

Exhibit E-Kennewick

Invalidity Chart for U.S. Patent No. 9,495,957 in view of U.S. Patent App. Pub. No. 2004/0044516 (“Kennewick”)


U.S. Patent App. Pub. No. 2004/0044516 (“Kennewick”) was published Mar. 4, 2004. Kennewick qualifies as prior art to U.S. Patent No. 9,495,957 (the “’957 Patent”) at least under one or more applicable provisions of 35 U.S.C. §§ 102(a), (b), (e), and/or (g). Kennewick discloses and/or renders obvious one or more of claims 7 and 8 of the ’957 Patent alone and/or in combination with other references, as set forth in the chart below. To the extent Kennewick does not disclose one or more limitations of the claims, such limitations would have been inherent and/or it would have been obvious to combine the disclosures of Kennewick with the knowledge of one of ordinary skill in the art and/or with one or more other references, including as identified herein and/or in the Cover Pleading, for the reasons identified herein and/or in the Cover Pleading, to render the claims at-issue in the ’957 Patent invalid.


Defendant Bank of America, N.A. (“Bank”) incorporates in this chart all applicable qualifications, clarifications, and other statements made in Bank’s Invalidity Contentions. This invalidity claim chart is based on Bank’s present understanding of claims 7, 8, and Plaintiff Dialect, LLC (“Dialect”) apparent construction of the claims, as set forth in Dialect’s Infringement Contentions. These invalidity assertions are not an admission that Bank concedes or acquiesces to any claim construction(s) implied or suggested by Dialect the associated infringement claim charts. Nor is Bank asserting any claim construction positions through these charts, including whether the preamble is a limitation. Bank also does not concede or acquiesce that any asserted claim satisfies the requirements of 35 U.S.C. §§ 112 or 101, and submits these invalidity contentions only to the extent Dialect’s assertions may be understood.

The excerpts cited herein are exemplary. For any claim limitation, Bank may rely on excerpts cited for any other limitation and/or additional excerpts not set forth fully herein to the extent necessary to provide a more comprehensive explanation for a reference’s disclosure of a limitation. Where an excerpt refers to or discusses a figure or figure items, that figure and any additional descriptions of that figure should be understood to be incorporated by reference as if set forth fully herein. To the extent any cell lacks citations to the charted reference(s), this should not be taken as an admission that the reference(s) do(es) not disclose the corresponding limitation but rather indicates that Bank does not presently have sufficient information regarding the charted reference(s) based on Bank’s present understanding of the claim limitation.

These invalidity contentions are not an admission by Bank that the accused products or components, including any current or past version of the products or components, are covered by, or infringe the asserted claims, particularly when these claims are properly construed and applied.

1. Claim 7

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
[p] A computer-implemented method of processing a natural language utterance, the method comprising:	To the extent the preamble is limiting, Kennewick discloses a “computer-implemented method of processing a natural language utterance[.]” <i>E.g.</i>, Kennewick at Fig. 1:  Kennewick at Fig. 5:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	FIG. 5 500  <pre> graph TD 502[INPUT UTTERANCE] --> 504[INTERPRET INPUT] 504 --> 506[SCORE THE INTERPRETATION] 506 --> 508{CONFIDENCE LEVEL SATISFACTORY?} 508 -- YES --> 516([CONTINUE WITH PROCESS]) 508 -- NO --> 510[REQUEST USER VERIFY INTERPRETATION] 510 --> 512{USER SATISFIED?} 512 -- YES --> 516 512 -- NO --> 514[HAVE USER REPHRASE UTTERANCE AND/OR PROVIDE MORE INFORMATION] 514 --> 502 </pre> Kennewick at Abstract: Systems and methods for receiving natural language queries and/or commands and execute the queries and/or commands. The systems and methods overcomes the deficiencies of prior art speech query and response systems through the application of a complete speech-based information query, retrieval, presentation and command environment. This environment makes significant use of context, prior information, domain knowledge, and user specific profile data to achieve a natural environment for one or more users making queries or commands in multiple domains. Through this integrated approach, a complete speech-based natural language query and response environment can be created. The systems and methods creates, stores and uses extensive

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	personal profile information for each user, thereby improving the reliability of determining the context and presenting the expected results for a particular question or command. Kennewick at [0032]: Although the system, according to another aspect of the invention, is intended to be able to accept most any natural language question or command, ambiguity can still be a problem. To assist users formulate concise questions and commands the system may support a voice query language. The language may help users clearly specify the keywords or contexts of the question or command along with the parameters or criteria. The system may provide built in training capabilities to help the user learn the best methods to formulate their questions and commands. Kennewick at [0010]: According to another aspect of the invention, the system comprises a speech unit interface device that receives spoken natural language queries, commands and/or other utterances from a user, and a computer device or system that receives input from the speech unit and processes the input (e.g., retrieves information responsive to the query, takes action consistent with the command and performs other functions as detailed herein). The system may further generate a natural language speech response in response to the user and/or generate a natural language speech message without any prompting from a user such as an alert message. Kennewick at [0151]: Alternatively, when a user uses a word that is not recognized at all or is not correctly recognized by the speech recognition engine 120 then the user may be asked to spell the word. The speech recognition engine 120 determines this condition based on confidence level for the scoring process. The word is looked up in the dictionary 112 and the pronunciation for the word is added to either the dictionary, the agent 106, or the user's profile 110. The word pronunciation can then be associated with the domain, the question, the context and the user. Though this process the

