



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))

TITLE OF APPLICATION

Low-Dimensional Real-Time Concatenative Speech Synthesizer

STATEMENT OF GOVERNMENT RIGHTS

5 The invention was made with US Government support under Contract No. DC002852 awarded by the National Institutes of Health. The Government has certain rights in the invention.

BACKGROUND

10 The present invention is related to the field of real-time speech synthesis, used for example as a substitute for voice-generated speech by persons who have lost control over their vocal apparatus due to disease.

 Different types of speech synthesis techniques and systems are known, including text-to-speech systems, articulatory synthesizers, and acoustic or "formant" synthesizers. Text-to-
15 speech systems are generally not suited for conversation-like speech, because of the requirement that users type out all words and the limitation of fixed or "canned" words and relative inexpressiveness. The requirement for typing may be especially problematic for users who are paralyzed or otherwise limited in the speed at which they can generate user input to the system. Both articulatory and formant synthesizers generally require adjustment of a high
20 number of parameters for accurate operation, i.e., to achieve desired speech quality, and thus are viewed as high-dimensional systems from a control perspective. If not properly adjusted, these systems may provide relatively low speech quality, for example due to difficulty with rendering consonants.

SUMMARY

25 It is desired to provide real-time speech synthesis that has high speech quality while requiring generally low-dimensional user input and control. With such characteristics, a synthesis technique may be more widely adopted and utilized, enhancing the ability of different populations of people to engage in quality conversational speech.

30 Presently disclosed methods and apparatus have two aspects that promote the above goals. In one aspect, usability is enhanced by a two-dimensional graphical interface for accepting user input, where the interface represents a complete set or "space" of phonemes from which speech can be synthesized. The two-dimensional layout reflects vocal tract

morphology and use and thus the interface may be learned relatively easily and used efficiently and effectively. In another aspect, a modified diphone synthesis technique is used that provides high quality synthesized speech with relatively low computing demand at rendering time. High quality is provided in part by enabling users to control duration of sounds as is done in vocal speech for emphasis or other effects.

More particularly, in one aspect a method is disclosed for operating a computerized device to provide real-time synthesis of speech based on user input. The method includes presenting a graphical user interface having a low-dimensional representation of a multi-dimensional phoneme space, for example a first dimension representing degree of vocal tract constriction and voicing, and a second dimension representing location in a vocal tract. The method further includes receiving user input via the interface and translating received user input into a sequence of phonemes, and rendering the sequence of phonemes on an audio output device.

In some embodiments, the first dimension further represents single versus diphthong vowel sounds.

In some embodiments, the low-dimensional representation has a disk-shaped layout, and the first and second dimensions are selected from a radial dimension and an angular dimension of the disk-shaped layout. In one specific example, the radial dimension represents the first dimension of the phonemic space, and the angular dimension represents the second dimension of the phonemic space.

In some embodiments, the computerized device uses a touch sensitive display and the user input is received as sets of swipes between segments of the interface that represent corresponding phonemes. Each set of swipes corresponds to a unitary utterance of one or more syllables (e.g., a word), with the user maintaining unbroken contact with the touch sensitive display for each set. The sets of swipes are separated by intervals corresponding to silence in which the user has broken contact with the touch sensitive display.

There may be alternative modes of rendering, which may be user selectable. In one mode, individual utterances are rendered upon completion of their corresponding sets of swipes. The durations of rendered phonemes may be fixed or may correspond to duration of user dwell on the corresponding segments of the low-dimensional representation (i.e., how long a user's finger or other pointer rests on each phoneme-specific segment).

In another mode, individual syllables are rendered upon completion of their corresponding individual swipes of the sets of swipes. The durations of rendered phonemes

may correspond to respective durations of user dwell on the corresponding segments of the low-dimensional representation.

In yet another mode, individual phonemes are rendered at respective times of user selection. The phonemes may be rendered at either the beginning or ending of user selection, i.e., when a user first lands on a segment or when a user leaves a segment. Alternatively, the phonemes may be rendered throughout respective durations of selection by the user as the user dwells on the respective segments of the low-dimensional representation.

In some embodiments, rendering is performed in a user-selected one of a plurality of available rendering modes, the rendering modes including (1) rendering individual utterances upon completion of their corresponding sets of swipes, (2) rendering individual syllables upon completion of their corresponding individual swipes of the sets of swipes, and (3) rendering individual phonemes at respective times of user selection by dwelling on corresponding segments of the low-dimensional representation.

In some embodiments, rendering includes accessing a stored library of post-processed speech samples including variable-length versions of diphones used to match durations specified by the user input.

