.'? Such queries will allow Cortana to look for weather reports pertaining to the said location and also display the identified results for you. 5. Ask for directions While searching for locations on maps, it is quite inconvenient to input or type-in the entire address. Such inconvenience is resolved by Cortana as it allows you to input voice commands and performs the tasks of locating, finding the route, and also giving you directions to the destination all by itself. You can simply say, 'Hey Cortana, how can I reach XYZ (location name)?' or 'Hey Cortana, which route should I take to reach XYZ?'. https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-cortana/

Element	Analysis for Cortana
	Blog / Mobile / Android / Post Cortana Picks Up To-Do List Skills Paul Thurrott NOV 17, 2016 8  Microsoft announced today that it has improved its Cortana personal digital assistant with a new “to do list” skill. https://www.thurrott.com/windows/windows-10/85172/cortana-picks-list-skills In addition, developers can build Cortana skills in different domains.  https://blogs.windows.com/windowsdeveloper/2017/05/10/cortana-skills-kit-empowers-developers-

Element	Analysis for Cortana																						
	build-intelligent-experiences-millions-users/ Microsoft describes different domains of Cortana in its 2016 paper: Table 1. Data discussion <table data-bbox="596 423 1266 797"> <tr> <th data-bbox="596 423 793 451">Domain</th><th data-bbox="793 423 1266 451">Supported Experiences</th></tr> <tr> <td data-bbox="596 459 793 483">Alarm</td><td data-bbox="793 459 1266 483">setting, querying, and deleting alarms</td></tr> <tr> <td data-bbox="596 488 793 513">Calendar</td><td data-bbox="793 488 1266 513">creating and querying calendar entries</td></tr> <tr> <td data-bbox="596 518 793 542">Communication</td><td data-bbox="793 518 1266 542">making phone calls, sending emails and SMS</td></tr> <tr> <td data-bbox="596 547 793 571">Device Settings</td><td data-bbox="793 547 1266 571">commands for changing local device settings</td></tr> <tr> <td data-bbox="596 576 793 600">Document</td><td data-bbox="793 576 1266 600">document search on the local filesystem</td></tr> <tr> <td data-bbox="596 605 793 630">Entertainment</td><td data-bbox="793 605 1266 630">music playback</td></tr> <tr> <td data-bbox="596 634 793 659">Local</td><td data-bbox="793 634 1266 659">traffic information, driving directions, search for businesses and attractions</td></tr> <tr> <td data-bbox="596 664 793 688">Reminder</td><td data-bbox="793 664 1266 688">setting, querying, and deleting reminders</td></tr> <tr> <td data-bbox="596 693 793 717">Weather</td><td data-bbox="793 693 1266 717">weather information</td></tr> <tr> <td data-bbox="596 722 793 747">Web</td><td data-bbox="793 722 1266 747">general web search, chit-chat, and targeted question-answering not covered above</td></tr> </table> AN OVERVIEW OF END-TO-END LANGUAGE UNDERSTANDING AND DIALOG MANAGEMENT FOR PERSONAL DIGITAL ASSISTANTS, Microsoft Corporation, 2016, available at https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/CortanaLUDialog-FromSLTproceedings.pdf On information and belief, Cortana selects a domain agent and transmits the request to the domain agent. For example, the question, "Who won today's football match between Liverpool and Barcelona?" prompts a sports skill that follows the rules of a pre-coded script to fill in slots for the type of sport, information requested, date and teams. "Will it rain this weekend?" prompts a weather skill and follows pre-scripted rules to get the weekend forecast. https://news.microsoft.com/source/features/ai/microsoft-build-future-of-natural-language/	Domain	Supported Experiences	Alarm	setting, querying, and deleting alarms	Calendar	creating and querying calendar entries	Communication	making phone calls, sending emails and SMS	Device Settings	commands for changing local device settings	Document	document search on the local filesystem	Entertainment	music playback	Local	traffic information, driving directions, search for businesses and attractions	Reminder	setting, querying, and deleting reminders	Weather	weather information	Web	general web search, chit-chat, and targeted question-answering not covered above
Domain	Supported Experiences																						
Alarm	setting, querying, and deleting alarms																						
Calendar	creating and querying calendar entries																						
Communication	making phone calls, sending emails and SMS																						
Device Settings	commands for changing local device settings																						
Document	document search on the local filesystem																						
Entertainment	music playback																						
Local	traffic information, driving directions, search for businesses and attractions																						
Reminder	setting, querying, and deleting reminders																						
Weather	weather information																						
Web	general web search, chit-chat, and targeted question-answering not covered above																						

Element	Analysis for Cortana
	How Cortana skills work  How Cortana skills tie into Cognitive Services (click to enlarge) A Cortana skill is essentially a channel of artificial intelligence accessed through Cortana. There are built-in skills like Search, setting reminders, or launching applications; and there are add-in skills built by third parties. Add-in skills are accessed via invocations, key phrases which are registered with Microsoft, so that when the user says, for example, "Ask <invocation Name> <something>", the question is passed to the registered service rather than being handled by built-in skills. There are currently 17 words you can use before the invocation name, though these are language specific. Unfortunately the only language available in the preview is US English, though if you are in the UK or elsewhere, you can easily set Cortana to US English in order to test a Cortana skill. https://www.theregister.com/2017/07/26/hands_on_with_cortana_skills/

Element	Analysis for Cortana
	5.4. Hypothesis Ranking and Selection HRS is the mechanism by which the platform ranks and selects the final response shown to the user. The processing is done in two stages, with an initial ranking of the hypotheses (HR) followed by selection of the best hypothesis according to rank score and other features (HS). One of the key decisions in dealing with open domain requests is the recognition of requests that are not handled by any specialized experience provider and letting such requests fall through to web search. Ranking is carried out over dialog hypotheses, where each dialog hypothesis is an assembly of the SF and an experience provider response (<i>e.g.</i> back-end knowledge plus suggested dialog act). Making the selection of the system response as late as possible has the benefit that the ML model has full access to all LU analyses, as well as the full session context and knowledge of whether a particular experience produced a response. An experience is allowed to not respond. HRS is contextual in that it should learn to favor the continuation of an existing multi-turn dialog unless the user's utterance indicates a strong likelihood of wishing to switch domains. AN OVERVIEW OF END-TO-END LANGUAGE UNDERSTANDING AND DIALOG MANAGEMENT FOR PERSONAL DIGITAL ASSISTANTS, Microsoft Corporation, 2016, available at https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/CortanaLUDialog-FromSLTproceedings.pdf
[11.5] determining whether a personalized cognitive model associated with the user includes sufficient information for predicting one or more subsequent actions associated with the user, wherein the personalized cognitive model is generated based on a	Cortana meets this element. On information and belief, Cortana uses a personalized cognitive model based on a pattern of interactions between the user and system, and the personalized cognitive model is able to predict user actions.


Element	Analysis for Cortana
tracking of a pattern of interactions between the user and the system, and wherein the one or more subsequent actions include one or more actions predicted to occur after receiving the first input; and	How Cortana learns In the initial stages, when you set up Cortana for your Microsoft account, it gathers your personal information such as name, email, age, and other facts that you provided for that account. Logging into your Microsoft account is needed to use Cortana. Secondly, the more you use Cortana, the better it performs. As stated earlier, Cortana learns and fine tunes itself when you tend to ask more and more questions. Also, allowing it to access your location, search history, photos, documents, music, emails, etc., gives it a chance to optimize its learning. https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-cortana/ Reddy says he is now working on helping make user conversations with Cortana more personal and contextual. Sometimes, people ask Cortana follow-up questions, and the team is working on making sure she handles those as skillfully as the initial query. https://news.microsoft.com/life/why-cortanas-awesome-and-she-knows-it/ (emphasis added) Furthermore, if given permission to access phone data, Cortana can read calendars, contacts, and email to improve its knowledge of context and connections. Heck, who plays classical trumpet in a local orchestra might receive a calendar update about a change in rehearsal time. Cortana would let him know about the change and alert him if the new time conflicts with another appointment. https://www.microsoft.com/en-us/research/blog/anticipating-more-from-cortana/

Element	Analysis for Cortana
	LU modeling thus must be performed in a contextual manner. Here, conversational session information from the history greatly reduces the ambiguity of the current turn. We built a set of models addressing contextual domain, intent and slot modeling [14, 15, 16]. Our current approach involves modeling contextual intent prediction as a sequential tagging problem with previous utterances as features that condition this turn's prediction. Contextual information from beyond LU, such as dialog state (where this is tracked by an experience), and PDA responses are also exploited as effective external features. Contextual modeling of queries reduces the likelihood of abrupt intent or domains switching leading to more coherent interaction during a multi-turn session. https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/CortanaLUDialog-FromSLTproceedings.pdf
[11.6] predicting the one or more subsequent actions based on a generalized cognitive model in response to a determination that the personalized cognitive model does not include the sufficient information, wherein the generalized cognitive model is generated based on a tracking of patterns of interactions between a plurality of users and the system	Cortana meets this element. On information and belief, Cortana can predict subsequent actions of a user based on a generalized cognitive model, which uses machine learning and is generated based on a tracking of patterns of interactions between a plurality of users and the system.

