




(22) **Date de dépôt/Filing Date:** 2011/02/02
(41) **Mise à la disp. pub./Open to Public Insp.:** 2011/08/08
(45) **Date de délivrance/Issue Date:** 2018/09/04
(30) **Priorité/Priority:** 2010/02/08 (US12/701,788)

(51) **Cl.Int./Int.Cl. G10L 15/06** (2013.01)
(72) **Inventeurs/Inventors:**
SHU, CHANG-QING, US;
SHU, HAN, US;
MERWIN, JOHN M., US
(73) **Propriétaire/Owner:**
ADACEL SYSTEMS, INC., US
(74) **Agent:** MACRAE & CO.

(54) **Titre : MODELE LINGUISTIQUE INTEGRE, AINSI QUE LES SYSTEMES ET PROCEDES CONNEXES**
(54) **Title: INTEGRATED LANGUAGE MODEL, RELATED SYSTEMS AND METHODS**



(57) **Abrégé/Abstract:**

An integrated language model includes an upper-level language model component and a lower-level language model component, with the upper-level language model component including a non-terminal and the lower-level language model component being applied to the non-terminal. The upper-level and lower-level language model components can be of the same or different language model formats, including finite state grammar (FSG) and statistical language model (SLM) formats. Systems and methods for making integrated language models allow designation of language model formats for the upper-level and lower-level components and identification of non-terminals. Automatic non-terminal replacement and retention criteria can be used to facilitate the generation of one or both language model components, which can include the modification of existing language models.

Abstract

An integrated language model includes an upper-level language model component and a lower-level language model component, with the upper-level language model component including a non-terminal and the lower-level language model component being applied to the non-terminal. The upper-level and lower-level language model components can be of the same or different language model formats, including finite state grammar (FSG) and statistical language model (SLM) formats. Systems and methods for making integrated language models allow designation of language model formats for the upper-level and lower-level components and identification of non-terminals. Automatic non-terminal replacement and retention criteria can be used to facilitate the generation of one or both language model components, which can include the modification of existing language models.

INTEGRATED LANGUAGE MODEL, RELATED SYSTEMS AND METHODS

Field of the Invention

[0001] The present invention relates to automatic speech recognition, and more particularly to systems and methods for making language models for speech recognition engines.

Background of the Invention

[0002] Referring to Figure 3, in a typical speech recognition engine 1000, a signal 1002 corresponding to speech 1004 is fed into a front end module 1006. The front end 1006 module extracts feature data 1008 from the signal 1002. The feature data 1008 is input to a decoder 1010, which the decoder 1010 outputs as recognized speech 1012. An application 1014 could, for example, take the recognized speech 1012 as an input to display to a user, or as a command that results in the performance of predetermined actions.

[0003] To facilitate speech recognition, an acoustic model 1018 and a language model 1020 also supply inputs to the decoder 1010. The acoustic model 1018 utilizes the decoder 1010 to segment the input speech into a set of speech elements and identify to what speech elements the received feature data 1008 most closely correlates.

[0004] The language model 1020 assists the operation of the decoder 1010 by supplying information about what a user is likely to be saying. There are two major formats for language models: the finite state grammar (FSG) and the statistical language model (SLM).

[0005] The FSG format typically includes a plurality of predetermined text element sequences. "Text element" as used herein can refer to words, phrases or any other subdivision of text, although words are the most common text elements. To apply an FSG format language model, the decoder 1010 compares the feature data 1008 (also utilizing input from the acoustic model 1018) to each of the text element sequences, looking for a best fit.

[0006] Provided the user actually is speaking one of the predetermined sequences, the FSG format offers relatively high accuracy. However, if the user does not speak one of the sequences, a decoder applying the FSG format will not yield the correct result. Additionally, compiling a suitable list of sequences for a given application can be time and labor intensive. Moreover, to yield acceptable results for complex applications, an FSG format language model must be extremely large, resulting in higher memory and processing demands.

[0007] An SLM format language model, sometimes referred to as an "n-gram" format, is built from a textual corpus by identifying, for each text element (e.g., each word), the probability that the element will be found in proximity with the other text elements. Typically, probabilities are determined for each group of two (bi-gram) or three (tri-gram) text elements, although other quantities can be used. A nominal probability is usually also assigned to text elements groups that do not actually occur in the textual corpus.

[0008] The SLM format allows for the potential recognition of a larger range of user utterances with a relatively small language model. However, the accuracy of the SLM format typically compares unfavorably with the FSG format.

