

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2021360822 B2**

(54) Title
SYSTEM AND/OR METHOD FOR SEMANTIC PARSING OF AIR TRAFFIC CONTROL AUDIO

(51) International Patent Classification(s)
G10L 15/26 (2006.01) **G10L 15/22** (2006.01)
G06F 16/24 (2019.01) **G10L 21/00** (2013.01)
G08G 5/00 (2006.01)

(21) Application No: **2021360822** (22) Date of Filing: **2021.10.13**

(87) WIPO No: **WO22/081669**

(30) Priority Data

(31) Number	(32) Date	(33) Country
63/090,898	2020.10.13	US

(43) Publication Date: **2022.04.21**

(44) Accepted Journal Date: **2024.03.14**

(71) Applicant(s)
MERLIN LABS, INC.

(72) Inventor(s)
PUST, Michael;BONDARYK, Joseph;GEORGE, Matthew

(74) Agent / Attorney
PIPERS, PO Box 5298, Auckland, 1141, NZ

(56) Related Art
US 5652897 A
US 2018/0182250 A1
CN 107315737 A



(51) International Patent Classification:

G10L 15/26 (2006.01) G10L 15/22 (2006.01)
G06F 16/24 (2019.01) G10L 21/00 (2013.01)
G08G 5/00 (2006.01)

(21) International Application Number:

PCT/US2021/054723

(22) International Filing Date:

13 October 2021 (13.10.2021)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/090,898 13 October 2020 (13.10.2020) US

(71) Applicant: **MERLIN LABS, INC.** [US/US]; 31 St. James Avenue, 6th Floor, Boston, MA 02116 (US).

(72) Inventors: **PUST, Michael**; 31 St. James Avenue, 6th Floor, Boston, MA 02116 (US). **BONDARYK, Joseph**; 31 St. James Avenue, 6th Floor, Boston, MA 02116 (US). **GEORGE, Matthew**; 31 St. James Avenue, 6th Floor, Boston, MA 02116 (US).

(74) Agent: **LIN, Diana**; 501 3rd Street, Suite 300, San Francisco, CA 94107 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: SYSTEM AND/OR METHOD FOR SEMANTIC PARSING OF AIR TRAFFIC CONTROL AUDIO

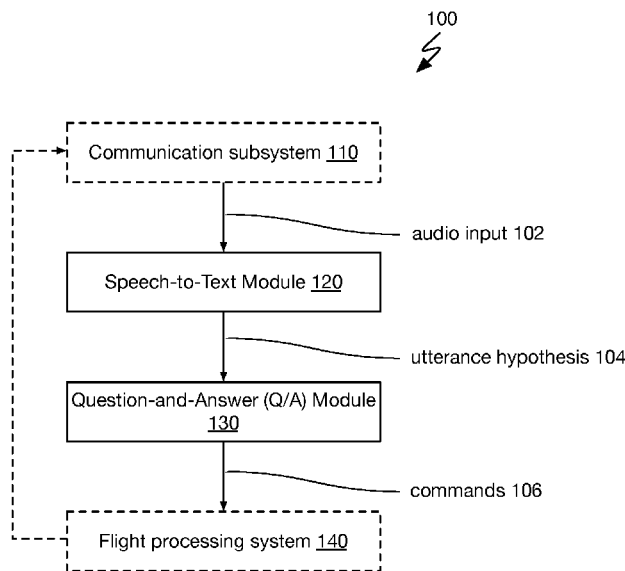


FIGURE 1

(57) Abstract: The method S200 can include: at an aircraft, receiving an audio utterance from air traffic control S210, converting the audio utterance to text, determining commands from the text using a question-and-answer model S240, and optionally controlling the aircraft based on the commands S250. The method functions to automatically interpret flight commands from the air traffic control (ATC) stream.



TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

SYSTEM AND/OR METHOD FOR SEMANTIC PARSING OF AIR TRAFFIC CONTROL
AUDIO

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 63/090,898, filed 13-OCT-2020, which is incorporated herein in its entirety by this reference.

TECHNICAL FIELD

[0002] This invention relates generally to the aviation field, and more specifically to a new and useful semantic parser in the aviation field.

BRIEF DESCRIPTION OF THE FIGURES

[0003] FIGURE 1 is a schematic representation of a variant of the system.

[0004] FIGURE 2 is a diagrammatic representation of a variant of the method.

[0005] FIGURE 3 is a diagrammatic representation of a variant of the method.

[0006] FIGURE 4 is a diagrammatic representation of an example of training an ASR model in a variant of the method.

[0007] FIGURE 5 is a diagrammatic representation of an example of training a language model in a variant of the method.

[0008] FIGURE 6 is a diagrammatic representation of an example of training a Question/Answer model in a variant of the method.

[0009] FIGURE 7 is a schematic representation of an example of the system.

[0010] FIGURE 8 is a graphical representation of an example of a domain expert evaluation tool in a variant of the method.

[0011] FIGURE 9 is a diagrammatic representation of a variant of the method.

[0012] FIGURES 10A-D are diagrammatic representations of a first, second, third, and fourth variant of the system, respectively.

[0013] FIGURES 11A-C are first, second, and third examples of tree-based query structures, respectively.

[0014] FIGURE 12 is a diagrammatic representation of a variant of the system and/or method.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0015] The following description of the preferred embodiments of the invention is not intended to limit the invention to these preferred embodiments, but rather to enable any person skilled in the art to make and use this invention.

1. Overview.

[0016] The method, an example of which is shown in FIGURE 2, can include performing inference using the system S200; and can optionally include training the system components S100. The method functions to automatically interpret flight commands from a stream of air traffic control (ATC) radio communications. The method can additionally or alternatively function to train and/or update a natural language processing system based on ATC communications.

[0017] The performing inference S200 can include: at an aircraft, receiving an audio utterance from air traffic control S210, converting the audio utterance into a predetermined format S215, determining commands using a question-and-answer model S240, and optionally controlling the aircraft based on the commands S250 (example shown in FIGURE 3). The method functions to automatically interpret flight commands from the air traffic control (ATC) stream. The flight commands can be: automatically used to control aircraft flight; presented to a user (e.g., pilot, a remote teleoperator); relayed to an auto-pilot system in response to a user (e.g., pilot) confirmation; and/or otherwise used.

[0018] In an illustrative example, the method can receive ATC audio stream, convert the ATC audio stream to ATC text, and provide the ATC text (as the reference text) and a predetermined set of queries (each associated with a different flight command parameter) to an ATC-tuned question and answer model (e.g., ATC-tuned BERT), which analyzes an ATC text for the query answers. The query answers (e.g., responses of the

question and answer model) can then be used to select follow-up queries and/or fill out a command parameter value, which can be used for direct or indirect aircraft control. The ATC audio stream can be converted to the ATC text using an ATC-tuned integrated sentence boundary detection and automatic speech recognition model (SBD/ASR model) and an ATC-tuned language model, wherein an utterance hypotheses (e.g., a sentence hypothesis, utterance by an individual speaker, etc.) can be selected for inclusion in the ATC text based on the joint score from the SBD/ASR model and the language model.

[0019] S200 can be performed using a system 100 including a Speech-to-Text module and a question and answer (Q/A) module (e.g., cooperatively forming a semantic parser). The system functions to interpret audio air traffic control (ATC) audio into flight commands, and can optionally control the aircraft based on the set of flight commands.

[0020] The system 100 is preferably mounted to, installed on, integrated into, and/or configured to operate with any suitable vehicle (e.g., the system can include the vehicle). Preferably, the vehicle is an aircraft, but can alternately be a watercraft, land-based vehicle, spacecraft, and/or any other suitable vehicle. The system can be integrated with any suitable aircraft, such as a rotorcraft (e.g., helicopter, multi-copter), fixed-wing aircraft (e.g., airplane), VTOL, STOL, lighter-than-air aircraft, multi-copter, and/or any other suitable aircraft. However, the vehicle can be an autonomous aircraft, unmanned aircraft (UAV), manned aircraft (e.g., with a pilot, with an unskilled operator executing primary aircraft control), semi-autonomous aircraft, and/or any other suitable aircraft. Hereinafter, the term ‘vehicle’ can refer to any suitable aircraft, and the term ‘aircraft’ can likewise refer to any other suitable vehicle.

[0021] The system is preferably equipped on an autonomous aircraft, which is configured to control the aircraft according to a set of flight commands using a flight processing system without user (e.g., pilot) intervention. Alternatively, the system can be equipped on a semi-autonomous vehicle and/or human-operated vehicle as a flight aid. In a first variant, the system can display ATC commands to a user (e.g., pilot) and/or relay ATC commands to an auto-pilot system in response to a user (e.g., pilot) confirmation.

[0022] The term “tuned,” as referenced in regard to neural networks, language models, or otherwise, can be understood to relate to tuning (e.g., adjusting) model

parameters using training data. Accordingly, an ATC-tuned network can be understood as having parameters tuned based on ATC audio and/or ATC-specific semantic training data (as opposed to a network dedicated to a specific radiofrequency band).

2. Benefits.

[0023] Variations of the technology can afford several benefits and/or advantages.

[0024] First, variants of the system and method can confer increased semantic parsing accuracy over conventional systems by utilizing a multiple-query (or repeated question-and-answer) approach, for example by neural network (e.g., BERT), since existing deep neural network models have high intrinsic accuracy in responding to these types of questions.

[0025] Second, variations of this technology utilizing a multiple-query approach which asks natural language questions (e.g., “message intended for DAL456?”; “topics?”; “heading values?”; etc.) of a neural network can improve the interpretability and/or auditability of the semantic parser. In such variants, a specific module/model/query of the system can be identified as a point of failure when a user rejects a command, which can be used to further train/improve the system. In some variants, the multi-query approach can additionally enable portions of the semantic parser to be trained based on partial and/or incomplete tagged responses (e.g., which can be sufficient to answer a subset of the queries used to extract a command from an ATC transcript). As an example, training data can be used when values and/or aircraft tail numbers are not identified and/or validated within a training dataset.

[0026] Third, variations of this technology can enable semantic parsing of ATC utterances without the use of grammar rules or syntax – which can be time intensive to develop, slow to execute, and yield inaccurate results (particularly when handling edge case scenarios or unusual speech patterns). In an example: as a conversation between ATC and an aircraft continues, the ATC controller and the pilot often shorten phrases and/or deviate from the standard speech template, which can severely impact the efficacy of grammar/syntax-based NLP approaches. In variants, the system and/or method can convert unformatted audio, syntactically inconsistent (non-standardized) audio, and/or non-uniform audio data or corresponding ATC transcript into a standardized/formatted

data input (e.g., as may be accepted/interpreted by a certified aircraft processor). In variants, standardized inputs can be utilized to certify aircraft systems in a deterministically testable manner. As an example, the technology can be used to convert an arbitrarily large number of audio signals into a substantially finite set of commands (e.g., with bounded ranges of values corresponding to a predetermined set of aircraft command parameters, which can be deterministically tested and/or repeatably demonstrated).

[0027] Fourth, variations of this technology can include an approach necessarily rooted in computer technology for overcoming a problem specifically arising in the realm of computer networks. In an example, the technology can automatically translate audio into a computer readable format which can be interpreted by an aircraft processor. In an example, the technology can enable control of a partially and/or fully autonomous system based on communications with ATC operators. In such examples, the system/method may act in place of an incapacitated pilot (e.g., for a manned aircraft) and/or replace an onboard pilot (e.g., for an unmanned aircraft).

[0028] Fifth, variations of this technology can enable high speed and/or high accuracy natural language processing (NLP) of air traffic control (ATC) utterances by leveraging neural network models that were pre-trained on other datasets (e.g., pretrained models), then tuned to ATC-specific semantics. These ATC-tuned models can improve the speed/accuracy of the system in the context of noisy, multi-speaker ATC channels. These ATC-tuned models can also retain the broad ‘common sense’ comprehension of the pre-existing model and avoid overly biasing the system towards conventional ATC language – thus enabling the system to effectively respond to edge case scenarios or speech patterns which infrequently occur in ATC communications.

[0029] However, variations of the technology can additionally or alternately provide any other suitable benefits and/or advantages.

3. System.

[0030] The system 100, an example of which is shown in FIGURE 1, can include: a Speech-to-Text module 120 and a question-and-answer (Q/A) module 130 (e.g., cooperatively the “semantic parser”). The system can optionally include a communication

subsystem 110 and a flight processing system 140. However, the system 100 can additionally or alternatively include any other suitable set of components. The system 100 functions to determine flight commands 106 from an audio input 102 (e.g., received ATC radio transmission) which can be used for vehicle guidance, navigation, and/or control.

[0031] The audio input 102 can include a unitary utterance (e.g., sentence), multiple utterances (e.g., over a predetermined window – such as 30 seconds, within a continuous audio stream, over a rolling window), periods of silence, a continuous audio stream (e.g., on a particular radio channel, such as based on a current aircraft location or dedicated ATC communication channel), and/or any other suitable audio input. In a first example, the audio input can be provided as a continuous stream. In a second example, a continuous ATC radiofrequency stream can be stored locally, and a rolling window of a particular duration (e.g., last 30 seconds, dynamic window sized based on previous utterance detections, etc.) can be analyzed from the continuous radiofrequency stream.

[0032] The audio input is preferably in the form of a digital signal (e.g., radio transmission passed through an A/D converter and/or a wireless communication chipset), however can be in any suitable data format. In a specific example, the audio input is a radio stream from an ATC station in a digital format. In variants, the system can directly receive radio communications from an ATC tower and translate the communications into commands which can be interpreted by a flight processing system. In a first ‘human in the loop’ example, a user (e.g., pilot in command, unskilled operator, remote moderator, etc.) can confirm and/or validate the commands before they are sent to and/or executed by the flight processing system. In a second ‘autonomous’ example, commands can be sent to and/or executed by the flight processing system without direct involvement of a human. However, the system 100 can otherwise suitably determine commands from an audio input.

[0033] The system 100 is preferably mounted to, installed on, integrated into, and/or configured to operate with any suitable vehicle (e.g., the system can include the vehicle). The system 100 is preferably specific to the vehicle (e.g., the modules are specifically trained for the vehicle, the module is trained on a vehicle-specific dataset), but can be generic across multiple vehicles. The vehicle is preferably an aircraft (e.g.,

cargo aircraft, autonomous aircraft, passenger aircraft, manually piloted aircraft, manned aircraft, unmanned aircraft, etc.), but can alternately be a watercraft, land-based vehicle, spacecraft, and/or any other suitable vehicle. In a specific example, the aircraft can include exactly one pilot/PIC, where the system can function as a backup or failsafe in the event the sole pilot/PIC becomes incapacitated (e.g., an autonomous co-pilot, enabling remote validation of aircraft control, etc.).

