

Description

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of the following patent applications, the contents of which are incorporated by reference herein:

U.S. Patent Application No. 13/874,523, entitled *Method of and System for Inferring User Intent in Search Input in a Conversational Interaction System*, filed on May 1, 2013;
 U.S. Patent Application No. 13/667,400, entitled *Method of and System for Inferring User Intent in Search Input in a Conversational Interaction System*, filed on November 2, 2012,
 U.S. Patent Application No. 13/667,388, entitled *Method of and System for Using Conversation State Information in a Conversational Interaction System*, filed on November 2, 2012;
 U.S. Provisional Patent Application No. 61/712,721, entitled *Method of and System for Content Search Based on Conceptual Language Clustering*, filed on October 11, 2012; and
 U.S. Provisional Patent Application No. 61/673,867, entitled *A Conversational Interaction System for Large Corpus Information Retrieval*, filed on July 20, 2012.

BACKGROUND OF THE INVENTION

Field of Invention

[0002] The invention generally relates to conversational interaction techniques, and, more specifically, to inferring user input intent based on resolving input ambiguities and/or inferring a change in conversational session has occurred.

Description of Related Art

[0003] Conversational systems are poised to become a preferred mode of navigating large information repositories across a range of devices: Smartphones, Tablets, TVs/STBs, multi-modal devices such as wearable computing devices such as "Goggles" (Google's sunglasses), hybrid gesture recognition/speech recognition systems like Xbox/Kinect, automobile information systems, and generic home entertainment systems. The era of touch based interfaces being center stage, as the primary mode of interaction, is perhaps slowly coming to an end, where in many daily life use cases, user would rather speak his intent, and the system understands and executes on the intent. This has also been triggered by the significant hardware, software and algorithmic advances making text to speech significantly effective compared to a few years ago.

[0004] While progress is being made towards pure

conversation interfaces, existing simple request response style conversational systems suffice only to addresses specific task oriented or specific information retrieval problems in small sized information repositories - these systems fail to perform well on large corpus information repositories.

[0005] Current systems that are essentially request response systems at their core, attempt to offer a conversational style interface such as responding to users question, as follows:

User: What is my checking account balance?
 System: It is \$2,459.34.
 User: And savings?
 System: It is \$6,209.012.
 User: How about the money market?
 System: It is \$14,599.33.

[0006] These are inherently goal oriented or task oriented request response systems providing a notion of continuity of conversation though each request response pair is independent of the other and the only context maintained is the simple context that it is user's bank account Other examples of current conversational systems are ones that walk user through a sequence of well-defined and often predetermined decision tree paths, to complete user intent (such as making a dinner reservation, booking a flight etc.)

[0007] Applicants have discovered that understanding user intent (even within a domain such as digital entertainment where user intent could span from pure information retrieval, to watching a show, or reserving a ticket for a show/movie), combined with understanding the semantics of the user utterance expressing the intent, so as to provide a clear and succinct response matching user intent is a hard problem that present systems in the conversation space fall short in addressing. Barring simple sentences with clear expression of intent, it is often hard to extract intent and the semantics of the sentence that expresses the intent, even in a single request/response exchange style interaction. Adding to this complexity, are intents that are task oriented without having well defined steps (such as the traversal of a predetermined decision tree). Also problematic are interactions that require a series of user requests and system responses to get to the completion of a task (e.g., like making a dinner reservation). Further still, rich information repositories can be especially challenging because user intent expression for an entity may take many valid and natural forms, and the same lexical tokens (words) may arise in relation to many different user intents.

[0008] When the corpus is large, lexical conflict or multiple semantic interpretations add to the complexity of satisfying user intent without a dialog to clarify these conflicts and ambiguities. Sometimes it may not even be possible to understand user intent, or the semantics of the sentence that expresses the intent - similar to what happens in real life conversations between humans. The abil-

ity of the system to ask the minimal number of questions (from the point of view of comprehending the other person in the conversation) to understand user intent, just like a human would do (on average where the participants are both aware of the domain being discussed), would define the closeness of the system to human conversations.

[0009] Systems that engage in a dialog or conversation, which go beyond the simple multi-step travel/dinner reservation making (e.g., where the steps in the dialog are well defined request / response subsequences with not much ambiguity resolution in each step), also encounter the complexity of having to maintain the state of the conversation in order to be effective. For example, such systems would need to infer implicit references to intents and entities (e.g., reference to people, objects or any noun) and attributes that qualify the intent in user's sentences (e.g., "show me the **latest** movies of Tom Hanks and not the **old** ones; "show me **more** action and **less** violence). Further still, applicants have discovered that it is beneficial to track not only references made by the user to entities, attributes, etc. in previous entries, but also to entities, attributes, etc. of multi-modal responses of the system to the user.

[0010] Further still, applicants have found that maintaining pronoun to object/subject associations during user / system exchanges enhances the user experience. For example, a speech analyzer (or natural language processor) that relates the pronoun "it" to its object/subject "Led Zeppelin song" in a complex user entry, such as, "The Led Zeppelin song in the original sound track of the recent Daniel Craig movie... Who performed it?" assists the user by not requiring the user to always use a particular syntax. However, this simple pronoun to object/subject association is ineffective in processing the following exchange:

Q1: Who acts as Obi-wan Kenobi in the new star wars?

A: Ewan McGregor.

Q2: How about his movies with Scarlet Johansson?

[0011] Here the "his" in the second question refers to the person in the response, rather than from the user input. A more complicated example follows:

Q1: Who played the lead roles in Kramer vs. Kramer?

A1: Meryl Streep and Dustin Hoffman.

Q2: How about more of his movies?

A2: Here are some of Dustin Hoffman movies... [list of Dustin Hoffman movies].

Q3: What about more of her movies?

[0012] Here the "his" in Q2 and "her" in Q3 refer back to the response A1. A natural language processor in isolation is ineffective in understanding user intent in these cases. In several of the embodiments described below, the language processor works in conjunction with a conversation state engine and domain specific information

indicating male and female attributes of the entities that can help resolve these pronoun references to prior conversation exchanges.

[0013] Another challenge facing systems that engage a user in conversation is the determination of the user's intent change, even if it is within the same domain. For example, user may start off with the intent of finding an answer to a question, e.g., in the entertainment domain. While engaging in the conversation of exploring more about that question, decide to pursue a completely different intent path. Current systems expect user to offer a clear cue that a new conversation is being initiated. If the user fails to provide that important clue, the system responses would be still be constrained to the narrow scope of the exploration path user has gone down, and will constrain users input to that narrow context, typically resulting undesirable, if not absurd, responses. The consequence of getting the context wrong is even more glaring (to the extent that the system looks comically inept) when user chooses to switch domains in the middle of a conversation. For instance, user may, while exploring content in the entertainment space, say, "I am hungry". If the system does not realize this as a switch to a new domain (restaurant/food domain), it may respond thinking "I am hungry" is a question posed in the entertainment space and offer responses in that domain, which in this case, would be a comically incorrect response.

[0014] A human, on the other hand, naturally recognizes such a drastic domain switch by the very nature of the statement, and responds accordingly (e.g., "Shall we order pizza?"). Even in the remote scenario where the transition to new domain is not so evident, a human participant may falter, but quickly recover, upon feedback from the first speaker ("Oh no. I mean I am hungry -1 would like to eat!"). These subtle, yet significant, elements of a conversation, that humans take for granted in conversations, are the ones that differentiate the richness of human-to-human conversations from that with automated systems.

[0015] In summary, embodiments of the techniques disclosed herein attempt to closely match user's intent and engage the user in a conversation not unlike human interactions. Certain embodiments exhibit any one or more of the following, non-exhaustive list of characteristics: a) resolve ambiguities in intent and/or description of the intent and, whenever applicable, leverage off of user's preferences (some implementations use computing elements and logic that are based on domain specific vertical information); b) maintain state of active intents and/or entities/attributes describing the intent across exchanges with the user, so as to implicitly infer references made by user indirectly to intents/entities/attributes mentioned earlier in a conversation; c) tailor responses to user, whenever applicable, to match user's preferences; d) implicitly determine conversation boundaries that start a new topic within and across domains and tailor a response accordingly; e) given a failure to understand user's intent (e.g., either because the intent cannot be found

or the confidence score of its best guess is below a threshold), engage in a minimal dialog to understand user intent (in a manner similar to that done by humans in conversations to understand intent). In some embodiments of the invention, the understanding of the intent may leverage off the display capacity of the device (e.g., like a tablet device) to graphically display intuitive renditions that user could interact with to offer clues on user intent.

BRIEF SUMMARY OF THE INVENTION

[0016] In an aspect of the invention, a method of and system for inferring user intent in search input in a conversational interaction system is disclosed.

[0017] In another aspect of the invention, a method of inferring user intent in a search input based on resolving ambiguous portions of the search input includes providing access to a set of content items. Each of the content items is associated with metadata that describes the corresponding content item. The method also includes providing a user preference signature. The user preference signature describes preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items. The method also includes receiving search input from the user. The search input is intended by the user to identify at least one desired content item. The method also includes determining that a portion of the search input contains an ambiguous identifier. The ambiguous identifier being intended by the user to identify, at least in part, the at least one desired content item. The method also includes inferring a meaning for the ambiguous identifier based on matching portions of the search input to the preferences of the user described by the user preference signature and selecting content items from the set of content items based on comparing the search input and the inferred meaning of the ambiguous identifier with metadata associated with the content items.

[0018] In a further aspect of the invention, the ambiguous identifier can be a pronoun, a syntactic expletive, an entertainment genre, and/or at least a portion of a name.

[0019] In still another aspect of the invention, the metadata associated with the content items includes a mapping of relationships between entities associated with the content items.

[0020] In yet a further aspect of the invention, the user preference signature is based on explicit preferences provided by the user and/or is based on analyzing content item selections by user conducted over a period of time. Optionally, the user preference signature describes preferences of the user for metadata associated with the content items, the metadata including entities preferred by the user.

[0021] In another aspect of the invention, a method of inferring user intent in a search input based on resolving ambiguous portions of the search input includes provid-

ing access to a set of content items. Each of the content items is associated with metadata that describes the corresponding content item. The method also includes receiving search input from the user. The search input is intended by the user to identify at least one desired content item.

5 The method also includes determining whether or not a portion of the search input contains an ambiguous identifier. The ambiguous identifier being intended by the user to identify, at least in part, the at least one desired content item. Upon a condition in which a portion of the search input contains an ambiguous identifier, the method includes: inferring a meaning for the ambiguous identifier based on matching portions of the search input to preferences of the user described by a user preference signature, selecting content items from the set of content items based on comparing the search input and the inferred meaning of the ambiguous identifier with metadata associated with the content item, and, upon a condition in which the search input does not contain an ambiguous identifier, selecting content items from the set of content items based on comparing the search input with metadata associated with the content items.

10 **[0022]** Any of the aspects listed above can be combined with any of the other aspects listed above and/or with the techniques disclosed herein.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

15 **[0023]** For a more complete understanding of various embodiments of the present invention, reference is now made to the following descriptions taken in connection with the accompanying drawings in which:

20 35 Figure 1 illustrates a user interface approach incorporated here for elucidative purposes.
Figure 2 illustrates a user interface approach incorporated here for elucidative purposes.
Figure 3 illustrates a user interface approach incorporated here for elucidative purposes.
Figure 4 illustrates a user interface approach incorporated here for elucidative purposes.
Figure 5 illustrates a user interface approach incorporated here for elucidative purposes.
Figure 6 illustrates an example of a graph that represents entities and relationships between entities.
Figure 7 illustrates an example of a graph that represents entities and relationships between entities.
Figure 8 illustrates an example of a graph that represents entities and relationships between entities.
Figure 9 illustrates an example of a graph that represents entities and relationships between entities.
Figure 10 illustrates an architecture that is an embodiment of the present invention.
Figure 11 illustrates a simplified flowchart of the operation of embodiments of the invention.
Figure 12 illustrates a control flow of the operation of embodiments of the invention.

