Microsoft Research

The Conversational Web

Larry Heck

Contributors:
Dilek Hakkani-Tür, Gokhan Tur, Andreas Stolcke, Ashley Fidler, Lisa Stifelman, Liz Shriberg, Madhu Chinthakunta, Malcolm Slaney, Patrick Pantel, Benoit Dumoulin

Microsoft

Outline

- Background
  - Evolution of the conversational systems
  - Parallel evolution of the web, information retrieval/search
  - Converging trends
- The Emergence of the Conversational Web
  - Motivations and Benefits
  - Major Research Challenges
  - Progress to Date
- Summary and Conclusions

Conversational Systems
Brief History

TV Voice Search
e.g., Bing on Xbox

Multi-modal systems
e.g., Microsoft MiPad, Pocket PC

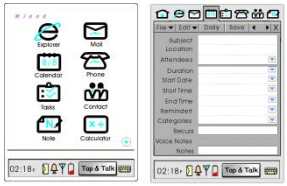

2012

Task-specific argument extraction
(e.g., Nuance, SpeechWorks)
User: "I want to fly from Boston to New York next week."

Early 2000s

Intent Determination
(e.g. Nuance's Emily™, AT&T HMIHY)
User: "Uh…we want to move…we want to change our phone line from this house to another house"

Virtual Personal Assistants:
e.g., Siri, Google now

DARPA
CALO Project

Early 1990s

Keyword Spotting
(e.g., AT&T)
System: "*Please say collect, calling card, person, third number, or operator*"

Web Search: A Parallel Evolution

Brief History[2]

- 1940's: The potential of information retrieval (IR) through computers was realized
- 1950's: The idea of using words as indexing units for documents and measuring word overlap as a criterion for retrieval was established
- 1960's: Systematic methods of evaluating efficacy of IR systems evolved
- 1990's: Advent of the internet and inception of the Text Retrieval Conference in 1992 made available large collections of text for validating/enhancing various IR techniques
- 2000's: Web search became the most ubiquitous application of IR scaling to information stored in unstructured form over billions of documents

[2] Modern Information Retrieval: A Brief Overview, Amit Singhal, Google

The Applification of the Web

Another Parallel Evolution

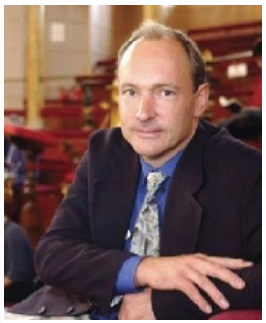

1991

World Wide Web Goes Online

With help of Robert Cailliau, Tim Berners-Lee writes the first world wide web server. It goes online world-wide via the Internet network in the summer of 1991.

1994

Web Directories

In the early days of the Web, companies like Yahoo! and Galaxy created online directories to help users find what they needed. Between 1994-1996, the Web grows from 60K documents to millions of documents.

Google™

2000

Searching the Web

Since 1997, Larry Page and Sergey Brin go on to revolutionize web search by counting site links and ranking accordingly.

1991 1994 2000 2007 2010 2012

2007

Appification of the Web Begins

Apple launches the iPhone and announces that it would support third-party applications, which users would access via the Internet. The first app, “OneTrip”, keeps track of a user’s shopping list.

2007-2012

Apps Directories

Marketplaces by Apple, Google and Microsoft offer directories to help users find apps. Between 2007-2012, the marketplaces grow from 0 to 600K+400K+70K apps, a $12B industry.

The Opportunity: Convergence

Implications for Spoken Language Technology

- The Limitations of Past Conversational Systems
  - Lack of Scale: adding domains is too manual
  - Lack of Ubiquity: spotty coverage of voice interfaces
  - Non-standard User Interfaces (UI)

- The Success of Web Search & Browse
  - Breadth: access to the world's information
  - Ubiquity: anywhere, anytime
  - Well understood, simple UI

- The Application of the Web
  - Adds Depth to Breadth of search & browse
  - Web of apps *and* documents
  - Convergence of web search and apps underway

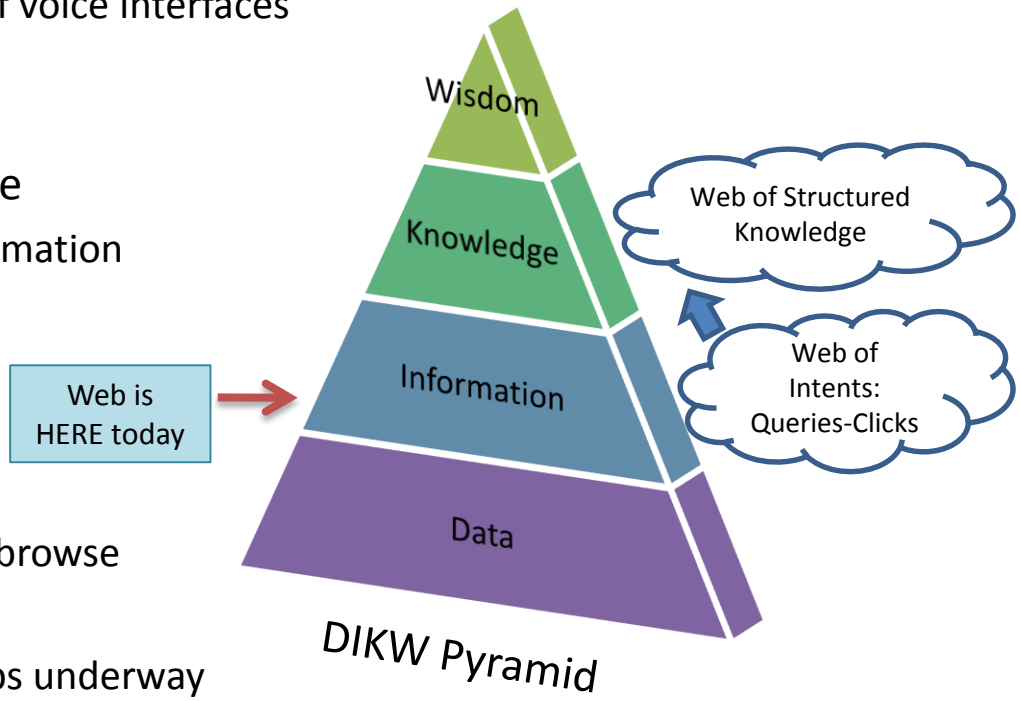

- Where is the natural convergence of speech and the web?

The Opportunity: Convergence

Implications for Spoken Language Technology

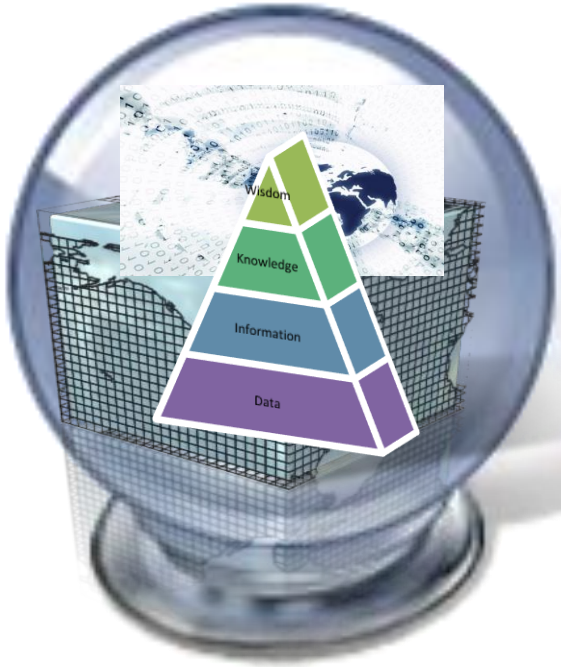

The Opportunity for SLT

Using the web should be as simple as having a conversation

- Web Technologies to Leverage
  - Web of Intents
  - Web of Structured Knowledge

