



(12) **United States Patent**  
**Galitsky**

(10) **Patent No.:** **US 11,562,135 B2**  
(45) **Date of Patent:** **Jan. 24, 2023**

(54) **CONSTRUCTING CONCLUSIVE ANSWERS FOR AUTONOMOUS AGENTS**

(71) Applicant: **Oracle International Corporation**,  
Redwood Shores, CA (US)

(72) Inventor: **Boris Galitsky**, San Jose, CA (US)

(73) Assignee: **Oracle International Corporation**,  
Redwood Shores, CA (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 511 days.

(21) Appl. No.: **16/654,258**

(22) Filed: **Oct. 16, 2019**

(65) **Prior Publication Data**  
US 2020/0117709 A1 Apr. 16, 2020

**Related U.S. Application Data**

(60) Provisional application No. 62/746,261, filed on Oct. 16, 2018.

(51) **Int. Cl.**  
**G06F 40/211** (2020.01)  
**G06N 5/00** (2006.01)  
**G06F 16/22** (2019.01)  
**G06F 16/953** (2019.01)  
**G06N 20/00** (2019.01)  
**G06F 40/279** (2020.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 40/211** (2020.01); **G06F 16/2246** (2019.01); **G06F 16/953** (2019.01); **G06F 40/279** (2020.01); **G06N 5/003** (2013.01); **G06N 20/00** (2019.01)

(58) **Field of Classification Search**  
CPC ..... G10L 15/1822; G06F 8/31; G06F 40/211; G06N 3/08; B60W 50/14  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,495,605 A \* 2/1996 Cadot ..... G06F 16/24542  
6,279,476 B1 8/2001 Ellis  
6,598,054 B2 7/2003 Schuetze et al.  
6,731,307 B1 \* 5/2004 Strubbe ..... H04N 21/466 715/764  
6,922,699 B2 7/2005 Schuetze et al.  
(Continued)

OTHER PUBLICATIONS

Boltuzic-Snajder, "Back up your Stance: Recognizing Arguments in Online Discussions", Proceedings of the First Workshop on Argumentation Mining, pp. 49-58, Baltimore, Maryland USA, Jun. 26, 2014. (Year: 2014).\*

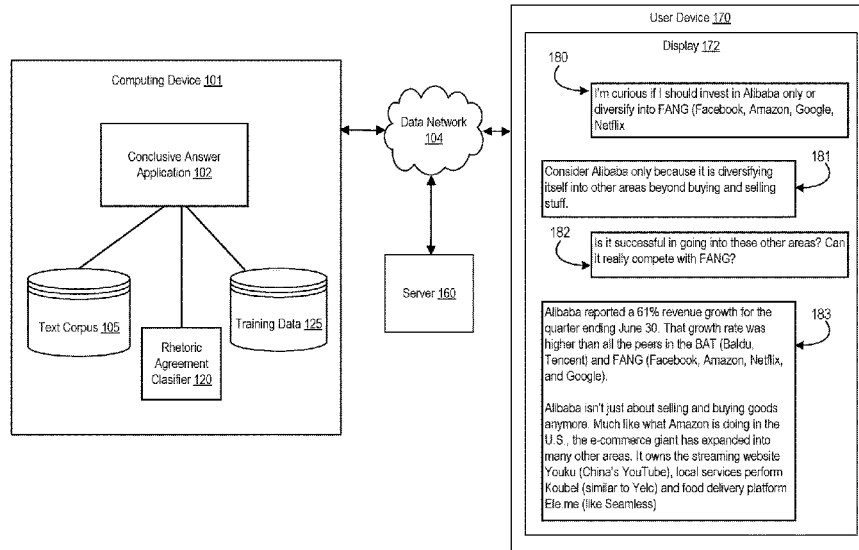
(Continued)

*Primary Examiner* — Anne L Thomas-Homescu  
(74) *Attorney, Agent, or Firm* — Kilpatrick Townsend & Stockton LLP

(57) **ABSTRACT**

Techniques are described herein for enabling autonomous agents to generate conclusive answers. An example of a conclusive answer is text that addresses concerns of a user who is interacting with an autonomous agent. For example, an autonomous agent interacts with a user device, answering user utterances, for example questions or concerns. Based on the interactions, the autonomous agent determines that a conclusive answer is appropriate. The autonomous agent formulates the conclusive answer, which addresses multiple user utterances. The conclusive answer provided to the user device.

**17 Claims, 23 Drawing Sheets**





(56)

## References Cited

## OTHER PUBLICATIONS

- Carstens-Toni, "Using Argumentation to Improve Classification in Natural Language Problems", *ACM Transactions on Internet Technology*, vol. 17, No. 3, Article 30, Publication date: Jul. 2017. (Year: 2017).\*
- Mozina et al., "Argument Based Machine Learning from Examples and Text", 2009 First Asian Conference on Intelligent Information and Database Systems. (Year: 2009).\*
- Cox et al., "Vicarious Learning From Dialogue and Discourse: A Controlled Comparison", *Instructional Science*, vol. 27, Nov. 1999, pp. 431-458.
- Craig et al., "Overhearing Dialogues and Monologues in Virtual Tutoring Sessions: Effects on Questioning and Vicarious Learning", *International Journal of Artificial Intelligence in Education*, vol. 11, Jan. 2000, pp. 242-253.
- Folino et al., "Proceedings of the 3rd Workshop on Biologically Inspired Algorithms for Distributed Systems", *BADS*, Jun. 14-18, 2011.
- Leouski et al., "An Evaluation of Techniques for Clustering Search Results", Available Online at: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.34.5627&rep=rep1&type=pdf>, 1996, 19 pages.
- Li et al., "Adversarial Learning for Neural Dialogue Generation", *Empirical Methods in Natural Language Processing*, Sep. 7-11, 2017, pp. 2157-2169.
- Li et al., "DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset", *Proceedings of the Eighth International Joint Conference on Natural Language Processing, Long Papers*, vol. 1, Dec. 1, 2017, pp. 986-995.
- Li et al., "Deep Reinforcement Learning for Dialogue Generation", *Empirical Methods in Natural Language Processing*, Available Online at: <https://www.aclweb.org/anthology/D16-1127>, Nov. 2016, pp. 1192-1202.
- Lowe et al., "On the Evaluation of Dialogue Systems with Next Utterance Classification", *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Available online at: <https://arxiv.org/pdf/1605.05414.pdf>, Jul. 23, 2016, 6 pages.
- Luan et al., "LSTM based Conversation Models", *arXiv:1603.09457v1*, Mar. 31, 2016, 5 pages.
- Manning et al., "An Introduction to Information Retrieval", Cambridge University Press, vol. 39, 2008, 581 pages.
- Mitkov et al., "A Computer-Aided Environment for Generating Multiple-Choice Test Items", *Natural Language Engineering*, vol. 12, No. 2, Jun. 2006, pp. 177-194.
- Piwek et al., "T2D: Generating Dialogues Between Virtual Agents Automatically from Text", *Intelligent Virtual Agents, LNAI*, vol. 4722, Sep. 2007, pp. 161-174.
- Tunkelang, "Search Results Clustering", Available Online at: <https://queryunderstanding.com/search-results-clustering-b2fa64c6c809>, Apr. 16, 2018, 3 pages.
- U.S. Appl. No. 16/240,232, Non-Final Office Action, dated Apr. 9, 2021, 13 pages.
- AI Marketing, Chatbots, and Your CMS, Available Online at <https://simplea.com/Articles/AI-Marketing-Chatbots-and-Your-CMS>, Accessed from Internet on: Nov. 19, 2019, 9 pages.
- Aleman-Meza et al., Context-Aware Semantic Association Ranking, *Proc. First Int'l Workshop Semantic Web and Databases*, Sep. 7-8, 2003, 18 pages.
- Alsinet et al., A Logic Programming Framework for Possibilistic Argumentation: Formalization and Logical Properties, *Fuzzy Sets and Systems*, vol. 159, Issue 10, May 16, 2008, pp. 1208-1228.
- Altinel et al., A Corpus-Based Semantic Kernel for Text Classification by Using Meaning Values of Terms, *Engineering Applications of Artificial Intelligence*, vol. 43, Aug. 2015, pp. 54-66.
- An Vo, FBK-HLT: A New Framework for Semantic Textual Similarity, *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, Jun. 2015, pp. 102-106.
- Antoniou et al., Representation Results for Defeasible Logic, *ACM Transactions on Computational Logic*, vol. 2, Issue 2, Apr. 2001, pp. 255-287.
- Banerjee et al., WikiWrite: Generating Wikipedia Articles Automatically, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, 2016, pp. 2740-2746.
- Banko et al., Open Information Extraction from the Web, In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, Jan. 6-12, 2007, pp. 2670-2676.
- Baralis et al., Generalized Association Rule Mining with Constraints, *Information Sciences—Informatics and Computer Science, Intelligent Systems, Applications: An International Journal*, vol. 194, Jul. 2012, pp. 68-84.
- Bar-Haim et al., Semantic Inference at the Lexical-Syntactic Level, *Proceedings of AAAI*, 2007, pp. 871-876.
- Baroni et al., CleanEval: A Competition for Cleaning Webpages, *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, May 2008, pp. 638-643.
- Blei et al., Latent Dirichlet Allocation, *Journal of Machine Learning Research*, vol. 3, Jan. 2003, pp. 993-1022.
- Bordini et al., A Survey of Programming Languages and Platforms for Multi-Agent Systems, *Informatica*, vol. 30, Issue 1, Jan. 2006, pp. 33-44.
- Bron et al., Algorithm 457: Finding All Cliques of an Undirected Graph, *Communications of the ACM*, vol. 16, No. 9, Sep. 1973, pp. 575-579.
- Brzezinski et al., Accuracy Updated Ensemble for Data Streams with Concept Drift, *International Conference on Hybrid Artificial Intelligence Systems*, vol. 6679, May 23-25, 2011, pp. 155-163.
- Bunke, Graph-Based Tools for Data Mining and Machine Learning, *Lecture Notes in Computer Science*, vol. 2734, Jan. 2003, pp. 7-19.
- Cai et al., Extracting Content Structure for Web Pages Based on Visual Representation, *LNCS*, vol. 2642, Springer, 2003, 12 pages.
- Cardie et al., Guest Editor's Introduction: Machine Learning and Natural Language, *Machine Learning*, vol. 1, No. 5, Feb. 1999, 5 pages.
- Carreras et al., Introduction to the CoNLL-2004 Shared Task: Semantic Role Labeling, *Proceedings of the Eighth Conference on Computational Natural Language Learning, Association for Computational Linguistics*, May 6-7, 2004, pp. 89-97.
- Chakrabarti et al., Graph Mining: Laws, Generators, and Algorithms, *ACM Computing Surveys*, vol. 38, Issue 1, Mar. 2006, pp. 69-123.
- Chesnevar et al., Empowering Recommendation Technologies Through Argumentation, *Argumentation in Artificial Intelligence*, 2009, pp. 403-422.
- Cumby et al., On Kernel Methods for Relational Learning, *ICML*, 2003, pp. 107-114.
- Cuzzocrea, Intelligent Knowledge-Based Models and Methodologies for Complex Information Systems, *Information Sciences*, vol. 194, Jul. 1, 2012.
- De Salvo Braz et al., An Inference Model for Semantic Entailment in Natural Language, *AAAI'05 Proceedings of the 20th National Conference on Artificial Intelligence*, vol. 3, Jul. 9-13, 2005, pp. 1043-1049.
- Ding et al., Swoogle: A Search and Metadata Engine for the Semantic Web, *CIKM '04 Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management*, Nov. 8-13, 2004, pp. 652-659.
- Van Durme et al., Towards Light Semantic Processing for Question Answering, *HLT-NAACL-TEXTMEANING '03 Proceedings of the HLT-NAACL 2003 Workshop on Text meaning*, vol. 9, 2003, pp. 54-61.
- "Exploring Dialog Management for Bots", *Chatbots Magazine*, Available online at: <https://chatbotmagazine.com/exploring-dialog-management-for-bots-cbb8665a2fd3>, Jul. 11, 2016, 7 pages.
- Erenel et al., Nonlinear Transformation of Term Frequencies for Term Weighting in Text Categorization, *Engineering Applications of Artificial Intelligence*, vol. 25, No. 7, Oct. 2012, pp. 1505-1514.
- Ferretti et al., An Application of Defeasible Logic Programming to Decision Making in a Robotic Environment, *International Conference on Logic Programming and Nonmonotonic Reasoning*, vol. 4483, May 2007, pp. 297-302.
- Galitsky et al., A Novel Approach for Classifying Customer Complaints Through Graphs Similarities in Argumentative Dialogues, *Decision Support Systems*, vol. 46, No. 3, Feb. 2009, 28 pages.

(56)

## References Cited

## OTHER PUBLICATIONS

- Galitsky, A Tool for Efficient Content Compilation, Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: System Demonstrations, C16-2042, Dec. 11-17, 2016, pp. 198-202.
- Galitsky et al., A Web Mining Tool for Assistance with Creative Writing, European Conference on Information Retrieval, vol. 7814, 2013, pp. 828-831.
- Galitsky et al., Chatbot with a Discourse Structure-Driven Dialogue Management, Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics Available online at: <https://pdfs.semanticscholar.org/439f/cbe14dbc84b694bae2ee42a914d743126d12.pdf>, Apr. 2017, pp. 87-90.
- Galitsky et al., Concept-Based Learning of Human Behavior for Customer Relationship Management, Information Sciences, vol. 181, No. 10, May 15, 2011, pp. 2016-2035.
- Galitsky, Finding a Lattice of Needles in a Haystack: Forming a Query From a Set of Items of Interest, FCA4AI'15 Proceedings of the 4th International Conference on What can FCA do for Artificial Intelligence?, vol. 1430, 2015, 8 pages.
- Galitsky et al., Inferring the Semantic Properties of Sentences by Mining Syntactic Parse Trees, Data & Knowledge Engineering, vols. 81-82, Nov.-Dec. 2012, 44 pages.
- Galitsky, Learning Parse Structure of Paragraphs and its Applications in Search, Engineering Applications of Artificial Intelligence, vol. 32, Jun. 2014, pp. 160-184.
- Galitsky, Machine Learning of Syntactic Parse Trees for Search and Classification of Text, Engineering Applications of Artificial Intelligence, vol. 26, No. 3, Mar. 2013, pp. 1072-1091.
- Galitsky, Natural Language Question Answering System, Technique of Semantic Headers, Advanced Knowledge International, vol. 2, Apr. 2003, 333 pages.
- Galitsky et al., Parse Thicket Representations for Answering Multi-Sentence Search, International Conference on Conceptual Structures, vol. 7735, 2013, pp. 153-172.
- Galitsky et al., Rhetoric Map of an Answer to Compound Queries, Proceedings of the 53rd Annual Meeting of the 20 Association for Computational Linguistics and the 7th International Joint Conference of Natural Language Processing, Jul. 26-31, 2015, pp. 681-686.
- Galitsky et al., Text Classification into Abstract Classes Based on Discourse Structure, Proceedings of Recent Advances in Natural Language Processing, Sep. 7-9, 2015, pp. 200-207.
- Galitsky, Transfer Learning of Syntactic Structures for Building Taxonomies for Search Engines, Engineering Applications of Artificial Intelligence, vol. 26, No. 10, Nov. 2013, pp. 2504-2515.
- Galitsky et al., Using Generalization of Syntactic Parse Trees for Taxonomy Capture on the Web, Proceedings of the 19th International Conference on Conceptual Structures for Discovering Knowledge, Jul. 25-29, 2011, pp. 104-117.
- Galitsky, "Learning Noisy Discourse Trees", Computational Linguistics and Intellectual Technologies, Proceedings of the International Conference "Dialogue 2017". Available online at: <http://www.dialog-21.ru/media/3911/galitskyb.pdf>, May 31-Jun. 3, 2017, 14 pages.
- Garcia et al., Defeasible Logic Programming: An Argumentative Approach, Theory and Practice of Logic Programming, vol. 4, Issue 2, Jan. 2004, pp. 95-138.
- Gartner, Gartner Says 25 Percent of Customer Service Operations Will Use Virtual Customer Assistants by 2020, Newsroom, Available Online at <https://www.gartner.com/newsroom/id/3858564>, Feb. 19, 2018, 3 pages.
- Get Search Results Faster, Available Online at <https://support.google.com/websearch/answer/106230>, Accessed from Internet on: Nov. 19, 2019, 2 pages.
- Gildea, Loosely Tree-Based Alignment for Machine Translation, Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics, vol. 1, Jul. 2003, pp. 80-87.
- Go et al., Twitter Sentiment Classification Using Distant Supervision, Technical Report, Jan. 2009, 6 pages.
- Gomez et al., CICBUAPlp: Graph-Based Approach for Answer Selection in Community Question Answering Task, Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015), Jun. 4-5, 2015, pp. 18-22.
- Gomez et al., Reasoning with Inconsistent Ontologies through Argumentation, Applied Artificial Intelligence, vol. 24, Issue 1-2, Feb. 2010, pp. 102-148.
- Hendriks et al., Procedural Content Generation for Games: A Survey, ACM Transactions on Multimedia Computing, Communications and Applications, vol. 9, No. 1, Feb. 2013, pp. 1:1-1:22.
- Iosif et al., Unsupervised Semantic Similarity Computation Between Terms Using Web Documents, IEEE Transactions on Knowledge and Data Engineering, vol. 22, Issue 11, Nov. 2010, pp. 1637-1647.
- Jaccard, The Distribution of the Flora in the Alpine Zone, New Phytologist, vol. 11, No. 2, Available Online at: [http://www.researchgate.net/profile/Paul\\_Jaccard/publication/230302439\\_The\\_distribution\\_of\\_the\\_flora\\_in\\_the\\_alpine\\_zone/links/02e7e51cb76619a0fa000000.pdf](http://www.researchgate.net/profile/Paul_Jaccard/publication/230302439_The_distribution_of_the_flora_in_the_alpine_zone/links/02e7e51cb76619a0fa000000.pdf), Feb. 1912, 15 pages.
- Janusz et al., Unsupervised Similarity Learning from Textual Data, Fundamenta Informaticae, vol. 119, 2012, pp. 319-336.
- Johnson et al., Procedural Generation of Linguistics, Dialects, Naming Conventions and Spoken Sentences, Proceedings of 1st International Joint Conference of DiGRA and FDG, 2016, 9 pages.
- Kapoor et al., Algorithms for Enumerating All Spanning Trees of Undirected and Weighted Graph, SIAM Journal on Computing, vol. 24, No. 2, Society for Industrial and Applied Mathematics, Apr. 1995, pp. 247-265.
- Kong et al., Improve Tree Kernel-Based Event Pronoun Resolution with Competitive Information, Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, vol. 3, Jul. 16-22, 2011, pp. 1814-1819.
- Kuncheva, Classifier Ensembles for Changing Environments, Published in: F. Roli, J. Kittier and T. Windeatt (Eds.), Proc. 5th Int. Workshop on Multiple Classifier Systems, Springer-Verlag, LNCS, vol. 3077, 2004, pp. 1-15.
- Krippendorff, Reliability in Content Analysis: Some Common Misconceptions and Recommendations, Human Communication Research, vol. 30, No. 3, 2004, 15 pages.
- Liapis et al., Sentient Sketchbook: Computer-Aided Game Level Authoring, In FDG, 2013, pp. 213-220.
- Levenshtein, et al., Binary Codes Capable of Correcting Deletions, Insertions, and Reversals Cybernetics and Control Theory, vol. 10, No. 8, pp. 707-710, Feb. 1966.
- Lin et al., DIRT—Discovery of Inference Rules from Text, Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Aug. 26-29, 2001, pp. 323-328.
- Makhalova et al., News Clustering Approach Based on Discourse Text Structure, Proceedings of the First Workshop on Computing News Storylines, Jul. 2015, pp. 16-20.
- Mavridis et al., Semantic Analysis of Web Documents for the Generation of Optimal Content, Engineering Applications of Artificial Intelligence, vol. 35, Oct. 2014, pp. 114-130.
- Moldovan et al., COGEX: A Logic Prover for Question Answering, Proceedings of HLT-NAACL, Main Papers, May-Jun. 2003, pp. 87-93.
- Moreda et al., Corpus-Based Semantic Role Approach in Information Retrieval, Data & Knowledge Engineering, vol. 61, Issue 3, Jun. 2007, pp. 467-483.
- Moschitti, Kernel Methods, Syntax and Semantics for Relational Text Categorization, In proceeding of ACM 17th Conference on Information and Knowledge Management (CIKM), Oct. 26-30, 2008, pp. 253-262.
- Nagarajan et al., Pivotal Sentiment Tree Classifier, International Journal of Scientific & Technology Research, vol. 3, No. 11, Nov. 15, 2014, pp. 290-295.
- Pak et al., Twitter as a Corpus for Sentiment Analysis and Opinion Mining, Proceedings of the International Conference on Language Resources and Evaluation, May 17-23, 2010, pp. 1320-1326.
- Pasternack et al., Extracting Article Text from the Web with Maximum Subsequence Segmentation, In WWW '09: Proceedings of the 18th international conference on World wide web, Apr. 20-24, 2009, pp. 971-980.

(56)

**References Cited**

## OTHER PUBLICATIONS

- Perrin et al., An Information-Theoretic Based Model for Large-Scale Contextual Text Processing, *Information Sciences*, vol. 116, Issue 2-4, Jan. 1999, pp. 229-252.
- Plotkin, A Note on Inductive Generalization, *Machine Intelligence 5, Chapters*, 1970, pp. 153-163.
- Poon et al., Unsupervised Semantic Parsing, *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, 2009, 10 pages.
- Rahwan et al., An Argumentation-Based Approach for Practical Reasoning, *AAMAS'06*, May 8-12, 2006, pp. 347-354.
- Redey, Conformal Text Representation, *Engineering Applications of Artificial Intelligence*, vol. 6, No. 1, Feb. 1993, pp. 65-71.
- Robinson, A Machine-Oriented Logic Based on the Resolution Principle, *Journal of the Association for Computing Machinery*, vol. 12, No. 1, Jan. 1965, pp. 23-41.
- Rosenthal et al., SemEval-2014 Task 9: Sentiment Analysis in Twitter, *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Aug. 23-24, 2014, pp. 73-80.
- Rubiolo et al., Knowledge Discovery Through Ontology Matching: An Approach Based on an Artificial Neural Network Model, *Information Sciences*, vol. 194, Jul. 2012, pp. 107-119.
- Sagui et al., Modeling News Trust: A Defeasible Logic Programming Approach, *Inteligencia Artificial*, vol. 12, No. 40, Nov. 2008, pp. 63-72.
- Sauper et al., Automatically Generating Wikipedia Articles: A Structure-Aware Approach, *Proceedings of ACL. Suntec, Singapore*, Aug. 2-7, 2009, pp. 208-216.
- Sjoera, "The Linguistics Behind Chat Bots", *iCapps*, Available online at: <http://www.icapps.com/the-linguistics-behind-chatbots/>, Feb. 22, 2017, 9 pages.
- Socher et al., Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank, *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP*, Oct. 2013, pp. 1631-1642.
- Stevenson et al., A Semantic Approach to IE Pattern Induction, *Proceedings of the 43rd Annual Meeting of the ACL*, Jun. 25-30, 2005, pp. 379-386.
- Suykens et al., *Advances in Learning Theory: Methods, Models and Applications*, IOS Press, NATO Science Series, III: Computer and Systems Sciences, vol. 190, May 2006, 440 pages.
- Tang et al., Coooolll: A Deep Learning System for Twitter Sentiment Classification, *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Aug. 23-24, 2014, pp. 208-212.
- Thompson et al., Learning to Parse NL Database Queries into Logical Form, *Proceedings of the ICML-97 Workshop on Automata Induction, Grammatical Inference, and Language Acquisition*, Jul. 1997, 9 pages.
- Tneogi, *Conversational Interfaces Need a Different Content Management System*, *Chatbot Magazine*, Available Online at <https://chatbotmagazine.com/conversational-interfaces-need-a-different-content-management-system-b105bb6f716>, Mar. 26, 2017, 3 pages.
- Wade, *5 Ways Chatbots Are Revolutionizing Knowledge Management*, Available Online at <https://blog.getbizzy.io/5-ways-chatbots-are-revolutionizing-knowledge-management-bd925db66e9>, Feb. 12, 2018, 8 pages.
- Wang et al., "Using Learning Analytics to Understand the Design of an Intelligent Language Tutor-Chatbot Lucy", *International Journal of Advanced Computer Science and Applications*, vol. 4, No. 11, Nov. 2013, pp. 124-131.
- Wenyan et al., A Short Text Modeling Method Combining Semantic and Statistical Information, *Information Sciences*, vol. 180, No. 20, Oct. 15, 2010, pp. 4031-4041.
- Yang et al., Extract Conceptual Graphs from Plain Texts in Patent Claims, *Engineering Applications of Artificial Intelligence*, vol. 25, No. 4, Jun. 2012, pp. 874-887.
- Zarella et al., MITRE: Seven Systems for Semantic Similarity in Tweets, *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, Jun. 4-5, 2015, pp. 12-17.
- Zhang et al., Exploring Syntactic Structured Features Over Parse Trees for Relation Extraction Using Kernel Methods, *Information Processing & Management*, vol. 44, No. 2, Mar. 2008, pp. 687-701.
- Zhang et al., Question Classification Using Support Vector Machines, In *SIGIR '03: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Jul. 28-Aug. 1, 2003, pp. 26-32.
- U.S. Appl. No. 16/240,232, Final Office Action, dated Oct. 21, 2021, 13 pages.
- U.S. Appl. No. 16/240,232, Non-Final Office Action, dated Jan. 4, 2022, 15 pages.
- U.S. Appl. No. 16/789,840, Non-Final Office Action, dated Oct. 1, 2021, 26 pages.
- U.S. Appl. No. 16/789,840, Notice of Allowance, dated Feb. 7, 2022, 17 pages.
- U.S. Appl. No. 16/789,840, Corrected Notice of Allowability dated Apr. 5, 2022, 2 pages.