In another aspect, a method is disclosed for operating a computerized device to provide real-time synthesis of speech using a particular real-time synthesis technique. The method includes maintaining a library of prerecorded samples of diphones, each diphone sample including a transition between a respective pair of phonemes, the diphone samples being organized into diphone groups at least some of which include multiple examples of transitions between a respective single pair of phonemes. The different examples include different durations. The method further includes continually receiving a time-stamped sequence of phonemes to be synthesized, and based on successive pairs of the phonemes selecting a sequence of respective diphone groups and corresponding diphone time stamps. The method further includes continually generating a sequence of production diphones from the sequence of diphone groups by (1) identifying a best diphone within each group, the best diphone being part of the optimal path most closely matching each individual diphone duration calculated from successive diphone time stamps as well as minimizing artifacts resulting from stitching dissimilar sound samples, and (2) placing each identified best diphone into a production buffer with a corresponding diphone time stamp. The diphones are sequentially rendered from the production buffer at respective times corresponding to the diphone time stamps.

In some embodiments, the groups of diphones are diphone lists, and continually generating the sequence of production diphones includes (1) placing successive diphones lists on a diphone lists queue, and (2) computing a desired trajectory of rendering across a diphone in the production buffer and diphone lists in the diphone lists queue, the computing
5 identifying the respective best diphones for the diphone lists.

In some embodiments, the diphones have respective GO/noGO flags used to selectively delay production of the diphones based on a mode selected for rendering, the mode being selected from a set of modes that differ in timing of initiating rendering.

In some embodiments, rendering includes switching to silence in cases including (1) a
10 missing diphone sample for which there is no recorded speech sample, (2) a current speech sample reaching an end of the current production buffer when a current production value is not already silence, signifying that a current phoneme production duration is longer than is available from recorded samples, and (3) absence of valid diphone transitions between two diphones, as occurs for a planned phoneme duration longer than is available from recorded
15 samples.

In some embodiments, the library includes post-processed samples of collected speech samples including linking points for linking diphones together and extended samples for longer-duration renderings of recorded diphone samples.

Also in some embodiments, receiving the sequence of phonemes includes (1)
20 presenting a graphical user interface to a user, the graphical user interface having a low-dimensional representation of a multi-dimensional phoneme space, a first dimension representing degree of vocal tract constriction and voicing, a second dimension representing location in a vocal tract, and (2) receiving user input via the interface and translating received user input into a sequence of phonemes.

25 Computerized devices are also disclosed that include respective processors, input/output circuitry, and memory containing instructions that are executed by the processors to cause the computerized devices to perform corresponding methods including methods as outlined above.

30 BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features and advantages will be apparent from the following description of particular embodiments of the invention, as illustrated in the

accompanying drawings in which like reference characters refer to the same parts throughout the different views.

Figure 1 is a block diagram of a speech synthesis system;

Figure 2 is a block diagram of a real-time synthesizer computer;

Figure 3 is a diagram of a graphical user interface with a representation of a phonemic space;

Figure 4 is a diagram of the graphical user interface with an example set of user-generated swipes corresponding to an utterance;

Figure 5 is a flow diagram of a user-interface aspect of operation;

Figure 6 is a schematic diagram illustrating real-time speech synthesis;

Figure 7 is a flow diagram for a real-time speech synthesis aspect of operation;

Figure 8 is a block diagram of a library generator computer;

Figure 9 is a flow diagram of a library generating aspect of operation;

Figures 10-12 are schematic diagrams illustrating aspects of library generating operation.

DETAILED DESCRIPTION

Figure 1 shows a system including two computers shown as a library generator computer 10 and a real-time synthesizer computer 12. Although shown as two separate computers, other embodiments may employ only one computer or may employ more than two computers. In this description, "computer" refers to components conventionally called "computers", such as personal computers, tablet computers, server computers, etc., as well as to computerized devices including computer technology and functionality, such as smart phones, smart televisions, and devices that may be purpose-built or purpose-configured as platforms for the functionality described herein. As one specific example, the real-time synthesizer computer 12 may be realized as a purpose-built consumer device tailored for use by impaired users, very young users, etc.

In operation, the library generator computer 10 receives sample speech 14 and performs processing to generate a library 16 of speech samples that is provided to the real-time synthesizer computer 12. The sample speech 14 may be received directly from a human speaker via a microphone, or it may be received in a pre-recorded form. In the present description it is assumed, without loss of generality, that the library generator computer 10 captures live spoken speech. The captured speech samples are post-processed in a manner

described below to generate the library 16 whose contents are particularly suited for use in the processing performed by the real-time synthesizer computer 12.