DETAILED DESCRIPTION

[0024] Preferred embodiments of the invention include methods of and systems for inferring user's intent and satisfying that intent in a conversational exchange. Certain implementations are able to resolve ambiguities in user input, maintain state of intent, entities, and/or attributes associated with the conversational exchange, tailor responses to match user's preferences, infer conversational boundaries that start a new topic (i.e. infer a change of a conversational session), and/or engage in a minimal dialog to understand user intent. The concepts that follow are used to describe embodiments of the invention.

Information repositories

[0025] Information repositories are associated with domains, which are groupings of similar types of information and/or certain types of content items. Certain types of information repositories include entities and relationships between the entities. Each entity / relationship has a type, respectively, from a set of types. Furthermore, associated with each entity / relationship are a set of attributes, which can be captured, in some embodiments, as a defined finite set of name-value fields. The entity / relationship mapping also serves as a set of metadata associated with the content items because the entity / relationship mapping provides information that describes the various content items. In other words, a particular entity will have relationships with other entities, and these "other entities" serve as metadata to the "particular entity". In addition, each entity in the mapping can have attributes assigned to it or to the relationships that connect the entity to other entities in the mapping. Collectively, this makes up the metadata associated with the entities / content items. In general, such information repositories are called *structured information repositories*. Examples of information repositories associated with domains follow below.

[0026] A **media entertainment** domain includes entities, such as, movies, TV-shows, episodes, crew, roles/characters, actors/personalities, athletes, games, teams, leagues and tournaments, sports people, music artists and performers, composers, albums, songs, news personalities, and/or content distributors. These entities have relationships that are captured in the information repository. For example, a movie entity is related via an "acted in" relationship to one or more actor/personality entities. Similarly, a movie entity may be related to an music album entity via an "original sound track" relationship, which in turn may be related to a song entity via a "track in album" relationship. Meanwhile, names, descriptions, schedule information, reviews, ratings, costs, URLs to videos or audios, application or content store handles, scores, etc. may be deemed attribute fields.

[0027] A **personal electronic mail (email)** domain includes entities, such as, emails, email-threads, contacts, senders, recipients, company names, departments/busi-

ness units in the enterprise, email folders, office locations, and/or cities and countries corresponding to office locations. Illustrative examples of relationships include an email entity related to its sender entity (as well as the to, cc, bcc, receivers, and email thread entities.) Meanwhile, relationships between a contact and his or her company, department, office location can exist. In this repository, instances of attribute fields associated with entities include contacts' names, designations, email handles, other contact information, email sent/received timestamp, subject, body, attachments, priority levels, an office's location information, and/or a department's name and description.

[0028] A **travel-related / hotels and sightseeing** domain includes entities, such as, cities, hotels, hotel brands, individual points of interest, categories of points of interest, consumer facing retail chains, car rental sites, and/or car rental companies. Relationships between such entities include location, membership in chains, and/or categories. Furthermore, names, descriptions, keywords, costs, types of service, ratings, reviews, etc. all amount of attribute fields.

[0029] An **electronic commerce** domain includes entities, such as, product items, product categories and subcategories, brands, stores, etc. Relationships between such entities can include compatibility information between product items, a product "sold by" a store, etc. Attribute fields in include descriptions, keywords, reviews, ratings, costs, and/or availability information.

[0030] An **address book** domain includes entities and information such as contact names, electronic mail addresses, telephone numbers, physical addresses, and employer.

[0031] The entities, relationships, and attributes listed herein are illustrative only, and are not intended to be an exhaustive list.

[0032] Embodiments of the present invention may also use repositories that are *not* structured information repositories as described above. For example, the information repository corresponding to network-based documents (e.g., the Internet / World Wide Web) can be considered a relationship web of linked documents (entities). However, in general, no directly applicable type structure can meaningfully describe, in a nontrivial way, all the kinds of entities and relationships and attributes associated with elements of the Internet in the sense of the structured information repositories described above. However, elements such as domain names, internet media types, filenames, filename extension, etc. can be used as entities or attributes with such information.

[0033] For example, consider a corpus consisting of a set of unstructured text documents. In this case, no directly applicable type structure can enumerate a set of entities and relationships that meaningfully describe the document contents. However, application of semantic information extraction processing techniques as a preprocessing step may yield entities and relationships that can partially uncover structure from such a corpus.

Illustrative examples of accessing information repositories under certain embodiments of the present invention

[0034] The following description illustrates examples of information retrieval tasks in the context of structured and unstructured information repositories as described above.

[0035] In some cases, a user is interested in one or more entities of some type - generally called intent type herein - which the user wishes to uncover by specifying only attribute field constraints that the entities must satisfy. Note that sometimes intent may be a (type, attribute) pair when the user wants some attribute of an entity of a certain type. For example, if the user wants the rating of a movie, the intent could be viewed as (type, attribute) = (movie, rating). Such query-constraints are generally called *attribute-only constraints* herein.

[0036] Whenever the user names the entity or specifies enough information to directly match attributes of the desired intent type entity, it is an attribute-only constraint. For example, when the user identifies a movie by name and some additional attribute (e.g., 'Cape Fear' made in the 60s), or when he specifies a subject match for the email he wants to uncover, or when he asks for hotels based on a price range, or when he specifies that he wants a 32GB, black colored iPod touch.

[0037] However, in some cases, a user is interested in one or more entities of the intent type by specifying not only attribute field constraints on the intent type entities but also by specifying attribute field constraints on or naming *other entities* to which the intent type entities are *connected* via relationships in some well-defined way. Such query-constraints are generally called *connection oriented constraints* herein.

[0038] An example of a connection oriented constraint is when the user wants a movie (an intent type) based on specifying two or more actors of the movie or a movie based on an actor and an award the movie won. Another example, in the context of email, is if the user wants emails (intent type) received from certain senders from a particular company in the last seven days. Similarly, a further example is if the user wants to book a hotel room (intent type) to a train station as well as a Starbucks outlet. Yet another example is if the user wants a television set (intent type) made by Samsung that is also compatible with a Nintendo Wii. All of these are instances of connection oriented constraints queries.

[0039] In the above connection-oriented constraint examples, the user explicitly describes or specifies the other entities connected to the intent entities. Such constraints are generally called *explicit connection oriented constraints* and such entities as *explicit entities* herein.

[0040] Meanwhile, other queries contain connection oriented constraints that include unspecified or implicit entities as part of the constraint specification. In such a situation, the user is attempting to identify a piece of information, entity, attribute, etc. that is not known through relationships between the unknown item and items the

user does now. Such constraints are generally called *implicit connection oriented constraints* herein and the unspecified entities are generally called *implicit entities* of the constraint herein.

[0041] For example, the user may wish to identify a movie she is seeking via naming two characters in the movie. However, the user does not recall the name of one of the characters, but she does recall that a particular actor played the character. Thus, in her query, she states one character by name and identifies the unknown character by stating that the character was played by the particular actor.

[0042] In the context of email repository, an example includes a user wanting to get the last email (intent) from an unspecified gentleman from a specified company 'Intel' to whom he was introduced via email (an attribute specifier) last week. In this case, the implicit entity is a contact who can be discovered by examining contacts from 'Intel', via an employee / company relationship, who was a first time common-email-recipient with the user last week.

[0043] Further examples of implicit connection oriented constraints are described in more detail below.

[0044] In the context of connection oriented constraints, it is useful to map entities and relationships of information repositories to nodes and edges of a graph. The motivation for specifically employing the graph model is the observation that relevance, proximity, and relatedness in natural language conversation can be modeled simply by notions such as link-distance and, in some cases, shortest paths and smallest weight trees. During conversation when a user dialog involves other entities related to the actually sought entities, a subroutine addressing information retrieval as a simple graph search problem effectively helps reduce dependence on deep unambiguous comprehension of sentence structure. Such an approach offers system implementation benefits. Even if the user intent calculation is ambiguous or inconclusive, so long as entities have been recognized in the user utterance, a graph-interpretation based treatment of the problem enables a system to respond in a much more intelligible manner than otherwise possible, as set forth in more detail below.

45 Attribute-only constraints

[0045] What follows are examples of information retrieval techniques that enable the user to specify attribute-only constraints. While some of these techniques are known in the art (where specified), the concepts are presented here to illustrate how these basic techniques can be used with the inventive techniques described herein to enhance the user experience and improve the quality of the search results that are returned in response to the user's input.

Examples of attributes-only constraints during information retrieval from a Movie/TV search interface

[0046] Figure 1 shows a search interface 100 for a search engine for movie and television content that is known in the art (i.e., the IMDb search interface). Figure 1 includes a pull-down control 105 that allows the user to expressly select an entity type or attribute. For example, *Title* means intent entity type is Movie or TV Show, *TV Episode* means the intent type is Episode, *Names* means intent type is Personality, *Companies* means the intent type is Company (e.g., Production house or Studio etc.), *Characters* means the intent type is Role. Meanwhile, *Keywords*, *Quotes*, and *Plots* specify attribute fields associated with intent entities of type Movie or TV Show or Episode that are sought to be searched. Meanwhile, the pull-down control 110 allows the user to only specify attributes for entities of type Movie, Episode, or TV Show.

[0047] Figure 2 shows the Advanced Title Search graphical user interface of the IMDB search interface (known in the art) 200. Here, the Title Type choice 205 amounts to selection of intent entity type. Meanwhile, Release Date 210, User Rating 215, and Number of Votes 220 are all attributes of entities of type movies, TV Shows, episodes, etc. If the number of Genres 225 and Title Groups 230 shown here is deemed small enough, then those genres and title groups can be deemed descriptive attributes of entities. So the genre and title groups section here is also a way of specifying attribute constraints. The Title Data 235 section is specifying the constraint corresponding to the data source attribute.

Examples of attributes-only constraints during information retrieval from an electronic-commerce search interface

[0048] Figure 3 illustrates a graphical user interface 300 for an electronic-commerce website's search utility that is known in the art. In previous examples, the user interface allowed users to specify sets of attribute constraints *before* initiating any search in the information repository. Meanwhile, Figure 3 shows the user interface *after* the user has first launched a text-only search query 'car stereo'. Based on features and attributes associated with the specific results returned by the text search engine for the text search query 305, the post-search user interface is constructed by dynamically picking a subset of attributes for this set of search results, which allows the user to specify further attribute constraints for them. As a result, the user is forced to follow the specific flow of first doing a text search or category filtering and then specifying the constraints on further attributes.

[0049] This 'hard-coded' flow - of first search followed by post-search attribute filters - results from a fundamental limitation of this style of graphical user interface because it simply cannot display all of the meaningful attributes up-front without having any idea of the product

the user has in mind. Such an approach is less efficient than the inventive techniques disclosed herein because the user may want to declare some of the attribute constraints he or she has in mind at the beginning of the search. This problem stems, in part, from the fact that even though the number of distinct attributes for each individual product in the database is a finite number, the collective set is typically large enough that a graphical user interface cannot display a sufficient number of the attributes, thereby leading to the hard coded flow.