\* cited by examiner

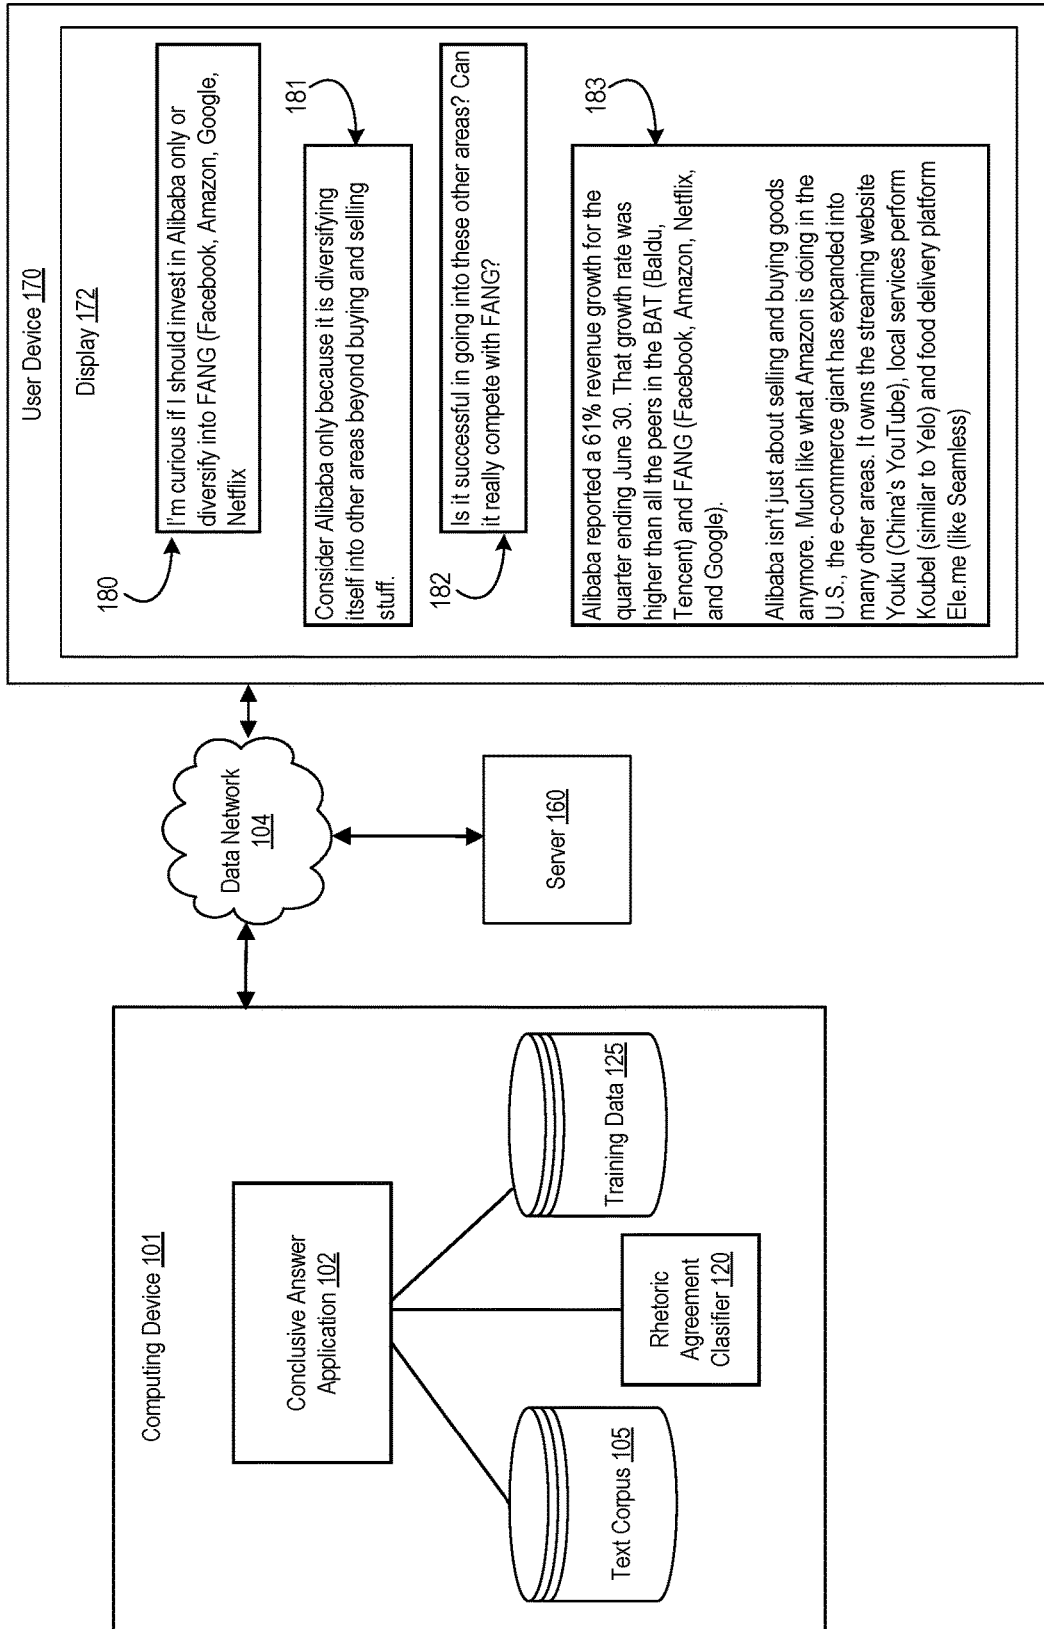


FIG. 1


**180** I'm curious if I should invest in Alibaba only or diversify into FANG (Facebook, Amazon, Google, Netflix)

**181** Consider Alibaba only because it is diversifying itself into other areas beyond buying and selling stuff.

**182** Is it successful in going into these other areas? Can it really compete with FANG?

**183** Alibaba reported a 61% revenue growth for the quarter ending June 30. That growth rate was higher than all the peers in the BAT (Baidu, Tencent) and FANG (Facebook, Amazon, Netflix, and Google).  
Alibaba isn't just about selling and buying goods anymore. Much like what Amazon is doing in the U.S., the e-commerce giant has expanded into many other areas. It owns the streaming website Youku (China's YouTube), local services perform Koubel (similar to Yelo) and food delivery platform Ele.me (like Seamless)