The real-time synthesizer computer 12 uses the contents of the library 16, responsive to user input 18, to generate synthesized speech 20. As described more below, the user's control is provided by a selector type of user input device. In one embodiment, a touch sensitive display is used. The user selects representations of a sequence of phonemes for words that the user wishes to have synthesized. The real-time synthesizer computer 12 interprets the user selections and performs the speech synthesis in a real-time manner using the contents of the library 16. Specifically, the real-time synthesizer computer 12 generates a stream of samples of sound pairs called "diphones", and these samples are converted to the audio signals that drive an audio output device. In the present description, the term "rendering" is used to describe generating sound from a stream of data representing sound samples such as diphones.

In the present description the terms "phoneme" and "diphone" may each refer to different aspects or representations depending on context. For example, "phoneme" generally refers to a "sound", i.e., an audible pressure disturbance in air or other medium. In the disclosed system and methods, both actual sounds and certain representations or analogs of the sounds are described. Thus the term "phoneme" may refer to an actual rendered phonemic sound, or to an analog waveform representing such a phonemic sound, or to a collection or unit of digital data representing analog samples of such a phonemic sound, or to data that merely identifies one phonemic sound as distinct from another (e.g., addresses or other unique identifiers of respective units of data in the library 16 for corresponding phonemic sounds). The term "diphone" is used in a similar manner. The actual meanings will be clear in context to those skilled in the art.

Figure 2 shows the organization of the real-time synthesizer computer 12. It includes a memory 30, processor(s) 32, and input/output (I/O) interface circuitry 34. The I/O interface circuitry 34 provides electrical and functional interfaces to a display (DISP) 36 and a speaker (SPKR) 38. As previously indicated, in one embodiment the display 36 is a touch sensitive display and thus serves as both a user input device as well as a display device for displaying graphical output to the user. Although not shown, the real-time synthesizer computer 12 generally also includes non-volatile secondary storage such as a disk drive, Flash memory array, etc.

The memory 30 stores computer program instructions that are executed by the processors 32 to cause the real-time synthesizer computer 12 to function in a corresponding manner. In particular, as generally known the instructions include instructions of an operating system (O/S) 40 and one or more application programs, in this case a real-time synthesizer application (R/T SYNTH APP) 42. In the illustrated embodiment, the real-time synthesizer application 42 includes two major components, a synthesizer graphical user interface (SYNTH GUI) 44 and a real-time concatenative synthesizer (R-T CONCAT SYNTH) 46. Figure 2 shows the program instructions in the memory 30 from which the instructions are retrieved and executed by the processors 32. The instructions may also be stored on a separate computer-readable storage medium from which they are retrieved and loaded into the memory 30, as generally known in the art.

In operation, the real-time synthesizer computer 12 performs user-directed speech synthesis by executing the instructions of the synthesizer GUI 44 and real-time concatenative synthesizer 46. The synthesizer GUI 44 is responsible for displaying a synthesizer interface pattern to the user and accepting and user touch input representing selection of phonemes. The real-time concatenative synthesizer 46 operates in response to the stream of user-selected phonemes to synthesize and render corresponding speech via the speaker 38. Details of these operations are provided more below.

Figure 3 is a graphical synthesizer interface 50 as presented to a user on the display 36 (Figure 2) and used as described below to provide user selection of a stream or sequence of phonemes that represent a sequence of spoken words or utterances. The synthesizer interface 50 is a disk-shaped graphic having an arrangement of English alphabet letters representing discrete sounds or phonemes. The synthesizer interface 50 provides a two-dimensional representation of the sound space or phoneme space that is available for speech synthesis. In the case of the disk shape as shown, the two graphical dimensions are the radial and angular dimensions, and the two corresponding phonemic dimensions are degree of vocal tract constriction (i.e., voicing and stops) and vocal tract location (i.e., shaping of oral cavity for distinct phonemes).

In the radial dimension, the synthesizer interface 50 is divided into four concentric rings, and each ring is divided at regular angles into wedge-shaped segments 52. The rings progress from diphthong vowels through hard consonants, and the angularly-separated segments 52 of each ring progress through different vocal tract shapes/locations for different phonemes of the ring-specific type. The innermost ring contains five segments 52

representing the vowel diphthongs *aw*, *ou*, *oi*, *ay*, *iy* as shown. The next ring contains nine segments 52 for single vowel sounds represented by *ah*, *uh*, *er*, ... as shown. The next outer ring contains twelve segments 52 for the relatively soft consonants *d*, *L*, *zsh*, *j*, etc., and the outermost ring contains twelve segments 52 for the corresponding harder consonants *t*, *y*, *sh*, *ch*, etc. It will be appreciated that many radially adjacent consonant sounds are similar except for differences in the stopping or voicing. As an example, the sounds *t* and *d* differ primarily by the voicing of *d*.

