



SDS: ASR, NLU, & VXML

Ling575
Spoken Dialog
April 14, 2016

Roadmap

- Dialog System components:
 - ASR: Noisy channel model
 - Representation
 - Decoding
 - NLU:
 - Call routing
 - Grammars for dialog systems
- Basic interfaces: VoiceXML

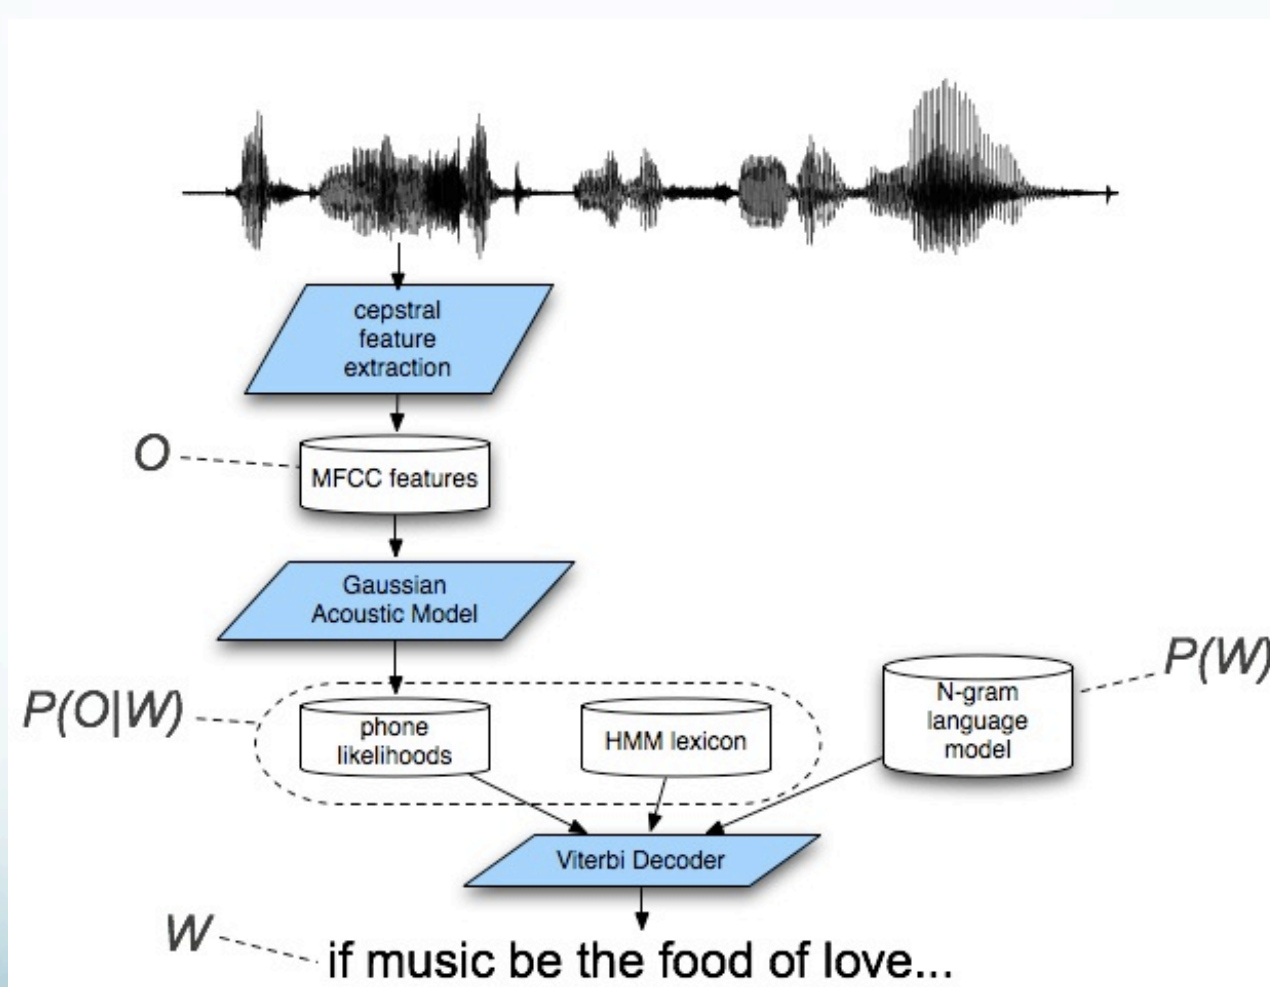
Why is conversational speech harder?

-  • A piece of an utterance without context
-  • The same utterance with more context

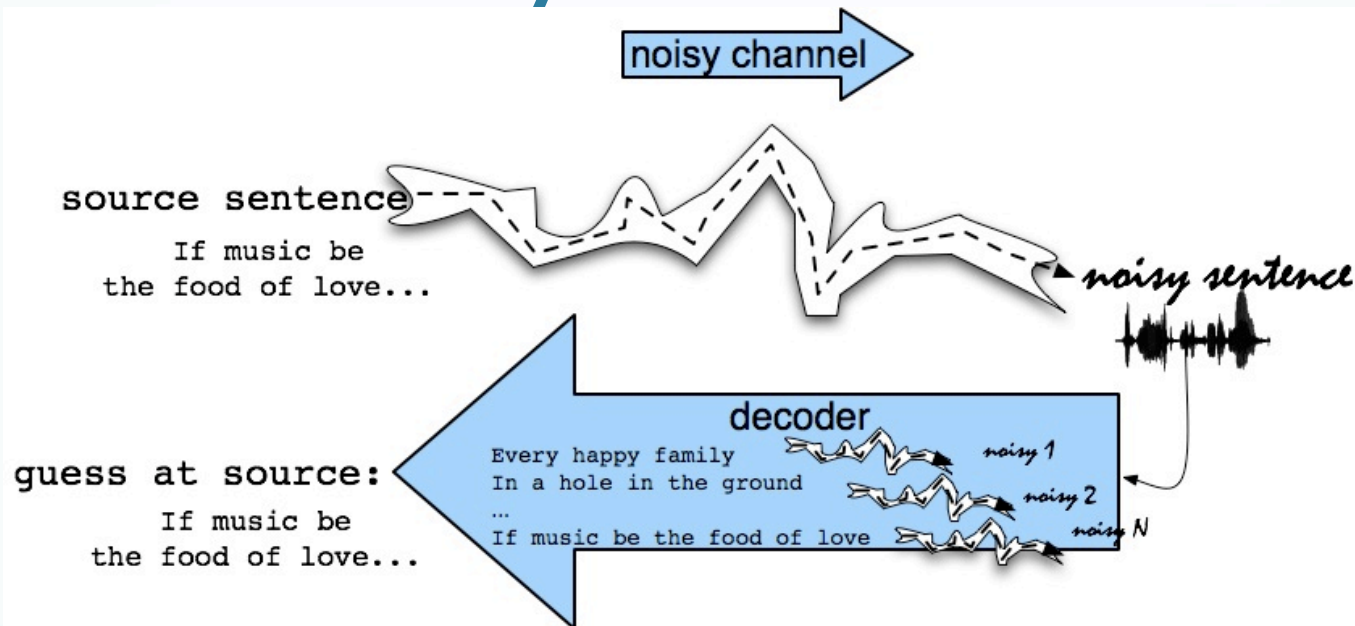
LVCSR Design Intuition

- Build a statistical model of the speech-to-words process
- Collect lots and lots of speech, and transcribe all the words.
- Train the model on the labeled speech
- Paradigm: Supervised Machine Learning + Search

Speech Recognition Architecture



The Noisy Channel Model



- Search through space of all possible sentences.
- Pick the one that is most probable given the waveform.

Decomposing Speech Recognition

- Q1: What speech sounds were uttered?
 - Human languages: 40-50 phones
 - Basic sound units: b, m, k, ax, ey, ...(arpabet)
 - Distinctions categorical to speakers
 - Acoustically continuous
 - Part of knowledge of language
 - Build per-language inventory
 - Could we learn these?

Decomposing Speech Recognition

- Q2: What words produced these sounds?
 - Look up sound sequences in dictionary
 - Problem 1: Homophones
 - Two words, same sounds: too, two
 - Problem 2: Segmentation
 - No “space” between words in continuous speech
 - “I scream”/”ice cream”, “Wreck a nice beach”/”Recognize speech”
- Q3: What meaning produced these words?
 - NLP (But that’s not all!)