[0034] The system 100 can include any suitable data processors and/or processing modules. Data processing for the various system and/or method elements preferably occurs locally onboard the aircraft, but can additionally or alternatively be distributed among remote processing systems (e.g., for primary and/or redundant processing operations), such as at a remote validation site, at an ATC data center, on a cloud computing system, and/or at any other suitable location. Data processing for the Speech-to-Text module and Q/A module can be centralized or distributed. In a specific example, the data processing for the Speech-to-Text module and the Q/A module can occur at a separate processing system from the flight processing system (e.g., are not performed by the FMS or FCS processing systems; the Speech-to-Text module and Q/A module can be decoupled from the FMS/FCS processing; an example is shown in FIGURE 12), but can additionally or alternatively be occur at the same compute node and/or within the same (certified) aircraft system. Data processing can be executed at redundant endpoints (e.g., redundant onboard/aircraft endpoints), or can be unitary for various instances of system/method. In a first variant, the system can include a first natural language processing (NLP) system, which includes the Speech-to-Text module and the Q/A module, which can be used with a second flight processing system, which includes the flight processing system and/or communication systems (e.g., ATC radio). In a second variant, an aircraft can include a unified 'onboard' processing system for all runtime/inference processing operations. In a third variant, remote (e.g., cloud) processing can be utilized for Speech-to-Text operations and/or Q/A response generation. However, the system 100 can include any other suitable data processing systems/operations.

[0035] The system 100 can optionally include a communication subsystem, which functions to transform an ATC communication (e.g., radio signal) into an audio input which can be processed by the ASR module. Additionally or alternately, the communication subsystem can be configured to communicate a response to ATC. The communication subsystem can include an antenna, radio receiver (e.g., ATC radio receiver), a radio transmitter, an A/D converter, filters, amplifiers, mixers, modulators/demodulators, detectors, a wireless (radiofrequency) communication chipset, and/or any other suitable components. The communication subsystem include: an ATC radio, cellular communications device, VHF/UHF radio, and/or any other suitable communication devices. In a specific example, the communication subsystem is configured to execute S210. However, the communication subsystem can include any other suitable components, and/or otherwise suitably establish communication with air traffic control (ATC).

[0036] The Speech-to-Text module of the system 100 functions to convert the audio input (e.g., ATC radio signal) into an utterance hypothesis 104, such as in the form of text (e.g., an ATC transcript) and/or alphanumeric characters. The utterance hypothesis is preferably a text stream (e.g., dynamic transcript), but can alternatively be a text document (e.g., static transcript), a string of alphanumeric characters (e.g., ASCII characters), or have any other suitable human-readable and/or machine-readable format. The Speech-to-Text module is preferably onboard the aircraft, but can additionally or alternatively be remote. The Speech-to-Text module is preferably an ATC-tuned Speech-to-Text module, which includes one or more models pre-trained on ATC audio data, but can additionally or alternatively include one or more generic models/networks and/or models/networks pre-trained on generalized training data (e.g., natural language utterances not associated with ATC communication).

[0037] The Speech-to-Text module can include: an integrated automatic speech recognition (ASR) module 122, a sentence boundary detection (SBD) module 124, a language module 126, and/or other modules, and/or combinations thereof. In a specific example, the Speech-to-Text module can include an integrated ASR/SBD module 125. The Speech-to-Text module (and/or submodules thereof) can include a neural network

(e.g., DNN, CNN, RNN, etc.), a cascade of neural networks, compositional networks, Bayesian networks, Markov chains, pre-determined rules, probability distributions, attention-based models, heuristics, probabilistic graphical models, or other models. The Speech-to-Text module (and/or submodules thereof) can be tuned versions of pretrained models (e.g., pretrained for another domain or use case, using different training data), be trained versions of previously untrained models, and/or be otherwise constructed.

[0038] In variants, a submodule(s) of the Speech-to-Text module (e.g., ASR module and/or SBD module) can ingest the audio input (e.g., audio stream, audio clip) and generate a set of linguistic hypotheses (e.g., weighted or unweighted), which can serve as an intermediate data format, such as may be used to audit the Speech-to-Text module, audit sub-modules/models therein, and/or select a unitary utterance hypothesis. The set of linguistic hypotheses can include overlapping/alternative hypotheses for segments of audio, or can be unitary (e.g., a single hypothesis for an individual audio segment or time period). The set of linguistic hypotheses can include: utterance hypotheses (e.g., utterance hypothesis candidates), letters, word-segment streams, phonemes, words, sentence segments (e.g., text format), word sequences (e.g., phrases), sentences, speaker changes, utterance breaks (e.g., starts, stops, etc.), and/or any other suitable hypotheses. In variants where the audio stream includes multiple speakers/utterances, the set of linguistic hypotheses can additionally include an utterance boundary hypothesis which can distinguish multiple speakers and/or identify the initiation and termination of an utterance, with an associated weight and/or a speaker hypothesis (e.g., tag identifying a particular speaker, tag identifying a particular aircraft/tower). Additionally or alternately, the utterance boundary hypothesis can identify utterance boundaries and/or change in speaker without identifying individual speaker(s). Each linguistic hypothesis preferably includes an associated weight/score associated with an utterance (and/or utterance boundary), assigned according to a relative confidence (e.g., statistical; such as determined using an ASR model, SBD model, and/or language model; etc.). The set of linguistic hypotheses is preferably ordered, sequential, and/or time-stamped in association with the receipt time, but can be otherwise suitably related.

[0039] However, the Speech-to-Text module can generate, store, and/or output any other suitable set of hypotheses. As an example, the linguistic hypotheses can include a plurality of utterance hypotheses, wherein a single utterance hypothesis can be selected based on the set of generated set of utterance hypotheses. As a second example, a subset (e.g., complete set) of linguistic hypotheses, with a corresponding weight/score, can be output by the Speech-to-Text module.

[0040] The Speech-to-Text module can include an ASR module which functions to extract linguistic hypotheses from the audio input. Using the audio input, the ASR module can determine a sequence of linguistic hypotheses, such as: letters, word-segment streams, phonemes, words, sentence segments (e.g., text format), word sequences (e.g., phrases), sentences, and/or any other suitable linguistic hypotheses (e.g., with a corresponding weight). The ASR module is preferably a neural network (e.g., Wav2Letter, Kaldi, Botium, etc.), but can alternatively be any other suitable model. In an example, a pretrained neural network can be tuned for ATC audio and/or trained using ATC audio (e.g., with an associated transcript). In a second example, the ASR module can include the ASR model trained by S110 and/or S120. In a specific example, the ASR module is configured to execute S220 of the method. The ASR module can optionally include an integrated SBD module. In variants where the ASR module outputs lower-level linguistic components (e.g., phonemes, phonetics, etc.), the system can optionally include auxiliary transformation modules (e.g., phoneme-to-word transformations) that convert the lower-level linguistic components to linguistic components compatible by the language module and/or other system modules.

[0041] The Speech-to-Text module can include an SBD module which functions to identify utterance boundaries and/or speaker changes for a multi-utterance audio inputs. Using the audio input, the SBD module can determine a sequence of linguistic hypotheses, such as: an utterance boundary hypothesis, a speaker hypothesis (e.g., tag identifying a particular speaker, tag identifying a particular aircraft/tower), and/or any other suitable hypotheses. The SBD module is preferably integrated with the ASR module (an example is shown in FIGURE 10A), but can otherwise be separate from the ASR module, such as operating sequentially with the ASR module (e.g., passing a single utterance input into

the ASR module, tagging outputs of the ASR module, etc.; examples are shown in FIGURES 10C-D) or in parallel with the ASR module (e.g., separately providing speaker change and/or utterance boundary annotations by way of time stamps, etc.; an example is shown in FIGURE 10B). The SBD module is preferably a neural network (e.g., Wav2Letter, Kaldi, Botium, etc.), but can alternatively be any other suitable model. In an example, a pretrained SBD neural network can be tuned for ATC audio and/or trained using ATC audio (e.g., with an associated transcript). In a second example, an SBD neural network can be trained separately from the ASR module (e.g., using a distinct training set, using a training set including periods of radio silence and/or audio artifacts, etc.). In a third example, the SBD model can be tuned for ATC audio and/or trained using ATC audio, such as trained to identify silence speakers and/or utterance boundary characters (e.g., transition speakers, transition audio artifacts). However, the Speech-to-Text module can include any other suitable SBD module(s).

[0042] The language module of the Speech-to-Text module functions to select an utterance hypothesis based on the set of linguistic hypotheses, which can then be passed into the Q/A module. The language module receives the set of linguistic hypotheses from the ASR module (e.g., phonemes, words, sentence subsets, etc.) and returns an utterance hypothesis associated with a single utterance (e.g., a sentence, a series of linguistic hypothesis, etc.). The language module preferably determines the utterance hypothesis purely from the linguistic hypotheses, but can alternatively or additionally ingest the audio input and/or other auxiliary data. Auxiliary data can include: an aircraft ID, contextual information (e.g., vehicle state, geographical position, ATC control tower ID and/or location, etc.), weather data, and/or any other suitable information. The utterance hypothesis is preferably text (e.g., a text string or utterance transcript), but can alternatively be a set of phoneme indexes, audio, or any suitable data format.

[0043] The language module preferably selects an utterance hypothesis from the set of linguistic hypotheses by weighting the likelihood of various 'sound-based' language interpretations in the context of the entire utterance and/or ATC language patterns. In a first variant, the language module assigns language weights/scores to each utterance hypothesis using a neural network language model (e.g., an LSTM network, a CNN,

FairSeq ConvLM, etc.) tuned for ATC language (e.g., neural network trained using ATC transcripts, etc.; such as a language model trained according to S140). In a second variant, the language module assigns language weights/scores according to a grammar-based language model (e.g., according to a set of heuristics, grammar rules, etc.). In a third variant, the language module can be tightly integrated with the ASR module. In examples, a language model(s) can be used during the search, during the first pass, and/or during reranking. However, the language module can assign weights / scores in any other suitable manner. In a specific example, the language module is configured to execute S230 of the method.

[0044] In an example, the Speech-to-Text module transforms an ATC audio stream into a natural language text transcript which is provided to the Q/A module, preserving the syntax as conveyed by the ATC speaker (e.g., arbitrary, inconsistent, non-uniform syntax).

[0045] Alternatively, the speech-to-text module can include a neural network trained (e.g., using audio data labeled with an audio transcript) to output utterance hypotheses (e.g., one or more series of linguistic components separated by utterance boundaries) based on an audio input. However, the speech-to-text module can include: only an automated speech recognition module, only a language module, and/or be otherwise constructed.

[0046] However, the system can include any other suitable Speech-to-Text module.

[0047] The system 100 can include a question-and-answer (Q/A) module (example shown in FIGURE 7), which functions to determine a set of commands from the selected hypothesis (e.g., text transcript) using a set of flight command queries. The Q/A module preferably receives an utterance hypothesis from the Speech-to-Text module in text, but can alternately receive audio and/or any other suitable inputs.

[0048] The Q/A module preferably includes one or more Q/A models (e.g., BERT, BERT tuned to ATC applications, etc.), but can additionally or alternatively include a classifier or other model. The Q/A model is preferably a pre-trained language model tuned for ATC transcripts but can be untrained or have another format. The Q/A model can be: a convolutional neural network, a (pre-trained) large neural language model,

bidirectional encoder representations from transformers (BERT), generative pre-trained transformer (GPT), and/or any other suitable language model. However, the Q/A module can include any other suitable neural language models.

[0049] The Q/A module preferably answers a set of flight command queries (e.g., natural language queries). The flight command queries are preferably predetermined (e.g., manually determined, extracted from a command template, etc.), but can be dynamically determined. Flight command queries are preferably semantic queries in a human-readable format, but can additionally or alternatively be provided in a machine-readable format. The command queries are preferably natural language (“reading comprehension”), but can alternatively be vectors, tensors, and/or have another format. The set of flight command queries is preferably organized in a hierarchical structure (e.g., with parent-child query relationships), but can alternatively be organized in a serial structure, or be otherwise organized. The flight command queries can be organized in a list, a tree, or otherwise organized. In variants, flight command queries can be provided as a sequence/series of chained nodes (examples are shown in FIGURES 11A-C), each node corresponding to a predetermined query, wherein the nodes include a set of independent nodes and a set of dependent nodes, each dependent node linked to a specific answer/response (e.g., specific answer value) of a broader/higher-level parent semantic query (e.g., where queries have a finite set of answers or a closed range of answers). Accordingly, dependent queries may be triggered in response to a determination of a predetermined answer at a higher-level linked node. Alternatively, the set of predetermined flight command queries can be provided synchronously or asynchronously in any suitable combination/permutation of series and/or parallel.

[0050] The command queries can be configured to have binary answers (e.g., “yes”, “no”, discrete answers (e.g., letters, integers, etc.), continuous answers (e.g., coordinate values, etc.), and/or any other suitable type of answer value. Different types of commands can have different query structures. For example, high-criticality queries, such as aircraft identifiers, can be structured as binary queries. In another example, attributes with multiple potential answers can be structured as open-ended questions (e.g., “topics?”) instead of binary questions (e.g., “Does the utterance include heading?” Does the

utterance include altitude?”). However, the queries can be otherwise structured. Examples of command queries include: whether the aircraft is the intended recipient of an utterance hypothesis, what or whether command parameters or topics (e.g., heading, altitude, etc.) are included in the utterance hypothesis, what or whether command parameter values (e.g., altitude direction, altitude level, etc.) are included in the utterance hypothesis, and/or other queries. In a first example, the Q/A module determines that the utterance is intended for the aircraft (e.g., Question: “Intended for DAL456?”; Answer: “yes”). In a second example, the Q/A module determines the topics of an utterance (e.g., Question: “Topics?”; Answer: “Heading, Altitude”). In a third example, the Q/A determines the values associated with a topic of the utterance (e.g., Question: “Altitude values?”; Answer: “Direction: down, Level: 2000”). In an example, the Q/A module can be configured to execute S240.

[0051] Based on the queries, the Q/A module outputs a set of flight commands, which can include guidance commands (e.g., navigational instructions; sequences of waypoints, approach landing site, etc.), vehicle state commands (e.g., instructions to modify vehicle state parameters, increase altitude to 5000 ft, etc.), effector state commands (e.g., effector instructions; deploy landing gear, etc.), flightpath commands (e.g., trajectory between waypoints, etc.), and/or any other suitable commands. The commands preferably output in a prescribed format based on the answers generated by the Q/A module, such as a standardized human-readable format (e.g., allowing human validation) and/or a machine-readable format (e.g., allowing human interpretation/validation of the commands). In a specific example, the commands can be provided as the union of the answers to the command parameter identification query and at least one command parameter value query (e.g., corresponding to the answer of the command parameter identification query). In a second example, the commands can be directly taken as a combination of each answer/response as generated by the Q/A module. Output commands are preferably text based and/or alphanumeric, but can be otherwise suitably provided (e.g., text-to-speech validation, etc.). In some variants, the commands can be post-processed according to any suitable heuristics, grammar rules, or formatting protocols, but can otherwise be provided to a pilot and/or flight processing system

directly as the output of the Q/A module. In a specific example, the Q/A module can convert an utterance hypothesis into a command in a standardized data format (e.g., as may be accepted/interpreted by a certified aircraft processor). In variants, the commands can include a substantially finite set of command parameters (e.g., altitude, heading, etc.) corresponding to a predetermined set of topics. Additionally, command parameters can be within substantially finite and/or bounded ranges (e.g., heading limited to compass directions, altitude limited by physical aircraft constraints, commands cooperatively limited by flight envelope, etc.). However, command parameters can additionally or alternatively be arbitrary, unbounded, and/or substantially unconstrained. However, the Q/A module can generate any other suitable commands.