[0050] Note that the conversational interface embodiments disclosed herein do not suffer from physical spatial limitations. Thus, a user can easily specify any attribute constraint in the first user input

Explicit Connection Oriented Constraints

[0051] What follows are examples of explicit connection oriented constraints employed in information retrieval systems. Graph model terminology of nodes and edges can also be used to describe connection oriented constraints as can the terminology of entities and relationships.

[0052] When using an attribute-only constraints interface, the user only specifies the type and attribute constraints on intent entities. Meanwhile, when using an explicit connected node constraints interface, the user can additionally specify the type and attribute constraints on other nodes connected to the intent nodes via specified kinds of edge connections. One example of an interface known in the art that employs explicit connected node constraints during information retrieval is a Movie/TV information search engine 400 shown in Figure 4.

[0053] Considering that the number of possible death and birth places 405 across all movie and TV personalities is a huge number, birth and death places are treated as nodes rather than attributes in the movie information repository graph. Thus, birth and death place specifications in the graphical user interface 400 are specifications for nodes connected to the intended personality node. The filmography filter 410 in the graphical user interface 400 allows a user to specify the name of a movie or TV show node, etc., which is again another node connected to the intended personality node. The other filters 500, shown in Figure 5, of the graphical user interface are specifiers of the attributes of the intended node.

[0054] In the first part of the graphical user interface 400, a user may specify two movie or TV show nodes when his intent is to get the personalities who collaborated on both these nodes. In the second part of the graphical UI above, a user may specify two personality nodes when his intent is to get movie or TV show nodes corresponding to their collaborations. In both case, the user is specifying connected nodes other than his intended nodes, thereby making this an explicit connected node constraint. However, the interfaces known in the art do not support certain types of explicit connected node constraints (explicit connection oriented constraints), as de-

scribed below.

[0055] Figure 6 illustrates a graph 600 of the nodes (entities) and edges (relationships) analyzed by the inventive techniques disclosed herein to arrive at the desired result when the user seeks a movie based on the fictional character Jack Ryan that stars Sean Connery. The user may provide the query, "What movie has Jack Ryan and stars Sean Connery?" The techniques herein interpret the query, in view of the structured information repositories as: Get the node of type Movie (intent) that is connected by an edge 605 to the explicit node of type Role named 'Jack Ryan' 610 and also connected via an 'Acted In' edge 615 to the explicit node of type Personality named 'Sean Connery' 620. The techniques described herein return the movie 'The Hunt for the Red October' 625 as a result.

[0056] Referring again to Figure 6, assume the user asks, "Who are all of the actors that played the character of Jack Ryan?" The disclosed techniques would interpret the query as:

Get nodes of type Personality (intent) connected by means of an 'Acted-as' edge 630 to the explicit node of type Role named 'Jack Ryan' 610. Embodiments of the inventive systems disclosed herein would return the actors 'Alec Baldwin' 635, 'Harrison Ford' 640, and 'Ben Affleck' 645.

[0057] A further example is a user asking for the name of the movie starring Tom Cruise based on a John Grisham book. Thus, the query becomes: Get the node of type Movie (intent) connected by an 'Acted In' edge to the explicit node of type Personality named Tom Cruise and connected by a 'Writer' edge to the explicit node of type Personality named 'John Grisham'. Embodiments of the inventive systems disclosed herein would return the movie 'The Firm'.

Implicit Connection Oriented Constraints

[0058] The following examples illustrate the implicit connection oriented constraints and implicit entities used for specific information retrieval goals. The first two examples used the terminology of entities and relationships.

[0059] In one example, the user wants the role (intent) played by a specified actor/personality (e.g., Michelle Pfeiffer) in an unspecified movie that is about a specified role (e.g., the character Tony Montana.) In this case, the user's constraint includes an unspecified or implicit entity. The implicit entity is the movie 'Scarface'. Figure 7 illustrates a graph 700 of the entities and relationships analyzed by the techniques disclosed herein to arrive at the desired result. The graph 700 is an illustrative visual representation of a structured information repository. Specifically, the implicit movie entity 'Scarface' 705 is arrived at via a 'Acted In' relationship 710 between the movie entity 'Scarface' 705 and the actor entity 'Michelle

Pfeiffer' 715 and a 'Character In' relationship 720 between the character entity 'Tony Montana' 725 and the movie entity 'Scarface' 705. The role entity 'Elvira Hancock' 730 played by 'Michelle Pfeiffer' is then discovered by the 'Acted by' relationship 735 to 'Michelle Pfeiffer' and the 'Character In' relationship 740 to the movie entity 'Scarface' 705.

[0060] In a further example, suppose that the user wants the movie (intent) starring the specified actor entity Scarlett Johansson and the unspecified actor entity who played the specified role of Obi-Wan Kenobi in a specified movie entity Star Wars. In this case, the implicit entity is the actor entity 'Ewan McGregor' and the resulting entity is the movie 'The Island' starring 'Scarlett Johansson' and 'Ewan McGregor'. Figure 8 illustrates a graph 800 of the entities and relationships analyzed by the techniques disclosed herein to arrive at the desired result. Specifically, the implicit actor entity Ewan McGregor 805 is arrived at via an Acted In relationship 810 with at least one movie entity Star Wars 815 and via a Character relationship 820 to a character entity Obi-Wan Kenobi 825, which in turn is related via a Character relationship 830 to the movie entity Star Wars 815. Meanwhile, the result entity The Island 835 is arrived at via an Acted In relationship 840 between the actor/personality entity Scarlett Johansson 845 and the movie entity The Island 835 and an Acted In relationship 850 between the implicit actor entity Ewan McGregor 805 and the movie entity The Island.

[0061] Figure 9 illustrates a graph 900 of the entities and relationships analyzed by the techniques disclosed herein to arrive at a desired result. This example uses the terminology of nodes and edges. The user knows that there is a band that covered a Led Zeppelin song for a new movie starring Daniel Craig. The user recalls neither the name of the covered song nor the name of the movie, but he wants to explore the other music (i.e., songs) of the band that did that Led Zeppelin cover. Thus, by specifying the known entities of Led Zeppelin (as the song composer) and Daniel Craig (as an actor in the movie), the interposing implied nodes are discovered to find the user's desired result. Thus, embodiments of the inventive techniques herein compose the query constraint as follows: Return the nodes of type Song (intent) connected by a 'Composer' edge 905 to an implicit node of type Band 910 (Trent Reznor) such that this Band node has a 'Cover Performer' edge 915 with an implicit node of type Song 920 (Immigrant Song) which in turn has a 'Composer' edge 925 with an explicit node of type Band named 'Led Zeppelin' 930 and also a 'Track in Album' edge 935 with an implicit node of type Album 940 (Girl with the Dragon Tattoo Original Sound Track) which has an 'Original Sound Track (OST)' edge 945 with an implicit node of type Movie 950 (Girl with the Dragon Tattoo Original Sound Track) that has an 'Acted In' edge 955 with the explicit node of type Personality named 'Daniel Craig'. 960.

[0062] As mentioned above, known techniques and

systems for information retrieval suffer from a variety of problems. Described herein are embodiments of an inventive conversational interaction interface. These embodiments enable a user to interact with an information retrieval system by posing a query and/or instruction by speaking to it and, optionally, selecting options by physical interaction (e.g., touching interface, keypad, keyboard, and/or mouse). Response to a user query may be performed by machine generated spoken text to speech and may be supplemented by information displayed on a user screen. Embodiments of the conversation interaction interface, in general, allow a user to pose his next information retrieval query or instruction in reaction to the information retrieval system's response to a previous query, so that an information retrieval session is a sequence of operations, each of which has the user first posing a query or instruction and the system the presenting a response to the user.

[0063] Embodiments of the present invention are a more powerful and expressive paradigm than graphical user interfaces for the query-constraints discussed herein. In many situations, especially when it comes to flexibly selecting from among a large number of possible attributes or the presence of explicit and implicit connected nodes, the graphical user interface approach does not work well or does not work at all. In such cases, embodiments of the conversational interaction interface of the present invention are a much more natural fit. Further, embodiments of the present invention are more scalable in terms of the number of distinct attributes a user may specify as well as the number of explicit connected node constraints and the number of implicit node constraints relative to graphical user interfaces.

Conversational System Architecture

[0064] Figure 10 represents the overall system architecture 1000 of an embodiment of the present invention. User 1001 speaks his or her question that is fed to a speech to text engine 1002. While the input could be speech, the embodiment does not preclude the input to be direct text input. The text form of the user input is fed to session dialog content module 1003. This module maintains state across a conversation session, a key use of which is to help in understanding user intent during a conversation, as described below.

[0065] The session dialog content module 1003, in conjunction with a Language Analyzer 1006, a Domain Specific Named Entity Recognizer 1007, a Domain Specific Context and Intent Analyzer 1008, a Personalization Based Intent Analyzer 1009, a Domain Specific Graph Engine 1010, and an Application Specific Attribute Search Engine 1011 (all described in more detail below) process the user input so as to return criteria to a Query Execution Engine 1004. The Query Execution Engine 1004 uses the criteria to perform a search of any available source of information and content to return a result set.

[0066] A Response Transcoding Engine 1005, dis-

patches the result set to the user for consumption, e.g., in the device through which user is interacting. If the device is a tablet device with no display constraints, embodiments of the present invention may leverage off the display to show a graphical rendition of connection similar in spirit to Figures 7, 6, 9, and 8, with which the user can interact with to express intent. In a display-constrained device such as a smartphone, the Response Transcoding Engine 1005 may respond with text and/or speech (using a standard text to speech engine).

[0067] While Figure 10 is a conversation architecture showing the modules for a specific domain, the present embodiment is a conversation interface that can take user input and engage in a dialog where user's intent can span domains. In an embodiment of the invention, this is accomplished by having multiple instances of the domain specific architecture shown in Figure 10, and scoring the intent weights across domains to determine user intent. This scoring mechanism is also used to implicitly determine conversation topic switching (for example, during an entertainment information retrieval session, a user could just say "I am hungry").

[0068] Figure 11 illustrates a simplified flowchart of the operation of embodiments of the invention. First, the user's speech input is converted to text by a speech recognition engine 1101. The input is then broken down into intent, entities, and attributes 1102. This process is assisted by information from the prior conversation state 1103. The breakdown into intents, entities, and attributes, enables the system to generate a response to the user 1104. Also, the conversation state 1103 is updated to reflect the modifications of the current user input and any relevant returned response information.

[0069] Figure 12 illustrates the control flow in more detail. First, the user's speech is input to the process as text 1201. Upon receiving the user input as text, query execution coordination occurs 1202. The query execution coordination 1202 oversees the breakdown of the user input to understand user's input. The query execution coordination 1202 makes use of language analysis 1203 that parses the user input and generates a parse tree. The query execution coordination 1202 also makes use of the maintenance and updating of the dialog state 1208. The parse tree and any relevant dialog state values are passed to modules that perform intent analysis 1204, entity analysis 1205, and attribute analysis 1206. These analysis processes work concurrently, because sequential processing of these three analysis steps may not be possible. For instance, in some cases of user input, the recognition of entities may require the recognition of intents and vice versa. These mutual dependencies can only be resolved by multiple passes on the input by the relevant modules, until the input is completely analyzed. Once the breakdown and analysis is complete, a response to the user is generated 1207. The dialog state is also updated 1208 to reflect the modifications of the current input and return of relevant results. In other words, certain linguistic elements (e.g., spoken / recog-

nized words and/or phrases) are associated with the present conversation session.