Element	Analysis for Cortana
	4.1. Domain, Intent, Slot Modeling Prior work has shown a variety of approaches can be used for language understanding [17, 8, 20, 21, 15]. Some early systems used rule based parsing, hand-authored by an expert for a specific domain [17]. Recent state-of-the-art systems use machine learned models [21, 20, 15]. Semantic meaning of a query is represented as a semantic frame (SF), which is a tuple of (DOMAIN, INTENT, SLOTS) in the semantic space. The slots are a list of key-value pairs: (SLOTNAME, SLOTVALUE). In our system, domain classification is done using machine learned (ML) classifiers such as SVMs [25], which generate a score for each domain. We use a separate binary one-against-all classifier for each domain. For each domain, we also train a domain specific intent and slot model. Intents and slots can be shared across different domains. Intent detection also uses SVM classification but is framed as a multi-class classification problem. Slot tagging is considered as a sequence classification problem. Conditional random fields (CRFs) [26] and more recently deep learning techniques [9, 31, 16] are used for slot tagging. All of the LU models use weighted entity lexicons [18, 19] and are additionally trained with artificial out-of-vocabulary (OOV) words to improve generalization to unseen entities. The slot values can be further resolved into entities (e.g. strongly typed object in some back-end data source) or canonicalized into a standard form (e.g. time/date values). To deal with the open domain nature of queries, each LU domain uses training data from all other domains as negative examples, including representative data for the web domain https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/CortanaLUDialog-FromSLTproceedings.pdf Cortana gathers data from user interactions and machine learning and adapts its responses, providing more accurate and personalized assistance over time. The more you use the app, the better it understands your needs and preferences. https://www.ninjaone.com/blog/what-is-cortana/

Element	Analysis for Cortana
	Fundamentally, Cortana operates on the foundations of machine learning and artificial intelligence. It is capable of understanding user preferences, likes, dislikes, and habits over time. Moreover, it indulges in intelligent conversations with users rather than just referring to a stored set of questions and answers. Cortana achieves seamless human-computer interaction by analyzing the semantics and contexts of the conversations. It primarily performs following tasks: • Keeps track of user behavior, habits, and interaction patterns • Uses search engines and cloud computing features • Analyzes unstructured data • Reads and learns from texts, emails, images, or videos To sum up, Cortana performs well when it knows its users better. https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-cortana/

Copilot

Element	Analysis for Copilot
[42.0] A method of processing natural language utterances, the method being implemented by a computer system that includes one or more processors executing one or more computer program instructions which, when executed, perform the method, the method comprising:	Copilot meets this element. 

Element	Analysis for Copilot
[42.1] receiving a first input of a user that comprises a natural language utterance;	Copilot meets this element. See element [42.0].
[42.2] generating an interpretation of the natural language utterance based on one or more recognized words of the natural language utterance;	Copilot meets this element. For example, Copilot uses automated speech recognition (ASR) technology and natural language processing (NLP) technology to generate an interpretation of natural language utterance. What's changed? Traditionally, when a copilot can't determine a user's intent, it asks the user to rephrase their question. If after two prompts, the copilot still can't determine the user's intent, the copilot escalates to a live agent, using the Escalate system topic. Today, before involving a live agent, the copilot uses natural language processing (NLP) to: • Parse what a user types to determine what they're asking. • Find, collate, and parse relevant information from a specified source. This source can be your company's website, or from multiple sources, including Sharepoint. • Summarize search results into plain language delivered to the copilot user. https://learn.microsoft.com/en-us/microsoft-copilot-studio/nlu-gpt-overview (emphasis added)

Element	Analysis for Copilot
	How does Copilot use NLU? Copilot leverages NLU to enhance user interactions by understanding the context and intent behind user inputs. When a user asks a question or makes a request, Copilot uses NLU to parse the text, identify key entities, and determine the user's intent. This allows Copilot to provide accurate and contextually relevant responses, as in the examples below: 1. Text prompts • Customer service assistant: When a user types, "I'm having trouble with my order," Copilot uses NLU to understand that the user is experiencing an issue and needs assistance, even though the word "help" is not explicitly mentioned. • Virtual assistant: If a user says, "Remind me to call mom tomorrow," Copilot interprets the intent to set a reminder for a specific task and time. • Sentiment analysis: When a user writes a review saying, "The product is amazing, but the delivery was slow," Copilot can identify the positive sentiment towards the product and the negative sentiment towards the delivery. 2. Image prompts • Image captioning: When provided with an image of a dog playing in the park, Copilot can generate a caption like, "A happy dog playing in the park," understanding the context and sentiment of the image. • Object recognition: Given an image of a cluttered desk, Copilot can identify and list objects such as a laptop, coffee mug, and notepad, understanding the context of a workspace. • Scene description: For an image of a beach at sunset, Copilot can describe the scene as, "A serene beach with the sun setting over the horizon," capturing the context and sentiment of the image. https://www.microsoft.com/en-us/bing/do-more-with-ai/what-is-natural-language-understanding?form=MA13KP For example, Copilot uses trigger phrases to interpret user input, and the interpretation includes the series of syllables mapped to the recognized stream of phonemes.

Element	Analysis for Copilot
	Learn / Microsoft Copilot Studio / ⊕ ⋮ Optimizing trigger phrases and natural language understanding Article • 07/02/2024 • 3 contributors 👍 Feedback In this article What are trigger phrases in Copilot Studio The importance of the triggering context On punctuation Creating new trigger phrases Show 3 more What are trigger phrases in Copilot Studio • Trigger phrases train your copilot's natural language understanding (NLU) model. • Trigger phrases are configured at the topic level and indicate to the copilot for what typical user utterances a specific topic should be triggered. • Trigger phrases typically capture the way an end-user would ask about a problem or issue. For example, <i>"problem with weeds in lawn"</i>

Element	Analysis for Copilot
	The importance of the triggering context Copilot Studio NLU behaves differently based on the conversation state, which can sometimes lead to different behaviors for the same user utterance. The following are the different conversation states: • Start of the conversation: the copilot has no context, so a user utterance is expected to either: trigger a topic directly (intent recognition), trigger a "did you mean" (Multiple Topics Matched) disambiguation question among intent candidates if there are multiple matching topics, or go to a fallback topic if the intent isn't recognized. • After a "did you mean" (Multiple Topics Matched) is triggered: NLU optimizes to match one of the suggested topics, with higher thresholds to move out of the presented options. • Switching out from a current topic: If the NLU is trying to fill a slot in a topic, and the user is giving a user query that could trigger another topic (topic switching). https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/trigger-phrases-best-practices

Element	Analysis for Copilot
	Using topic triggers Article • 10/11/2024 • 6 contributors  Feedback In this article - Prerequisites - Trigger types - Changing the trigger for a topic - Trigger condition Show 2 more By default, topics use the Phrases trigger, which activates when one or more of their trigger phrases match closely with an incoming message from a user. There are other trigger types available that determine when a topic should be executed. https://learn.microsoft.com/en-us/microsoft-copilot-studio/authoring-triggers For example, Copilot uses entities and intents to interpret user input, and the interpretation includes the series of syllables mapped to the recognized stream of phonemes.



Element	Analysis for Copilot
	Use entities and slot filling in agents Article • 11/19/2024 • 5 contributorsFeedback In this article <a data-bbox="640 487 819 527">Prebuilt entities<a data-bbox="640 527 819 568">Custom entities<a data-bbox="640 568 966 609">Use entities in a conversation<a data-bbox="640 609 756 665">Slot filling A significant part of agent conversations in Copilot Studio is natural language understanding, which is the ability for the AI to understand a user's intent. For example, the user might say "I tried to use my gift card but it doesn't work" and the agent knows to route the user to the topic related to gift cards not working, even if that exact phrase isn't listed as a trigger phrase. One fundamental aspect of natural language understanding is to identify <i>entities</i> in a user dialog. Watch the video to get started with entities.

