



(51) International Patent Classification:
G06N 3/08 (2006.01)

(21) International Application Number:
PCT/US2017/057251

(22) International Filing Date:
18 October 2017 (18.10.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/410,773 20 October 2016 (20.10.2016) US

(71) Applicant: **GOOGLE LLC** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).

(72) Inventors: **LAO, Ni**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **KAISER, Lukasz, Mieczyslaw**; 1600 Amphitheatre Parkway, Mountain View, CA

94043 (US). **GUPTA, Nitin**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **MOHIUDDIN, Afroz**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **POPAT, Preyas**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).

(74) Agent: **PORTNOV, Michael**; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, MN 55440-1022 (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, 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,

(54) Title: ANSWER TO QUESTION NEURAL NETWORKS




FIG. 1

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on computer storage media, for identifying answers to questions using neural networks. One of the methods includes receiving an input text passage and an input question string; processing the input text passage using an encoder neural network to generate a respective encoded representation for each passage token in the input text passage; at each time step: processing a decoder input using a decoder neural network to update the internal state of the decoder neural network; and processing the respective encoded representations and a preceding output of the decoder neural network using a matching vector neural network to generate a matching vector for the time step; and generating an answer score that indicates how well the input text passage answers a question posed by the input question string.

SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, 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, 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))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

ANSWER TO QUESTION NEURAL NETWORKS

BACKGROUND

[0001] This specification relates to identifying answers to questions using neural networks.

[0002] Neural networks are machine learning models that employ one or more layers of models to predict an output for a received input. Some neural networks include one or more hidden layers in addition to an output layer. The output of each hidden layer is used as input to the next layer in the network, i.e., the next hidden layer or the output layer. Each layer of the network generates an output from a received input in accordance with current values of a respective set of parameters.

[0003] Some neural networks are recurrent neural networks. A recurrent neural network is a neural network that receives an input sequence and generates an output sequence from the input sequence. In particular, a recurrent neural network can use some or all of the internal state of the network from a previous time step in computing an output at a current time step. An example of a recurrent neural network is a long short term (LSTM) neural network that includes one or more LSTM memory blocks. Each LSTM memory block can include one or more cells that each include an input gate, a forget gate, and an output gate that allow the cell to store previous states for the cell, e.g., for use in generating a current activation or to be provided to other components of the LSTM neural network.

SUMMARY

[0004] In general, one innovative aspect of the subject matter described in this specification can be embodied in a method comprising: receiving (i) an input text passage comprising a plurality of passage tokens and (ii) an input question string comprising a respective question token at each of a plurality of time steps; processing the input text passage using an encoder neural network, wherein the encoder neural network is configured to receive the input text passage and to process the input text passage to generate a respective encoded representation for each passage token in the input text passage; at each time step: processing a decoder input for the time step using a decoder neural network, wherein the decoder input comprises at least a preceding token in the input question string and a preceding internal state of the decoder neural network, and wherein the decoder neural network is configured to, for each time step, process the decoder input to at least update the internal state of the decoder neural network; and processing the respective encoded representations and a preceding output of the decoder neural network using a matching vector neural network, wherein the matching vector neural

network is configured to receive the preceding output of the decoder neural network and the encoded representations and to process the preceding output and the encoded representations to generate a matching vector for the time step; and generating an answer score that indicates how well the input text passage answers a question posed by the input question string based at least on the matching vectors for the time steps.