[0070] Referring again to Figure 10, in one illustrative embodiment, the Session Dialog Content Module 1003, in conjunction with a Language Analyzer 1006, and the other recognizer module, analyzer modules, and/or engines described in more detail below, perform the analysis steps mentioned in connection with Figure 12 and break down the sentence into its constituent parts. The Language Analyzer 1006 creates a parse tree from the text generated from the user input, and the other recognizer module, analyzer modules, and/or engines operate on the parse tree to determine the constituent parts. Those parts can be broadly categorized as (1) intents - the actual intent of the user (such as "find a movie", "play a song", "tune to a channel", "respond to an email", etc.), (2) entities - noun or pronoun phrases describing or associated with the intent, and (3) attributes - qualifiers to entities such as the "latest" movie, "less" violence, etc. Other constituent part categories are within the scope of the invention.

[0071] In the context of the goal of providing an intelligent and meaningful conversation, the intent is among the most important of all three categories. Any good search engine can perform an information retrieval task fairly well just by extracting the entities from a sentence - without understanding the grammar or the intent. For instance, the following user question, "Can my daughter watch pulp fiction with me" - most search engines would show a link for pulp fiction, which may suffice to find the rating that is most likely available from traversing that link. But in a conversational interface, the expectation is clearly higher - the system must ideally understand the (movie, rating) intent corresponding to the expected response of the rating of the movie and the age group it is appropriate for. A conversational interface response degenerating to that of a search engine is tantamount to a failure of the system from a user perspective. Intent determination and, even more importantly, responding to user's question that appears closer to a human's response when the intent is not known or clearly discernible is an important aspect for a conversational interface that strives to be closer to human interaction than to a search engine.

[0072] In this example, although the user never used the word "rating", the system infers that user is looking for rating, from the words "can ... watch" based on a set of rules and/or a naive Bayes classifier, described in more details below. Meanwhile, "my daughter" could be recognized as an attribute. In order for the daughter to watch a program, several criteria must be met: the show timing, the show availability, and "watchability" or rating. This condition may be triggered by other attributes too such as "son", "girl", "boy" etc. These could be rules-based domain specific intents or naive Bayes classifier scoring based on domain specific training sets to look for ratings and show timings in this case. There could also be weightage factor for the satisfiability of these conditions

that is driven by the entity that is being watched.

[0073] Intent Analyzer 1008 is a domain specific module that analyzes and classifies intent for a domain and works in conjunction with other modules - domain specific entity recognizer 1007, personalization based intent analyzer 1009 that classifies intent based on user's personal preferences, and the domain specific graph engine 1010. The attribute specific search engine 1011 assists in recognizing attributes and their weights influence the entities they qualify.

[0074] The intent analyzer 1008, in an embodiment of the invention, is a rules driven intent recognizer and/or a naive Bayes classifier with supervised training. The rules and/or training set capture how various words and word sets relate to user intent. It takes as input a parse tree, entity recognizer output, and attribute specific search engine output (discussed above and below). In some implementations, user input may go through multiple entity recognition, the attribute recognition, and intent recognition steps, until the input is fully resolved. The intent recognizer deciphers the intent of a sentence, and also deciphers the differences in nuances of intent. For instance, given "I would like to see the movie Top Gun" versus "I would like to see a movie like Top Gun", the parse trees would be different. This difference assists the intent recognizer to differentiate the meaning of "like". The rules based recognition, as the very name implies, recognizes sentences based on predefined rules. Predefined rules are specific to a domain space, for example, entertainment. The naive Bayes classifier component, however, just requires a training data set to recognize intent.

[0075] The entity recognizer 1007, using the inputs mentioned above, recognizes entities in user input. Examples of entities are "Tom cruise" in "can I watch a Tom Cruise movie", or "Where Eagles Dare" in "when was Where Eagles Dare released". In certain implementations, the entity recognizer 1007 can be rules driven and/or a Bayes classifier. For example, linguistic elements such as nouns and gerunds can be designated as entities in a set of rules, or that association can arise during a supervised training process for the Bayes classifier. Entity recognition can, optionally, involve error correction or compensation for errors in user input (such as errors in speech to text recognition). When an input matches two entities phonetically, e.g., newman, and neuman, both are picked as likely candidates. In some embodiments, the resolution between these two comes from the information gleaned from the rest of user input, where relationships between entities may weed out one of the possibilities. The classifying of a subset of user input as an entity is only a weighting. There could be scenarios in which an input could be scored as both an entity and as an attribute. These ambiguities are resolved in many cases as the sentence semantics become clearer with subsequent processing of the user input. In certain embodiments, a component used for resolution is the entity relationship graph. In certain implementations, an output of the entity recognizer 1007 is a probability score

for subsets of input to be entities.

[0076] The application specific attribute search engine 1011 recognizes attributes such as "latest", "recent", "like" etc. Here again, there could be conflicts with entities. For example "Tomorrow Never Dies" is an entity (a movie), and, when used in a sentence, there could be an ambiguity in interpreting "tomorrow" as an attribute. The scoring of tomorrow as an attribute may be lower than the scoring of "tomorrow" as part of "Tomorrow Never Dies" as determined by entity relationship graph (which may depend on other elements of the input, e.g., the words "movie", "show", "actor", etc.). The output of the attribute search engine 1011 is a probability score for input words similar to that of the output of entity recognizer 1007.

[0077] The language analyzer 1006 is a pluggable module in the architecture to enable the system to support multiple languages. While understanding the semantics of user input is not constrained to the language analyzer 1006 alone, the core modules of the architecture such as dialog context module 1003 or graph engine 1010 are language independent. As mentioned earlier, the language module alone cannot do much more than analysis of a sentence and performing tasks such as relating a pronoun to its subject/object etc. ("The Led Zeppelin song in the OST of the recent Daniel Craig movie... Who performed it?"), it is ineffective in isolation to associate pronouns across exchanges. It is the interaction with the session dialog context module 1003, that enables resolution of pronouns across exchanges as in the following:

Q1: Who acts as obi-wan Kenobi in the new star wars?
 A1: Ewan McGregor
 Q2: How about his movies with Scarlet Johansson

[0078] While it may seem, at first glance, that dialog session context is a simple state maintenance of the last active entity or entities, the following examples show the lurking complexity in dialog session context:

Q1: Who played the lead roles in Kramer vs. Kramer?
 A1: Meryl Streep and Dustin Hoffman
 Q2: How about more of his movies
 A2: Here are some of Dustin Hoffman movies... [list of Dustin Hoffman movies]
 Q3: What about more of her movies?
 A3: [list of movies if any]
 Q4: What about just his early movies?
 A4: [list of movies if any]
 Q5: What about her recent movies?
 A5: [list of movies if any]
 Q6: Have they both acted again in the recent past?
 A6: [list of movies if any]
 Q7: Have they both ever acted again at all?

[0079] In the example above, the entities Meryl Streep and Dustin Hoffman are indirectly referred to in six ques-

tions, sometimes together and sometimes separately. The above example also illustrates a distinction of embodiments of the present invention from simple request response systems that engage in an exploratory exchange around a central theme. While the present embodiments not only resolve ambiguities in an exchange, they simultaneously facilitate an exploratory exchange with implicit references to entities and/or intents mentioned much earlier in a conversation - something that is naturally done in rich human interactions. In certain embodiments, this is done through the recognition of linguistic linking elements, which are words and/or phrases that link the present user input to a previous user input and/or system response. Referring to the example provided above, the pronouns "his", "hers", and "they" are words that link the present user input to a previous user input and/or system response. Other pronouns, as well as syntactic expletives, can act as linguistic linking elements.

[0080] Whether a particular word or phrase used by the user in a later question is a suitable or appropriate link to an entity mentioned in an earlier input (or some other part of the earlier input) is determined by examining the attributes of the earlier entity and the attributes of the potential linking element. For example, "his" is a suitable link to Dustin Hoffman in the example above because Dustin Hoffman is male, and "his" is a male gender pronoun. Moreover, "his" is a possessive pronoun, which is appropriate because the user is referring to movies in which Dustin Hoffman appears.

[0081] The following example illustrates user referring to an entity who is not even explicitly specified in a prior exchange (an implicit connection oriented constraint). In other words, the user never identifies the entity that is saved as part of the dialog state.

Q1: Which show had that physically challenged scientist alluding to the possibility of non-carbon based life form?
 A1: That was Stephen Hawking's discovery channel program on aliens.
 Q2: Was he in another show that David Attenborough produced?

[0082] In another example of dialog state maintenance not being restricted to just entities and their attributes, the system can maintain state of intents across conversation exchanges, as shown as follows:

Q1: "Can my daughter watch pulp fiction with me?"
 A1: Pulp fiction by Quentin Tarantino is rated R for graphic violence and nudity.
 Q2: What about his other movies?
 A2: They are all rated R - Reservoir Dogs, Jackie Brown, Kill Bill, Death Proof.

[0083] In this example in addition to maintain state of the entity "Quentin Tarantino" which enables the system to understand the pronoun reference to him (in the form

of "his") in Q2, the system also keeps track of user intent across the exchanges - the user intent being the "rating". Again, the system's decision to maintain both "Quentin Tarantino" and the "rating" intent stems from the rules and/or Bayes classifier training sets. Thus, the techniques disclosed herein enable the preservation and use of multiple intents. In such an implementation, the set intents would be passed as a collection of intents with weights. Depending on the output of the rules and/or Bayes classifier, the system may elect to save all intents during a session (and/or entities, attributes, etc.), but may only use the one intent that is scored highest for a particular input. Thus, it is possible that an intent that accrued relatively earlier in the dialog exchange applies much later in the conversation. Maintaining the state in this way facilitates a succinct and directed response as in A2, almost matching a human interaction.

[0084] The directed responses illustrated above are possible with the intent analyzer 1008 and entity recognizer 1009 working in close concert with the personalization based intent analyzer 1009. These modules are all assisted by an application specific attribute search engine 1011 that assists in determining relevant attributes (e.g., latest, less of violence, more of action) and assigns weights to them. So a user input exchange would come from the speech to text engine 1002, would be processed by the modules, analyzers, and engines working in concert (with the query execution engine 1004 playing a coordinating role), and would yield one or more candidate interpretations of the user input. For instance the question, "Do you have the Kay Kay Menon movie about the Bombay bomb blasts?", the system may have two alternative candidate representations wherein one has "Bombay" as an entity (there is a movie called Bombay) with "bomb blast" being another and the other has "Bombay bomb blast" as a single entity in another. The system then attempts to resolve between these candidate representations by engaging in a dialog with the user, on the basis of the presence of the other recognized entity Kay Kay Menon who is an actor. In such a case, the question(s) to formulate depend on the ambiguity that arises. In this example, the actor entity is known, it is the associated movie entities that are ambiguous. Thus, the system would ask questions concerning the movie entities. The system has a set of forms that are used as a model to form questions to resolve the ambiguities.