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	speech recognition engine learns with time and improves accuracy. To assist users in spelling words an individualized phonetic alphabet can be used. Each user can modify the standard phonetic alphabets with words, which they can remember more easily. Kennewick at [0153]: The parser 118 uses a scoring system to determine the mostly likely context or domain for a user's question and/or command. The score is determined from weighting a number of factors including, the user profile 110, the domain agent's data content and previous context. Based on this scoring, the system 90 invokes the correct agent. If the confidence level of the score is not high enough to ensure a reliable response, the system 90 may ask the user to verify whether the question and/or command is correctly understood. In general, the question that is asked by the system 90 may be phrased to indicate the context of the question including all criteria or parameters. For example, the question can be in the form of: "Did I understand that you want such-and-such" If the user confirms that the question is correct the system proceeds to produce a response. Otherwise, the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system can ask one or more questions to attempt to resolve the ambiguity. Kennewick at [0167-171]: Based on the continuous scoring processes, the domain agent 156 may determine if a single best answer can be extracted. For most questions the desired result has a set of tokens that must be found to formulate an answer. Once a value has been found for each of these tokens the results is ready for presentation to the user. For example, for a question on weather, the tokens can include the date, day of week, predicted high temperature, predicted low temperature, chance of precipitation, expected cloud cover, expected type of precipitation and other tokens. Results processed in this manner include error messages. For subjective questions this determination is made by determining a most likely answer or answers, extracted by matching of the results received. If no satisfactory answer can be inferred from the results of the query the agent can do one of the following:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	1. Ask the user for more information, typically through the speech interface, and based on the results obtained formulate new queries. This approach is applied when an irresolvable ambiguity arises in the formulation of a response. 2. Formulate new queries based on the results received from the first set of queries. This approach is typically applied in cases where the responses received do not contain all the required information. Information sources to query can be inferred from the results already obtained (i.e., links in an HTML document) or from other sources. Using this approach one or more sets of queries and responses can be chained without the need for action by the user. 3. Wait for additional queries to return results. In any case, the domain agent 156 may continue to make queries and evaluate results until a satisfactory response is constructed. In doing so, the agent can start several overlapping query paths or threads of inquiry, typically mediated by the event manager 100. This technique, combined with the use of asynchronous queries from multiple data sources provides the real-time response performance required for a natural interaction with the user. Kennewick at [0185]: FIG. 3 is a high level process 300 for receiving natural language speech-based queries and/or commands and generating a response according to an embodiment of the invention. The process 300 comprises of individual steps, several of which may be repetitive in order to accommodate partial failures. The process 300 may begin when a user's utterance is received and processed at 302. The user utterance may be a query and/or a command and may be of a natural language speech-based form. Kennewick at [0187]: FIG. 4B is a process 450 for receiving natural language speech-based commands in order to locally or remotely control functions of the system 90 or for other devices according to another embodiment of the invention. The process 450 may begin when a user's utterance (i.e., user query and/or command) is captured through accurate speech recognition operating in a variety of real-

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	world environments at 452. Once the utterance is captured, it is parse and interpret to determine the command that is contained in the utterance at 454. Kennewick at [0188]: Some of the steps depicting in FIG. 4A may actually require multiple steps that may also be repetitive. For instance, in step 404, the user's utterance is parsed and interpreted in order to determine the meaning of the utterance. The system 90 may make an initial interpretation of the utterance based the data stored in, for example, the dictionary and phrases module 124, user profiles, 110, agents 106 and the databases 104. A real-time scoring system or other techniques may be applied to the interpretation results in order to generate a domain or a context score. If the confidence level of the domain or context score is not high enough to ensure a reliable response, the system 90 can request that the user verify the question or command is correctly understood. In general, the question may be phrased to indicate the context of the question including all criteria or parameters. If the user confirms that the question is correct, the system 90 may proceed to produce a response. Otherwise, either the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system may ask one or more questions to attempt to resolve the ambiguity or other actions may taken. FIG. 5 is a process 500 for correctly interpreting a user's utterance according to one embodiment of the invention. The process 500 generally represents steps 402 and 404 of process 400. Initially the user's utterance is inputted into the system 90 at 502. The input is then interpreted using, for example, various stored data such as user profiles, agent data, dictionary and phrases, and other relevant data. The interpretation is then scored using, for instance, a weighted scoring system as described previously at 506. Once a score is obtained, determine whether the confidence level of the interpretation is satisfactory at 508. In other words, whether the score given to the interpretation exceeds a certain value. If the confidence level is determined to be unsatisfactory then a request may be submitted to the user requesting that the user verify the interpretation at 510 and 512. If the user is unsatisfied with the interpretation, he/she may be asked to rephrase the utterance and/or provide additional information at 514. Once the user provides the rephrased utterance and/or additional information, the process 500 returns to the beginning at 502. If, on the other hand, the user is satisfied with the


Dialect, LLC v Bank of America, N.A.
Eastern District of Texas No. 2:24-cv-00207-JRG


Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	interpretation, then the process 400 for receiving a natural language speech-based queries and/or commands and generating a response may continue at 516. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[a] generating a context stack comprising context information that corresponds to a plurality of prior utterances, wherein the context stack includes a plurality of context entries;	Kennewick discloses “generating a context stack comprising context information that corresponds to a plurality of prior utterances, wherein the context stack includes a plurality of context entries[.]” <i>E.g.</i>, Kennewick at Fig. 1:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	FIG. 1 Kennewick at [0108]: 1. Content packages include questions or commands. Each command or question or group of commands or questions includes contexts used for creation of one or more queries. The agent 156 passes a regular grammar expression to the parser 118 for evaluation of a context or question. An initial or default context is typically supplied for each command or question. The command or question includes a grammar for the management and evaluation of the context stack. Kennewick at [0145]: The voice query language may be sensitive to the contents of the context stack. Thus, a following-on question can be asked using an abbreviated grammar, since key words and criteria can be

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	inherited from the stack. For example, the user can simply asked about another keyword if the criteria of the question remain constant. Kennewick at [0149]: According to another embodiment of the invention, the system 90 may employ a speech recognition engine 124 seeding for improved word recognition accuracy, using data from the dictionary and phrase tables 112, user profiles 110, and the agents 106. At the same time, the fuzzy set possibilities or prior probabilities for the words in the dictionary and phrase tables may be dynamically updated to maximize the probability of correct recognition at each stage of the dialog. The probabilities or possibilities may be dynamically updated based on a number of criteria including the application domain, the questions or commands, contexts, the user profile and preferences, user dialog history, the recognizer dictionary and phrase tables, and word spellings. Kennewick at [0152]: Once the words and phrases have been recognized by the speech recognition engine 120, the tokens and user identification is passed to the parser 118. The parser 118 examines the tokens for the questions or commands, context and criteria. The parser 118 determines a context for an utterance by applying prior probabilities or fuzzy possibilities to keyword matching, user profile 110, and dialog history. The context of a question or command determines the domain and thereby, the domain agent 156, if any, to be evoked. For example, a question with the keywords “temperature” implies a context value of weather for the question. The parser dynamically receives keyword and associated prior probability or fuzzy possibility updates from the system agent 150 or an already active domain agent 156. Based on these probabilities or possibilities the possible contexts are scored and the top one or few are used for further processing. Kennewick at [0166]: The domain agent 156 evaluates the results of the one or more queries as they arrive. The domain agent 156 scores the relevance of the results based on results already received, the context, the criteria, the history of the dialog, the user profile 110 and domain specific information using