Element	Analysis for Copilot
	Slot filling <i>Slot filling</i> is a natural language understanding concept that means saving an extracted entity to an object. However, in Copilot Studio, slot filling means placing the extracted entity value into a variable. We'll continue to use the camping activities topic as an example, which is triggered by typing "I want something to do" in the test chat. The topic is successfully triggered and the agent asks for the type of activity, also showing the button choices specified when authoring the Ask a question node. In the dialog tree, tracking between topics also shows the agent is running to the question node you just edited. https://learn.microsoft.com/en-us/microsoft-copilot-studio/advanced-entities-slot-filling
[11.3] generating a request based on the interpretation;	Copilot meets this element. For example, Copilot uses keywords such as trigger phrases to parse user input and generate a request

Element	Analysis for Copilot
	Learn / Microsoft Copilot Studio / Optimizing trigger phrases and natural language understanding Article • 07/02/2024 • 3 contributors <a data-bbox="1236 396 1320 414" href="#">Feedback In this article - What are trigger phrases in Copilot Studio - The importance of the triggering context - On punctuation - Creating new trigger phrases Show 3 more What are trigger phrases in Copilot Studio • Trigger phrases train your copilot's natural language understanding (NLU) model. • Trigger phrases are configured at the topic level and indicate to the copilot for what typical user utterances a specific topic should be triggered. • Trigger phrases typically capture the way an end-user would ask about a problem or issue. For example, <i>"problem with weeds in lawn"</i>


Element	Analysis for Copilot
	The importance of the triggering context Copilot Studio NLU behaves differently based on the conversation state, which can sometimes lead to different behaviors for the same user utterance. The following are the different conversation states: • Start of the conversation: the copilot has no context, so a user utterance is expected to either: trigger a topic directly (intent recognition), trigger a "did you mean" (Multiple Topics Matched) disambiguation question among intent candidates if there are multiple matching topics, or go to a fallback topic if the intent isn't recognized. • After a "did you mean" (Multiple Topics Matched) is triggered: NLU optimizes to match one of the suggested topics, with higher thresholds to move out of the presented options. • Switching out from a current topic: If the NLU is trying to fill a slot in a topic, and the user is giving a user query that could trigger another topic (topic switching). https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/trigger-phrases-best-practices For example, Copilot uses the orchestrator to generate a request based on the interpretation.  <pre> graph LR subgraph Apps A1[Plugins & declarative copilots] A2[Graph connectors] A3[Custom engine copilots] A4[Microsoft Copilot] A5[Your copilots] end subgraph AI_Orchestration [AI Orchestration] O[AI Orchestration] end subgraph Microsoft_Cloud [Microsoft Cloud] C1[Foundation models] C2[Your data] C3[AI infrastructure] end Apps --> AI_Orchestration AI_Orchestration --> Microsoft_Cloud </pre>

Element	Analysis for Copilot
	Copilot orchestrator https://learn.microsoft.com/en-us/microsoft-365-copilot/extensibility/orchestrator (emphasis added)
[11.4] transmitting the request to a domain agent for processing;	Copilot meets this element. Copilot's orchestrator selects a right skill based on the interpretation of the user input.


Element	Analysis for Copilot
	Copilot orchestrator Between the user's natural language input and Copilot's natural language output, the Copilot orchestrator works behind the scenes to select and run the right skills from the right plugins for the user's task. The orchestration layer represents the interface between foundation Large Language Models (LLMs) and the many ways you can extend, enrich, and customize Copilot for the way your customers work.  The following diagram illustrates how the Microsoft 365 Copilot orchestrator selects the right plugin, with the right skill, at the right time, even when there are multiple options to choose from. 


Element	Analysis for Copilot
	3. Reasoning: The Copilot orchestrator formulates a plan comprised of multiple actions that it performs in an attempt to respond to the user's prompt. a. Context and tool selection: The orchestrator retrieves the user's conversation context from the context store and integrates data from Microsoft Graph to refine the context. It then adjusts the initial query based on this updated context and forwards it to the LLM (large language model) to guide the next steps. The LLM might proceed to generating a response using Copilot's built-in capabilities, or it might determine that additional data is necessary. If more information is needed, the orchestrator does a search for the plugins (tools) with the right skill for the task from the user's enabled plugins based on the descriptions of the plugins and their functions. b. Function matching and parameter determination: The orchestrator formulates a new prompt that incorporates the user's initial query, the updated context, and the selected plugins, and presents it to the LLM. The LLM evaluates the input and specifies the optimal plugin and function within that plugin to address the task. It then provides the orchestrator with the necessary function details and parameters required to gather the needed information.

Element	Analysis for Copilot
	How Copilot's orchestrator matches plugins to user queries When a user submits a query to Copilot, the orchestrator searches its full catalog of skills (<i>functions</i>) from installed plugins to identify up to five skills that best match the query. The orchestrator first tries to match on exact words (lexical match) and expands its search scope as needed to include matches on descriptive meanings (semantic match), working from specific function names to general plugin descriptions, until all five function candidate slots are filled. Specifically, the following list shows the hierarchy of matching mechanisms for Copilot plugin function selection: 1. Lexical match on function name. 2. Semantic match on function description. 3. Lexical match on plugin name (adds all plugin functions to candidate list). 4. Semantic match on plugin name (adds all plugin functions to candidate list). The orchestrator works through this list until all five function candidate slots are filled. To best ensure that the Copilot orchestrator chooses your plugin to provide the right skill for users' queries, apply the Explore best practices for high-quality plugins. https://learn.microsoft.com/en-us/microsoft-365-copilot/extensibility/orchestrator (emphasis added) In addition, Copilot uses Copilot agents to process queries and tasks in a given domain or context. What are agents? In essence, agents are a scoped version of Copilot. They are designed to automate and execute specific business processes on your behalf. Each agent is an expert in one or two specific tasks, such as retrieving information, summarizing data, and even taking actions like sending emails or updating records. Agents can be customized to meet specific business needs, such as: • Human Resources agent • Analytics agent • Project management agent • Legal agent • Image creation agent https://learn.microsoft.com/ko-kr/microsoft-365-copilot/extensibility/


Element	Analysis for Copilot
	New Copilot agents: Supercharge Microsoft 365 Copilot  ClintWilliams  MICROSOFT Oct 31, 2024 Are you ready to supercharge Microsoft 365 Copilot? Following the recent Wave 2 announcements and the general availability of Copilot agents, we are excited to introduce a curated set of new Copilot agents developed by Microsoft. These agents are designed to help you and your team optimize and reinvent your business processes. https://techcommunity.microsoft.com/blog/microsoft365copilotblog/new-copilot-agents-supercharge-microsoft-365-copilot/4266995 Explore agents Agents can be found in the app store in Microsoft Teams and the Microsoft 365 app. From Copilot in Teams or the Microsoft 365 app, select Get Copilot agents on the side panel to the right to explore agents in the app store. A screenshot of the Microsoft 365 Copilot app interface. The top bar shows the Copilot logo. Below it, the word 'Agents' is displayed. A list of agents is shown, including 'Your Copilot' and 'Microsoft Support Assistant'. At the bottom, there are two buttons: 'Get Copilot agents' and 'Create agents'. The 'Get Copilot agents' button is circled in blue. https://support.microsoft.com/en-us/topic/get-started-with-agents-for-microsoft-365-copilot-


Element	Analysis for Copilot
	169469d7-328d-4d37-9090-bfc2058a39bd https://learn.microsoft.com/sr-cyrl-rs/microsoft-365-copilot/extensibility/decision-guide On information and belief, Copilot uses orchestrator use select the right Copilot agents to handle user requests.