"North Star", Major Research Themes, Progress-to-Date

THE CONVERSATIONAL WEB

Conversational Web

Seamless Conversational Experiences Across Your Life

Conversational Systems
Component View

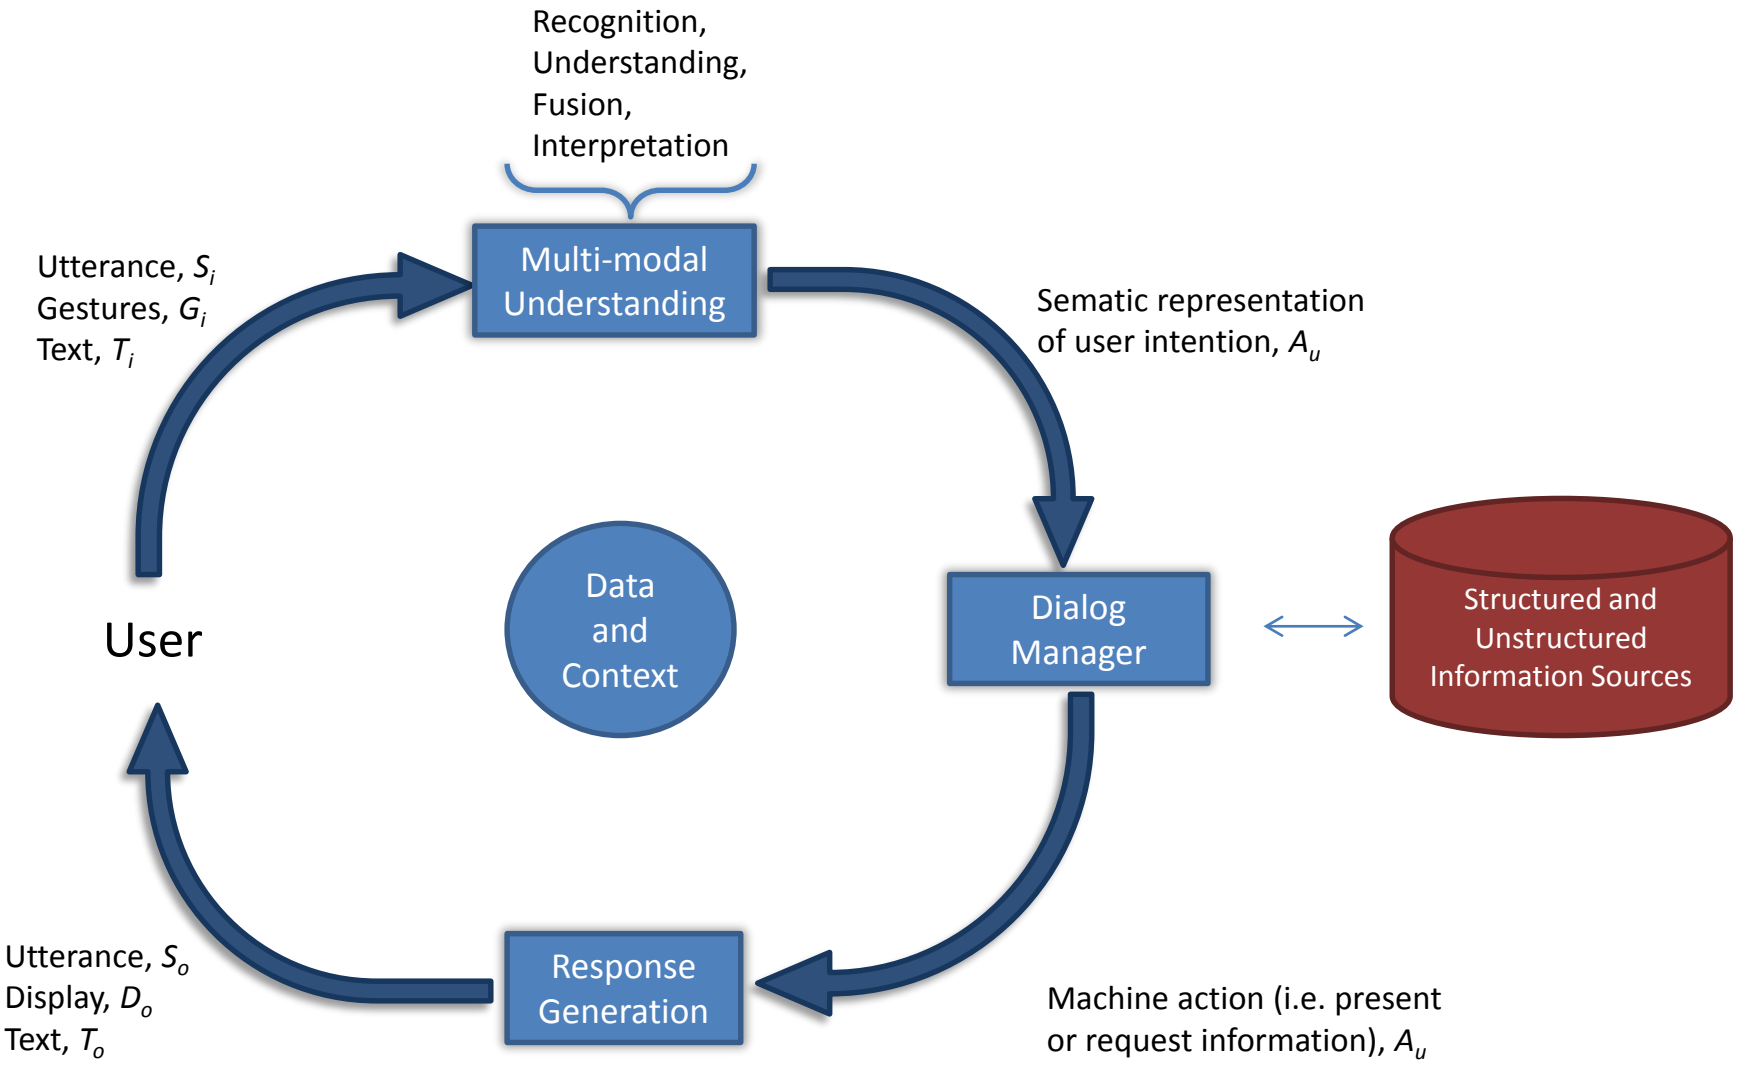

Conversational Web

Major Research Themes: Scaling to Breadth & Depth

- Theme 1: *Adapting to **Context** (visual content, personal, dialog)*
- Theme 2: *Leveraging the **Taxonomy of the Conversational Web***
- Theme 3: *Learning/Training with Big Data - the **Web of Intents***
- Theme 4: *Exploiting the **Web of Structured Knowledge***

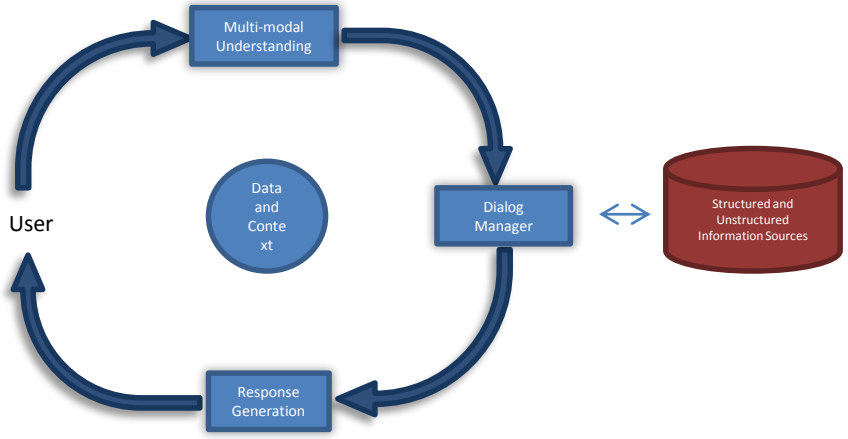

Conversational Web

Major Research Themes: Scaling to Breadth & Depth

- Theme 1: *Adapting to **Context** (visual content, personal, dialog)*
- Theme 2: *Leveraging the **Taxonomy of the Conversational Web***
- Theme 3: *Learning/Training with Big Data - the **Web of Intents***
- Theme 4: *Exploiting the **Web of Structured Knowledge***

Conversational Web

*Theme 1: Adapting to **Context** (visual content, personal, dialog)*

Goals

- Conversational Search & Browsing
  - Web of **Documents** and **Apps**
  - UX Approach: guide user through the open web (**not walled-garden UX**)
- As Natural as a Conversation
  - Natural Combination of **Conversational Speech** and **Gesture**
  - Say Anything, Anywhere, Anytime
  - Open Microphone, Open Grammars (~0.5M words)
- NUI That Scales to the Web
  - Scales: **breadth** of the web & **depth** of key domains
  - Dynamically adapts on-the-fly to the screen and user: grammars, concepts, entities/actions

Conversational Web

*Theme 1: Adapting to **Context** (visual content, personal, dialog)*

Natural Dialog | Cross-Screen | Depth & Breadth | Scalable to New Domains | Cloud-Based

Demo: Conversational Browser for Home TV

Conversation domains: movies, restaurants, weather, hotels& flights, music, videos & web, shopping

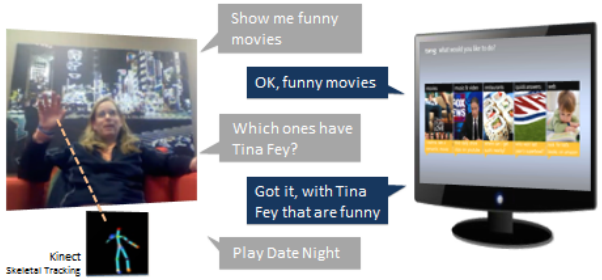

Demo: Personal Assistant for WP7

Other Screens:

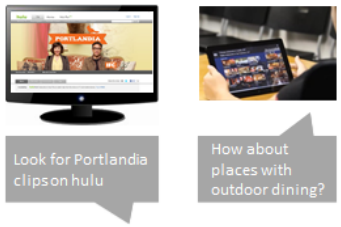

System Structure and Research Areas

Speaker diarization / ID

Joint feature modeling
- e.g., joint speech + gesture

Multimodal feature extraction
- Speech recognition
- Speech prosody
- Gesture, gaze, touch

Addressee detection

Dialog state tracking

Policy implementation

Reasoning

Personalization

Response generation
- Text and other visual feedback
- Spoken feedback

Semantic parsing
- Knowledge discovery & mining
- Domain detection
- Intent determination
- Named entity extraction
- Slot filling
- Question detection