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	probabilistic or fuzzy scoring techniques. Part of the dialog history is maintained in a context stack. The weight of each context for the scoring is based on the relevance of one context to another and the age of the contexts. Other scoring variables can be associated through the context stack. Contexts can also be exclusive, so that previous contexts have no weight in the scoring. Kennewick at [0187]: Next, review the command and determine the domain for the command and context, invoking the proper resources including agents at 456. The original command submitted by the user will often require that multiple queries and/or commands be generated. For instance, suppose a user is interested in recorded his favorite television program. The user may utter “please record my favorite TV program.” The system 90 may review this request together with stored data such as the user's profile and determine keywords such as “record” and “my favorite TV program.” The system 90 may then generate queries to determine the name, the channel, and time for the user's favorite TV program. These queries may then be sent to a agent domain, which may access a database containing the user's profile and determine the answers to these questions. A command may then be sent to a video recorder ordering the video recorder to record the selected TV program. Thus, at 458 formulate one or more queries that are to be sent to one or more local and/or network data sources and/or send appropriate commands to local or remote devices or the system itself. Next, route the generated commands to the appropriate system[s] and/or external devices at 460. Once the commands have been executed, results including errors may be received and processed at 462. The results of the executed commands may be optionally presented to the user at 464. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[b] receiving, by one or more computer processors, the natural language utterance, wherein the natural language utterance is associated with a command or is associated with a request;	Kennewick discloses “receiving, by one or more computer processors, the natural language utterance, wherein the natural language utterance is associated with a command or is associated with a request[.]” <i>E.g.</i>, Kennewick at Fig. 1:  Kennewick at Fig. 5:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	FIG. 5 500  <pre> graph TD 502[INPUT UTTERANCE] --> 504[INTERPRET INPUT] 504 --> 506[SCORE THE INTERPRETATION] 506 --> 508{CONFIDENCE LEVEL SATISFACTORY?} 508 -- YES --> 516([CONTINUE WITH PROCESS]) 508 -- NO --> 510[REQUEST USER VERIFY INTERPRETATION] 510 --> 512{USER SATISFIED?} 512 -- YES --> 516 512 -- NO --> 514[HAVE USER REPHRASE UTTERANCE AND/OR PROVIDE MORE INFORMATION] 514 --> 502 </pre> Kennewick at Abstract: Systems and methods for receiving natural language queries and/or commands and execute the queries and/or commands. The systems and methods overcomes the deficiencies of prior art speech query and response systems through the application of a complete speech-based information query, retrieval, presentation and command environment. This environment makes significant use of context, prior information, domain knowledge, and user specific profile data to achieve a natural environment for one or more users making queries or commands in multiple domains. Through this integrated approach, a complete speech-based natural language query and response environment can be created. The systems and methods creates, stores and uses extensive

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	personal profile information for each user, thereby improving the reliability of determining the context and presenting the expected results for a particular question or command. Kennewick at [0032]: Although the system, according to another aspect of the invention, is intended to be able to accept most any natural language question or command, ambiguity can still be a problem. To assist users formulate concise questions and commands the system may support a voice query language. The language may help users clearly specify the keywords or contexts of the question or command along with the parameters or criteria. The system may provide built in training capabilities to help the user learn the best methods to formulate their questions and commands. Kennewick at [0010]: According to another aspect of the invention, the system comprises a speech unit interface device that receives spoken natural language queries, commands and/or other utterances from a user, and a computer device or system that receives input from the speech unit and processes the input (e.g., retrieves information responsive to the query, takes action consistent with the command and performs other functions as detailed herein). The system may further generate a natural language speech response in response to the user and/or generate a natural language speech mes Kennewick at [0151]: Alternatively, when a user uses a word that is not recognized at all or is not correctly recognized by the speech recognition engine 120 then the user may be asked to spell the word. The speech recognition engine 120 determines this condition based on confidence level for the scoring process. The word is looked up in the dictionary 112 and the pronunciation for the word is added to either the dictionary, the agent 106, or the user's profile 110. The word pronunciation can then be associated with the domain, the question, the context and the user. Though this process the speech recognition engine learns with time and improves accuracy. To assist users in spelling

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	words an individualized phonetic alphabet can be used. Each user can modify the standard phonetic alphabets with words, which they can remember more easily. Kennewick at [0153]: The parser 118 uses a scoring system to determine the mostly likely context or domain for a user's question and/or command. The score is determined from weighting a number of factors including, the user profile 110, the domain agent's data content and previous context. Based on this scoring, the system 90 invokes the correct agent. If the confidence level of the score is not high enough to ensure a reliable response, the system 90 may ask the user to verify whether the question and/or command is correctly understood. In general, the question that is asked by the system 90 may be phrased to indicate the context of the question including all criteria or parameters. For example, the question can be in the form of: "Did I understand that you want such-and-such" If the user confirms that the question is correct the system proceeds to produce a response. Otherwise, the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system can ask one or more questions to attempt to resolve the ambiguity. Kennewick at [0167-171]: Based on the continuous scoring processes, the domain agent 156 may determine if a single best answer can be extracted. For most questions the desired result has a set of tokens that must be found to formulate an answer. Once a value has been found for each of these tokens the results is ready for presentation to the user. For example, for a question on weather, the tokens can include the date, day of week, predicted high temperature, predicted low temperature, chance of precipitation, expected cloud cover, expected type of precipitation and other tokens. Results processed in this manner include error messages. For subjective questions this determination is made by determining a most likely answer or answers, extracted by matching of the results received. If no satisfactory answer can be inferred from the results of the query the agent can do one of the following:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	1. Ask the user for more information, typically through the speech interface, and based on the results obtained formulate new queries. This approach is applied when an irresolvable ambiguity arises in the formulation of a response. 2. Formulate new queries based on the results received from the first set of queries. This approach is typically applied in cases where the responses received do not contain all the required information. Information sources to query can be inferred from the results already obtained (i.e., links in an HTML document) or from other sources. Using this approach one or more sets of queries and responses can be chained without the need for action by the user. 3. Wait for additional queries to return results. In any case, the domain agent 156 may continue to make queries and evaluate results until a satisfactory response is constructed. In doing so, the agent can start several overlapping query paths or threads of inquiry, typically mediated by the event manager 100. This technique, combined with the use of asynchronous queries from multiple data sources provides the real-time response performance required for a natural interaction with the user. Kennewick at [0185]: FIG. 3 is a high level process 300 for receiving natural language speech-based queries and/or commands and generating a response according to an embodiment of the invention. The process 300 comprises of individual steps, several of which may be repetitive in order to accommodate partial failures. The process 300 may begin when a user's utterance is received and processed at 302. The user utterance may be a query and/or a command and may be of a natural language speech-based form. Kennewick at [0187]: FIG. 4B is a process 450 for receiving natural language speech-based commands in order to locally or remotely control functions of the system 90 or for other devices according to another embodiment of the invention. The process 450 may begin when a user's utterance (i.e., user query and/or command) is captured through accurate speech recognition operating in a variety of real-


Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	world environments at 452. Once the utterance is captured, it is parse and interpret to determine the command that is contained in the utterance at 454. Kennewick at [0188]: Some of the steps depicting in FIG. 4A may actually require multiple steps that may also be repetitive. For instance, in step 404, the user's utterance is parsed and interpreted in order to determine the meaning of the utterance. The system 90 may make an initial interpretation of the utterance based the data stored in, for example, the dictionary and phrases module 124, user profiles, 110, agents 106 and the databases 104. A real-time scoring system or other techniques may be applied to the interpretation results in order to generate a domain or a context score. If the confidence level of the domain or context score is not high enough to ensure a reliable response, the system 90 can request that the user verify the question or command is correctly understood. In general, the question may be phrased to indicate the context of the question including all criteria or parameters. If the user confirms that the question is correct, the system 90 may proceed to produce a response. Otherwise, either the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system may ask one or more questions to attempt to resolve the ambiguity or other actions may taken. FIG. 5 is a process 500 for correctly interpreting a user's utterance according to one embodiment of the invention. The process 500 generally represents steps 402 and 404 of process 400. Initially the user's utterance is inputted into the system 90 at 502. The input is then interpreted using, for example, various stored data such as user profiles, agent data, dictionary and phrases, and other relevant data. The interpretation is then scored using, for instance, a weighted scoring system as described previously at 506. Once a score is obtained, determine whether the confidence level of the interpretation is satisfactory at 508. In other words, whether the score given to the interpretation exceeds a certain value. If the confidence level is determined to be unsatisfactory then a request may be submitted to the user requesting that the user verify the interpretation at 510 and 512. If the user is unsatisfied with the interpretation, he/she may be asked to rephrase the utterance and/or provide additional information at 514. Once the user provides the rephrased utterance and/or additional information, the process 500 returns to the beginning at 502. If, on the other hand, the user is satisfied with the

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	interpretation, then the process 400 for receiving a natural language speech-based queries and/or commands and generating a response may continue at 516. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[c] determining one or more words of the natural language utterance by performing speech recognition on the natural language utterance;	Kennewick discloses “determining one or more words of the natural language utterance by performing speech recognition on the natural language utterance[.]” <i>E.g.</i>, Kennewick at Fig. 1:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	FIG. 1 The diagram illustrates a system 90 for speech processing. A Speech Unit 128 is connected to a Main Unit 98 via a Transceiver 130. The Speech Unit 128 includes a Speaker 136, Transceiver 130, Speech Coder 138, Filter 132, and Array Microphone 134. The Main Unit 98 includes a Text to Speech Engine 124, Transceiver 122, Speech Recognition Engine 120, Parser 118, Network Interface 116, Graphical User Interface 114, Event Manager 100, Dictionary and Phrases 112, User Profile 110, Personality 108, Agents 106, Update Manager 104, and Databases 102. The system is connected to the Internet 136. Kennewick at [0086]: The main user interface for the system 90, according to an embodiment of the invention, is through one or more speech units 128. The speech unit 128 includes one or more microphones, for example array microphone 134, to receive the utterances of the user. The speech received at the microphone 134 may be processed by filter 132 and passed to the speech coder 138 for encoding and compression. In one embodiment, a transceiver module 130 transmits the coded speech to the main unit 98. Coded speech received from the main unit 98 is detected by the transceiver 130, then decoded and decompressed by the speech coder 138 and annunciated by the speaker 136. Kennewick at [0088]:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	Coded speech received at the transceiver 126 on the main unit 98 is passed to the speech coder 122 for decoding and decompression. The decoded speech is processed by the speech recognition engine 120 using data in the dictionary and phrases module 112 and received from the agents 106. The recognized words and phrases are processed by the parser 118, which transforms them into complete commands and questions using data supplied by an agent[s] 106. The agent[s] 106 then process the commands or questions. The agent[s] 106 create queries to local databases 102 or through the network interface 116 to data sources on the Internet 136 or other networks. Commands typically result in actions taken by the system 90 itself (i.e., pause or stop), or to a remote device or data source (i.e., download data or program, or control a remote device), through the network interface to the Internet or other data interface. Kennewick at [0147]: The output of the speech coder 122 may be fed to the speech recognition engine 120. The speech recognition engine 120 recognizes words and phrases, using information in the dictionary and phrase tables 112, and passes these to the parser 118 for interpretation. The speech recognition engine 120 may determine the user's identity by voice and name for each utterance. Recognized words and phrases may be tagged with this identity in all further processing. Thus, as multiple users engage in overlapping sessions, the tags added by the speech recognition engine 120 to each utterance allows other components of the system 90 to tie that utterance to the correct user and dialog. The user recognition capability may further be used as a security measure for applications, such as auctions or online shopping, where this is required. Voice characteristics of each user may be contained in the user profile 110. Kennewick at [0150-151]: For uncommon words or new vocabulary words, a user may be given the option to spell the words. The spelling may be done by saying the names or the letters or using a phonetic alphabet. The phonetic alphabet can be a default one or one of the user's choosing.