Element	Analysis for Copilot
	Agents come in all shapes and sizes. They help you retrieve information from grounding data and reason over it to summarize or answer questions. More capable agents take actions when asked and the most advanced agents are autonomous, operating independently to create and perform plans, orchestrate other agents, and learn when to escalate to an employee for help.  The diagram illustrates six Copilot agents arranged around a central horizontal bar. The bar has a color gradient from red on the left to blue on the right, with a white circle in the middle. The left side of the bar is labeled 'Simple' and the right side is labeled 'Advanced'. The agents are connected to the bar by lines. Each agent has a title, a description of its function, and a play button icon. IT Helpdesk agent: How do I connect to the corporate network? Device Refresh agent: Request a new laptop and send approvals via IT Service tool. Lead Gen agent: The agent has identified and researched 15 new leads for you to review. Project Tracker agent: What is the status of phase 2 for project X and the remaining budget? Budget Management agent: Review outstanding open purchase orders and begin financial planning. Customer Support agent: The agent has identified new support issues and triaged to other agents. https://www.microsoft.com/en-us/microsoft-copilot/blog/copilot-studio/unveiling-copilot-agents-built-with-microsoft-copilot-studio-to-supercharge-your-business/ For example, Copilot uses the orchestrator to send a query or command to the selected skill or agent.




Element	Analysis for Copilot
	Copilot orchestrator https://learn.microsoft.com/en-us/microsoft-365-copilot/extensibility/orchestrator Copilot can also communicate user requests to Copilot agents. For example, Copilot can send a query to one of the prompt-and-response agents.


Element	Analysis for Copilot
	Introducing agents—automate and execute business processes to reduce cost  Agents are AI assistants designed to automate and execute business processes, working with or for humans. They range in capability from simple, prompt-and-response agents to agents that replace repetitive tasks to more advanced, fully autonomous agents. You can use pre-built agents, or create them yourself. Examples include an FAQ manager, a meeting scheduler, or customer service. Learn how agents work Try in BizChat <i>See also, Microsoft 365 Copilot Copilot Studio agent builder, https://www.youtube.com/watch?v=uo-vCFL96yQ</i> For example, Copilot communicates a weather-related request to a weather domain agent to generate a response.

Element	Analysis for Copilot
	 how is the weather in New York? The current weather in New York City is sunny with a temperature of 52°F (11°C) ¹. It's a pleasant day, so you might want to enjoy some outdoor activities if you're in the area! Is there anything else you'd like to know about New York or anything else? ¹ www.timeanddate.com what about the weekend? The weekend weather in New York City looks pretty nice ¹! • Saturday: Expect partly cloudy skies with highs around 66°F (19°C) and lows around 46°F (8°C) ². • Sunday: It will be mostly sunny with highs reaching 57°F (14°C) and lows around 40°F (4°C) ². It's a great time to enjoy some outdoor activities and maybe even catch some fall foliage! Do you have any plans for the weekend? ¹ www.msn.com ² weather.com
[11.5] determining whether a personalized cognitive model associated with the user includes sufficient information for predicting one or more subsequent actions associated with the user, wherein the personalized cognitive model is generated based on a	Copilot meets this element. On information and belief, Copilot uses a personalized cognitive model based on a pattern of interactions between the user and system, and the personalized cognitive model is able to predict user actions.

Element	Analysis for Copilot
tracking of a pattern of interactions between the user and the system, and wherein the one or more subsequent actions include one or more actions predicted to occur after receiving the first input; and	Maintain your identity in meetings optimized for Copilot and meeting recap The most essential input for Copilot in Teams is the identity of each speaker. Copilot needs a meeting transcript, with attribution for every speaker, to deliver meeting summaries, insights, and action items. In a hybrid meeting, without speaker recognition, the video, and audio feed for people in the room would be attributed to the space (for example, Conference Room 1), not the individuals speaking, making it difficult to query individuals' contributions, summarize everyone's perspectives, and tackle those to-do items. Teams Rooms devices use advanced technology called speaker recognition to analyze the distinct vocal characteristics of each speaker, such as pitch, tone, and speaking style, to create a voiceprint for each participant, akin to a fingerprint for their voice. https://learn.microsoft.com/en-us/microsoftteams/rooms/voice-recognition (emphasis added) For example, on information and belief, Copilot parses user input based on information of the user. Intuitive design and helpful responses The refreshed Copilot mobile app is here—faster, simpler, and more personal. https://www.microsoft.com/en-us/microsoft-copilot/personal-ai-assistant We value your privacy Microsoft and our third-party vendors use cookies and similar technologies to deliver, maintain, and improve our services and ads. If you agree, we will use this data for ads personalization and associated analytics. You can select 'Accept' to consent to these uses, 'Reject' to decline these uses or click on 'More options' to review your options. You can change your selection under 'Manage Cookie Preferences' in the settings menu. Privacy Statement Accept Reject More options https://copilot.microsoft.com/onboarding

Element	Analysis for Copilot
	How to access and control your personal data You can also make choices about the collection and use of your data by Microsoft. You can control your personal data that Microsoft has obtained, and exercise your data protection rights, by contacting Microsoft or using various tools we provide. In some cases, your ability to access or control your personal data will be limited, as required or permitted by applicable law. How you can access or control your personal data will also depend on which products you use. For example, you can: • Control the use of your data for personalised advertising from Microsoft, including Xandr, by visiting our opt-out page. • Choose whether you wish to receive promotional emails, SMS messages, telephone calls and post from Microsoft. • Access and clear some of your data through the Microsoft privacy dashboard. Not all personal data processed by Microsoft can be accessed or controlled via the tools above. If you want to access or control personal data processed by Microsoft that is not available via the tools above or directly through the Microsoft products you use, you can always contact Microsoft at the address in the How to contact us section or by using our web form. We provide aggregate metrics about user requests to exercise their data protection rights via the Microsoft Privacy Report. Microsoft Copilot capabilities. Microsoft Copilot is Microsoft's everyday AI companion, and is designed to help you achieve more through a single experience that runs across devices, understanding relevant context on the web, on your PC and across apps to bring you the right skills at the right time. With the help of Copilot, users can start a draft of a new Word document, generate a PowerPoint presentation, quickly find the answers to complex search queries online, find relevant documents or other personal content, or be inspired to create new songs, stories, images or other content, among other tasks. Copilot is a family of services, and Microsoft's collection and use of data may differ depending on the service and the intended functionality in a given scenario.

Element	Analysis for Copilot
	The Copilot website and app (available on iOS and Android) is the core of the consumer Copilot experience. Within this core experience, users can search the web, create text, images, songs, or other outputs, or engage with other features, such as plugins. On the website and in the app, users enter “prompts” that provide instructions to Copilot (e.g. “Give me recommendations for a restaurant that accommodates parties of 10 near me”). In order to provide a relevant response, Copilot will use this prompt, along with the user’s location, language and similar settings, to formulate a helpful response. In some markets, authenticated users can choose to allow Copilot to have access to prior prompt history to better personalise the product. The consumer Copilot product uses the data collected to provide and improve the Copilot services, including to provide relevant advertising. Users who are signed-in to their account can manage their prompt history in product and on the Microsoft Privacy Dashboard, and can adjust their location, language, and other settings in the product https://www.microsoft.com/en-ca/privacy/privacystatement (emphasis added) 

Element	Analysis for Copilot
	https://copilot.microsoft.com/  Dive deeper Find out more about what you can do with Copilot—your research assistant, personal planner, and creative partner, all in one.   Do more with AI See how to use AI for your daily life, work, or school. Find prompting tips and explore ways to get more done with Copilot as your personal AI assistant.  Learn more  https://www.microsoft.com/en-us/microsoft-copilot/personal-ai-assistant (emphasis added) For example, Copilot gives a personalized clothing recommendation based on its interaction with a user.

Element	Analysis for Copilot
	what clothes should i bring to my villa? For a villa getaway, consider packing the following clothing essentials: Warm Weather Villa • Lightweight Clothing: Breathable fabrics like cotton and linen to stay cool. • Swimwear: Don't forget swimsuits for the pool or beach. • Sun Protection: Hats, sunglasses, and a light cover-up for sun protection. • Evening Wear: Casual, yet slightly dressy outfits for dining out. • Comfortable Shoes: Sandals for the day and casual shoes for the evening. Cooler Weather Villa • Layers: Light sweaters, jackets, and long sleeves for fluctuating temperatures. • Warm Accessories: Hats, scarves, and gloves if it gets chilly. • Comfortable Shoes: Walking shoes for exploring and cozy slippers for indoors. Other Essentials • Activewear: If you plan on exercising or doing outdoor activities. • Casual Clothes: Comfortable outfits for lounging around the villa. • Personal Care Items: Don't forget toiletries, sunscreen, and any medications. What activities are you planning during your villa stay? I can offer more tailored recommendations based on that!