Contextual understanding
- Session modeling
- Click intent detection

Semantic interpretation

Deep linking
- Bing, Office
- 3rd party apps

Microsoft Research

Conversational Web

An Early Prototype

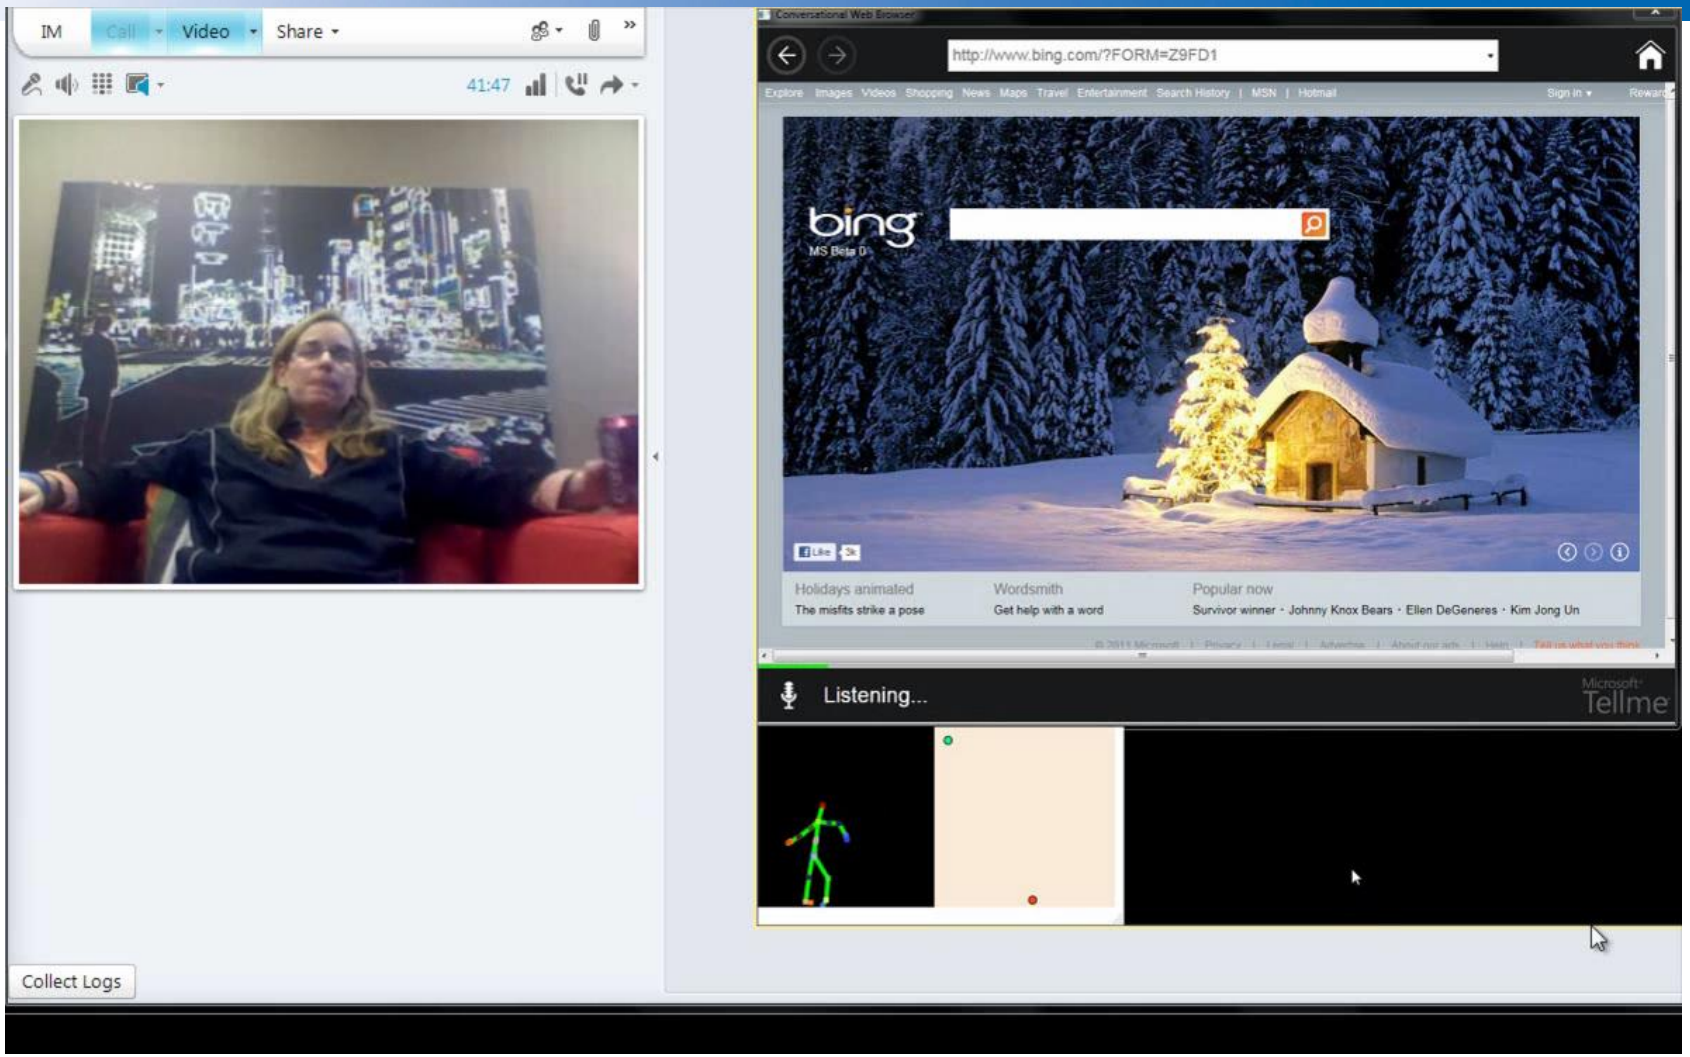

Theme 1: Adapting to **Context**

Conversational Browsing of Web Pages AND Apps

- Dynamic Page/App content affects the SLU
  - Browsing to a new page or App adds 100s of **click intent** actions to ***static SLU***
  - Entities automatically extracted
  - Relevant actions/intents can be retrieved (e.g., Semantic Web)
  - Multi-tiered logic determines final intent

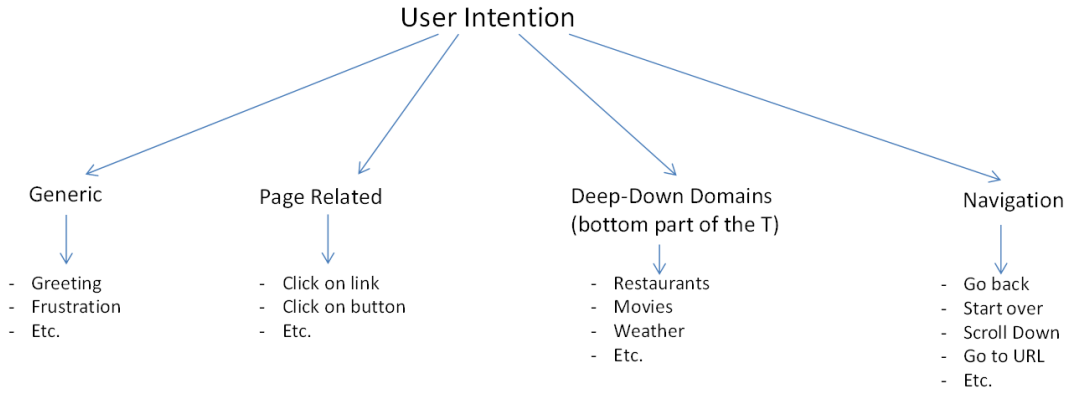

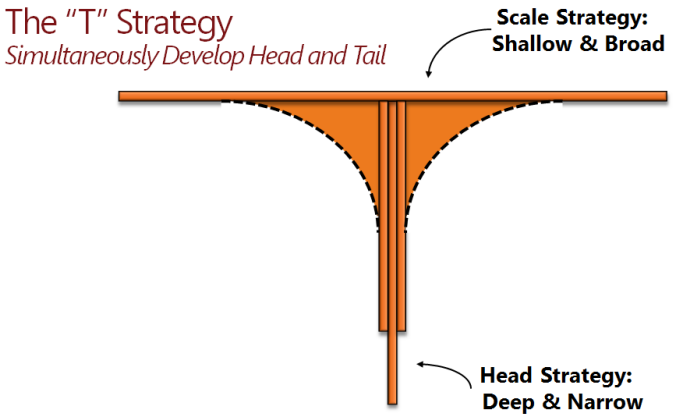

Theme 1: Adapting to **Context**

Deep and Narrow SLU

Spoken Language Understanding (SLU): convert automatic speech recognizer (ASR) output into pre-determined semantic output format

DOMAIN = movies

"when was james cameron's avatar released"

INTENT: Find_release_date
MOVIE NAME: avatar
DIRECTOR NAME: james cameron

Intents	Slots
Find movie	Movie genre
Find showtime	Movie award
Find theater	Theater location
Buy tickets	Number of tickets
...	...

DOMAIN = company

"show me media companies in california"

INTENT: Find_company
LOCATION: california
INDUSTRY: media

Intents	Slots
Find company	Company name
Find revenue	Company address
Find founder	Company revenue
Find contact	Company industry
...	...

Conversational Web

Major Research Themes: Scaling to Breadth & Depth

- Theme 1: *Adapting to **Context** (visual content, personal, dialog)*
- Theme 2: *Leveraging the **Taxonomy of the Conversational Web***
- Theme 3: *Learning/Training with Big Data - the **Web of Intents***
- Theme 4: *Exploiting the **Web of Structured Knowledge***