[0052] However, the system can include any other suitable Q/A module.

[0053] The system 100 can optionally include and/or be used with a flight processing system, which functions to control various effectors of the aircraft according to the commands. The flight processing system can include an aircraft flight management system (FMS), a flight control system (FCS), flight guidance/navigation systems, and/or any other suitable processors and/or control systems. The flight processing system can control flight effectors/actuators during normal operation of the vehicle, takeoff, landing, and/or sustained flight. Alternatively, the flight processing system can be configured to implement conventional manual flight controls in a flight-assistive configuration. The system can include a single flight processing system, multiple (e.g., three) redundant flight processing systems, and/or any other suitable number of flight processing systems. The flight processing system(s) can be located onboard the aircraft, distributed between the aircraft and a remote system, remote from the aircraft, and/or otherwise suitably distributed. In a specific example, the flight processing system is configured to execute S250.

[0054] In variants, the flight processing system can be configured (e.g., certified) to accept only a predetermined set of command input and/or inputs having a predetermined format, where the outputs of the Q/A model are provided in the predetermined format and/or are a subset of the predetermined set of commands.

[0055] However, the system can include any other suitable components and/or be otherwise suitably configured to execute S200 of the method.

4. Method.

[0056] The method, an example of which is shown in FIGURE 2, can optionally include training the system components S100; and performing inference using the system S200. The method functions to automatically interpret flight commands from a stream of air traffic control (ATC) radio communications. The method can additionally or alternatively function to train and/or update a natural language processing system based on ATC communications.

4.1 Training

[0057] Training the system components S100 (example shown in FIGURE 9) functions to generate an ATC-tuned system capable of interpreting ATC audio signals into flight commands. S100 can include training a Speech-to-Text model and training a question-and-answer (Q/A) model S150. S100 can optionally include generating augmented ATC transcripts S130. However, training the semantic parser S100 can include any other suitable elements. S100 is preferably performed offline and/or by a remote computing system, but can alternatively be performed onboard the aircraft (e.g., locally, during flight, asynchronously with aircraft flight).

[0058] Training the Speech-to-Text model functions to generate a transcription model that is specific to ATC communications, accounting for ATC-specific grammar, lexicon, speech patterns, and other idiosyncrasies. Training the Speech-to-Text model can include training an ASR model S110, training an SBD model S120, training a language model S140, and/or any other suitable elements. Training can include: tuning the network weights, determining weights de novo, and/or otherwise training the network. Training (and/or inference) can leverage: gradient-based methods (e.g., stochastic gradient descent), belief propagation (e.g., sum-product message passing; max product message passing, etc.), and/or any other suitable training method.

[0059] Training an automatic speech recognition (ASR) module S110 functions to train a neural network to recognize natural language in ATC communications. The ASR model is preferably trained (e.g., using supervised training, semi-supervised training)

from a pre-existing ASR model (e.g., Wav2Letter), and can be ‘tuned’ by providing the neural network a mix (e.g., 50/50, 60/40, 70/30, pre-determined mix, 100/0, etc.) of ATC training audio with corresponding ATC transcripts and the original training data (e.g., from the pre-existing model). An example is shown in FIGURE 4. The ATC training audio with transcripts is preferably manually determined (e.g., by a human, by a domain expert), but can be verified/audited ATC communication audio/transcripts (e.g., generated from an existing ASR model), and/or otherwise determined. The ATC training audio can include a single utterance, multiple utterances, a stream of radio communication over an ATC communications channel, and/or any other suitable training audio. Preferably, utterances (e.g., statements from an individual speaker, sentences, etc.) are individually associated with a transcript as part of the training data. However, the ASR model can be otherwise trained for ATC speech recognition.

[0060] Training a sentence boundary detection (SBD) module S120 functions to train the Speech-to-Text module to identify utterance boundaries (e.g., sentence segment boundaries, sentence boundaries). S120 can optionally train the Speech-to-Text module to differentiate unique utterances and/or utterances from different speakers/entities. S120 can train an existing ASR model (e.g., as determined in S110, which generates an integrated ASR/SBD model) or a separate model to generate the SBD module. Preferably, the SBD model can be trained using time-length concatenated audio, which includes a series of multiple utterances and periods of silence (e.g., periods of no speaking) therebetween, and the associated multi-utterance training transcripts. The ATC audio and transcripts used to train the SBD model can be the same as the ASR model and/or different from the ASR model.

[0061] Multi-utterance training transcripts preferably include boundary annotations (e.g., with a unique boundary character or other identifier; using a ‘/’ or ‘%’ character; etc.) which can delineate unique speakers, unique utterances, breaks between utterances, periods of silence, audio artifacts (e.g., the “squelch” when the ATC speaker starts and/or starts broadcasting), and/or any other appropriate boundaries. Boundary annotations are preferably automatically added during transcript concatenation, but can be inserted manually, be determined from the audio, and/or otherwise added.

[0062] In a specific example, the ASR model is trained by assigning a unique ‘silence speaker’ and/or a unique ‘transition speaker’ in the audio and/or transcript – which can be particularly advantageous in SBD for ATC radio communications, commonly exhibit a characteristic radio “squelch” sound prior to an utterance. By assigning these segments of audio to a unique ‘transition speaker’ (or a ‘squelch speaker’) the SBD model can more accurately differentiate between back-to-back utterances (e.g., with minimal intervening silence), which commonly occurs in noisy ATC radio channels.

[0063] However, an SBD model can be otherwise trained.

[0064] Training a language model S140 functions to train a language model to distinguish ATC linguistic patterns. In variants, the language model can determine whether a transcript is contextually correct/logical (e.g., syntactically correct, based on ATC grammar, etc.), determine a language/syntax score for a transcript, and/or otherwise determine whether a transcript makes sense. Preferably, S140 tunes a pre-existing language model (e.g., convolutional neural network, FairSeq ConvLM, etc.), but can alternately train an untrained language model. An existing language model can be tuned based on ATC transcripts, which can be single utterance ATC transcripts, multi-utterance ATC transcripts, and/or boundary annotated ATC transcripts (e.g., such as those used to train the SBD model in S120), however the language model can be trained using any suitable ATC transcripts. S140 preferably does not train on the ATC audio, but can alternatively train on the ATC audio. In variants, the language model can be trained using entity-tagged ATC transcripts, which identify ATC specific entities within the transcript. Tagged entities can include: carriers, aircraft, waypoints, airports, numbers, directions, and/or any other suitable entities. Entity tags can be assigned manually, automatically (e.g., unsupervised), with a semi-supervised HMM tagger (e.g., using a domain expert evaluation tool, etc.), and/or in any other suitable manner. A single word or phrase appearing in a transcript can be assigned to multiple entities depending on the context in which it appears (i.e., the entity tag lexicon can include multiple phonetically and/or lexicographically conflicting entities which are pronounced and/or spelled substantially identically). In an example, “Southwest” can be tagged as (and/or communicate) a direction or a carrier depending on the context in which it appears. Likewise, in a second

example, “delta” can be tagged as part of an aircraft name (e.g., DAL456 = “delta alpha lima four five six”), a carrier, and/or untagged (e.g., referring to a change in value or parameter) depending on the context in which it appears. In a third example, “Lima” can be an airport, a waypoint, part of an aircraft name, and/or otherwise tagged. In a fourth example, waypoints can be pronounced substantially identically (e.g., “ocean”) while corresponding to different waypoint entities depending on the context in which they appear. However, the language model can be trained with any other suitable transcripts and/or information.

[0065] In variants, a portion of the training text provided to train the language model is the same as that used to originally train the pre-existing language model (e.g., FairSeq ConvLM). Accordingly, the language model can be ‘tuned’ by providing the neural network a mix (e.g., 50/50, 60/40, 70/30, pre-determined mix, etc.) of ATC training transcripts and the original training data (e.g., from the pre-existing model). However, a language model can be otherwise trained for ATC linguistic patterns.

[0066] S100 can optionally include generating augmented ATC transcripts S130 (e.g., synthetic transcripts), which functions to expand the number/quantity of ATC training transcripts available to train the language model in S140, an example of which is shown in FIGURE 5. In variants, this can be beneficial in order to provide training transcripts specific to areas/regions where entities are known (e.g., airport names, waypoints, carriers, etc.), but from which ATC transcripts are unavailable. Additionally or alternately, S130 can improve the accuracy of the language model by increasing a size of the training dataset (e.g., number of available utterance transcripts). S130 preferably substitutes the values of tagged entities (e.g., within the entity-tagged ATC transcripts) with different entity values from an ATC entity lexicon. The ATC entity lexicon can be manually generated, generated by a domain expert (e.g., pilot), randomly generated (e.g., number substitution), generated using: historical flight logs, aircraft databases, airport databases, randomly generated, and/or otherwise generated. In variants, the augmented ATC transcripts can preferentially (e.g., at a higher rate; with greater frequency; occurring with greater than a threshold number of instances – such as 3 or more within the training set) substitute phonetically and/or lexicographically conflicting entity names (e.g., which

are identified by multiple tags in different contexts), such as “southwest” and “delta.” The augmented ATC transcripts can then be used to train the language model in S140 and/or question-and-answer model in S150 (e.g., an example of training an ATC-tuned language model is shown in FIGURE 5).

[0067] However, ATC transcripts can be otherwise generated. Alternatively, the system (and/or neural network models therein) can be trained entirely with real ATC communication transcripts.

[0068] S100 can include training a question-and-answer (Q/A) module S150, which functions to train a model to answer ATC-specific queries. S150 preferably includes tuning a pre-trained language model, but can include training an untrained model. The language model can be trained using: an ATC transcript, the associated parsed meaning (e.g., reference outputs; answers to the queries; values for command parameters determined from the ATC transcript, etc.), the set of command queries, and/or other data. In variants, S150 can also provide the language model contextual information pertaining to a particular utterance— such as a tail number or carrier for a particular aircraft, a flight plan for the aircraft, a set of utterance transcripts preceding the particular utterance, and/or any other suitable contextual information.

[0069] The text transcripts used to train the Q/A model can be the same ATC transcripts used to train the ASR and/or SBD model, the same ATC transcripts (and/or augmented ATC transcripts) used to train the language model, the utterance hypotheses output by the Speech-to-Text module, and/or other transcripts. However, the Q/A model can be trained using any suitable ATC transcripts.

[0070] The parsed meaning used to train the Q/A model can be: manually determined, manually audited by a domain expert, provided by a grammatical semantic parser (e.g., SEMPRE, a lower-accuracy parser than the system, a previous iteration of the system, etc.; an example is shown in FIGURE 6) referencing ATC grammar (e.g., manually determined, iteratively determined, learned, etc.), and/or otherwise suitably determined.

[0071] In a specific example, a grammatical semantic parser parses the command parameter values from the ATC transcripts, wherein the parsed values (e.g., command

hypotheses), source transcript, optionally ATC audio, and/or other data are presented on a domain evaluation tool (an example is shown in FIGURE 8) to domain experts. The domain expert can: label to the model output (e.g., as “correct,” “incomplete,” “incorrect,” etc.), correct the parsed values, and/or otherwise interact with the parser output. In variants, reference outputs labelled as “incorrect” and/or “incomplete” can be reviewed and used to update or improve grammar rules of a grammatical semantic parser. In variants, reference outputs labelled “incorrect” are not used to train the Q/A model, but can alternately be used to train the Q/A model (e.g., the “incorrect” label serving to train by counterexample). In variants, reference outputs which are labelled as “correct” and/or “incomplete” can be passed into the Q/A model during S150. In variants, incomplete label data can be used to train a subset of queries associated with a particular utterance (e.g., based on the correctly labelled portions of the transcript). As an example, where the parameter values may be unlabelled and the topics are identified, the topics may be used to train a command identification (e.g., “topics?”) query. Likewise, where the aircraft tail number is tagged/identified, incomplete label data can be used to train the plane-specific speaker identification query(ies). However, the labels can be otherwise used, and model outputs can be otherwise suitably determined.

[0072] However, a question-and-answer model can be otherwise suitably trained.

[0073] In variants, the ASR model, SBD model, language model, and/or Q/A model can be optionally retrained and/or updated based on pilot/PIC validation with any suitable update frequency. The models can be updated/retrained independently, synchronously, asynchronously, periodically (e.g., with a common update frequency, with different frequencies), never (e.g., which may be desirable in instances where the deterministic model(s) are certified), based on auditing of the intermediate outputs, and/or can be otherwise suitably updated or trained. The models can be updated locally, onboard the aircraft, periodically via remote/cloud (push) updates, and/or can be otherwise suitably updated/retrained.

[0074] In variants, the model(s) can be audited based on a pilot rejection of the final output parameters in order to locate error origin(s) within the data pipeline (e.g., as part of a root cause analysis), which can be used as a training input to improve the

network. As an example: an erroneous intermediate parameter (such as in the utterance hypothesis or linguistic hypothesis) can result in an incorrect output of the Q/A module even in cases where the Q/A module performs correctly. In variants, the outputs of each model/module can additionally be audited against a formatting template prescribed to each step (e.g., to enable certification compliance of the system). However, the system and/or various subcomponents can be otherwise suitably audited.

[0075] However, the system components can be otherwise suitable trained.

4.2 Runtime/Inference

[0076] S200 can include: at an aircraft, receiving an audio utterance from air traffic control S210, converting the audio utterance into a predetermined format S215, determining commands using a question-and-answer model S240, and controlling the aircraft based on the commands S250. However, the method S200 can additionally or alternatively include any other suitable elements. S200 functions to automatically interpret flight commands from the air traffic control (ATC) stream. The flight commands can be automatically used to control aircraft flight; presented to a user (e.g., pilot, a remote teleoperator); relayed to an auto-pilot system in response to a user (e.g., pilot) confirmation; and/or otherwise used.

[0077] All or portions of S200 can be performed continuously, periodically, sporadically, in response to transmission of a radio receipt, during aircraft flight, in preparation for and/or following flight, at all times, and/or with any other timing. S200 can be performed in real- or near-real time, or asynchronously with aircraft flight or audio utterance receipt. S200 is preferably performed onboard the aircraft, but can alternatively be partially or entirely performed remotely.

[0078] Receiving an audio utterance from air traffic control S210 functions to receive a communication signal at the aircraft and/or convert the communication signal into an audio input, which can be processed by the ASR module. In a specific example, S210 transforms an analog radio signal into a digital signal using an A/D converter (and/or other suitable wireless communication chipset), and sends the digital signal to the ASR module (e.g., via a wired connection) as the audio input. S210 preferably monitors a single radio channel (e.g., associated with the particular aircraft), but can

alternately sweep multiple channels (e.g., to gather larger amounts of ATC audio data). However, S210 can otherwise suitably receive an utterance.

[0079] Converting the audio utterance into a predetermined format S215 functions to generate a transcript from the ATC audio. This can be performed by the Speech-to-Text module or other system component. Converting the audio utterance to into a predetermined (e.g., text) format can include: determining a set of utterance hypotheses for an utterance S220 and selecting an utterance hypothesis from the set of utterance hypotheses S230; however, the ATC audio can be otherwise converted.