The Noisy Channel Model (II)

- What is the most likely sentence out of all sentences in the language L given some acoustic input O ?
- Treat acoustic input O as sequence of individual observations
 - $O = o_1, o_2, o_3, \dots, o_t$
- Define a sentence as a sequence of words:
 - $W = w_1, w_2, w_3, \dots, w_n$

Noisy Channel Model (III)

- Probabilistic implication: Pick the highest prob $S = W$:

$$\hat{W} = \operatorname{argmax}_{W \in L} P(W | O)$$

- We can use Bayes rule to rewrite this:

$$\hat{W} = \operatorname{argmax}_{W \in L} \frac{P(O | W)P(W)}{P(O)}$$

- Since denominator is the same for each candidate sentence W , we can ignore it for the argmax:

$$\hat{W} = \operatorname{argmax}_{W \in L} P(O | W)P(W)$$

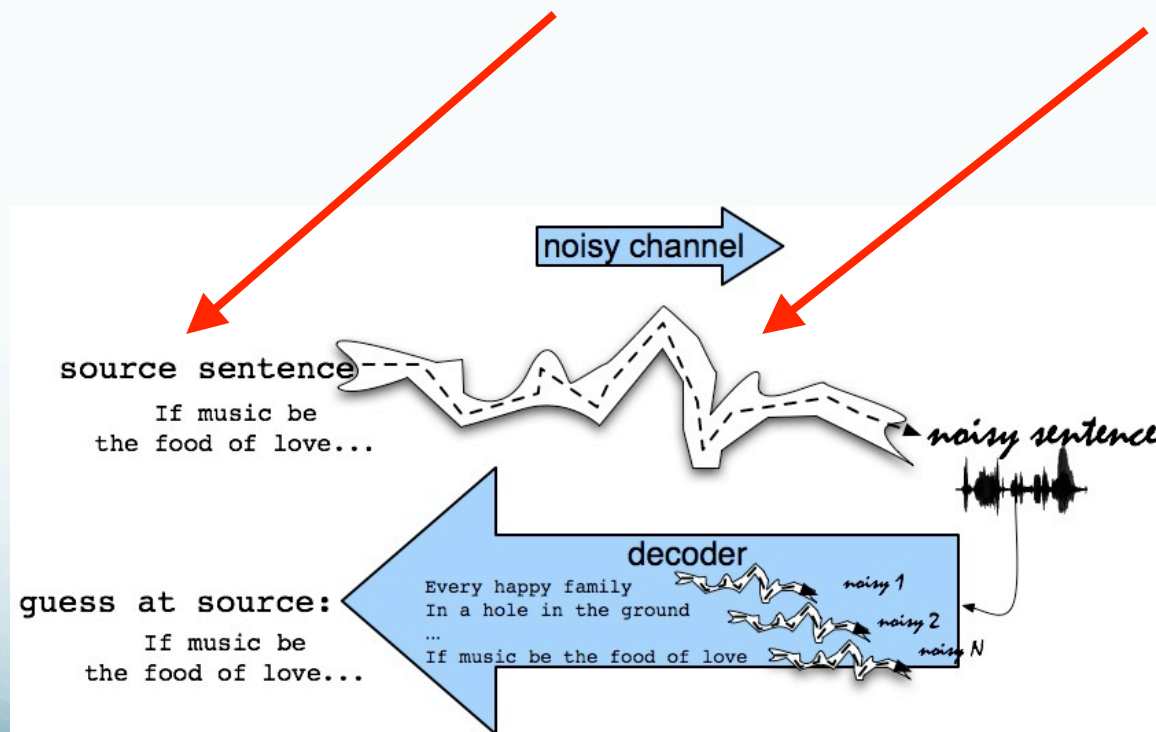
Noisy channel model

$$\hat{W} = \arg \max_{W \in L} P(O | W) P(W)$$

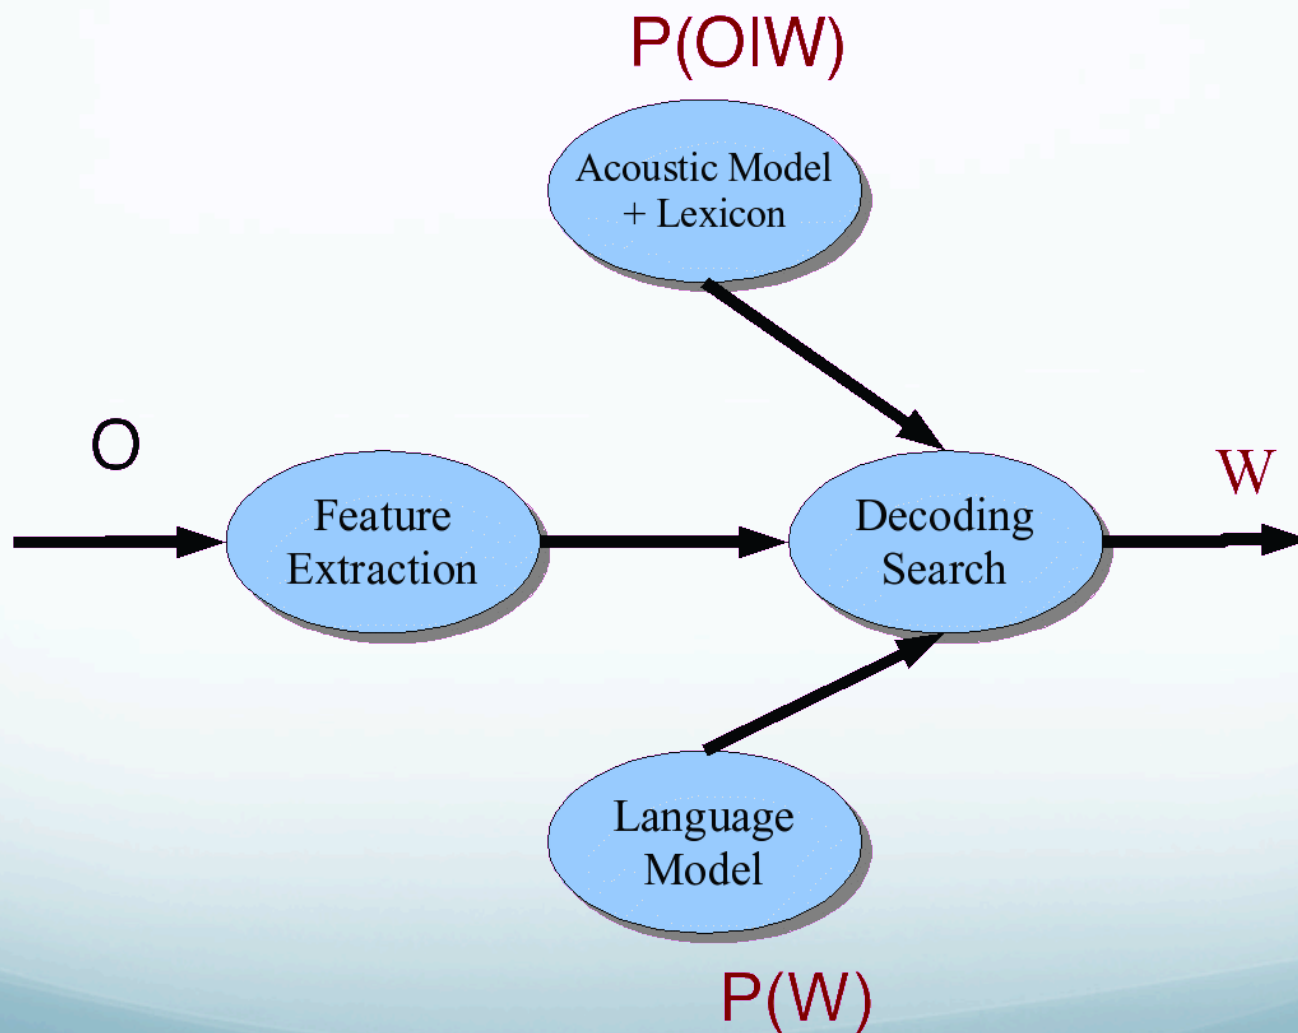
likelihood prior

The noisy channel model

- Ignoring the denominator leaves us with two factors:
 $P(\text{Source})$ and $P(\text{Signal}|\text{Source})$



Speech Architecture meets Noisy Channel



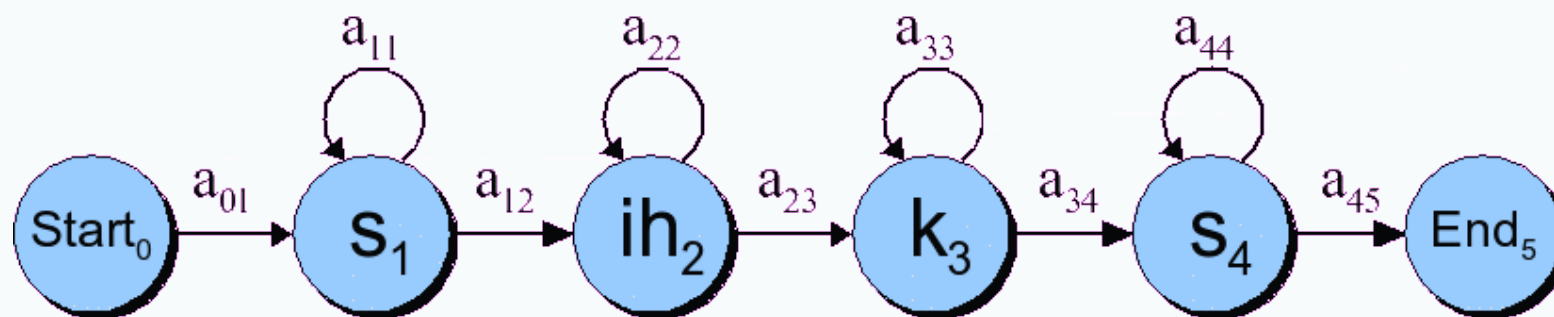
ASR Components

- Lexicons and Pronunciation:
 - Hidden Markov Models
- Feature extraction
- Acoustic Modeling
- Decoding
- Language Modeling:
 - Ngram Models