200



**FIG. 2**

300




FIG. 3



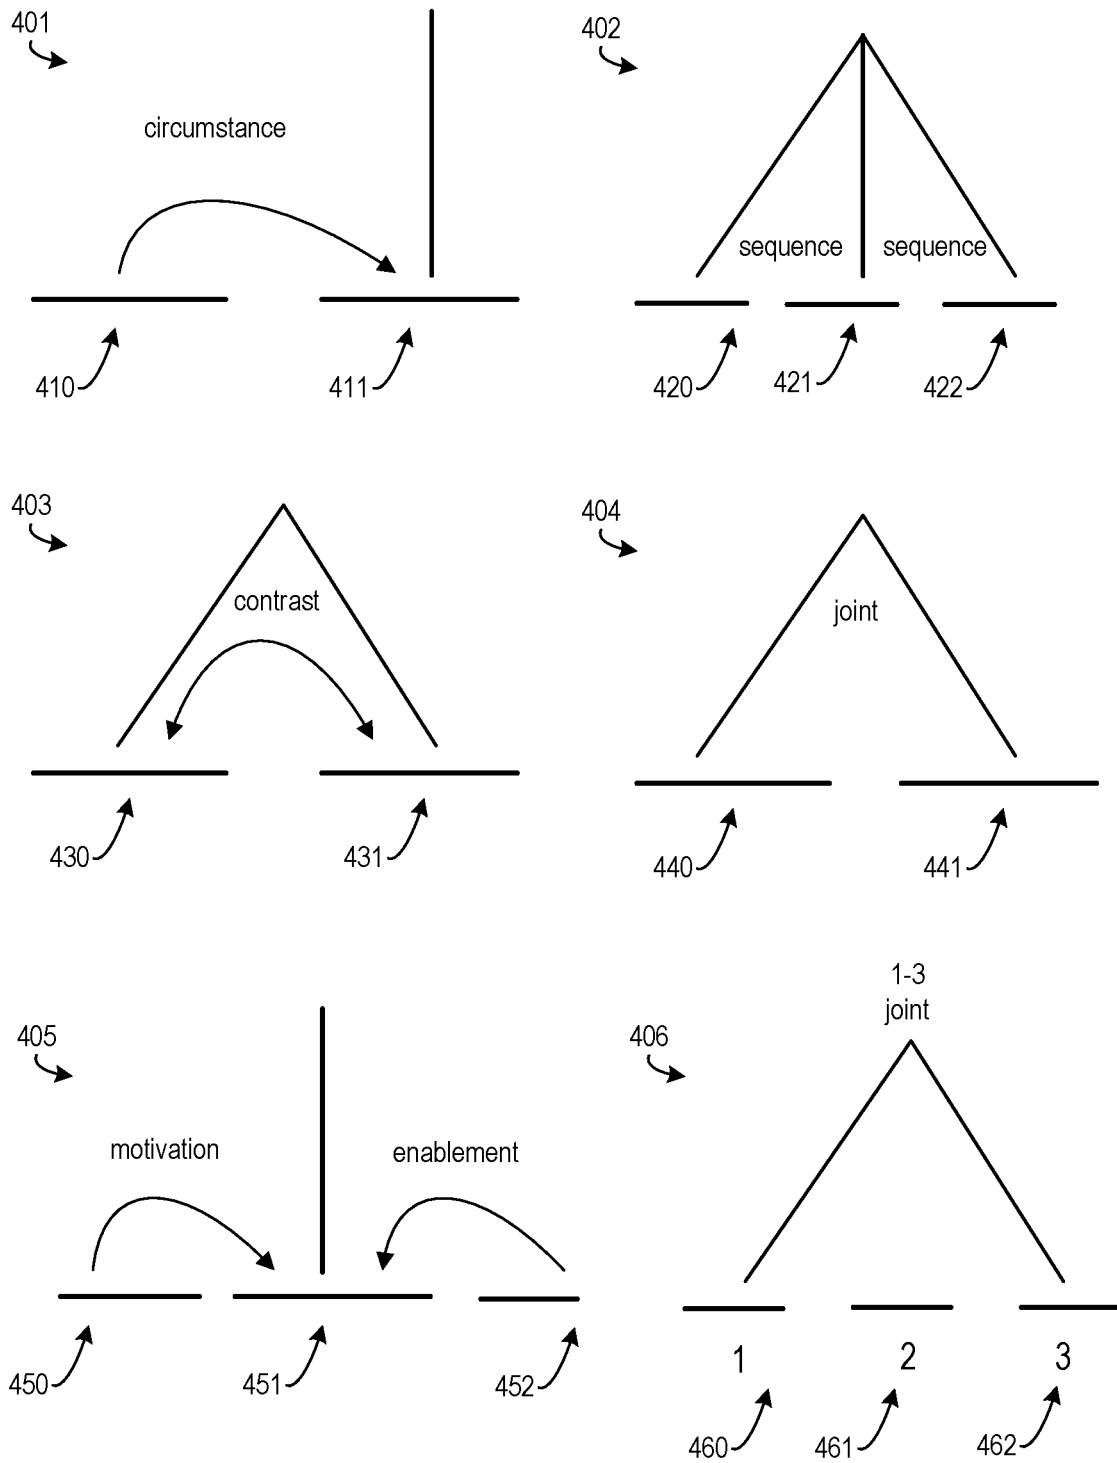


FIG. 4

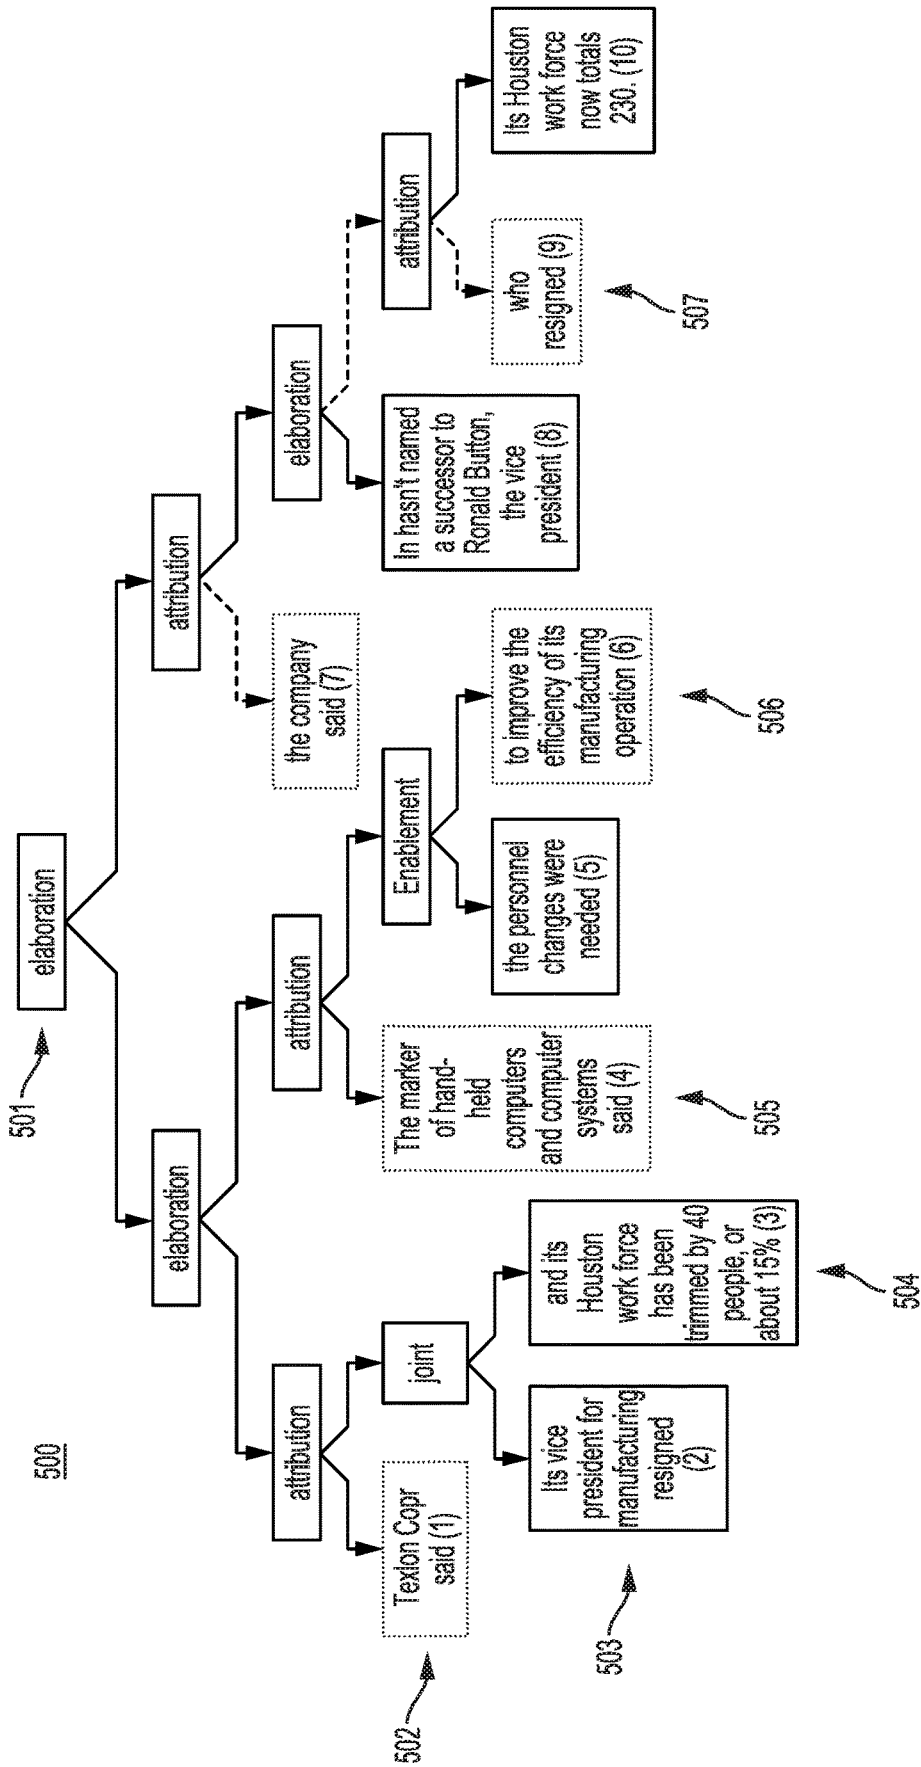


FIG. 5

600

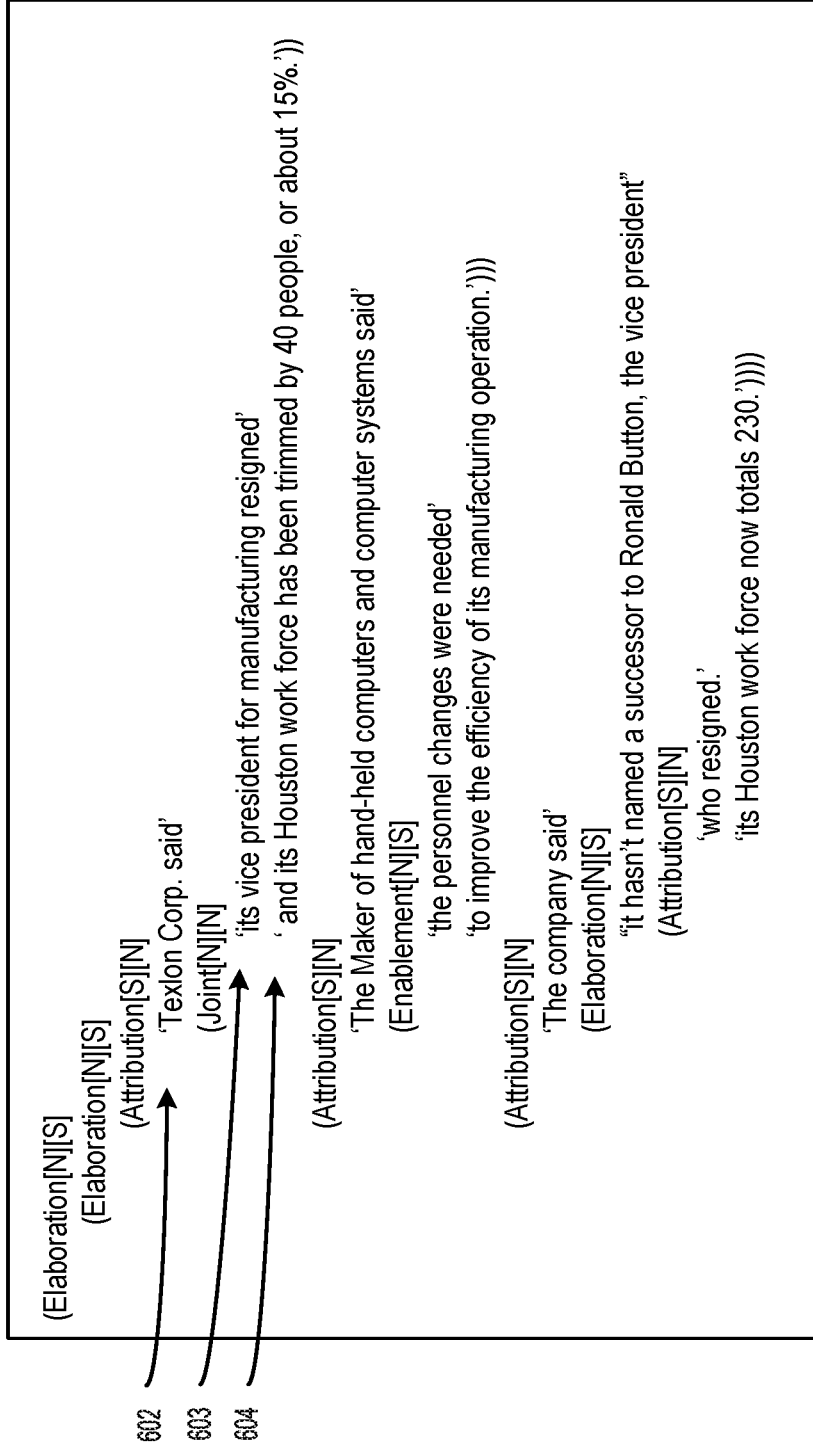


FIG. 6

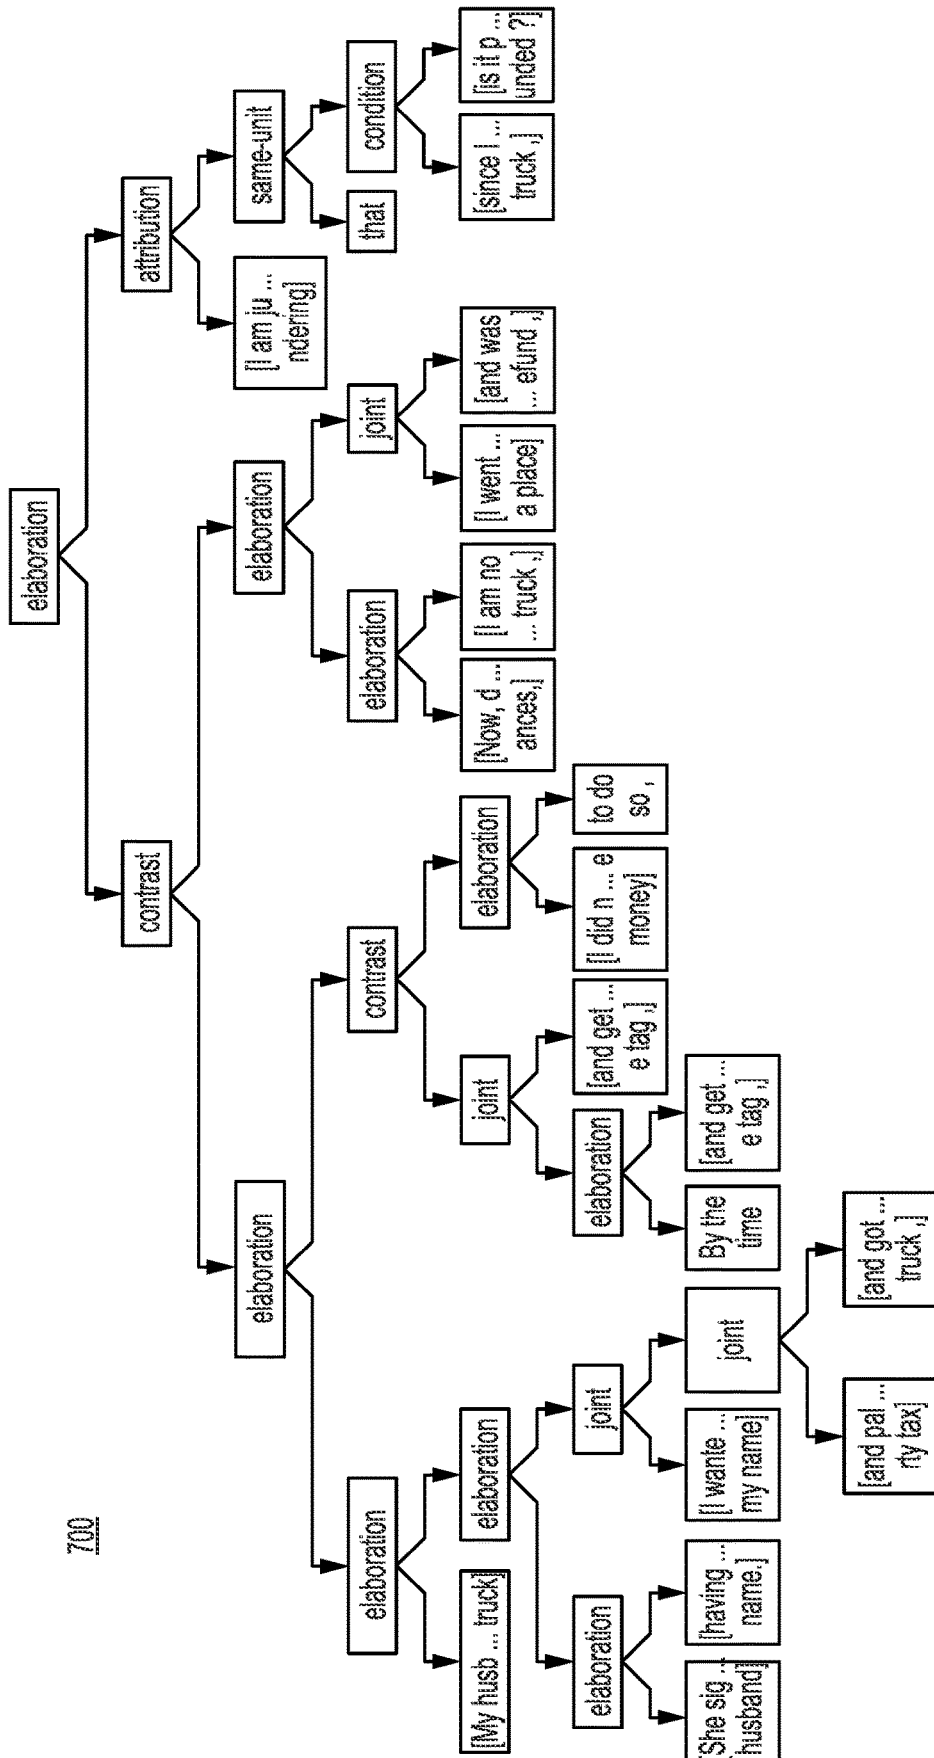


FIG. 7



900

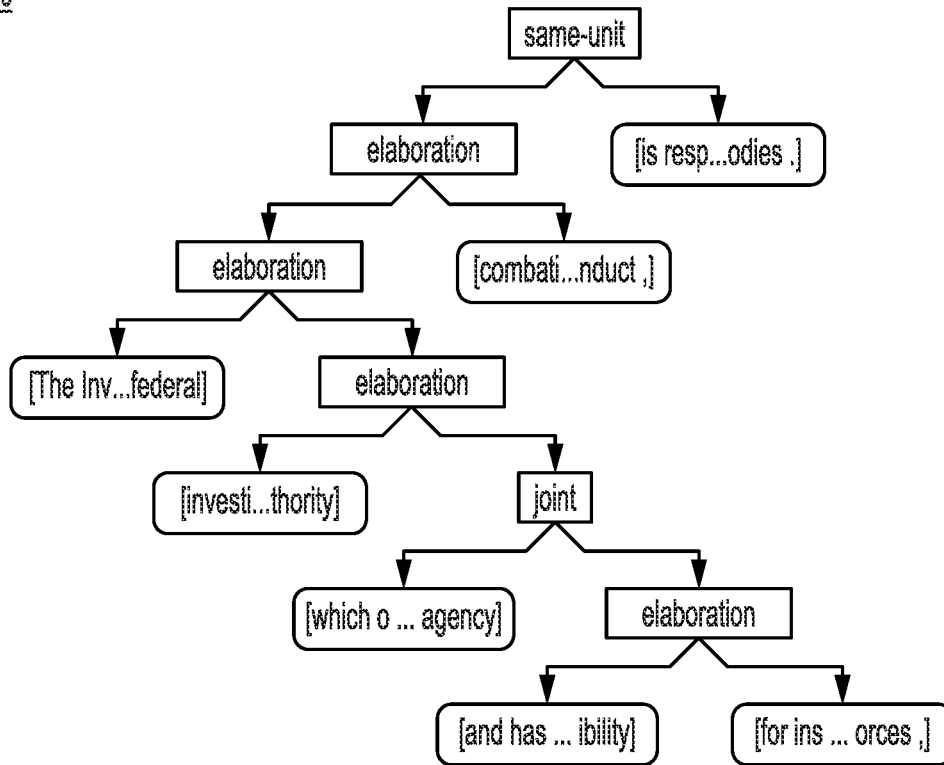


FIG. 9

1000




FIG. 10

1100

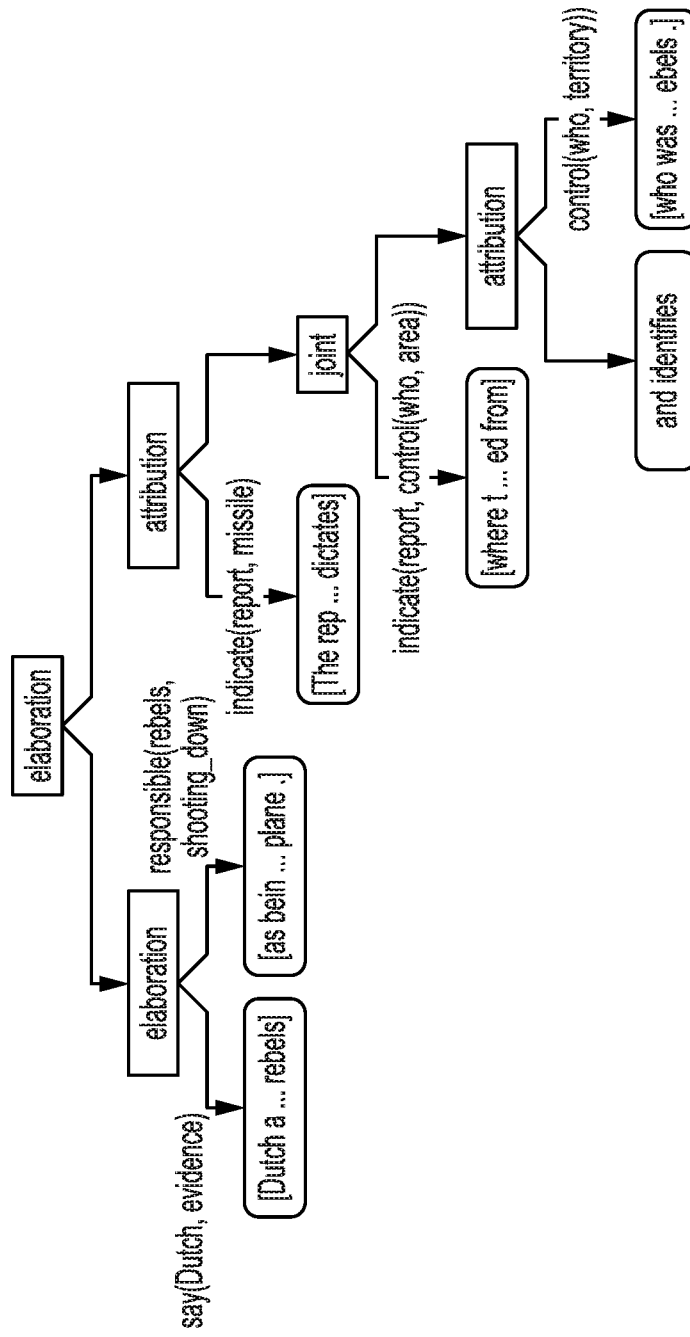


FIG. 11



1200

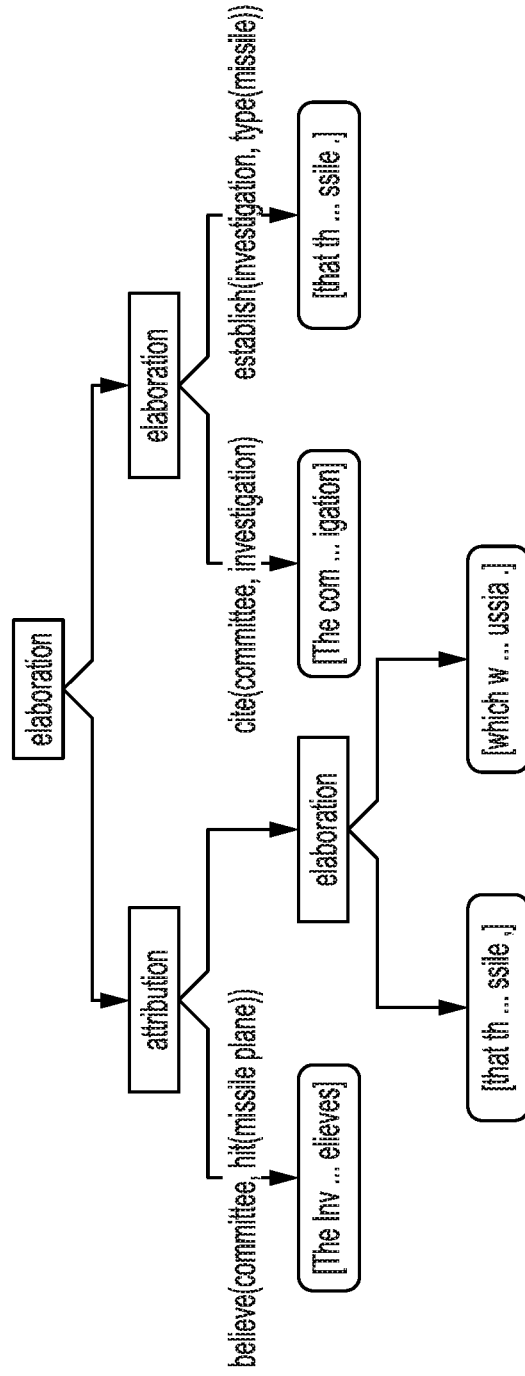


FIG. 12

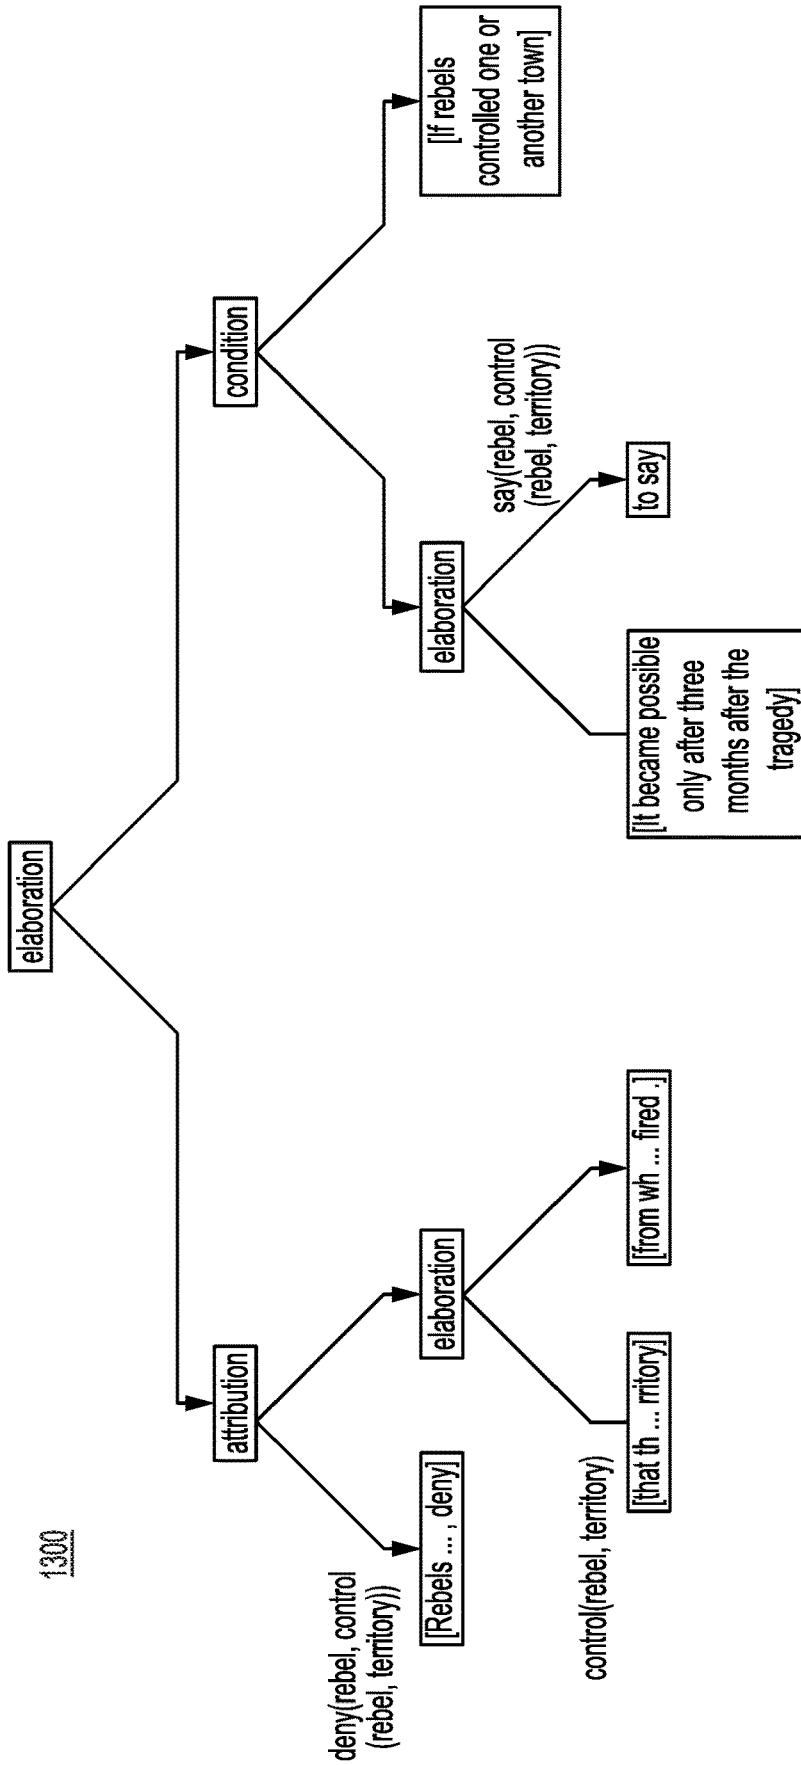
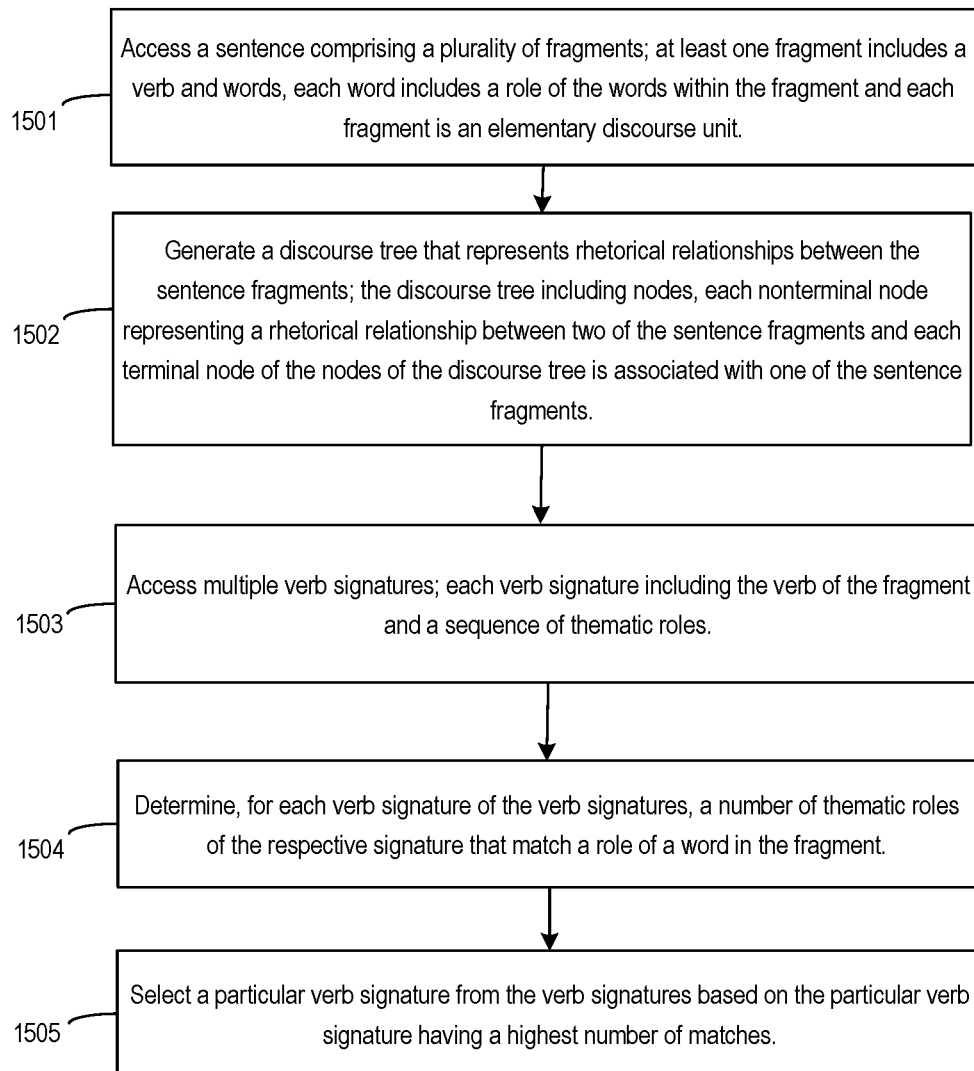


FIG. 13



1500**FIG. 15**

1600

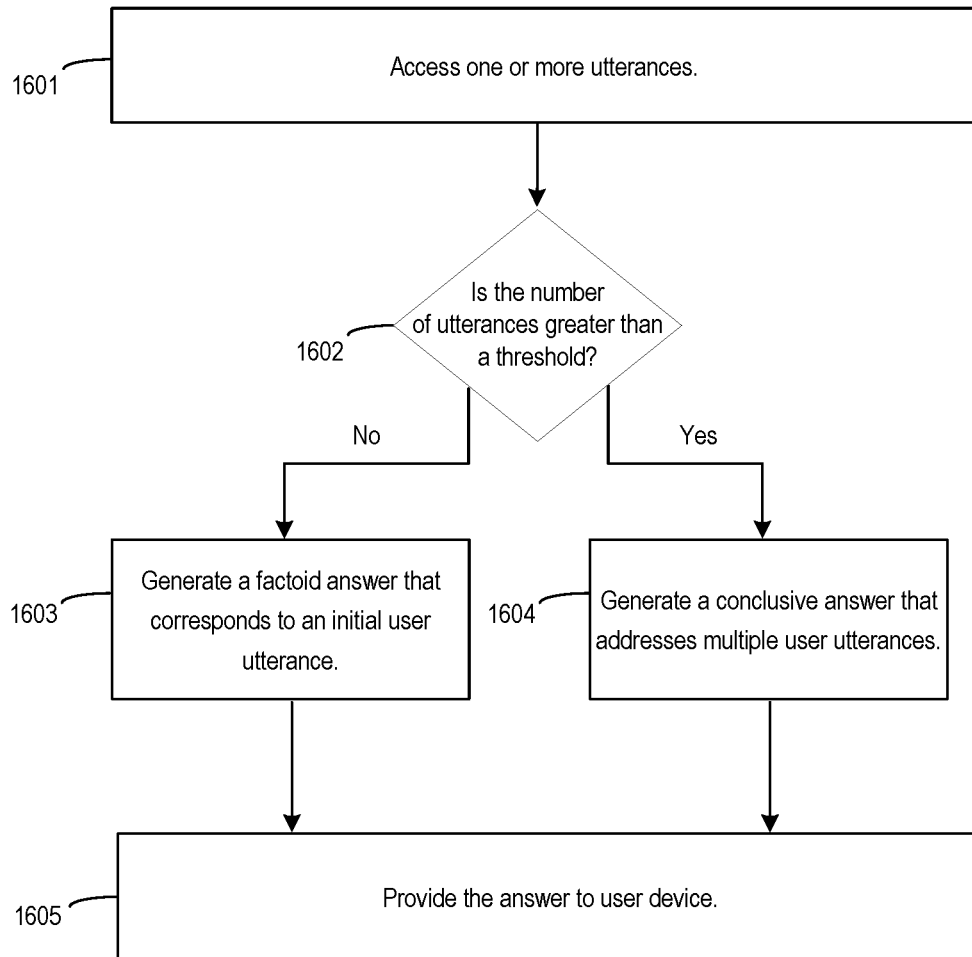


FIG. 16

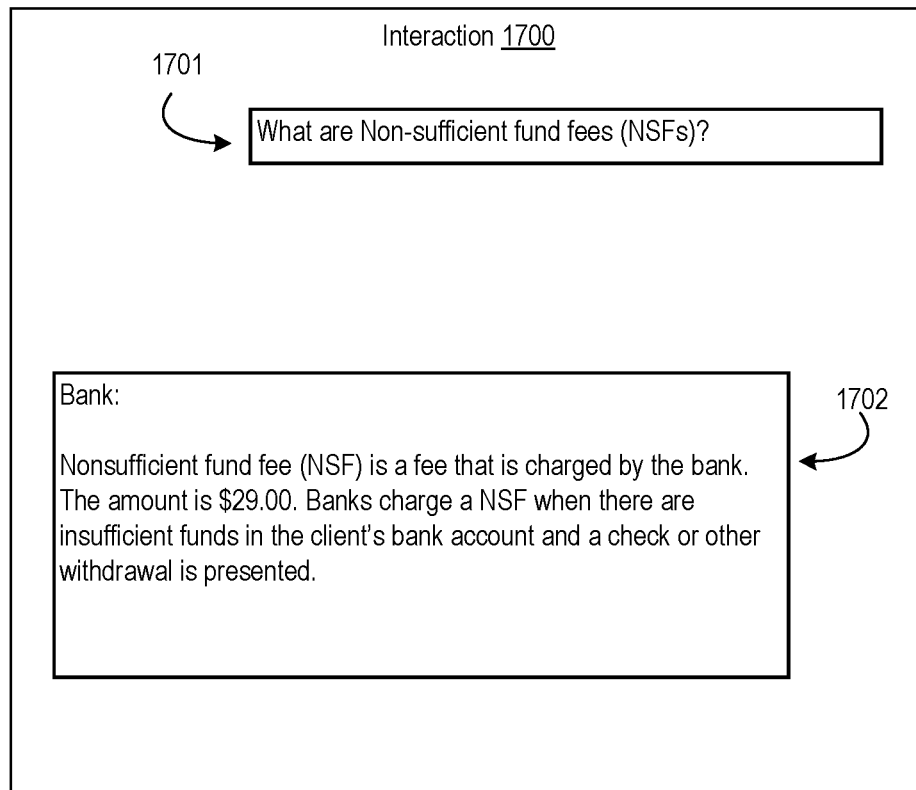


FIG. 17

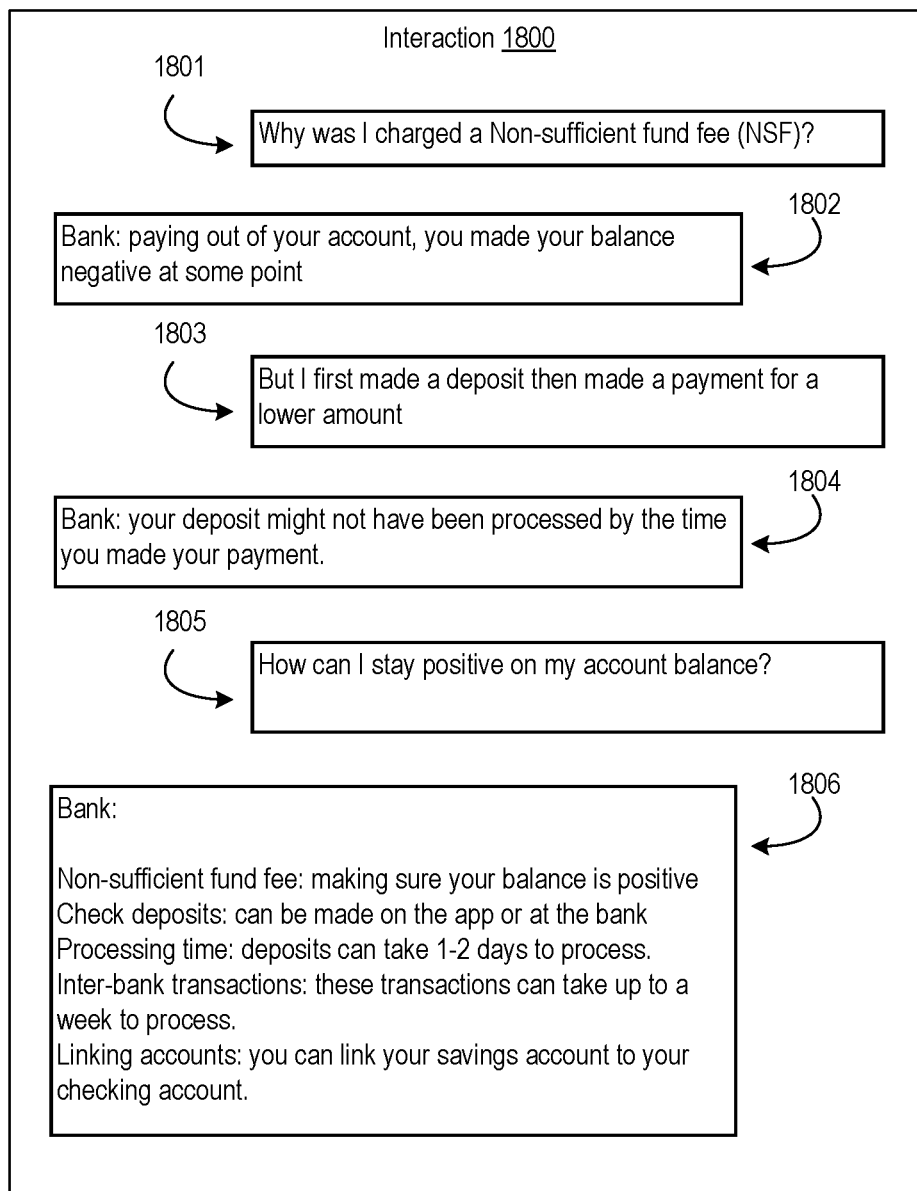
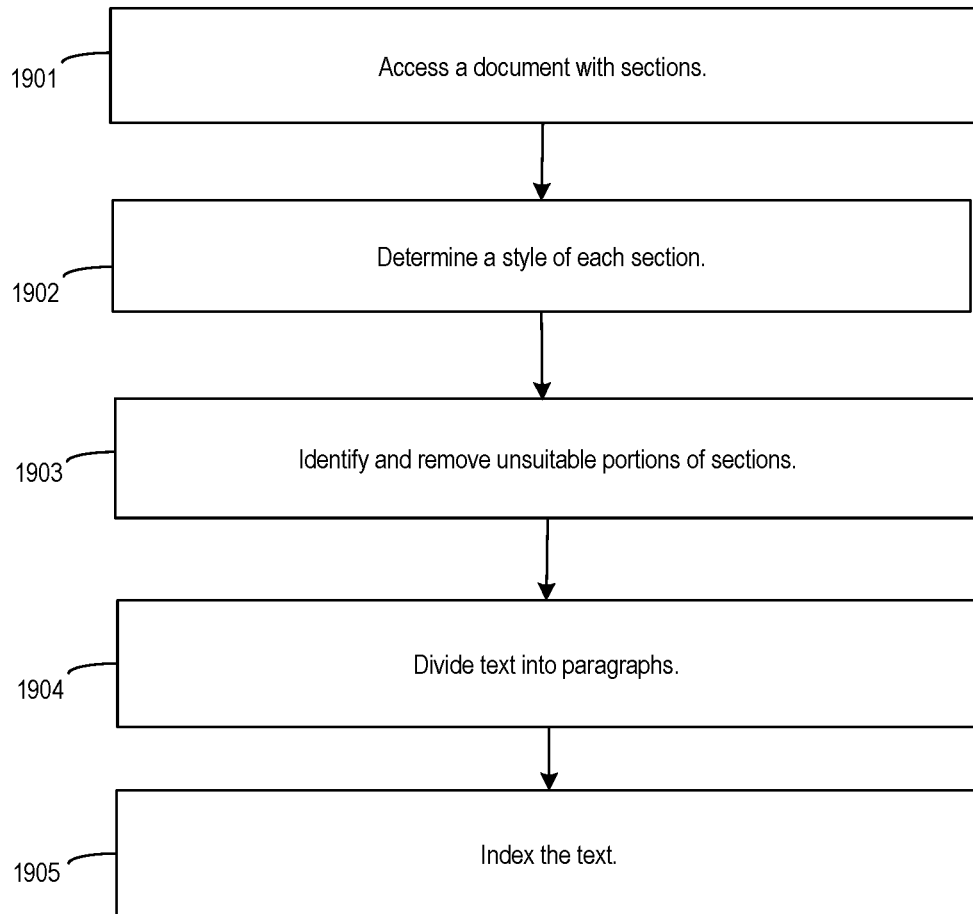