In use, a user first places a selector (e.g., a finger or stylus on a touch sensitive display, or a user-controlled electronic pointer) on the segment 52 for the beginning sound of an utterance, and then traces or "swipes" among the segments 52 for the different phonemes of the utterance, at which point the user lifts the selector. Each swipe has a radial component for vowel-consonant transitions and an angular component for the location or sound-shaping transition. It will be appreciated that many syllables are specified by single swipes having either inward or outward movement, and others are specified by a pair of swipes, one from a consonant to a vowel and then another from the vowel to a consonant. Examples are given below. Additionally, there may be different ways of performing real-time speech rendering based on the pattern and especially timing of user swiping activity, and some embodiments may support multiple modes and some manner of selecting among the modes. These are also described more below.

Figure 4 illustrates an example of use, specifically a pattern used to specify the word *sop*. A user starts at the *s* located at about 11:00 on the outermost ring, swipes inwardly to the *ah* sound, then back outwardly and around to the *p* sound at about 7:00 on the outermost ring. The user maintains contact with the touch sensitive display while swiping, and breaks contact when the utterance is complete, which in this example is upon reaching the *p*. In the case of a user-controlled electronic selector such as an on-screen pointer, some other way of indicating termination of an utterance may be used such as dragging the selector to an area of no segments 50, for example.

Figure 5 presents a general flow of operation of the real-time synthesizer computer 12 with respect to the control of speech synthesis via the user interface 50 as described above. It will be appreciated that this operation is provided by execution of the synthesizer GUI 44 of the real-time synthesizer application 42 (Figure 2).

At 60, the real-time synthesizer computer 12 presents a graphical user interface to the user, the graphical user interface having a low-dimensional representation of a multi-

dimensional phoneme space, a first dimension representing a degree of constriction of the vocal tract and voicing, and a second dimension representing location in the vocal tract. As explained above, the disk-shaped interface 50 is an example of such a low-dimensional representation, using radial and angular graphical dimensions.

At 62, the real-time synthesizer computer 12 receives user input via the graphical user interface and translates the received user input into a sequence of phonemes. In the touch sensitive display embodiment described herein, user input is provided as an ordered set of points on the display 36 which map to segments of the interface 50 and their corresponding phonemes. The timing of user movement among the points is also preserved by assigning corresponding time stamps. Successive sets of points specify diphones for which dihone samples can be obtained from the library 16, as described more below.

At 64, the real-time synthesizer computer 12 renders the user-specified sequence of phonemes on an audio output device. This process involves smoothly stitching together successive diphones and using the time stamps to selectively shorten or lengthen syllables, which can help in providing emphasis or other cues. Details of this processing are given below.

Figure 6 is a schematic depiction of the processing involved in steps 62-64 of Figure 5. The processing flow is right to left. In the illustrated example the user is specifying the word *sop* by swiping between the *s*, *ah*, and *p* sounds at respective times t_1 , t_2 and t_3 .

Phonemes and diphones are shown with respective GO and tx fields:

tx: Timestamp of target phoneme

GO/noGO: A flag or signal for whether the phoneme is to be rendered. Use of GO allows for delayed production, i.e., a listed or queued phoneme is only processed for production when its GO flag is set. This can be useful for different modes as described below.

When a new target phoneme 66 is selected (e.g., /p/ at t_3), a new dihone 68 is identified by pairing the new dihone 68 with the immediately preceding phoneme (e.g., /ah-/p/). A new corresponding dihone list 70 of available recorded samples of this the dihone 68 is then created. The dihone list 70 is added to a dihone lists queue 72, which operates in a first-in, first-out (FIFO) fashion. The dihone list 70 includes multiple versions of the dihone 68 having different durations.

Also at this point, an optimal path or trajectory is computed at 74 across the current production 76 (the diphone $/\#-s/_1$ at t_1) and the elements within the diphone lists queue 72 (versions of $/s-ah/$ at t_2 and $/ah-p/$ at t_3). Each possible path between the current production 76 and each of the elements within the diphone list queue 72 is assigned a cost that captures both its similarity to the desired duration of syllables, as specified by the user's dwelling on segments 52, as well as a smooth stitching together of successive sound samples from the library 16 to minimize artifacts like audible differences between the same phoneme productions in two different audio samples as well as pops or other similar staccato noise. The optimal path is selected that minimizes the estimated cost. As shown, stitching uses the concept of a "connecting sample", i.e., a portion of a current sound that closely matches a portion of a next sound, so that transitioning playback between the two sounds at the connecting sample minimizes undesirable abrupt noises. When the current sample reaches the connecting sample, and there is at least one GO element in the queue 72, then the first diphone list is pulled from the queue 72, and its "winning" element is entered into the current production buffer 76. The winning element is the one that is part of the optimal path as computed at 74. In the illustrated example, the winning element for the $/s-ah/$ list is $/s-ah/_1$, and the winning element for the $/ah-p/$ list is $/ah-p/_2$. Also at this point, the current sample and connecting sample values are reset to those determined by the current optimal trajectory.