[0009] The concept has been advanced to combine features of both language model formats to mitigate the disadvantages and capitalize on the advantages of each. An example of efforts in this direction can be found in U.S. Patent No. 7,286,978. However, there have been only limited practical attempts at such combinations, and further enhancements and improvements are possible.

Summary of the Invention

[0010] In view of the foregoing, it is an object of the present invention to provide integrated language models, and related systems and methods for making and using the integrated language models.

[0011] According to an embodiment of the present invention, a system for making an integrated language model comprises at least one processor and machine-readable memory. The processor and memory are configured to

execute a language model integration control module adapted to receive user inputs regarding language model integration options and to generate language model modification rules and application rules based thereon, and a language model generation module adapted to generate new and/or modify existing language models based upon the language model generation rules to generate upper-level and lower-level language model components for the integrated language model.

[0012] According to a method aspect of the present invention, a method of making an integrated language model for a speech recognition engine includes identifying a first language model format for an upper-level language model component, identifying a plurality of text elements to be represented by a non-terminal in the upper-level language model component and generating the upper-level language model component including the non-terminal. "Non-terminal," as used herein, generally indicates a term or other marker determined as a result of decoder operation that does not represent an actual text element in a hypothesis, but rather is an intermediate marker that indicates further operations must be applied thereto to determine the corresponding text element or elements. The method of making an integrated language model further includes identifying a second language model format for a lower-level language model component to be applied to the non-terminal of the upper-level language model component and generating the lower-level language model component.

[0013] According to another method aspect of the present invention, a method for identifying text elements to be represented by non-terminals in an integrated language model for a speech recognition engine includes determining a text element replacement criterion allowing automatic identification of the text elements to be represented by the non-terminals within an existing language model or textual corpus, and applying the text element replacement criterion to the existing language model or textual corpus.

[0014] These and other objects, aspects and advantages of the present invention will be better understood in view of the drawings and following detailed description of preferred embodiments.

Brief Description of the Drawings

- [0015]** Figure 1 is a schematic overview of a system for making an integrated language model, according to an embodiment of the present invention;
- [0016]** Figure 2 is a flow diagram of a method for making an integrated language model, according to a method aspect of the present invention; and
- [0017]** Figure 3 is a schematic overview of a typical speech recognition engine.

Detailed Description of Preferred Embodiments

[0018] Referring to Figure 1, according to an embodiment of the present invention, a system 10 for making an integrated language model includes a language model integration control module 12 and a language model generation module 14. Existing statistical language model (SLM) format language models 16 and finite state grammar (FSG) format language models 18 are accessible by the language model generation module 14. An integrated language model 20, including an upper-level language model component 26 having one or more non-terminals, and one or more lower-level language model components 28 for application to the non-terminal(s), is output by the language model generation module 14 and available as an input to the decoder 40 of a speech recognition engine.

[0019] It will be appreciated that speech recognition engines are inherently machine processor-based. Accordingly, the systems and methods herein are realized by at least one processor executing machine-readable code and that inputs to and outputs from the system or method are stored, at least temporarily, in some form of machine-readable memory. However, the present invention is not necessarily limited to particular processor types, numbers or designs, to particular code formats or languages, or to particular hardware or software memory media.

[0020] The language model integration control module 12 is adapted to receive user inputs regarding language model integration options; for instance, which language model format to be applied globally as the upper-level language model component 26, which language model format(s) to be applied to non-terminals as the one or more lower-level language model components 28, and how the upper-level and lower-level language model components are to be generated for application globally and to non-terminals. Based on the user inputs, the language model integration control module 12 determines generation rules for the language model generation module 14 and application rules for the decoder 40.

[0021] The language model generation module 14, based on the generation rules from the language model integration control module 12, makes new and/or modifies the existing SLM and/or FSG format language models 16, 18 to generate the upper-level and lower-level language model components 26, 28 that form the integrated language model 20.

[0022] The integrated language model 20 is applied by the decoder 40 based on the application rules supplied by the language model integration control module 12, such that the upper-level language model component 26 is applied globally and the one or more lower-level language model components 28 are applied to non-terminals.

[0023] It will be appreciated that systems for making an integrated language model falling within the scope of the present invention will not necessarily require the exact components described above. Additionally, it is not necessarily required that all of the various functional modules be executed by the same machine or within a particular time period. For instance, generation and application rules could be generated by a language model integration control module 12 on one computer and then stored and loaded onto a second computer with the language model generation module 14.

[0024] Also, it will be appreciated that the present invention can also include speech recognition engines and related systems that use integrated language models such as those produced by the system 10.