[0080] Determining a set of utterance hypotheses for an utterance S220 functions to identify audio patterns (e.g., such as letters, phonemes, words, short phrases, etc.) within the utterance. In a specific example, S220 can be performed by the Speech-to-Text module, an ASR module (and/or ASR model therein), an integrated ASR/SBD module (e.g., with an integrated ASR/SBD model therein), a language module, and/or combinations thereof. S220 can optionally include assigning a weight or score to each audio pattern (a.k.a. linguistic hypothesis) using the ASR module and/or other modules. An utterance hypothesis can be: a linguistic hypothesis, a series of linguistic hypotheses, and/or any other suitable hypothesis.

[0081] In a first variation, an ASR and/or integrated SBD/ASR module generates a set of linguistic hypotheses, wherein a language module receives the linguistic hypotheses and generates a score (e.g., ASR score; same or different from language weight/score) for each string or sequence of linguistic hypotheses. One or more linguistic hypothesis sets can be generated from the same audio clip. The SBD/ASR module can also output a score (ASR score or ASR weight) for each linguistic hypothesis, sequence of hypotheses, and/or set of linguistic hypotheses. However, the set of utterance hypotheses can be otherwise determined.

[0082] Selecting an utterance hypothesis from the set of utterance hypotheses S230 functions to detect language patterns from the set of linguistic hypotheses in the context of the entire utterance. Additionally or alternately, S230 can function to select the highest probability string/sequence of linguistic hypotheses as the utterance hypothesis. S230 can be performed by the language module, the Q/A module, and/or another module.

[0083] In a first variation, the language module can select the string or sequence of linguistic hypotheses which has the highest combined language weight (or score) and ASR weight (or score) as the utterance hypothesis.

[0084] In a second variation, multiple modules' outputs are cooperatively used to select the utterance hypothesis. For example, the utterance hypothesis with the highest combined hypothesis score and/or maximum hypothesis weight cooperatively determined by the language model and the integrated ASR/SBD model is selected. In a first example, the utterance hypothesis which maximizes the language weight multiplied by the ASR weight for an utterance is selected. In a second example, the hypothesis which maximizes the sum of the language score and the ASR score for an utterance.

[0085] However, the utterance hypothesis can be otherwise selected.

[0086] Determining commands from the utterance hypothesis using a question-and-answer model S240 functions to extract flight commands from the utterance hypothesis, which can be interpreted and/or implemented by a flight processing system. S240 is preferably performed by one or more instances of the Q/A module, but can be performed by another component. S240 is preferably performed using the set of flight command queries and the utterance hypothesis, but can be otherwise performed.

[0087] S240 can include providing the Q/A module with a set of command queries in addition to the utterance hypothesis as an input, wherein the Q/A module answers the command queries using the utterance hypothesis as a reference text. In a first embodiment, the queries are provided serially, wherein the successive query is determined based on the prior answer. The query series can be determined from the command query set structure (e.g., list, tree, etc.), randomly determined, or otherwise determined. In a specific example, S240 includes querying for topic presence within the utterance hypothesis, then only querying for values for the topics confirmed to be within the utterance. In a second specific example, S240 includes initially determines if the aircraft (and/or pilot) is the intended recipient of the utterance (associated with the utterance hypothesis), and only querying further if the utterances are intended for the aircraft/pilot (e.g., utterances not intended for the aircraft/pilot are ignored and/or any commands therein are not passed to the flight processing system; utterances

corresponding to a transition speaker detections can be neglected; etc.). Alternatively, the Q/A model (or different versions or instances thereof) can be queried with multiple queries in parallel or can be otherwise queried.

[0088] In a second variant, the Q/A module includes pre-embedded queries, wherein the Q/A module answers a predetermined set of questions based on the utterance hypothesis. For example, the Q/A module can be a multi-class classifier that outputs values, determined from the utterance hypothesis, for each of a set of “classes,” wherein each class represents a command parameter. However, S240 can otherwise suitably determine command parameter values.

[0089] S200 can optionally include controlling the aircraft based on the commands S250, which functions to modify the aircraft state according to the utterance (e.g., ATC directives). In a specific example, S250 autonomously controls the effectors and/or propulsion systems of the aircraft according to the commands (e.g., to achieve the commanded values). In a second example, the flight processing system can change waypoints and/or autopilot inputs based on the commands. In variants, S200 can include providing the commands to a flight processing system (e.g., FCS) in a standardized format (e.g., a standardized machine-readable format).

[0090] However, S250 can otherwise suitably control the aircraft based on the commands. Alternatively, the system can be used entirely in an assistive capacity (e.g., without passing commands to an aircraft processor or controlling the aircraft, such as to enable control of an aircraft by a hearing-impaired pilot), and/or can be otherwise used.

[0091] However, S200 can include any other suitable elements.

[0092] Alternative embodiments implement the above methods and/or processing modules in non-transitory computer-readable media, storing computer-readable instructions. The instructions can be executed by computer-executable components integrated with the computer-readable medium and/or processing system. The computer-readable medium may include any suitable computer readable media such as RAMs, ROMs, flash memory, EEPROMs, optical devices (CD or DVD), hard drives, floppy drives, non-transitory computer readable media, or any suitable device. The computer-executable component can include a computing system and/or processing system (e.g.,

including one or more collocated or distributed, remote or local processors) connected to the non-transitory computer-readable medium, such as CPUs, GPUs, TPUS, microprocessors, or ASICs, but the instructions can alternatively or additionally be executed by any suitable dedicated hardware device.

[0093] Embodiments of the system and/or method can include every combination and permutation of the various system components and the various method processes, wherein one or more instances of the method and/or processes described herein can be performed asynchronously (e.g., sequentially), concurrently (e.g., in parallel), or in any other suitable order by and/or using one or more instances of the systems, elements, and/or entities described herein.

[0094] As a person skilled in the art will recognize from the previous detailed description and from the figures and claims, modifications and changes can be made to the preferred embodiments of the invention without departing from the scope of this invention defined in the following claims.

CLAIMS

We claim:

1. A system for an aircraft for semantic parsing of air traffic control (ATC) utterances, the system comprising:
 - a communication system onboard the aircraft, the communication subsystem configured to receive an ATC radio signal and transform the ATC radio signal into an audio signal;
 - a first processing system connected to the communication system, comprising:
 - a speech-to-text module configured to determine an utterance hypothesis from the audio signal; and
 - a question-and-answer (Q/A) module configured to determine aircraft commands based on the utterance hypothesis using a plurality of natural language queries, the aircraft commands determined by querying a pre-trained neural network model tuned to ATC audio according to a structured sequence of the natural language queries; and
 - a second processing system connected to the first processing system and configured to control the aircraft based on the aircraft commands, wherein the ATC utterances are syntactically non-standardized, wherein the first processing system is configured to automatically provide the commands to the second processing system in a standardized format.
2. The system of Claim 1, wherein the first processing system is configured to convert the syntactically non-standardized utterances into the standardized format.
3. The system of Claim 1, wherein the structured sequence comprises a tree-based sequence with a plurality of dependencies linking one or more natural language queries to a determination that the aircraft is an intended recipient of an utterance corresponding to the utterance hypothesis.
4. The system of Claim 3, wherein the utterance hypothesis comprises a speaker identification, wherein the determination that the aircraft is the intended recipient is based on the speaker identification.
5. The system of Claim 1, wherein each command comprises a command parameter and a set of values corresponding to the command parameter, wherein the command parameter is selected from a predetermined set of command parameters,

wherein the set of values and the command parameter are determined via distinct natural language queries of the structured sequence.

6. The system of Claim 5, wherein the structured sequence comprises a tree-based sequence, wherein the values corresponding to the command parameter are queried in response to a determination that the command parameter is a topic of the utterance.

7. The system of Claim 1, wherein the plurality of natural language queries is selected from a predetermined set.

8. A system for an aircraft for semantic parsing of air traffic control (ATC) utterances, the system comprising:

- a communication system onboard the aircraft, the communication subsystem configured to receive an ATC radio signal and transform the ATC radio signal into an audio signal; and
- a first processing system connected to the communication system, comprising:
 - a speech-to-text module configured to determine an utterance hypothesis from the audio signal, the speech-to-text module comprising an integrated automatic speech recognition (ASR) and sentence boundary detection (SBD) module, wherein the integrated ASR/SBD module comprises a pretrained neural network tuned for ATC audio; and
 - a question-and-answer (Q/A) module configured to determine aircraft commands based on the utterance hypothesis using a plurality of natural language queries.

9. The system of Claim 8, wherein each utterance hypothesis comprises a boundary hypothesis, wherein the speech-to-text module comprises a sentence boundary detection (SBD) model configured to tag entities within the audio signal and generate the boundary hypothesis based on the tagged entities.

10. The system of Claim 9, wherein the entities comprise a transition speaker, wherein the SBD model is a neural network pre-trained to identify the transition speaker within multi-utterance ATC audio based on audio artifact annotations.

11. The system of Claim 8, wherein the speech-to-text module further comprises an ATC-tuned language model, wherein determining the utterance hypothesis comprises:

- with the integrated ASR/SBD module, generating a plurality of linguistic hypotheses for each utterance;

- using the ATC-tuned language model, determining a language score for each the plurality of linguistic hypotheses; and
 - selecting an utterance hypothesis from the plurality of the linguistic hypotheses based on the corresponding language score.
12. The system of Claim 11, wherein the integrated ASR/SBD module is configured to assign a phonetic score to each of the plurality of linguistic hypotheses, wherein the utterance hypothesis is selected based on a combination of the corresponding language and phonetic scores.
13. The system of Claim 8, wherein determining the utterance hypothesis comprises:
- with the speech-to-text module, generating a plurality of utterance hypotheses for an utterance within the audio signal;
 - using a language model, selecting an utterance hypothesis of the plurality.
14. The system of Claim 13, wherein the language model comprises a neural network pre-trained using entity tagged ATC transcripts.
15. The system of Claim 14, wherein the entity tagged ATC transcripts comprises tags corresponding to phonetically conflicting entities.
16. The system of Claim 8, wherein the utterance hypothesis comprises a text transcript.
17. A method comprising:
- receiving an audio utterance;
 - converting the audio utterance into a predetermined format, comprising:
 - determining a set of utterance hypotheses; and
 - selecting an utterance hypothesis from the set; and
 - based on the selected utterance hypothesis, determining a command using a question-and-answer model,
- wherein the method is performed using the system of any one of Claims 1 – 16.