Actual sound rendering is performed by transferring sounds from the current production buffer 76 to a sound buffer 78, and then converting the sound samples to an analog waveform which is played out on an audio device such as the speaker 38. The transfer can be done in response to the size of the sound buffer decreasing below a pre-specified threshold. A new sound segment is pulled from the current production buffer 76 and placed in the sound buffer 78, then a current sample counter is increased.

There may be separate special processing for special cases. One type of special processing is "switch to silence", which includes modifying the diphone lists queue 72 by inserting in an appropriate place a silence phoneme, and re-computing the optimal trajectory. Switch to silence may be performed in the following cases:

1. There are no recorded samples for one diphone (missing diphone sample)
2. The current sample reaches the end of the current production buffer and the current production is not silence already (effective current phoneme production duration longer than available from samples)

3. There are no valid diphone transitions between two diphones (e.g. planned phoneme duration longer than available from samples)

Figure 7 presents the above-described processing as a flow of operations.

At 80, the real-time synthesizer computer 12 maintains a library of prerecorded samples of diphones, each diphone sample including a transition between a respective pair of phonemes, the diphone samples being organized into diphone groups at least some of which include multiple examples of transitions between a respective single pair of phonemes. In one embodiment the diphone groups are realized as diphone lists 70 as shown in Figure 6 and described above.

At 82, the real-time synthesizer computer 12 continually receives a time-stamped sequence of phonemes to be synthesized, and based on successive pairs of the phonemes, it selects a sequence of respective diphone groups having corresponding diphone time stamps. This corresponds to items 66, 68 and 70 in the processing of Figure 6.

At 84, the real-time synthesizer computer 12 continually generates a sequence of production diphones from the sequence of diphone groups by (1) identifying a best diphone within each group, the best diphone being part of the optimal path most closely matching each diphone duration calculated from successive diphone time stamps as well as minimizing artifacts like audible differences between the same phoneme productions in two different audio samples as well as pops or other similar staccato noise resulting from stitching dissimilar sound samples between successive diphones, and (2) placing each identified best diphone into a production buffer with a corresponding diphone time stamp. This operation corresponds to items 72 through 76 in the processing of Figure 6.

At 80, the real-time synthesizer computer 12 sequentially renders the diphones from the production buffer at respective times corresponding to the diphone time stamps. This operation corresponds to item 78 in the processing Figure 6.

Figure 8 shows the organization of the library generator computer 10. It includes a memory 90, processor(s) 92, and input/output (I/O) interface circuitry 94. The I/O interface circuitry 94 provides electrical and functional interfaces to a microphone (MIC) 96 and a storage device (ST DEV) 98 on which the library 16 is stored.

The memory 90 stores computer program instructions that are executed by the processors 92 to cause the library generator computer 10 to function in a corresponding manner. In particular, as generally known the instructions include instructions of an operating

system (O/S) 100 and one or more application programs, in this case a library generator application (LIB GEN APP) 102. In the illustrated embodiment, the library generator application 102 includes two major components, a sample recording component 104 and a post processing component 106.

Figure 8 shows the program instructions in the memory 90 from which the instructions are executed. The instructions may also be stored on a separate computer-readable storage medium from which they are retrieved and loaded into the memory 90, as generally known in the art.

In operation, the library generator computer 10 collects samples of speech via the microphone 96, stores the samples in association with data describing the phonemic content of the samples, and performs certain post-processing to generate the library 16 having content that is used in real-time concatenative speech synthesis such as described above with reference to Figure 6. Details of these operations are provided below.

A complete usable library minimally includes: a) samples for each CV (consonant-vowel) and VC (vowel-consonant) phoneme sequence, typically produced in a CVC or CVCVC sequence; and b) samples for each individual phoneme produced in isolation, typically a long utterance of each phoneme (this is not required for phonemes like 'w' which are never produced in isolation). Ideally the library may also include several utterances for each sample, as well as other commonly used transitions in the target language (e.g. consonant clusters like /F/-/L/, /B/-/R/, etc.) Multiple utterances provide the system with more options for generating smooth transitions during real-time synthesis.