[0085] In some instances, resolution of ambiguity can be done, without engaging in a dialog, by knowing user's preferences. For instance, the user may ask "Is there a sox game tonight?" While this question has an ambiguous portion - the ambiguity of the team being the Boston Red Sox or the Chicago White Sox - if the system is aware that user's preference is Red Sox, then the response can be directed to displaying a Red Sox game schedule if there is one that night. In instances where there are multiple matches across domains, the domain match resulting in the higher overall confidence score will win. Personalization of results can also be done,

when applicable, based on the nature of the query. For instance, if the user states "show me movies of Tom Cruise tonight", this query should not apply personalization but just return latest movies of Tom Cruise. However

5 if user states "show me sports tonight", system should apply personalization and display sports and games that are known to be of interest to the user based on his explicit preferences or implicit actions captured from various sources of user activity information.

10 **[0086]** A user preference signature can be provided by the system using known techniques for discovering and storing such user preference information. For example, the methods and systems set forth in U.S. Patent No. 7,774,294, entitled *Methods and Systems for Selecting and Presenting Content Based on Learned Periodicity of User Content Selections*, issued August 10, 2010, U.S. Patent No. 7,835,998, entitled *Methods and Systems for Selecting and Presenting Content on a First System Based on User Preferences Learned on a Second System*, issued November 16, 2010, U.S. Patent No.

15 7,461,061, entitled *User Interface Methods and Systems for Selecting and Presenting Content Based on User Navigation and Selection Actions Associated with the Content*, issued December 2, 2008, and U.S. Patent No. 8, 112,454, entitled *Methods and Systems for Ordering Content Items According to Learned User Preferences*, issued February 7, 2012, each of which is incorporated by reference herein, can be used with the techniques disclosed herein. However, the use of user's preference

20 signatures and/or information is not limited to the techniques set forth in the incorporated applications.

25 **[0087]** The relationship or connection engine 1010 is one of the modules that plays a role in comprehending user input to offer a directed response. The relationship engine could be implemented in many ways, a graph data structure being one instance so that we may call the relationship engine by the name graph engine. The graph engine evaluates the user input in the backdrop of known weighted connections between entities.

30 **[0088]** One embodiment showing the importance of the graph engine is illustrated by the following example in which user intent is clearly known. If the user simply queries 'what is the role played by Michelle Pfeiffer in the Tony Montana movie', the system knows the user intent

35 (the word role and its usage in the sentence may be used to deduce that the user wants to know the character that Michelle Pfeiffer has played somewhere) and has to grapple with the fact that the named entity Tony Montana could be the actor named Tony Montana or the name of

40 the leading character of the movie Scarface. The graph engine in this instance is trivially able to disambiguate since a quick analysis of the path between the two Tony Montana entities respectively and the entity of Michelle Pfeiffer quickly reveals that the actor Tony Montana never

45 collaborated with Michelle Pfeiffer, whereas the movie Scarface (about the character Tony Montana) starred Michelle Pfeiffer. Thus, the system will conclude that it can safely ignore the actor Tony Montana and that the

user wants to know the name of the character played by Michelle Pfeiffer in the movie Scarface.

[0089] In another embodiment, the graph engine 1010 assists when the system is unable to determine the user intent despite the fact that the entity recognizer 1007 has computed the entities specified by the user. This is illustrated by the following examples in which the user intent cannot be inferred or when the confidence score of the user intent is below a threshold. In such a scenario, two illustrative strategies could be taken by a conversation system to get the user's specific intent. In some embodiments, the system determined the most important keywords from the user utterance and treats each result candidate as a document, calculates a relevance score of each document based on the each keyword's relevance, and presents the top few documents to the user for him to peruse. This approach is similar to the web search engines. In other embodiments, the system admits to the user that it cannot process the user request or that the information it gathered is insufficient, thereby prompting the user to provide more information or a subsequent query.

[0090] However, neither approach is entirely satisfactory when one considers the response from the user's perspective. The first strategy, which does blind keyword matches, can often look completely mechanical. The second approach attempts to be human-like when it requests the user in a human-like manner to furnish more information to make up for the fact that it could not compute the specific user-intent. However, in the cases that the user clearly specifies one or more other entities related to the desired user intent, the system looks incapable if the system appear to not attempt an answer using the clearly specified entities in the user utterance.

[0091] In certain implementations, a third strategy is employed so long as entity recognition has succeeded (even when the specific user intent calculation has failed). Note that entity recognition computation is successful in a large number of cases, especially when the user names or gives very good clues as to the entities in his utterance, which is usually the case.

[0092] The strategy is as follows:

1. Consider the entity relationship graph corresponding to the information repository in question. Entities are nodes and relationships are edges in this graph. This mapping involving entities / nodes and relationships / edges can involve one-to-one, one-to-many, and many-to-many mapping based on information and metadata associated with the entities being mapped.
2. Entities / nodes have types from a finite and well-defined set of types.
3. Since entity recognition is successful (e.g., from an earlier interaction), we consider the following cases:

a. Number of presently recognized entities is 0:

In this case, the system gives one from a fixed set of responses based on response templates using the information from the user that is recognized. The template selections is based on rules and/or Bayes classifier determinations.

b. Number of recognized entities is 1: Suppose that the entity identifier is A and the type of the entity is B and we know the finite set S of all the distinct edge/relationship types A can be involved in. In this case, a system employing the techniques set forth herein ("the IR system") speaks and/or displays a *human-consumption multi-modal response template* $T(A, B, S)$ that follows from an applicable template response based on A, B and S. The template response is selected from a set of manually constructed template responses based on *a priori* knowledge of all possible node types and edge types, which form finite well-defined sets. The response and IR system is designed to allow the user to select, using a touch interface or even vocally, more information and entities related to A.

c. Number of recognized edge types is 2: In this case, let the two entity nodes respectively have identifiers A, A', types B, B' and have edge-type sets S, S'.

If the edge distance between the two entity nodes is greater than some previously decided threshold k , then the IR system appropriately employs and delivers (via speech and/or display) the corresponding two independent human-consumption multi-modal response templates $T(A, B, S)$ and $T(A', B', S')$.

If the edge distance is no more than k , then the IR system selects a shortest edge length path between A and A'. If there are clues available in user utterance, the IR system may prefer some shortest paths to others. Let there be k' nodes in the selected shortest path denoted $A=A_1, A_2, A_3, \dots, A_k=A'$ where $k' < k+1$ and for each i , where i goes from 1 to k' , the i th entity node of the path is represented by the 3-tuple A_i, B_i, E_i where A_i is the entity identifier, B_i is the entity type and E_i is a list of one or two elements corresponding to the one or two edges connected to A_i that are present in the selected shortest path. In this case the IR system then delivers to the user an appropriate response based on an intelligent composition of the sequence of human-consumption multi-modal response templates $T(A_i, B_i, E_i)$ where i goes from 1 to k' .

d. Number of recognized edge types is $R \geq 3$: In this case the IR system simply calculates K maximal components C_1, C_2, \dots, C_k where each component C_i is such each entity node A in C_i is at a distance of no more than k edges away from at least one other node A' of C_i . For each C_i the IR System selects an appropriate repre-

sentative sequence of human-consumption multi-modal response template sequences, similar to c. above and composes a response based on the response template sequences for each component

[0093] This method to generate a response is suggested to be more human in that it has the ability to demonstrate to the user that, with the help of the entities recognized, it presented to the user a response which made it potentially easier compared to the two earlier strategies vis-à-vis his goal of retrieving information. Figures 7, 6, 9, and 8, illustrate examples implementations of the disclosed techniques.

[0094] The techniques set forth above are also used, in certain implementations, to reset all or part of the conversation state values. For example, assume a system has retained certain entities and/or attributes from user input and system responses. When the user provides subsequent input, the techniques disclosed herein enable the new input to be evaluated against the retained values. Speaking in terms of a graph model, if linguistic elements of the subsequent input are found in an entity / relationship graph to be too far removed from the retained information (also in the graph), it can be inferred that the user's subsequent intent has changed from the previous one. In such a case, the earlier retained information can be reset and/or disregarded when performing the subsequent search.

[0095] Further still, embodiments of the invention can recognize that a user has provided subsequent input that lacks entities, attributes, or relationship information, but the input is an appropriate response to an earlier system response. For example, a system implementing the techniques set forth herein may present a set of movies as a response to a first user input. The user may then respond that she is not interested in any of the movies presented. In such a case, the system would retain the various conversation state values and make a further attempt to satisfy the user's previous request (by, e.g., requesting additional information about the type of movie desired or requesting additional information to better focus the search, such as actor names, genre, etc.).

[0096] In the foregoing description, certain steps or processes can be performed on particular servers or as part of a particular engine. These descriptions are merely illustrative, as the specific steps can be performed on various hardware devices, including, but not limited to, server systems and/or mobile devices. Similarly, the division of where the particular steps are performed can vary, it being understood that no division or a different division is within the scope of the invention. Moreover, the use of "analyzer", "module", "engine", and/or other terms used to describe computer system processing is intended to be interchangeable and to represent logic or circuitry in which the functionality can be executed.

[0097] The techniques and systems disclosed herein may be implemented as a computer program product for

use with a computer system or computerized electronic device. Such implementations may include a series of computer instructions, or logic, fixed either on a tangible medium, such as a computer readable medium (e.g., a diskette, CD-ROM, ROM, flash memory or other memory or fixed disk) or transmittable to a computer system or a device, via a modem or other interface device, such as a communications adapter connected to a network over a medium.

5 **[0098]** The medium may be either a tangible medium (e.g., optical or analog communications lines) or a medium implemented with wireless techniques (e.g., Wi-Fi, cellular, microwave, infrared or other transmission techniques). The series of computer instructions embodies at least part of the functionality described herein with respect to the system. Those skilled in the art should appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems.

10 **[0099]** Furthermore, such instructions may be stored in any tangible memory device, such as semiconductor, magnetic, optical or other memory devices, and may be transmitted using any communications technology, such as optical, infrared, microwave, or other transmission technologies.

15 **[0100]** It is expected that such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation (e.g., shrink wrapped software), preloaded with a computer system (e.g., on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the network (e.g., the Internet or World Wide Web). Of course, some embodiments of the invention may be implemented as a combination of both software (e.g., a computer program product) and hardware. Still other embodiments of the invention are implemented as entirely hardware, or entirely software (e.g., a computer program product).