Conversational Web

*Theme 2: Leveraging the **Taxonomy of the Conversational Web***

Windows 8 Metro Voice Demo

Conversational Web

*Theme 2: Leveraging the **Taxonomy of the Conversational Web***

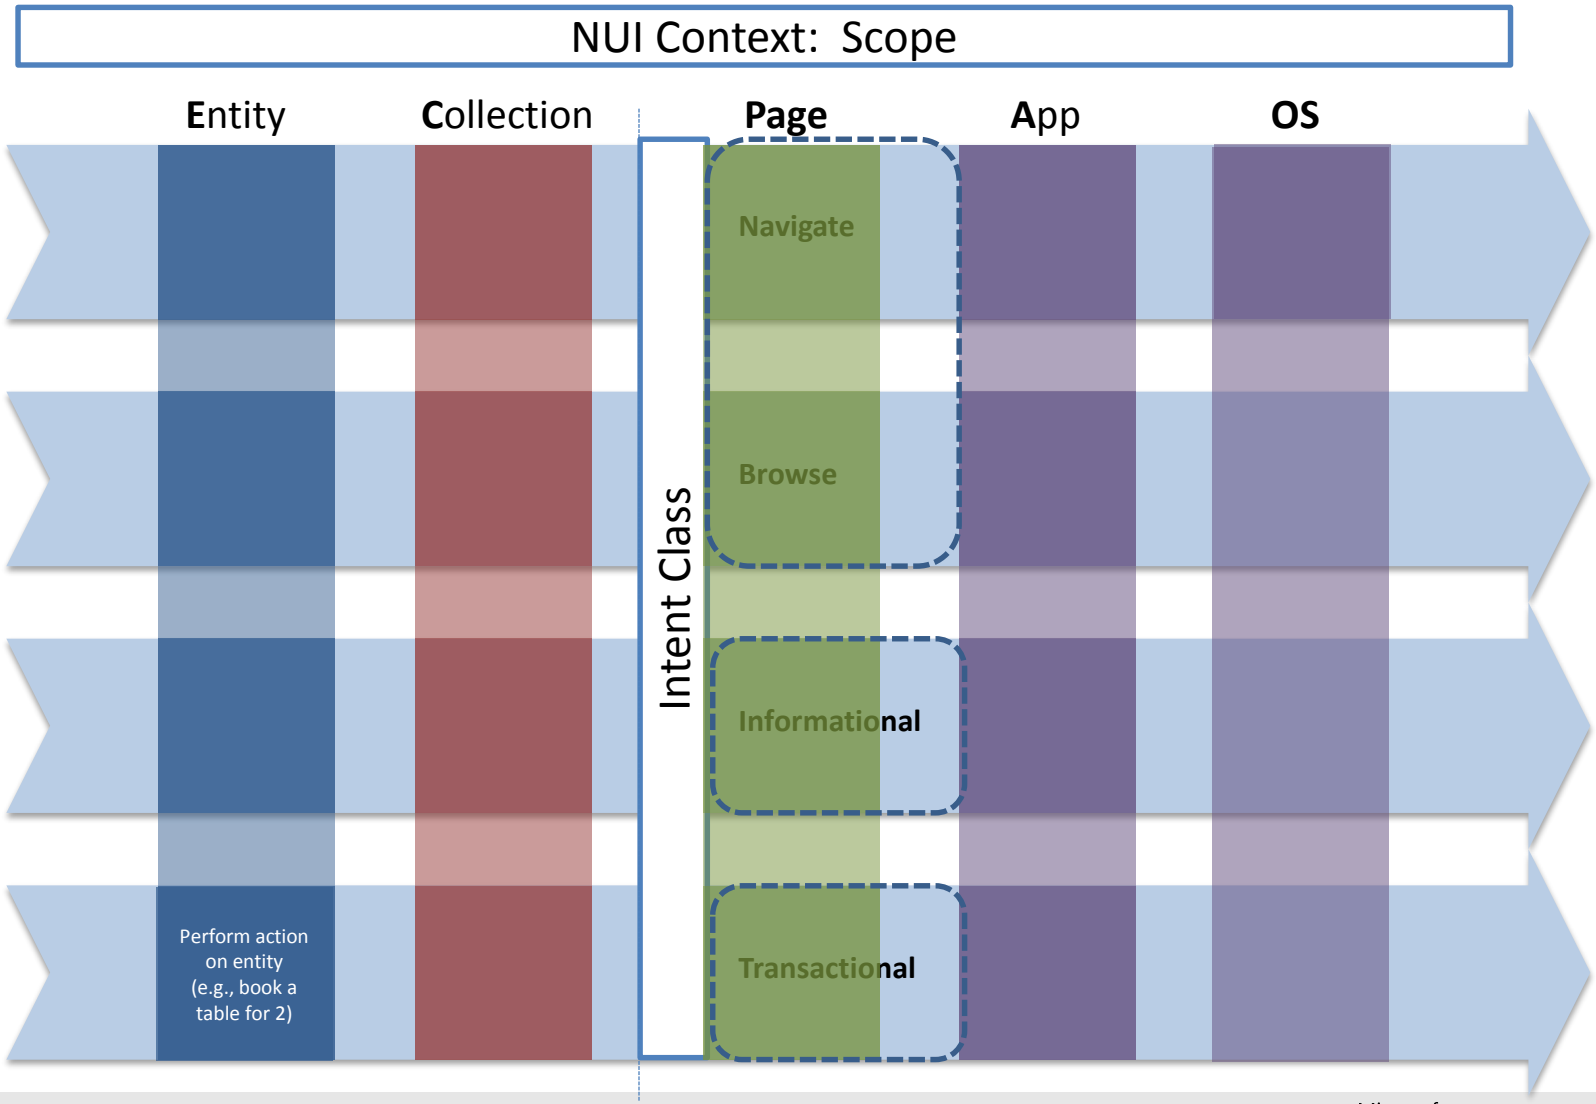

Microsoft Research

Conversational Web

*Theme 2: Leveraging the **Taxonomy of the Conversational Web***

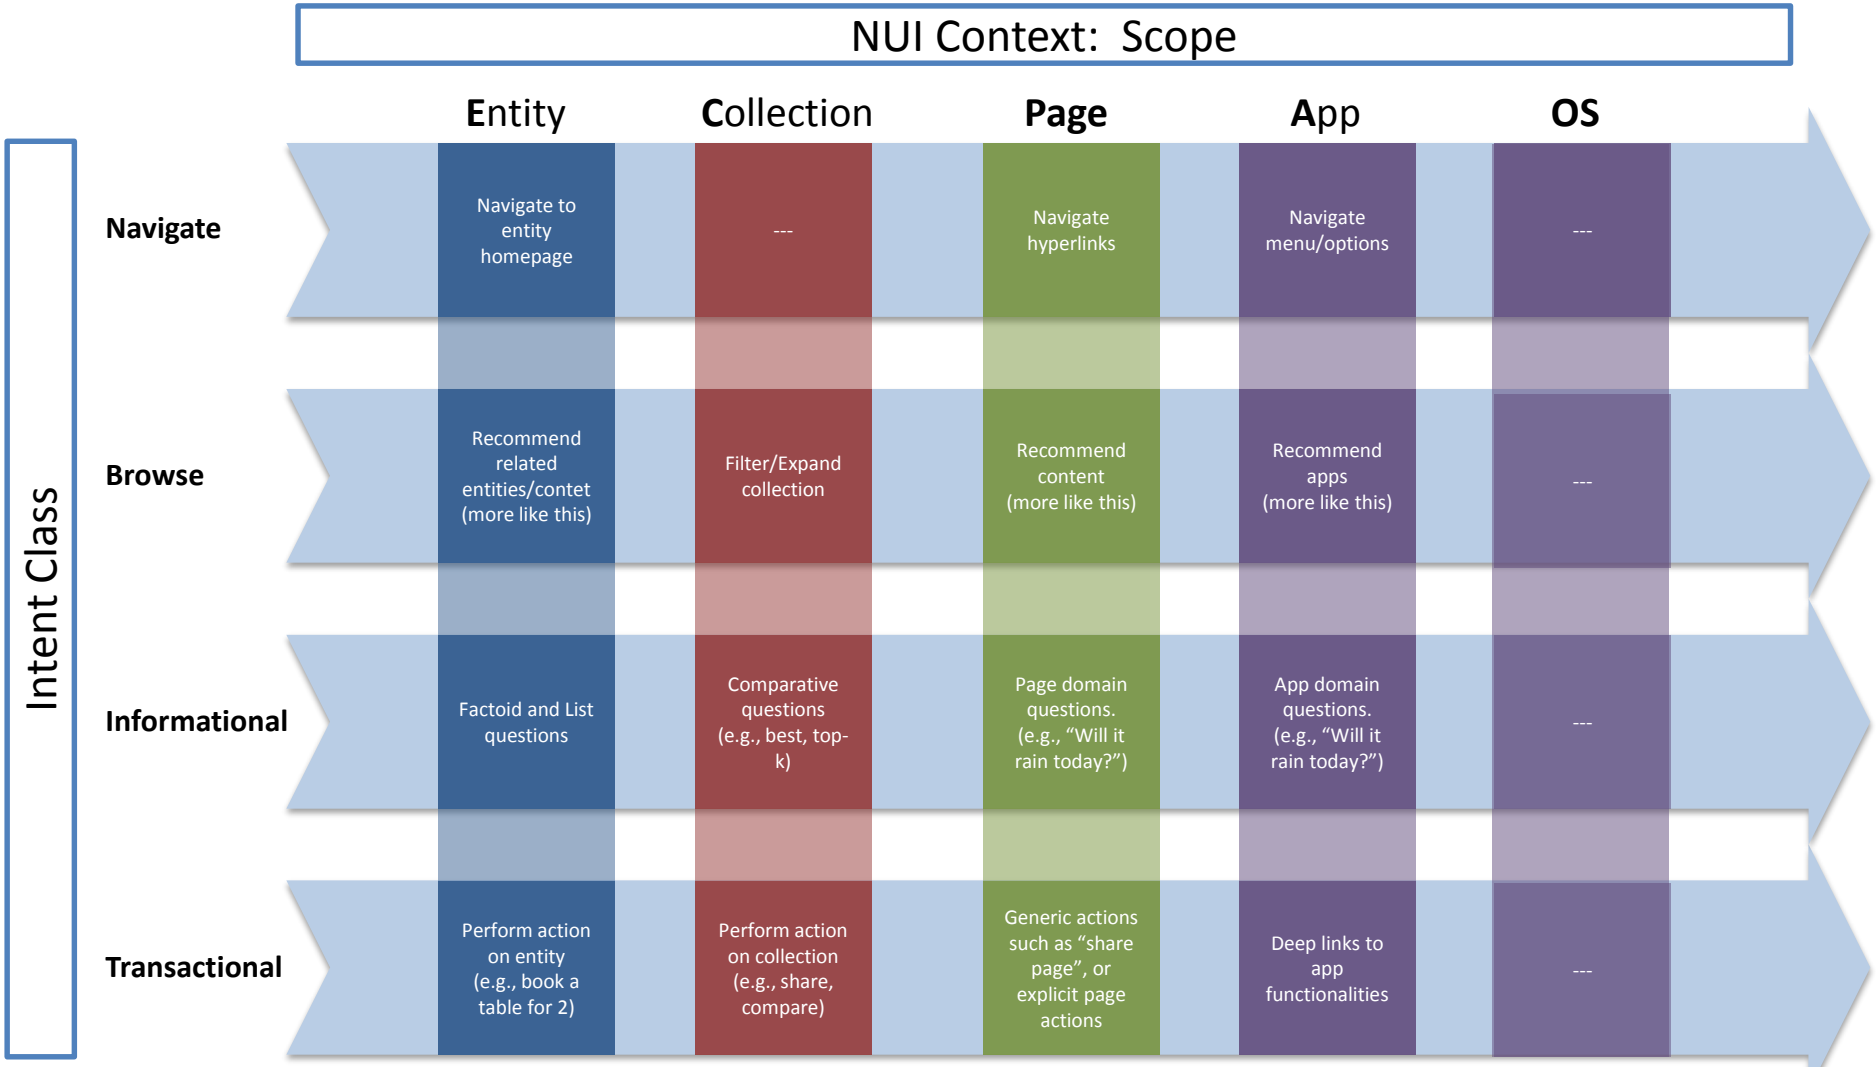

Microsoft Research

Conversational Web

*Theme 2: Leveraging the **Taxonomy of the Conversational Web***

Graph Representation

- **Graphs**: common and natural representation for major components of conversational system
- Algorithms for building, transforming, classifying, and optimizing **graphs** exist
- Weights provide capability to represent uncertainty and relational probabilities in a unified framework