[0005] In a development of the method, the method may further comprise the step of: at each time step: processing the respective encoded representations and the preceding output of the decoder neural network using an attention vector neural network, wherein the attention vector neural network is configured to receive the preceding output of the decoder neural network and the encoded representations and to process the preceding output and the encoded representations to generate an attention vector for the time step, wherein the attention vector includes a respective value for each passage token. In a further development of the method, the attention vector may have the same dimensionality as the matching vector. In a further development of the method, the attention vector neural network may have the same neural network architecture as the matching vector neural network but different parameter values. In some embodiments, at each time step the decoder input may further comprise the attention vector for the time step. In some embodiments, generating the answer score may comprise generating a matching vector score by determining an average or a weighted average of the entries of the matching vectors for the time steps. In a development of the method, the decoder neural network may be further configured to, at each time step, generate a set of question token scores that includes a respective score for each question token in a vocabulary of possible question tokens. In a further development of the method, generating the answer score may comprise: determining a perplexity score from an average perplexity of the question tokens in the input question string as reflected by the question token scores generated by the decoder neural network for the question tokens in the input question string; and combining the perplexity score and the matching vector score to generate the answer score. In some embodiments, the encoder neural network and the decoder neural network may be recurrent neural networks. In some embodiments, the preceding output may be an output of one of the layers of the decoder neural network at a preceding time step. In a development of the method, the method may further comprise: determining, based on the answer score, that the input passage answers the question posed by the input passage string; and in response, outputting data derived from the input passage. In some embodiments, the encoder neural network, decoder neural network and matching vector neural network may be trained based upon first training data, wherein the first training data comprises: a plurality of

training passages and, for each training passage: a training question string for which the training passage has been determined to be a high-quality answer, and a training answer score for each training passage – training question string pair; wherein training the encoder neural network, decoder neural network and matching vector neural network comprises minimizing a loss function that is a weighted combination of (i) an average token perplexity loss function and (ii) a pseudo Huber loss function. In some embodiments, the average token perplexity loss function may measure the average per token perplexity for the question tokens in the training question string generated by the decoder neural network. In some embodiments, the method may further comprise: receiving second training data for training the encoder neural network, decoder neural network and matching vector neural network, wherein the second training data comprises: a plurality of training passages and, for each training passage: a training question string; and a training answer score for the training passage – training question string pair; wherein the second training data includes training passage – training question string pairs of varying quality; and wherein the weight assigned to the average per token perplexity loss function is reduced relative to the weight assigned to the pseudo Huber loss function when training the encoder neural network, decoder neural network and matching vector neural network based upon the second training data.

[0006] Other embodiments of these aspects include corresponding computer systems, apparatus, and computer programs recorded on one or more computer storage devices, each configured to perform the actions of the methods. A system of one or more computers can be configured to perform particular operations or actions by virtue of software, firmware, hardware, or any combination thereof installed on the system that in operation may cause the system to perform the actions. One or more computer programs can be configured to perform particular operations or actions by virtue of including instructions that, when executed by data processing apparatus, cause the apparatus to perform the actions.

[0007] Another innovative aspect of the subject matter described in this specification can be embodied in a computer system that includes an encoder neural network, a decoder neural network, a matching vector neural network. The encoder neural network is configured to receive an input text passage and to generate a respective encoded representation for each passage token in the input text passage. The decoder neural network is configured to, for each of a plurality of time steps, receive as input at least a preceding token in the input question string and a preceding internal state of the decoder neural network and to process the input to update the internal state of the decoder neural network and, optionally, generate a set of question token scores for the time step. The matching vector neural network is configured

to, at each time step, receive the preceding internal state of the decoder neural network and the encoded representations and to generate a matching vector for the time step. The system also includes a subsystem that is configured to generate an answer score that indicates how well the input text passage answers the question posed by the input question string based at least on the matching vectors for the time step.

[0008] Other embodiments of this aspect include corresponding apparatus, and computer programs recorded on one or more computer storage devices, each configured to implement the computer system. A system of one or more computers can be configured to perform particular operations or actions by virtue of software, firmware, hardware, or any combination thereof installed on the system that in operation may cause the system to perform the actions. One or more computer programs can be configured to perform particular operations or actions by virtue of including instructions that, when executed by data processing apparatus, cause the apparatus to perform the actions.