20 **[0101]** The present disclosure includes at least the following:

25 Item 1. A method of inferring a change of a conversation session during continuous user interaction with an interactive content providing system having a processor, the method comprising:

30 providing access to a set of content items, each of the content items having associated metadata stored in an electronically readable medium that describes the corresponding content item; receiving at the processor a first input from a user, the first input including linguistic elements intended by the user to identify at least one desired content item; 35 associating by the processor at least one linguistic element of the first input with a first conversation session; 40 providing by the processor a first response based on the first input and based on the meta-

45

50

55

<p>data associated with the content items; receiving at the processor a second input from the user; inferring by the processor whether or not the second input from the user is related to the at least one linguistic element associated with the first conversation session; upon a condition in which the second input is inferred to relate to the at least one linguistic element associated with the first conversation session, providing by the processor a second response based on the metadata associated with the content items, the second input, and the at least one linguistic element of the first input associated with the first conversation session; and upon a condition in which the second input is inferred to not relate to the at least one linguistic element associated with the first conversation session, providing by the processor a second response based on the metadata associated with the content items and the second input.</p>	5 10 15 20	<p>latedness between (i) linguistic elements of the second input and (ii) the at least one linguistic element associated with the first conversation session based on the metadata associated with the content items; upon a condition in which the measure of relatedness is equal to or greater than a threshold value, the inferring by the processor concludes the second input is related to the at least one linguistic element associated with the first conversation session; and upon a condition in which the measure of relatedness is less than the threshold value, the inferring by the processor concludes the second input is not related to the at least one linguistic element associated with the first conversation session.</p>
<p>Item 2. The method of item 1, wherein the inferring by the processor whether or not the second input from the user is related to the at least one linguistic element associated with the first conversation session is based on a determination of whether or not the second input contains a linguistic linking element.</p>	25	<p>Item 9. The method of item 8, the metadata associated with the content items including a mapping of relationships between entities associated with the content items, and the determining by the processor the measure of relatedness including analyzing by the processor the mapping of relationships.</p>
<p>Item 3. The method of item 2, wherein the linguistic linking element is at least one of a pronoun and a syntactic expletive.</p>	30	<p>Item 10. The method of item 1, wherein the inferring by the processor whether or not the second input from the user is related to the at least one linguistic element associated with the first conversation session includes:</p>
<p>Item 4. The method of item 2, wherein the inferring by the processor concludes the second input from the user is related upon a condition in which the second input contains a linguistic linking element,</p>	35	<p>identifying by the processor a linguistic element associated with the first conversation session that identifies at least one entity; identifying by the processor a linguistic linking element of the second input; determining by the processor whether or not the linguistic linking element of the second input is a suitable link to the linguistic element associated with the first conversation session that identifies at least one entity; upon a condition in which the linguistic linking element of the second input is a suitable link, concluding by the processor the second input from the user is related to the at least one linguistic element associated with the first conversation session; and upon a condition in which the linguistic linking element of the second input is not a suitable link, concluding by the processor the second input from the user is not related to the at least one linguistic element associated with the first conversation session.</p>
<p>Item 5. The method of item 4, wherein the linguistic linking element is at least one of a pronoun and a syntactic expletive.</p>	40	<p>Item 11. The method of item 10, wherein the linguistic element associated with the first conversation session that identifies at least one entity is at least one</p>
<p>Item 6. The method of item 2, wherein the inferring by the processor concludes the second input from the user is not related upon a condition in which the second input does not contain a linguistic linking element.</p>	45	
<p>Item 7. The method of item 6, wherein the linguistic linking element is at least one of a pronoun and a syntactic expletive.</p>	50	
<p>Item 8. The method of item 2, further comprising, upon a condition in which the second input does not contain a linguistic linking element, the following:</p>	55	
<p>determining by the processor a measure of re-</p>		

of a noun and a gerund. Item 12. The method of item 10, wherein the linguistic linking element is at least one of a pronoun and a syntactic expletive.	5	Item 17. The method of item 16, wherein, upon the condition in which the second input is inferred to relate to the at least one linguistic element associated with the first conversation session, the providing by the processor the second response includes substituting the at least one linguistic element associated with the first conversation session in place of at least one linguistic element of the second input,
Item 13. The method of item 1, wherein the inferring by the processor whether or not the second input from the user is related to the at least one linguistic element associated with the first conversation session includes:	10	Item 18. The method of item 1, the metadata associated with the content items including a mapping of relationships between entities associated with the content items.
determining by the processor whether or not the second input includes a linguistic element that identifies at least one entity; upon a condition in which the second input does not include a linguistic element that identifies at least one entity, determining by the processor whether or not the second input is an appropriate response to the first response; upon a condition in which the second input is an appropriate response to the first response, concluding by the processor the second input from the user is related to the at least one linguistic element associated with the first conversation session; and upon a condition in which the second input is not an appropriate response to the first response, concluding by the processor the second input from the user is not related to the at least one linguistic element associated with the first conversation session.	15 20 25 30	Item 19. The method of item 1, further comprising: determining by the processor that a portion of at least one of the first input from the user and the second input from the user contains an ambiguous identifier, the ambiguous identifier intended by the user to identify, at least in part, the at least one desired content item; and inferring by the processor a meaning for the ambiguous identifier based on matching portions of the at least one of the first input from the user and the second input from the user to preferences of the user described by a user preference signature; the providing by the processor a second response including selecting content items from the set of content items based on comparing the inferred meaning of the ambiguous identifier with the metadata associated with the content items.
Item 14. The method of item 13, wherein the linguistic element included in the second input that identifies at least one entity is at least one of a noun and a gerund.	35	Item 20. A method of inferring user intent in a search input based on resolving ambiguous portions of the search input, the method comprising:
Item 15. The method of item 1, wherein, upon the condition in which the second input is inferred to relate to the at least one linguistic element associated with the first conversation session, the providing by the processor the second response includes substituting by the processor the at least one linguistic element associated with the first conversation session in place of at least one linguistic element of the second input.	40 45	providing access to a set of content items, each of the content items being associated with metadata that describes the corresponding content item, the metadata associated with the content items including a mapping of relationships between entities associated with the content items; receiving search input from the user, the search input being intended by the user to identify at least one desired content item; determining that a portion of the search input contains at least one specified entity and a reference to at least one unspecified entity related to the at least one desired content item; inferring a possible meaning for the at least one unspecified entity based on the at least one specified entity and the mapping of relationships between entities; and selecting at least one content item from the set of content items based on comparing the possi-
Item 16. The method of item 1, further comprising:		
associating by the processor at least one linguistic element of the first response with the first conversation session; and upon a condition in which the second input is inferred to relate to the at least one linguistic element associated with the first conversation session, further basing the second response on the at least one linguistic element of the first response associated with the first conversation session.	50 55	

ble meaning for the at least one unspecified entity with the metadata associated with the content items of the set of content items.		items and (ii) metadata associated with the content items; and ordering the at least one content item based on the preferences of the user described by the user preference signature.
Item 21. The method of item 20, the content items including media entertainment items and the metadata including at least one of crew, characters, actors, teams, leagues, tournaments, athletes, composers, music artists, performers, albums, songs, news personalities, and content distributors. 5	10	Item 29. A method of inferring user intent in a search input based on resolving ambiguous portions of the search input, the method comprising:
Item 22. The method of item 20, the content items including electronic mail items and the metadata including at least one of electronic mail threads, contacts, senders, recipients, company names, business departments, business units, electronic mail folders, and office location information. 15	15	providing access to a set of content items, each of the content items being associated with metadata that describes the corresponding content item, the metadata associated with the content items including a mapping of relationships between entities associated with the content items; receiving search input from the user, the search input being intended by the user to identify at least one desired content item; determining whether or not a portion of the search input contains at least one specified entity and a reference to at least one unspecified entity related to the at least one desired content item;
Item 23. The method of item 20, the content items including travel-related items and the metadata including at least one of cities, hotels, hotel brands, individual points of interest, categories of points of interest, retail chains, car rental websites, and car rental company names. 20	20	upon a condition in which a portion of the search input contains at least one unspecified entity:
Item 24. The method of item 20, the content items including electronic commerce items and the metadata including at least one of product items, product categories, product subcategories, product brand, and retail stores. 25	25	selecting at least one content item from the set of content items based on inferring a possible meaning for the at least one unspecified entity based on the at least one specified entity and the mapping of relationships between entities and comparing the possible meaning for the at least one unspecified entity with metadata associated with the content items of the set of content items; and
Item 25. The method of item 20, the content items including network-based documents and the metadata including at least one of domain names, internet media types, filenames, 30	35	upon a condition in which the search input does not contain at least one unspecified entity, selecting at least one content item from the set of content items based on comparing the search input with metadata associated with the content items.
Item 26. The method of item 20, the content items including address book items and the metadata including at least one of contact names, electronic mail address, telephone number, address, and employer. 40	40	Item 30. The method of item 29, the content items including media entertainment items and the metadata including at least one of crew, characters, actors, teams, leagues, tournaments, athletes, composers, music artists, performers, albums, songs, news personalities, and content distributors.
Item 27. The method of item 20, further comprising providing a user preference 45	45	Item 31. The method of item 29, the content items including electronic mail items and the metadata including at least one of electronic mail threads, contacts, senders, recipients, company names, business departments, business units, electronic mail folders, and office location information.
signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items, wherein the inferring the possible meaning for the at least one unspecified entity is further based on comparing portions of the search input to the preferences of the user described by the user preference signature. 50	50	
Item 28. The method of item 20, further comprising: 55		
providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content		

Item 32. The method of item 29, the content items including travel-related items and the metadata including at least one of cities, hotels, hotel brands, individual points of interest, categories of points of interest, retail chains, car rental websites, and car rental company names.	5	receiving search input from the user, the search input being intended by the user to identify at least one desired content item;
Item 33. The method of item 29, the content items including electronic commerce items and the metadata including at least one of product items, product categories, product subcategories, product brand, and retail stores.	10	determining that a portion of the search input contains an ambiguous identifier, the ambiguous identifier intended by the user to identify, at least in part, the at least one desired content item;
Item 34. The method of item 29, the content items including network-based documents and the metadata including at least one of domain names, internet media types, filenames, directories, and filename extensions.	15	inferring a meaning for the ambiguous identifier based on matching portions of the search input to the preferences of the user described by the user preference signature; and
Item 35. The method of item 29, the content items including address book items and the metadata including at least one of contact names, electronic mail address, telephone number, address, and employer.	20	selecting content items from the set of content items based on comparing the search input and the inferred meaning of the ambiguous identifier with metadata associated with the content items.
Item 36. The method of item 29, further comprising providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items, wherein the inferring the possible meaning for the at least one unspecified entity is further based on comparing portions of the search input to the preferences of the user described by the user preference signature.	25	Item 39. The method of item 38, the ambiguous identifier being a pronoun or syntactic expletive.
Item 37. The method of item 29, further comprising:	35	Item 40. The method of item 38, the ambiguous identifier being an identifier of an entertainment genre,
providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items; and	40	Item 41. The method of item 38, the ambiguous identifier being at least a portion of a name.
ordering the at least one content item based on the preferences of the user described by the user preference signature.	45	Item 42. The method of item 38, the metadata associated with the content items including a mapping of relationships between entities associated with the content items.
Item 38. A method of inferring user intent in a search input based on resolving ambiguous portions of the search input, the method comprising:	50	Item 43. The method of item 38, the user preference signature being based on explicit preferences provided by the user.
providing access to a set of content items, each of the content items being associated with metadata that describes the corresponding content item;	55	Item 44. The method of item 38, the user preference signature being based on analyzing content item selections by user conducted over a period of time.
providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items;		Item 45. The method of item 38, the user preference signature describing preferences of the user for metadata associated with the content items, the metadata including entities preferred by the user.
		Item 46. The method of item 38, the content items including media entertainment items and the metadata including at least one of crew, characters, actors, teams, leagues, tournaments, athletes, composers, music artists, performers, albums, songs, news personalities, and content distributors.
		Item 47. The method of item 38, the content items including electronic mail items and the metadata including at least one of electronic mail threads, contacts, senders, recipients, company names, business departments, business units, electronic mail folders, and office location information.