Element	Analysis for Copilot
	 The screenshot shows a Copilot chat window. At the top, a user input bubble contains the text "my villa is near whistler blackcomb". Below this, the Copilot response is displayed in a dark-themed box. The response begins with "Thanks for letting me know! The current weather near Whistler Blackcomb is mostly cloudy with temperatures around 1°C (34°F) during the day and freezing temperatures at night ¹. There's a 30% chance of flurries or rain showers today, so it's a good idea to bundle up and be prepared for some wintry conditions ¹." This is followed by the question "Do you have any plans for your time at the villa?". At the bottom of the response box, a footnote ¹ is provided: "www.whistlerblackcomb.com".
[11.6] predicting the one or more subsequent actions based on a generalized cognitive model in response to a determination that the personalized cognitive model does not include the sufficient information, wherein the generalized cognitive model is generated based on a tracking of patterns of interactions between a plurality of users and the system	Copilot meets this element. On information and belief, Copilot can predict subsequent actions of a user based on a generalized cognitive model, which uses machine learning and is generated based on a tracking of patterns of interactions between a plurality of users and the system. AI systems like Bing and Microsoft Copilot (web) are as good as they are because they continuously learn and improve from people's interactions. Since the early 2000s, user clicks on search result pages have fueled the continuous improvements of search engines. Recently, reinforcement learning from human feedback (RLHF) brought step-function improvements to response quality of generative AI models. Bing has a rich history of success in improving its AI offerings by learning from user interactions. For example, Bing pioneered the idea of improving search ranking² and personalizing search using short- and long-term user behavior data². https://www.microsoft.com/en-us/research/blog/learning-from-interaction-with-microsoft-copilot-web/(emphasis added)

Element	Analysis for Copilot
	How does Copilot use AI machine learning? Copilot harnesses the power of AI machine learning to elevate the user experience, enabling the following capabilities: Contextual assistance Copilot employs machine learning to offer contextual assistance that adapts to your needs. Whether you're drafting an email, writing poetry, planning a trip, or researching a topic, Copilot can help suggest relevant information. From brainstorming gift ideas to learning a new skill, use Copilot can help you gather information and resources quickly. For enhanced performance during peak usage, you can upgrade to Copilot Pro so you have assistance when it's most important. Text and images to order Copilot can produce text and images based on your text input, also known as your prompt. Simply enter a descriptive prompt asking Copilot to generate text or an image related to any subject or style. Using machine learning and natural language processing, Copilot will then produce writing or visuals based on your request. Because Copilot is conversational, you can keep asking for tweaks until you receive the output you're after. Predictive text enhancement Machine learning within Copilot enhances its predictive text capabilities, enabling it to anticipate your writing style and intent. Whether you're composing an email or a document, Copilot can help suggest word choices and sentence structures that align with your unique voice, saving you time and ensuring your content is tailored to you. Thanks to machine learning, the more you use Copilot, the better its results. That's one way it's always evolving to enhance efficiency and productivity for users. Try Copilot and the Copilot mobile app today for AI assistance anytime, anywhere. https://www.microsoft.com/en-us/bing/do-more-with-ai/what-is-machine-learning?form=MA13KP

Element	Analysis for Copilot
	<u>The Copilot System: Enterprise-ready AI</u> Microsoft is uniquely positioned to deliver enterprise-ready AI with the Copilot System. Copilot is more than OpenAI's ChatGPT embedded into Microsoft 365. It's a sophisticated processing and orchestration engine working behind the scenes to combine the power of LLMs, including GPT-4, with the Microsoft 365 apps and your business data in the Microsoft Graph — now accessible to everyone through natural language. Grounded in <i>your</i> business data. AI-powered LLMs are trained on a large but limited corpus of data. The key to unlocking productivity in business lies in connecting LLMs to your business data — in a secure, compliant, privacy-preserving way. Microsoft 365 Copilot has real-time access to both your <i>content and context</i> in the Microsoft Graph. This means it generates answers anchored in your business <i>content</i> — your documents, emails, calendar, chats, meetings, contacts and other business data — and combines them with your working <i>context</i> — the meeting you're in now, the email exchanges you've had on a topic, the chat conversations you had last week — to deliver accurate, relevant, contextual responses. https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/ (emphasis added)


Element	Analysis for Copilot
	Data stored about user interactions with Microsoft 365 Copilot When a user interacts with Microsoft 365 Copilot (using apps such as Word, PowerPoint, Excel, OneNote, Loop, or Whiteboard), we store data about these interactions. The stored data includes the user's prompt and Copilot's response, including citations to any information used to ground Copilot's response. We refer to the user's prompt and Copilot's response to that prompt as the "content of interactions" and the record of those interactions is the user's Copilot activity history. For example, this stored data provides users with Copilot activity history in Business Chat and meetings in Microsoft Teams. This data is processed and stored in alignment with contractual commitments with your organization's other content in Microsoft 365. The data is encrypted while it's stored and isn't used to train foundation LLMs, including those used by Microsoft 365 Copilot. To view and manage this stored data, admins can use Content search or Microsoft Purview. Admins can also use Microsoft Purview to set retention policies for the data related to chat interactions with Copilot. For more information, see the following articles: https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy(emphasis added)

Element	Analysis for Copilot
	• We will soon start using consumer data from Copilot, Bing, and Microsoft Start (MSN) (including interactions with advertisements) to help train the generative AI models in Copilot. We are making this change as real-world consumer interactions provide greater breadth and diversity in training data. We believe this will help us build more inclusive, relevant products, and improve the experience for all users. For example, our AI models can learn from an aggregated set of thumbs up/thumbs down selections to provide better, more useful responses in the future. Our AI models can learn colloquial phrases or local references from Copilot conversations. Our AI models can also learn from Microsoft advertising segments what, broadly, is most interesting and compelling to consumers, helping the models deliver more relevant content in the future. https://www.microsoft.com/en-us/microsoft-copilot/blog/2024/08/16/transparency-and-control-in-consumer-data-use/ (emphasis added) Learning from interaction with Microsoft Copilot (web) Published March 27, 2024 By Scott Counts, Senior Principal Research Manager; Jennifer Neville, Partner Research Manager; Mengting Wan, Principal Applied Research Scientist; Ryen W. White, General Manager and Deputy Lab Director; Longqi Yang, Principal Applied Research Manager Share this page      With the introduction of Microsoft Copilot (web), the way that people interact with AI systems has fundamentally changed from searching to conversing and from simple actions to complex workflows. Today, we are excited to share three technical reports on how we are <i>starting</i> to leverage new types of user interactions to understand and improve Copilot (web) for our consumer customers. ^[1]

Element	Analysis for Copilot
	https://www.microsoft.com/en-us/research/blog/learning-from-interaction-with-microsoft-copilot-web/ For example, when Copilot has not learned about the location of the user’s villa, Copilot gives generic clothing recommendations for a trip to the villa.

Element	Analysis for Copilot
	what clothes should i bring to my villa? For a villa getaway, consider packing the following clothing essentials: Warm Weather Villa • Lightweight Clothing: Breathable fabrics like cotton and linen to stay cool. • Swimwear: Don't forget swimsuits for the pool or beach. • Sun Protection: Hats, sunglasses, and a light cover-up for sun protection. • Evening Wear: Casual, yet slightly dressy outfits for dining out. • Comfortable Shoes: Sandals for the day and casual shoes for the evening. Cooler Weather Villa • Layers: Light sweaters, jackets, and long sleeves for fluctuating temperatures. • Warm Accessories: Hats, scarves, and gloves if it gets chilly. • Comfortable Shoes: Walking shoes for exploring and cozy slippers for indoors. Other Essentials • Activewear: If you plan on exercising or doing outdoor activities. • Casual Clothes: Comfortable outfits for lounging around the villa. • Personal Care Items: Don't forget toiletries, sunscreen, and any medications. What activities are you planning during your villa stay? I can offer more tailored recommendations based on that!

Azure AI and Azure OpenAI


Element	Analysis for Azure AI and Azure OpenAI
[42.0] A method of processing natural language utterances, the method being implemented by a computer system that includes one or more processors executing one or more computer program instructions which, when executed, perform the method, the method comprising:	Azure AI and Azure OpenAI meet this element.  https://learn.microsoft.com/en-us/azure/ai-services/speech-service/tutorial-voice-enable-your-bot-speech-sdk On information and belief, Azure AI and Azure OpenAI uses Language Understanding (LUIS) and Conversational Language Understanding technology to understand and respond to user natural language inputs.