[0009] The subject matter described in this specification can be implemented in particular embodiments so as to realize one or more of the following advantages. An answer to question system can effectively generate an answer score that indicates how well a passage answers a given question. The answer to question system can also effectively generate a question that is answered by an input passage. By training the neural network-based components of the answer to question system as described in this specification, the answer to question system can not only effectively determine the best question for a given input passage, but can also effectively score input questions to reflect how well the input passage answers the input question, i.e., to discern bad answers from mediocre answers, mediocre answers from good answers, bad answers from good answers, and so on. The neural network can therefore be trained using given relationships between text inputs such that the neural network is then able to determine corresponding relationships based upon unseen input. Additionally, because the system is configured to and is trained to determine questions that are answered by a given passage rather than determining passages that answer a given question, the amount of data needed to train the system and that needs to be analyzed by the system to generate an output is greatly reduced. In particular, the system does not need to analyze a large body of electronic documents to generate an output and can instead analyze multiple questions, which tend to be shorter and more computationally efficient to process than electronic documents or passages extracted from electronic documents, against a single answer, rather than analyzing a single question against multiple answers, which are generally longer and less computationally efficient to process. In other words, because the system

encodes a single relatively longer passage and can then score multiple relatively shorter possible questions rather than encoding a single short question and then scoring multiple relatively longer possible passages when generating an output, the system is much more computationally-efficient and can generate more accurate answer scores. In particular, a system according to embodiments is such that it may be effectively trained using a relatively small body of text.

[0010] The use of the answer to question system can improve the operation of the computer or computers on which the system is implemented. For example, because the answer to question system allows answer scores for input passage – input answer pairs to be effectively computed in advance, less compute time, bandwidth, battery use, and so on is required to determine appropriate answers to questions when the questions are actually posed. As a particular example, when a search query is received that has been classified as asking a question, a passage that effectively answers the question posed by the search query can be quickly identified and presented as part of the response to the search query. As another example, because the answer to question system allows answer scores for input passage – input question pairs to be determined without needing to use templates or other external data, the amount of computing resources needed to effectively determine the question score are reduced, i.e., the amount of processing resources and memory needed to determine the answer score are reduced. That is, because the described system scores an input passage – question pair without needing to take any other external data, e.g., templates, as input, the system can generate the scores using fewer computing resources.

[0011] The details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] FIG. 1 shows an example answer to question system.

[0013] FIG. 2 shows an example attention, matching, and decoding subsystem.

[0014] FIG. 3 is a flow diagram of an example process for selecting a passage that provides an answer to a question.

[0015] FIG. 4 is a flow diagram of an example process for generating a question that is answered by a passage.

[0016] FIG. 5 is a flow diagram of an example process for training an encoder neural network and an attention, matching, and decoding subsystem.

[0017] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0018] FIG. 1 shows an example answer to question system 100. The answer to question system 100 is an example of a system implemented as computer programs on one or more computers in one or more locations, in which the systems, components, and techniques described below can be implemented.

[0019] The answer to question system 100 can (i) receive an input text passage and generate a score for each of one or more input question strings that indicates how well the input text passage answers the question posed by the input question string or (ii) receive an input text passage and generate an output question string that is answered by the input text passage.

[0020] For example, the answer to question system 100 can obtain an input passage 102 and generate a corresponding question string 110 that is answered by the input passage 102. As another example, the answer to question system 100 can obtain an input passage 102 and an input question string 112 and generate a corresponding answer score 114 that indicates how well the input passage 102 answers the input question string 112. In other words, the answer score 114 indicates how well the input passage 102 satisfies the informational need of a user posing the input question string 112.

[0021] An input passage is a sequence of tokens, referred to in this specification as passage tokens, that each represent n -grams, i.e., characters, words, sub-words, and so on. For example, the passage 102 may be a phrase, a sentence, a paragraph, or another text segment extracted from a source document.

[0022] A question string is a sequence of tokens, referred to in this specification as question tokens, that each represent n -grams, such that the question string represents a natural language question or search query that explicitly or implicitly poses a question.

[0023] The answer to question system 100 includes an encoder neural network 104 and an attention, matching, and decoding subsystem 108.