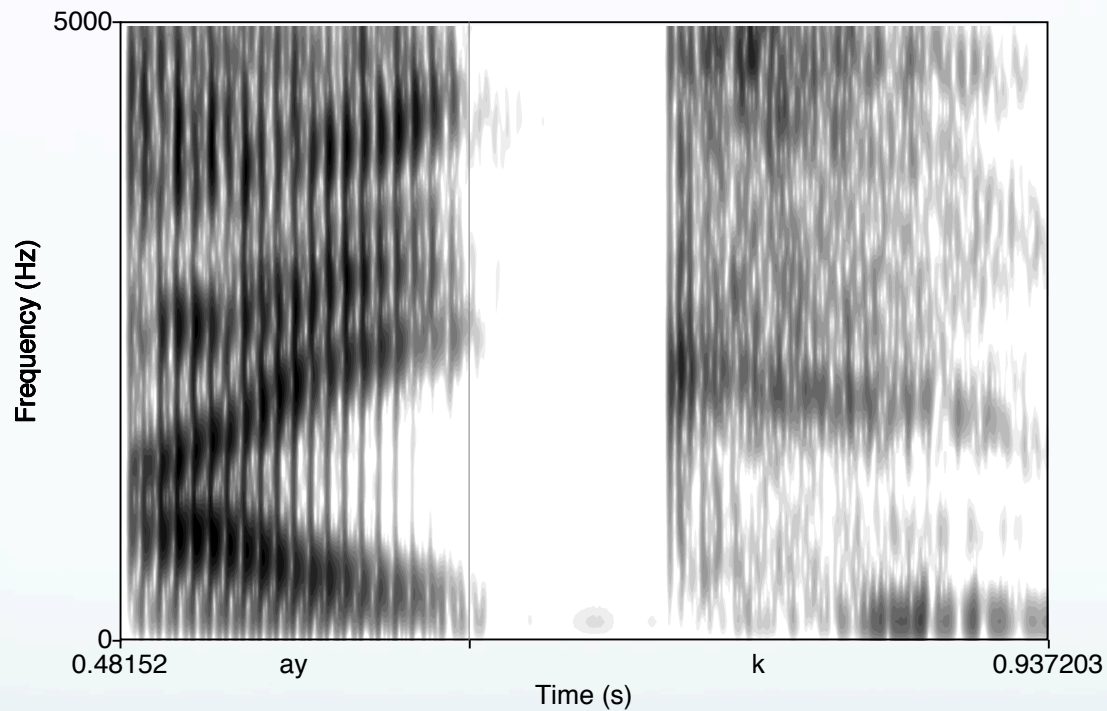
Lexicon

- A list of words
- Each one with a pronunciation in terms of phones
- We get these from on-line pronunciation dictionary
- CMU dictionary: 127K words
 - <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>
- We'll represent the lexicon as an HMM

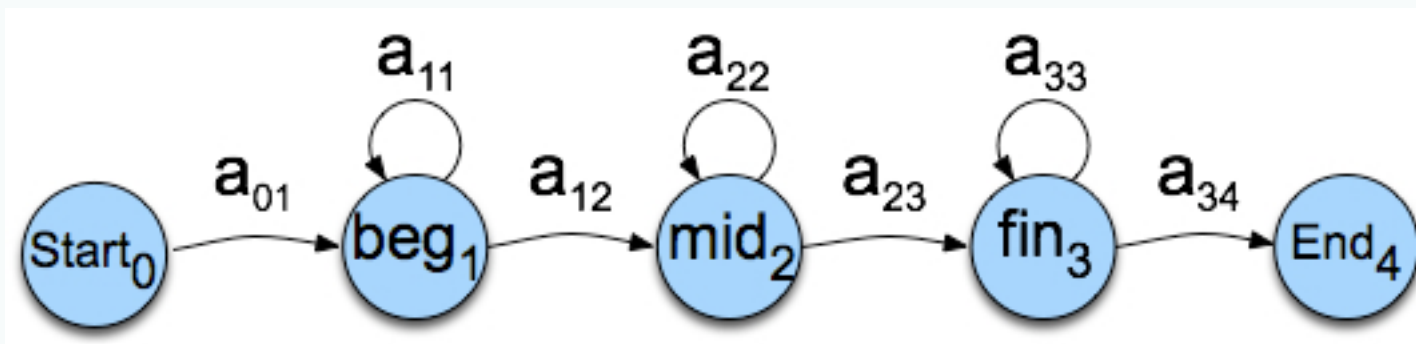
HMMs for speech: the word “six”



Phones are not homogeneous!

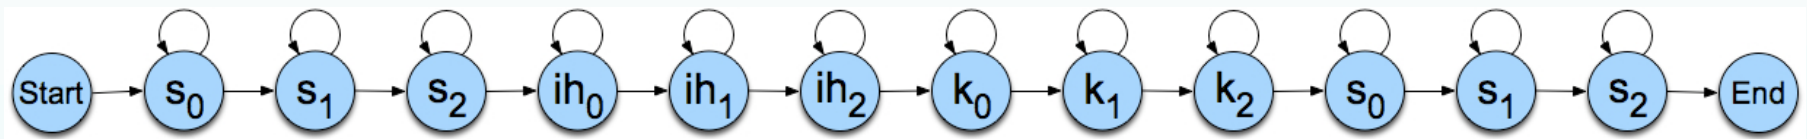


Each phone has 3 subphones



HMM word model for “six”

- Resulting model with subphones



HMMs for speech

$$Q = q_1 q_2 \dots q_N$$

$$A = a_{01} a_{02} \dots a_{n1} \dots a_{nn}$$

$$B = b_i(o_t)$$

a set of **states** corresponding to **subphones**

a **transition probability matrix** A , each a_{ij} representing the probability for each subphone of taking a **self-loop** or going to the next subphone. Together, Q and A implement a **pronunciation lexicon**, an HMM state graph structure for each word that the system is capable of recognizing.

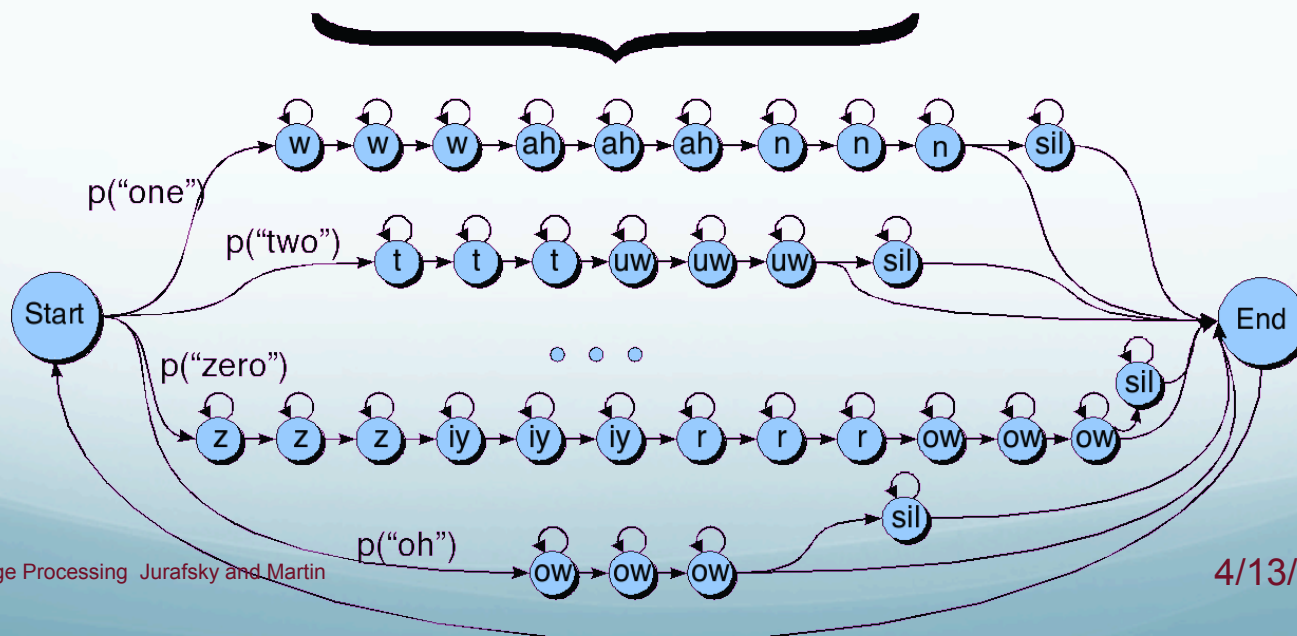
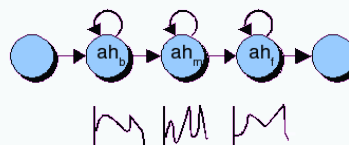
A set of **observation likelihoods**:, also called **emission probabilities**, each expressing the probability of a cepstral feature vector (observation o_t) being generated from subphone state i .