Pivot on Entities

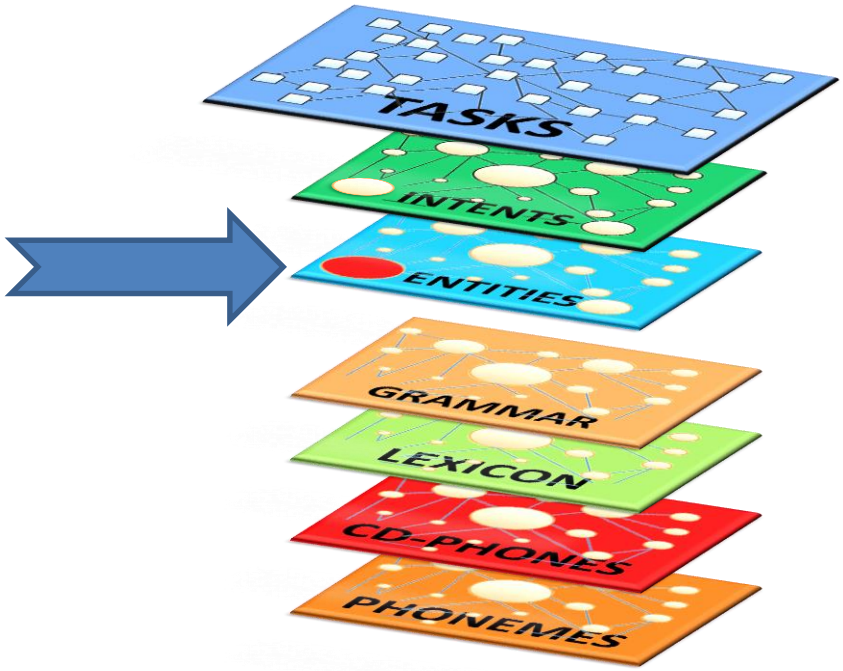

T	Tasks to Intents
I	Intents/Actions to Entities
E	Entities to Grammar
G	Grammar to Lexicon
L	Lexicon to CD Phones
C	CD to CI Phones
H	CI Phones to features (HMMs)

Conversational Web

*Theme 2: Leveraging the **Taxonomy of the Conversational Web***

The User says "Add this to my play list for tonight"

- Touches a picture of the movie Titanic, specifying the **scope context** of the knowledge graph
  - Entity: **Titanic the movie** (extracted from picture of movie box cover)
  - Relations: **Genre (Drama), Director(James Cameron), Release Year (1997), Award(Oscar, Best Director), Cast (Kate Winslet)**
- in the **general context** of
  - Dialog state: **Previous turn "What are the Oscar winning movies from the 90s?"**
  - Personal
    - ontologies of humans in the conversation (all entities-relations): **Pr(like drama | User) = 0.8**
    - histories of interactions: **List is sorted based on user's history(drama, James Cameron movies)**
  - Geospatial
    - Proximity to people, places, things, events
    - Environment
    - Business meeting
    - Social event
    - In-car
    - In-office
    - **At-home**
  - Natural/available modalities: **voice + touch**
  - Time-of-day: **understands that "tonight" means @8pm (patterns)**
  - Season
- with a specific **intent**
  - Information
  - **Transaction: "add this…"**
  - Navigation/Browse

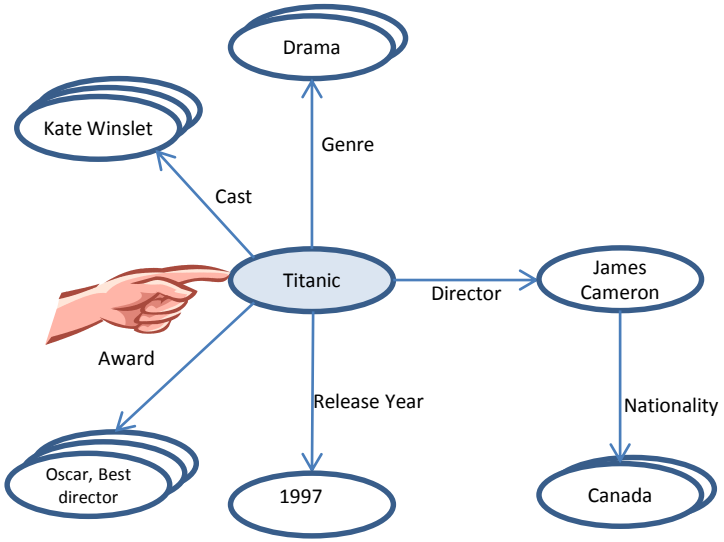

Conversational Web

Major Research Themes: Scaling to Breadth & Depth

- Theme 1: *Adapting to **Context** (visual content, personal, dialog)*
- Theme 2: *Leveraging the **Taxonomy of the Conversational Web***
- Theme 3: *Learning/Training with Big Data - the **Web of Intents***
- Theme 4: *Exploiting the **Web of Structured Knowledge***

Conversational Web

*Theme 3: Learning/Training with Big Data - the **Web of Intents***

- "Web of Intents" – massive (100M queries-clicks/day) data with implicit semantics related to surface forms

- Highly leveraged for machine learning-based web search relevance
  - Queries
  - URLs returned by the search engines and clicked by the users
  - Page Navigation
  - Dwell Times

- Possible to extract information from
  - Multiple users' behavior: high-quality query-click pairs
  - Search sessions
    - Users' reformulation of their queries.
    - Modeling interactions, sequencing of intents.

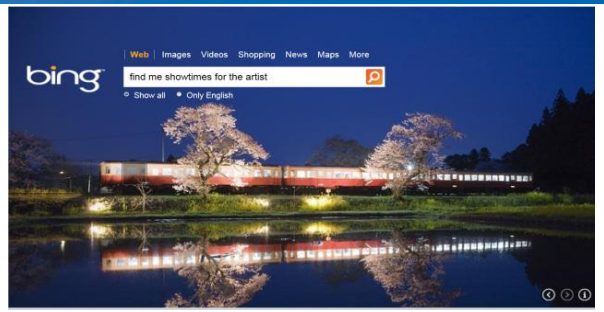

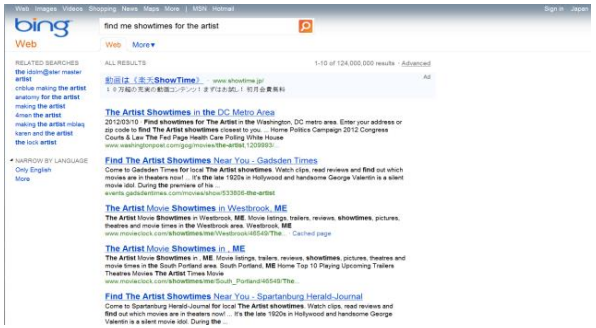

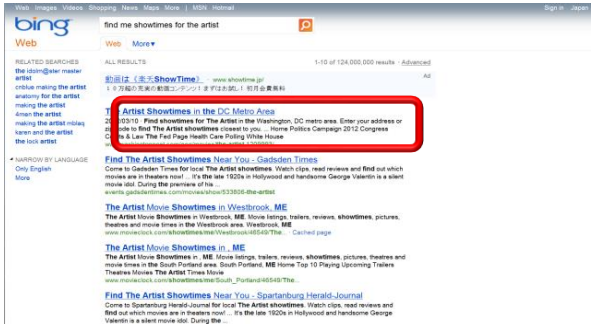

Microsoft Research

Theme 3: Learning/Training with Big Data - the **Web of Intents**

SLU Components to Train

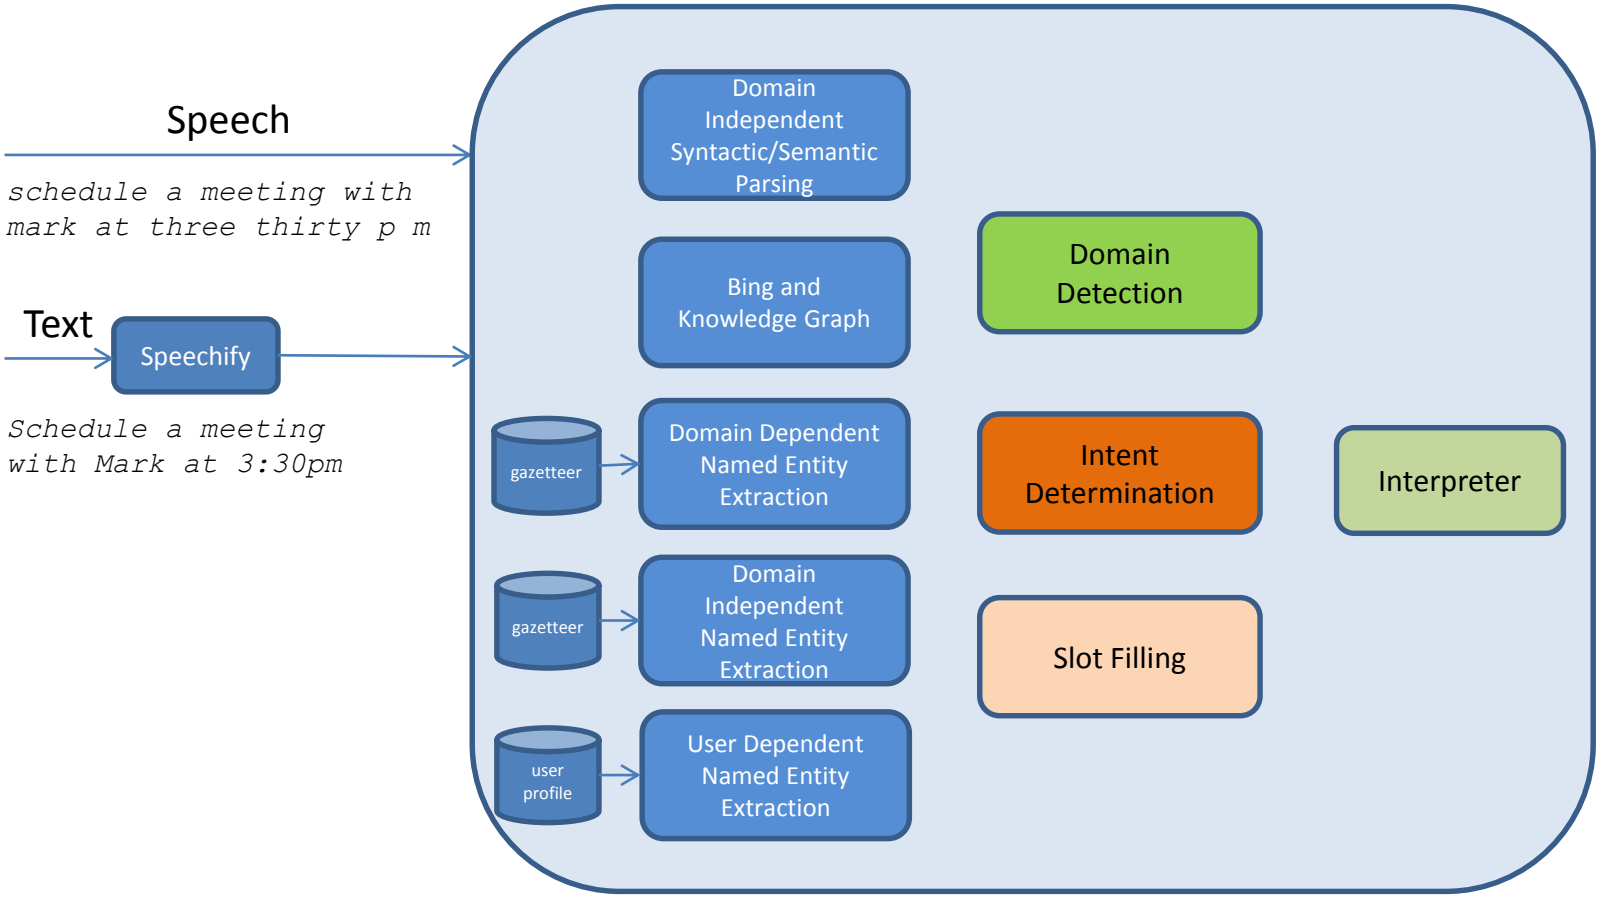


```
<calendar parse="schedule a meeting with mark at three thirty p m">
  <intent parse="add to calendar" />
  <cal_start_time parse="three thirty p m" norm=3:30pm/>
  <cal_attendees parse="mark" norm="mark smith"/>  </calendar>
```