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	Alternatively, when a user uses a word that is not recognized at all or is not correctly recognized by the speech recognition engine 120 then the user may be asked to spell the word. The speech recognition engine 120 determines this condition based on confidence level for the scoring process. The word is looked up in the dictionary 112 and the pronunciation for the word is added to either the dictionary, the agent 106, or the user's profile 110. The word pronunciation can then be associated with the domain, the question, the context and the user. Though this process the speech recognition engine learns with time and improves accuracy. To assist users in spelling words an individualized phonetic alphabet can be used. Each user can modify the standard phonetic alphabets with words, which they can remember more easily. Kennewick at [0185]: After receiving the natural language query and/or command, the query and/or command may be processed so that the data is in a more machine understandable format such as in a digital format. Once the utterance has been converted into an understandable format, the utterance may be processed so that at least an approximate meaning of the utterance can be determined at 304. This may be accomplished by, for example, the use of stored dictionaries, phrases, user profiles, domain agent data, and other data. Based on the approximate meaning of the utterance, one or more correctly formatted queries and/or commands may be generated. A user created query and/or command may require that multiple queries and/or commands may be needed in order to generate the desired response or action. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
[d] identifying, from among the plurality of context entries, one or more context entries that correspond to the one or more words, wherein the context information includes the one or more context entries, wherein identifying the one or more context entries comprises:	within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. Kennewick discloses “identifying, from among the plurality of context entries, one or more context entries that correspond to the one or more words, wherein the context information includes the one or more context entries[.]” <i>E.g.</i>, Kennewick at Fig. 1:  Kennewick at [0108]: 1. Content packages include questions or commands. Each command or question or group of commands or questions includes contexts used for creation of one or more queries. The agent 156 passes a regular grammar expression to the parser 118 for evaluation of a context or question. An

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	initial or default context is typically supplied for each command or question. The command or question includes a grammar for the management and evaluation of the context stack. Kennewick at [0145]: The voice query language may be sensitive to the contents of the context stack. Thus, a following-on question can be asked using an abbreviated grammar, since key words and criteria can be inherited from the stack. For example, the user can simply asked about another keyword if the criteria of the question remain constant. Kennewick at [0149]: According to another embodiment of the invention, the system 90 may employ a speech recognition engine 124 seeding for improved word recognition accuracy, using data from the dictionary and phrase tables 112, user profiles 110, and the agents 106. At the same time, the fuzzy set possibilities or prior probabilities for the words in the dictionary and phrase tables may be dynamically updated to maximize the probability of correct recognition at each stage of the dialog. The probabilities or possibilities may be dynamically updated based on a number of criteria including the application domain, the questions or commands, contexts, the user profile and preferences, user dialog history, the recognizer dictionary and phrase tables, and word spellings. Kennewick at [0152]: Once the words and phrases have been recognized by the speech recognition engine 120, the tokens and user identification is passed to the parser 118. The parser 118 examines the tokens for the questions or commands, context and criteria. The parser 118 determines a context for an utterance by applying prior probabilities or fuzzy possibilities to keyword matching, user profile 110, and dialog history. The context of a question or command determines the domain and thereby, the domain agent 156, if any, to be evoked. For example, a question with the keywords “temperature” implies a context value of weather for the question. The parser dynamically receives keyword and associated prior probability or fuzzy possibility updates from the system

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	agent 150 or an already active domain agent 156. Based on these probabilities or possibilities the possible contexts are scored and the top one or few are used for further processing. Kennewick at [0166]: The domain agent 156 evaluates the results of the one or more queries as they arrive. The domain agent 156 scores the relevance of the results based on results already received, the context, the criteria, the history of the dialog, the user profile 110 and domain specific information using probabilistic or fuzzy scoring techniques. Part of the dialog history is maintained in a context stack. The weight of each context for the scoring is based on the relevance of one context to another and the age of the contexts. Other scoring variables can be associated through the context stack. Contexts can also be exclusive, so that previous contexts have no weight in the scoring. Kennewick at [0187]: Next, review the command and determine the domain for the command and context, invoking the proper resources including agents at 456. The original command submitted by the user will often require that multiple queries and/or commands be generated. For instance, suppose a user is interested in recorded his favorite television program. The user may utter “please record my favorite TV program.” The system 90 may review this request together with stored data such as the user's profile and determine keywords such as “record” and “my favorite TV program.” The system 90 may then generate queries to determine the name, the channel, and time for the user's favorite TV program. These queries may then be sent to a agent domain, which may access a database containing the user's profile and determine the answers to these questions. A command may then be sent to a video recorder ordering the video recorder to record the selected TV program. Thus, at 458 formulate one or more queries that are to be sent to one or more local and/or network data sources and/or send appropriate commands to local or remote devices or the system itself. Next, route the generated commands to the appropriate system[s] and/or external devices at 460. Once the commands have been executed, results including errors may be received and processed at 462. The results of the executed commands may be optionally presented to the user at 464.

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[e] comparing the plurality of context entries to the one or more words;	Kennewick discloses “comparing the plurality of context entries to the one or more words[.]” <i>E.g.</i>, Kennewick at [0108]: 1. Content packages include questions or commands. Each command or question or group of commands or questions includes contexts used for creation of one or more queries. The agent 156 passes a regular grammar expression to the parser 118 for evaluation of a context or question. An initial or default context is typically supplied for each command or question. The command or question includes a grammar for the management and evaluation of the context stack. Kennewick at [0145]: The voice query language may be sensitive to the contents of the context stack. Thus, a following-on question can be asked using an abbreviated grammar, since key words and criteria can be inherited from the stack. For example, the user can simply asked about another keyword if the criteria of the question remain constant. Kennewick at [0149]: According to another embodiment of the invention, the system 90 may employ a speech recognition engine 124 seeding for improved word recognition accuracy, using data from the