HMM for the digit recognition task

Lexicon

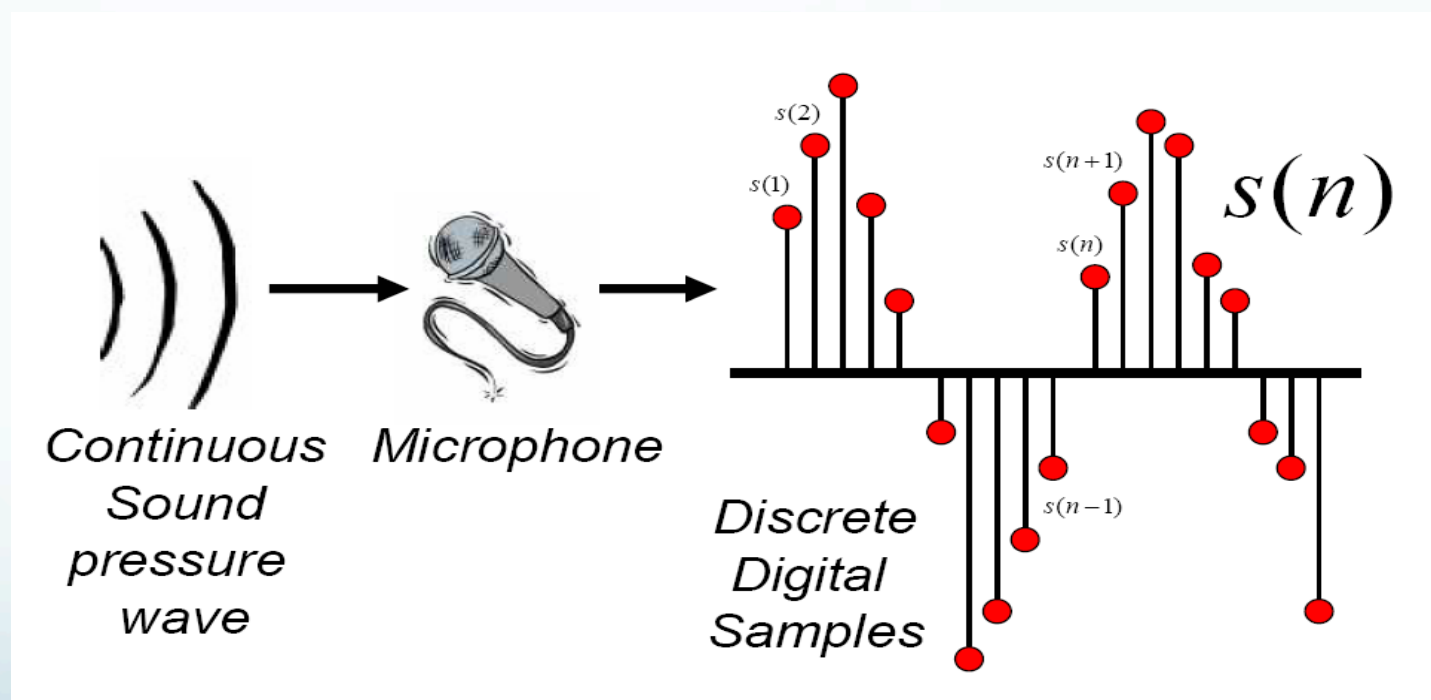
one	w ah n
two	t uw
three	th r iy
four	f ao r
five	f ay v
six	s ih k s
seven	s eh v ax n
eight	ey t
nine	n ay n
zero	z iy r ow
oh	ow

Phone HMM



Discrete Representation of Signal

- Represent continuous signal into discrete form.

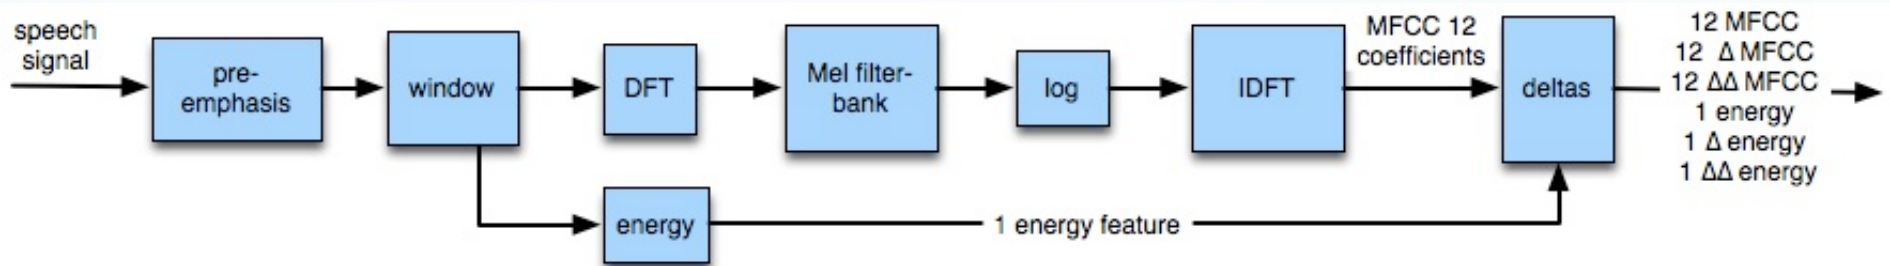


Digitizing the signal (A-D)

■ **Sampling:**

- measuring amplitude of signal at time t
- 16,000 Hz (samples/sec) Microphone (“Wideband”):
- 8,000 Hz (samples/sec) Telephone
- Why?
 - Need at least 2 samples per cycle
 - max measurable frequency is half sampling rate
 - Human speech $< 10,000$ Hz, so need max 20K
 - Telephone filtered at 4K, so 8K is enough

MFCC: Mel-Frequency Cepstral Coefficients



Typical MFCC features

- Window size: 25ms
- Window shift: 10ms
- Pre-emphasis coefficient: 0.97
- MFCC:
 - 12 MFCC (mel frequency cepstral coefficients)
 - 1 energy feature
 - 12 delta MFCC features
 - 12 double-delta MFCC features
 - 1 delta energy feature
 - 1 double-delta energy feature
- Total 39-dimensional features

Why is MFCC so popular?

- Efficient to compute
- Incorporates a perceptual Mel frequency scale
- Separates the source and filter
- Fits well with HMM modelling

Decoding

- In principle:

$$\hat{W} = \operatorname{argmax}_{W \in \mathcal{L}} \overbrace{P(O|W)}^{\text{likelihood}} \overbrace{P(W)}^{\text{prior}}$$

- In practice:

$$\hat{W} = \operatorname{argmax}_{W \in \mathcal{L}} P(O|W)P(W)^{LMSF}$$

$$\hat{W} = \operatorname{argmax}_{W \in \mathcal{L}} P(O|W)P(W)^{LMSF} WIP^N$$

$$\hat{W} = \operatorname{argmax}_{W \in \mathcal{L}} \log P(O|W) + LMSF \times \log P(W) + N \times \log WIP$$

Why is ASR decoding hard?

[ay d ih s hh er d s ah m th ih ng ax b aw m uh v ih ng r ih s en l ih]

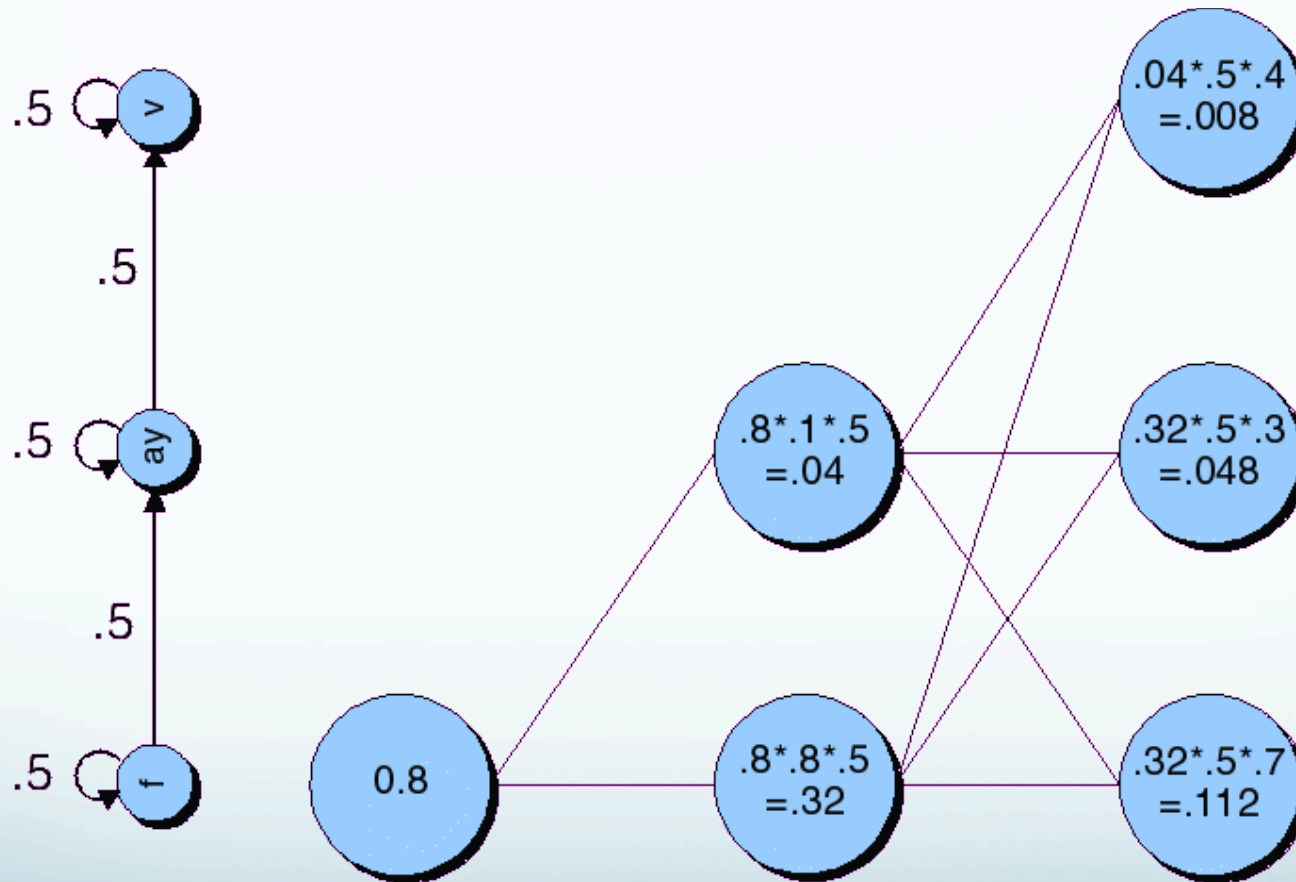
The Evaluation (forward) problem for speech

- The observation sequence O is a series of MFCC vectors
- The hidden states W are the phones and words
- For a given phone/word string W , our job is to evaluate $P(O|W)$
- Intuition: how likely is the input to have been generated by just that word string W

Evaluation for speech: Summing over all different paths!