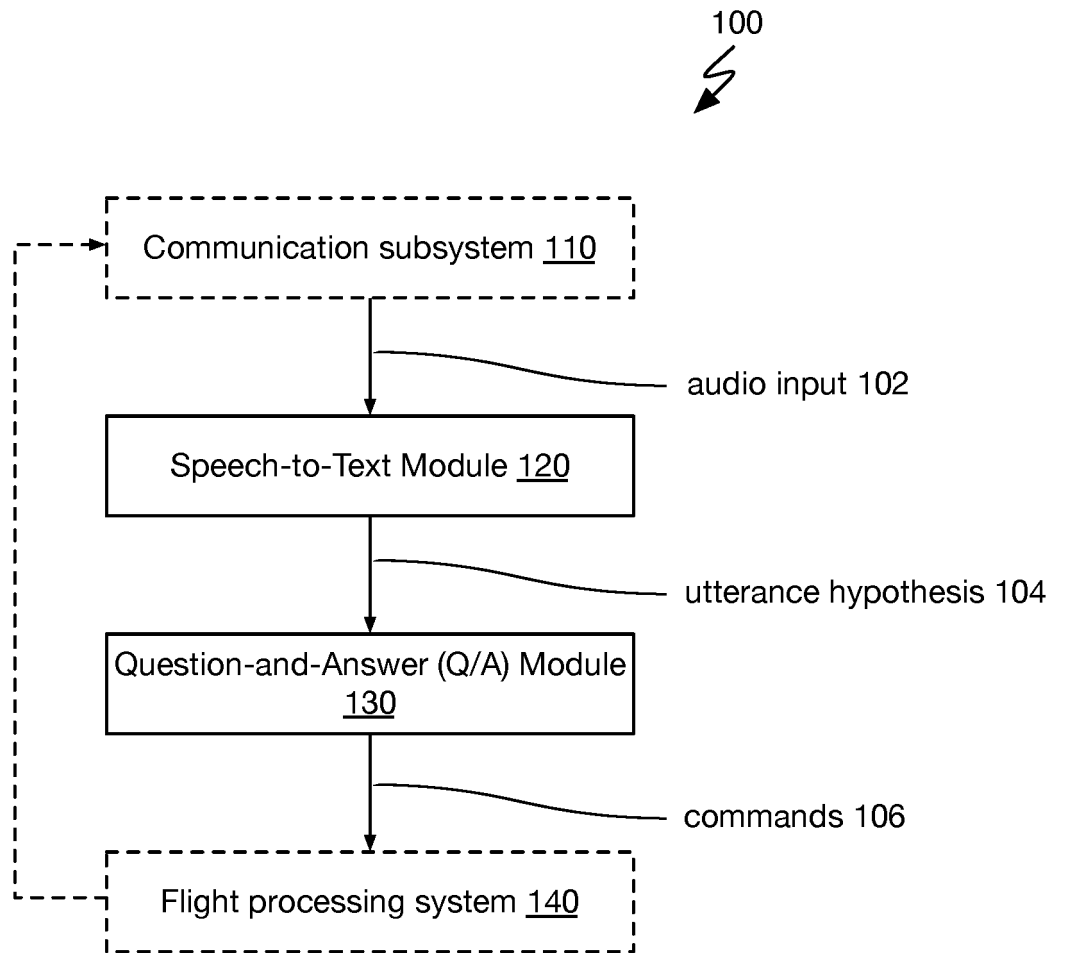


FIGURE 1




FIGURE 2

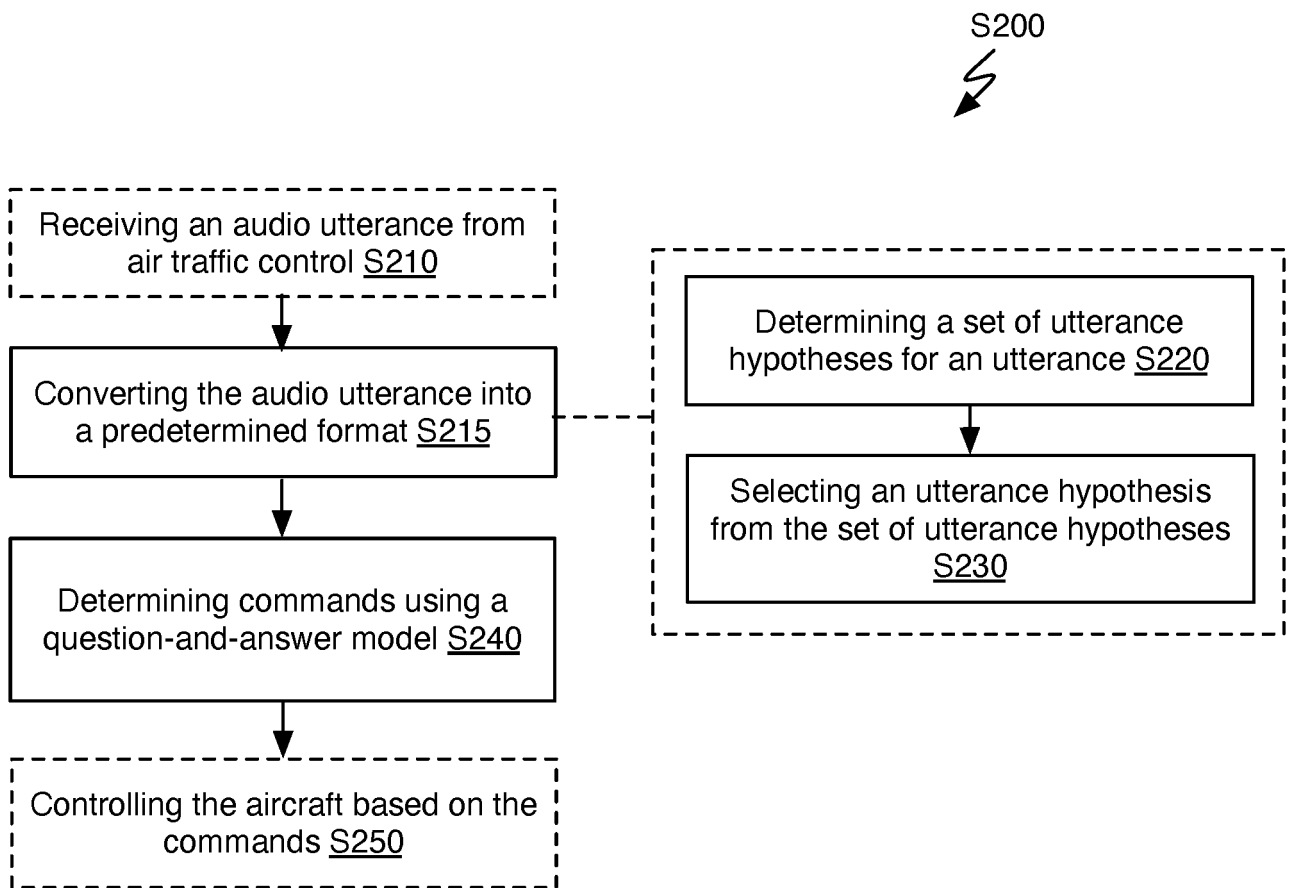


FIGURE 3

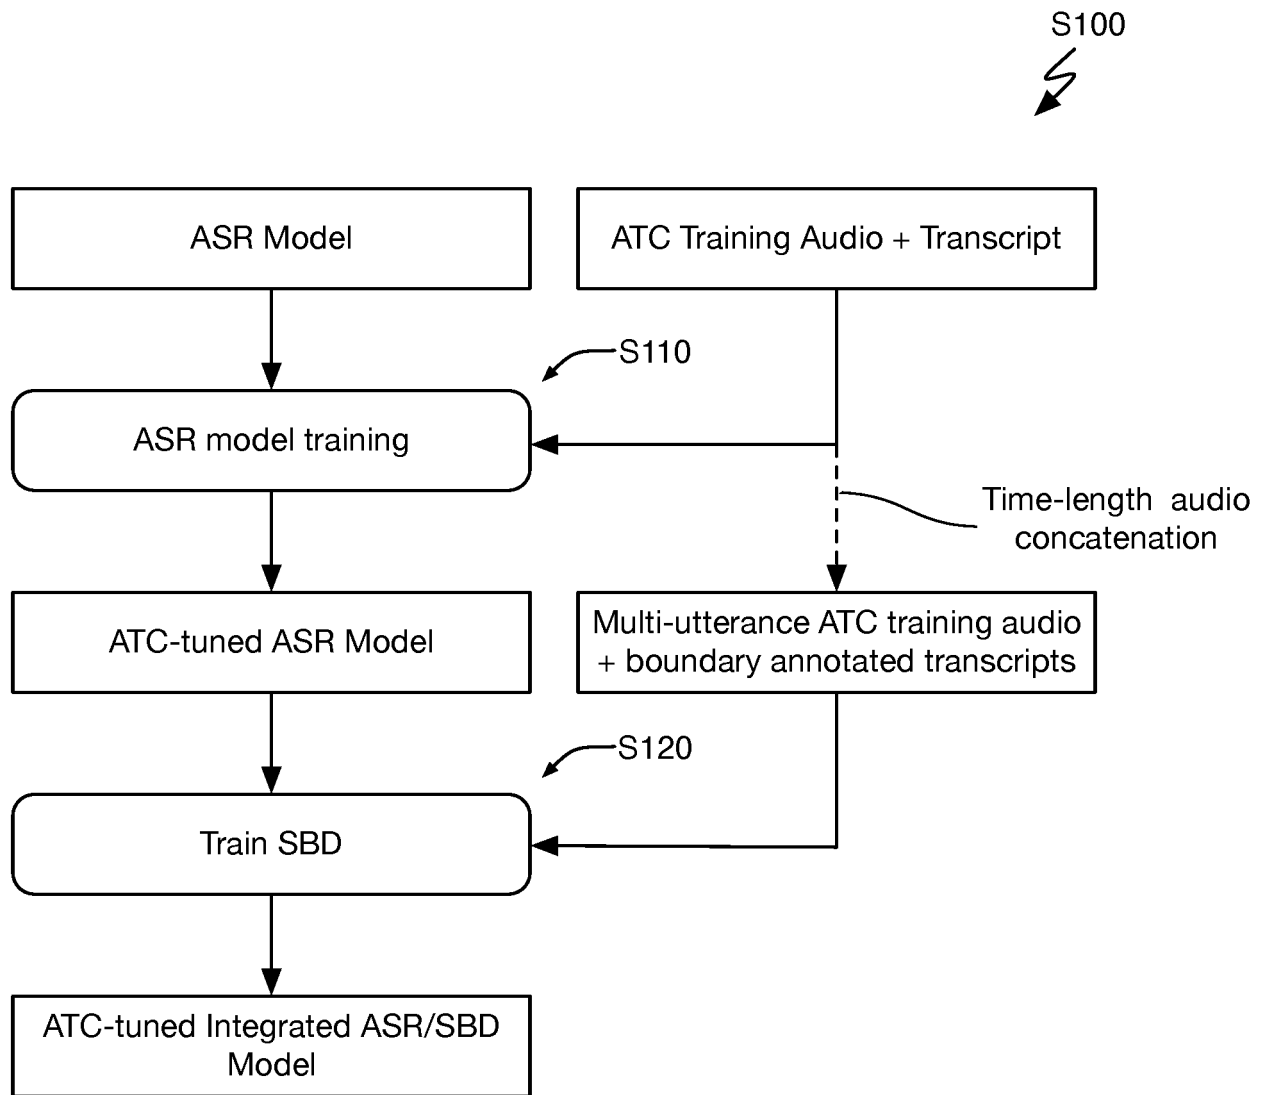


FIGURE 4

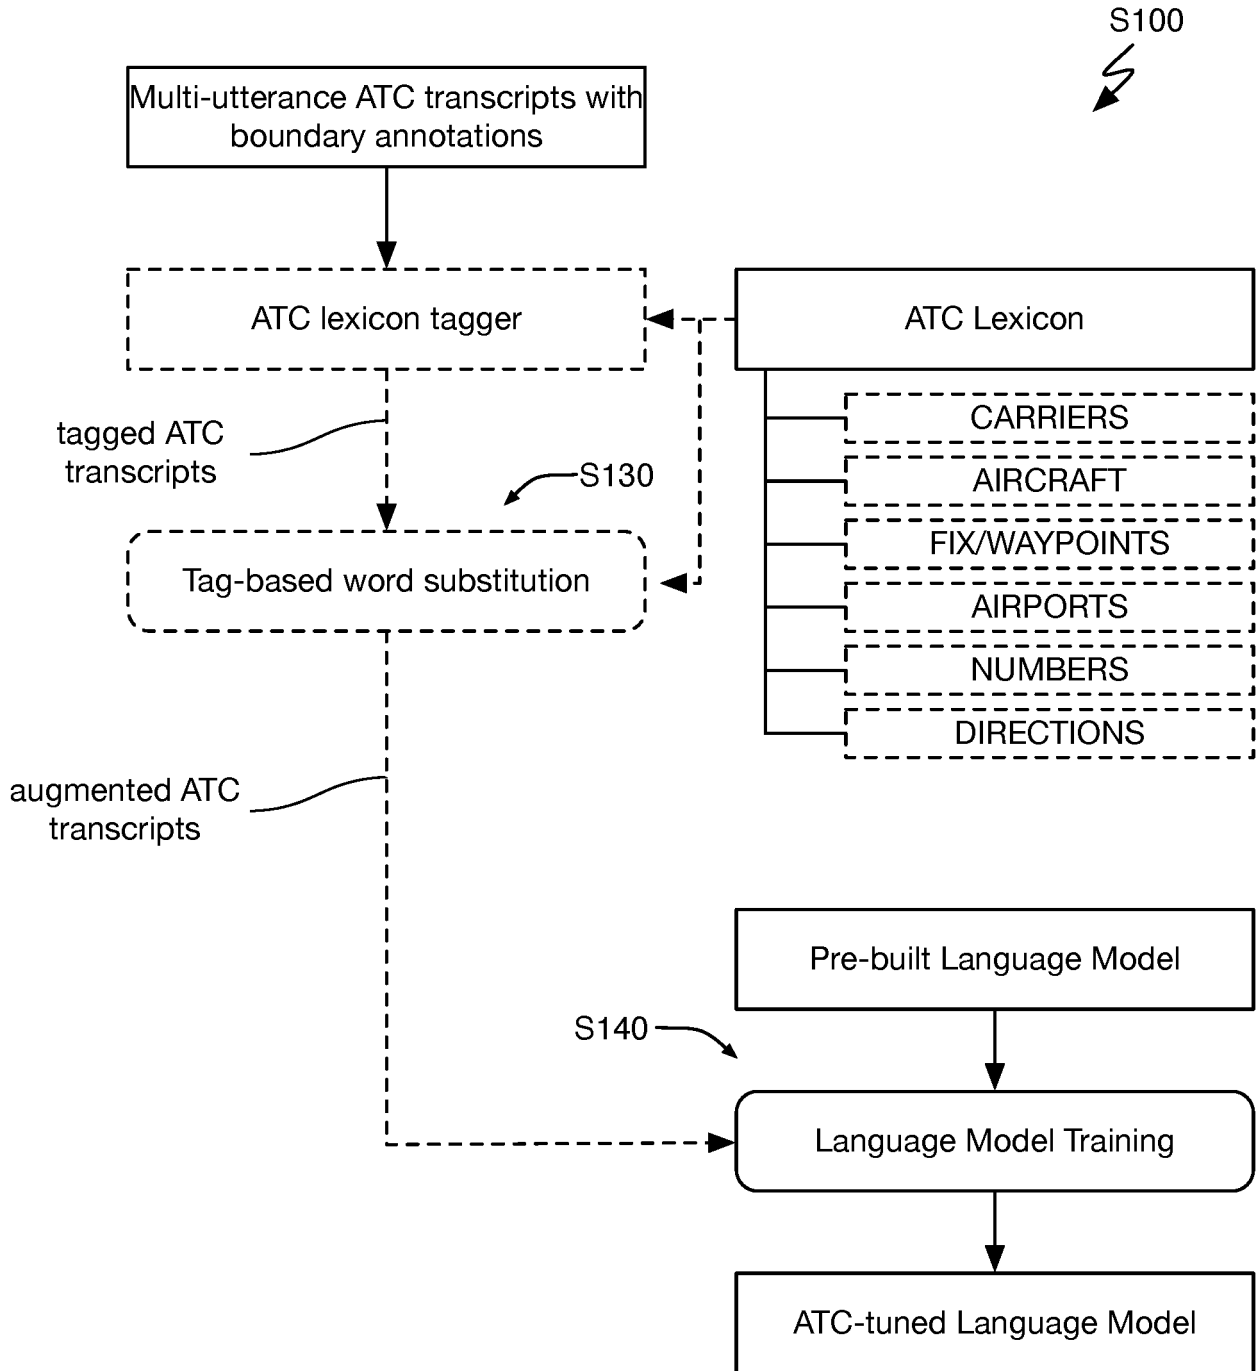


FIGURE 5

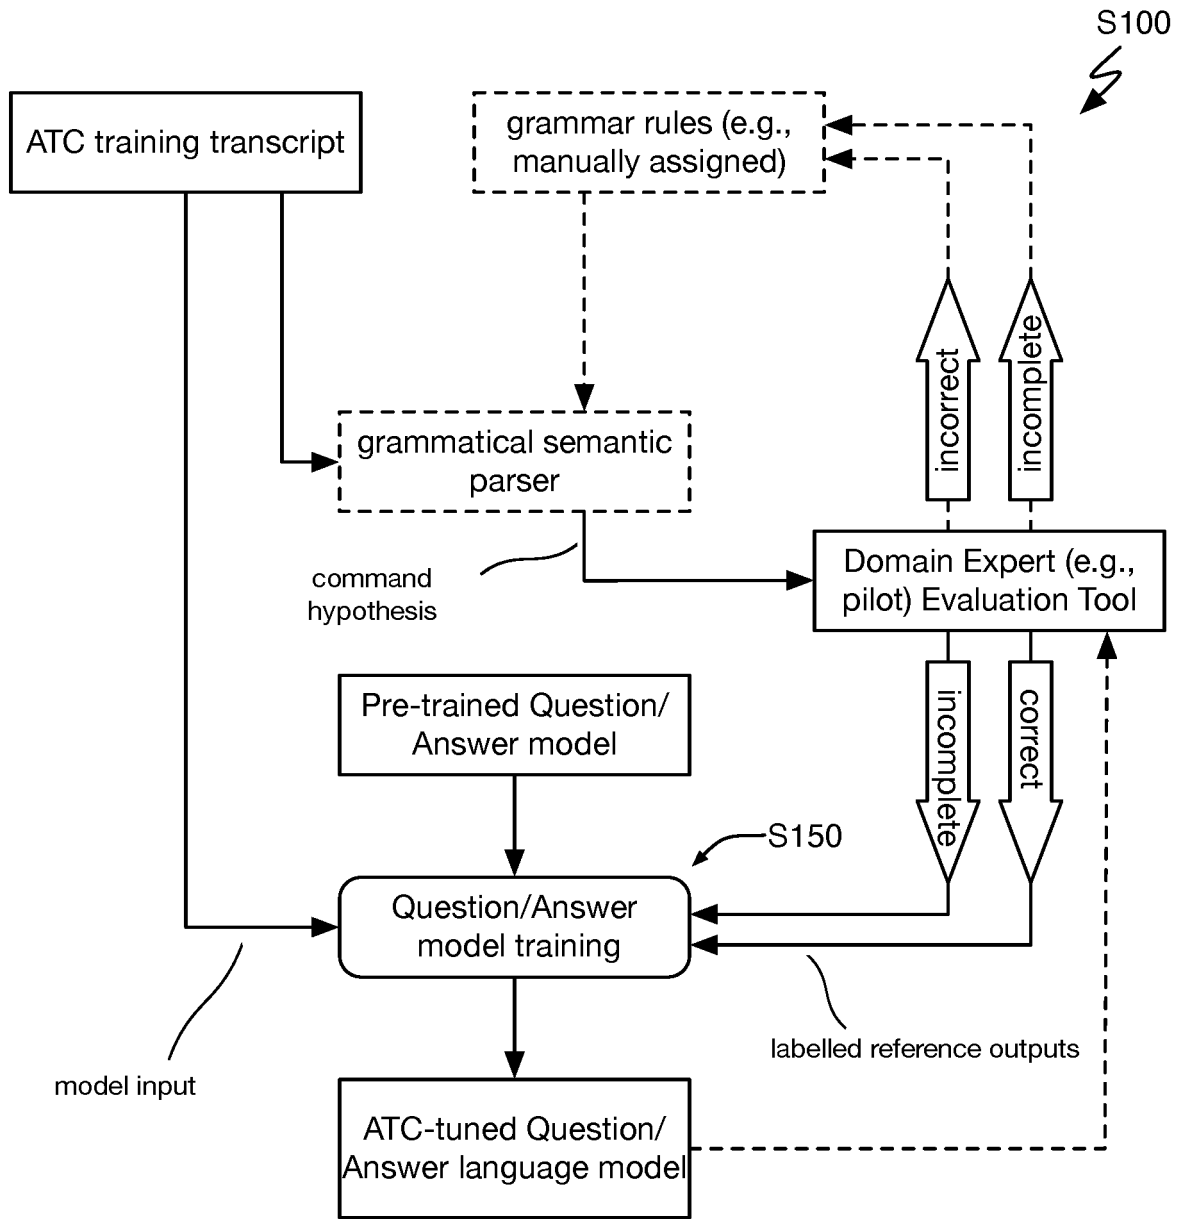


FIGURE 6

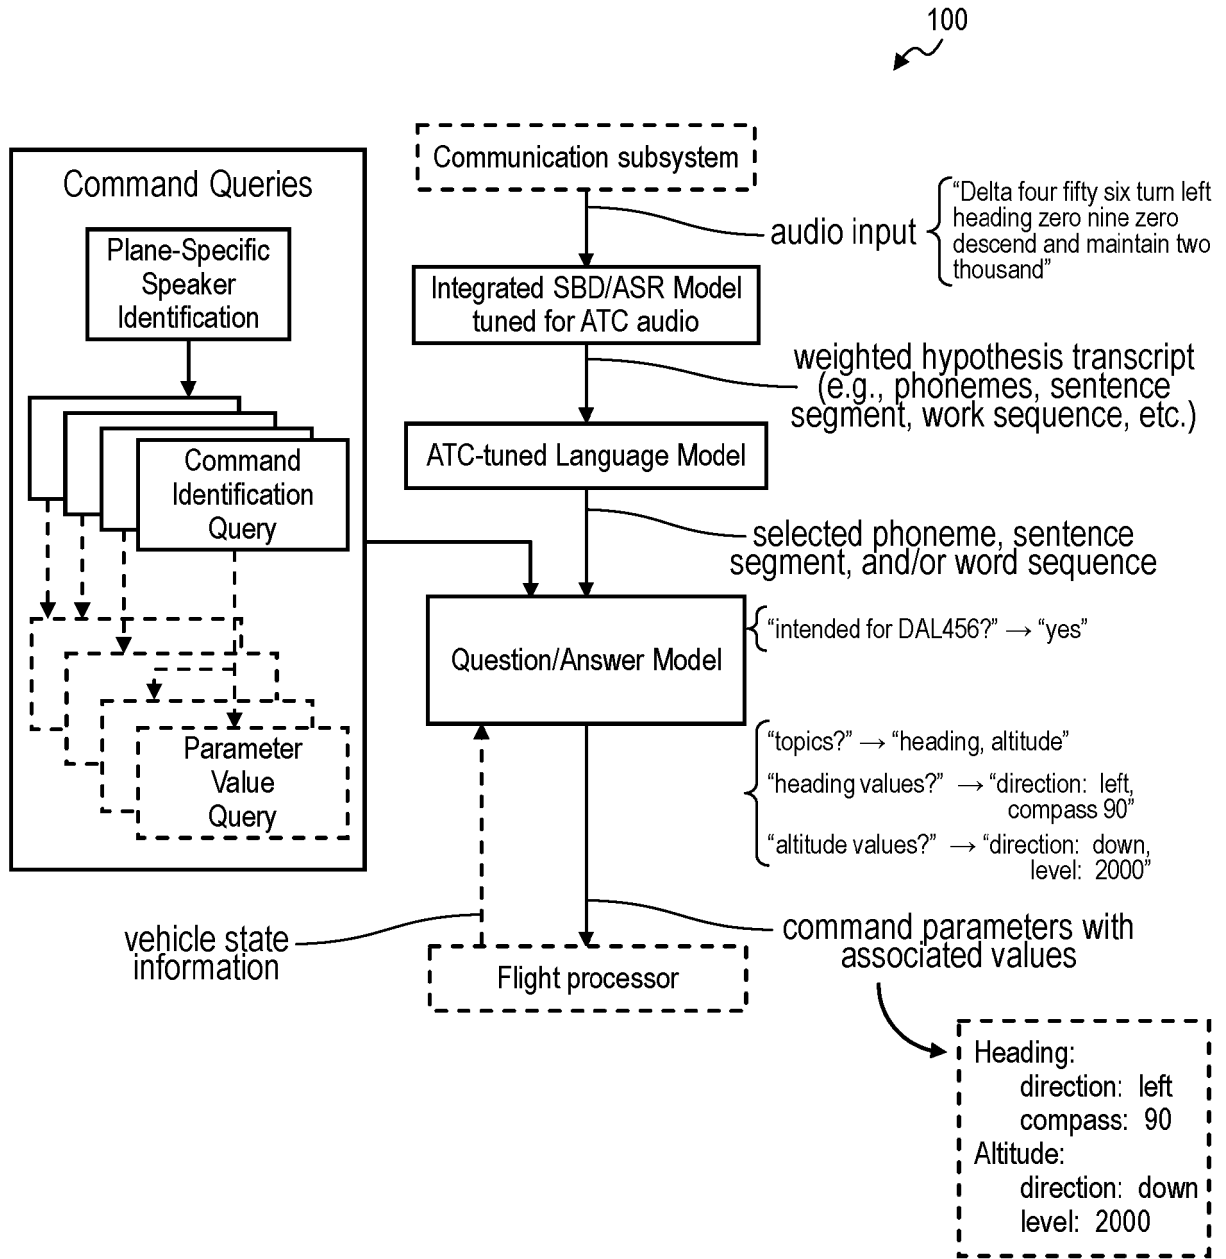


FIGURE 7

Domain Expert
Evaluation Tool
↙

from: F1-2 **to:** USA1524
message: u s air fifteen twenty four boston approach

meaning:
conversation_state: begin

judgement:
 correct
 incomplete
 incorrect

reviewer comments:

from: USA1524 **to:** F1-2
message: u s air fifteen twenty four checking in ah at six thousand in a right turn to south

from: F1-2 **to:** USA1524
message: u s air fifteen twenty four boston approach turn right two five zero intercept two two loc ah two seven localizer maintain four thousand and what is your speed right now

meaning:
altitude:
 direction: ''
 level: 4000
heading:
 compass: 250
 direction: right
request:
 mod: current
 type: airspeed

judgement:
 correct
 incomplete
 incorrect

reviewer comments:

FIGURE 8

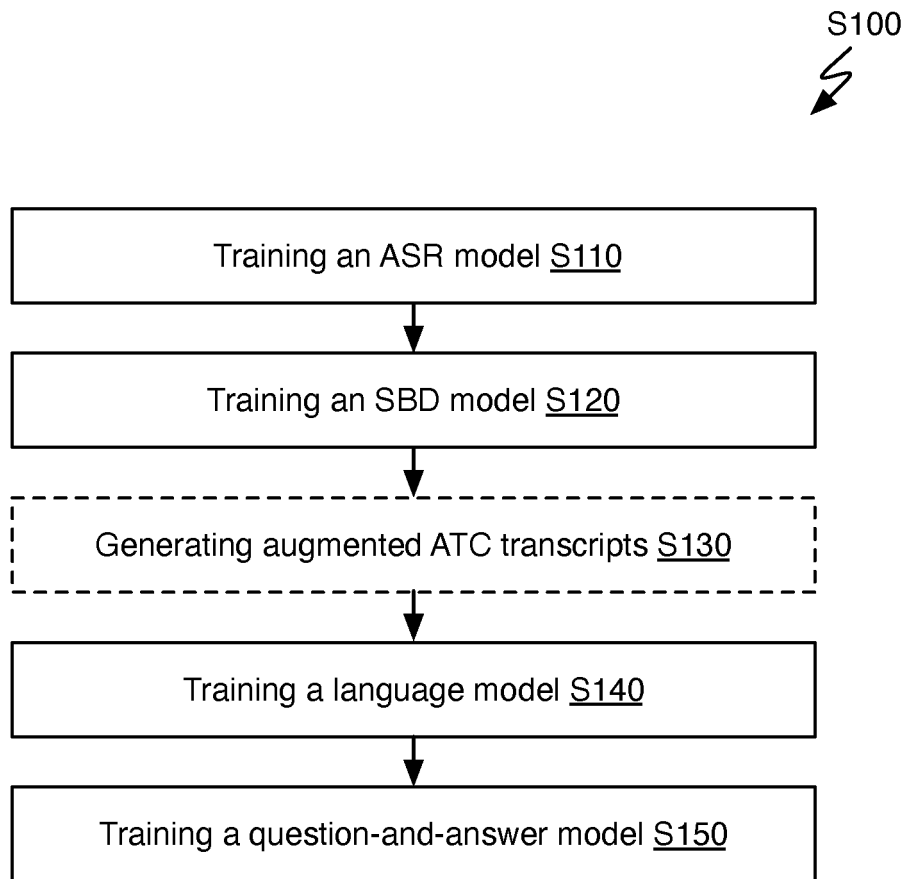


FIGURE 9

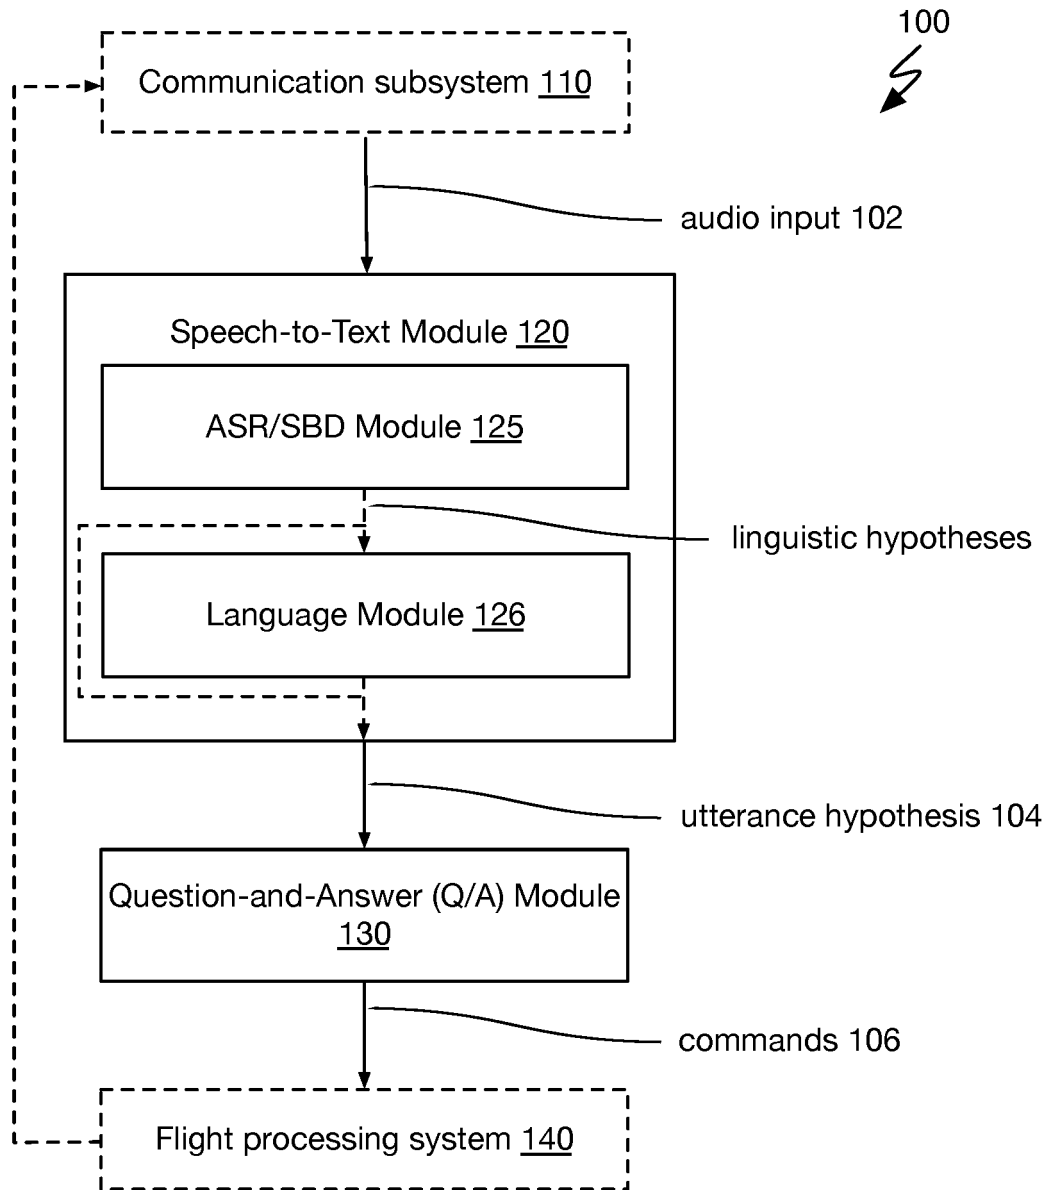


FIGURE 10A

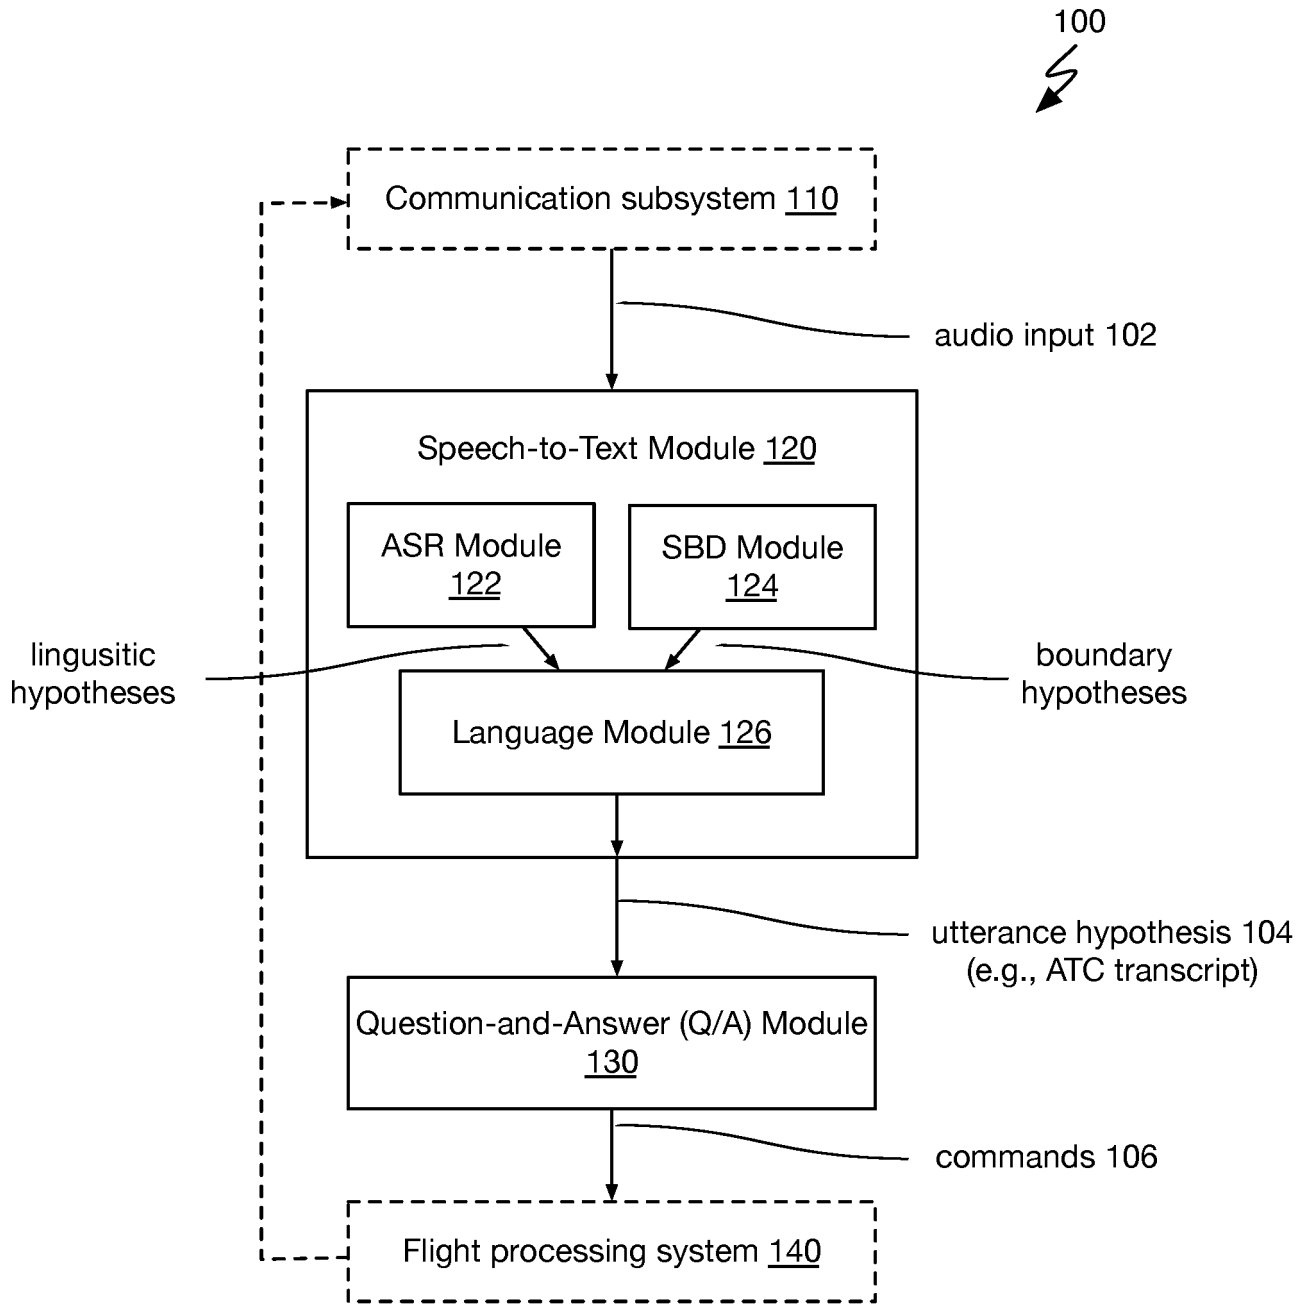


FIGURE 10B

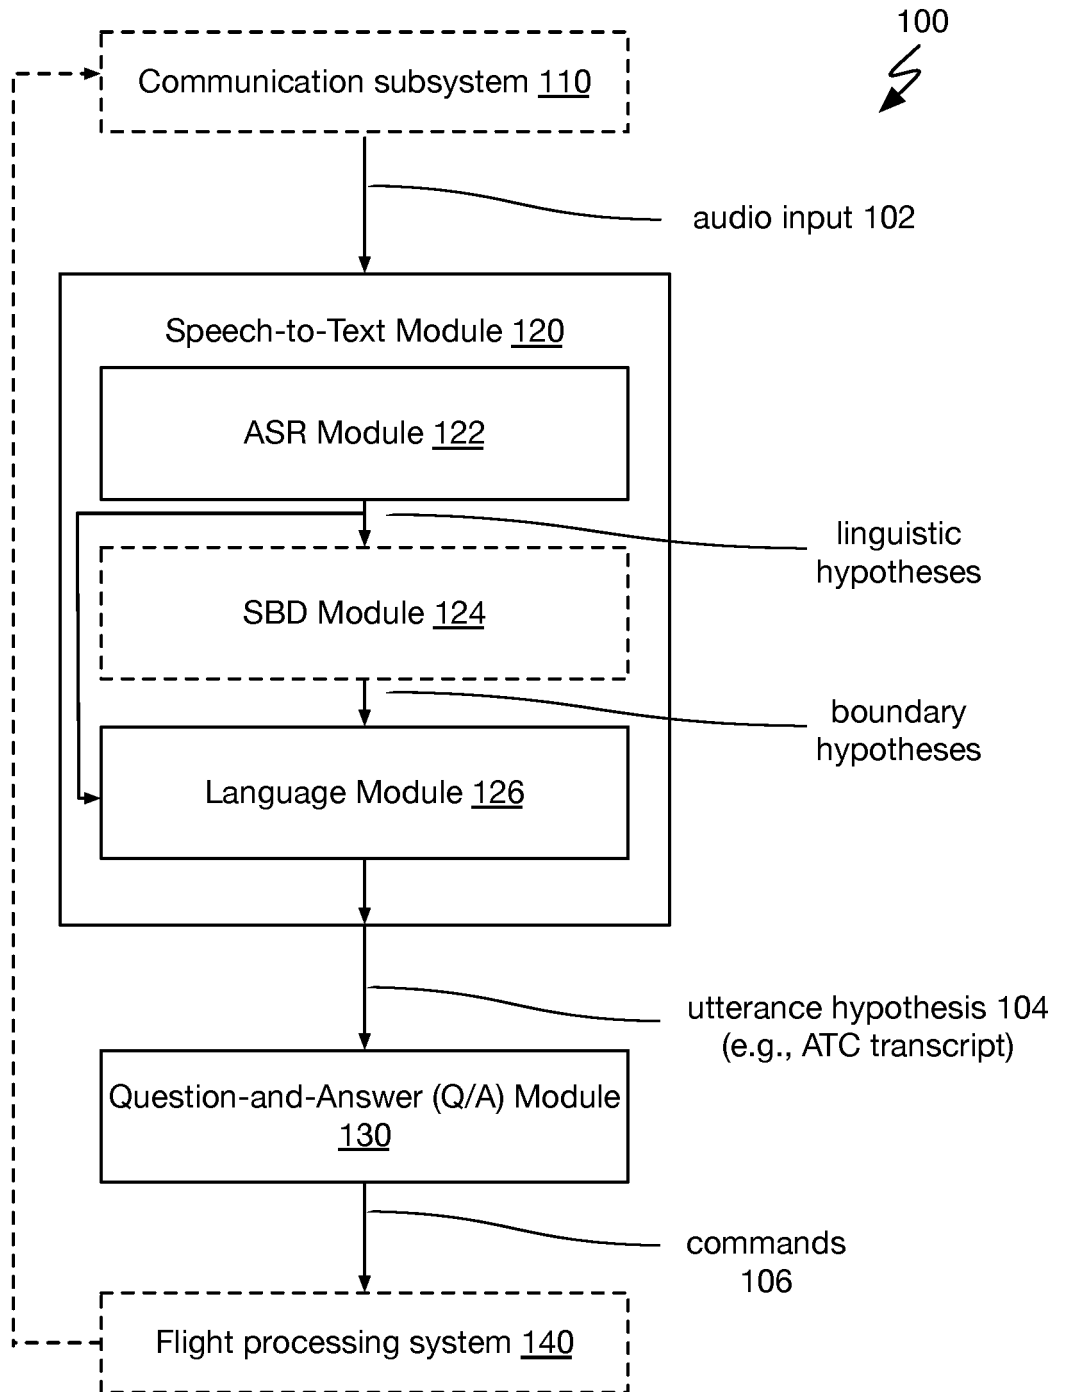


FIGURE 10C

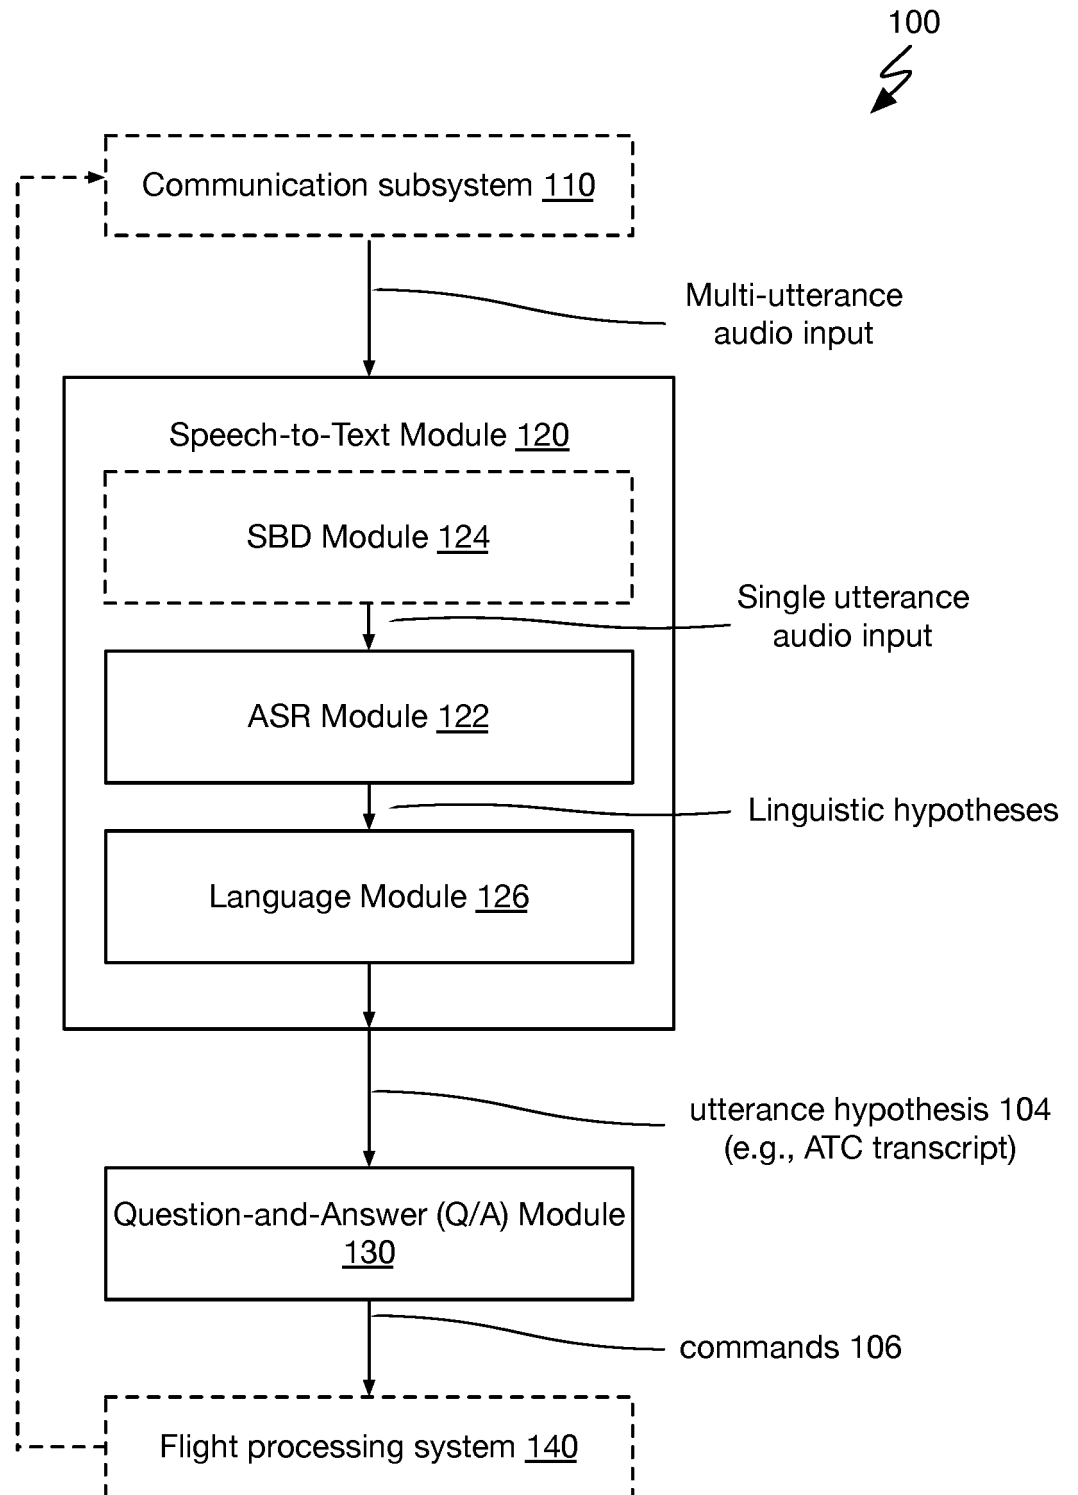