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	dictionary and phrase tables 112, user profiles 110, and the agents 106. At the same time, the fuzzy set possibilities or prior probabilities for the words in the dictionary and phrase tables may be dynamically updated to maximize the probability of correct recognition at each stage of the dialog. The probabilities or possibilities may be dynamically updated based on a number of criteria including the application domain, the questions or commands, contexts, the user profile and preferences, user dialog history, the recognizer dictionary and phrase tables, and word spellings. Kennewick at [0152]: Once the words and phrases have been recognized by the speech recognition engine 120, the tokens and user identification is passed to the parser 118. The parser 118 examines the tokens for the questions or commands, context and criteria. The parser 118 determines a context for an utterance by applying prior probabilities or fuzzy possibilities to keyword matching, user profile 110, and dialog history. The context of a question or command determines the domain and thereby, the domain agent 156, if any, to be evoked. For example, a question with the keywords “temperature” implies a context value of weather for the question. The parser dynamically receives keyword and associated prior probability or fuzzy possibility updates from the system agent 150 or an already active domain agent 156. Based on these probabilities or possibilities the possible contexts are scored and the top one or few are used for further processing. Kennewick at [0166]: The domain agent 156 evaluates the results of the one or more queries as they arrive. The domain agent 156 scores the relevance of the results based on results already received, the context, the criteria, the history of the dialog, the user profile 110 and domain specific information using probabilistic or fuzzy scoring techniques. Part of the dialog history is maintained in a context stack. The weight of each context for the scoring is based on the relevance of one context to another and the age of the contexts. Other scoring variables can be associated through the context stack. Contexts can also be exclusive, so that previous contexts have no weight in the scoring. Kennewick at [0187]:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	Next, review the command and determine the domain for the command and context, invoking the proper resources including agents at 456. The original command submitted by the user will often require that multiple queries and/or commands be generated. For instance, suppose a user is interested in recorded his favorite television program. The user may utter “please record my favorite TV program.” The system 90 may review this request together with stored data such as the user's profile and determine keywords such as “record” and “my favorite TV program.” The system 90 may then generate queries to determine the name, the channel, and time for the user's favorite TV program. These queries may then be sent to a agent domain, which may access a database containing the user's profile and determine the answers to these questions. A command may then be sent to a video recorder ordering the video recorder to record the selected TV program. Thus, at 458 formulate one or more queries that are to be sent to one or more local and/or network data sources and/or send appropriate commands to local or remote devices or the system itself. Next, route the generated commands to the appropriate system[s] and/or external devices at 460. Once the commands have been executed, results including errors may be received and processed at 462. The results of the executed commands may be optionally presented to the user at 464. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[f] generating, based on the comparison, one or more rank scores for individual context	Kennewick discloses “generating, based on the comparison, one or more rank scores for individual context entries of the plurality of context entries[.]” <i>E.g.</i>, Kennewick at [0151-153]:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
entries of the plurality of context entries; and	Alternatively, when a user uses a word that is not recognized at all or is not correctly recognized by the speech recognition engine 120 then the user may be asked to spell the word. The speech recognition engine 120 determines this condition based on confidence level for the scoring process. The word is looked up in the dictionary 112 and the pronunciation for the word is added to either the dictionary, the agent 106, or the user's profile 110. The word pronunciation can then be associated with the domain, the question, the context and the user. Though this process the speech recognition engine learns with time and improves accuracy. To assist users in spelling words an individualized phonetic alphabet can be used. Each user can modify the standard phonetic alphabets with words, which they can remember more easily. Once the words and phrases have been recognized by the speech recognition engine 120, the tokens and user identification is passed to the parser 118. The parser 118 examines the tokens for the questions or commands, context and criteria. The parser 118 determines a context for an utterance by applying prior probabilities or fuzzy possibilities to keyword matching, user profile 110, and dialog history. The context of a question or command determines the domain and thereby, the domain agent 156, if any, to be evoked. For example, a question with the keywords "temperature" implies a context value of weather for the question. The parser dynamically receives keyword and associated prior probability or fuzzy possibility updates from the system agent 150 or an already active domain agent 156. Based on these probabilities or possibilities the possible contexts are scored and the top one or few are used for further processing. The parser 118 uses a scoring system to determine the mostly likely context or domain for a user's question and/or command. The score is determined from weighting a number of factors including, the user profile 110, the domain agent's data content and previous context. Based on this scoring, the system 90 invokes the correct agent. If the confidence level of the score is not high enough to ensure a reliable response, the system 90 may ask the user to verify whether the question and/or command is correctly understood. In general, the question that is asked by the system 90 may be phrased to indicate the context of the question including all criteria or parameters. For example, the question can be in the form of: "Did I understand that you want such-and-such" If the user confirms that the question is correct the system proceeds to produce a response. Otherwise, the

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system can ask one or more questions to attempt to resolve the ambiguity. Kennewick at [0166]: The domain agent 156 evaluates the results of the one or more queries as they arrive. The domain agent 156 scores the relevance of the results based on results already received, the context, the criteria, the history of the dialog, the user profile 110 and domain specific information using probabilistic or fuzzy scoring techniques. Part of the dialog history is maintained in a context stack. The weight of each context for the scoring is based on the relevance of one context to another and the age of the contexts. Other scoring variables can be associated through the context stack. Contexts can also be exclusive, so that previous contexts have no weight in the scoring. Kennewick at [0188]: FIG. 5 is a process 500 for correctly interpreting a user's utterance according to one embodiment of the invention. The process 500 generally represents steps 402 and 404 of process 400. Initially the user's utterance is inputted into the system 90 at 502. The input is then interpreted using, for example, various stored data such as user profiles, agent data, dictionary and phrases, and other relevant data. The interpretation is then scored using, for instance, a weighted scoring system as described previously at 506. Once a score is obtained, determine whether the confidence level of the interpretation is satisfactory at 508. In other words, whether the score given to the interpretation exceeds a certain value. If the confidence level is determined to be unsatisfactory then a request may be submitted to the user requesting that the user verify the interpretation at 510 and 512. If the user is unsatisfied with the interpretation, he/she may be asked to rephrase the utterance and/or provide additional information at 514. Once the user provides the rephrased utterance and/or additional information, the process 500 returns to the beginning at 502. If, on the other hand, the user is satisfied with the interpretation, then the process 400 for receiving a natural language speech-based queries and/or commands and generating a response may continue at 516. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim,