[0024] The encoder neural network 104 processes a given input passage and generates a respective encoded representation for each passage token in the input passage, e.g., encoded representations 106 for the passage tokens in the input passage 102. The encoded

representation for each passage token is a vector or other ordered collection of numeric values that represents the passage token (or the n -gram represented by the passage token) in the context of the input passage. Generally, the encoder neural network 104 is a recurrent neural network, e.g., a long short-term memory neural network (LSTM), that has been configured through training to generate the encoded representations. Example encoder recurrent neural networks that can be configured through training to generate appropriate encoded representations are described in Wu, Yonghui, et al, *Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation*, available at <https://arxiv.org/abs/1609.08144>.

[0025] When the answer to question system 100 is generating a question score for an input question, the subsystem 108 receives the encoded representations 106 and the input question string 108 and generates a matching output that defines the answer score for the input question string 108 that indicates how well the input passage 102 answers the question posed by the input question string 108. A question scoring engine 104 in the answer to question system 100 then receives the matching output generated by the subsystem 108 and generates the answer score from at least the matching output. The operations performed by the subsystem 108 to generate a matching score are described in more detail below with reference to FIG. 2.

[0026] When the answer to question system 100 is generating an output question string, the subsystem 108 receives the encoded representations 106 and generates an output question string that is answered by the input passage 102. The operations performed by the subsystem 108 to generate a question string are described in more detail below with reference to FIG. 2.

[0027] FIG. 2 shows an example attention, matching, and decoding subsystem 108 as described above with reference to FIG. 1.

[0028] Generally, when generating an output question string, the subsystem 108 generates the output question string token by token, i.e., by outputting a respective set of question token scores 214 time step by time step until the output question string has been completed. Each set of question token scores is a score distribution over tokens in a vocabulary of tokens that may be included in an output question string. The vocabulary of tokens can include multiple tokens that each represent n -grams, i.e., characters, words, sub-words, and so on. The vocabulary of tokens also includes one or more special tokens, e.g., an end-of-sequence token, a start-of-sequence token, and so on.

[0029] The subsystem 108 includes a decoder Recurrent Neural Network (RNN) 202, an attention vector neural network 204, and a matching vector neural network 220.

[0030] The decoder RNN 202 is a recurrent neural network that receives an input sequence and maintains an internal state while processing the inputs in the input sequence. In particular, the RNN uses some or all of the internal hidden state of the network from processing a preceding input in the input sequence in combination with the preceding output in the output sequence to update the internal hidden state of the recurrent neural network and to generate a decoder output. In this case, the decoder RNN 202 has been configured to, for each time step, generate a decoder output that is a score distribution over tokens in the vocabulary of tokens that may be included in an output question string. Generally, the decoder RNN includes multiple recurrent neural network layers, e.g., multiple LSTM layers, and an output layer.

[0031] At each time step, the decoder RNN 202 receives as input a preceding question token 208 and an attention vector 210 generated by the attention vector neural network 204 for the time step and processes the preceding question token 208 and the preceding attention vector 210 to generate as output an updated hidden state of the decoder RNN for the time step, i.e., to update the internal hidden states of the recurrent layers, and a set of question token scores 214 that define a score distribution over the question tokens in the vocabulary of possible question tokens.

[0032] The token that is used as the preceding question token 208 for a given time step is dependent on the task that the answer to question system is currently performing.

[0033] When the task is to generate an output question string that is answered by the input passage, the preceding token is the question token that has been selected by the system as the question token at the immediately preceding time step in the output question string based on the question token scores generated at the preceding time step. Selecting a question token for a time step based on question token scores is described further below.

[0034] When the task is to generate a question score for an input question string, the preceding token is the question token that is in the input question string at the preceding time step. That is, in this task, the system generates question token scores for as many time steps as there are question tokens in the input question string, so that each question token in the input question string corresponds to a respective time step. For a given time step, the system selects the question token in the input question string that corresponds to the immediately preceding time step as the preceding question token.