Figure 9 presents a high-level flow of processing performed by the library generator computer 10 to generate the contents of the library 16.

At 110, the phonemic segments from the speech samples are labelled. Each recorded sample is automatically pitch-marked, and individual phonemes are manually labeled (context-independent segments only). The labeled segments representing the 'stable' portion of each phoneme (for glides the 'stable' portion is considered the beginning of the production; for stops the 'stable' portion is considered the closure just before the release). Context-independent segments are expected to be relatively similar across all the instances of each phoneme.

At 112, selected phonemic segments are extended in duration. Samples that represent individual phonemes produced in isolation (C or V productions) are used to extend the duration of the rest of the utterances of the corresponding phonemes as long as possible. Each

diphone sample is similarly extended by using the extended phoneme for the second phoneme in the diphone (e.g. /s-ah/ uses the original /s/ phoneme and the extended /ah/ phoneme).

At 114, optimal diphone transitions are pre-computed and stored along with the diphones. For each pair of matching extended segments (those sharing an intermediate phoneme, e.g. /s-ah/ and /ah-p/), and for each target time t , this pre-computing identifies an optimal way to link these two diphones that results in an intermediate phoneme (e.g. /ah/) with total duration t . This optimization is also computed constraining the linking point (n_1) to be beyond any arbitrary value (n) in order to be able to link two diphones in real time when some portion of the first diphone may already have been produced.

Figures 10-12 provide graphical examples of the operations of Figure 9 using sound waveforms representing phonemes and diphones.

Figure 10 illustrates the labelling of step 110.

Figure 11 illustrates the phonemic segment extension of step 112. Mathematically, this operation can be described as finding n_1, n_2 that minimize a value $D(n_1, n_2) - \lambda_1 n_1 + \lambda_2 n_2$. $D(n_1, n_2)$ is a difference (RMS) between the n_1 pitch-synchronous waveform of the first sample and the n_2 pitch-synchronous waveform of the second sample (after amplitude normalization up to a 1:2 factor). The lambda factors attempt to result in relatively long 'extended diphone'.

Figure 12 illustrates the pre-computation at step 114. Mathematically, the operation at 114 can be described as finding, for each t and n , n_1, n_2 that minimize a value $D(n_1, n_2) + \lambda_0(t_0 - t)$, subject to $t_0 > t$ & $n_1 > n$. This particular form of optimization cost function may be used because it allows a fast and efficient computation of all of the desired solutions using a sorted list of optima.

While various embodiments of the invention have been particularly shown and described, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the scope of the invention as defined by the appended claims.

CLAIMS

What is claimed is:

1. A method of operating a computerized device to provide real-time synthesis of speech based on user input, comprising:

presenting a graphical user interface having a low-dimensional representation of a multi-dimensional phoneme space, a first dimension representing degree of vocal tract constriction and voicing, a second dimension representing location in a vocal tract;

receiving user input via the interface and translating received user input into a sequence of phonemes; and

rendering the sequence of phonemes on an audio output device.

2. The method of claim 1, wherein the first dimension further represents single versus diphthong vowel sounds.

3. The method of claim 1, wherein the low-dimensional representation has a disk-shaped layout, and the first and second dimensions are selected from a radial dimension and an angular dimension of the disk-shaped layout.

4. The method of claim 3, wherein the first dimension is the radial dimension, and the second dimension is the angular dimension.

5. The method of claim 1, wherein the computerized device uses a touch sensitive display and the user input is received as sets of swipes between segments of the low-dimensional representation representing corresponding phonemes, each set of swipes corresponding to a unitary utterance of one or more syllables and having unbroken user contact with the touch sensitive display, the sets of swipes separated by periods corresponding to silence and having broken user contact with the touch sensitive display.

6. The method of claim 5, wherein individual utterances are rendered upon completion of their corresponding sets of swipes.