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Domain/Intent Classification: Collaborative Filtering

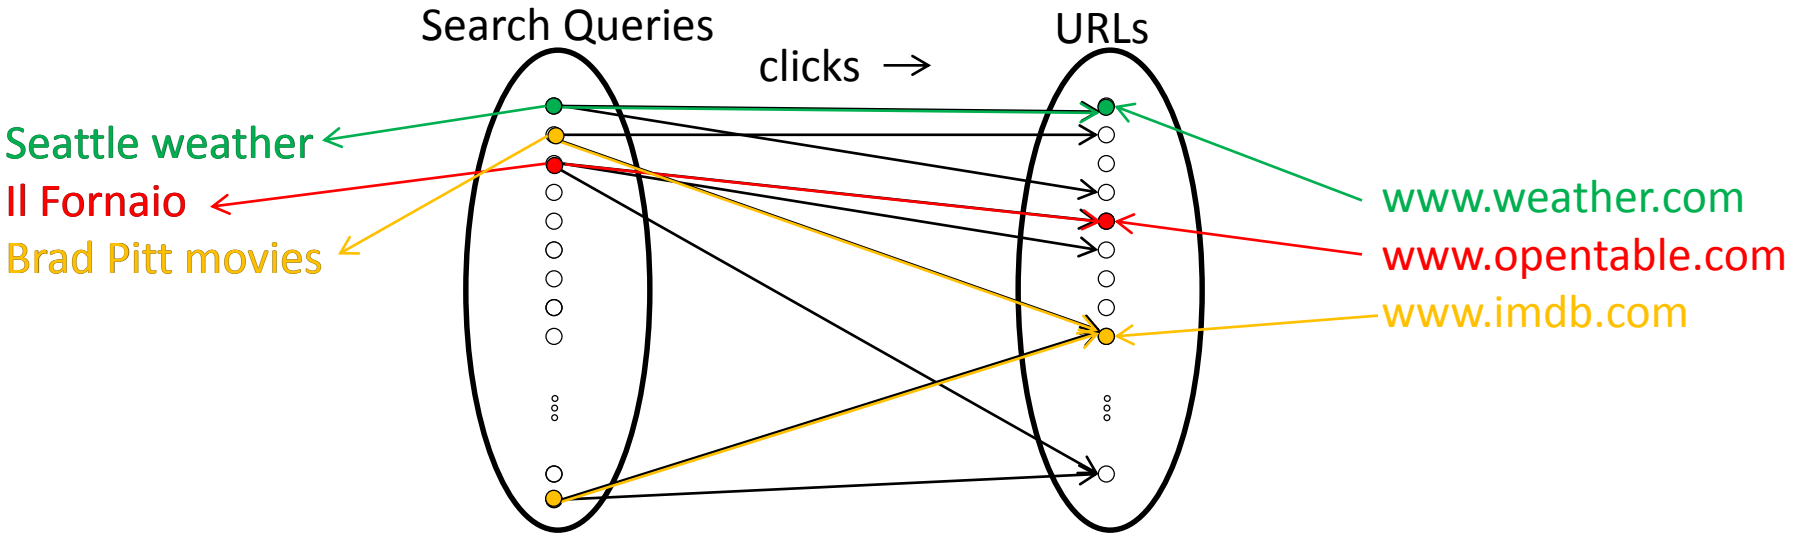

Added to the labeled training set.

Manually labeled SLU queries:

e.g.: *x* = "I wanna reserve a table at Zucca."

y = "restaurant"

OR

x = "What will the weather be like in Menlo Park tomorrow?"

y = "weather"

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Domain/Intent Classification: Collaborative Filtering

Approach: Web search logs (clicks/queries) for unsupervised learning of new domain

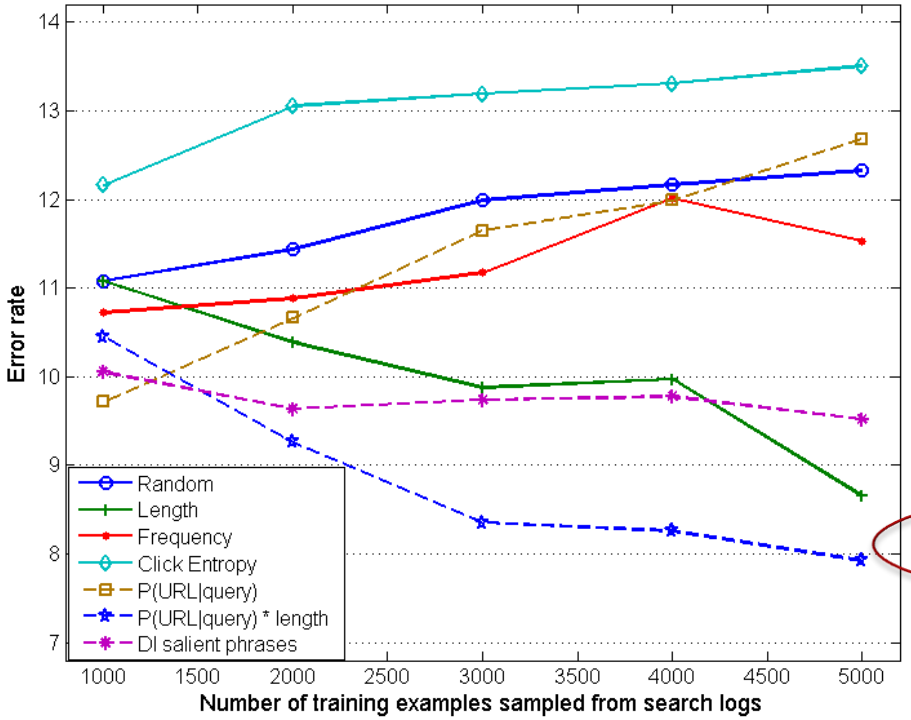

Experiment	Error Rate
Baseline: No domain data	27.5%
All web queries	18.9%
Supervised	6.2%
Unsupervised Method Web search log-based features	**8.0%**

Dilek Hakkani-Tur, Gokhan Tur, Larry Heck, and Elizabeth Shriberg, Bootstrapping Domain Detection Using Query Click Logs for New Domains, August 2011

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Domain/Intent Classification: Label Propagation Approach