[0035] In either task, for the first time step of decoding, the preceding question token is a predetermined start of sequence token.

[0036] To generate the attention vector 210 for the time step, the attention vector neural network 204 receives a preceding RNN output 216 and the encoded representations 106, i.e., the encoded representations of the passage 102 described above with reference to FIG. 1. The preceding RNN output 216 is an output generated by one of the layers of the decoder RNN during the processing for the preceding time step. For example, the preceding RNN output 216 may be the output of the last layer in the decoder RNN for the preceding time step or the output of the first layer in the decoder RNN for the preceding time step. Depending on the architecture of the decoder neural network, the preceding RNN output 216 may be the same as or different from the updated hidden state of the layer of the decoder RNN after the processing for the preceding time step.

[0037] The attention vector neural network 204 processes the preceding RNN output 216 of the decoder RNN 202 and the encoded representation 106 to generate as output the attention vector 210 for the time step. The attention vector 210 includes a respective attention value for each token in the input text passage.

[0038] An example architecture for the attention vector neural network 204 and the decoder RNN 202 and an example configuration for how the attention vector neural network 204 can be connected to the decoder RNN 202 are described in more detail in Wu, Yonghui, et al, *Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation*, available at <https://arxiv.org/abs/1609.08144>.

[0039] In some cases, to generate an output sequence from the question token scores, for each time step, the system samples from the score distribution over the question tokens in the vocabulary or selects the highest-scoring question token as the question token in the output question string for the time step. The system continues to add question tokens to the output question string until an end-of-sequence token is selected or sampled.

[0040] In some other cases, the system maintains multiple partial output question strings using a left to right beam search decoding. For example, during inference, the system may maintain a beam, i.e., a set of partial hypotheses, i.e., partial output question strings. At each time step, each partial hypothesis in the beam may be expanded with every possible question token and only a predetermined number of likely beams are kept. When an end-of-sequence token is encountered, the partial hypothesis may be removed from the beam, i.e., the system stops adding additional outputs to the partial hypothesis, and added to the set of complete hypotheses. The system can repeat these steps until each partial hypothesis in the beam has

been finalized. The system can then select the final hypothesis from the set of complete hypotheses having the highest sequence score as the hypothesis for the input passage or may output multiple of the finalized hypotheses as possible questions to be answered by the input passage. An example beam search technique is described in Wu, Yonghui, et al, *Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation*, available at <https://arxiv.org/abs/1609.08144>.

[0041] The matching neural network 220 is used by the subsystem 108 when the task being performed is to evaluate how well the input passage answers the input question string. In particular, like the attention vector neural network 204, at each time step the matching vector neural network 220 also receives the preceding RNN output 216 and the encoded representations 106 and processes the preceding RNN output 216 and the encoded representations 106 to generate as output a matching vector 216 for the time step. Unlike the attention vectors, however, the matching vectors generated by the matching neural network 220 are not provided as input to the decoder vector RNN and are instead used by the question scoring engine 104

[0042] In some implementations, the matching vector neural network 220 has the same neural network architecture as the attention neural network 204. Thus, the matching vectors generated by the matching vector neural network 220 will have the same dimensionality as the attention vectors generated by the attention vector neural network 204, i.e., will also have a respective value/entry for each passage token in the input passage. However, because the attention vectors and the matching vectors are used for different purposes by the answer to question system when generating outputs, after training of the subsystem the values of the parameters of the matching vector neural network 220 will differ from those of the attention vector neural network 204. Accordingly, the matching vectors generated by the matching vector neural network 220 will differ from the attention vectors generated by the attention vector neural network 204.

[0043] In particular, the question scoring engine 104 generates the answer score for a given question from the matching vectors over all of the time steps for which the decoder RNN 202 generates an output while scoring the given question. In particular, the question scoring engine 104 can generate the answer score by combining the entries in the matching vectors over all of the time steps. For example, the question scoring engine 104 can compute an average of the entries in the matching vectors over all of the time steps, a weighted average of entries in the matching vectors over all of the time steps, e.g., with entries from later time

steps being given greater weight, or a sum of entries in the matching vectors over all of the time steps.