- f ay ay ay ay v v v v
- f f ay ay ay ay v v v
- f f f f ay ay ay ay v
- f f ay ay ay ay ay ay v
- f f ay ay ay ay ay ay ay ay v
- f f ay v v v v v v v

Viterbi trellis for “five”



Viterbi trellis for “five”

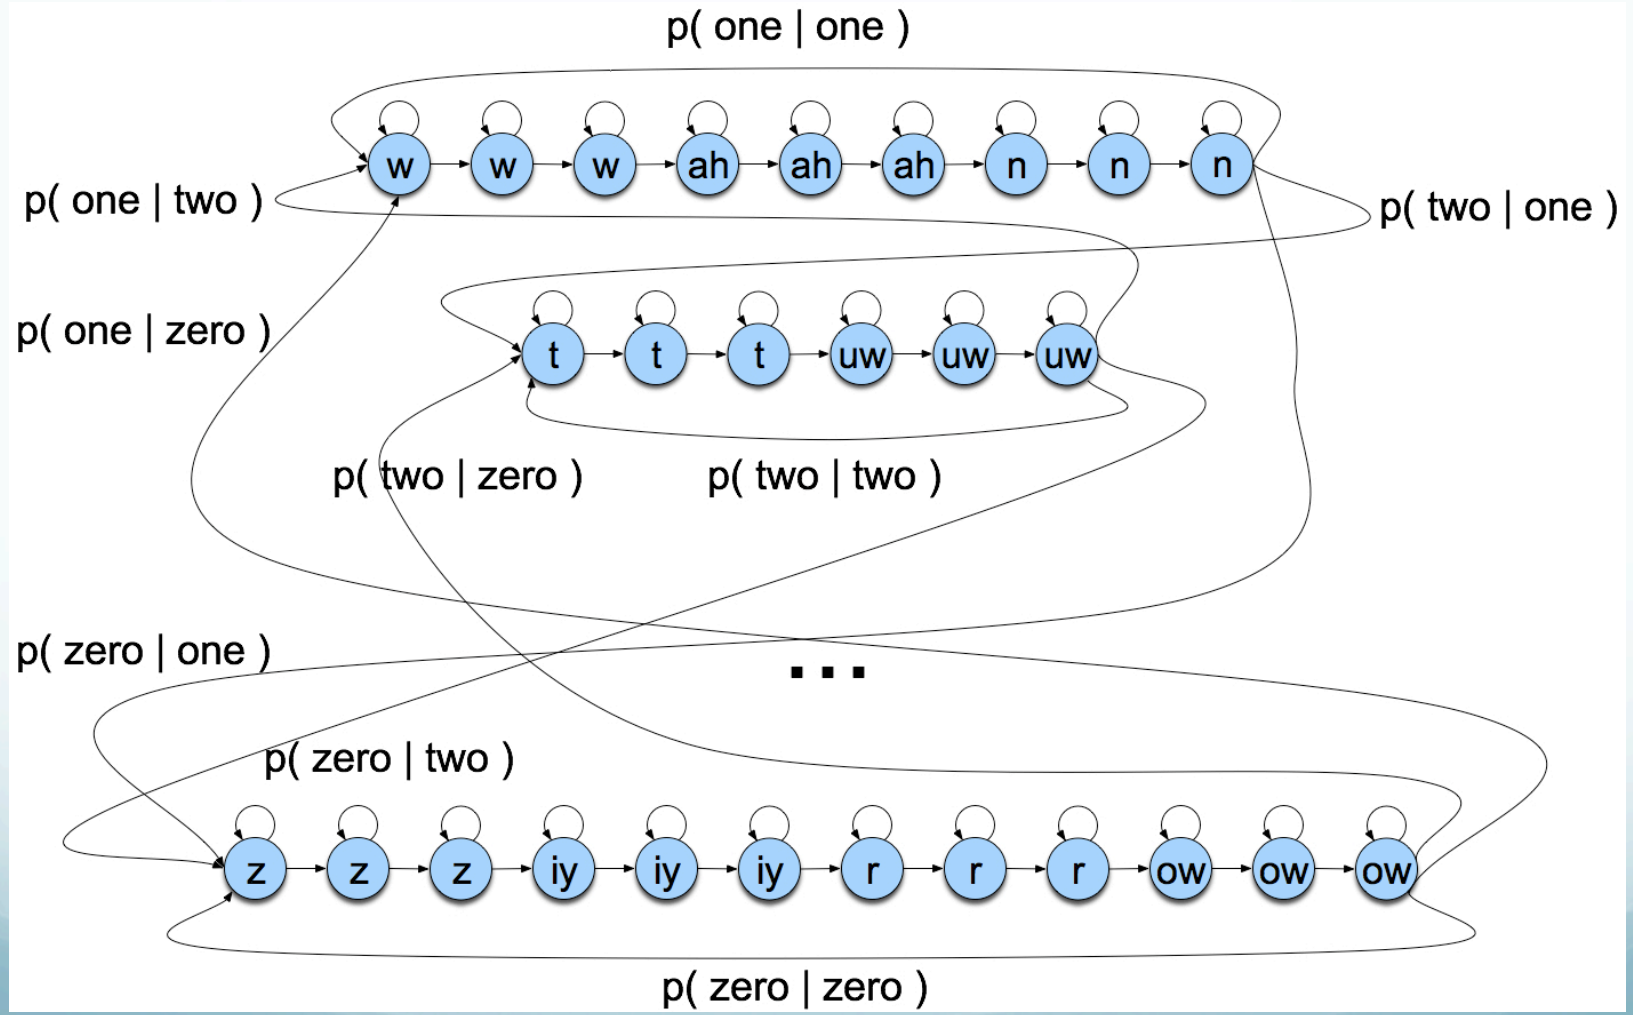
V	0	0	0.008	0.0072	0.00672	0.00403	0.00188	0.00161	0.000667	0.000493
AY	0	0.04	0.048	0.0448	0.0269	0.0125	0.00538	0.00167	0.000428	8.78e-05
F	0.8	0.32	0.112	0.0224	0.00448	0.000896	0.000179	4.48e-05	1.12e-05	2.8e-06
Time	1	2	3	4	5	6	7	8	9	10
B	<i>f</i> 0.8	<i>f</i> 0.8	<i>f</i> 0.7	<i>f</i> 0.4	<i>f</i> 0.4	<i>f</i> 0.4	<i>f</i> 0.4	<i>f</i> 0.5	<i>f</i> 0.5	<i>f</i> 0.5
	<i>ay</i> 0.1	<i>ay</i> 0.1	<i>ay</i> 0.3	<i>ay</i> 0.8	<i>ay</i> 0.8	<i>ay</i> 0.8	<i>ay</i> 0.8	<i>ay</i> 0.6	<i>ay</i> 0.5	<i>ay</i> 0.4
	<i>v</i> 0.6	<i>v</i> 0.6	<i>v</i> 0.4	<i>v</i> 0.3	<i>v</i> 0.3	<i>v</i> 0.3	<i>v</i> 0.3	<i>v</i> 0.6	<i>v</i> 0.8	<i>v</i> 0.9
	<i>p</i> 0.4	<i>p</i> 0.4	<i>p</i> 0.2	<i>p</i> 0.1	<i>p</i> 0.1	<i>p</i> 0.1	<i>p</i> 0.1	<i>p</i> 0.1	<i>p</i> 0.3	<i>p</i> 0.3
	<i>iy</i> 0.1	<i>iy</i> 0.1	<i>iy</i> 0.3	<i>iy</i> 0.6	<i>iy</i> 0.6	<i>iy</i> 0.6	<i>iy</i> 0.6	<i>iy</i> 0.5	<i>iy</i> 0.5	<i>iy</i> 0.4