Item 48. The method of item 38, the content items including travel-related items and the metadata including at least one of cities, hotels, hotel brands, individual points of interest, categories of points of interest, retail chains, car rental websites, and car rental company names.	5	data associated with the content items.
Item 49. The method of item 38, the content items including electronic commerce items and the metadata including at least one of product items, product categories, product subcategories, product brand, and retail stores.	10	Item 53. The method of item 52, the ambiguous identifier being a pronoun or syntactic expletive.
Item 50. The method of item 38, the content items including network-based documents and the metadata including at least one of domain names, internet media types, filenames, directories, and filename extensions.	15	Item 54. The method of item 52, the ambiguous identifier being an identifier of an entertainment genre.
Item 51. The method of item 38, the content items including address book items and the metadata including at least one of contact names, electronic mail address, telephone number, address, and employer.	20	Item 55. The method of item 52, the ambiguous identifier being at least a portion of a name.
Item 52. A method of inferring user intent in a search input based on resolving ambiguous portions of the search input, the method comprising: providing access to a set of content items, each of the content items being associated with metadata that describes the corresponding content item; receiving search input from the user, the search input being intended by the user to identify at least one desired content item; determining whether or not a portion of the search input contains an ambiguous identifier, the ambiguous identifier intended by the user to identify, at least in part, the at least one desired content item; upon a condition in which a portion of the search input contains an ambiguous identifier, performing the following: inferring a meaning for the ambiguous identifier based on matching portions of the search input to preferences of the user described by a user preference signature; and selecting content items from the set of content items based on comparing the search input and the inferred meaning of the ambiguous identifier with metadata associated with the content items; and	25 30 35 40 45 50	Item 56. The method of item 52, the metadata associated with the content items including a mapping of relationships between entities associated with the content items. Item 57. The method of item 52, the user preference signature being based on explicit preferences provided by the user. Item 58. The method of item 52, the user preference signature being based on analyzing content item selections by user conducted over a period of time. Item 59. The method of item 52, the user preference signature describing preferences of the user for metadata associated with the content items, the metadata including entities preferred by the user. Item 60. The method of item 52, the user preference signature describing preferences of the user for particular content items. Item 61. The method of item 52, the content items including media entertainment items and the metadata including at least one of crew, characters, actors, teams, leagues, tournaments, athletes, composers, music artists, performers, albums, songs, news personalities, and content distributors.
upon a condition in which the search input does not contain an ambiguous identifier, selecting content items from the set of content items based on comparing the search input with meta-	55	Item 62. The method of item 52, the content items including electronic mail items and the metadata including at least one of electronic mail threads, contacts, senders, recipients, company names, business departments, business units, electronic mail folders, and office location information. Item 63. The method of item 52, the content items including travel-related items and the metadata including at least one of cities, hotels, hotel brands, individual points of interest, categories of points of interest, retail chains, car rental websites, and car rental company names. Item 64. The method of item 52, the content items including electronic commerce items and the metadata including at least one of product items, product categories, product subcategories, product brand,

and retail stores.

Item 65. The method of item 52, the content items including network-based documents and the metadata including at least one of domain names, internet media types, filenames, directories, and filename extensions.

Item 66. The method of item 52, the content items including address book items and the metadata including at least one of contact names, electronic mail address, telephone number, address, and employer.

Claims

1. A computer-implemented method of inferring user intent in a search input, the method comprising:

receiving information for a set of content items, each of the content items being associated with metadata that describes the corresponding content item, and each of the content items being associated with at least one information domain; receiving the search input from a user at one or more of the processors, the search input being intended by the user to identify at least one desired content item; inferring by the one or more of the processors that a first portion of the search input contains a phrase of at least one word that has a plurality of possible meanings, wherein each possible meaning of the plurality of possible meanings is associated with a corresponding information domain; determining by the one or more of the processors at least one information domain relevant to the search input based in part on a second portion of the search input and not based on the first portion of the search input; selecting by the one or more of the processors one possible meaning for the phrase from the plurality of possible meanings based on the search input and the at least one information domain determined to be relevant to the search input; and selecting by the one or more of the processors a result for the search input based on the search input, the one possible meaning for the phrase, and metadata associated with the content items, wherein the result comprises at least one content item from the set of content items.

2. The method of claim 1, the at least one information domain including media entertainment and the metadata including at least one of crew, characters, actors, teams, leagues, tournaments, athletes, composers, music artists, performers, albums, songs,

news personalities, and content distributors.

3. The method of claim 1 or 2, further comprising providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items, wherein the selecting the one possible meaning for the phrase is further based on comparing portions of the search input to the preferences of the user described by the user preference signature.

4. The method of any one of claim 1-3, further comprising: providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items; and ordering the at least one content item based on the preferences of the user described by the user preference signature.

5. The method of any one of claims 1-4, the metadata associated with the content items including a mapping of relationships between entities associated with the content items, and the determining by the one or more processors the at least one information domain relevant to the search input being based in part on the mapping of relationships and the search input.

6. The method of any one of claims 1-4, wherein the one possible meaning for the phrase from the plurality of possible meanings is selected by one or more of the processors based in part on a previous search input.

7. A system of inferring user intent in a search input, the method comprising:

means for receiving information for a set of content items, each of the content items being associated with metadata that describes the corresponding content item, and each of the content items being associated with at least one information domain; means for receiving the search input from a user at one or more of the processors, the search input being intended by the user to identify at least one desired content item; means for inferring by the one or more of the processors that a first portion of the search input contains a phrase of at least one word that has a plurality of possible meanings, wherein each possible meaning of the plurality of possible meanings is associated with a corresponding information domain; means for determining by the one or more of the processors at least one information domain relevant to the search input based in part on a second portion of the search input and not based on the first portion of the search input; and means for selecting by the one or more of the processors one possible meaning for the phrase from the plurality of possible meanings based on the search input and the at least one information domain determined to be relevant to the search input; and means for selecting by the one or more of the processors a result for the search input based on the search input, the one possible meaning for the phrase, and metadata associated with the content items, wherein the result comprises at least one content item from the set of content items.

event to the search input based in part on a second portion of the search input and not based on the first portion of the search input;
means for selecting by the one or more of the processors one possible meaning for the phrase 5
from the plurality of possible meanings based on the search input and the at least one information domain determined to be relevant to the search input; and
means for selecting by the one or more of the processors a result for the search input based 10
on the search input, the one possible meaning for the phrase, and metadata associated with the content items, wherein the result comprises at least one content item from the set of content items. 15

8. The system of claim 7, the at least one information domain including media entertainment and the metadata including at least one of crew, characters, actors, teams, leagues, tournaments, athletes, composers, music artists, performers, albums, songs, news personalities, and content distributors. 20
9. The system of claim 7 or 8, further comprising providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items, wherein the selecting the one possible meaning for the phrase is further based on comparing portions of the search input to the preferences of the user described by the user preference signature. 25
10. The system of any one of claim 7-9, further comprising: providing a user preference signature, the user preference signature describing preferences of the user for at least one of (i) particular content items and (ii) metadata associated with the content items; and ordering the at least one content item based on 30
the preferences of the user described by the user preference signature. 35
11. The system of any one of claims 7-10, the metadata associated with the content items including a mapping of relationships between entities associated with the content items, and the determining by the one or more processors the at least one information domain relevant to the search input being based in part on the mapping of relationships and the search input. 40
50
12. The system of any one of claims 7-11, wherein the one possible meaning for the phrase from the plurality of possible meanings is selected by one or more of the processors based in part on a previous search input. 55

100

105

110

Advanced Search

Find Movies, TV shows, Celebrities and more...

Movies TV News Trailers Community

All All Titles TV Episodes Names Companies Keywords Characters Videos Quotes Bios Plots Advanced Search >>

Welcome to Advanced Search, which gives you instant access to our 1,633,300 titles and 3,649,234 names.

Advanced Title Search

Want to get a list of comedies from the 1970s that have at least 1000 votes and an average rating of 7.5 or higher? Use [Advanced Title Search](#).

Advanced Name Search

Want a list of males in the database who are Virgos and over 6 feet tall? Use [Advanced Name Search](#).

Collaborations and Overlaps

Want a list of titles in which both Brad Pitt and George Clooney appeared? Or a list of people who worked on both Forrest Gump and Apollo 13? Try searching [Collaborations and Overlaps](#).

Browse/Search Videos

If you'd like to browse our full list of available content, you can look at our [Browse Videos](#) page. Alternately, you can search our video catalog at [Search Videos](#) and narrow down your results using the facets on the right hand side of the page.

Title Text Search

Please select a section to search within (plot summary, quotes, trivia, etc) and enter a word to search for (e.g. "horses").

Plot

Plot

Trivia
Quotes
Gags
Crazy Credits
Filming Locations
Soundtracks
Literature
Versions

search

tion to search within (mini biography, trivia, quotes) and enter a word to

ested").

FIG. 1

200



e.g. *The Godfather*

205

Title Type

Feature Film TV Movie TV Series TV Episode
 TV Special Mini-Series Documentary Video Game
 Short Film Video Unknown Work

Release Date 

210

 to

Format: YYYY-MM-DD, YYYY-MM, or YYYY

215

User Rating
 - to -

220

Number of Votes 
 to
Genres

225

Action Adventure Animation Biography
 Comedy Crime Documentary Drama
 Family Fantasy Film-Noir Game-Show
 History Horror Music Musical
 Mystery News Reality-TV Romance
 Sci-Fi Sport Talk-Show Thriller
 War Western

Title Groups

230

<input type="checkbox"/> Oscar-Winning	<input type="checkbox"/> Best Picture-Winning	<input type="checkbox"/> Best Director-Winning
<input type="checkbox"/> Oscar-Nominated	<input type="checkbox"/> Emmy Award-Winning	<input type="checkbox"/> Emmy Award-Nominated
<input type="checkbox"/> Golden Globe-Winning	<input type="checkbox"/> Golden Globe-Nominated	<input type="checkbox"/> Razzie-Winning
<input type="checkbox"/> Razzie-Nominated	<input type="checkbox"/> National Film Board Preserved	<input type="checkbox"/> IMDb "Top 100"
<input type="checkbox"/> IMDb "Top 250"	<input type="checkbox"/> IMDb "Top 1000"	<input type="checkbox"/> IMDb "Bottom 100"
<input type="checkbox"/> IMDb "Bottom 250"	<input type="checkbox"/> IMDb "Bottom 1000"	<input type="checkbox"/> Now-Playing

235

Title Data 

Alternate Versions	   
Awards	
Blu-ray at Amazon.ca	
Blu-ray at Amazon.com	

FIG. 2

300

Car Stereos

Shop Car Stereos 1 - 15 of 208 [View: 15 per page](#)

CD Decks Compare up to 4 items [COMPARE](#)

Narrow Your Results

Customer Reviews Top-Rated (17)

Brand Kenwood (28) Pioneer (25) JVC (23) Pyle (19) Dual (19) Sony (13) Boss (12) Alpine (12) [See all...](#)

Price Range Less than \$100 (86) \$100 - \$199.99 (65) \$200 - \$299.99 (24) \$300 - \$399.99 (11) \$400 - \$499.99 (4) \$500 - \$599.99 (2) \$600 - \$699.99 (5) \$700 - \$799.99 (6) [See all...](#)

Features Built-in HD Radio (24) Color Display (56) HD Capable (48) iPod Capable (96) MP3 Playback Capability (122)

Faceplate Style [?](#) Detachable (64) Detachable flip-down (6) Detachable motorized (6)

Pioneer - 50W x 4 MOSFET In-Dash CD Deck with MP3 Playback and Detachable Faceplate

Model: DEH-1300MP SKU: 1149632

Front auxiliary input; compatible with CD, CD-R/RW, MP3 and WMA formats; remote

Customer Reviews:  3.9 of 5 (36 reviews)

[Check Shipping & Availability >](#)

Kenwood - 50W x 4 MOSFET Apple® iPod®-Ready In-Dash CD Deck

Model: KDC-252U SKU: 4880712

Detachable faceplate; auxiliary input; supports MP3/WMA playback; Pandora control for Apple iPhone®