[0025] Referring to Figure 2, according to a method aspect of the present invention, a method for making an integrated language model begins at block 100. At blocks 102 and 104, determinations are made regarding language model formats to be used for the upper-level and lower-level language model components. The following table illustrates basic possible combinations of language model formats.

TABLE 1	
UPPER-LEVEL	LOWER-LEVEL (non-terminals)
FSG	FSG
FSG	SLM
SLM	SLM
SLM	FSG

[0026] It will be appreciated that further combinations can fall within the scope of the present invention other than those above. For example, multiple language model components, of the same or different formats could be applied to non-terminals within the integrated language model. For instance, an FSG format lower-level language model component could be applied to one or more non-terminals and an SLM format lower-level language model component applied to other non-terminals.

[0027] Alternately, different FSG format lower-level language model components could be applied to different non-terminals. Also, multiple layers of lower-level language model components could be applied to non-terminals. For instance, the lower-level language model component applied to non-terminals of the upper-level language model component could have further non-terminals therein, to which another lower-level language model component is applied. The following table illustrates some more complex combinations that are possible, although many additional combinations are possible.

TABLE 2

UPPER-LEVEL	LOWER-LEVEL 1 (non-terminals of Upper Level)	LOWER-LEVEL 2 (non-terminals of Lower Level 1)
FSG	FSG	SLM
FSG	SLM1	FSG1
	SLM2	FSG2
SLM	FSG	FSG
SLM	FSG	-
	SLM	SLM

[0028] At block 106, an upper-level language model component, of the format determined at block 102, is generated for use in the integrated language model. This can include replacing text elements in an existing language model, or vocabulary thereof, with non-terminals. The language model is then compiled or re-compiled, or otherwise re-structured as necessary, with the non-terminals instead of the text elements. This process will generally result in a reduction in size of the upper-level language model component, as well as a reduction in the overall size of the integrated language model.

[0029] For example, “n” different words are represented by a single non-terminal. When each instance of the “n” words is replaced by the non-terminal, it occurs “m” times within a grammar. Accounting for these words without the non-terminal would require a corresponding language model to account for (n x m) word instances, whereas accounting for these words with the non-terminal potentially requires only accounting for (n + m) word instances.

[0030] Various mechanisms can be employed to identify the text elements within the upper-level language model component that are to be replaced with non-terminals. Advantageously, text element replacement criteria can be specified that allow text elements to be automatically identified and replaced with non-terminals. For instance, text element sequences having a definable length can be automatically identified and replaced. For example, if domestic telephone numbers were to be replaced with a non-terminal, the following criteria

could be used to identify telephone numbers within the language model vocabulary:

[0031] <telephone_number> = Loop{<zero>, <one>, ..., <nine>; length==7 or length ==10}.

[0032] Also, a definable text element sequence having a definable value range can be automatically identified and replaced. For example, if compass headings and angle measurements were to be replaced with a non-terminal, the following criteria could be used:

[0033] <direction_degree> = Loop{<zero>, <one>, ..., <nine>; Value>=0 and Value<=360}.

[0034] Additionally, a user could provide a list of text elements to be associated with a given non-terminal rather than, or in addition to, automatic replacement criteria, and the system could search for and replace those text elements with the given non-terminal. Moreover, whether using automatic replacement criteria or a specific list, certain text elements might be replaced with non-terminals that should remain within the language model. Lists of such text elements can be provided by a user, or a set of automatic retention criteria can be applied to automatically identify such text elements. Also, the automatic replacement criteria can further include text element context information to help prevent undesired non-terminal replacements.

[0035] In general, the selection of the upper-level and lower-level language model component formats for the integrated language model, as well as the non-terminals to be employed, will be driven by the nature of the speech recognition application, with a goal of achieving an optimal combination of reduction in grammar size and processing time and increasing recognition accuracy for a given speech recognition application. Predictive modeling and/or experimentation can be used to help measure the achievement of this goal.

[0036] At block 108, lower-level language model components, of the format(s) determined at block 104, are made or modified for lower-level application to the non-terminals. As discussed above, multiple lower-level language model components can be used, and such lower-level language model

components can, themselves, include non-terminals to which one or more additional lower-level language model components are applied.

[0037] The lower-level language model components to be applied to non-terminals can be purpose-built or generated by modifying existing language models. For instance, automatic criteria or list-based mechanisms, similar to those used to introduce non-terminals into the upper-level language model component, can be used in reverse to eliminate text elements that are non-applicable to a given non-terminal.