FIG. 18

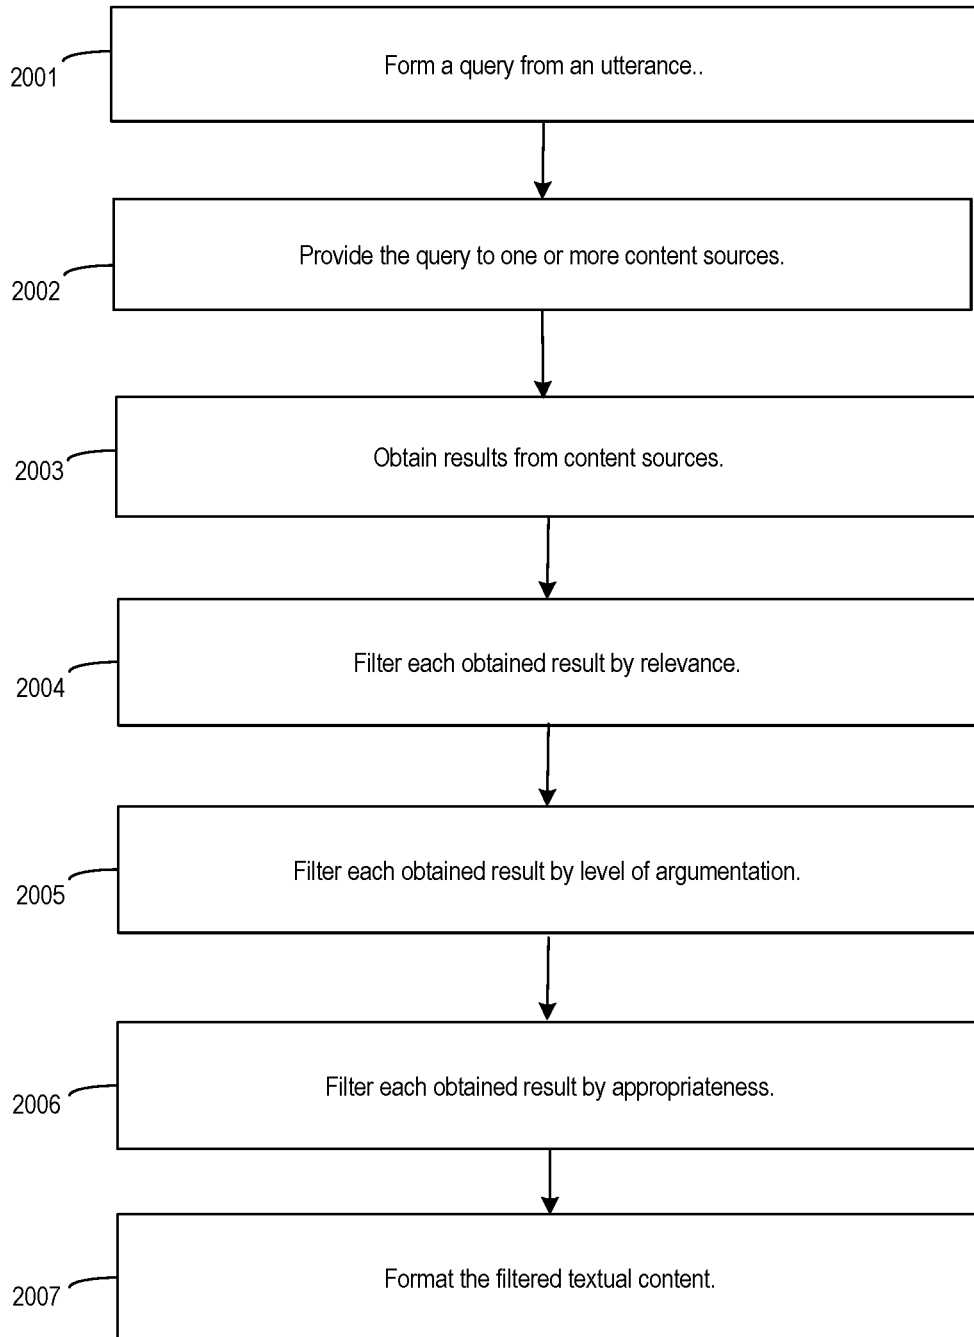
1900



**FIG. 19**



2000



**FIG. 20**

2100




FIG. 21

2200

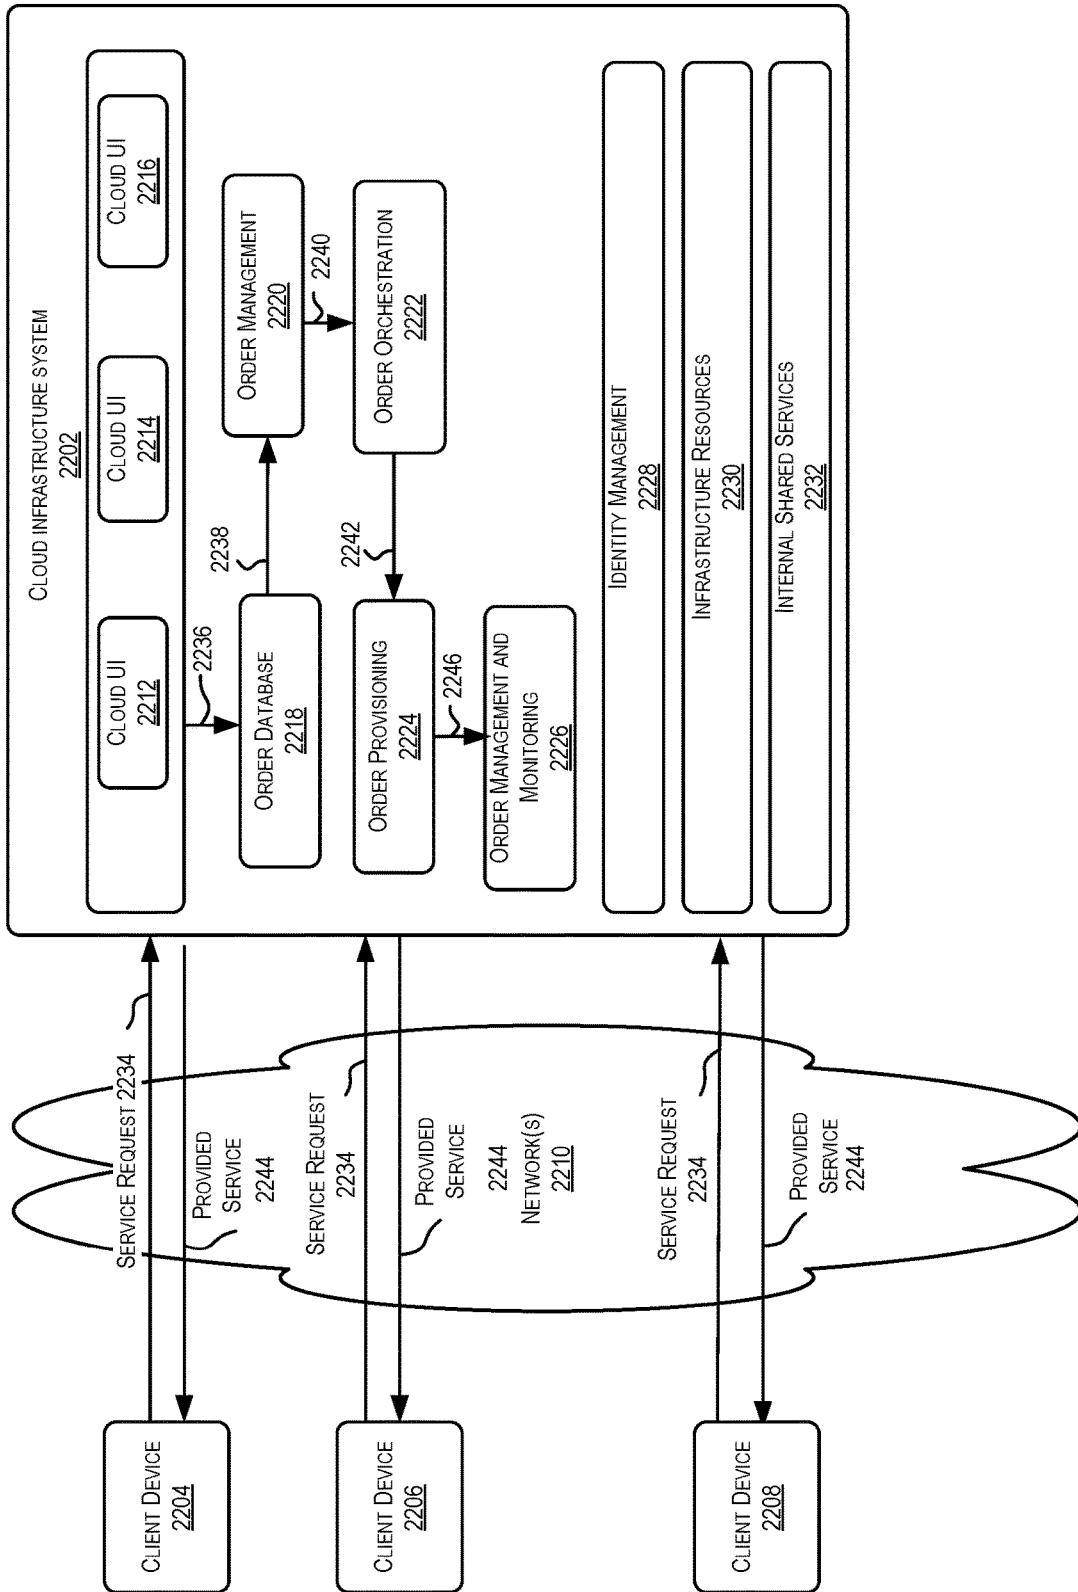


FIG. 22




FIG. 23

## CONSTRUCTING CONCLUSIVE ANSWERS FOR AUTONOMOUS AGENTS

### CROSS-REFERENCES TO RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application No. 62/746,261, filed Oct. 16, 2018, which is incorporated by reference in its entirety.

### TECHNICAL FIELD

This disclosure is generally concerned with linguistics. More specifically, within an interactive session between an autonomous agent and a user, certain embodiments formulate and provide a conclusive answer.

### BACKGROUND

Linguistics is the scientific study of language. One aspect of linguistics is the application of computer science to human natural languages such as English. Due to the greatly increased speed of processors and capacity of memory, computer applications of linguistics are on the rise. For example, computer-enabled analysis of language discourse facilitates numerous applications such as automated agents that can answer questions from users. The use autonomous agents to answer questions, facilitate discussion, manage dialogues, and provide social promotion is increasingly popular.

Following an interactive session with an autonomous agent, a comprehensive, conclusive answer of the session may be desired. But existing solutions are unable to create such an answer. In contrast, these solutions are only able to provide a short reply of text that is derived from a traditional search index. These texts fail to be comprehensive.

As such, solutions are needed that can automatically generate a comprehensive answer for use in a session with an autonomous agent.

### SUMMARY

Techniques are described herein for enabling autonomous agents to generate conclusive answers. An example of a conclusive answer is text that addresses concerns of a user who is interacting with an autonomous agent. For example, an autonomous agent interacts with a user device, answering user utterances, for example questions or concerns. Based on the interactions, the autonomous agent determines that a conclusive answer is appropriate. The autonomous agent formulates the conclusive answer, which addresses multiple user utterances. The conclusive answer provided to the user device.

In an example, a method of computationally fortifying an answer using syntactic parse trees is provided. The method includes accessing a seed sentence including a first set of text fragments. The method further includes syntactically parsing the seed sentence to generate a first syntactic parse tree. The method further includes obtaining a search result by providing, to a search engine, at least one of the first set of text fragments, wherein the search result includes a second text fragment. The method further includes syntactically parsing the search result to generate a second syntactic parse tree. The method further includes calculating, a relevancy metric for the search result by computing a maximum common subgraph between the first syntactic parse tree and the second syntactic parse tree. The method

further includes identifying the search result as an additional fragment based on a determination that the relevancy metric is greater than a first threshold. The method further includes constructing a paragraph from the additional fragment. The method further includes providing the paragraph to a user device.

In another example, a method of computationally fortifying an answer using syntactic parse trees is provided. The method includes accessing a first seed sentence including text fragments and a second seed sentence including text fragments. The method includes generating, from the first seed sentence, a first syntactic parse tree. The method includes generating, from the second seed sentence, a second syntactic parse tree. The method includes identifying, from the first syntactic parse tree, a first entity. The method includes identifying, from the second syntactic parse tree, the first entity and a second entity. The method includes obtaining search results by providing the second entity to a search engine. The method includes generating, from the search result, a second syntactic parse tree. The method includes calculating a relevancy metric for the search result by computing a maximum common subgraph between the second syntactic parse tree and the second syntactic parse tree. The method includes identifying the search result as an additional fragment based on (i) a determination that the relevancy metric is greater than a first threshold. The method includes constructing a paragraph from fragments of the first seed sentence, fragments of the second seed sentence and additional fragment. The method includes providing the paragraph to a user device.

In yet another example, a system includes a computer-readable medium storing non-transitory computer-executable program instructions; and a processing device communicatively coupled to the computer-readable medium for executing the non-transitory computer-executable program instructions. Executing the non-transitory computer-executable program instructions configures the processing device to perform operations. The operations include accessing a seed sentence including a first set of text fragments. The operations include syntactically parsing the seed sentence to generate a first syntactic parse tree. The operations include obtaining a search result by providing at least one of the first set of text fragments to a search engine, wherein the search result includes a second text fragment. The operations include syntactically parsing the search result to generate a second syntactic parse tree. The operations include calculating, a relevancy metric for the search result by computing a maximum common subgraph between the first syntactic parse tree and the second syntactic parse tree. The operations include identifying the search result as an additional fragment based on a determination that the relevancy metric is greater than a first threshold. The operations include constructing a paragraph from the additional fragment; and providing the paragraph to a user device.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows an exemplary conclusive answer generating environment, in accordance with an aspect of the present disclosure.

FIG. 2 depicts an example of a discourse tree in accordance with an aspect.

FIG. 3 depicts a further example of a discourse tree in accordance with an aspect.

FIG. 4 depicts illustrative schemas, in accordance with an aspect.

FIG. 5 depicts a node-link representation of the hierarchical binary tree in accordance with an aspect.

FIG. 6 depicts an exemplary indented text encoding of the representation in FIG. 5 in accordance with an aspect.

FIG. 7 depicts an exemplary DT for an example request about property tax in accordance with an aspect.

FIG. 8 depicts an exemplary response for the question represented in FIG. 7.

FIG. 9 illustrates a discourse tree for an official answer in accordance with an aspect.

FIG. 10 illustrates a discourse tree for a raw answer in accordance with an aspect.

FIG. 11 illustrates a communicative discourse tree for a claim of a first agent in accordance with an aspect.

FIG. 12 illustrates a communicative discourse tree for a claim of a second agent in accordance with an aspect.

FIG. 13 illustrates a communicative discourse tree for a claim of a third agent in accordance with an aspect.

FIG. 14 illustrates parse thickets in accordance with an aspect.

FIG. 15 illustrates an exemplary process for building a communicative discourse tree in accordance with an aspect.

FIG. 16 illustrates an exemplary process for determining an appropriate type of answer to be provided by an autonomous agent, in accordance with an aspect.

FIG. 17 illustrates an example of a conversation flow with a factoid answer, in accordance with an aspect.

FIG. 18 illustrates an example of a conversation flow with a conclusive answer, in accordance with an aspect.

FIG. 19 illustrates an exemplary process for converting textual content for use in forming a conclusive answer, in accordance with an aspect.

FIG. 20 illustrates an exemplary process for building a conclusive answer, in accordance with an aspect.

FIG. 21 depicts a simplified diagram of a distributed system for implementing one of the aspects

FIG. 22 is a simplified block diagram of components of a system environment by which services provided by the components of an aspect system may be offered as cloud services in accordance with an aspect.

FIG. 23 illustrates an exemplary computer system, in which various aspects of the present invention may be implemented.

### DETAILED DESCRIPTION

Aspects disclosed herein provide technical improvements to the area of computer-implemented linguistics, autonomous agents (chatbots), identifying a need for and generating conclusive answers. A conclusive answer is text that addresses concerns of a user who is interacting with an autonomous agent. In an example, a user device is interacting with an autonomous agent in an interactive session. During the session, the user identifies one or more issues. The autonomous agent determines that a conclusive answer should be provided. The agent determines components to be included in the conclusive answer, identifies and formats text that addresses is responsive to the necessary components, and generates the conclusive answer. The conclusive answer can in turn be provided to a display such as a display of a mobile device.

Generating a conclusive answer can involve determining a suitable structure of the conclusive answer and creating the conclusive answer based on the determined structure. The structure of the conclusive answer can involve identifying one or more entities present in an initial user utterance and/or subsequent user utterances. More specifically, certain

aspects obtain content from a text corpus, verify that the obtained text is relevant, and verify that the obtained text is appropriate (e.g., in style).

Different technical approaches can be used. For example, machine learning, keyword-based approaches, and/or search engineering techniques can be used. Additionally or alternatively, some aspects represent rhetorical relationships between one or more sentences in communicative discourse trees. “Communicative discourse trees” or “CDTs” include discourse trees that are supplemented with communicative actions. Communicative discourse trees combine rhetoric information with communicative actions. A communicative action is a cooperative action undertaken by individuals based on mutual deliberation and argumentation. By incorporating labels that identify communicative actions, learning of communicative discourse trees can occur over a richer features set than simply rhetoric relations and syntax of elementary discourse units (EDUs). With such a feature set, additional techniques such as classification can be used to determine a level of rhetoric agreement between questions and answers or request-response pairs, thereby enabling improved automated agents. Representing text as communicative discourse trees can facilitate different analysis such as detection of argumentation and verification that two sentences are of a similar style.

Technical advantages of some aspects include improved autonomous agents and improved search engine performance over traditional statistical-based approaches. As discussed, existing solutions are unable to create conclusive answer. Instead, such solutions provide short replies by text indexed by a traditional search index. Statistical learning and especially deep learning-based agents attempt to use learning to tailor answers to a user session, but only brief texts can be obtained as a result. These texts can be meaningful but fail to be comprehensive.

#### Certain Definitions

As used herein, an “entity” has an independent and distinct existence. Examples includes objects, places, and persons. An entity can also be a subject or topic such as “electric cars” or “brakes.”

As used herein, a named entity is an entity that is identified by a name. Examples of named entities are “France,” “John Doe,” and “Russell Square” are named entities

As used herein, “rhetorical structure theory” is an area of research and study that provided a theoretical basis upon which the coherence of a discourse could be analyzed.

As used herein, “discourse tree” or “DT” refers to a structure that represents the rhetorical relations for a sentence of part of a sentence.

As used herein, a “rhetorical relation,” “rhetorical relationship,” or “coherence relation” or “discourse relation” refers to how two segments of discourse are logically connected to one another. Examples of rhetorical relations include elaboration, contrast, and attribution.

As used herein, a “sentence fragment,” or “fragment” is a part of a sentence that can be divided from the rest of the sentence. A fragment is an elementary discourse unit. For example, for the sentence “Dutch accident investigators say that evidence points to pro-Russian rebels as being responsible for shooting down the plane,” two fragments are “Dutch accident investigators say that evidence points to pro-Russian rebels” and “as being responsible for shooting down the plane.” A fragment can, but need not, include a verb.

As used herein, “signature” or “frame” refers to a property of a verb in a fragment. Each signature can include one or

more thematic roles. For example, for the fragment “Dutch accident investigators say that evidence points to pro-Russian rebels,” the verb is “say” and the signature of this particular use of the verb “say” could be “agent verb topic” where “investigators” is the agent and “evidence” is the topic.

As used herein, “thematic role” refers to components of a signature used to describe a role of one or more words. Continuing the previous example, “agent” and “topic” are thematic roles.

As used herein, “nuclearity” refers to which text segment, fragment, or span, is more central to a writer’s purpose. The nucleus is the more central span, and the satellite is the less central one.

As used herein, “coherency” refers to the linking together of two rhetorical relations.

As used herein, “communicative verb” is a verb that indicates communication. For example, the verb “deny” is a communicative verb.

As used herein, “communicative action” describes an action performed by one or more agents and the subjects of the agents.

Turning now to the Figures, FIG. 1 shows an exemplary conclusive answer generating environment, in accordance with an aspect of the present disclosure. FIG. 1 depicts computing device 101, data network 104, server 160, and user device 170. In the example depicted by FIG. 1, user device 170 interacts with computing device 101 over data network 104. As depicted by FIG. 1, a user requested a recommendation about investment, received it and expressed her doubts. The chatbot provides the comprehensive conclusive answer entitled ‘Why a 61% revenue growth is not enough for Alibaba’ with detailed information on competitiveness including the charts.

User device 170 can be any mobile device such as a mobile phone, smart phone, tablet, laptop, smart watch, and the like. User device 170 communicates via data network 104 to server 160 and/or computing device 101. Data network 104 can be any public or private network, wired or wireless network, Wide Area Network, Local Area Network, or the Internet.

In an example, user device 170 can provide questions, e.g., questions 180 and 180, which can be from a user, to computing device 101. Computing device 101 can determine a suitable answer, e.g., answer 181 which is responsive to question 180, or conclusive answer 183, and provide the answer to user device 170. In particular, computing device 101 can determine that a conclusive answer is appropriate, generate the conclusive answer, and transmit the conclusive answer to user device 170.

As depicted, user device 170 includes a display 172. The display 172 depicts a chat history, or an interactive session with the autonomous agent implemented by computing device 101. As can be seen, a user initially asks a question 180, which reads “I’m curious if I should invest in Alibaba only or diversify into FANG (Facebook, Amazon, Google, Netflix).” In response, computing device 101 determines a suitable answer 181, which reads “Consider Alibaba only because it is diversifying itself into other areas beyond buying and selling stuff” Computing device 101 then receives user question 182, which reads “Is it successful in going into these other areas? Can it really compete with FANG?” In response, computing device 101 generates conclusive answer 183, which reads “Alibaba reported a 61% revenue growth for the quarter ending June 30. That growth rate was higher than all the peers in the BAT (Baldu, Tencent) and FANG (Facebook, Amazon, Netflix, and

Google) . . . Alibaba isn’t just about selling and buying goods anymore. Much like what Amazon is doing in the U.S., the e-commerce giant has expanded into many other areas. It owns the streaming website Youku (China’s YouTube), local services perform Koubel (similar to Yelo) and food delivery platform Ele.me (like Seamless).” As can be seen, conclusive answer 183 addresses both question 180 and question 182, because conclusive answer 183 addresses a main entity in question 180 “Alibaba” and a specific request about Alibaba’s success as present in question 182.

Computing device 101 implements an autonomous agent via conclusive answer application 102, text corpus 105, rhetoric agreement classifier 120, and training data 125. Conclusive answer application 102 can implement any of the functionality described herein, including analyzing text and questions, determining answers, and determining conclusive answers. Text corpus 105 includes text such as electronic documents, webpages, customer support conversation logs, internal issue resolution logs, or other correspondence. Conclusive answer application 102 can index the content of text corpus 105 for quicker access. Rhetoric agreement classifier 120 can include one or more machine learning models that are trained to perform one or more different actions. For example, rhetoric agreement classifier 120 can determine a level of rhetoric agreement and/or style between a question and a candidate answer. In this manner, conclusive answer application 102 can provide the most suitable answer to the mobile device 160.

Server 160 can be a public or private internet server, such as a public database of user questions and answers. Server 160 can implement any of the functionality described herein, including the functionality associated with computing device 101. In some cases, server 160 can complement the functionality of computing device 101, for example, by providing additional databases, corpuses of text, etc. Examples of additional functionality that can be performed by computing device 101 and/or server 160 are discussed with respect to FIGS. 15, 16, 19, and 20. For example, computing device 101 can perform process 1600 described with respect to FIG. 16 to determine whether a conclusive answer is necessary. If a conclusive answer is necessary, then computing device 101 can perform process 2000 described with respect to FIG. 20.

Rhetoric Structure Theory and Discourse Trees

Linguistics is the scientific study of language. For example, linguistics can include the structure of a sentence (syntax), e.g., subject-verb-object, the meaning of a sentence (semantics), e.g. dog bites man vs. man bites dog, and what speakers do in conversation, i.e., discourse analysis or the analysis of language beyond the sentence.

The theoretical underpinnings of discourse, Rhetoric Structure Theory (RST), can be attributed to Mann, William and Thompson, Sandra, “Rhetorical structure theory: A Theory of Text organization,” *Text-Interdisciplinary Journal for the Study of Discourse*, 8(3):243-281, 1988. Similar to how the syntax and semantics of programming language theory helped enable modern software compilers, RST helped enabled the analysis of discourse. More specifically RST posits structural blocks on at least two levels, a first level such as nuclearity and rhetorical relations, and a second level of structures or schemas. Discourse parsers or other computer software can parse text into a discourse tree.

Rhetoric Structure Theory models logical organization of text, a structure employed by a writer, relying on relations between parts of text. RST simulates text coherence by forming a hierarchical, connected structure of texts via discourse trees. Rhetoric relations are split into the classes of

coordinate and subordinate; these relations hold across two or more text spans and therefore implement coherence. These text spans are called elementary discourse units (EDUs). Clauses in a sentence and sentences in a text are logically connected by the author. The meaning of a given sentence is related to that of the previous and the following sentences. This logical relation between clauses is called the coherence structure of the text. RST is one of the most popular theories of discourse, being based on a tree-like discourse structure, discourse trees (DTs). The leaves of a DT correspond to EDUs, the contiguous atomic text spans. Adjacent EDUs are connected by coherence relations (e.g., Attribution, Sequence), forming higher-level discourse units. These units are then also subject to this relation linking. EDUs linked by a relation are then differentiated based on their relative importance: nuclei are the core parts of the relation, while satellites are peripheral ones. As discussed, in order to determine accurate request-response pairs, both topic and rhetorical agreement are analyzed. When a speaker answers a question, such as a phrase or a sentence, the speaker's answer should address the topic of this question. In the case of an implicit formulation of a question, via a seed text of a message, an appropriate answer is expected not only maintain a topic, but also match the generalized epistemic state of this seed.

Rhetoric Relations

As discussed, aspects described herein use communicative discourse trees. Rhetorical relations can be described in different ways. For example, Mann and Thompson describe twenty-three possible relations. C. Mann, William & Thompson, Sandra. (1987) ("Mann and Thompson"). Rhetorical Structure Theory: A Theory of Text Organization. Other numbers of relations are possible.

Relation Name	Nucleus	Satellite
Antithesis	ideas favored by the author	ideas disfavored by the author
Background	text whose understanding is being facilitated	text for facilitating understanding
Circumstance	text expressing the events or ideas occurring in the interpretive context	an interpretive context of situation or time
Concession	situation affirmed by author	situation which is apparently inconsistent but also affirmed by author
Condition	action or situation whose occurrence results from the occurrence of the conditioning situation	conditioning situation
Elaboration Enablement	basic information an action	additional information intended to aid the reader in performing an action
Evaluation	a situation	an evaluative comment about the situation
Evidence	a claim	information intended to increase the reader's belief in the claim
Interpretation	a situation	an interpretation of the situation
Justify	text	information supporting the writer's right to express the text
Motivation	an action	information intended to increase the reader's desire to perform the action

-continued

Relation Name	Nucleus	Satellite
5 Non-volitional Cause	a situation	another situation which causes that one, but not by anyone's deliberate action
Non-volitional Result	a situation	another situation which is caused by that one, but not by anyone's deliberate action
10 Otherwise (anti conditional)	action or situation whose occurrence results from the lack of occurrence of the conditioning situation	conditioning situation
Purpose	an intended situation	the intent behind the situation
15 Restatement	a situation	a reexpression of the situation
Solutionhood	a situation or method supporting full or partial satisfaction of the need	a request, problem, or other expressed need
Summary	text	a short summary of that text
Volitional Cause	a situation	another situation which causes that one, by someone's deliberate action
20 Volitional Result	a situation	another situation which is caused by that one, by someone's deliberate action

25 Some empirical studies postulate that the majority of text is structured using nucleus-satellite relations. See Mann and Thompson. But other relations do not carry a definite selection of a nucleus. Examples of such relations are shown below.

Relation Name	Span	Other Span
35 Contrast	One alternate	The other alternate
Joint	(unconstrained)	(unconstrained)
List	An item	A next item
Sequence	An item	A next item

FIG. 2 depicts an example of a discourse tree, in accordance with an aspect. FIG. 2 includes discourse tree 200. Discourse tree includes text span 201, text span 202, text span 203, relation 210 and relation 228. The numbers in FIG. 2 correspond to the three text spans. FIG. 3 corresponds to the following example text with three text spans numbered 1, 2, 3:

1. Honolulu, Hi. will be site of the 2017 Conference on Hawaiian History
2. It is expected that 200 historians from the U.S. and Asia will attend
3. The conference will be concerned with how the Polynesians sailed to Hawaii

For example, relation 210, or elaboration, describes the relationship between text span 201 and text span 202. Relation 228 depicts the relationship, elaboration, between text span 203 and 204. As depicted, text spans 202 and 203 elaborate further on text span 201. In the above example, given a goal of notifying readers of a conference, text span 1 is the nucleus. Text spans 2 and 3 provide more detail about the conference. In FIG. 2, a horizontal number, e.g., 1-3, 1, 2, 3 covers a span of text (possibly made up of further spans); a vertical line signals the nucleus or nuclei; and a curve represents a rhetoric relation (elaboration) and the direction of the arrow points from the satellite to the nucleus. If the text span only functions as a satellite and not as a nuclei, then deleting the satellite would still leave a coherent text. If from FIG. 2 one deletes the nucleus, then text spans 2 and 3 are difficult to understand.



FIG. 3 depicts a further example of a discourse tree in accordance with an aspect. FIG. 3 includes components 301 and 302, text spans 305-307, relation 310 and relation 328. Relation 310 depicts the relationship 310, enablement, between components 306 and 305, and 307, and 305. FIG. 3 refers to the following text spans:

1. The new Tech Report abstracts are now in the journal area of the library near the abridged dictionary.

2. Please sign your name by any means that you would be interested in seeing.

3. Last day for sign-ups is 31 May.

As can be seen, relation 328 depicts the relationship between entity 307 and 306, which is enablement. FIG. 3 illustrates that while nuclei can be nested, there exists only one most nuclear text span.