7. The method of claim 6, wherein the durations of rendered phonemes corresponds to respective durations of user dwell on the corresponding segments of the low-dimensional representation.
8. The method of claim 5, wherein individual syllables are rendered upon completion of their corresponding individual swipes of the sets of swipes.
9. The method of claim 8, wherein the durations of rendered phonemes corresponds to respective durations of user dwell on the corresponding segments of the low-dimensional representation.
10. The method of claim 5, wherein individual phonemes are rendered at respective times of user selection by dwelling on corresponding segments of the low-dimensional representation.
11. The method of claim 10, wherein the phonemes are rendered at respective end times of selection as respective new swipes away from the phonemes are initiated.
12. The method of claim 10, wherein the phonemes are rendered throughout respective durations of selection by the user as the user dwells on the respective segments of the low-dimensional representation.
13. The method of claim 5, wherein rendering is performed in a user-selected one of a plurality of available rendering modes, the rendering modes including (1) rendering individual utterances upon completion of their corresponding sets of swipes, (2) rendering individual syllables upon completion of their corresponding individual swipes of the sets of swipes, and (3) rendering individual phonemes at respective times of user selection by dwelling on corresponding segments of the low-dimensional representation.
14. The method of claim 1, wherein rendering includes accessing a stored library of post-processed speech samples including variable-length versions of diphones used to match durations specified by the user input.

15. A method of operating a computerized device to provide real-time synthesis of speech, comprising:

maintaining a library of prerecorded samples of diphones, each diphone sample including a transition between a respective pair of phonemes, the diphone samples being organized into diphone groups at least some of which include multiple examples of transitions between a respective single pair of phonemes;

continually receiving a time-stamped sequence of phonemes to be synthesized, and based on successive pairs of the phonemes selecting a sequence of respective diphone groups and corresponding diphone time stamps;

continually generating a sequence of production diphones from the sequence of diphone groups by (1) identifying a best diphone within each group, the best diphone being part of the optimal path most closely matching each individual diphone duration calculated from successive diphone time stamps as well as minimizing artifacts resulting from stitching dissimilar sound samples, and (2) placing each identified best diphone into a production buffer with a corresponding diphone time stamp; and

sequentially rendering the diphones from the production buffer at respective times corresponding to the diphone time stamps.

16. The method of claim 15, wherein the groups of diphones are diphone lists, and wherein continually generating the sequence of production diphones includes:

placing successive diphones lists on a diphone lists queue; and

computing a desired trajectory of rendering across a diphone in the production buffer and diphone lists in the diphone lists queue, the computing identifying the respective best diphones for the diphone lists.

17. The method of claim 15, wherein the diphones have respective GO/noGO flags used to selectively delay production of the diphones based on a mode selected for rendering, the mode being selected from a set of modes that differ in timing of initiating rendering relative to respective groups of diphones for syllables or multi-syllable utterances.

18. The method of claim 15, wherein rendering includes switching to silence in cases including (1) a missing diphone sample for which there is no recorded speech sample, (2) a current speech sample reaches an end of the current production buffer and a current

production value is not already silence, signifying that a current phoneme production duration is longer than is available from recorded samples, and (3) absence of valid diphone transitions between two diphones, as occurs for a planned phoneme duration longer than is available from recorded samples.

19. The method of claim 15, wherein the library includes post-processed samples of collected speech including linking points for linking diphones together and extended samples for longer-duration renderings of recorded diphone samples.

20. The method of claim 15, wherein the receiving the sequence of phonemes includes (1) presenting a graphical user interface to a user, the graphical user interface having a low-dimensional representation of a multi-dimensional phoneme space, a first dimension representing degree of vocal tract constriction and voicing, a second dimension representing location in a vocal tract, and (2) receiving user input via the interface and translating received user input into a sequence of phonemes.

21. A computerized device, comprising:

one or more processors;

input/output circuitry for connecting the computerized device to a user interface device; and

memory coupled to the processors and the input/output circuitry, the memory storing computer program instructions that, when executed by the processors, cause the computerized device to perform a method for providing real-time synthesis of speech based on user input, the method including:

presenting a graphical user interface having a low-dimensional representation of a multi-dimensional phoneme space, a first dimension representing degree of vocal tract constriction and voicing, a second dimension representing location in a vocal tract;

receiving user input via the interface and translating received user input into a sequence of phonemes; and

rendering the sequence of phonemes on an audio output device.

22. A computerized device, comprising:

one or more processors;
input/output circuitry for connecting the computerized device to an audio output device; and

memory coupled to the processors and the input/output circuitry, the memory storing computer program instructions that, when executed by the processors, cause the computerized device to perform a method for providing real-time synthesis and rendering of speech, the method including:

maintaining a library of prerecorded samples of diphones, each diphone sample including a transition between a respective pair of phonemes, the diphone samples being organized into diphone groups at least some of which include multiple examples of transitions between a respective single pair of phonemes;

continually receiving a time-stamped sequence of phonemes to be synthesized, and based on successive pairs of the phonemes selecting a sequence of respective diphone groups and corresponding diphone time stamps;

continually generating a sequence of production diphones from the sequence of diphone groups by (1) identifying a best diphone within each group, the best diphone being part of the optimal path most closely matching each individual diphone duration calculated from successive diphone time stamps as well as minimizing artifacts resulting from stitching dissimilar sound samples, and (2) placing each identified best diphone into a production buffer with a corresponding diphone time stamp; and

sequentially rendering the diphones from the production buffer, via the audio output device, at respective times corresponding to the diphone time stamps.