Language Model

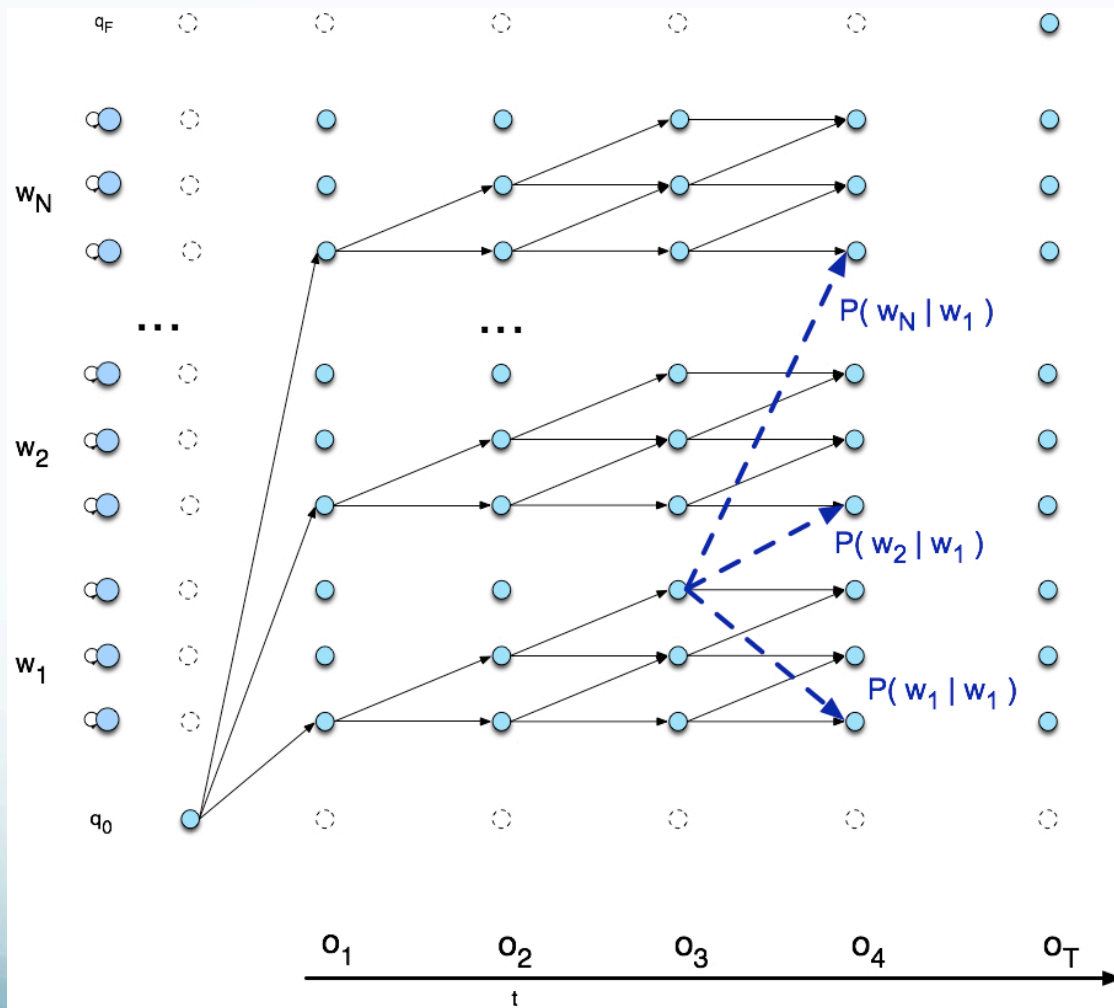
- Idea: some utterances more probable
- Standard solution: “n-gram” model
 - Typically tri-gram: $P(w_i | w_{i-1}, w_{i-2})$
 - Collect training data from large side corpus
 - Smooth with bi- & uni-grams to handle sparseness
 - Product over words in utterance:

$$P(w_1^n) \approx \prod_{k=1}^n P(w_k | w_{k-1}, w_{k-2})$$

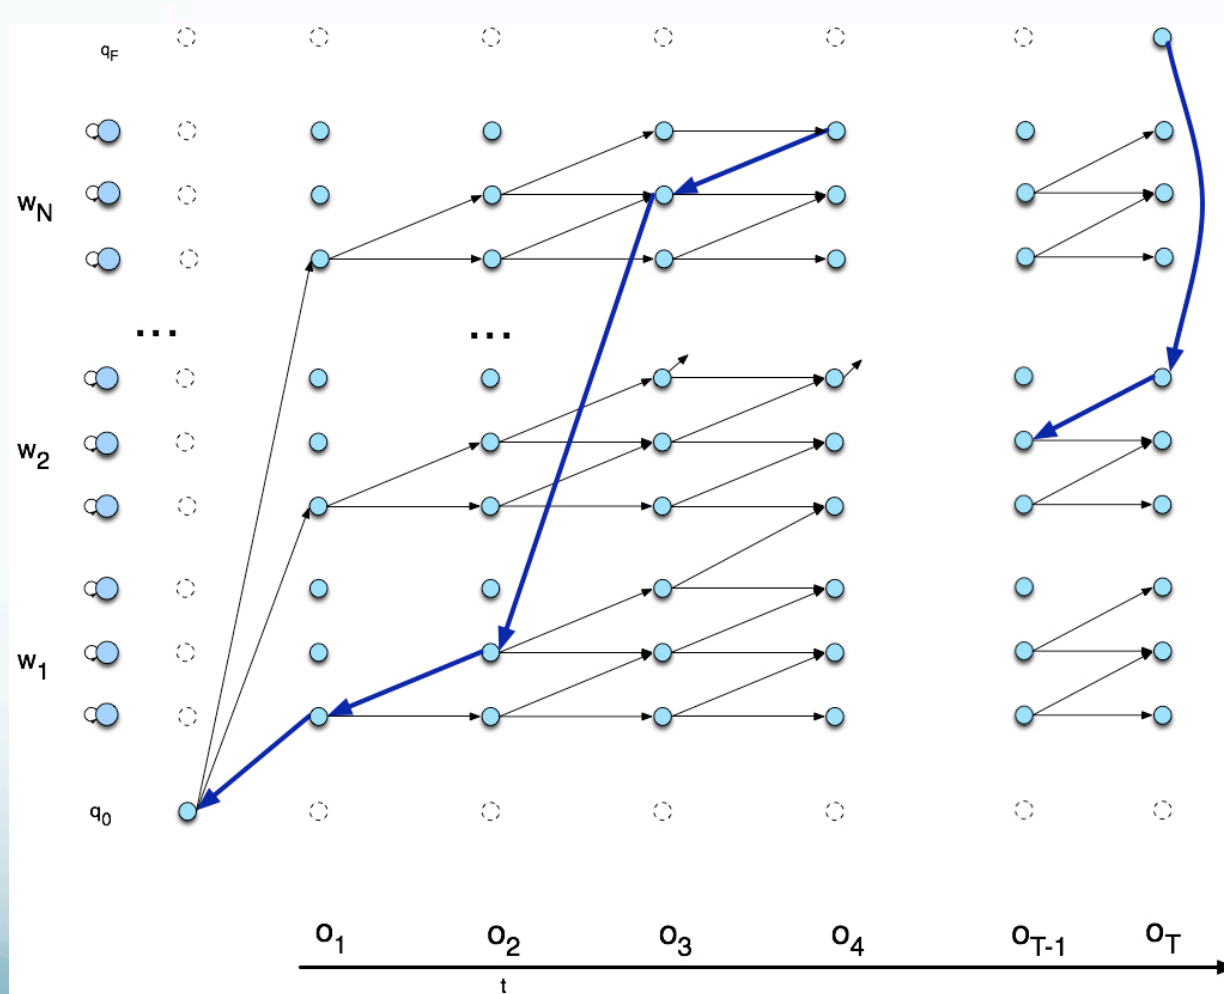
Search space with bigrams



Viterbi trellis

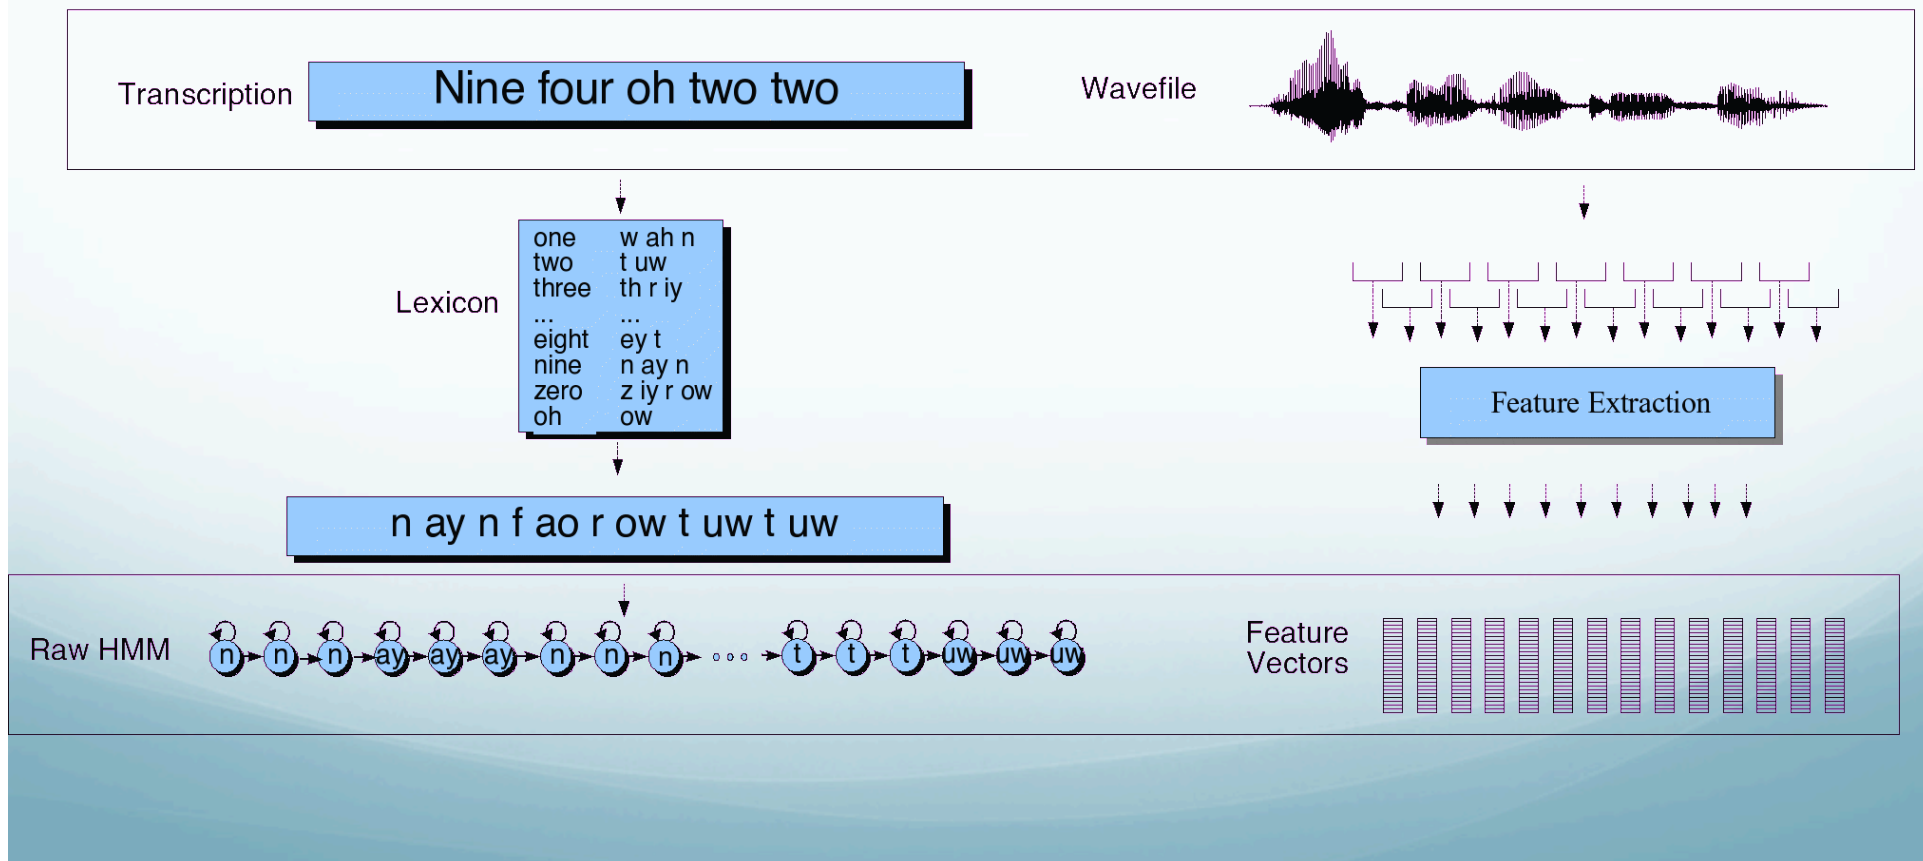


Viterbi backtrace



Training

- Trained using Baum-Welch algorithm



Summary: ASR Architecture

- Five easy pieces: ASR Noisy Channel architecture
 - 1) Feature Extraction:
 - 39 “MFCC” features
 - 2) Acoustic Model:
 - Gaussians for computing $p(o|q)$
 - 3) Lexicon/Pronunciation Model
 - HMM: what phones can follow each other
 - 4) Language Model
 - N-grams for computing $p(w_i|w_{i-1})$
 - 5) Decoder
 - Viterbi algorithm: dynamic programming for combining all these to get word sequence from speech!

Deep Neural Networks for ASR

- Since ~2012, yielded significant improvements
- Applied to two stages of ASR
 - Acoustic modeling for tandem/hybrid HMM:
 - DNNs replace GMMs to compute phone class probabilities
 - Provide observation probabilities for HMM
 - Language modeling:
 - Continuous models often interpolated with n-gram models

DNN Advantages for Acoustic Modeling

- Support improved acoustic features
 - GMMs use MFCCs rather than raw filterbank ones
 - MFCCs advantages are compactness and decorrelation
 - BUT lose information
 - Filterbank features are correlated, too expensive for GMM
 - DNNs:
 - Can use filterbank features directly
 - Can also effectively incorporate longer context
- Modeling:
 - GMMs more local, weak on non-linear; DNNs more flexible
 - GMMs model single component; (D)NNs can be multiple
 - DNNs can build richer representations

Why the post-2012 boost?

- Some earlier NN/MLP tandem approaches
 - Had similar modeling advantages
- However, training was problematic and expensive
- Newer approaches have:
 - Better strategies for initialization
 - Better learning methods for many layers
 - See “vanishing gradient”
 - GPU implementations support faster computation
 - Parallelism at scale

Word Error Rate

- Word Error Rate =

100 (Insertions+Substitutions + Deletions)

Total Word in Correct Transcript

Alignment example:

REF: portable **** PHONE UPSTAIRS last night so

HYP: portable FORM OF STORES last night so

Eval I S S

$$\text{WER} = 100 (1+2+0)/6 = 50\%$$

NIST sctk-1.3 scoring software: Computing WER with sclite

- <http://www.nist.gov/speech/tools/>
- Sclite aligns a hypothesized text (HYP) (from the recognizer) with a correct or reference text (REF) (human transcribed)

id: (2347-b-013)

Scores: (#C #S #D #I) 9 3 1 2

REF: was an engineer SO I i was always with **** * MEN UM and they

HYP: was an engineer ** AND i was always with THEM THEY ALL THAT and they

Eval: D S I I S S

Better metrics than WER?

- WER has been useful
- But should we be more concerned with meaning (“semantic error rate”)?
 - Good idea, but hard to agree on
 - Has been applied in dialogue systems, where desired semantic output is more clear

Accents: An experiment

- A word by itself



- The word in context



Challenges for the Future

- Doing more with more
 - More applications:
 - From Siri, in-car navigation, call-routing
 - To full voice search, voice-based personal assistants, ubiquitous computing
 - More speech types:
 - Accented speech
 - Speech in noise
 - Overlapping speech
 - Child speech
 - Speech pathology



NLU for Dialog Systems

Natural Language Understanding

- Generally:
 - Given a string of words representing a natural language utterance, produce a meaning representation
- For well-formed natural language text (see ling571),
 - Full parsing with a probabilistic context-free grammar
 - Augmented with semantic attachments in FOPC
 - Producing a general lambda calculus representation
- What about spoken dialog systems?

NLU for SDS

- Few SDS fully exploit this approach
- Why not?
 - Examples of travel air speech input (due to A. Black)
 - Eh, I wanna go, wanna go to Boston tomorrow
 - If its not too much trouble I'd be very grateful if one might be able to aid me in arranging my travel arrangements to Boston, Logan airport, at sometime tomorrow morning, thank you.
 - Boston, tomorrow

NLU for SDS

- Analyzing speech vs text
 - Utterances:
 - ill-formed, disfluent, fragmentary, desultory, rambling
 - Vs well-formed
 - Domain:
 - Restricted, constrains interpretation
 - Vs. unrestricted
 - Interpretation:
 - Need specific pieces of data
 - Vs. full, complete representation
 - Speech recognition:
 - Error-prone, perfect full analysis difficult to obtain

NLU for Spoken Dialog

- Call routing (aka call classification):
 - (Chu-Carroll & Carpenter, 1998, Al-Shawi 2003)
 - Shallow form of NLU
 - Goal:
 - Given a spoken utterance, assign to class c , in finite set C
 - Banking Example:
 - Open prompt: **"How may I direct your call?"**
 - Responses: may I have consumer lending?,
 - I'd like my checking account balance, or
 - "ah I'm calling 'cuz ah a friend gave me this number and ah she told me ah with this number I can buy some cars or whatever but she didn't know how to explain it to me so I just called you you know to get that information."

Call Routing

- General approach:
 - Build classification model based on labeled training data, e.g. manually routed calls
 - Apply classifier to label new data
- Vector-based call routing:
 - Model: Vector of word unigram, bigrams, trigrams
 - Filtering: by frequency
 - Exclude high frequency stopwords, low frequency rare words
 - Weighting: term frequency * inverse document frequency
 - (Dimensionality reduction by singular value decomposition)
 - Compute cosine similarity for new call & training examples

Natural Language Understanding

- Most systems use frame-slot semantics
Show me morning flights from Boston to SFO on Tuesday
- SHOW:
- FLIGHTS:
 - ORIGIN:
 - CITY: Boston
 - DATE:
 - DAY-OF-WEEK: Tuesday
 - TIME:
 - PART-OF-DAY: Morning
 - DEST:
 - CITY: San Francisco

Another NLU Example

- Sagae et 2009
- Utterance (speech): we are prepared to give you guys generators for electricity downtown
- ASR (NLU input): we up apparently give you guys generators for a letter city don town
- Frame (NLU output):
 - **<s>.mood declarative**
 - **<s>.sem.agent kirk**
 - **<s>.sem.event deliver**
 - **<s>.sem.modal.possibility can**
 - **<s>.sem.speechact.type offer**
 - **<s>.sem.theme power-generator**
 - **<s>.sem.type event**

Question

- Given an ASR output string, how can we tractably and robustly derive a meaning representation?
- Many approaches:
 - Shallow transformation:
 - Terminal substitution
 - Integrated parsing and semantic analysis
 - E.g. semantic grammars
 - Classification or sequence labeling approaches
 - HMM-based, MaxEnt-based

Grammars

- Formal specification of strings in a language
- A 4-tuple:
 - A set of terminal symbols: Σ
 - A set of non-terminal symbols: N
 - A set of productions P : of the form $A \rightarrow \alpha$
 - A designated start symbol S
- In regular grammars:
 - A is a non-terminal and α is of the form $\{N\} \Sigma^*$
- In context-free grammars:
 - A is a non-terminal and α in $(\Sigma \cup N)^*$