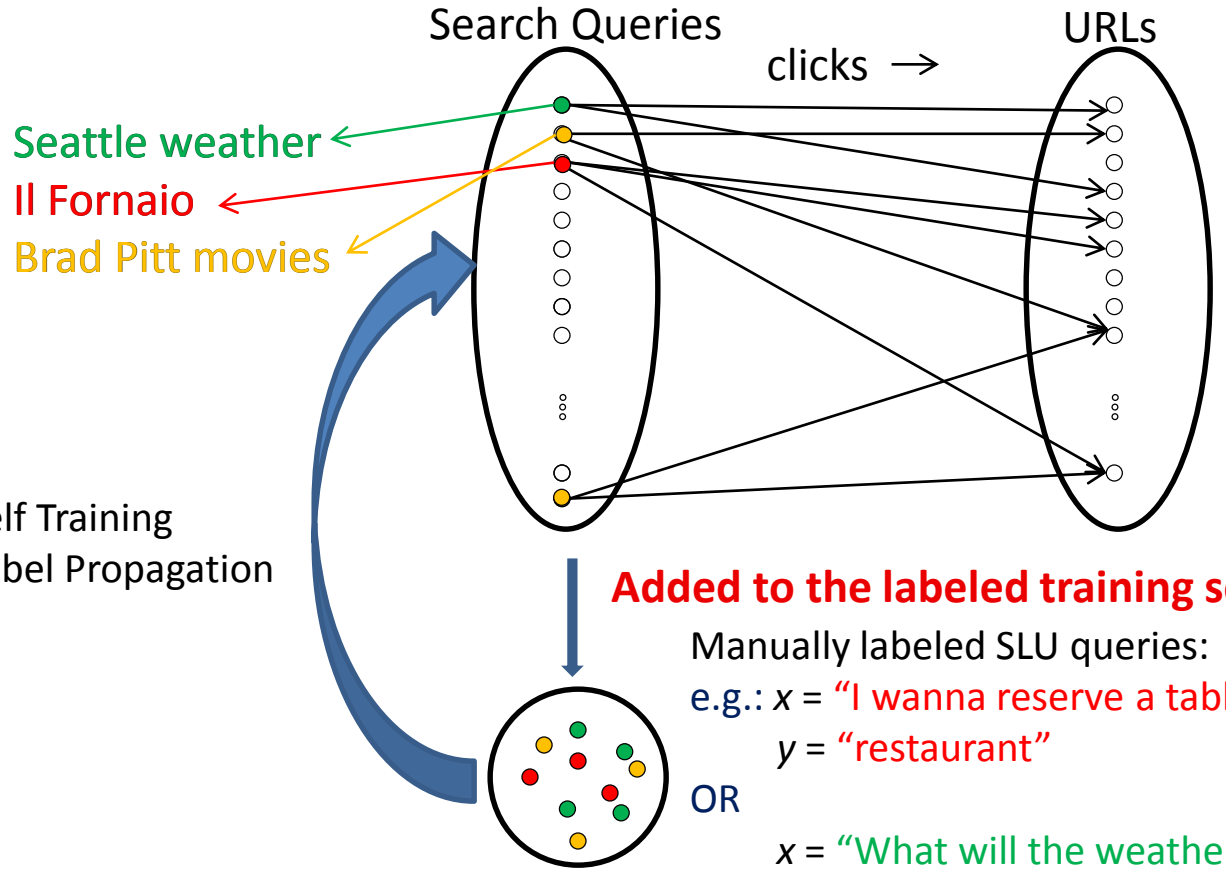

Added to the labeled training set.

Manually labeled SLU queries:

e.g.: *x* = "I wanna reserve a table at Zucca."

y = "restaurant"

OR

x = "What will the weather be like in Menlo Park tomorrow?"

y = "weather"

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Domain/Intent Classification: Label Propagation Approach

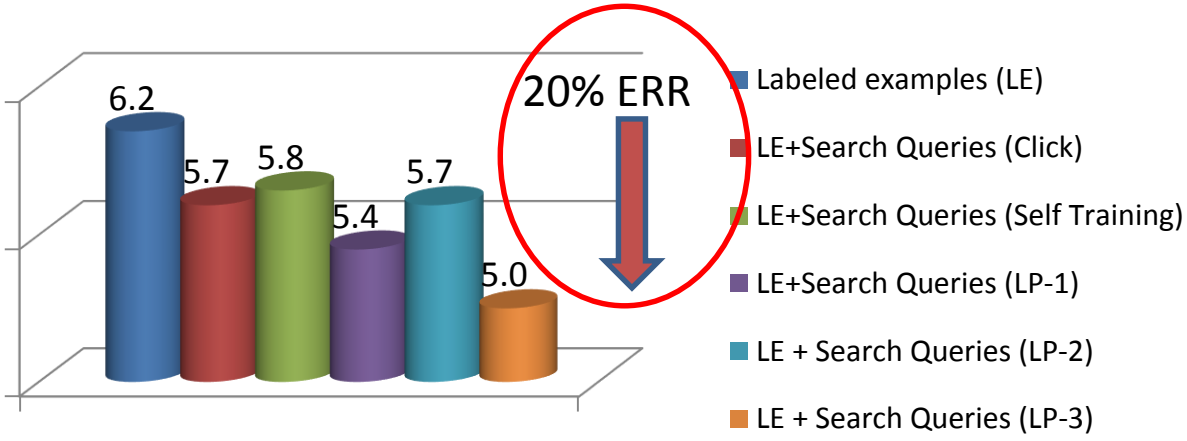

- Inferred domain from clicked URL (Click), Self Training, and Label Propagation (LP)
- Adding Web Search Queries improved error rates for all methods
- Label propagation with noisy supervision had lowest error rate (20% reduction)

Dilek Hakkani-Tur, Larry Heck, and Gokhan Tur, Exploiting Query Click Logs for Utterance Domain Detection in Spoken Language Understanding, in *Proceedings of the ICASSP*, Prague, Czech Republic, May 2011

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Creating "Keyword" from NL Training Queries

Clusters:

show me a resume sample
show me a resume
show me a sample resume
...
resume samples

when is passover celebrated
when is pass over
when is passover 2011
when is passover this year
when is passover
when is passover 2010
when is the passover
...
passover dates

what are my bonds worth
what are my savings bonds worth
what are savings bonds worth
what are bonds worth
what are my series ee bonds worth
what are us savings bonds worth
what are ee bonds worth
what are my us savings bonds worth
what are my ee bonds worth
what are series ee bonds worth
what are savings bonds
what's my savings bond worth
...
savings bond calculator

what are the symptoms of flu
what are symptoms of the flu
what are the flu symptoms
what are flu symptoms
what are the symptoms of the flu
what are flu like symptoms
what are symptoms of flu
what are the symptoms of the flu
what are the symptoms of flu
what are the flu symptoms
...
flu symptoms

NL Search Query (*DISP are in italics)*	Keyword Query
what are the signs of throat cancer	throat cancer symptoms
how many calories do I need in a day	calories per day
what are the biggest us companies	fortune 500 companies
are there any diet pills that actually work	diet pills that work
how do I know if I am anemic	anemic

Similarity between NL search query and the keyword query allows for sorting and filtering pairs.

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Creating "Keyword" from NL Training Queries

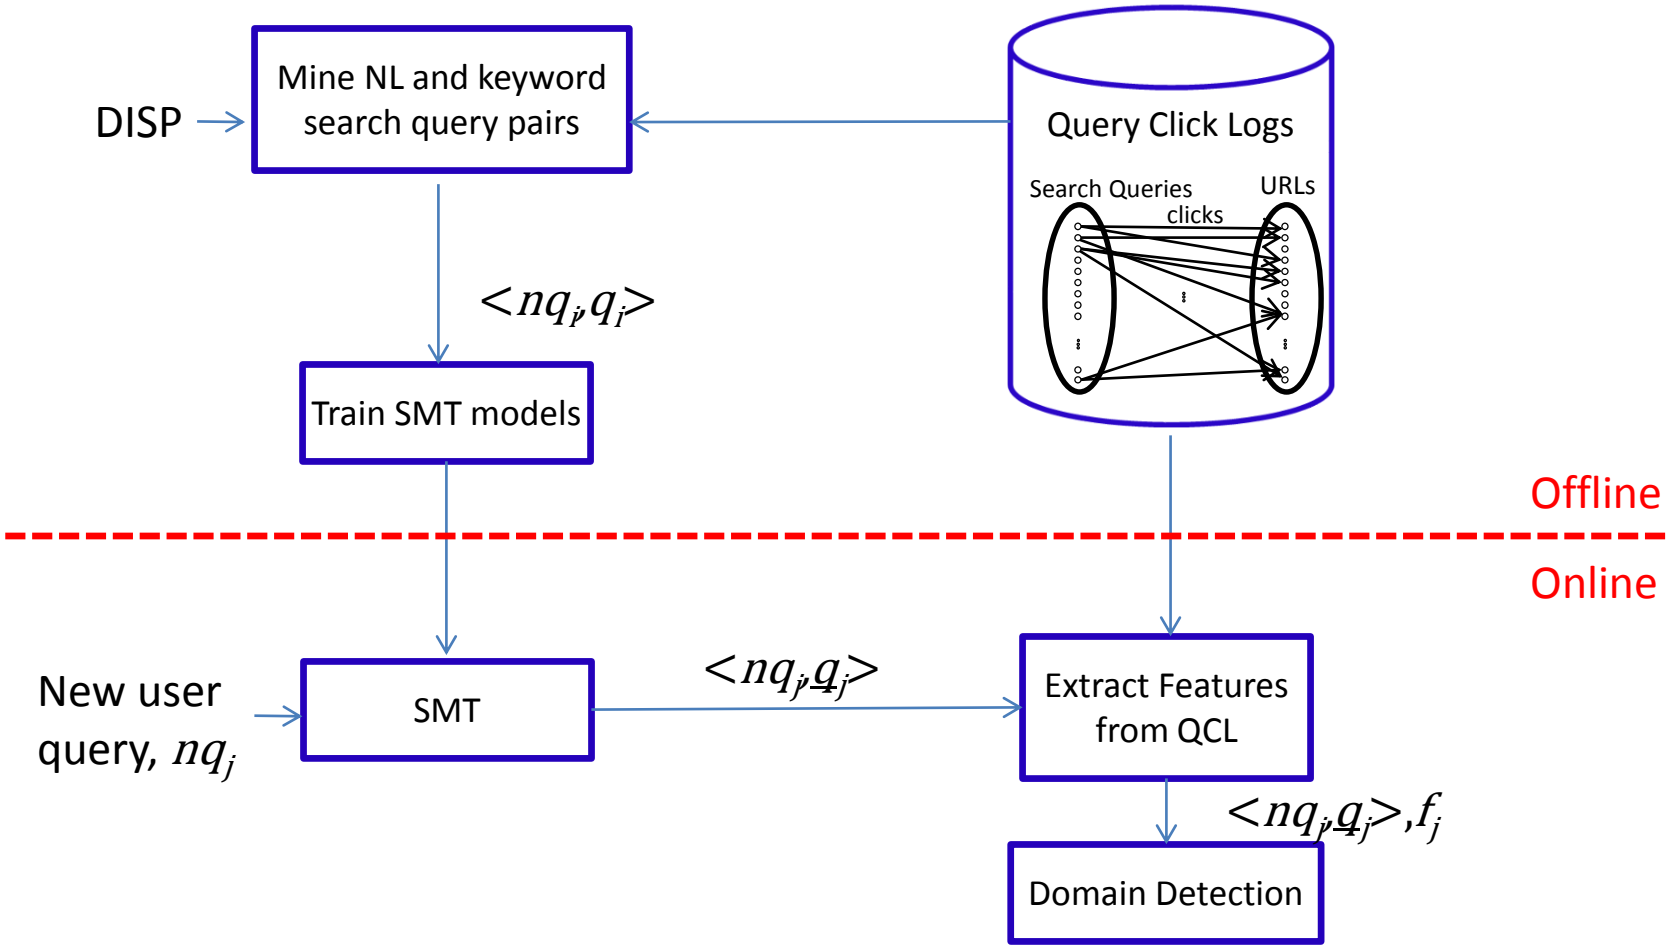

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Creating "Keyword" from NL Training Queries