#### Constructing a Discourse Tree

Discourse trees can be generated using different methods. A simple example of a method to construct a DT bottom up is:

- (1) Divide the discourse text into units by:
  - (a) Unit size may vary, depending on the goals of the analysis
  - (b) Typically, units are clauses
- (2) Examine each unit, and its neighbors. Is there a relation holding between them?
- (3) If yes, then mark that relation.
- (4) If not, the unit might be at the boundary of a higher-level relation. Look at relations holding between larger units (spans).
- (5) Continue until all the units in the text are accounted for.

Mann and Thompson also describe the second level of building block structures called schemas applications. In RST, rhetoric relations are not mapped directly onto texts; they are fitted onto structures called schema applications, and these in turn are fitted to text. Schema applications are derived from simpler structures called schemas (as shown by FIG. 4). Each schema indicates how a particular unit of text is decomposed into other smaller text units. A rhetorical structure tree or DT is a hierarchical system of schema applications. A schema application links a number of consecutive text spans, and creates a complex text span, which can in turn be linked by a higher-level schema application. RST asserts that the structure of every coherent discourse can be described by a single rhetorical structure tree, whose top schema creates a span encompassing the whole discourse.

FIG. 4 depicts illustrative schemas, in accordance with an aspect. FIG. 4 shows a joint schema is a list of items consisting of nuclei with no satellites. FIG. 4 depicts schemas 401-406. Schema 401 depicts a circumstance relation between text spans 410 and 428. Schema 402 depicts a sequence relation between text spans 420 and 421 and a sequence relation between text spans 421 and 422. Schema 403 depicts a contrast relation between text spans 430 and 431. Schema 404 depicts a joint relationship between text spans 440 and 441. Schema 405 depicts a motivation relationship between 450 and 451, and an enablement relationship between 452 and 451. Schema 406 depicts joint relationship between text spans 460 and 462. An example of a joint scheme is shown in FIG. 4 for the three text spans below:

1. Skies will be partly sunny in the New York metropolitan area today.

2. It will be more humid, with temperatures in the middle 80's.

3. Tonight will be mostly cloudy, with the low temperature between 65 and 70.

While FIGS. 2-4 depict some graphical representations of a discourse tree, other representations are possible.

FIG. 5 depicts a node-link representation of the hierarchical binary tree in accordance with an aspect. As can be seen from FIG. 5, the leaves of a DT correspond to contiguous non-overlapping text spans called Elementary Discourse Units (EDUs). Adjacent EDUs are connected by relations (e.g., elaboration, attribution . . .) and form larger discourse units, which are also connected by relations. "Discourse analysis in RST involves two sub-tasks: discourse segmentation is the task of identifying the EDUs, and discourse parsing is the task of linking the discourse units into a labeled tree." See Joty, Shafiq R and Giuseppe Carenini, Raymond T Ng, and Yashar Mehdad. 2013. Combining intra- and multi-sentential rhetorical parsing for document-level discourse analysis. In *ACL* (1), pages 486-496.

FIG. 5 depicts text spans that are leaves, or terminal nodes, on the tree, each numbered in the order they appear in the full text, shown in FIG. 6. FIG. 5 includes tree 500. Tree 500 includes, for example, nodes 501-507. The nodes indicate relationships. Nodes are non-terminal, such as node 501, or terminal, such as nodes 502-507. As can be seen, nodes 503 and 504 are related by a joint relationship. Nodes 502, 505, 506, and 508 are nuclei. The dotted lines indicate that the branch or text span is a satellite. The relations are nodes in gray boxes.

FIG. 6 depicts an exemplary indented text encoding of the representation in FIG. 5 in accordance with an aspect. FIG. 6 includes text 600 and text sequences 602-604. Text 600 is presented in a manner more amenable to computer programming. Text sequence 602 corresponds to node 502, sequence 603 corresponds to node 503, and sequence 604 corresponds to node 504. In FIG. 6, "N" indicates a nucleus and "S" indicates a satellite.

#### Examples of Discourse Parsers

Automatic discourse segmentation can be performed with different methods. For example, given a sentence, a segmentation model identifies the boundaries of the composite elementary discourse units by predicting whether a boundary should be inserted before each particular token in the sentence. For example, one framework considers each token in the sentence sequentially and independently. In this framework, the segmentation model scans the sentence token by token, and uses a binary classifier, such as a support vector machine or logistic regression, to predict whether it is appropriate to insert a boundary before the token being examined. In another example, the task is a sequential labeling problem. Once text is segmented into elementary discourse units, sentence-level discourse parsing can be performed to construct the discourse tree. Machine learning techniques can be used.

In one aspect of the present invention, two Rhetorical Structure Theory (RST) discourse parsers are used: CoreNLPProcessor which relies on constituent syntax, and FastNLPProcessor which uses dependency syntax. See Surdeanu, Mihai & Hicks, Thomas & Antonio Valenzuela-Escarcega, Marco. Two Practical Rhetorical Structure Theory Parsers. (2015).

In addition, the above two discourse parsers, i.e., CoreNLPProcessor and FastNLPProcessor use Natural Language Processing (NLP) for syntactic parsing. For example, the Stanford CoreNLP gives the base forms of words, their parts of speech, whether they are names of companies, people, etc., normalize dates, times, and numeric quantities,

mark up the structure of sentences in terms of phrases and syntactic dependencies, indicate which noun phrases refer to the same entities. Practically, RST is a still theory that may work in many cases of discourse, but in some cases, it may not work. There are many variables including, but not limited to, what EDU's are in a coherent text, i.e., what discourse segmenters are used, what relations inventory is used and what relations are selected for the EDUs, the corpus of documents used for training and testing, and even what parsers are used. So for example, in Surdeanu, et al., "Two Practical Rhetorical Structure Theory Parsers," paper cited above, tests must be run on a particular corpus using specialized metrics to determine which parser gives better performance. Thus unlike computer language parsers which give predictable results, discourse parsers (and segmenters) can give unpredictable results depending on the training and/or test text corpus. Thus, discourse trees are a mixture of the predictable arts (e.g., compilers) and the unpredictable arts (e.g., like chemistry where experimentation is needed to determine what combinations will give you the desired results).

In order to objectively determine how good a Discourse analysis is, a series of metrics are being used, e.g., Precision/Recall/F1 metrics from Daniel Marcu, "The Theory and Practice of Discourse Parsing and Summarization," MIT Press, (2000). Precision, or positive predictive value is the fraction of relevant instances among the retrieved instances, while recall (also known as sensitivity) is the fraction of relevant instances that have been retrieved over the total amount of relevant instances. Both precision and recall are therefore based on an understanding and measure of relevance. Suppose a computer program for recognizing dogs in photographs identifies eight dogs in a picture containing 12 dogs and some cats. Of the eight dogs identified, five actually are dogs (true positives), while the rest are cats (false positives). The program's precision is  $\frac{5}{8}$  while its recall is  $\frac{5}{12}$ . When a search engine returns 30 pages only 20 of which were relevant while failing to return 40 additional relevant pages, its precision is  $\frac{20}{30}=\frac{2}{3}$  while its recall is  $\frac{20}{60}=\frac{1}{3}$ . Therefore, in this case, precision is 'how useful the search results are', and recall is 'how complete the results are.'" The F1 score (also F-score or F-measure) is a measure of a test's accuracy. It considers both the precision and the recall of the test to compute the score:  $F1=2 \times ((\text{precision} \times \text{recall}) / (\text{precision} + \text{recall}))$  and is the harmonic mean of precision and recall. The F1 score reaches its best value at 1 (perfect precision and recall) and worst at 0.

#### Autonomous Agents or Chatbots

A conversation between Human A and Human B is a form of discourse. For example, applications exist such as Facebook® Messenger, WhatsApp®, Slack®, SMS, etc., a conversation between A and B may typically be via messages in addition to more traditional email and voice conversations. A chatbot (which may also be called intelligent bots or virtual assistant, etc.) is an "intelligent" machine that, for example, replaces human B and to various degrees mimics the conversation between two humans. An example ultimate goal is that human A cannot tell whether B is a human or a machine (the Turing test, developed by Alan Turing in 1950). Discourse analysis, artificial intelligence, including machine learning, and natural language processing, have made great strides toward the long-term goal of passing the Turing test. Of course, with computers being more and more capable of searching and processing vast repositories of data and performing complex analysis on the data to include predictive analysis, the long-term goal is the chatbot being human-like and a computer combined.

For example, users can interact with the Intelligent Bots Platform through a conversational interaction. This interaction, also called the conversational user interface (UI), is a dialog between the end user and the chatbot, just as between two human beings. It could be as simple as the end user saying "Hello" to the chatbot and the chatbot responding with a "Hi" and asking the user how it can help, or it could be a transactional interaction in a banking chatbot, such as transferring money from one account to the other, or an informational interaction in a HR chatbot, such as checking for vacation balance, or asking an FAQ in a retail chatbot, such as how to handle returns. Natural language processing (NLP) and machine learning (ML) algorithms combined with other approaches can be used to classify end user intent. An intent at a high level is what the end user would like to accomplish (e.g., get account balance, make a purchase). An intent is essentially, a mapping of customer input to a unit of work that the backend should perform. Therefore, based on the phrases uttered by the user in the chatbot, these are mapped that to a specific and discrete use case or unit of work, for e.g. check balance, transfer money and track spending are all "use cases" that the chatbot should support and be able to work out which unit of work should be triggered from the free text entry that the end user types in a natural language.

The underlying rationale for having an AI chatbot respond like a human is that the human brain can formulate and understand the request and then give a good response to the human request much better than a machine. Thus, there should be significant improvement in the request/response of a chatbot, if human B is mimicked. So an initial part of the problem is how does the human brain formulate and understand the request? To mimic, a model is used. RST and DT allow a formal and repeatable way of doing this.

At a high level, there are typically two types of requests: (1) A request to perform some action; and (2) a request for information, e.g., a question. The first type has a response in which a unit of work is created. The second type has a response that is, e.g., a good answer, to the question. The answer could take the form of, for example, in some aspects, the AI constructing an answer from its extensive knowledge base(s) or from matching the best existing answer from searching the internet or intranet or other publically/private available data sources.

Communicative Discourse Trees and The Rhetoric Classifier Aspects of the present disclosure build communicative discourse trees and use communicative discourse trees to analyze whether the rhetorical structure of a request or question agrees with an answer. More specifically, aspects described herein create representations of a request-response pair, learns the representations, and relates the pairs into classes of valid or invalid pairs. In this manner, an autonomous agent can receive a question from a user, process the question, for example, by searching for multiple answers, determine the best answer from the answers, and provide the answer to the user.

More specifically, to represent linguistic features of text, aspects described herein use rhetoric relations and speech acts (or communicative actions). Rhetoric relations are relationships between the parts of the sentences, typically obtained from a discourse tree. Speech acts are obtained as verbs from a verb resource such as VerbNet. By using both rhetoric relations and communicative actions, aspects described herein can correctly recognize valid request-response pairs. To do so, aspects correlate the syntactic structure of a question with that of an answer. By using the structure, a better answer can be determined.

For example, when an autonomous agent receives an indication from a person that the person desires to sell an item with certain features, the autonomous agent should provide a search result that not only contains the features but also indicates an intent to buy. In this manner, the autonomous agent has determined the user's intent. Similarly, when an autonomous agent receives a request from a person to share knowledge about a particular item, the search result should contain an intent to receive a recommendation. When a person asks an autonomous agent for an opinion about a subject, the autonomous agent shares an opinion about the subject, rather than soliciting another opinion.

Analyzing Request and Response Pairs

FIG. 7 depicts an exemplary DT for an example request about property tax in accordance with an aspect. The node labels are the relations and the arrowed line points to the satellite. The nucleus is a solid line. FIG. 7 depicts the following text.

Request: "My husbands' grandmother gave him his grandfather's truck. She signed the title over but due to my husband having unpaid fines on his license, he was not able to get the truck put in his name. I wanted to put in my name and paid the property tax and got insurance for the truck. By the time it came to sending off the title and getting the tag, I didn't have the money to do so. Now, due to circumstances, I am not going to be able to afford the truck. I went to the insurance place and was refused a refund. I am just wondering that since I am not going to have a tag on this truck, is it possible to get the property tax refunded?"

Response: "The property tax is assessed on property that you own. Just because you chose to not register it does not mean that you don't own it, so the tax is not refundable. Even if you have not titled the vehicle yet, you still own it within the boundaries of the tax district, so the tax is payable. Note that all states give you a limited amount of time to transfer title and pay the use tax. If you apply late, there will be penalties on top of the normal taxes and fees. You don't need to register it at the same time, but you absolutely need to title it within the period of time stipulated in state law."

As can be seen in FIG. 7, analyzing the above text results in the following. "My husbands' grandmother gave him his grandfather's truck" is elaborated by "She signed the title over but due to my husband" elaborated by "having unpaid fines on his license, he was not able to get the truck put in his name." which is elaborated by "I wanted to put in my name," "and paid the property tax", and "and got insurance for the truck."

"My husbands' grandmother gave him his grandfather's truck. She signed the title over but due to my husband having unpaid fines on his license, he was not able to get the truck put in his name. I wanted to put in my name and paid the property tax and got insurance for the truck." is elaborated by;

"I didn't have the money" elaborated by "to do so" contrasted with "By the time" elaborated by "it came to sending off the title"

"and getting the tag"

"My husbands' grandmother gave him his grandfather's truck. She signed the title over but due to my husband having unpaid fines on his license, he was not able to get the truck put in his name. I wanted to put in my name and paid the property tax and got insurance for the truck. By the time it came to sending off the title and getting the tag, I didn't have the money to do so" is contrasted with

"Now, due to circumstances," elaborated with "I am not going to be able to afford the truck." which is elaborated with

"I went to the insurance place"

"and was refused a refund"

"My husbands' grandmother gave him his grandfather's truck. She signed the title over but due to my husband having unpaid fines on his license, he was not able to get the truck put in his name. I wanted to put in my name and paid the property tax and got insurance for the truck. By the time it came to sending off the title and getting the tag, I didn't have the money to do so. Now, due to circumstances, I am not going to be able to afford the truck. I went to the insurance place and was refused a refund." is elaborated with

"I am just wondering that since I am not going to have a tag on this truck, is it possible to get the property tax refunded?"

"I am just wondering" has attribution to

"that" is the same unit as "is it possible to get the property tax refunded?" which has condition "since I am not going to have a tag on this truck"

As can be seen, the main subject of the topic is "Property tax on a car". The question includes the contradiction: on one hand, all properties are taxable, and on the other hand, the ownership is somewhat incomplete. A good response has to address both topic of the question and clarify the inconsistency. To do that, the responder is making even stronger claim concerning the necessity to pay tax on whatever is owned irrespectively of the registration status. This example is a member of positive training set from our Yahoo! Answers evaluation domain. The main subject of the topic is "Property tax on a car". The question includes the contradiction: on one hand, all properties are taxable, and on the other hand, the ownership is somewhat incomplete. A good answer/response has to address both topic of the question and clarify the inconsistency. The reader can observe that since the question includes rhetoric relation of contrast, the answer has to match it with a similar relation to be convincing. Otherwise, this answer would look incomplete even to those who are not domain experts.

FIG. 8 depicts an exemplary response for the question represented in FIG. 7, according to certain aspects of the present invention. The central nucleus is "the property tax is assessed on property" elaborated by "that you own". "The property tax is assessed on property that you own" is also a nucleus elaborated by "Just because you chose to not register it does not mean that you don't own it, so the tax is not refundable. Even if you have not titled the vehicle yet, you still own it within the boundaries of the tax district, so the tax is payable. Note that all states give you a limited amount of time to transfer title and pay the use tax."

The nucleus "The property tax is assessed on property that you own. Just because you chose to not register it does not mean that you don't own it, so the tax is not refundable. Even if you have not titled the vehicle yet, you still own it within the boundaries of the tax district, so the tax is payable. Note that all states give you a limited amount of time to transfer title and pay the use tax." is elaborated by "there will be penalties on top of the normal taxes and fees" with condition "If you apply late," which in turn is elaborated by the contrast of "but you absolutely need to title it within the period of time stipulated in state law." and "You don't need to register it at the same time".

Comparing the DT of FIG. 7 and DT of FIG. 8, enables a determination of how well matched the response (FIG. 8) is to the request (FIG. 7). In some aspects of the present

invention, the above framework is used, at least in part, to determine the DTs for the request/response and the rhetoric agreement between the DTs.

In another example, the question “What does The Investigative Committee of the Russian Federation do” has at least two answers, for example, an official answer or an actual answer.

FIG. 9 illustrates a discourse tree for an official answer in accordance with an aspect. As depicted in FIG. 9, an official answer, or mission statement states that “The Investigative Committee of the Russian Federation is the main federal investigating authority which operates as Russia’s Anti-corruption agency and has statutory responsibility for inspecting the police forces, combating police corruption and police misconduct, is responsible for conducting investigations into local authorities and federal governmental bodies.”

FIG. 10 illustrates a discourse tree for a raw answer in accordance with an aspect. As depicted in FIG. 10, another, perhaps more honest, answer states that “Investigative Committee of the Russian Federation is supposed to fight corruption. However, top-rank officers of the Investigative Committee of the Russian Federation are charged with creation of a criminal community. Not only that, but their involvement in large bribes, money laundering, obstruction of justice, abuse of power, extortion, and racketeering has been reported. Due to the activities of these officers, dozens of high-profile cases including the ones against criminal lords had been ultimately ruined.”

The choice of answers depends on context. Rhetoric structure allows differentiating between “official”, “politically correct”, template-based answers and “actual”, “raw”, “reports from the field”, or “controversial” answers. (See FIG. 9 and FIG. 10). Sometimes, the question itself can give a hint about which category of answers is expected. If a question is formulated as a factoid or definitional one, without a second meaning, then the first category of answers is suitable. Otherwise, if a question has the meaning “tell me what it really is”, then the second category is appropriate. In general, after extracting a rhetoric structure from a question, selecting a suitable answer that would have a similar, matching, or complementary rhetoric structure is easier.

The official answer is based on elaboration and joints, which are neutral in terms of controversy a text might contain (See FIG. 9). At the same time, the raw answer includes the contrast relation. This relation is extracted between the phrase for what an agent is expected to do and what this agent was discovered to have done.

#### Classification of Request-Response Pairs

Rhetoric classification application 102 can determine whether a given answer or response, such as an answer obtained from answer database 105 or a public database, is responsive to a given question, or request. More specifically, rhetoric classification application 102 analyzes whether a request and response pair is correct or incorrect by determining one or both of (i) relevance or (ii) rhetoric agreement between the request and the response. Rhetoric agreement can be analyzed without taking into account relevance, which can be treated orthogonally.

Rhetoric classification application 102 can determine similarity between question-answer pairs using different methods. For example, rhetoric classification application 102 can determine level of similarity between an individual question and an individual answer. Alternatively, rhetoric classification application 102 can determine a measure of similarity between a first pair including a question and an answer, and a second pair including a question and answer.

For example, rhetoric classification application 102 uses rhetoric agreement classifier 120 trained to predict matching or non-matching answers. Rhetoric classification application 102 can process two pairs at a time, for example  $\langle q1, a1 \rangle$  and  $\langle q2, a2 \rangle$ . Rhetoric classification application 102 compares  $q1$  with  $q2$  and  $a1$  with  $a1$ , producing a combined similarity score. Such a comparison allows a determination of whether an unknown question/answer pair contains a correct answer or not by assessing a distance from another question/answer pair with a known label. In particular, an unlabeled pair  $\langle q2, a2 \rangle$  can be processed so that rather than “guessing” correctness based on words or structures shared by  $q2$  and  $a2$ , both  $q2$  and  $a2$  can be compared with their corresponding components  $q1$  and  $a2$  of the labeled pair  $\langle q2, a2 \rangle$  on the grounds of such words or structures. Because this approach targets a domain-independent classification of an answer, only the structural cohesiveness between a question and answer can be leveraged, not ‘meanings’ of answers.

In an aspect, rhetoric classification application 102 uses training data 125 to train rhetoric agreement classifier 120. In this manner, rhetoric agreement classifier 120 is trained to determine a similarity between pairs of questions and answers. This is a classification problem. Training data 125 can include a positive training set and a negative training set. Training data 125 includes matching request-response pairs in a positive dataset and arbitrary or lower relevance or appropriateness request-response pairs in a negative dataset. For the positive dataset, various domains with distinct acceptance criteria are selected that indicate whether an answer or response is suitable for the question.

Each training data set includes a set of training pairs. Each training set includes a question communicative discourse tree that represents a question and an answer communicative discourse tree that represents an answer and an expected level of complementarity between the question and answer. By using an iterative process, rhetoric classification application 102 provides a training pair to rhetoric agreement classifier 120 and receives, from the model, a level of complementarity. Rhetoric classification application 102 calculates a loss function by determining a difference between the determined level of complementarity and an expected level of complementarity for the particular training pair. Based on the loss function, rhetoric classification application 102 adjusts internal parameters of the classification model to minimize the loss function.

Acceptance criteria can vary by application. For example, acceptance criteria may be low for community question answering, automated question answering, automated and manual customer support systems, social network communications and writing by individuals such as consumers about their experience with products, such as reviews and complaints. RR acceptance criteria may be high in scientific texts, professional journalism, health and legal documents in the form of FAQ, professional social networks such as “stackoverflow.”

#### Communicative Discourse Trees (CDTs)

Rhetoric classification application 102 can create, analyze, and compare communicative discourse trees. Communicative discourse trees are designed to combine rhetoric information with speech act structures. CDTs include with arcs labeled with expressions for communicative actions. By combining communicative actions, CDTs enable the modeling of RST relations and communicative actions. A CDT is a reduction of a parse thicket. See Galitsky, B, Ilvovsky, D. and Kuznetsov S O. Rhetoric Map of an Answer to Compound Queries Knowledge Trail Inc. ACL 2015, 681-

686. (“Galitsky 2015”). A parse thicket is a combination of parse trees for sentences with discourse-level relationships between words and parts of the sentence in one graph. By incorporating labels that identify speech actions, learning of communicative discourse trees can occur over a richer features set than just rhetoric relations and syntax of elementary discourse units (EDUs).

In an example, a dispute between three parties concerning the causes of a downing of a commercial airliner, Malaysia Airlines Flight 17 is analyzed. An RST representation of the arguments being communicated is built. In the example, three conflicting agents, Dutch investigators, The Investigative Committee of the Russian Federation, and the self-proclaimed Donetsk People’s Republic exchange their opinions on the matter. The example illustrates a controversial conflict where each party does all it can to blame its opponent. To sound more convincing, each party does not just produce its claim but formulates a response in a way to rebuff the claims of an opponent. To achieve this goal, each party attempts to match the style and discourse of the opponents’ claims.

FIG. 11 illustrates a communicative discourse tree for a claim of a first agent in accordance with an aspect. FIG. 11 depicts communicative discourse tree 100, which represents the following text: “Dutch accident investigators say that evidence points to pro-Russian rebels as being responsible for shooting down plane. The report indicates where the missile was fired from and identifies who was in control of the territory and pins the downing of MH17 on the pro-Russian rebels.”

As can be seen from FIG. 11, non-terminal nodes of CDTs are rhetoric relations, and terminal nodes are elementary discourse units (phrases, sentence fragments) which are the subjects of these relations. Certain arcs of CDTs are labeled with the expressions for communicative actions, including the actor agent and the subject of these actions (what is being communicated). For example, the nucleus node for elaboration relation (on the left) are labeled with say (Dutch, evidence), and the satellite with responsible (rebels, shooting down). These labels are not intended to express that the subjects of EDUs are evidence and shooting down but instead for matching this CDT with others for the purpose of finding similarity between them. In this case just linking these communicative actions by a rhetoric relation and not providing information of communicative discourse would be too limited way to represent a structure of what and how is being communicated. A requirement for an RR pair to have the same or coordinated rhetoric relation is too weak, so an agreement of CDT labels for arcs on top of matching nodes is required.

The straight edges of this graph are syntactic relations, and curvy arcs are discourse relations, such as anaphora, same entity, sub-entity, rhetoric relation and communicative actions. This graph includes much richer information than just a combination of parse trees for individual sentences. In addition to CDTs, parse thickets can be generalized at the level of words, relations, phrases and sentences. The speech actions are logic predicates expressing the agents involved in the respective speech acts and their subjects. The arguments of logical predicates are formed in accordance to respective semantic roles, as proposed by a framework such as VerbNet. See Karin Kipper, Anna Korhonen, Neville Ryant, Martha Palmer, A Large-scale Classification of English Verbs, Language Resources and Evaluation Journal, 42(1), pp. 21-40, Springer Netherland, 2008. and/or Karin Kipper Schuler, Anna Korhonen, Susan W. Brown, VerbNet

overview, extensions, mappings and apps, Tutorial, NAACL-HLT 2009, Boulder, Colo.

FIG. 12 illustrates a communicative discourse tree for a claim of a second agent in accordance with an aspect. FIG. 12 depicts communicative discourse tree 1200, which represents the following text: “The Investigative Committee of the Russian Federation believes that the plane was hit by a missile, which was not produced in Russia. The committee cites an investigation that established the type of the missile.”

FIG. 13 illustrates a communicative discourse tree for a claim of a third agent in accordance with an aspect. FIG. 13 depicts communicative discourse tree 1300, which represents the following text: “Rebels, the self-proclaimed Donetsk People’s Republic, deny that they controlled the territory from which the missile was allegedly fired. It became possible only after three months after the tragedy to say if rebels controlled one or another town.”

As can be seen from communicative discourse trees 1100-1300, a response is not arbitrary. A response talks about the same entities as the original text. For example, communicative discourse trees 1200 and 1300 are related to communicative discourse tree 1100. A response backs up a disagreement with estimates and sentiments about these entities, and about actions of these entities.

More specifically, replies of involved agent need to reflect the communicative discourse of the first, seed message. As a simple observation, because the first agent uses Attribution to communicate his claims, the other agents have to follow the suite and either provide their own attributions or attack the validity of attribution of the proponent, or both. To capture a broad variety of features for how communicative structure of the seed message needs to be retained in consecutive messages, pairs of respective CDTs can be learned.

To verify the agreement of a request-response, discourse relations or speech acts (communicative actions) alone are often insufficient. As can be seen from the example depicted in FIGS. 11-13, the discourse structure of interactions between agents and the kind of interactions are useful. However, the domain of interaction (e.g., military conflicts or politics) or the subjects of these interactions, i.e., the entities, do not need to be analyzed.

Representing Rhetoric Relations and Communicative Actions

In order to compute similarity between abstract structures, two approaches are frequently used: (1) representing these structures in a numerical space, and express similarity as a number, which is a statistical learning approach, or (2) using a structural representation, without numerical space, such as trees and graphs, and expressing similarity as a maximal common sub-structure. Expressing similarity as a maximal common sub-structure is referred to as generalization.

Learning communicative actions helps express and understand arguments. Computational verb lexicons help support acquisition of entities for actions and provide a rule-based form to express their meanings. Verbs express the semantics of an event being described as well as the relational information among participants in that event, and project the syntactic structures that encode that information. Verbs, and in particular communicative action verbs, can be highly variable and can display a rich range of semantic behaviors. In response, verb classification helps a learning systems to deal with this complexity by organizing verbs into groups that share core semantic properties.