[0044] In some implementations, the question scoring engine 104 also considers the perplexity for the tokens in the input question string as reflected by the question token scores generated by the decoder RNN 202. For example, the question scoring engine 104 can compute a perplexity score from an average perplexity of the tokens in the input question string as reflected by the scores generated by the decoder RNN 202 for the tokens and can combine, e.g., sum, average, or multiply, the perplexity score with the score determined from the matching vectors to generate the final answer score for the input question string.

[0045] FIG. 3 is a flow diagram of an example process 300 for selecting a passage that provides an answer to a question. For convenience, the process 300 will be described as being performed by a system of one or more computers located in one or more locations.

[0046] The system receives a question (step 302). For example, the question may have been submitted as a search query to an Internet search engine.

[0047] The system receives candidate passages (step 304). For example, the candidate passages may have been extracted from resources identified by search results for the search query, e.g., a predetermined number of highest-ranked search results as determined by the Internet search engine.

[0048] The system determines a respective answer score for each candidate passage (step 306). The answer score for a given candidate passage indicates how well the candidate passage answers the question. In particular, to determine the answer scores, the system provides each candidate passage as input to the answer to question system along with a question string derived from the received question. The answer to question system processes the candidate passage and the question string to generate an answer score for the candidate passage as described above with reference to FIGS. 1 and 2.

[0049] In some implementations, the system has already pre-processed the candidate passages and the question string using the answer to question system and has stored the answer scores for the candidate passages. In these implementations, the system accesses the pre-computed answer scores from memory.

[0050] The system selects a candidate passage as the passage that provides the best answer to the question (step 308) using the answer scores for the candidate passages. For example, the system can select the highest-scoring candidate passage according to the answer scores or can combine the answer scores with scores generated based on other signals to generate a

respective final answer score for each passage and can select the highest-scoring candidate passage according to the final answer scores.

[0051] In some implementations, the system outputs the selected candidate passage to a user as part of a response to the search query. For example, the system can provide the selected candidate passage for presentation to the user as part of a response to the search query, i.e., along with the search results for the search query. For example, the system can provide a formatted presentation of the passage that identifies the passage as an answer to the question along with a presentation of search results for the search query. As another example, the system can receive the question as a voice query and can verbalize all of or part of the selected candidate passage in response to the voice query, e.g., using a text-to-speech system.

[0052] FIG. 4 is a flow diagram of an example process 400 for generating a question that is answered by a passage. For convenience, the process 400 will be described as being performed by a system of one or more computers located in one or more locations.

[0053] The system receives an input passage (step 402).

[0054] The system processes the input passage using the answer to question system to generate an output question string that represents a question that is answered by the input passage (step 404), e.g., as described above with reference to FIGS. 1 and 2.

[0055] FIG. 5 is a flow diagram of an example process 500 for training an encoder neural network and an attention, matching, and decoding subsystem. For convenience, the process 500 will be described as being performed by a system of one or more computers located in one or more locations. For example, an answer to question system, e.g., the answer to question system 100 of FIG. 1, appropriately programmed, can perform the process 500.

[0056] The system receives first training data for training the encoder neural network and the attention, matching, and decoding subsystem (step 502). The first training data includes multiple training passages and, for each training passage, a training question string for which the training passage has been determined to be a high-quality answer, and a training answer score for the training passage – training question string pair.

[0057] The system trains the encoder neural network and the attention, matching, and decoding subsystem on the first training data (step 504). In particular, the system trains the neural network and the subsystem on the first training data to minimize a loss function that is a weighted combination of (i) an average token perplexity loss function and (ii) a pseudo Huber loss function.