1/9




FIG. 1




FIG. 2




FIG. 3




FIG. 4

4/9




FIG. 5




FIG. 6

6/9




FIG. 7

7/9




FIG. 8




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/033525

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

See supplemental page

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☒ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/033525

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G10L 13/08; G10L 13/06; G10L 19/00; G10L 15/00 (2016.01)

CPC - G10L 13/08; G10L 13/06; G10L 19/0018; G10L 15/22 (2016.05)

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)

IPC - G10L 13/08; G10L 13/06; G10L 19/00; G10L 15/00

CPC - G10L 13/08; G10L 13/06; G10L 19/0018; G10L 15/22

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

US: 704/261; 704/258; 704/231 (keyword delimited)

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

PatBase, Google Patents, ProQuest, Google Scholar

Search terms used: speech synthesis, phonemes, vocal tract, constriction, diphones, timestamp

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,729,694 A (HOLZRICHTER et al) 17 March 1998 (17.03.1998) entire document	1-22
A	WO 2014/130571 A1 (THE REGENTS OF THE UNIVERSITY OF CALIFORNIA) 28 August 2014 (28.08.2014) entire document	1-22
A	US 5,774,854 A (SHARMAN) 30 June 1998 (30.06.1998) entire document	1-22
A	US 2010/0076762 A1 (COSATTO et al) 25 March 2010 (25.03.2010) entire document	1-22
A	US 7,216,076 B2 (ATAL) 08 May 2007 (08.05.2007) entire document	1-22
A	US 2006/0100878 A1 (CONKIE) 11 May 2006 (11.05.2006) entire document	1-22
A	DUSAN. Statistical Estimation of Articulatory Trajectories from the Speech Signal using Dynamical and Phonological Constraints. 2000. [retrieved on 2016-09-13]. Retrieved from the Internet:<URL: http://www.collectionscanada.gc.ca/obj/s4/f2/dsk1/tape4/PQDD_0022/NQ51191.pdf >. entire document.	1-22
A	US 6,041,297 A (GOLDBERG) 21 March 2000 (21.03.2000) entire document	1-22
A	KATSIR. Artificial Bandwidth Extension of Band Limited Speech Based on Vocal Tract Shape Estimation. 2012. [retrieved on 2016-09-13]. Retrieved from the Internet:<URL: http://webee.technion.ac.il/Sites/People/IsraelCohen/Info/Graduates/PDF/ItaiKatsir_MSc_2012.pdf >. entire document.	1-22
A	US 2012/0109654 A1 (NURMINEN et al) 03 May 2012 (03.05.2012) entire document	1-22



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

14 September 2016

Date of mailing of the international search report

07 OCT 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/033525

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2005/0008179 A1 (QUINN) 13 January 2005 (13.01.2005) entire document	1-22
A	US 2002/0143543 A1 (SIRIVARA) 03 October 2002 (03.10.2002) entire document	1-22

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/033525

Continued from Box No. III Observations where unity of invention is lacking

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I, claims 1-14,21, drawn to operating a computerized device to provide real-time synthesis of speech.

Group II, claims 15-20,22, drawn to operating a computerized device to provide real-time synthesis of speech comprising maintaining a library of prerecorded samples of diphones.

The inventions listed as Groups I-II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons: the special technical feature of the Group I invention: a first dimension representing degree of vocal tract constriction and voicing, a second dimension representing location in a vocal tract as claimed therein is not present in the invention of Group II. The special technical feature of the Group II invention: placing each identified best diphone into a production buffer with a corresponding diphone time stamp; and sequentially rendering the diphones from the production buffer at respective times corresponding to the diphone time stamps as claimed therein is not present in the invention of Group I.

Groups I and II lack unity of invention because even though the inventions of these groups require the technical feature of operating a computerized device to provide real-time synthesis of speech based on user input; sequence of phonemes, this technical feature is not a special technical feature as it does not make a contribution over the prior art.

Specifically, US 2006/0100878 A1 (CONKIE) 11 May 2006 (11.05.2006) teaches operating a computerized device to provide real-time synthesis of speech based on user input (Paras. 16 and 21); sequence of phonemes (Paras. 16 and 21).

Since none of the special technical features of the Group I or II inventions are found in more than one of the inventions, unity of invention is lacking.