VerbNet is one such lexicon, which identifies semantic roles and syntactic patterns characteristic of the verbs in

each class and makes explicit the connections between the syntactic patterns and the underlying semantic relations that can be inferred for all members of the class. See Karin Kipper, Anna Korhonen, Neville Ryant and Martha Palmer, *Language Resources and Evaluation*, Vol. 42, No. 1 (March 2008), at 21. Each syntactic frame, or verb signature, for a class has a corresponding semantic representation that details the semantic relations between event participants across the course of the event.

For example, the verb amuse is part of a cluster of similar verbs that have a similar structure of arguments (semantic roles) such as amaze, anger, arouse, disturb, and irritate. The roles of the arguments of these communicative actions are as follows: Experiencer (usually, an animate entity), Stimulus, and Result. Each verb can have classes of meanings differentiated by syntactic features for how this verb occurs in a sentence, or frames. For example, the frames for amuse are as follows, using the following key noun phrase (NP), noun (N), communicative action (V), verb phrase (VP), adverb (ADV):

NP V NP. Example: “The teacher amused the children.”  
Syntax: Stimulus V Experiencer. Clause: amuse(Stimulus, E, Emotion, Experiencer), cause(Stimulus, E), emotional\_state(result(E), Emotion, Experiencer).

NP V ADV-Middle. Example: “Small children amuse quickly.” Syntax: Experiencer V ADV. Clause: amuse(Experiencer, Prop): property(Experiencer, Prop), adv(Prop).

NP V NP-PRO-ARB. Example “The teacher amused.”  
Syntax Stimulus V. amuse(Stimulus, E, Emotion, Experiencer):. cause(Stimulus, E), emotional\_state(result(E), Emotion, Experiencer).

NP.cause V NP. Example “The teacher’s dolls amused the children.” syntax Stimulus <+genitive> (’s) V Experiencer. amuse(Stimulus, E, Emotion, Experiencer): cause(Stimulus, E), emotional\_state(during(E), Emotion, Experiencer).

NP V NP ADJ. Example “This performance bored me totally.” syntax Stimulus V Experiencer Result. amuse(Stimulus, E, Emotion, Experiencer). cause(Stimulus, E), emotional\_state(result(E), Emotion, Experiencer), Pred(result(E), Experiencer).

Communicative actions can be characterized into clusters, for example:

Verbs with Predicative Complements (Appoint, characterize, dub, declare, conjecture, masquerade, orphan, captain, consider, classify), Verbs of Perception (See, sight, peer).  
Verbs of Psychological State (Amuse, admire, marvel, appeal), Verbs of Desire (Want, long).

Judgment Verbs (Judgment), Verbs of Assessment (Assess, estimate), Verbs of Searching (Hunt, search, stalk, investigate, rummage, ferret), Verbs of Social Interaction (Correspond, marry, meet, battle), Verbs of Communication (Transfer(message), inquire, interrogate, tell, manner(speaking), talk, chat, say, complain, advise, confess, lecture, overstate, promise). Avoid Verbs (Avoid), Measure Verbs, (Register, cost, fit, price, bill), Aspectual Verbs (Begin, complete, continue, stop, establish, sustain).

Aspects described herein provide advantages over statistical learning models. In contrast to statistical solutions, aspects use a classification system can provide a verb or a verb-like structure which is determined to cause the target feature (such as rhetoric agreement). For example, statistical machine learning models express similarity as a number, which can make interpretation difficult.

Representing Request-Response Pairs

Representing request-response pairs facilitates classification based operations based on a pair. In an example, request-response pairs can be represented as parse thickets.

A parse thicket is a representation of parse trees for two or more sentences with discourse-level relationships between words and parts of the sentence in one graph. See Galitsky 2015. Topical similarity between question and answer can be expressed as common sub-graphs of parse thickets. The higher the number of common graph nodes, the higher the similarity.

FIG. 14 illustrates parse thickets in accordance with an aspect. FIG. 14 depicts parse thicket 1400 including a parse tree for a request 1401, and a parse tree for a corresponding response 1402.

Parse tree 1401 represents the question “I just had a baby and it looks more like the husband I had my baby with. However it does not look like me at all and I am scared that he was cheating on me with another lady and I had her kid. This child is the best thing that has ever happened to me and I cannot imagine giving my baby to the real mom.”

Response 1402 represents the response “Marital therapists advise on dealing with a child being born from an affair as follows. One option is for the husband to avoid contact but just have the basic legal and financial commitments. Another option is to leave the wife fully involved and have the baby fully integrated into the family just like a child from a previous marriage.”

FIG. 14 represents a greedy approach to representing linguistic information about a paragraph of text. The straight edges of this graph are syntactic relations, and curvy arcs are discourse relations, such as anaphora, same entity, sub-entity, rhetoric relation and communicative actions. The solid arcs are for same entity/sub-entity/anaphora relations, and the dotted arcs are for rhetoric relations and communicative actions. Oval labels in straight edges denote the syntactic relations. Lemmas are written in the boxes for the nodes, and lemma forms are written on the right side of the nodes.

Parse thicket 1400 includes much richer information than just a combination of parse trees for individual sentences. Navigation through this graph along the edges for syntactic relations as well as arcs for discourse relations allows to transform a given parse thicket into semantically equivalent forms for matching with other parse thickets, performing a text similarity assessment task. To form a complete formal representation of a paragraph, as many links as possible are expressed. Each of the discourse arcs produces a pair of thicket phrases that can be a potential match.

Topical similarity between the seed (request) and response is expressed as common sub-graphs of parse thickets. They are visualized as connected clouds. The higher the number of common graph nodes, the higher the similarity. For rhetoric agreement, common sub-graph does not have to be large as it is in the given text. However, rhetoric relations and communicative actions of the seed and response are correlated and a correspondence is required.

Generalization for Communicative Actions

A similarity between two communicative actions  $A_1$  and  $A_2$  is defined as an abstract verb which possesses the features which are common between  $A_1$  and  $A_2$ . Defining a similarity of two verbs as an abstract verb-like structure supports inductive learning tasks, such as a rhetoric agreement assessment. In an example, a similarity between the following two common verbs, agree and disagree, can be generalized as follows: agree<sup>^</sup>A disagree=verb(Interlocutor, Proposed\_action, Speaker), where Interlocution is the person who proposed the Proposed\_action to the Speaker and to whom the Speaker communicates their response. Proposed\_action is an action that the Speaker would perform if they were to accept or refuse the request or offer, and The Speaker

is the person to whom a particular action has been proposed and who responds to the request or offer made.

In a further example, a similarity between verbs agree and explain is represented as follows: agree' explain=verb(Interlocutor, \*, Speaker). The subjects of communicative actions are generalized in the context of communicative actions and are not generalized with other "physical" actions. Hence, aspects generalize individual occurrences of communicative actions together with corresponding subjects.

Additionally, sequences of communicative actions representing dialogs can be compared against other such sequences of similar dialogs. In this manner, the meaning of an individual communicative action as well as the dynamic discourse structure of a dialogue is (in contrast to its static structure reflected via rhetoric relations) is represented. A generalization is a compound structural representation that happens at each level. Lemma of a communicative action is generalized with lemma, and its semantic role are generalized with respective semantic role.

Communicative actions are used by text authors to indicate a structure of a dialogue or a conflict. See Searle, J. R. 1969, *Speech acts: an essay in the philosophy of language*. London: Cambridge University Press. Subjects are generalized in the context of these actions and are not generalized with other "physical" actions. Hence, the individual occurrences of communicative actions together are generalized with their subjects, as well as their pairs, as discourse "steps."

Generalization of communicative actions can also be thought of from the standpoint of matching the verb frames, such as VerbNet. The communicative links reflect the discourse structure associated with participation (or mentioning) of more than a single agent in the text. The links form a sequence connecting the words for communicative actions (either verbs or multi-words implicitly indicating a communicative intent of a person).

Communicative actions include an actor, one or more agents being acted upon, and the phrase describing the features of this action. A communicative action can be described as a function of the form: verb (agent, subject, cause), where verb characterizes some type of interaction between involved agents (e.g., explain, confirm, remind, disagree, deny, etc.), subject refers to the information transmitted or object described, and cause refers to the motivation or explanation for the subject.

A scenario (labeled directed graph) is a sub-graph of a parse thicket  $G=(V, A)$ , where  $V=\{\text{action}_1, \text{action}_2, \dots, \text{action}_n\}$  is a finite set of vertices corresponding to communicative actions, and  $A$  is a finite set of labeled arcs (ordered pairs of vertices), classified as follows:

Each arc  $\text{action}_i, \text{action}_j \in A_{\text{sequence}}$  corresponds to a temporal precedence of two actions  $v_i, ag_i, s_i, c_i$  and  $v_j, ag_j, s_j, c_j$  that refer to the same subject, e.g.,  $s_j=s_i$  or different subjects. Each arc  $\text{action}_i, \text{action}_j \in A_{\text{cause}}$  corresponds to an attack relationship between  $\text{action}_i$  and  $\text{action}_j$ , indicating that the cause of action, in conflict with the subject or cause of action.

Subgraphs of parse thickets associated with scenarios of interaction between agents have some distinguishing features. For example, (1) all vertices are ordered in time, so that there is one incoming arc and one outgoing arc for all vertices (except the initial and terminal vertices), (2) for  $A_{\text{sequence}}$  arcs, at most one incoming and only one outgoing arc are admissible, and (3) for  $A_{\text{cause}}$  arcs, there can be many outgoing arcs from a given vertex, as well as many incoming arcs. The vertices involved may be associated with different agents or with the same agent (i.e., when he contradicts

himself). To compute similarities between parse thickets and their communicative action, induced subgraphs, the subgraphs of the same configuration with similar labels of arcs and strict correspondence of vertices are analyzed.

The following similarities exist by analyzing the arcs of the communicative actions of a parse thicket: (1) one communicative action from with its subject from T1 against another communicative action with its subject from T2 (communicative action arc is not used), and (2) a pair of communicative actions with their subjects from T1 compared to another pair of communicative actions from T2 (communicative action arcs are used).

Generalizing two different communicative actions is based on their attributes. See (Galitsky et al 2013). As can be seen in the example discussed with respect to FIG. 14, one communicative action from T1, cheating(husband, wife, another lady) can be compared with a second from T2 avoid(husband, contact(husband, another lady)). A generalization results in  $\text{communicative\_action}(\text{husband}, *)$  which introduces a constraint on A in the form that if a given agent (=husband) is mentioned as a subject of CA in Q, he/she should also be a subject of (possibly, another) CA in A. Two communicative actions can always be generalized, which is not the case for their subjects: if their generalization result is empty, the generalization result of communicative actions with these subjects is also empty.

Generalization of RST Relations

Some relations between discourse trees can be generalized, such as arcs that represent the same type of relation (presentation relation, such as antithesis, subject matter relation, such as condition, and multinuclear relation, such as list) can be generalized. A nucleus or a situation presented by a nucleus is indicated by "N." Satellite or situations presented by a satellite, are indicated by "S." "W" indicates a writer. "R" indicates a reader (hearer). Situations are propositions, completed actions or actions in progress, and communicative actions and states (including beliefs, desires, approve, explain, reconcile and others). Generalization of two RST relations with the above parameters is expressed as:

$$\text{rst1}(N1, S1, W1, R1) \wedge \text{rst2}(N2, S2, W2, R2) = (\text{rst1} \wedge \text{rst2}) \\ (N1 \wedge N2, S1 \wedge S2, W1 \wedge W2, R1 \wedge R2).$$

The texts in N1, S1, W1, R1 are subject to generalization as phrases. For example,  $\text{rst1} \wedge \text{rst2}$  can be generalized as follows: (1) if  $\text{relation\_type}(\text{rst1}) \neq \text{relation\_type}(\text{rst2})$  then a generalization is empty. (2) Otherwise, the signatures of rhetoric relations are generalized as sentences:  $\text{sentence}(N1, S1, W1, R1) \wedge \text{sentence}(N2, S2, W2, R2)$ . See Iruskieta, Mikel, Iria da Cunha and Maite Taboada. A qualitative comparison method for rhetorical structures: identifying different discourse structures in multilingual corpora. *Lang Resources & Evaluation*. June 2015, Volume 49, Issue 2.

For example, the meaning of  $\text{rst} \text{--} \text{background} \wedge \text{rst} \text{--} \text{enablement} = (S \text{ increases the ability of } R \text{ to comprehend an element in } N) \wedge (R \text{ comprehending } S \text{ increases the ability of } R \text{ to perform the action in } N) = \text{increase-VB the-DT ability-NN of-IN R-NN to-IN}$ .

Because the relations  $\text{rst} \text{--} \text{background} \wedge \text{rst} \text{--} \text{enablement}$  differ, the RST relation part is empty. The expressions that are the verbal definitions of respective RST relations are then generalized. For example, for each word or a placeholder for a word such as an agent, this word (with its POS) is retained if the word the same in each input phrase or remove the word if the word is different between these phrases. The resultant

expression can be interpreted as a common meaning between the definitions of two different RST relations, obtained formally.

Two arcs between the question and the answer depicted in FIG. 14 show the generalization instance based on the RST relation “RST-contrast”. For example, “I just had a baby” is a RST-contrast with “it does not look like me,” and related to “husband to avoid contact” which is a RST-contrast with “have the basic legal and financial commitments.” As can be seen, the answer need not have to be similar to the verb phrase of the question but the rhetoric structure of the question and answer are similar. Not all phrases in the answer must match phrases in question. For example, the phrases that do not match have certain rhetoric relations with the phrases in the answer which are relevant to phrases in question.

Building a Communicative Discourse Tree

FIG. 15 illustrates an exemplary process for building a communicative discourse tree in accordance with an aspect. Rhetoric classification application 102 can implement process 1500. As discussed, communicative discourse trees enable improved search engine results.

At block 1501, process 1500 involves accessing a sentence including fragments. At least one fragment includes a verb and words and each word includes a role of the words within the fragment, and each fragment is an elementary discourse unit. For example, rhetoric classification application 102 accesses a sentence such as “Rebels, the self-proclaimed Donetsk People’s Republic, deny that they controlled the territory from which the missile was allegedly fired” as described with respect to FIG. 13.

Continuing the example, rhetoric classification application 102 determines that the sentence includes several fragments. For example, a first fragment is “rebels . . . deny.” A second fragment is “that they controlled the territory.” A third fragment is “from which the missile was allegedly fired.” Each fragment includes a verb, for example, “deny” for the first fragment and “controlled” for the second fragment. Although, a fragment need not include a verb.

At block 1502, process 1500 involves generating a discourse tree that represents rhetorical relationships between the sentence fragments. The discourse tree including nodes, each nonterminal node representing a rhetorical relationship between two of the sentence fragments and each terminal node of the nodes of the discourse tree is associated with one of the sentence fragments.

Continuing the example, rhetoric classification application 102 generates a discourse tree as shown in FIG. 13. For example, the third fragment, “from which the missile was allegedly fired” elaborates on “that they controlled the territory.” The second and third fragments together relate to attribution of what happened, i.e., the attack cannot have been the rebels because they do not control the territory.

At block 1503, process 1500 involves accessing multiple verb signatures. For example, rhetoric classification application 102 accesses a list of verbs, e.g., from VerbNet. Each verb matches or is related to the verb of the fragment. For example, the for the first fragment, the verb is “deny.” Accordingly, rhetoric classification application 102 accesses a list of verb signatures that relate to the verb deny.

As discussed, each verb signature includes the verb of the fragment and one or more of thematic roles. For example, a signature includes one or more of noun phrase (NP), noun (N), communicative action (V), verb phrase (VP), or adverb (ADV). The thematic roles describing the relationship between the verb and related words. For example “the teacher amused the children” has a different signature from

“small children amuse quickly.” For the first fragment, the verb “deny,” rhetoric classification application 102 accesses a list of frames, or verb signatures for verbs that match “deny.” The list is “NP V NP to be NP,” “NP V that S” and “NP V NP.”

Each verb signature includes thematic roles. A thematic role refers to the role of the verb in the sentence fragment. Rhetoric classification application 102 determines the thematic roles in each verb signature. Example thematic roles include actor, agent, asset, attribute, beneficiary, cause, location destination source, destination, source, location, experiencer, extent, instrument, material and product, material, product, patient, predicate, recipient, stimulus, theme, time, or topic.

At block 1504, process 1500 involves determining, for each verb signature of the verb signatures, a number of thematic roles of the respective signature that match a role of a word in the fragment. For the first fragment, rhetorical classification application 102 determines that the verb “deny” has only three roles, “agent”, “verb” and “theme.”

At block 1505, process 1500 involves selecting a particular verb signature from the verb signatures based on the particular verb signature having a highest number of matches. For example, referring again to FIG. 13, deny in the first fragment “the rebels deny . . . that they control the territory” is matched to verb signature deny “NP V NP”, and “control” is matched to control (rebel, territory). Verb signatures are nested, resulting in a nested signature of “deny (rebel, control(rebel, territory)).”

Concluding a Question Answering Session

As discussed, aspects of the present disclosure construct a comprehensive answer for presentation an interactive session between user and autonomous agent. In some cases, a short and concise answer such as account balance or person name suffices. This type of answer can be referred to as a factoid answer. In other cases, a longer answer that addresses subsequent user utterances is appropriate. This answer, a conclusive answer, is a comprehensive answer that gives a user a chance to get a deep understanding of an entity or topic. An example of a process for determining whether a factoid or conclusive answer is appropriate is discussed further with respect to FIG. 16.

FIG. 16 illustrates an exemplary process 1600 for determining an appropriate type of answer to be provided by an autonomous agent, in accordance with an aspect. Process 1600 can be performed by conclusive answer application 102. Process 1600 can be executed after one or more utterances have been received by conclusive answer application 102.

At block 1601, process 1600 involves accessing one or more utterances. After an initial user utterance, for example, a question, a number of clarification steps usually follow. For example, user device 170 transmits one or more utterances to computing device 101. Computing device 101 receives the utterance and proceeds to block 1602.

At block 1602, process 1600 involves determining whether a number of utterances is greater than a threshold. Conclusive answer application 102 can decide which kind of answer is most suitable for a given session. Using a threshold number of utterances is one way to determine whether a conclusive answer is appropriate. In an example, the threshold is four utterances. In this example, if four or more utterances are present, then a conclusive answer is generated. Based on a threshold number of occurrences being met, then conclusive answer application 102 can determine that a conclusive answer should be provided.



Other methods can be used. For example, conclusive answer application 102 can determine that a particular entity is present in multiple user utterances, indicating that the user is focused on a particular issue. In an aspect, conclusive answer application 102 can detect that a user accepted information related to an entity. For example, if conclusive answer application 102 responds to a user question about an entity and the user accepts the response, then the conclusive answer application 102 can determine that a conclusive answer is appropriate. In another aspect, conclusive answer application 102 can detect a mental state from an utterance and determine that a conclusive answer is appropriate based on the mental state being detected. Mental states can be detected using trained machine learning models, using the techniques described herein.

If conclusive answer application 102 determines that a factoid answer is appropriate, then process 1600 proceeds to block 1603. If conclusive answer application 102 determines that a conclusive answer is appropriate, then process 1600 proceeds to block 1604.

At block 1603, process 1600 involves generating a factoid answer that corresponds to an initial user utterance. A factoid answer can be constructed based on a structure of how entity and its attributes are introduced. Conclusive answer application 102 can determine a value of the parameter or attribute that corresponds to the entity mentioned by the user. For a simple user utterance that requests information about one entity (e.g., the cost of a non-sufficient funds fee), a factoid answer is appropriate. Continuing the example, conclusive answer application 102 determines that the cost (the attribute) of a non-sufficient funds fee is \$29 (the value) in this particular case.

FIG. 17 illustrates an example of a conversation flow with a factoid answer, in accordance with an aspect. FIG. 17 depicts interaction 1700, which includes question 1701 and agent answer 1702. As can be seen, the user's question is a relatively factual one and can be addressed by a simple factual answer.

Returning to FIG. 16, at block 1604, process 1600 involves generating a conclusive answer that addresses multiple user utterances. The goal of the conclusive answer is that sufficient information is provided that a user is satisfied with a session with an autonomous agent. If further questions based on this answer arise, the user can start a new session keeping in mind a specific focus. The answer should be as easy to perceive and as intuitive as possible. Therefore the addition of images, videos and audio files can be beneficial. The answer compilation method should be domain independent and well-presented. In some cases, a table of contents and references are generated. Multiple sources can be used to assure an unbiased, objective description. In the case that the conclusive answer is opinionated, however, multiple opinions from a broad spectrum of perspectives must be compiled in a coherent manner.

A structure of a conclusive answer can be based on questions, disagreements, and/or misunderstandings about an entity occurring in an utterance. For example, if a user is upset about a non-sufficient funds charge, then a conclusive answer may be appropriate because the user desires a more comprehensive answer that explains why the fee was charged. FIG. 18 illustrates one such case.

FIG. 18 illustrates an example of a conversation flow with a conclusive answer, in accordance with an aspect. FIG. 18 depicts interaction 1800, which includes user questions 1801, 1803, and 1805 and agent responses 1802, 1804, and 1806. As depicted by FIG. 18, the user has some issues with

non-sufficient funds fees that go well beyond a simple factual question (as illustrated in the example provided by FIG. 17).

As can be seen, because the user is frustrated about the non-sufficient funds fee and is trying to understand why it happened and how to avoid it, a generic answer about an entity would probably frustrate the user because the user appears to know general information about non-sufficient funds fees. Therefore the conclusive answer should focus on a specific user issue/misunderstanding exposed in the previous utterances of a dialogue. Here, conclusive answer application provides a conclusive answer, agent response 1806. As can be seen, agent response 1806 addresses the user's multiple, but related issues, such as the issue raised by the user that a deposit was made (question 1803) and how to stay positive (question 1805).

At block 1605, process 1600 involves providing the conclusive answer to a user device. For example, conclusive answer application 102 provides the conclusive answer to user device 170. User device 170 can present the answer to a user via display 171.

Preparing Textual Content for Conclusive Answer Generation

In some cases, textual content is prepared for use by conclusive answer application 102. Preparation can help speed up access to text and/or remove extraneous or irrelevant content from the text. Conclusive answer application 102 can populate and use text corpus 105 with material gathered from other sources. Examples of sources include customer support logs, various forms of corresponding with customers or internal issue logs. Determining whether a text is suitable can be done by means of discourse-level analysis or in a string-based manner. Once chunks of text are extracted from various sources, the text can be indexed for faster searching.

FIG. 19 illustrates an exemplary process for converting textual content for use in forming a conclusive answer, in accordance with an aspect. Process 1900 can be performed by conclusive answer application 102.

At block 1901, process 1900 involves accessing a document with sections. Various document sources can be used, including the written documents and web pages explaining entities, their attributes and specifying business rules. Conclusive answer application 102 can convert the text into a unified form. In an example, the text can adhere (1) contain paragraph-size text (e.g., two to six sentences, 60-150 words) and (2) be self-contained. For example, conclusive answer application 102 can remove any references to previous or subsequent paragraphs.

At block 1902, process 1900 involves determining a style of each section. Style can be determined by using communicative discourse trees. For example, a communicative discourse tree can be created for a candidate section of text (e.g., by using process 1500 described in FIG. 15). The first communicative discourse tree can be provided to a machine learning model that is trained to detect whether the communicative discourse tree is similar in style to a reference (or desired style). Based on a threshold similarity level being reached, the candidate text is integrated into the text corpus 105.

At block 1903, process 1900 involves identifying and removing unsuitable portions of sections. Text can be extracted from a proper area of a webpage or a proper section of a document, and cleaned. For example, for web-based content, conclusive answer application 102 determines whether an article includes desired content (e.g. by matching keywords from the target content and the

document). If the document does include the desired content, then conclusive answer application 102 finds a contiguous block of HTML in the webpage starting with the first word in the article and ending with the last.

Additionally, conclusive answer application can remove everything other than the desired text. For example, HTML tags, advertisements, and other irrelevant content can be removed. When the first word or last word of desired content is nested within one or more pairs of HTML tags, conclusive answer application can append the relevant opening and ending to the beginning and ending of the extracted block. If the content is not nested, then one or more pairs of tags can be left open, disrupting the article text's formatting. Such cases can be ignored.

At block 1904, process 1900 involves dividing text into paragraphs. Conclusive answer application 102 can divide sentences when a newline character is detected and/or maintain a minimum and/or maximum paragraph length (number of sentences).

At block 1905, process 1900 involves indexing the text. Standard indexing techniques can be used.

Building a Structure of a Factoid Answer

The logical structure of an answer reflects the structure of preceding dialogue. For example, if the user has asked questions and/or received answers from the autonomous agent after the initial utterance, then these questions or answers is reflected in the structure of the conclusive answer.

Determining the structure of an answer involves determining one or more entities present in the user utterances. An entity is a noun that refers to a person, place, or thing. An entity can have an attribute. For example, if an entity is "non-sufficient funds fee," then the attribute might be "\$29" indicating that the fee is \$29. Once a syntactic parse tree is constructed for a user utterance, entities can be identified. For example, noun phrases are first identified. If a noun phrase includes no more than three words, then it can be considered to form a named entity.

For a factoid answer, a default structure can be provided. This structure can be mined from the general web sources such as Wikipedia and domain-specific sources such as Investopedia.com. For example, the TOC for the topic Adjusted Gross Margin would use the section structure from the respective Investopedia.com page such as the main definitions, treatment in depth, associated topics and others. In this case it is possible to build TOC in a hierarchical manner. Table 1 below illustrates an example of a dialog flow in which a user asks about a particular entity.

TABLE 1

An example of a format for a factoid answer.	
Section	Contents
Section 1	What is Entity E (the topic of the initial user utterance)
Section 2	E has its attribute A: the first user clarification request
Section 3	E and its attribute A1 and is it similar to A.
Section 4	E and how it is similar to another entity E1.

When determining the content of the answer, conclusive answer application 102 uses the determined structure as a guide. For each section, conclusive answer application 102 obtains content that corresponds to the entities and attributes determined in the user utterances.

Building a Structure of Conclusive Answer

A conclusive answer includes a section structure that reflects the logical flow of the user utterances as received so far. Conclusive answer application 102 can generate table of contents (TOC) that matches the structure.

For example, if a user has a specific concern about an entity, such as 'Why banks can increase APR without advance notice', then conclusive answer application 102 builds a structure to address this concern. Relevant documents can be identified by their respective section titles, which are selected to correspond to the contents of the user utterances. As compared to a factoid answer, the selected documents are identified to be relevant not just to the main entity but to the why certain entities are related and/or have certain attributes.

TABLE 2

An example of a format for a conclusive answer.	
Section	Contents
Section 1	What is Entity E (the topic of the initial user utterance)
Section 2	Why E has its attribute A: the first user clarification request
Section 3	E and its attribute A1 and is it similar to A: the second user clarification request
Section 4	E and how it is similar to another entity E1: the user expressed her concern about E2

For example, to form the document structure for the example considered in FIG. 18, the following phrases from the user questions can be as queries to establish the section structure of the conclusive answer: (1) Non-sufficient fund fee (NSF); (2) Why was I charged; (3) Make a deposit; and (4) make a payment for a lower amount. These phrases (extended with synonyms) should match some section structures of certain documents about NSF and banking customer support logs: they will form a skeleton of the resultant answer.