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[g] identifying, based on the one or more rank scores, the one or more context entries from among the plurality of context entries; and	Kennewick discloses “identifying, based on the one or more rank scores, the one or more context entries from among the plurality of context entries[.]” <i>E.g.</i>, Kennewick at [0153]: The parser 118 uses a scoring system to determine the mostly likely context or domain for a user's question and/or command. The score is determined from weighting a number of factors including, the user profile 110, the domain agent's data content and previous context. Based on this scoring, the system 90 invokes the correct agent. If the confidence level of the score is not high enough to ensure a reliable response, the system 90 may ask the user to verify whether the question and/or command is correctly understood. In general, the question that is asked by the system 90 may be phrased to indicate the context of the question including all criteria or parameters. For example, the question can be in the form of: “Did I understand that you want such-and-such” If the user confirms that the question is correct the system proceeds to produce a response. Otherwise, the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system can ask one or more questions to attempt to resolve the ambiguity. Kennewick at [0166]: The domain agent 156 evaluates the results of the one or more queries as they arrive. The domain agent 156 scores the relevance of the results based on results already received, the context, the criteria, the history of the dialog, the user profile 110 and domain specific information using probabilistic or fuzzy scoring techniques. Part of the dialog history is maintained in a context stack. The weight of each context for the scoring is based on the relevance of one context to

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	another and the age of the contexts. Other scoring variables can be associated through the context stack. Contexts can also be exclusive, so that previous contexts have no weight in the scoring. Kennewick at [0167]: Based on the continuous scoring processes, the domain agent 156 may determine if a single best answer can be extracted. For most questions the desired result has a set of tokens that must be found to formulate an answer. Once a value has been found for each of these tokens the results is ready for presentation to the user. For example, for a question on weather, the tokens can include the date, day of week, predicted high temperature, predicted low temperature, chance of precipitation, expected cloud cover, expected type of precipitation and other tokens. Results processed in this manner include error messages. For subjective questions this determination is made by determining a most likely answer or answers, extracted by matching of the results received. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[h] determining based on the determined one or more words and the context information, the command or the request associated with the natural language utterance	Kennewick discloses “determining based on the determined one or more words and the context information, the command or the request associated with the natural language utterance[.]” <i>E.g.</i>, Kennewick at [0152]:

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	Once the words and phrases have been recognized by the speech recognition engine 120, the tokens and user identification is passed to the parser 118. The parser 118 examines the tokens for the questions or commands, context and criteria. The parser 118 determines a context for an utterance by applying prior probabilities or fuzzy possibilities to keyword matching, user profile 110, and dialog history. The context of a question or command determines the domain and thereby, the domain agent 156, if any, to be evoked. For example, a question with the keywords “temperature” implies a context value of weather for the question. The parser dynamically receives keyword and associated prior probability or fuzzy possibility updates from the system agent 150 or an already active domain agent 156. Based on these probabilities or possibilities the possible contexts are scored and the top one or few are used for further processing. Kennewick at [0154-155]: Once the context for the question or command has been determined, the parser 118 can invoke the correct agent 156, 150. To formulate a question or command in the regular grammar used by agents, the parser 118 may determine required and optional values for the criteria or parameters. These criteria may have been explicitly supplied by the user or may need to be inferred. The parser 118 makes use of the criteria handlers 152 supplied by the system agent 150. The criteria handlers 152 provide context sensitive procedures for extracting the criteria or parameters from the user's question or command. Some criteria are determined by executing algorithms in the agent, while others may be determined by applying probabilistic or fuzzy reasoning to tables of possible values. Prior probabilities or fuzzy possibilities and associated values are received from a number of sources including the history of the dialog, the user profile 110, and the agent. Based on user responses, the prior probabilities or fuzzy possibilities are updated as the system learns the desired behavior. For a weather context, examples of criteria include, location, date and time. Other criteria can include command criteria (i.e., yes/no, on/off, pause, stop), and spelling. Special criteria handlers are available from the system agent for processing lists, tables, barge-in commands, long strings of text and system commands. The criteria handlers 152 operate iteratively or recursively on the criteria extracted to eliminate ambiguity. This processing helps reduce the ambiguity in the user's question or command. For

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	example, if the user has a place name (or other proper noun) in their utterance the parser 118 can use services of the domain agent 156 to look up tables in the databases 102 for place names or can attempt to determine which word is the proper noun from the syntax of the utterance. In another example, the user asks “what about flight one hundred and twenty too?”. The parser and domain agent use flight information in the database and network information along with context to determine the most plausible interpretation among; flight 100 and flight 20 also, flight 100 and flight 22, flight 122, and the like. Kennewick at [0161-164]: The agents 150, 156 may receive a command or question once the parser 118 has placed it in the required standard format. Based on the context, the parser 118 evokes the correct agent to process the question or command. Commands can be directed to the system 90 or to an external entity. System commands are generally directed to the system agent 150. Commands for external entities are generally processed by a domain agent 156, which includes the command context and behavior for the external entity. Specific questions are generally directed to one of the domain agents 156. Based on the question or context and the parameters or criteria, the domain agent creates one or more queries to one or more local or external information sources. Questions can be objective or subjective in nature. Results for objective questions can often be obtained by structured queries to one or more local or network information sources. Even for objective questions, the system 90 may need to apply probabilistic or fuzzy set analysis to deal with cases of conflicting information or incomplete information. Information to answer subjective questions is generally obtained by one or more ad-hoc queries to local or network data sources, followed by probabilistic or fuzzy set evaluation of the one results to determine a best answer. Once the domain agent 156 has formulated the one or more queries, they may be sent to local and/or network information sources. The queries are performed in an asynchronous manner to