[0038] At block 110, the upper-level and lower-level language model components are combined in the integrated language model. This does not necessarily, and preferably does not, require a merger of the different language model components, as the upper-level and lower-level language model components will generally be applied separately. The combination can include, however, the naming or otherwise identifying of the components as upper-level and lower-level, and further identifying the non-terminal(s) associated with each lower-level language model component, so that the decoder can properly access the language model components during operation of the speech recognition engine.

[0039] At block 112, application rules for the decoder are generated to support the use of the integrated language model. For example, where the decoder supports selective use of either FSG or SLM format language models as alternate options, the decoder should be modified to allow application of both FSG and SLM format language models, including language models of either format having non-terminals which direct the decoder to language models of either format corresponding to the non-terminals. When using an integrated language model according to the present invention, the decoder should be configured to treat non-terminals as text elements during searching, back tracing and the like, with the upper-level language model component being applied outside of the non-terminals and the lower-level language model component applied within.

[0040] At block 114, the method ends, although the method can be repeated as necessary to create or modify further integrated language models.

Additionally, it will be appreciated that all the method steps enumerated above are not necessary for every execution of the method for making an integrated language model. Also, the steps are not necessarily limited to the sequence described, and many steps can be performed in other orders, in parallel, or iteratively. Furthermore, the present invention also encompasses methods of speech recognition engine use including the application of language models like those generated by the method for making an integrated language model.

[0041] In general, the foregoing description is provided for exemplary and illustrative purposes; the present invention is not necessarily limited thereto.

Rather, those skilled in the art will appreciate that additional modifications, as well as adaptations for particular circumstances, will fall within the scope of the invention as herein shown and described and the claims appended hereto.

What is claimed is:

1. A method of making an integrated language model for a speech recognition engine, the method comprising:
 - identifying a first language model format for an upper-level language model component;
 - identifying a plurality of text elements to be represented by a non-terminal in the upper-level language model component;
 - generating the upper-level language model component including the non-terminal;
 - identifying a second language model format for a lower-level language model component to be applied to the non-terminal of the upper-level language model component; and
 - generating the lower-level language model component;wherein generating the upper-level language model component includes modifying an existing language model; and
wherein identifying the plurality of text elements to be represented by the non-terminal in the upper-level language model component includes applying an automatic text element replacement criterion to the existing language model.
2. The method of claim 1, wherein the first language model format is a finite state grammar format.
3. The method of claim 2, wherein the second language model format is a statistical language model format.
4. The method of any one of claims 1 to 3, wherein a plurality of non-terminals are included in the upper-level language model component.

5. The method of claim 4, wherein a plurality of lower-level language model components are applied to the plurality of non-terminals.
6. The method of claim 5, wherein the plurality of lower-level language model components include at least two language model components having different language model formats.
7. The method of claim 5, further comprising identifying to which of the plurality of non-terminals each of the plurality of lower-level language model components is to be applied during operation of the speech recognition engine.
8. The method of claim 1, wherein the lower-level language model component also includes a non-terminal, and the method further comprises identifying a third language model format for an additional lower-level language model component to be applied to the non-terminal of the lower-level language model component, and generating the additional lower-level language model component.
9. The method of claim 1, wherein the existing language model is a finite state grammar format language model.
10. The method of claim 1, wherein the existing language model is a statistical language model format language model.
11. The method of claim 1, wherein the text element replacement criterion is to replace text element sequences having a definable length.
12. The method of claim 11, wherein the definable length is a number of digits.

13. The method of claim 1, wherein the text element replacement criterion is to replace text element sequences having a definable value range.
14. The method of claim 1, wherein a plurality of automatic text element replacement criteria are applied.
15. The method of claim 1, wherein identifying the plurality of text elements to be represented by the non-terminal in the upper-level language model component further includes applying an automatic text element retention criterion to the existing language model.
16. The method of any one of claims 1 to 8, wherein generating the lower-level language model component includes modifying the existing language model.
17. The method of claim 16, wherein modifying the existing language model includes automatically eliminating text elements that are determined not to be relevant to the non-terminal of the lower-level language model, the determination being based on a non-terminal replacement criterion applied to identify the plurality of text elements to be represented by the non-terminal of the lower-level language model.
18. The method of any one of claims 1 to 17, further comprising generating instructions for a speech recognition engine decoder to apply the upper-level and lower-level language model components during operation.
19. The method of any one of claims 1 to 18, wherein the text elements are words.

1/3




FIG. 1

10




FIG. 2




FIG. 3 (Prior Art)