Simple Air Travel Grammar

- LIST -> show me | I want | can I see|...
- DEPARTTIME -> (after|around|before) HOUR| morning | afternoon | evening
- HOUR -> one|two|three...|twelve (am|pm)
- FLIGHTS -> (a) flight|flights
- ORIGIN -> from CITY
- DESTINATION -> to CITY
- CITY -> Boston | San Francisco | Denver | Washington

Shallow Semantics

- Terminal substitution
 - Employed by some speech toolkits, e.g. CSLU
- Rules convert terminals in grammar to semantics
 - LIST -> show me | I want | can I see|...
 - e.g. show -> LIST
 - see -> LIST
 - I -> ϵ
 - can -> ϵ
 - * Boston -> Boston
- Simple, but...
 - VERY limited, assumes direct correspondence

Semantic Grammars

- Domain-specific semantic analysis
- Syntactic structure:
 - Context-free grammars (CFGs) (typically)
 - Can be parsed by standard CFG parsing algorithms
 - e.g. Earley parsers or CKY
- Semantic structure:
 - Some designated non-terminals correspond to slots
 - Associate terminal values to corresponding slot
- Frames can be nested
- Widely used: Phoenix NLU (CU, CMU), vxml grammars

Show me morning flights from Boston to SFO on Tuesday

- LIST -> show me | I want | can I see|...
- DEPARTTIME -> (after|around|before) HOUR| morning | afternoon | evening
- HOUR -> one|two|three...| twelve (am|pm)
- FLIGHTS -> (a) flight|flights
- ORIGIN -> from CITY
- DESTINATION -> to CITY
- CITY -> Boston | San Francisco | Denver | Washington
- SHOW:
- FLIGHTS:
 - ORIGIN:
 - CITY: Boston
 - DATE:
 - DAY-OF-WEEK: Tuesday
 - TIME:
 - PART-OF-DAY: Morning
 - DEST:
 - CITY: San Francisco

Semantic Grammars: Issues

- Issues:
 - Generally manually constructed
 - Can be expensive, hard to update/maintain
 - Managing ambiguity:
 - Can associate probabilities with parse & analysis
 - Build rules manually, then train probabilities w/data
 - Domain- and application-specific
 - Hard to port



VoiceXML

VoiceXML

- W3C standard for voice interfaces
 - XML-based ‘programming’ framework for speech systems
 - Provides recognition of:
 - Speech, DTMF (touch tone codes)
 - Provides output of synthesized speech, recorded audio
 - Supports recording of user input
 - Enables interchange between voice interface, web-based apps
 - Structures voice interaction
 - Can incorporate Javascript/PHP/etc for functionality

Capabilities

- Interactions:
 - Default behavior is FST-style, system initiative
 - Can implement frame-based mixed initiative
 - Support for sub-dialog call-outs

Speech I/O

- ASR:
 - Supports speech recognition defined by
 - Grammars
 - Trigrams
 - Domain managers: credit card nos etc
- TTS:
 - <ssml> markup language
 - Allows choice of: language, voice, pronunciation
 - Allows tuning of: timing, breaks

Simple VoiceXML Example

- Minimal form:

```
<form>
  <field name="transporttype">
    <prompt>
      Please choose airline, hotel, or rental car.
    </prompt>
    <grammar type="application/x=nuance-gsl">
      [airline hotel "rental car"]
    </grammar>
  </field>
  <block>
    <prompt>
      You have chosen <value expr="transporttype">.
    </prompt>
  </block>
</form>
```

Basic VXML Document

- Main body: `<form></form>`
 - Sequence of fields: `<field></field>`
 - Correspond to variable storing user input
 - `<field name="transporttype">`
 - Prompt for user input
 - `<prompt> Please choose airline, hotel, or rental car.</prompt>`
 - Can include URL for recorded prompt, backs off
 - Specify grammar to recognize/interpret user input
 - `<grammar>[airline hotel "rental car"]</grammar>`

Other Field Elements

- Context-dependent help:
 - `<help>Please select activity.</help>`
- Action to be performed on input:
 - `<filled>`
 - `<prompt>You have chosen <value exp="transporttype">.`
 - `</prompt></filled>`

Control Flow

- Default behavior:
 - Step through elements of form in document order
- Goto allows jump to:
 - Other form: `<goto next="weather.xml">`
 - Other position in form: `<goto next="#departdate">`
- Conditionals:
 - `<if cond="varname=='air'">....</if>`
- Guards:
 - Default: Skip field if slot value already entered

General Interaction

- ‘Universals’:
 - Behaviors used by all apps, specify particulars
 - Pick prompts for conditions
- <noinput>:
 - No speech timeout
- <nomatch>:
 - Speech, but nothing valid recognized
- <help>:
 - General system help prompt

Complex Interaction

- Preamble, grammar:

```
<noinput>    I'm sorry, I didn't hear you. <reprompt/> </noinput>
<nomatch> I'm sorry, I didn't understand that. <reprompt/> </nomatch>

<form>
  <grammar type="application/x=nuance-gsl">
    <![CDATA[
      Flight ( ?[
        (i [wanna (want to)] [fly go])
        (i'd like to [fly go])
        ((i wanna)(i'd like a)] flight)
      ]
      [
        ( [from leaving departing] City:x) {<origin $x>}
        ( [(?going to)(arriving in)] City:x) {<destination $x>}
        ( [from leaving departing] City:x
          [(?going to)(arriving in)] City:y) {<origin $x> <destination $y>}
        ]
      ?please
    )
    City [ [(san francisco) (s f o)] {return( "san francisco, california")}
          [(denver) (d e n)] {return( "denver, colorado")}
          [(seattle) (s t x)] {return( "seattle, washington")}
        ]
    ]> </grammar>

  <initial name="init">
    <prompt> Welcome to the consultant. What are your travel plans? </prompt>
  </initial>
```

Mixed Initiative

- With guard defaults

```
<field name="origin">
  <prompt> Which city do you want to leave from? </prompt>
  <filled>
    <prompt> OK, from <value expr="origin"> </prompt>
  </filled>
</field>
<field name="destination">
  <prompt> And which city do you want to go to? </prompt>
  <filled>
    <prompt> OK, to <value expr="destination"> </prompt>
  </filled>
</field>
<block>
  <prompt> OK, I have you are departing from <value expr="origin">
    to <value expr="destination">. </prompt>
  send the info to book a flight...
</block>
</form>
```


Complex Interaction

- Preamble, external grammar:

```
<?xml version="1.0"?>
<vxml version = "2.0">

<form id="F1">

  <field name="F_1">
    <grammar src="NameGram.xml"
type="application/grammar-xml" />
    <prompt>
      Please tell me your full name so I can verify you
    </prompt>
  </field>

  <filled mode="all" namelist="F_1">
    <prompt>
      Your name is <value expr="F_1"/>
      <break strength="medium"/>
    </prompt>
  </filled>
</form>
</vxml>
```

Multi-slot Grammar

- ```
<?xml version= "1.0"?>
 <grammar xml:lang="en-US" root = "TOPLEVEL">
 <rule id="TOPLEVEL" scope="public">
 <item>
<!-- FIRST NAME RETURN .. >
 <item repeat="0-1">
 <ruleref uri="#FIRSTNAME"/>
 <tag>out.firstNameSlot=rules.FIRSTNAME.firstNameSubslot;</tag>
 </item>
<!-- MIDDLE NAME RETURN ..>
 <item repeat="0-1">
 <ruleref uri="#MIDDLENAME"/>
 <tag>out.middleNameSlot=rules.MIDDLENAME.middleNameSubslot;</tag>
 </item>
<!-- LAST NAME RETURN .. >
 <ruleref uri="#LASTNAME"/>
 <tag>out.lastNameSlot=rules.LASTNAME.lastNameSubslot;</tag>
 </item>
<!-- TOP LEVEL RETURN..>
 <tag> out.F_1= out.firstNameSlot + out.middleNameSlot + out.lastNameSlot; </tag>
 </rule>
```

# Multi-slot Grammar II

- ```
<rule id="FIRSTNAME" scope="public">
  <one-of>
    <item> matt<tag>out.firstNameSubslot="matthew";</tag></item>
    <item> dee <tag> out.firstNameSubslot="dee ";</tag></item>
    <item> jon <tag> out.firstNameSubslot="jon ";</tag></item>
    <item> george <tag>out.firstNameSubslot="george ";</tag></item>
    <item> billy <tag> out.firstNameSubslot="billy ";</tag></item>
  </one-of>
</rule>
<rule id="MIDDLENAME" scope="public">
  <one-of>
    <item> bon <tag>out.middleNameSubslot="bon ";</tag></item>
    <item> double ya <tag> out.middleNameSubslot="w ";</tag></item>
    <item> dee <tag> out.middleNameSubslot="dee ";</tag></item>
  </one-of>
</rule>
<rule id="LASTNAME" scope="public">
  <one-of>
    <item> henry <tag> out.lastNameSubslot="henry "; </tag></item>
    <item> ramone <tag> out.lastNameSubslot="dee "; </tag></item>
    <item> jovi <tag> out.lastNameSubslot="jovi "; </tag></item>
    <item> bush <tag> out.lastNameSubslot=""bush "; </tag></item>
    <item> williams <tag> out.lastNameSubslot="williams "; </tag></item>
  </one-of>
</rule>
```

Augmenting VoiceXML

- Don't write XML directly
 - Use php or other system to generate VoiceXML
 - Used in 'Let's Go Dude' bus info system
- Pass input to other web services
 - i.e. to RESTful services
- Access web-based audio for prompts