Customer Reviews:  5 of 5 (3 reviews)

Not Available for Shipping

[Check Shipping & Availability >](#)

JVC - 50W x 4 MOSFET In-Dash CD Deck with MP3 Playback

Model: KDR330 SKU: 4001816

Front and rear auxiliary inputs; Bluetooth compatible, plays back CD, CD-R/RW, WMA and MP3 formats

Customer Reviews:  4.6 of 5 (22 reviews)

[Check Shipping & Availability >](#)

Pioneer - 50W x 4 MOSFET Apple® iPod®-Ready In-Dash CD Deck

Model: DEH-3400UB SKU: 4039973

Detachable faceplate; front auxiliary input; compatible with CD, CD-R, MP3 and WMA formats

[Check Shipping & Availability >](#)

Sort by: [Best Selling](#) [Price Low to High](#) [Price High to Low](#) [Brand A-Z](#) [Brand Z-A](#) [Best Selling](#) [Free Shipping](#) [New Arrivals](#) [Highest Rated](#) [See How](#)

\$79.99 **\$89.99** **\$69.99** **\$99.99**

Reg. Price: \$99.99
You Save: \$10.00
• Free Shipping
• Get 4% Back in Rewards:
See How

[ADD TO CART](#) [ADD TO CART](#) [ADD TO CART](#) [ADD TO CART](#)

305

FIG. 3

400



Name	<input type="text"/>	Visit Now						
e.g. <i>Audrey Hepburn</i>		ad feedback						
Birth Date	<input type="text"/> to <input type="text"/>							
Format: YYYY-MM-DD, YYYY-MM, or YYYY								
Name Groups								
<input type="checkbox"/> Best Actress-Nominated <input type="checkbox"/> Best Actor-Nominated <input type="checkbox"/> Best Actress-Winning <input type="checkbox"/> Best Actor-Winning <input type="checkbox"/> Best Supporting Actress-Nominated <input type="checkbox"/> Best Supporting Actor-Nominated <input type="checkbox"/> Best Supporting Actress-Winning <input type="checkbox"/> Best Supporting Actor-Winning <input type="checkbox"/> Best Director-Nominated <input type="checkbox"/> Best Director-Winning <input type="checkbox"/> Oscar-Winning								
Star Sign								
<input type="checkbox"/> Aquarius <input type="checkbox"/> Pisces <input type="checkbox"/> Aries <input type="checkbox"/> Taurus <input type="checkbox"/> Gemini <input type="checkbox"/> Cancer <input type="checkbox"/> Leo <input type="checkbox"/> Virgo <input type="checkbox"/> Libra <input type="checkbox"/> Scorpio <input type="checkbox"/> Sagittarius <input type="checkbox"/> Capricorn								
405	Birth Place	<input type="text"/>						
405	Death Date	<input type="text"/> to <input type="text"/>						
Format: YYYY-MM-DD, YYYY-MM, or YYYY								
405	Death Place	<input type="text"/>						
Gender								
<input type="checkbox"/> Males <input type="checkbox"/> Females								
Height								
<input type="text"/> - to <input type="text"/> -								
410	Filmography	<input type="text"/>						
Biographies								
<small>Search for words that might appear in the Mini-Biography</small> <input type="text"/>								
Lists								
<small>You must login or register to use this feature.</small>								
Name Data <table border="1" style="display: inline-table; vertical-align: middle;"> <tr><td>Awards</td></tr> <tr><td>Biography</td></tr> <tr><td>Birth Date</td></tr> <tr><td>Date of Death</td></tr> <tr><td>Height Info</td></tr> <tr><td>Quotes</td></tr> </table>			Awards	Biography	Birth Date	Date of Death	Height Info	Quotes
Awards								
Biography								
Birth Date								
Date of Death								
Height Info								
Quotes								

FIG. 4

500
↓

Common Cast / Crew Between Two Titles

To find a list of people who worked together on two titles, start typing the first title and select it from the drop-down. Then, advance to the second box and follow the same steps. Once you have your two titles selected, click the search button to continue.

Search

Two People Working Together

To begin, start typing one person's name in the first box and select the name from the drop-down list of possible matches. Then, advance to the second box and follow the same steps. Once you have your two names selected, click the search button to continue.

Search

FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12



EUROPEAN SEARCH REPORT

Application Number

EP 17 21 0590

5

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
10 X	US 2012/016678 A1 (GRUBER THOMAS ROBERT [US] ET AL) 19 January 2012 (2012-01-19) * paragraphs [0012], [0127], [0219], [0410] - [0419], [0476], [0585]; figures 8,28-33 *	1-12	INV. G06F17/27 G06F17/30 G06F17/28 G06F17/21
15 X	----- US 2009/144609 A1 (LIANG JISHENG [US] ET AL) 4 June 2009 (2009-06-04) * Abstract, Fig. 7 and associated passages, claims 1-10 *	1-5,7-12	
20 X	----- ASHWANI KUMAR ET AL: "Reference resolution as a facilitating process towards robust multimodal dialogue management: A cognitive grammar approach", INTERNATIONAL SYMPOSIUM ON REFERENCE RESOLUTION AND ITS APPLICATION TO QUESTION ANSWERING AND SUMMARIZATION, 1 January 2003 (2003-01-01), pages 1-8, XP055372479, * the whole document *	1-12	
25 A	----- Daniel Jurafsky ET AL: "Chapter 21 : Computational Discourse" In: "Computational Linguistics (Online)", 2007, MIT Press, Cambridge, MA, USA, XP055274595, ISSN: 1530-9312 vol. 26 DOI: 10.1162/089120100750105975, * the whole document *	1-12	TECHNICAL FIELDS SEARCHED (IPC)
30			G06F
35 T	----- N.A. MOHAIDEEN ET AL: "A Natural Language Interface for Search and Recommendations of Digital Entertainment Media", IBC 2015 CONFERENCE, 2015, pages 11 .-11 ., XP055274592, DOI: 10.1049/ibc.2015.0043 ISBN: 978-1-78561-185-8 * the whole document *	-----	
40			
45			
50 1	The present search report has been drawn up for all claims		
55	Place of search EPO FORM 1503 03-82 (P04C01) Berlin	Date of completion of the search 5 June 2018	Examiner Alt, Susanne
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO. [REDACTED]**

EP 17 21 0590

5 This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

05-06-2018

Patent document cited in search report		Publication date	Patent family member(s)			Publication date
US	2012016678	A1	19-01-2012	AU	2011205426 A1	23-08-2012
				BR	112012017826 A2	26-09-2017
				BR	122012028965 A2	27-03-2018
				CA	2787351 A1	21-07-2011
				CA	2791791 A1	21-07-2011
				CA	2792412 A1	21-07-2011
				CA	2792442 A1	21-07-2011
				CA	2792570 A1	21-07-2011
				CA	2793002 A1	21-07-2011
				CA	2793118 A1	21-07-2011
				CA	2793248 A1	21-07-2011
				CA	2793741 A1	21-07-2011
				CA	2793743 A1	21-07-2011
				CA	2954559 A1	21-07-2011
				CA	3000109 A1	21-07-2011
				CN	102792320 A	21-11-2012
				CN	105808200 A	27-07-2016
				EP	2526511 A2	28-11-2012
				EP	3131023 A1	15-02-2017
				GB	2490444 A	31-10-2012
				JP	5948372 B2	06-07-2016
				JP	5956511 B2	27-07-2016
				JP	5957038 B2	27-07-2016
				JP	5973500 B2	23-08-2016
				JP	6027052 B2	16-11-2016
				JP	6175413 B2	02-08-2017
				JP	6193181 B2	06-09-2017
				JP	2013517566 A	16-05-2013
				JP	2014222509 A	27-11-2014
				JP	2014222510 A	27-11-2014
				JP	2014222511 A	27-11-2014
				JP	2014222512 A	27-11-2014
				JP	2014222513 A	27-11-2014
				JP	2014222514 A	27-11-2014
				JP	2014222515 A	27-11-2014
				JP	2014222516 A	27-11-2014
				JP	2014222517 A	27-11-2014
				JP	2017224300 A	21-12-2017
				KR	20120120316 A	01-11-2012
				KR	20120136417 A	18-12-2012
				KR	20120137424 A	20-12-2012
				KR	20120137425 A	20-12-2012
				KR	20120137434 A	20-12-2012
				KR	20120137435 A	20-12-2012
				KR	20120137440 A	20-12-2012
				KR	20120138826 A	26-12-2012

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 17 21 0590

5 This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

05-06-2018

10	Patent document cited in search report	Publication date	Patent family member(s)	Publication date
15			KR 20120138827 A	26-12-2012
			KR 20130000423 A	02-01-2013
			KR 20160105995 A	08-09-2016
			KR 20170104006 A	13-09-2017
20			MX 338784 B	02-05-2016
			MX 342072 B	13-09-2016
			MX 348250 B	05-06-2017
			RU 2012135502 A	27-02-2014
25			RU 2012144605 A	27-04-2014
			RU 2012144606 A	10-05-2014
			RU 2012144637 A	10-05-2014
			RU 2012144639 A	10-05-2014
			RU 2012144640 A	10-05-2014
30			RU 2012144643 A	10-05-2014
			RU 2012144644 A	10-05-2014
			RU 2012144647 A	10-05-2014
			RU 2012144648 A	10-05-2014
35			RU 2015120954 A	27-12-2016
			US 2012016678 A1	19-01-2012
			US 2012245944 A1	27-09-2012
			US 2013110505 A1	02-05-2013
40			US 2013110515 A1	02-05-2013
			US 2013110518 A1	02-05-2013
			US 2013110519 A1	02-05-2013
			US 2013110520 A1	02-05-2013
			US 2013111348 A1	02-05-2013
45			US 2013111487 A1	02-05-2013
			US 2013117022 A1	09-05-2013
			US 2013185074 A1	18-07-2013
			US 2013185081 A1	18-07-2013
			US 2017178626 A1	22-06-2017
50			WO 2011088053 A2	21-07-2011
	-----		US 2009144609 A1 04-06-2009	23-04-2009
			AU 2008312487 A1	23-04-2009
			CA 2703007 A1	23-04-2009
			EP 2206057 A1	14-07-2010
			US 2009144609 A1	04-06-2009
55			US 2014142922 A1	22-05-2014
			US 2017262412 A1	14-09-2017
			WO 2009052277 A1	23-04-2009

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

page 2 of 2

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 87452313 A [0001]
- US 66740012 A [0001]
- US 66738812 A [0001]
- US 61712721 A [0001]
- US 61673867 A [0001]
- US 7774294 B [0086]
- US 7835998 B [0086]
- US 7461061 B [0086]
- US 8112454 B [0086]

摘要：

本发明公开涉及在对话交互系统中推断搜索输入中的用户意图的方法和系统。公开了在对话交互系统中推断搜索输入中的用户意图的方法。推断搜索输入中的用户意图的方法包括提供描述用户偏好的用户偏好签名、从用户接收要被用户用于识别至少一个期望的项的搜索输入，并且确定搜索输入的一部分包含歧义标识符。该歧义标识符要被用户用于至少部分地识别期望的项。该方法还包括基于搜索输入的部分与由用户偏好签名描述的用户偏好匹配而推断歧义标识符的含义并且基于将搜索输入及所推断的歧义标识符的含义与和内容项相关联的元数据的比较而从一组内容项中选择项。