Element	Analysis for Azure AI and Azure OpenAI
	Combine LUIS and question answering capabilities Article • 08/28/2024 • 1 contributor Feedback In this article - When to use each feature - Create an orchestration project - Set up orchestration between Azure AI services features - Change a LUIS resource to a language resource - Next steps https://learn.microsoft.com/en-us/azure/ai-services/luis/how-to/orchestration-projects What is conversational language understanding? Article • 08/28/2024 • 4 contributors Feedback In this article - Example usage scenarios - Project development lifecycle - Reference documentation and code samples - Responsible AI - Next steps Conversational language understanding is one of the custom features offered by Azure AI Language. It is a cloud-based API service that applies machine-learning intelligence to enable you to build natural language understanding component to be used in an end-to-end conversational application. https://learn.microsoft.com/en-us/azure/ai-services/language-service/conversational-language-understanding/overview


Element	Analysis for Azure AI and Azure OpenAI
[42.1] receiving a first input of a user that comprises a natural language utterance;	Azure AI and Azure OpenAI meet this element. See element [42.0].
[42.2] generating an interpretation of the natural language utterance based on one or more recognized words of the natural language utterance;	Azure AI and Azure OpenAI meet this element. For example, Azure AI uses automated speech recognition (ASR) technology and natural language processing (NLP) technology to generate an interpretation of natural language utterance. Out of the box, speech recognition utilizes a Universal Language Model as a base model that is trained with Microsoft-owned data and reflects commonly used spoken language. The base model is pretrained with dialects and phonetics representing various common domains. When you make a speech recognition request, the most recent base model for each supported language is used by default. The base model works well in most speech recognition scenarios. Custom speech allows you to tailor the speech recognition model to better suit your application's specific needs. This can be particularly useful for: • Improving recognition of domain-specific vocabulary: Train the model with text data relevant to your field. • Enhancing accuracy for specific audio conditions: Use audio data with reference transcriptions to refine the model. https://learn.microsoft.com/en-au/azure/ai-services/speech-service/speech-to-text For example, Azure AI uses pattern matching to interpret user utterances.

Element	Analysis for Azure AI and Azure OpenAI
	Pattern matching The Speech SDK provides an embedded pattern matcher that you can use to recognize intents in a strict way. This is useful for when you need a quick offline solution. This works especially well when the user is going to be trained in some way or can be expected to use specific phrases to trigger intents. For example: "Go to floor seven", or "Turn on the lamp" etc. It's recommended to start here and if it no longer meets your needs, switch to using CLU or a combination of the two. Use pattern matching if: • You're only interested in matching strictly what the user said. These patterns match more aggressively than conversational language understanding (CLU). • You don't have access to a CLU model, but still want intents. https://learn.microsoft.com/en-us/azure/ai-services/speech-service/intent-recognition

Element	Analysis for Azure AI and Azure OpenAI
	Patterns vs. Exact Phrases There are two types of strings used in the pattern matcher: "exact phrases" and "patterns". It's important to understand the differences. Exact phrases are strings of the exact words that you want to match. For example: "Take me to floor seven". A pattern is a phrase that contains a marked entity. Entities are marked with "{}" to define the place inside the pattern and the text inside the "{}" references the entity ID. Given the previous example perhaps you would want to extract the floor name in an entity named "floorName". You would do so with a pattern like this: "Take me to floor {floorName}" <a data-bbox="598 912 1793 941" href="https://learn.microsoft.com/en-us/azure/ai-services/speech-service/pattern-matching-overview">https://learn.microsoft.com/en-us/azure/ai-services/speech-service/pattern-matching-overview For example, Azure AI uses keywords to interpret user utterances.

Element	Analysis for Azure AI and Azure OpenAI
	Keyword verification and speech to text When keyword verification is used, it's always in combination with speech to text. Both services run in parallel, which means audio is sent to both services for simultaneous processing.  The diagram illustrates the data flow for keyword verification and speech-to-text processing. It starts with an 'AUDIO' input on the left, represented by a horizontal arrow. This arrow points into a dashed blue box labeled 'EDGE'. Inside the 'EDGE' box is a solid blue rectangle labeled 'On-device keyword recognition'. An arrow from the right side of the 'EDGE' box splits into two arrows, each pointing to a separate box within a larger dashed blue box labeled 'CLOUD'. The top box in the 'CLOUD' is labeled 'Keyword Verification', and the bottom box is labeled 'Speech-to-Text (incl. Custom Speech)'. https://learn.microsoft.com/en-us/azure/ai-services/speech-service/keyword-recognition-overview
[42.3] generating a request based on the interpretation;	Azure AI and Azure OpenAI meet this element. For example, Azure AI uses orchestrator to generate a request based on the interpretation.

Element	Analysis for Azure AI and Azure OpenAI
	Learn / Azure / AI Services / AI Bot Service / Use multiple LUIS and QnA models with Orchestrator If a bot uses multiple Language Understanding (LUIS) models and QnA Maker knowledge bases, you can use Bot Framework Orchestrator to determine which LUIS model or QnA Maker knowledge base best matches the user input. You can use the <code>bf orchestrator</code> CLI command to create an Orchestrator snapshot file, then use the snapshot file to route user input to the correct model at run time. This article describes how to use an <i>existing</i> QnA Maker knowledge base with Orchestrator. • For new bots, consider using the question answering and orchestration workflow features of Azure AI Language. • For more information about Orchestrator, see Intent recognition with Orchestrator in Composer. • For more information about the <code>bf orchestrator</code> command, see the Bot Framework CLI README. https://learn.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-orchestrator?view=azure-bot-service-4.0&tabs=cs

Element	Analysis for Azure AI and Azure OpenAI
	About this sample This sample is based on a predefined set of LUIS and QnA Maker projects. However, to use QnA Maker in your bot, you need an existing knowledge base in the QnA Maker portal. Your bot then can use the knowledge base to answer the user's questions. For new bot development, consider using Copilot Studio. If you need to create a new knowledge base for a Bot Framework SDK bot, see the following Azure AI services articles: • What is question answering? • Create an FAQ bot • Azure Cognitive Language Services Question Answering client library for .NET C# JavaScript  <pre> graph TD OnMessageActivityAsync[OnMessageActivityAsync] --> DispatchToTopIntentAsync[DispatchToTopIntentAsync] DispatchToTopIntentAsync --> ProcessSampleQnAAsync[ProcessSampleQnAAsync] DispatchToTopIntentAsync --> ProcessWeatherAsync[ProcessWeatherAsync] DispatchToTopIntentAsync --> ProcessHomeAutomationAsync[ProcessHomeAutomationAsync] </pre> <code>OnMessageActivityAsync</code> is called for each user input received. This module finds the top scoring user intent and passes that result on to <code>DispatchToTopIntentAsync</code>. <code>DispatchToTopIntentAsync</code>, in turn, calls the appropriate app handler. • <code>ProcessSampleQnAAsync</code> - for bot FAQ questions. • <code>ProcessWeatherAsync</code> - for weather queries. • <code>ProcessHomeAutomationAsync</code> - for home lighting commands. The handler calls the LUIS or QnA Maker service and returns the generated result back to the user.


Element	Analysis for Azure AI and Azure OpenAI
	Call the services from your bot For each input from your user, the bot logic passes in user input to Orchestrator Recognizer, finds the top returned intent, and uses that information to call the appropriate service for the input. C# JavaScript In the DispatchBot.cs file whenever the <code>OnMessageActivityAsync</code> method is called, we check the incoming user message and get the top intent from Orchestrator Recognizer. We then pass the <code>topIntent</code> and <code>recognizerResult</code> on to the correct method to call the service and return the result. bots\DispatchBot.cs C# (archived) https://learn.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-orchestrator?view=azure-bot-service-4.0&tabs=cs For example, Azure AI communicates a weather-related request to a weather domain agent to generate a response.

Element	Analysis for Azure AI and Azure OpenAI
	Step 3: Integrating the Weather API ↖ 1. Add Weather API Integration: 1. In your bot's main dialog (e.g., WeatherDialog.cs), create a method to call the weather API. 2. Integrate CLU 1. Modify the JSON file that defines intents, entities, and utterances of the weather. 3. Implement the Weather API Call: 1. Use RestService library to send a request to the weather API and parse the response. 4. Handle Weather Response: 1. Create a modal class WeatherData.cs to parse the resulting API into the object.