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	account for the fact that sources respond at different speeds or may fail to respond at all. Duplicate queries are sent to different information sources to ensure that at least one source responds with a useful result in a timely manner. Further, if multiple results are received in a timely manner, they can be scored by the system to determine which data is most reliable or appropriate. Examples of data sources accommodated include, HTTP data sources, sources with meta-data in various formats including XML, entertainment audio, video and game files including MP3, databases using query languages and structured responses such as SQL, and other data sources. Kennewick at [0185]: Further, depending upon the context (e.g., who is the user, what is the subject the query and/or command directed to, at what time or location is the query and/or command being directed to, and other parameters used in defining the query and/or command) in which the query and/or command has been submitted, a specific format for the query and/or command may be required. For instance, each of the domain agents 156 may require that queries and/or commands be formatted in a specific manner. Thus at 306, one or more properly formatted queries and/or commands may be generated. Steps 304 and 306 may be operations which may have to be repeated in order to obtain the correct interpretation of the user utterance and result in the desired response and/or action. Once, the properly formatted queries and/or commands are generated, the appropriate action may be taken at 308. This may require that the properly formatted queries and/or commands be sent to specific domain agent[s], information source[s], device[s], or other appropriate destination[s] that can fulfill the requirements of the query and/or command. Once the appropriate action[s] have been executed, this event may be recorded to, for example, the user's profile, database and/or one or more agents at 310. Such data may be useful for future user inquiries and commands. After the action[s] have been performed, a response, if need be, may be generated and forwarded to the user and/or third parties at 312. In the case of a query for retrieving data, the response would contain the requested information. In the case of a command, the response may be a confirmation that a specific action[s] has been executed. The response may be in the form of a natural language format. The response may also be formatted to reflect a particular personality or tone to the response in order to make the response more "human." The response

Claim 7 of the '957 Patent	Exemplary Citations to Kennewick
	may relayed to the user and/or third parties as an audio message and/or a visual message displayed on a user interface. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 8

Claim 8 of the '957 Patent	Exemplary Citations to Kennewick
[a] The method of claim 7, further comprising: providing a prompt regarding the command or request, wherein the prompt is a prompt for additional information or a prompt for confirmation; and	Kennewick discloses “providing a prompt regarding the command or request, wherein the prompt is a prompt for additional information or a prompt for confirmation[.]” <i>E.g.</i>, Kennewick at [0153]: The parser 118 uses a scoring system to determine the mostly likely context or domain for a user's question and/or command. The score is determined from weighting a number of factors including, the user profile 110, the domain agent's data content and previous context. Based on this scoring, the system 90 invokes the correct agent. If the confidence level of the score is not high enough to ensure a reliable response, the system 90 may ask the user to verify whether the question and/or command is correctly understood. In general, the question that is asked by the system 90 may be phrased to indicate the context of the question including all criteria or parameters. For example, the question can be in the form of: “Did I understand that you want such-and-such” If the user

Claim 8 of the '957 Patent	Exemplary Citations to Kennewick
	confirms that the question is correct the system proceeds to produce a response. Otherwise, the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system can ask one or more questions to attempt to resolve the ambiguity. Kennewick at [0145]: The voice query language may be sensitive to the contents of the context stack. Thus, a following-on question can be asked using an abbreviated grammar, since key words and criteria can be inherited from the stack. For example, the user can simply asked about another keyword if the criteria of the question remain constant. Kennewick at [0167-168]: Based on the continuous scoring processes, the domain agent 156 may determine if a single best answer can be extracted. For most questions the desired result has a set of tokens that must be found to formulate an answer. Once a value has been found for each of these tokens the results is ready for presentation to the user. For example, for a question on weather, the tokens can include the date, day of week, predicted high temperature, predicted low temperature, chance of precipitation, expected cloud cover, expected type of precipitation and other tokens. Results processed in this manner include error messages. For subjective questions this determination is made by determining a most likely answer or answers, extracted by matching of the results received. If no satisfactory answer can be inferred from the results of the query the agent can do one of the following: 1. Ask the user for more information, typically through the speech interface, and based on the results obtained formulate new queries. This approach is applied when an irresolvable ambiguity arises in the formulation of a response. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of

Claim 8 of the '957 Patent	Exemplary Citations to Kennewick
	ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[b] receiving, in response to providing the prompt, a non-speech input, wherein the command or request is determined further based on the non-speech input	Kennewick discloses “receiving, in response to providing the prompt, a non-speech input, wherein the command or request is determined further based on the non-speech input[.]” <i>E.g.</i>, Kennewick at [0153]: The parser 118 uses a scoring system to determine the mostly likely context or domain for a user's question and/or command. The score is determined from weighting a number of factors including, the user profile 110, the domain agent's data content and previous context. Based on this scoring, the system 90 invokes the correct agent. If the confidence level of the score is not high enough to ensure a reliable response, the system 90 may ask the user to verify whether the question and/or command is correctly understood. In general, the question that is asked by the system 90 may be phrased to indicate the context of the question including all criteria or parameters. For example, the question can be in the form of: “Did I understand that you want such-and-such” If the user confirms that the question is correct the system proceeds to produce a response. Otherwise, the user can rephrase the original question, perhaps adding additional information to remove ambiguity, or the system can ask one or more questions to attempt to resolve the ambiguity. Kennewick at [0145]: The voice query language may be sensitive to the contents of the context stack. Thus, a following-on question can be asked using an abbreviated grammar, since key words and criteria can be inherited from the stack. For example, the user can simply asked about another keyword if the criteria of the question remain constant.

Claim 8 of the '957 Patent	Exemplary Citations to Kennewick
	Kennewick at [0167-168]: Based on the continuous scoring processes, the domain agent 156 may determine if a single best answer can be extracted. For most questions the desired result has a set of tokens that must be found to formulate an answer. Once a value has been found for each of these tokens the results is ready for presentation to the user. For example, for a question on weather, the tokens can include the date, day of week, predicted high temperature, predicted low temperature, chance of precipitation, expected cloud cover, expected type of precipitation and other tokens. Results processed in this manner include error messages. For subjective questions this determination is made by determining a most likely answer or answers, extracted by matching of the results received. If no satisfactory answer can be inferred from the results of the query the agent can do one of the following: 1. Ask the user for more information, typically through the speech interface, and based on the results obtained formulate new queries. This approach is applied when an irresolvable ambiguity arises in the formulation of a response. To the extent it is found that Kennewick does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Kennewick does not anticipate this claim, Kennewick renders it obvious, standing alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with one or more other prior art references identified in the Cover Pleading or herein, for the reasons identified herein or in the Cover Pleading. A person of ordinary skill in the art would have been motivated to combine Kennewick with any of the other prior art references identified in the Cover Pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.