FIGURE 10D

Tree-based query structure

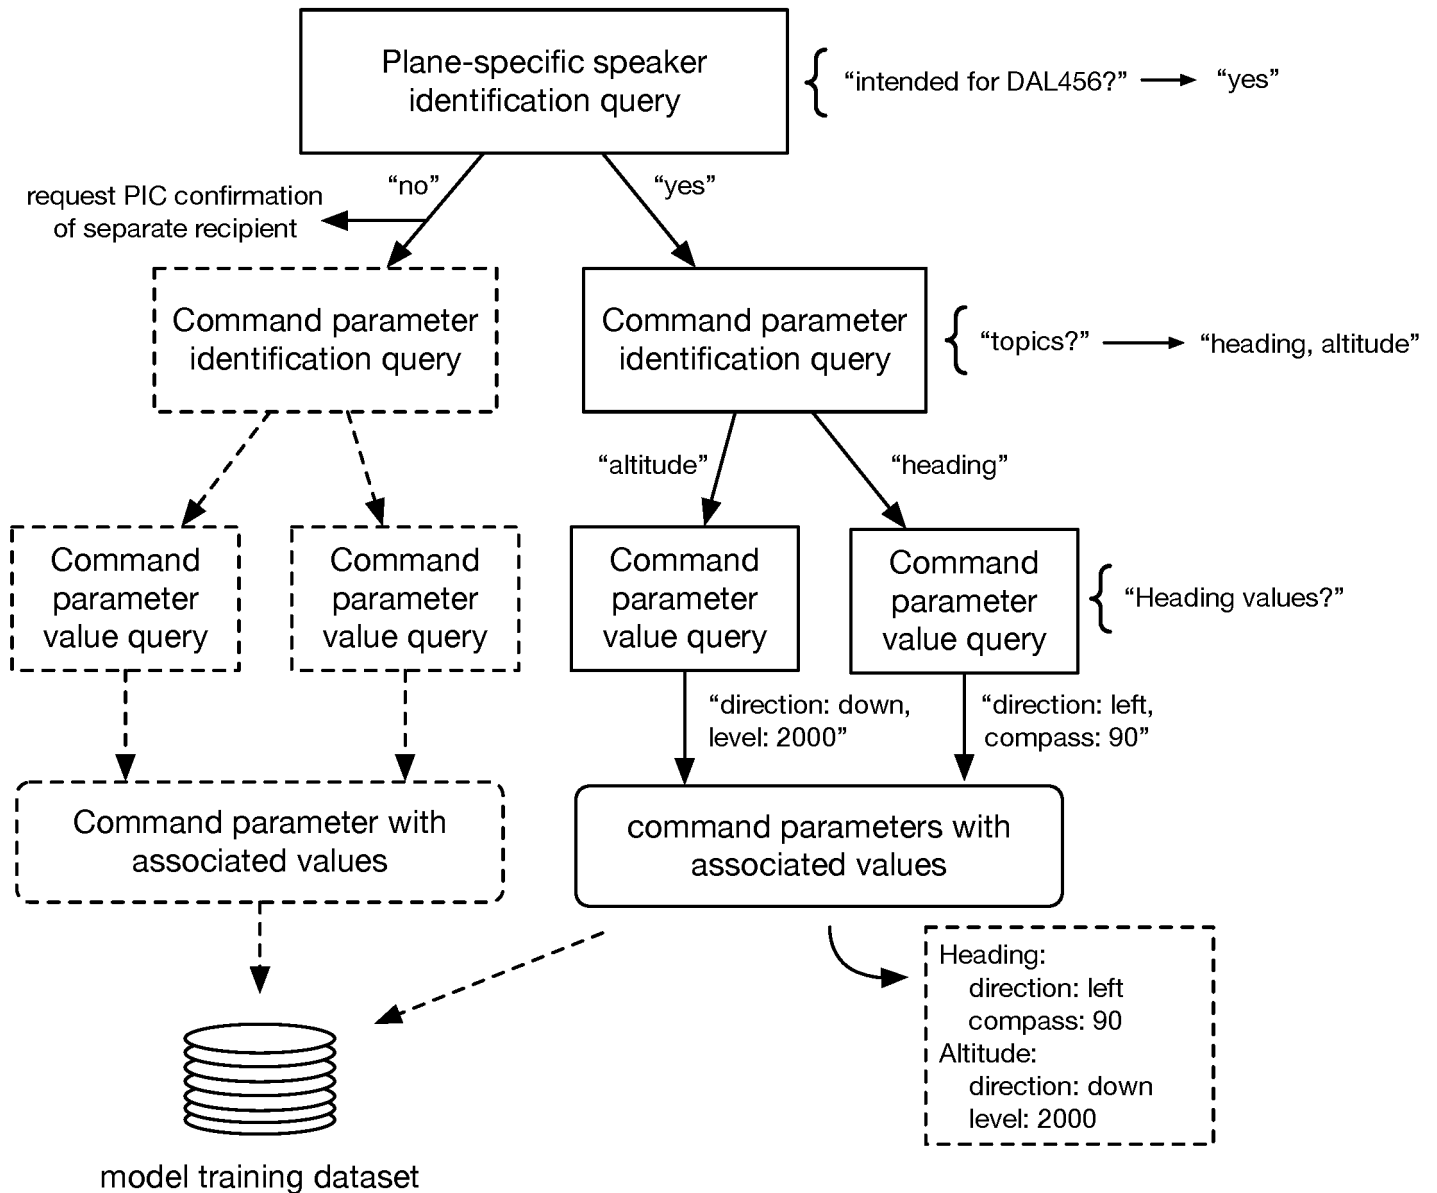


FIGURE 11A

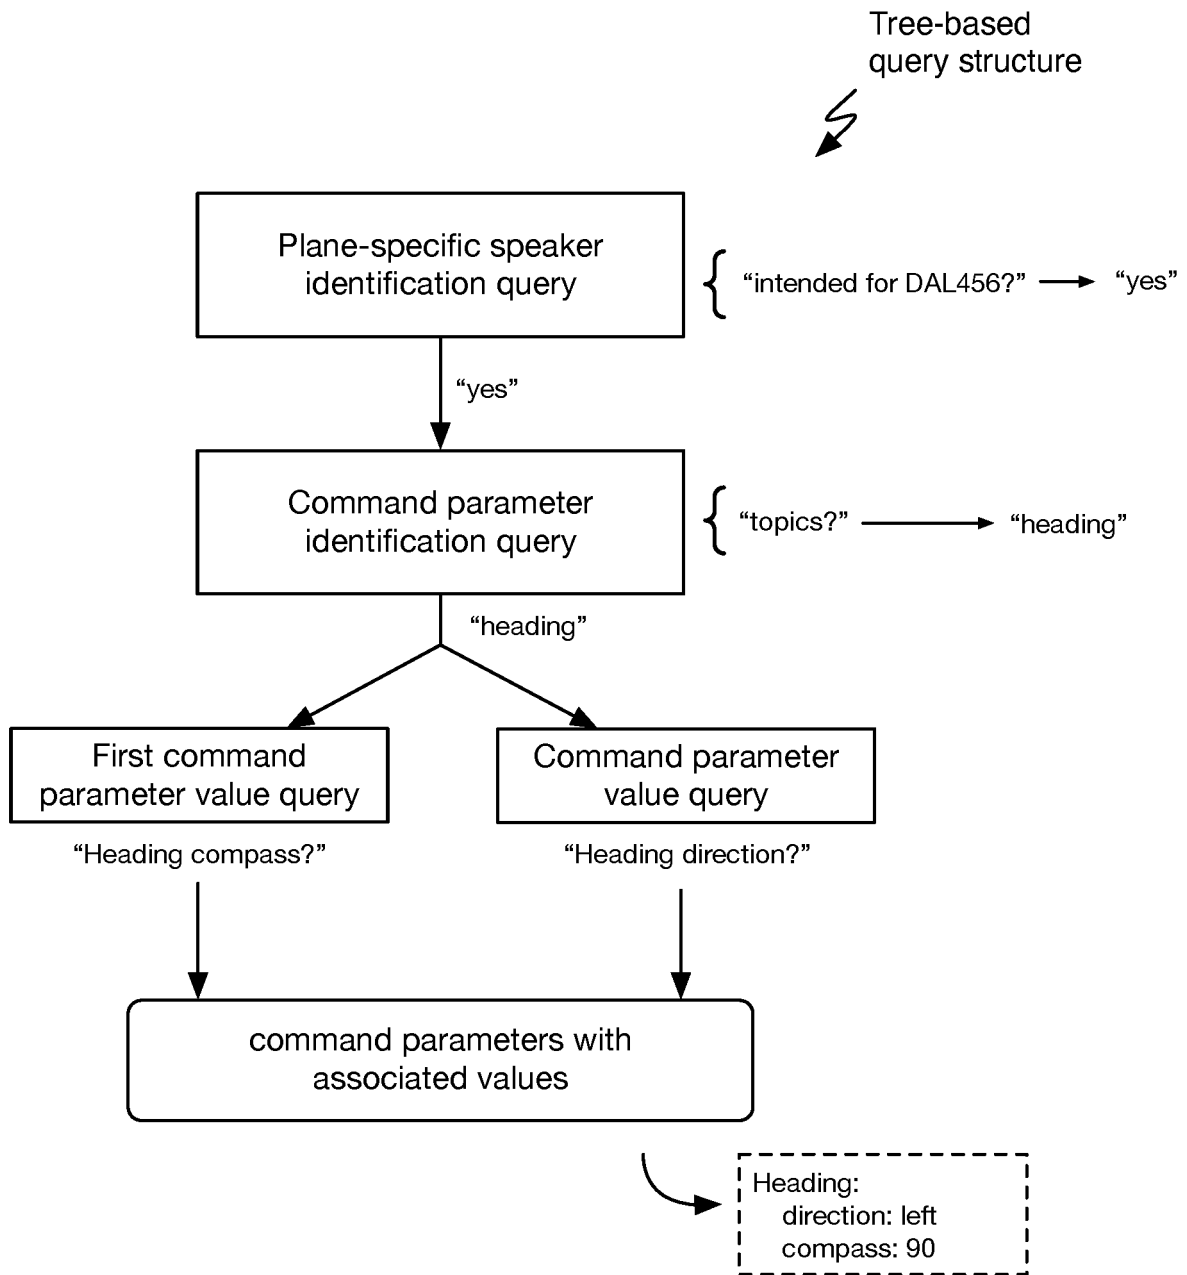


FIGURE 11B

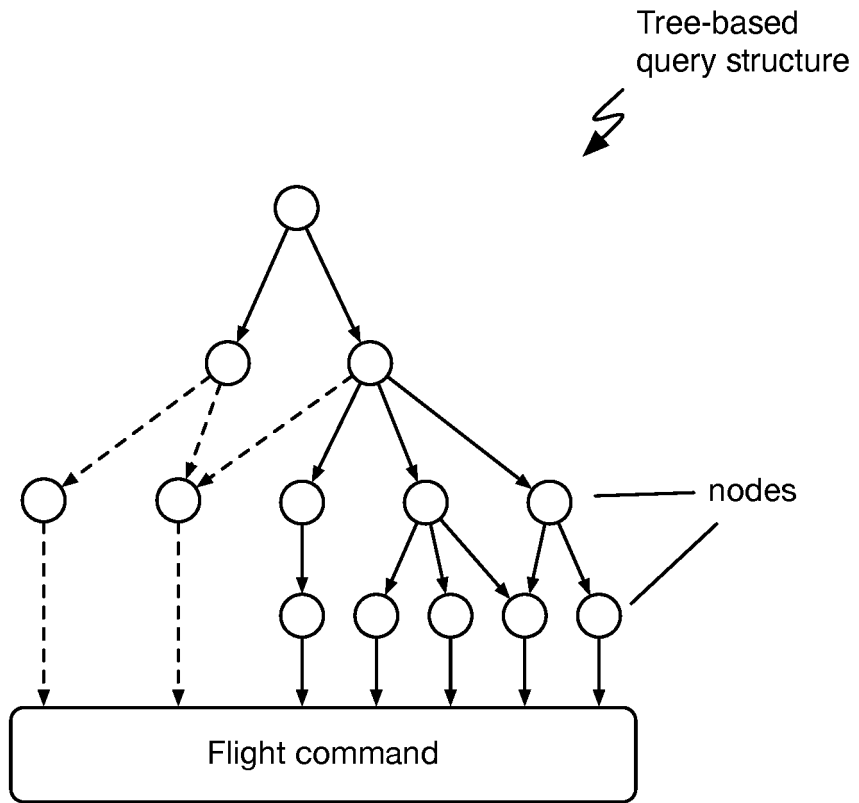


FIGURE 11C

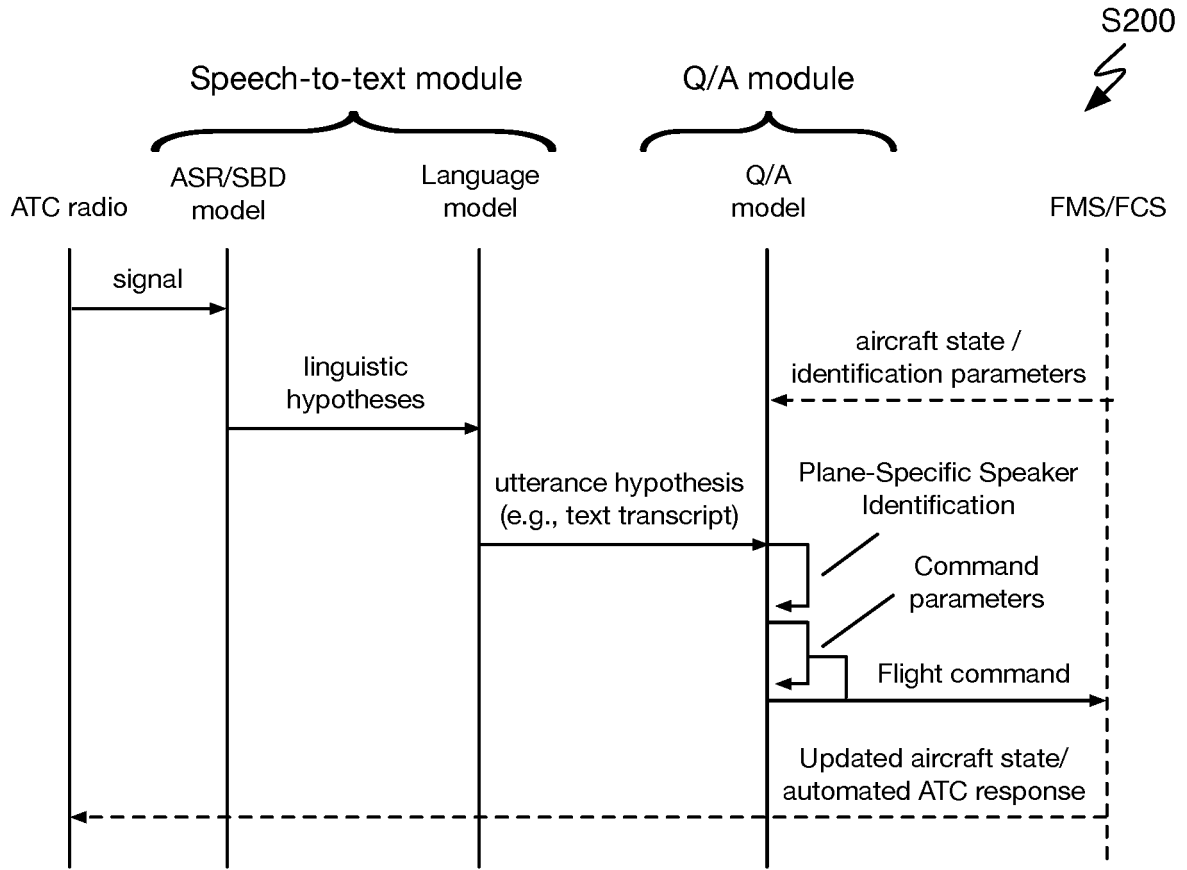


FIGURE 12