Element	Analysis for Azure AI and Azure OpenAI
	• Implement Weather API Call <pre> public const string UrlTimeLine = https://weather.visualcrossing.com/VisualCrossingWebServices/rest/services/t imeline"; public static string WeatherApiKey = ""; private string GetWeather(string location, string dateTime) { if (string.IsNullOrEmpty(location) string.IsNullOrEmpty(dateTime)) { return "Sorry, I didn't get that. Please try asking in a different way"; } var restService = new RestService(); var url = \$"{UrlTimeLine}/{location}/{dateTime:yyyy-MM-ddTHH:mm:ss}? key={WeatherApiKey}"; var results = restService.GetWeatherData(url); if (results != null) { var description = results.Days[0].Description; var temperature = results.Days[0].Temp; return \$"{description}. The temperature of {temperature}°F"; } return "Sorry, I didn't get that. Please try asking in a different way"; } </pre> https://www.visualcrossing.com/resources/documentation/weather-data-tutorials/how-to-build-an-artificially-intelligent-chatbot-using-the-microsoft-bot-framework-and-the-visual-crossing-weather-api/ (emphasis added)
[42.4] transmitting the request to a domain agent for	Azure AI and Azure OpenAI meet this element. For example, Azure AI uses orchestrator to communicate a request to an Azure AI bot or Azure AI

Element	Analysis for Azure AI and Azure OpenAI
processing;	skill. Learn / Azure / AI Services / AI Bot Service / Use multiple LUIS and QnA models with Orchestrator If a bot uses multiple Language Understanding (LUIS) models and QnA Maker knowledge bases, you can use Bot Framework Orchestrator to determine which LUIS model or QnA Maker knowledge base best matches the user input. You can use the <code>bf orchestrator</code> CLI command to create an Orchestrator snapshot file, then use the snapshot file to route user input to the correct model at run time. This article describes how to use an <i>existing</i> QnA Maker knowledge base with Orchestrator. • For new bots, consider using the question answering and orchestration workflow features of Azure AI Language. • For more information about Orchestrator, see Intent recognition with Orchestrator in Composer. • For more information about the <code>bf orchestrator</code> command, see the Bot Framework CLI README. https://learn.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-orchestrator?view=azure-bot-

Element	Analysis for Azure AI and Azure OpenAI
	<u>service-4.0&tabs=cs</u>

Element	Analysis for Azure AI and Azure OpenAI
	About this sample This sample is based on a predefined set of LUIS and QnA Maker projects. However, to use QnA Maker in your bot, you need an existing knowledge base in the QnA Maker portal. Your bot then can use the knowledge base to answer the user's questions. For new bot development, consider using Copilot Studio. If you need to create a new knowledge base for a Bot Framework SDK bot, see the following Azure AI services articles: • What is question answering? • Create an FAQ bot • Azure Cognitive Language Services Question Answering client library for .NET C# JavaScript  <pre> graph TD OnMessageActivityAsync[OnMessageActivityAsync] --> DispatchToTopIntentAsync[DispatchToTopIntentAsync] DispatchToTopIntentAsync --> ProcessSampleQnAAsync[ProcessSampleQnAAsync] DispatchToTopIntentAsync --> ProcessWeatherAsync[ProcessWeatherAsync] DispatchToTopIntentAsync --> ProcessHomeAutomationAsync[ProcessHomeAutomationAsync] </pre> <code>OnMessageActivityAsync</code> is called for each user input received. This module finds the top scoring user intent and passes that result on to <code>DispatchToTopIntentAsync</code>. <code>DispatchToTopIntentAsync</code>, in turn, calls the appropriate app handler. • <code>ProcessSampleQnAAsync</code> - for bot FAQ questions. • <code>ProcessWeatherAsync</code> - for weather queries. • <code>ProcessHomeAutomationAsync</code> - for home lighting commands. The handler calls the LUIS or QnA Maker service and returns the generated result back to the user.

Element	Analysis for Azure AI and Azure OpenAI
	Call the services from your bot For each input from your user, the bot logic passes in user input to Orchestrator Recognizer, finds the top returned intent, and uses that information to call the appropriate service for the input. C# JavaScript In the DispatchBot.cs file whenever the <code>OnMessageActivityAsync</code> method is called, we check the incoming user message and get the top intent from Orchestrator Recognizer. We then pass the <code>topIntent</code> and <code>recognizerResult</code> on to the correct method to call the service and return the result. bots\DispatchBot.cs C# (archived) https://learn.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-orchestrator?view=azure-bot-service-4.0&tabs=cs For example, Azure AI communicates a weather-related request to a weather domain agent to generate a response.

Element	Analysis for Azure AI and Azure OpenAI
	Step 3: Integrating the Weather API ⏏ 1. Add Weather API Integration: 1. In your bot's main dialog (e.g., WeatherDialog.cs), create a method to call the weather API. 2. Integrate CLU 1. Modify the JSON file that defines intents, entities, and utterances of the weather. 3. Implement the Weather API Call: 1. Use RestService library to send a request to the weather API and parse the response. 4. Handle Weather Response: 1. Create a modal class WeatherData.cs to parse the resulting API into the object.

Element	Analysis for Azure AI and Azure OpenAI
	• Implement Weather API Call <pre> public const string UrlTimeLine = https://weather.visualcrossing.com/VisualCrossingWebServices/rest/services/t imeline"; public static string WeatherApiKey = ""; private string GetWeather(string location, string dateTime) { if (string.IsNullOrEmpty(location) string.IsNullOrEmpty(dateTime)) { return "Sorry, I didn't get that. Please try asking in a different way"; } var restService = new RestService(); var url = \$"{UrlTimeLine}/{location}/{dateTime:yyyy-MM-ddTHH:mm:ss}? key={WeatherApiKey}"; var results = restService.GetWeatherData(url); if (results != null) { var description = results.Days[0].Description; var temperature = results.Days[0].Temp; return \$"{description}. The temperature of {temperature}°F"; } return "Sorry, I didn't get that. Please try asking in a different way"; } </pre> https://www.visualcrossing.com/resources/documentation/weather-data-tutorials/how-to-build-an-artificially-intelligent-chatbot-using-the-microsoft-bot-framework-and-the-visual-crossing-weather-api/ (emphasis added).
[42.5] determining whether a personalized cognitive model associated with the user	Azure AI and Azure OpenAI meet this element. For example, on information and belief, Azure AI uses a personalized cognitive model based on a

Element	Analysis for Azure AI and Azure OpenAI
includes sufficient information for predicting one or more subsequent actions associated with the user, wherein the personalized cognitive model is generated based on a tracking of a pattern of interactions between the user and the system, and wherein the one or more subsequent actions include one or more actions predicted to occur after receiving the first input; and	pattern of interactions between the user and system, and the personalized cognitive model is able to predict user actions. For example, Azure OpenAI enables access to chat history to streamline customer support and for data analytics.

Element	Analysis for Azure AI and Azure OpenAI
	Benefits of enabling chat history With Azure OpenAI Studio, Developers can build a chatbot with cutting-edge models that draws on your own data for informed and custom responses to customers' questions. In addition, you can incorporate multimodality - enabling your app to see, hear, and speak by pairing Azure OpenAI Service with Speech and Vision models. Streamline customer support: Chat history serves as a powerful ally for streamlining customer support services. By referencing past chat logs, support teams gain the ability to quickly find solutions for customers. This enhances the efficiency of issue resolution while enabling support agents to manage request volumes effectively leading to improved customer satisfaction. Data Analytics: Analyzing past interactions provides valuable insights into user behavior, preferences, and recurring issues. Armed with this data, you can make informed decisions to optimize user experiences, tailor content, and refine your application's performance. The analytics derived from chat history pave the way for data-driven strategies, ensuring your application evolves in tune with user needs and expectations. https://techcommunity.microsoft.com/blog/azure-ai-services-blog/enable-chat-history-on-azure-openai-studio-with-azure-cosmos-db/4051636

Element	Analysis for Azure AI and Azure OpenAI
	For example, Azure AI understands and processes user input by tracking the contexts for interactions with the user. DialogContext class Reference Feedback Package: botbuilder-dialogs The context for the current dialog turn with respect to a specific DialogSet. Remarks This includes the turn context, information about the dialog set, and the state of the dialog stack. From code outside of a dialog in the set, use DialogSet.createContext to create the dialog context. Then use the methods of the dialog context to manage the progression of dialogs in the set. When you implement a dialog, the dialog context is a parameter available to the various methods you override or implement. https://learn.microsoft.com/en-us/javascript/api/botbuilder-dialogs/dialogcontext?view=botbuilder-ts-latest On information and belief, Azure AI also uses user profile data to understand and process user input.