Conclusive answer application 102 can determine attributes that correspond to a particular entity is to form a structure of a document is an auto-complete feature for web search. Auto-complete results from a search engine can be the queries to the text corpus 105 and/or serve as section titles for the table-of-contents. For example, if an entity in the preceding dialogue is "Edison invented" then the final concluding document can include 'Edison invented the light bulb', and "Edison invented the phonograph."

Content Compilation for Answers

Once the structure of the conclusive answer has been determined, for example, by using process 1900, conclusive answer application 102 can determine relevant content. Content can be determined on a section basis. For example, referring back to Table 1 and Table 2 above, content can be determined for the first section, then the second section, and so on. Accordingly, conclusive answer application 102 can execute multiple times, once for each section.

FIG. 20 illustrates an exemplary process 2000 for building a conclusive answer, in accordance with an aspect. Process 2000 can be performed once, for a single seed sentence, or multiple times, once for each seed sentence, according the appropriate structure of the conclusive answer, for example, as determined by process 1600. A single phrase or multiple phrases of the seed can form one or more paragraphs about the respective topics.

At block 2001, process 2000 involves forming a query from an utterance. The utterance includes text fragments.

Each fragment can correspond to an elementary discourse unit. To find relevant sentences on the web for a seed sentence, conclusive answer application 102 develops a query.

For example, conclusive answer application 102 can identify any entities within the fragments and then uses the entities as a basis to query content sources. Identification of one or more entities can be determined generating syntactic parse trees from the utterance. In an example, significant noun phrases have three or more keywords, e.g., two or more modifiers for a noun, or an entity, such as a proper noun. A main entity can be identified by traversing a syntactic parse tree for the utterance from the top (root) down and identifying the first noun phrase in the syntactic parse tree.

In an example, a user utterance is a seed. The utterance is “Give me a break, there is no reason why you can’t retire in ten years if you had been a rational investor and not a crazy trader.” Conclusive answer application 102 identifies the entity in this utterance by creating a syntactic parse tree and identifying a noun phrase that is closest to the root (or the top) of the parse tree. Here, conclusive answer application 102 identifies the main entity as “rational investor.” The other candidates for the main entity are rejected since they are too broad (such as retire, a single-word concept), or occur with a negation not a crazy trader.

Conclusive answer application 102 identifies a main entity as retirement in the form of the verb retire. This main entity is constrained by the noun phrase that follows rational investor. To form the second query, rational investor and the next noun phrase not a crazy trader are combined. In some cases, a query with four to five keywords is used. Continuing the example, the following queries are formed for search engine API:

(Q1) +retire +rational +investor

(Q2) +rational +investor not +crazy +trader

At block 2002, process 2000 involves providing the query to one or more content sources. The content sources can include text corpus 105 or other sources. Examples of content sources include such sources as Wikipedia, Bing, Yahoo API or Google, as well as respective news content.

Continuing the example, conclusive answer application 102 determines a page for the search “rational investor” by searching Wikipedia. The page redirects to “Homo economicus,” which contains the following sections (1) the history of the term; (2) Model; (3) Criticisms; (4) Responses; (5) Perspectives; and (6) Homo sociologicus. In some cases, conclusive answer application 102 uses each section title. In some cases, if the queries do not result in enough relevant sentences, the whole sentence can be used as the query.

At block 2003, process 2000 involves obtaining results. Conclusive answer application 102 obtains the results from the search queries. In some cases, the results are divided into sentences and markers are inserted for any missing information, which can be substituted by text from original web pages or documents.

In some cases, only partial results are returned from the search. If only a fragment of sentence is present, then conclusive answer application 102 visits the original page, locates the sentence, and downloads the sentence.

Continuing the example, the following snippet is selected as a candidate to be included in a conclusive answer, since it contains all keywords from Q1. “How to Make Rational Investing Decisions|Sound Mind Investing . . . Nov. 1, 2014—How to Make Rational Investing Decisions . . . pleasant and you’ll probably have more money to spend in retirement and leave to your heirs.” Conclusive answer

application 102 downloads this webpage, extracts text from it and find a paragraph which corresponds to the above snippet. This is continued for all search results which contains all keywords from the query. Conclusive answer application 102 considers two text fragments from the search results:

(A1a) If you take the time to understand the psychology of rational investing, you’ll make your life more pleasant and you’ll probably have more money to spend in retirement and leave to your heirs.

(A1b) One needs many years of relevant data before deciding if a fund manager is truly skilled in rational investing or just lucky. Hence, by the time you have enough statistically relevant data to rely on, the manager is likely nearing retirement.

At block 2004, process 2000 involves filtering each obtained result by relevance. Conclusive answer application 102 determines a similarity between each search result of the search results and the seed sentence (e.g., as provided in block 2001). If a determined similarity is low, then conclusive answer application 102 compute a similarity for a preceding or consecutive sentence from the search results.

Relevance can be determined by using syntactic generalization. Conclusive answer application 102 can create a syntactic parse tree for each obtained result. Each of these syntactic parse trees can be compared to a syntactic parse tree generated from the utterance (e.g., at block 2001). The bag-of-words approach is extended towards extracting commonalities between the syntactic parse trees of the seed sentence and the sentence(s) obtained mined from another source (e.g., on the web). Syntactic generalization allows a domain-independent semantic measure of topical similarity between a pair of sentences, without it combination of sentences mined on the web would not form a meaningful text.

Common entities between seed and obtained sentence can be verified and an appropriateness metric obtained. The metric can include a syntactic generalization score (the cardinality of maximal common system of syntactic subtrees). For two words of the same part of speech (POS), their generalization is the same word with the POS. If the lemmas for the two words are different but the POS is the same, then the POS remains in the result. If lemmas are the same but POS is different, lemma stays in the result. A lemma represents a word without the related part-of-speech information.

To illustrate this concept, consider an example of two natural language expressions. The meanings of the expressions are represented by logic formulas. The unification and anti-unification of these formulas are constructed. Some words (entities) are mapped to predicates, some are mapped into their arguments, and some other words do not explicitly occur in logic form representation but indicate the above instantiation of predicates with arguments.

Consider the following two sentences “camera with digital zoom” and “camera with zoom for beginners.” To express the meanings, the following logic predicates are used:

camera(name\_of\_feature, type\_of\_users) and  
zoom(type\_of\_zoom).

Note that this is a simplified example, and as such, may have a reduced number of arguments as compared to more typical examples. Continuing the example, the above expressions can be represented as:

camera(zoom(digital), AnyUser),  
camera(zoom(AnyZoom), beginner)

According to the notation, variables (non-instantiated values, not specified in NL expressions) are capitalized.

Given the above pair of formulas, unification computes their most general specialization camera(zoom(digital), beginner), and anti-unification computes their most specific generalization, camera(zoom(AnyZoom), AnyUser).

At the syntactic level, the expressions are subjected to a generalization (``) of two noun phrases as: {NN-camera, PRP-with, [digital], NN-zoom [for beginners]}. The expressions in square brackets are eliminated because they occur in one expression but not occur in the other. As a result, obtain{NN-camera, PRP-with, NN-zoom}, which is a syntactic analog of semantic generalization, is obtained.

The purpose of an abstract generalization is to find commonality between portions of text at various semantic levels. Generalization operation occurs on the one or more levels. Examples of levels are paragraph level, sentence level, phrase level, and word level.

At each level (except word-level), individual words, the result of generalization of two expressions is a set of expressions. In such set, for each pair of expressions so that one is less general than other, the latter is eliminated. Generalization of two sets of expressions is a set of sets which are the results of pair-wise generalization of these expressions.

Only a single generalization exists for a pair of words: if words are the same in the same form, the result is a node with this word in this form. To involve word2vec models (Mikolov et al., 2015), compute generalization of two different words, the following rule is used. If subject1=subject2, then subject1^subject2=<subject1, POS(subject1), 1>. Otherwise, if they have the same part-of-speech, subject1^subject2=<\*,POS(subject1), word2vecDistance(subject1^subject2)>. If part-of-speech is different, generalization is an empty tuple. It cannot be further generalized.

For a pair of phrases, generalization includes all maximum ordered sets of generalization nodes for words in phrases so that the order of words is retained. In the following example,

“To buy digital camera today, on Monday.”

“Digital camera was a good buy today, first Monday of the month.”

Generalization is {<JJ-digital, NN-camera>, <NN-today, ADV, Monday>}, where the generalization for noun phrases is followed by the generalization for an adverbial phrase. Verb buy is excluded from both generalizations because it occurs in a different order in the above phrases. Buy-digital-camera is not a generalization phrase because buy occurs in different sequence with the other generalization nodes.

Therefore, a sentence similarity assessment can be expressed via a generalization operator.

A^A1a=RST-Condition (VP (. . . , NP rational investing), \*-retire)

A^A1b=NP rational investing), \*-retire.

In the first search result A1a retire and rational investing are connected in the similar way to the seed S: relational investing is connected by the rhetorical relation Condition to the phrase including retire. In A1b syntactic matching part is the same but these phrases occur in two different sentences and are related in a much more complex indirect way than in the seed. Hence A1a is a good fragment to include in the conclusive answer and A1b is not so good.

Returning to FIG. 20, at block 2005, process 2000 involves filtering each obtained result by a presence of argumentation. Conclusive answer application 102 can determine whether the obtained text is opinionated or argumentative. Communicative discourse trees can be used. For example, by forming a communicative discourse tree for

obtained text and providing the obtained text to a trained machine learning model trained to detect opinionated text, conclusive answer application 102 can determine whether the text contains an opinion. The trained machine learning model uses on mental states and/or communicative actions identified in the communicative discourse tree.

For example, conclusive answer application 102 can determine a presence of argumentation in a sentence by forming a communicative discourse tree for the sentence and applying a trained machine learning model (e.g., rhetoric agreement classifier 120) to the communicative discourse tree. The communicative discourse tree includes a root node. Examples are discussed herein with respect to FIGS. 13 and 15. Conclusive answer application 102 applies the trained machine learning model to the communicative discourse tree. The rhetoric agreement classifier 120 uses standard machine learning techniques. For example, conclusive answer application 102 can train rhetoric agreement classifier 120 with positive and negative classes communicative discourse trees or tree pairs. The positive class includes request-response pairs from text that includes argumentation and the negative class includes pairs from text that does not include argumentation. Examples of text with argumentation include product and service reviews.

At block 2006, process 2000 involves filtering each obtained result by appropriateness (or cohesiveness). Determining appropriateness can be accomplished by applying grammar rules. For example, results that include verbs in the imperative form should be excluded because such results are more likely to reflect advertisements or sales pitches. Selected fragments to be included into a conclusive answer should include opinionated sentences and avoid auxiliary comments, elaborations on topics which are not central to the topic of the seed and other supplementary parts of a document. These roles of each sentence can be determined by CDT of each document paragraph.

At block 2007, process 2000 involves formatting the filtered textual content. In some cases, conclusive answer application 102 can reformat the filtered content to better fit together, and joined in paragraphs. Once text fragments for a section are obtained, conclusive answer application 102 finds an optimal order to form the section text. For example, if both above text fragments are accepted (not just the first one), the second should follow the first since it contains the conclusion . . . Hence . . . And both these fragments are related to the same main entity. Still, the resultant text would not read well since there is a strong deviation of topics towards finding an account manager, which is not the main topic of this section. Given an unordered set of text fragments or paragraph, cohesiveness of the resultant text cannot be ensured. Instead, an optimal order for these fragments can be found, which minimizes the disturbance of content flow and coherence of the resultant text.

Process 2000 can be repeated for each utterance, according to the determined structure of the answer. For example, conclusive answer application 102 can iterate through each utterance. Conclusive answer application 102 can also combine sections in documents as appropriate and add reference sections for each section.

#### Modeling the Content Structure of Texts

Certain aspects can model a content structure of texts within a specific domain in terms of the attributes of an entity this texts expresses and the order in which these topics appear. Some research intended to characterize texts in terms of domain-independent rhetorical elements, such as schema items (McKeown, 1985) or rhetorical relations (Mann and Thompson, 1988; Marcu, 1997). Conversely, (Barzilay and

Lee 2004) focus on content, domain-dependent dimension of the structure of text. They present an effective knowledge-lean method for learning content models from unannotated documents, utilizing a novel adaptation of algorithms for Hidden Markov Models. Certain aspects perform two complementary tasks: information ordering and extractive summarization. The experiments showed that incorporating content models in these applications gives a substantial improvement.

In general, the flow of text is determined by the topic change: how attributes of an entity evolve. (Barzilay and Lee 2004) designed a model that can specify, for example, that articles about mountains typically contain information about height, climate, assents, and climbers. Instead of manually determining the evolution of attributes (the topics for a given domain) a distributional view can be taken. It is possible to machine learn these patterns of attribute evolution directly from un-annotated texts via analysis of word distribution patterns. (Harris 1982) wrote that a number of word recurrence patterns is correlated with various types of discourse structure type.

Advantages of a distributional perspective include both drastic reduction in human effort and recognition of “topics” that might not occur to a human expert and yet, when explicitly modeled, aid in applications. Of course, the success of the distributional approach depends on the existence of recurrent patterns. In arbitrary document collections, such recurrent patterns might be too variable to be easily detected by statistical means. However, research has shown that texts from the same domain tend to exhibit high similarity (Wray, 2002). At the same time, from the cognitive science perspective, this similarity is not random and is instead systematic, since text structure facilitates a text comprehension by readers and their capability of recall (Bartlett, 1932).

We assume that text chunks convey information about a single attribute of an entity (a single topic). Specifying the length of text chunks can defines the granularity of the induced attribute/topic: we select the average paragraph length.

We will build a content model as a Hidden-Markov Model in which each state  $s$  corresponds to a distinct topic and generates sentences relevant to that topic according to a state-specific language model  $p_s$ . Note that standard  $n$ -gram language models can therefore be considered to be degenerate (single-state) content models. State transition probabilities give the probability of changing from a given topic to another, thereby capturing constraints attribute evolution (topic shift).

In our implementation, we use bigram language models, so that the probability of an  $n$ -word sentence  $x=w_1 w_2 \dots w_n$  being generated by a state  $s$

$$p_s(x)=\prod_{i=1}^n p_s(w_i|w_{i-1}).$$

We will now describe state bigram probabilities  $p_s(w_i|w_{i-1})$ . To initialize a set of attributes by partitioning all of the paragraphs (or text chunks) from the documents in a given domain-specific collection into clusters. First, we create clusters via complete-link clustering, measuring sentence similarity by the cosine metric using word bigrams as features. Then, given our knowledge that documents may sometimes discuss new and/or irrelevant content as well, we create an AUX cluster by merging together all clusters containing  $\#$  paragraphs  $< t$  (selected threshold). We rely on the assumption that such clusters consist of “outlier” sentences.

Given a set  $c=c_1, c_2, \dots, c_m$  of  $m$  clusters, where  $c_m$  is the AUX cluster, we construct a content model with corresponding states  $s_1, s_2, \dots, s_m$ . we refer to  $s_m$  as the insertion state.

For each state  $s_i$ ,  $i < m$  bigram probabilities (which induce the state’s sentence-emission probabilities) are estimated using smoothed counts from the corresponding cluster

$$p_{s_i}(w' | w) \stackrel{\text{def}}{=} \frac{f_{c_i}(ww') + \delta_1}{f_{c_i}(w) + \delta_1|V|},$$

where  $f_{c_i}(y)$  is the frequency with which word sequence  $y$  occurs within the sentences in cluster  $c_i$ , and  $V$  is the vocabulary.

We want the insertion state  $s_m$  to simulate digressions or unseen attributes. We ignore the content of AUX cluster and force the language model to be complementary to those of the other states by setting

$$p_{s_m}(w' | w) \stackrel{\text{def}}{=} \frac{1 - \max_{i:i < m} p_{s_i}(w' | w)}{\sum_{u \in V} (1 - \max_{i:i < m} p_{s_i}(u | w))}.$$

Our state-transition probability estimates arise from considering how the paragraphs from the same document are distributed across the clusters. For two clusters  $c$  and  $c'$  we define  $D(c, c')$  as the number of documents in which a paragraph from  $c$  immediately precedes one from  $c'$ .  $D(c)$  is the number of documents containing paragraphs from  $c$ . For any two states  $s_i$  and  $s_j$ ,  $i, j < m$ , we rely on the following smooth estimate of the probability of transitioning from  $s_i$  to  $s_j$ :

$$p(s_j | s_i) = \frac{D(c_i, c_j) + \delta_2}{D(c_i) + \delta_2 m}.$$

Building Answer Document Based on Similarity and Compositional Semantics

Certain aspects can model a document in a vector representation using a paragraph vector model (Le and Mikolov 2014) that computes continuous distributed vector representations of varying-length texts. The source documents’ section that are semantically close (or similar) to the desired document is identified in this vector space using cosine similarity. The structure of similar articles can then be emulated, the important sections identified and assign relevant web-content or intranet content assigned to the sections.

The entire Wikipedia to obtain  $D$ -dimensional representations of words/entities as well as documents using the paragraph vector distributed memory model (Le and Mikolov, 2014). Similar articles are identified using cosine similarity between the vector representations of the missing entity and representations of the existing entities (entities that have corresponding articles). Content from the similar articles are used to train multi-class classifiers that can assign web-retrieved content on the red-linked entity to relevant sections of the article. The paragraph vector distributed memory model is used to identify similar documents to rely upon on one hand and also an inference of vector representations of new paragraphs retrieved from the web on the other hand.

We take a sequence of words from a similar document and approach the last word that can be reused. Then we attempt

to predict the next word using PV-DM. The PV-DM model is based on the principle that several contexts sampled from the paragraph can be used to predict the next word. Given a sequence of  $T$  words ( $w_1, w_2, \dots, w_T$ ), the task is to maximize the average log probability. In the top equation,  $c$  is the size of the context (number of words before and after the current word to be used for training). The conditional probability of  $w_{t+j}$  given  $w_t$  is given by the softmax function (Bridle 1990) in bottom equation, where  $v_{w_{t+j}}$  and  $v_w$  refers to the output and the input vector representations of the word  $w$ , respectively.  $W$  refers to the total number of words in the vocabulary.

$$F = \frac{1}{T} \sum_{t=1}^{t=T} \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{t+j} | w_t)$$

$$p(w_{t+j} | w_t) = \frac{\exp(v_{w_{t+j}}^T v_{w_t})}{\sum_{w=1}^W \exp(v_w^T v_{w_t})}$$

#### Exemplary Computing Systems

FIG. 21 depicts a simplified diagram of a distributed system 2100 for implementing one of the aspects. In the illustrated aspect, distributed system 2100 includes one or more client computing devices 2102, 2104, 2106, and 2108, which are configured to execute and operate a client application such as a web browser, proprietary client (e.g., Oracle Forms), or the like over one or more network(s) 2110. Server 2112 may be communicatively coupled with remote client computing devices 2102, 2104, 2106, and 2108 via network 2110.

In various aspects, server 2112 may be adapted to run one or more services or software applications provided by one or more of the components of the system. The services or software applications can include nonvirtual and virtual environments. Virtual environments can include those used for virtual events, tradeshow, simulators, classrooms, shopping exchanges, and enterprises, whether two- or three-dimensional (3D) representations, page-based logical environments, or otherwise. In some aspects, these services may be offered as web-based or cloud services or under a Software as a Service (SaaS) model to the users of client computing devices 2102, 2104, 2106, and/or 2108. Users operating client computing devices 2102, 2104, 2106, and/or 2108 may in turn utilize one or more client applications to interact with server 2112 to utilize the services provided by these components.

In the configuration depicted in the figure, the software components 2118, 2120 and 2122 of system 2100 are shown as being implemented on server 2112. In other aspects, one or more of the components of system 2100 and/or the services provided by these components may also be implemented by one or more of the client computing devices 2102, 2104, 2106, and/or 2108. Users operating the client computing devices may then utilize one or more client applications to use the services provided by these components. These components may be implemented in hardware, firmware, software, or combinations thereof. It should be appreciated that various different system configurations are possible, which may be different from distributed system 2100. The aspect shown in the figure is thus one example of a distributed system for implementing an aspect system and is not intended to be limiting.

Client computing devices 2102, 2104, 2106, and/or 2108 may be portable handheld devices (e.g., an iPhone®, cellular telephone, an iPad®, computing tablet, a personal digital assistant (PDA)) or wearable devices (e.g., a Google Glass® head mounted display), running software such as Microsoft Windows Mobile®, and/or a variety of mobile operating systems such as iOS, Windows Phone, Android, BlackBerry 10, Palm OS, and the like, and being Internet, e-mail, short message service (SMS), Blackberry®, or other communication protocol enabled. The client computing devices can be general purpose personal computers including, by way of example, personal computers and/or laptop computers running various versions of Microsoft Windows®, Apple Macintosh®, and/or Linux operating systems. The client computing devices can be workstation computers running any of a variety of commercially-available UNIX® or UNIX-like operating systems, including without limitation the variety of GNU/Linux operating systems, such as for example, Google Chrome OS. Alternatively, or in addition, client computing devices 2102, 2104, 2106, and 2108 may be any other electronic device, such as a thin-client computer, an Internet-enabled gaming system (e.g., a Microsoft Xbox gaming console with or without a Kinect® gesture input device), and/or a personal messaging device, capable of communicating over network(s) 2110.

Although exemplary distributed system 2100 is shown with four client computing devices, any number of client computing devices may be supported. Other devices, such as devices with sensors, etc., may interact with server 2112.

Network(s) 2110 in distributed system 2100 may be any type of network familiar to those skilled in the art that can support data communications using any of a variety of commercially-available protocols, including without limitation TCP/IP (transmission control protocol/Internet protocol), SNA (systems network architecture), IPX (Internet packet exchange), AppleTalk, and the like. Merely by way of example, network(s) 2110 can be a local area network (LAN), such as one based on Ethernet, Token-Ring and/or the like. Network(s) 2110 can be a wide-area network and the Internet. It can include a virtual network, including without limitation a virtual private network (VPN), an intranet, an extranet, a public switched telephone network (PSTN), an infra-red network, a wireless network (e.g., a network operating under any of the Institute of Electrical and Electronics (IEEE) 802.21 suite of protocols, Bluetooth®, and/or any other wireless protocol); and/or any combination of these and/or other networks.

Server 2112 may be composed of one or more general purpose computers, specialized server computers (including, by way of example, PC (personal computer) servers, UNIX® servers, mid-range servers, mainframe computers, rack-mounted servers, etc.), server farms, server clusters, or any other appropriate arrangement and/or combination. Server 2112 can include one or more virtual machines running virtual operating systems, or other computing architectures involving virtualization. One or more flexible pools of logical storage devices can be virtualized to maintain virtual storage devices for the server. Virtual networks can be controlled by server 2112 using software defined networking. In various aspects, server 2112 may be adapted to run one or more services or software applications described in the foregoing disclosure. For example, server 2112 may correspond to a server for performing processing described above according to an aspect of the present disclosure.

Server 2112 may run an operating system including any of those discussed above, as well as any commercially available server operating system. Server 2112 may also run any

of a variety of additional server applications and/or mid-tier applications, including HTTP (hypertext transport protocol) servers, FTP (file transfer protocol) servers, CGI (common gateway interface) servers, JAVA® servers, database servers, and the like. Exemplary database servers include without limitation those commercially available from Oracle, Microsoft, Sybase, IBM (International Business Machines), and the like.

In some implementations, server **2112** may include one or more applications to analyze and consolidate data feeds and/or event updates received from users of client computing devices **2102**, **2104**, **2106**, and **2108**. As an example, data feeds and/or event updates may include, but are not limited to, Twitter® feeds, Facebook® updates or real-time updates received from one or more third party information sources and continuous data streams, which may include real-time events related to sensor data applications, financial tickers, network performance measuring tools (e.g., network monitoring and traffic management applications), click-stream analysis tools, automobile traffic monitoring, and the like. Server **2112** may also include one or more applications to display the data feeds and/or real-time events via one or more display devices of client computing devices **2102**, **2104**, **2106**, and **2108**.

Distributed system **2100** may also include one or more databases **2114** and **2116**. Databases **2114** and **2116** may reside in a variety of locations. By way of example, one or more of databases **2114** and **2116** may reside on a non-transitory storage medium local to (and/or resident in) server **2112**. Alternatively, databases **2114** and **2116** may be remote from server **2112** and in communication with server **2112** via a network-based or dedicated connection. In one set of aspects, databases **2114** and **2116** may reside in a storage-area network (SAN). Similarly, any necessary files for performing the functions attributed to server **2112** may be stored locally on server **2112** and/or remotely, as appropriate. In one set of aspects, databases **2114** and **2116** may include relational databases, such as databases provided by Oracle, that are adapted to store, update, and retrieve data in response to SQL-formatted commands.

FIG. **22** is a simplified block diagram of one or more components of a system environment **2200** by which services provided by one or more components of an aspect system may be offered as cloud services, in accordance with an aspect of the present disclosure. In the illustrated aspect, system environment **2200** includes one or more client computing devices **2204**, **2206**, and **2208** that may be used by users to interact with a cloud infrastructure system **2202** that provides cloud services. The client computing devices may be configured to operate a client application such as a web browser, a proprietary client application (e.g., Oracle Forms), or some other application, which may be used by a user of the client computing device to interact with cloud infrastructure system **2202** to use services provided by cloud infrastructure system **2202**.

It should be appreciated that cloud infrastructure system **2202** depicted in the figure may have other components than those depicted. Further, the aspect shown in the figure is only one example of a cloud infrastructure system that may incorporate an aspect of the invention. In some other aspects, cloud infrastructure system **2202** may have more or fewer components than shown in the figure, may combine two or more components, or may have a different configuration or arrangement of components.

Client computing devices **2204**, **2206**, and **2208** may be devices similar to those described above for **2102**, **2104**, **2106**, and **2108**.

Although exemplary system environment **2200** is shown with three client computing devices, any number of client computing devices may be supported. Other devices such as devices with sensors, etc. may interact with cloud infrastructure system **2202**.

Network(s) **2210** may facilitate communications and exchange of data between clients **2204**, **2206**, and **2208** and cloud infrastructure system **2202**. Each network may be any type of network familiar to those skilled in the art that can support data communications using any of a variety of commercially-available protocols, including those described above for network(s) **2110**.

Cloud infrastructure system **2202** may comprise one or more computers and/or servers that may include those described above for server **2122**.

In certain aspects, services provided by the cloud infrastructure system may include a host of services that are made available to users of the cloud infrastructure system on demand, such as online data storage and backup solutions, Web-based e-mail services, hosted office suites and document collaboration services, database processing, managed technical support services, and the like. Services provided by the cloud infrastructure system can dynamically scale to meet the needs of its users. A specific instantiation of a service provided by cloud infrastructure system is referred to herein as a “service instance.” In general, any service made available to a user via a communication network, such as the Internet, from a cloud service provider’s system is referred to as a “cloud service.” Typically, in a public cloud environment, servers and systems that make up the cloud service provider’s system are different from the customer’s own on-premises servers and systems. For example, a cloud service provider’s system may host an application, and a user may, via a communication network such as the Internet, on demand, order and use the application.