[0058] The average token perplexity loss function measures the average per token perplexity for the question tokens in the training question string as reflected by the corresponding question token scores generated by the decoder neural network during the processing of the corresponding question passage. That is, for each time step, the perplexity is measured based on the question token score generated by the decoder neural network for the question token in the training question string at that time step.

[0059] The pseudo Huber loss function measures the loss between the training answer score for the training passage - training question pair and the answer score generated by the question scoring engine for the training passage – training question pair, i.e., the score generated based on the matching vectors generated by the matching vector neural network.

[0060] As described above, in the first training data, each training passage or most of the training passages are associated with a training question for which the training passage has been determined to be a high-quality answer. Accordingly, while training on the first training data, the weight assigned to the average per token perplexity loss function is significantly greater than the weight assigned to the pseudo Huber loss function.

[0061] The system receives second training data (step 506) and trains the encoder neural network and the attention, matching, and decoding subsystem on the second training data (step 508). Like the first training data, the second training data also includes multiple training passages and, for each training passage, a training question and a training answer score for the training passage – training question pair. However, unlike the first training data, the second training data includes training passage – training question pairs that are of varying quality, i.e., some of the training passage are not good answers to the questions they are paired with (as reflected by the training score associated with the pair). Accordingly, so that the neural networks learn to discern between good answers and not-so-good or poor answers, while training on the second training data, the weight assigned to the average per token perplexity loss function is significantly reduced relative to the weight assigned to the pseudo Huber loss function.

[0062] The system can train the encoder neural network and the attention, matching, and decoding subsystem on the first and second training data using conventional machine learning training techniques, e.g., a backpropagation-through-time based training techniques. Once the encoder neural network and the attention, matching, and decoding subsystem have been trained to determine trained values of the parameters of the encoder neural network and the attention, matching, and decoding subsystem, the system can use the encoder neural

network and the subsystem to accurately score input questions or to generate question strings for a given input passage.

[0063] Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non transitory program carrier for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. The computer storage medium is not, however, a propagated signal.

[0064] The term “data processing apparatus” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit). The apparatus can also include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0065] A computer program (which may also be referred to or described as a program, software, a software application, a module, a software module, a script, or code) can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a stand alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub programs, or portions of code. A computer program can be deployed to be

executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0066] As used in this specification, an “engine,” or “software engine,” refers to a software implemented input/output system that provides an output that is different from the input. An engine can be an encoded block of functionality, such as a library, a platform, a software development kit (“SDK”), or an object. Each engine can be implemented on any appropriate type of computing device, e.g., servers, mobile phones, tablet computers, notebook computers, music players, e-book readers, laptop or desktop computers, PDAs, smart phones, or other stationary or portable devices, that includes one or more processors and computer readable media. Additionally, two or more of the engines may be implemented on the same computing device, or on different computing devices.

[0067] The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0068] Computers suitable for the execution of a computer program include, by way of example, can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0069] Computer readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks;

and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0070] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

[0071] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), e.g., the Internet.

[0072] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0073] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination.

Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0074] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0075] Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

CLAIMS

WHAT IS CLAIMED IS:

1. A method comprising:
 - receiving (i) an input text passage comprising a plurality of passage tokens and (ii) an input question string comprising a respective question token at each of a plurality of time steps;
 - processing the input text passage using an encoder neural network, wherein the encoder neural network is configured to receive the input text passage and to process the input text passage to generate a respective encoded representation for each passage token in the input text passage;
 - at each time step:
 - processing a decoder input for the time step using a decoder neural network, wherein the decoder input comprises at least a preceding token in the input question string and a preceding internal state of the decoder neural network, and wherein the decoder neural network is configured to, for each time step, process the decoder input to at least update the internal state of the decoder neural network; and
 - processing the respective encoded representations and a preceding output of the decoder neural network using a matching vector neural network, wherein the matching vector neural network is configured to receive the preceding output of the decoder neural network and the encoded representations and to process the preceding output and the encoded representations to generate a matching vector for the time step; and
 - generating an answer score that indicates how well the input text passage answers a question posed by the input question string based at least on the matching vectors for the time steps.
2. The method of claim 1, further comprising:
 - at each time step:
 - processing the respective encoded representations and the preceding output of the decoder neural network using an attention vector neural network, wherein the attention vector neural network is configured to receive the preceding output of the decoder neural network and the encoded representations and to process the preceding output and the encoded representations to generate an attention vector for the time step, wherein the attention vector includes a respective value for each passage token.