Element	Analysis for Azure AI and Azure OpenAI
	User profile attributes Article • 01/11/2024 • 15 contributors Feedback In this article - Microsoft Entra user resource type - Required attributes - Display name attribute - Identities attribute - Show 5 more Your Azure Active Directory B2C (Azure AD B2C) directory user profile comes with a set of built-in attributes, such as given name, surname, city, postal code, and phone number. You can extend the user profile with your own application data without requiring an external data store. Microsoft Graph API supports most of the attributes that you can use with Azure AD B2C. This article describes user profile attributes that Azure AD B2C supports. It also notes those attributes that Microsoft Graph doesn't support, and Microsoft Graph API attributes that Azure AD B2C shouldn't use. https://learn.microsoft.com/en-us/azure/active-directory-b2c/user-profile-attributes

Element	Analysis for Azure AI and Azure OpenAI
	Azure AI Personalizer is an AI service that your applications make smarter decisions at scale using reinforcement learning. Personalizer processes information about the state of your application, scenario, and/or users (<i>contexts</i>), and a set of possible decisions and related attributes (<i>actions</i>) to determine the best decision to make. Feedback from your application (<i>rewards</i>) is sent to Personalizer to learn how to improve its decision-making ability in near-real time. Personalizer can determine the best actions to take in a variety of scenarios: • E-commerce: What product should be shown to customers to maximize the likelihood of a purchase? • Content recommendation: What article should be shown to increase the click-through rate? • Content design: Where should an advertisement be placed to optimize user engagement on a website? • Communication: When and how should a notification be sent to maximize the chance of a response? To get started with the Personalizer, follow the quickstart guide, or try Personalizer in your browser with this interactive demo [↗]. https://learn.microsoft.com/en-us/azure/ai-services/personalizer/what-is-personalizer On information and belief, Azure AI uses a personalized cognitive model in a similar way as Copilot. For example, Copilot gives a personalized clothing recommendation based on its interaction with a user.

Element	Analysis for Azure AI and Azure OpenAI
	what clothes should i bring to my villa? For a villa getaway, consider packing the following clothing essentials: Warm Weather Villa • Lightweight Clothing: Breathable fabrics like cotton and linen to stay cool. • Swimwear: Don't forget swimsuits for the pool or beach. • Sun Protection: Hats, sunglasses, and a light cover-up for sun protection. • Evening Wear: Casual, yet slightly dressy outfits for dining out. • Comfortable Shoes: Sandals for the day and casual shoes for the evening. Cooler Weather Villa • Layers: Light sweaters, jackets, and long sleeves for fluctuating temperatures. • Warm Accessories: Hats, scarves, and gloves if it gets chilly. • Comfortable Shoes: Walking shoes for exploring and cozy slippers for indoors. Other Essentials • Activewear: If you plan on exercising or doing outdoor activities. • Casual Clothes: Comfortable outfits for lounging around the villa. • Personal Care Items: Don't forget toiletries, sunscreen, and any medications. What activities are you planning during your villa stay? I can offer more tailored recommendations based on that!

Element	Analysis for Azure AI and Azure OpenAI
	 The screenshot shows a chat interface with a dark background. At the top, a user input in a light blue bubble reads: "my villa is near whistler blackcomb". Below this, the AI response in a dark blue bubble reads: "Thanks for letting me know! The current weather near Whistler Blackcomb is mostly cloudy with temperatures around 1°C (34°F) during the day and freezing temperatures at night ¹. There's a 30% chance of flurries or rain showers today, so it's a good idea to bundle up and be prepared for some wintry conditions ¹. Do you have any plans for your time at the villa?". At the bottom, a footnote in a light blue bubble reads: "¹ www.whistlerblackcomb.com".
[42.6] predicting the one or more subsequent actions based on a generalized cognitive model in response to a determination that the personalized cognitive model does not include the sufficient information, wherein the generalized cognitive model is generated based on a tracking of patterns of interactions between a plurality of users and the system	Azure AI and Azure OpenAI meet this element. For example, on information and belief, Azure AI can predict subsequent actions of a user based on a generalized cognitive model, which uses machine learning and is generated based on a tracking of patterns of interactions between a plurality of users and the system.

Element	Analysis for Azure AI and Azure OpenAI
	What data does the Azure OpenAI Service process? Azure OpenAI processes the following types of data: • Prompts and generated content. Prompts are submitted by the user, and content is generated by the service, via the completions, chat completions, images, and embeddings operations. • Uploaded data. You can provide your own data for use with certain service features (e.g., fine-tuning, assistants API, batch processing) using the Files API or vector store. • Data for stateful entities. When you use certain optional features of Azure OpenAI service, such as the Threads feature of the Assistants API, the service will create a data store to persist message history and other content, in accordance with how you configure the feature. • Augmented data included with or via prompts. When using data associated with stateful entities, the service retrieves relevant data from a configured data store and augments the prompt to produce generations that are grounded with your data. Prompts may also be augmented with data retrieved from a source included in the prompt itself, such as a URL. • Training & validation data. You can provide your own training data consisting of prompt-completion pairs for the purposes of fine-tuning an OpenAI model. https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?WT.mc_id=academic-97358-cacaste&context=%2Fazure%2Fcognitive-services%2Fopenai%2Fcontext%2Fcontext&tabs=azure-portal Data Analytics: Analyzing past interactions provides valuable insights into user behavior, preferences, and recurring issues. Armed with this data, you can make informed decisions to optimize user experiences, tailor content, and refine your application's performance. The analytics derived from chat history pave the way for data-driven strategies, ensuring your application evolves in tune with user needs and expectations. https://techcommunity.microsoft.com/blog/azure-ai-services-blog/enable-chat-history-on-azure-openai-studio-with-azure-cosmos-db/4051636

Element	Analysis for Azure AI and Azure OpenAI
	In trying to make a video recommendation, for instance, what someone prefers to watch might be driven by what time of day it is, whether they're sitting at home or moving around, or how much battery their device has left. Personalizer learns from the choices or actions that customers with similar characteristics have made. Agents that use reinforcement learning have the potential to better anticipate behaviors and react to nuances to enable effective collaboration with human players who are creative and unpredictable and have different styles of play, said Katja Hofmann, a principal researcher who leads a team that focuses on deep reinforcement learning in gaming and other application areas at Microsoft Research's Cambridge-UK lab. Bots developed with current technologies struggle to navigate those complexities and simply don't react in quite the same way that people do. https://news.microsoft.com/source/features/ai/reinforcement-learning/

Element	Analysis for Azure AI and Azure OpenAI
	What is reinforcement learning? In reinforcement learning, the goal is to train an agent <i>policy</i> that outputs actions based on the agent's observations of its environment. Actions result in further observations and <i>rewards</i> for taking the actions. In reinforcement learning, the full reward for policy actions may take many steps to obtain. Learning a policy involves many trial-and-error runs of the agent interacting with the environment and improving its policy. https://techcommunity.microsoft.com/blog/machinelearningblog/introducing-reinforcement-learning-on-azure-machine-learning/1403028 On information and belief, Azure AI and Azure OpenAI use a personalized cognitive model and a general cognitive model in a similar way as Copilot. For example, when Copilot has not learned about the location of the user's villa, Copilot gives generic clothing recommendations for a trip to the villa.

Element	Analysis for Azure AI and Azure OpenAI
	what clothes should i bring to my villa? For a villa getaway, consider packing the following clothing essentials: Warm Weather Villa • Lightweight Clothing: Breathable fabrics like cotton and linen to stay cool. • Swimwear: Don't forget swimsuits for the pool or beach. • Sun Protection: Hats, sunglasses, and a light cover-up for sun protection. • Evening Wear: Casual, yet slightly dressy outfits for dining out. • Comfortable Shoes: Sandals for the day and casual shoes for the evening. Cooler Weather Villa • Layers: Light sweaters, jackets, and long sleeves for fluctuating temperatures. • Warm Accessories: Hats, scarves, and gloves if it gets chilly. • Comfortable Shoes: Walking shoes for exploring and cozy slippers for indoors. Other Essentials • Activewear: If you plan on exercising or doing outdoor activities. • Casual Clothes: Comfortable outfits for lounging around the villa. • Personal Care Items: Don't forget toiletries, sunscreen, and any medications. What activities are you planning during your villa stay? I can offer more tailored recommendations based on that!