In some examples, a service in a computer network cloud infrastructure may include protected computer network access to storage, a hosted database, a hosted web server, a software application, or other service provided by a cloud vendor to a user, or as otherwise known in the art. For example, a service can include password-protected access to remote storage on the cloud through the Internet. As another example, a service can include a web service-based hosted relational database and a script-language middleware engine for private use by a networked developer. As another example, a service can include access to an email software application hosted on a cloud vendor’s web site.

In certain aspects, cloud infrastructure system **2202** may include a suite of applications, middleware, and database service offerings that are delivered to a customer in a self-service, subscription-based, elastically scalable, reliable, highly available, and secure manner. An example of such a cloud infrastructure system is the Oracle Public Cloud provided by the present assignee.

Large volumes of data, sometimes referred to as big data, can be hosted and/or manipulated by the infrastructure system on many levels and at different scales. Such data can include data sets that are so large and complex that it can be difficult to process using typical database management tools or traditional data processing applications. For example, terabytes of data may be difficult to store, retrieve, and process using personal computers or their rack-based counterparts. Such sizes of data can be difficult to work with using most current relational database management systems and desktop statistics and visualization packages. They can require massively parallel processing software running thousands of server computers, beyond the structure of com-

monly used software tools, to capture, curate, manage, and process the data within a tolerable elapsed time.

Extremely large data sets can be stored and manipulated by analysts and researchers to visualize large amounts of data, detect trends, and/or otherwise interact with the data. Tens, hundreds, or thousands of processors linked in parallel can act upon such data in order to present it or simulate external forces on the data or what it represents. These data sets can involve structured data, such as that organized in a database or otherwise according to a structured model, and/or unstructured data (e.g., emails, images, data blobs (binary large objects), web pages, complex event processing). By leveraging an ability of an aspect to relatively quickly focus more (or fewer) computing resources upon an objective, the cloud infrastructure system may be better available to carry out tasks on large data sets based on demand from a business, government agency, research organization, private individual, group of like-minded individuals or organizations, or other entity.

In various aspects, cloud infrastructure system **2202** may be adapted to automatically provision, manage and track a customer's subscription to services offered by cloud infrastructure system **2202**. Cloud infrastructure system **2202** may provide the cloud services via different deployment models. For example, services may be provided under a public cloud model in which cloud infrastructure system **2202** is owned by an organization selling cloud services (e.g., owned by Oracle) and the services are made available to the general public or different industry enterprises. As another example, services may be provided under a private cloud model in which cloud infrastructure system **2202** is operated solely for a single organization and may provide services for one or more entities within the organization. The cloud services may also be provided under a community cloud model in which cloud infrastructure system **2202** and the services provided by cloud infrastructure system **2202** are shared by several organizations in a related community. The cloud services may also be provided under a hybrid cloud model, which is a combination of two or more different models.

In some aspects, the services provided by cloud infrastructure system **2202** may include one or more services provided under Software as a Service (SaaS) category, Platform as a Service (PaaS) category, Infrastructure as a Service (IaaS) category, or other categories of services including hybrid services. A customer, via a subscription order, may order one or more services provided by cloud infrastructure system **2202**. Cloud infrastructure system **2202** then performs processing to provide the services in the customer's subscription order.

In some aspects, the services provided by cloud infrastructure system **2202** may include, without limitation, application services, platform services and infrastructure services. In some examples, application services may be provided by the cloud infrastructure system via a SaaS platform. The SaaS platform may be configured to provide cloud services that fall under the SaaS category. For example, the SaaS platform may provide capabilities to build and deliver a suite of on-demand applications on an integrated development and deployment platform. The SaaS platform may manage and control the underlying software and infrastructure for providing the SaaS services. By utilizing the services provided by the SaaS platform, customers can utilize applications executing on the cloud infrastructure system. Customers can acquire the application services without the need for customers to purchase separate licenses and support. Various different SaaS services may be pro-

vided. Examples include, without limitation, services that provide solutions for sales performance management, enterprise integration, and business flexibility for large organizations.

In some aspects, platform services may be provided by the cloud infrastructure system via a PaaS platform. The PaaS platform may be configured to provide cloud services that fall under the PaaS category. Examples of platform services may include without limitation services that enable organizations (such as Oracle) to consolidate existing applications on a shared, common architecture, as well as the ability to build new applications that leverage the shared services provided by the platform. The PaaS platform may manage and control the underlying software and infrastructure for providing the PaaS services. Customers can acquire the PaaS services provided by the cloud infrastructure system without the need for customers to purchase separate licenses and support. Examples of platform services include, without limitation, Oracle Java Cloud Service (JCS), Oracle Database Cloud Service (DBCS), and others.

By utilizing the services provided by the PaaS platform, customers can employ programming languages and tools supported by the cloud infrastructure system and also control the deployed services. In some aspects, platform services provided by the cloud infrastructure system may include database cloud services, middleware cloud services (e.g., Oracle Fusion Middleware services), and Java cloud services. In one aspect, database cloud services may support shared service deployment models that enable organizations to pool database resources and offer customers a Database as a Service in the form of a database cloud. Middleware cloud services may provide a platform for customers to develop and deploy various business applications, and Java cloud services may provide a platform for customers to deploy Java applications, in the cloud infrastructure system.

Various different infrastructure services may be provided by an IaaS platform in the cloud infrastructure system. The infrastructure services facilitate the management and control of the underlying computing resources, such as storage, networks, and other fundamental computing resources for customers utilizing services provided by the SaaS platform and the PaaS platform.

In certain aspects, cloud infrastructure system **2202** may also include infrastructure resources **2230** for providing the resources used to provide various services to customers of the cloud infrastructure system. In one aspect, infrastructure resources **2230** may include pre-integrated and optimized combinations of hardware, such as servers, storage, and networking resources to execute the services provided by the PaaS platform and the SaaS platform.

In some aspects, resources in cloud infrastructure system **2202** may be shared by multiple users and dynamically re-allocated per demand. Additionally, resources may be allocated to users in different time zones. For example, cloud infrastructure system **2202** may enable a first set of users in a first time zone to utilize resources of the cloud infrastructure system for a specified number of hours and then enable the re-allocation of the same resources to another set of users located in a different time zone, thereby maximizing the utilization of resources.

In certain aspects, a number of internal shared services **2232** may be provided that are shared by different components or modules of cloud infrastructure system **2202** and by the services provided by cloud infrastructure system **2202**. These internal shared services may include, without limitation, a security and identity service, an integration service, an enterprise repository service, an enterprise manager ser-



vice, a virus scanning and white list service, a high availability, backup and recovery service, service for enabling cloud support, an email service, a notification service, a file transfer service, and the like.

In certain aspects, cloud infrastructure system **2202** may provide comprehensive management of cloud services (e.g., SaaS, PaaS, and IaaS services) in the cloud infrastructure system. In one aspect, cloud management functionality may include capabilities for provisioning, managing and tracking a customer's subscription received by cloud infrastructure system **2202**, and the like.

In one aspect, as depicted in the figure, cloud management functionality may be provided by one or more modules, such as an order management module **2220**, an order orchestration module **2222**, an order provisioning module **2224**, an order management and monitoring module **2226**, and an identity management module **2228**. These modules may include or be provided using one or more computers and/or servers, which may be general purpose computers, specialized server computers, server farms, server clusters, or any other appropriate arrangement and/or combination.

In exemplary operation **2234**, a customer using a client device, such as client device **2204**, **2206** or **2208**, may interact with cloud infrastructure system **2202** by requesting one or more services provided by cloud infrastructure system **2202** and placing an order for a subscription for one or more services offered by cloud infrastructure system **2202**. In certain aspects, the customer may access a cloud User Interface (UI), cloud UI **2222**, cloud UI **2214** and/or cloud UI **2216** and place a subscription order via these UIs. The order information received by cloud infrastructure system **2202** in response to the customer placing an order may include information identifying the customer and one or more services offered by the cloud infrastructure system **2202** that the customer intends to subscribe to.

After an order has been placed by the customer, the order information is received via the cloud UIs, **2222**, **2214** and/or **2216**.

At operation **2236**, the order is stored in order database **2218**. Order database **2218** can be one of several databases operated by cloud infrastructure system **2202** and operated in conjunction with other system elements.

At operation **2238**, the order information is forwarded to an order management module **2220**. In some instances, order management module **2220** may be configured to perform billing and accounting functions related to the order, such as verifying the order, and upon verification, booking the order.

At operation **2240**, information regarding the order is communicated to an order orchestration module **2222**. Order orchestration module **2222** may utilize the order information to orchestrate the provisioning of services and resources for the order placed by the customer. In some instances, order orchestration module **2222** may orchestrate the provisioning of resources to support the subscribed services using the services of order provisioning module **2224**.

In certain aspects, order orchestration module **2222** enables the management of business processes associated with each order and applies business logic to determine whether an order should proceed to provisioning. At operation **2242**, upon receiving an order for a new subscription, order orchestration module **2222** sends a request to order provisioning module **2224** to allocate resources and configure those resources needed to fulfill the subscription order. Order provisioning module **2224** enables the allocation of resources for the services ordered by the customer. Order provisioning module **2224** provides a level of abstraction between the cloud services provided by cloud infrastructure

system **2200** and the physical implementation layer that is used to provision the resources for providing the requested services. Order orchestration module **2222** may thus be isolated from implementation details, such as whether or not services and resources are actually provisioned on the fly or pre-provisioned and only allocated/assigned upon request.

At operation **2244**, once the services and resources are provisioned, a notification of the provided service may be sent to customers on client devices **2204**, **2206** and/or **2208** by order provisioning module **2224** of cloud infrastructure system **2202**.

At operation **2246**, the customer's subscription order may be managed and tracked by an order management and monitoring module **2226**. In some instances, order management and monitoring module **2226** may be configured to collect usage statistics for the services in the subscription order, such as the amount of storage used, the amount data transferred, the number of users, and the amount of system up time and system down time.

In certain aspects, cloud infrastructure system **2200** may include an identity management module **2228**. Identity management module **2228** may be configured to provide identity services, such as access management and authorization services in cloud infrastructure system **2200**. In some aspects, identity management module **2228** may control information about customers who wish to utilize the services provided by cloud infrastructure system **2202**. Such information can include information that authenticates the identities of such customers and information that describes which actions those customers are authorized to perform relative to various system resources (e.g., files, directories, applications, communication ports, memory segments, etc.) Identity management module **2228** may also include the management of descriptive information about each customer and about how and by whom that descriptive information can be accessed and modified.

FIG. **23** illustrates an exemplary computer system **2300**, in which various aspects of the present invention may be implemented. The system **2300** may be used to implement any of the computer systems described above. As shown in the figure, computer system **2300** includes a processing unit **2304** that communicates with a number of peripheral subsystems via a bus subsystem **2302**. These peripheral subsystems may include a processing acceleration unit **2306**, an I/O subsystem **2308**, a storage subsystem **2318** and a communications subsystem **2324**. Storage subsystem **2318** includes tangible computer-readable storage media **2322** and a system memory **2310**.

Bus subsystem **2302** provides a mechanism for letting the various components and subsystems of computer system **2300** communicate with each other as intended. Although bus subsystem **2302** is shown schematically as a single bus, alternative aspects of the bus subsystem may utilize multiple buses. Bus subsystem **2302** may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. For example, such architectures may include an Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus, which can be implemented as a Mezzanine bus manufactured to the IEEE P2386.1 standard.

Processing unit **2304**, which can be implemented as one or more integrated circuits (e.g., a conventional microprocessor or microcontroller), controls the operation of computer system **2300**. One or more processors may be included

in processing unit **2304**. These processors may include single core or multicore processors. In certain aspects, processing unit **2304** may be implemented as one or more independent processing units **2332** and/or **2334** with single or multicore processors included in each processing unit. In other aspects, processing unit **2304** may also be implemented as a quad-core processing unit formed by integrating two dual-core processors into a single chip.

In various aspects, processing unit **2304** can execute a variety of programs in response to program code and can maintain multiple concurrently executing programs or processes. At any given time, some or all of the program code to be executed can be resident in processor(s) **2304** and/or in storage subsystem **2318**. Through suitable programming, processor(s) **2304** can provide various functionalities described above. Computer system **2300** may additionally include a processing acceleration unit **2306**, which can include a digital signal processor (DSP), a special-purpose processor, and/or the like.

I/O subsystem **2308** may include user interface input devices and user interface output devices. User interface input devices may include a keyboard, pointing devices such as a mouse or trackball, a touchpad or touch screen incorporated into a display, a scroll wheel, a click wheel, a dial, a button, a switch, a keypad, audio input devices with voice command recognition systems, microphones, and other types of input devices. User interface input devices may include, for example, motion sensing and/or gesture recognition devices such as the Microsoft Kinect® motion sensor that enables users to control and interact with an input device, such as the Microsoft Xbox® 360 game controller, through a natural user interface using gestures and spoken commands. User interface input devices may also include eye gesture recognition devices such as the Google Glass® blink detector that detects eye activity (e.g., ‘blinking’ while taking pictures and/or making a menu selection) from users and transforms the eye gestures as input into an input device (e.g., Google Glass®). Additionally, user interface input devices may include voice recognition sensing devices that enable users to interact with voice recognition systems (e.g., Siri® navigator), through voice commands.

User interface input devices may also include, without limitation, three dimensional (3D) mice, joysticks or pointing sticks, gamepads and graphic tablets, and audio/visual devices such as speakers, digital cameras, digital camcorders, portable media players, webcams, image scanners, fingerprint scanners, barcode reader 3D scanners, 3D printers, laser rangefinders, and eye gaze tracking devices. Additionally, user interface input devices may include, for example, medical imaging input devices such as computed tomography, magnetic resonance imaging, position emission tomography, medical ultrasonography devices. User interface input devices may also include, for example, audio input devices such as MIDI keyboards, digital musical instruments and the like.

User interface output devices may include a display subsystem, indicator lights, or non-visual displays such as audio output devices, etc. The display subsystem may be a cathode ray tube (CRT), a flat-panel device, such as that using a liquid crystal display (LCD) or plasma display, a projection device, a touch screen, and the like. In general, use of the term “output device” is intended to include all possible types of devices and mechanisms for outputting information from computer system **2300** to a user or other computer. For example, user interface output devices may include, without limitation, a variety of display devices that visually convey text, graphics and audio/video information

such as monitors, printers, speakers, headphones, automotive navigation systems, plotters, voice output devices, and modems.

Computer system **2300** may comprise a storage subsystem **2318** that comprises software elements, shown as being currently located within a system memory **2310**. System memory **2310** may store program instructions that are loadable and executable on processing unit **2304**, as well as data generated during the execution of these programs.

Depending on the configuration and type of computer system **2300**, system memory **2310** may be volatile (such as random access memory (RAM)) and/or non-volatile (such as read-only memory (ROM), flash memory, etc.) The RAM typically contains data and/or program modules that are immediately accessible to and/or presently being operated and executed by processing unit **2304**. In some implementations, system memory **2310** may include multiple different types of memory, such as static random access memory (SRAM) or dynamic random access memory (DRAM). In some implementations, a basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within computer system **2300**, such as during start-up, may typically be stored in the ROM. By way of example, and not limitation, system memory **2310** also illustrates application programs **2312**, which may include client applications, Web browsers, mid-tier applications, relational database management systems (RDBMS), etc., program data **2314**, and an operating system **2316**. By way of example, operating system **2316** may include various versions of Microsoft Windows®, Apple Macintosh®, and/or Linux operating systems, a variety of commercially-available UNIX® or UNIX-like operating systems (including without limitation the variety of GNU/Linux operating systems, the Google Chrome® OS, and the like) and/or mobile operating systems such as iOS, Windows® Phone, Android® OS, BlackBerry® 10 OS, and Palm® OS operating systems.

Storage subsystem **2318** may also provide a tangible computer-readable storage medium for storing the basic programming and data constructs that provide the functionality of some aspects. Software (programs, code modules, instructions) that when executed by a processor provide the functionality described above may be stored in storage subsystem **2318**. These software modules or instructions may be executed by processing unit **2304**. Storage subsystem **2318** may also provide a repository for storing data used in accordance with the present invention.

Storage subsystem **2318** may also include a computer-readable storage media reader **2320** that can further be connected to computer-readable storage media **2322**. Together and, optionally, in combination with system memory **2310**, computer-readable storage media **2322** may comprehensively represent remote, local, fixed, and/or removable storage devices plus storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information.

Computer-readable storage media **2322** containing code, or portions of code, can also include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to, volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information. This can include tangible, non-transitory computer-readable storage media such as RAM, ROM, electronically erasable programmable ROM (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disk (DVD), or other optical

storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or other tangible computer readable media. When specified, this can also include nontangible, transitory computer-readable media, such as data signals, data transmissions, or any other medium which can be used to transmit the desired information and which can be accessed by computing system **2300**.

By way of example, computer-readable storage media **2322** may include a hard disk drive that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, non-volatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD ROM, DVD, and Blu-Ray® disk, or other optical media. Computer-readable storage media **2322** may include, but is not limited to, Zip® drives, flash memory cards, universal serial bus (USB) flash drives, secure digital (SD) cards, DVD disks, digital video tape, and the like. Computer-readable storage media **2322** may also include, solid-state drives (SSD) based on non-volatile memory such as flash-memory based SSDs, enterprise flash drives, solid state ROM, and the like, SSDs based on volatile memory such as solid state RAM, dynamic RAM, static RAM, DRAM-based SSDs, magnetoresistive RAM (MRAM) SSDs, and hybrid SSDs that use a combination of DRAM and flash memory based SSDs. The disk drives and their associated computer-readable media may provide non-volatile storage of computer-readable instructions, data structures, program modules, and other data for computer system **2300**.

Communications subsystem **2324** provides an interface to other computer systems and networks. Communications subsystem **2324** serves as an interface for receiving data from and transmitting data to other systems from computer system **2300**. For example, communications subsystem **2324** may enable computer system **2300** to connect to one or more devices via the Internet. In some aspects, communications subsystem **2324** can include radio frequency (RF) transceiver components for accessing wireless voice and/or data networks (e.g., using cellular telephone technology, advanced data network technology, such as 3G, 4G or EDGE (enhanced data rates for global evolution), WiFi (IEEE 802.11 family standards, or other mobile communication technologies, or any combination thereof), global positioning system (GPS) receiver components, and/or other components. In some aspects, communications subsystem **2324** can provide wired network connectivity (e.g., Ethernet) in addition to or instead of a wireless interface.

In some aspects, communications subsystem **2324** may also receive input communication in the form of structured and/or unstructured data feeds **2326**, event streams **2328**, event updates **2323**, and the like on behalf of one or more users who may use computer system **2300**.

By way of example, communications subsystem **2324** may be configured to receive unstructured data feeds **2326** in real-time from users of social media networks and/or other communication services such as Twitter® feeds, Facebook® updates, web feeds such as Rich Site Summary (RSS) feeds, and/or real-time updates from one or more third party information sources.

Additionally, communications subsystem **2324** may also be configured to receive data in the form of continuous data streams, which may include event streams **2328** of real-time events and/or event updates **2323**, that may be continuous or unbounded in nature with no explicit end. Examples of applications that generate continuous data may include, for example, sensor data applications, financial tickers, network performance measuring tools (e.g. network monitoring and

traffic management applications), clickstream analysis tools, automobile traffic monitoring, and the like.

Communications subsystem **2324** may also be configured to output the structured and/or unstructured data feeds **2326**, event streams **2328**, event updates **2323**, and the like to one or more databases that may be in communication with one or more streaming data source computers coupled to computer system **2300**.

Computer system **2300** can be one of various types, including a handheld portable device (e.g., an iPhone® cellular phone, an iPad® computing tablet, a PDA), a wearable device (e.g., a Google Glass® head mounted display), a PC, a workstation, a mainframe, a kiosk, a server rack, or any other data processing system.

Due to the ever-changing nature of computers and networks, the description of computer system **2300** depicted in the figure is intended only as a specific example. Many other configurations having more or fewer components than the system depicted in the figure are possible. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, firmware, software (including applets), or a combination. Further, connection to other computing devices, such as network input/output devices, may be employed. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various aspects.

In the foregoing specification, aspects of the invention are described with reference to specific aspects thereof, but those skilled in the art will recognize that the invention is not limited thereto. Various features and aspects of the above-described invention may be used individually or jointly. Further, aspects can be utilized in any number of environments and applications beyond those described herein without departing from the broader spirit and scope of the specification. The specification and drawings are, accordingly, to be regarded as illustrative rather than restrictive.

What is claimed is:

1. A method of computationally fortifying an answer using syntactic parse trees, the method comprising:
  - accessing a seed sentence comprising a first plurality of text fragments;
  - syntactically parsing the seed sentence to generate a first syntactic parse tree comprising the syntactically parsed seed sentence;
  - obtaining a search result by providing, to a search engine, at least one of the first plurality of text fragments, wherein the search result comprises a second text fragment;
  - syntactically parsing the search result to generate a second syntactic parse tree comprising the syntactically parsed search result;
  - calculating a relevancy metric for the search result by:
    - identifying a common entity between the first syntactic parse tree and the second syntactic parse tree,
    - creating a generalized fragment comprising text that is associated with the common entity, and
    - deriving the relevancy metric from the generalized fragment;
  - identifying the search result as an additional fragment based on a determination that the relevancy metric is greater than a first threshold;
  - constructing a paragraph from the additional fragment; and
  - providing the paragraph to a user device.

47

2. The method of claim 1, wherein computing the relevancy metric comprises applying an additional machine learning model to the first syntactic parse tree and the second syntactic parse tree.

3. The method of claim 1, further comprising:

forming a table of contents comprising the additional fragment; and

providing the table of contents to the user device.

4. The method of claim 1, further comprising determining that the search result does not contain argumentation by applying an additional machine learning model to the first plurality of text fragments and the additional fragment, wherein the additional machine learning model is trained to determine a presence or an absence of argumentation in text.

5. The method of claim 4, wherein the additional machine learning model is trained by:

accessing a set of training data comprising a training pair, the training pair comprising a first communicative discourse tree that represents text comprising argumentation and a second communicative discourse tree that represents text without argumentation; and

providing one of the training pairs to the additional machine learning model;

receiving, from the additional machine learning model, a determined presence of argumentation;

calculating a loss function by calculating a difference between the determined presence of argumentation and an expected presence of argumentation; and

adjusting internal parameters of the additional machine learning model to minimize the loss function.

6. The method of claim 1, wherein identifying the search result as an additional fragment comprises determining that no imperative-form verbs exist in the search result.

7. The method of claim 1, further comprising calculating a metric that indicates a mental state expressed in text by applying an additional machine learning model to the first plurality of text fragments and the additional fragment, wherein the additional machine learning model is trained to determine a mental state in text, and wherein the identifying comprises determining that the metric is within a tolerance.

8. A method of computationally fortifying an answer using syntactic parse trees, the method comprising:

accessing a first seed sentence comprising text fragments and a second seed sentence comprising text fragments; generating, from the first seed sentence, a first syntactic parse tree;

generating, from the second seed sentence, a second syntactic parse tree;

identifying, from the first syntactic parse tree, a first entity;

identifying, from the second syntactic parse tree, the first entity and a second entity;

obtaining a plurality of search results by providing the first entity and the second entity to a search engine;

generating, from the search result, a third syntactic parse tree;

calculating a relevancy metric for the search result by: identifying a common entity between the first syntactic parse tree, the second syntactic parse tree, and the third syntactic parse tree,

creating a generalized fragment comprising text that is associated with the common entity, and

deriving the relevancy metric from the generalized fragment;

identifying the search result as an additional fragment based on a determination that the relevancy metric is greater than a first threshold;

48

constructing a paragraph from fragments of the first seed sentence, fragments of the second seed sentence and additional fragment; and

providing the paragraph to a user device.

9. The method of claim 8, wherein computing the relevancy metric comprises applying an additional machine learning model to the first syntactic parse tree, the second syntactic parse tree, and the third syntactic parse tree.

10. The method of claim 8, further comprising determining that search result does not contain argumentation by applying an additional machine learning model to the first syntactic parse tree, the second syntactic parse tree, and the third syntactic parse tree, wherein the additional machine learning model is trained to determine a presence or an absence of argumentation in text.

11. A system comprising:

a non-transitory computer-readable medium storing computer-executable program instructions; and

a processing device communicatively coupled to the non-transitory computer-readable medium for executing the computer-executable program instructions, wherein executing the computer-executable program instructions configures the processing device to perform operations comprising:

accessing a seed sentence comprising a first plurality of text fragments;

syntactically parsing the seed sentence to generate a first syntactic parse tree comprising the syntactically parsed seed sentence;

obtaining a search result by providing at least one of the first plurality of text fragments to a search engine, wherein the search result comprises a second text fragment;

syntactically parsing the search result to generate a second syntactic parse tree comprising the syntactically parsed search result;

calculating, a relevancy metric for the search result by: identifying a common entity between the first syntactic parse tree and the second syntactic parse tree,

creating a generalized fragment comprising text that is associated with the common entity, and

deriving the relevancy metric from the generalized fragment;

identifying the search result as an additional fragment based on a determination that the relevancy metric is greater than a first threshold;

constructing a paragraph from the additional fragment; and

providing the paragraph to a user device.

12. The system of claim 11, wherein executing the computer-executable program instructions further configures the processing device to perform operations comprising:

generating a first discourse tree for the seed sentence and a second discourse tree for the search result, wherein a discourse tree represents rhetorical relationships between text fragments, comprises a plurality of nodes, each nonterminal node representing a rhetorical relationship between two text fragments, each terminal node of the nodes of the discourse tree is associated with one of the text fragments;

generating, from the first discourse tree, a first communicative discourse tree, and from the second discourse tree, a second communicative discourse tree, wherein generating a communicative discourse tree comprises matching each fragment in a discourse tree that has a verb to a verb signature;

calculating a cohesiveness score by applying a machine learning model to the first communicative discourse tree and the second communicative discourse tree, wherein the machine learning model is trained to classify text as cohesive or incoherent; and wherein the identification is further based on a determination that the cohesiveness score is greater than a second threshold.

13. The system of claim 11, wherein computing the relevancy metric comprises applying an additional machine learning model to the first syntactic parse tree and the second syntactic parse tree.

14. The system of claim 11, wherein executing the computer-executable program instructions further configures the processing device to perform operations comprising:

forming a table of contents comprising the additional fragment; and

providing the table of contents to the user device.

15. The system of claim 11, wherein executing the computer-executable program instructions further configures the processing device to perform operations comprising determining that search result does not contain argumentation by applying an additional machine learning model to the first

plurality of text fragments and the additional fragment, wherein the additional machine learning model is trained to determine a presence or an absence of argumentation in text.

16. The system of claim 11, wherein executing the computer-executable program instructions further configures the processing device to perform operations comprising calculating a metric that indicates a mental state expressed in text by applying an additional machine learning model to the first plurality of text fragments and the additional fragment, wherein the additional machine learning model is trained to determine a mental state in text, and wherein the identifying comprises determining that the metric is within a tolerance.

17. The system of claim 11, wherein executing the computer-executable program instructions further configures the processing device to perform operations comprising calculating a cohesiveness score by applying a machine learning model to the first plurality of text fragments and the additional fragment, wherein the machine learning model is trained to classify text as cohesive or not cohesive and wherein the identifying further comprises determining that the cohesiveness score is greater than a threshold.

\* \* \* \* \*