3. The method of claim 2, wherein the attention vector has the same dimensionality as the matching vector.
4. The method of claim 3, wherein the attention vector neural network has the same neural network architecture as the matching vector neural network but different parameter values.
5. The method of any one of claims 2-4, wherein at each time step the decoder input further comprises the attention vector for the time step.
6. The method of any one of claims 1-5, wherein generating the answer score comprises generating a matching vector score by determining an average or a weighted average of the entries of the matching vectors for the time steps.
7. The method of claim 6, wherein the decoder neural network is further configured to, at each time step, generate a set of question token scores that includes a respective score for each question token in a vocabulary of possible question tokens.
8. The method of claim 7, wherein generating the answer score comprises:
 - determining a perplexity score from an average perplexity of the question tokens in the input question string as reflected by the question token scores generated by the decoder neural network for the question tokens in the input question string; and
 - combining the perplexity score and the matching vector score to generate the answer score.
9. The method of any one of claims 1-8, wherein the encoder neural network and the decoder neural network are recurrent neural networks.
10. The method of any one of claims 1-9, wherein the preceding output is an output of one of the layers of the decoder neural network at a preceding time step.
11. The method of any one of claims 1-10, further comprising:
 - determining, based on the answer score, that the input passage answers the question posed by the input passage string; and
 - in response, outputting data derived from the input passage.

12. The method of any preceding claim, wherein the encoder neural network, decoder neural network and matching vector neural network are trained based upon first training data, wherein the first training data comprises:

a plurality of training passages and, for each training passage:

a training question string for which the training passage has been determined to be a high-quality answer, and

a training answer score for each training passage – training question string pair;

wherein training the encoder neural network, decoder neural network and matching vector neural network comprises minimizing a loss function that is a weighted combination of (i) an average token perplexity loss function and (ii) a pseudo Huber loss function.

13. The method of claim 12, wherein the average token perplexity loss function measures the average per token perplexity for the question tokens in the training question string generated by the decoder neural network.

14. The method of claim 12 or 13, further comprising:

training the encoder neural network, decoder neural network and matching vector neural network on second training data, wherein the second training data comprises:

a plurality of training passages and, for each training passage:

a training question string; and


a training answer score for the training passage – training question string pair;

wherein the second training data includes training passage – training question string pairs of varying quality; and

wherein the weight assigned to the average per token perplexity loss function is reduced relative to the weight assigned to the pseudo Huber loss function when training the encoder neural network, decoder neural network and matching vector neural network based upon the second training data.

15. One or more computer storage media storing instructions that when implemented by one or more computers cause the one or more computers to perform the operations of the respective method of any one of claims 1-14.

16. A system comprising one or more computers and one or more storage devices storing instructions that when implemented by the one or more computers cause the one or more computers to perform the operations of the respective method of any one of claims 1-14.

**FIG. 1**




FIG. 2




FIG. 3




FIG. 4




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/057251

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06N3/08
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EP0-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	JUN YIN ET AL: "Neural Generative Question Answering", PROCEEDINGS OF THE WORKSHOP ON HUMAN-COMPUTER QUESTION ANSWERING, 22 April 2016 (2016-04-22), pages 36-42, XP055404542, Stroudsburg, PA, USA DOI: 10.18653/v1/W16-0106 abstract figures 1-3 page 1, line 1 - page 6, line 36 -----	1-16

☐

Further documents are listed in the continuation of Box C.

☐

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

3 April 2018

Date of mailing of the international search report

09/04/2018

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Appeltant, Lennert