Approach	Overall ER	ER on NL subset	ER on query-like subset	ER on subset with DISP	ER on subset without DISP
1: Word 1,2,3-grams (n-grams)	10.6%	11.3%	9.3%	9.9%	10.8%
2: n-grams + syntax + SMT-1	9.4%	10.7%	6.8%	10.1%	9.1%
3: n-grams + SMT-2	9.3%	10.9%	6.2%	10.3%	8.9%
4: n-grams + SMT-1+2	**8.5%**	**9.9%**	**5.8%**	**9.2%**	**8.2%**

- Statistical Machine Translation (SMT) of NL to "keyword" queries is promising!
- Primary source of gain: leverage the learning "flywheel" of web searching/browsing

Dilek Hakkani-Tur, Gokhan Tur, Rukmini Iyer, and Larry Heck, Translating Natural Language Utterances to Search Queries for SLU Domain Detection Using Query Click Logs, IEEE ICASSP) March 2012

Conversational Systems
Component View

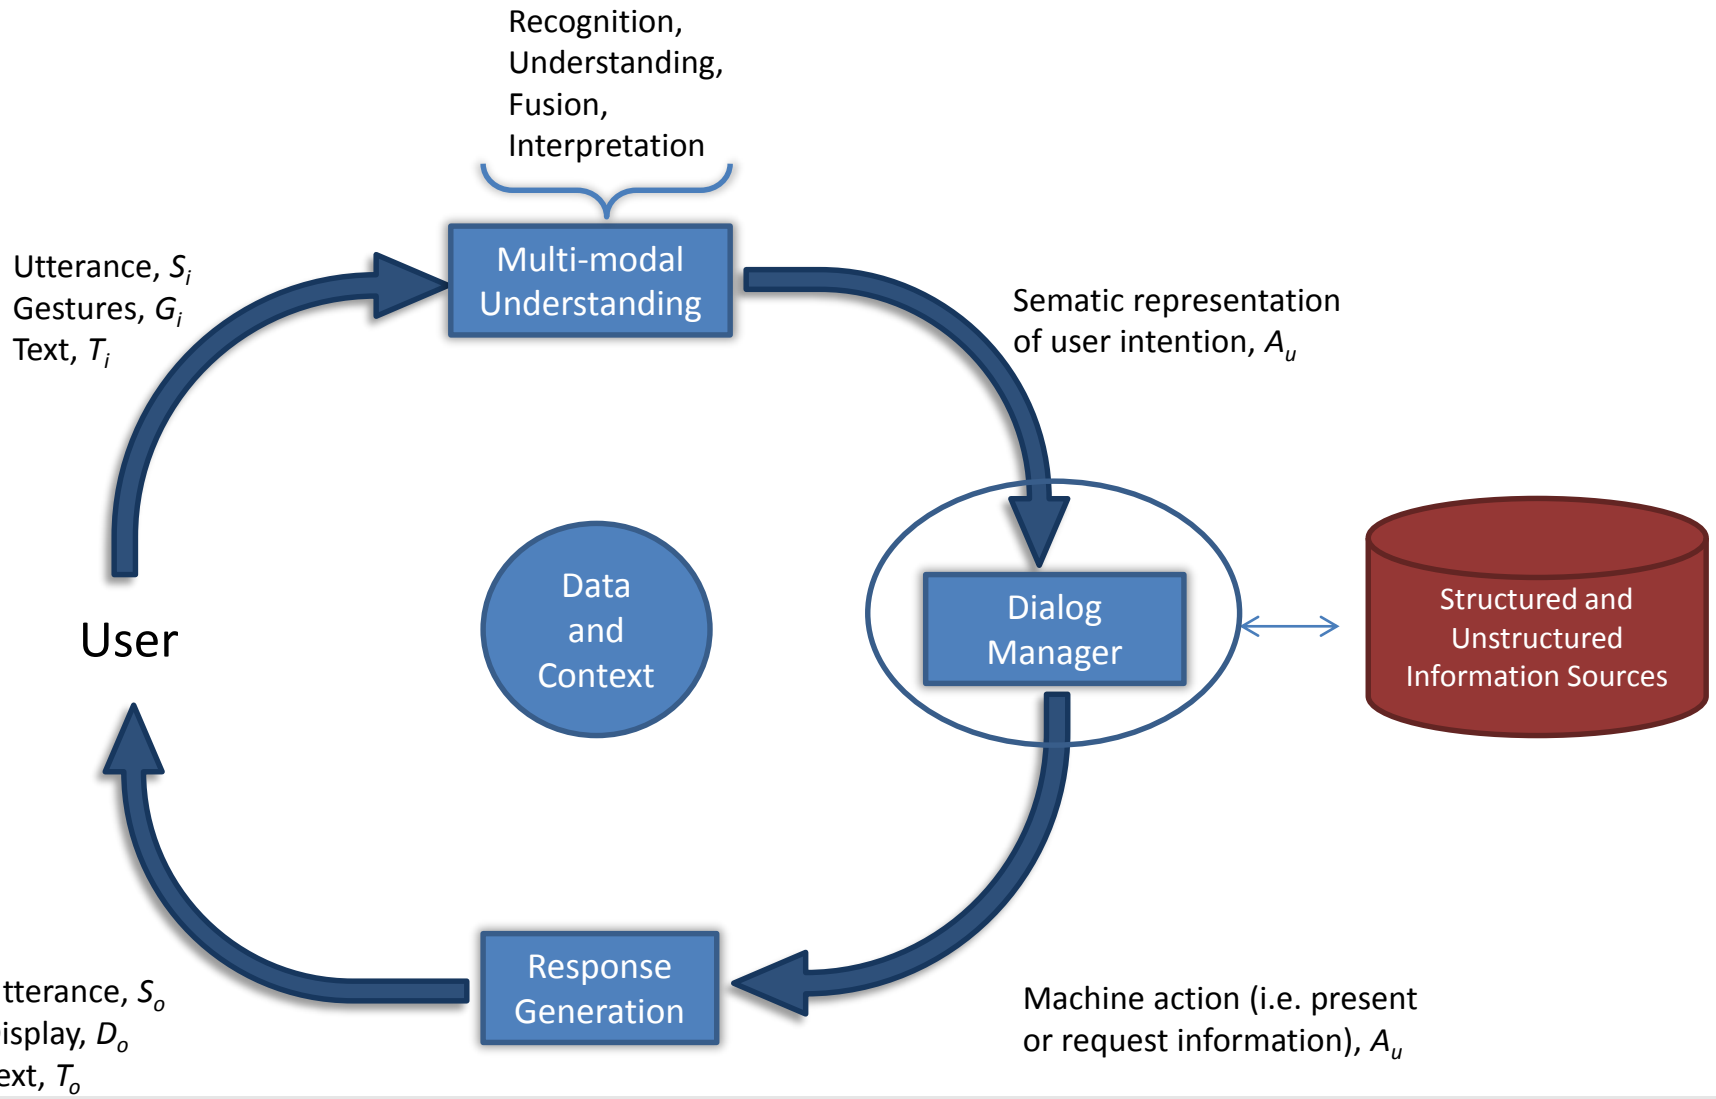

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Statistical Dialog Managers

- Statistical methods for dialog managers is active research topic[1,2] (e.g., POMDP)
- Key Technical Challenge: significant amount of annotated dialogs required for training[2]
- Conversational Web Approach ➔ Leverage **Web of Intents for training data**
  - Users complete tasks through search & browse
  - Advantages:
    - massive volume of interactions > 100M queries/day, Billions of users
    - Complexity of user interactions and breadth in variety of user goals
  - Clicks/Queries can potentially be used to infer:
    - Session success/failure
    - Domains, intents, slots
    - Belief states, User actions

[1] J. Williams, S. Young, B. Thomson, Statistical approaches to dialogue systems, *Interspeech Tutorial, 2009*
[2] P. Crook, "Statistical Dialogue Management for Conversational Spoken Interfaces: How, Why and Scaling-up, *SLT Tutorial, 2012*

Theme 3: Learning/Training with Big Data - the **Web of Intents**

Statistical Dialog Managers

QUERY and TIMESTAMP	GOAL #	MISSION #	DESCRIPTION
hiking; san francisco Tue Apr 17 23:43:17 2007 (4m 17s)	1	1	MISSION 1: Find info on hiking opportunities in and around San Francisco
hiking; san francisco bay area Tue Apr 17 23:47:34 2007 (4m 59s)	1	1	GOAL 1: Find info on hiking trails in San Francisco and the Bay Area
ano nuevo state reserve Tue Apr 17 23:52:33 2007 (7m 54s)	2	1	GOAL 2: Navigate to Ano Nuevo State Reserve and ↓nd out about distances
ano nuevo state reserve; miles Wed Apr 18 00:00:27 2007 (3m 34s)	2	1	
nature trails; san francisco Wed Apr 18 00:04:01 2007 (16m 15s)	1	1	
lobos creek trail Wed Apr 18 00:20:16 2007 (0m 3s)	3	1	GOAL 3: Navigate to Lobos Creek Trail
china camp state park; san rafael Wed Apr 18 00:20:19 2007 (2m 35s)	4	1	GOAL 4: Navigate to China Camp, San Rafael and ↓nd out about distances
china camp; miles Wed Apr 18 00:22:54 2007 (20m 2s)	4	1	
hike; san francisco Wed Apr 18 00:42:56 2007 (3m 19s)	1	1	
fort funston Wed Apr 18 00:46:15 2007 (1h 51m 26s)	5	1	GOAL 5: Navigate to Fort Funston

How to learn spoken dialogs from text/click search and browser sessions is not clear

- Goals/tasks/sequencing similar via search/browse and conversational systems
- Mismatch between the "language" of conversational speech and keywords/clicks

R. Jones and K.L. Klinkner. Beyond the session timeout: Automatic hierarchical segmentation of search topics in query logs. In *Proceedings of CIKM 2008.*

Ahmed Hassan, Rosie Jones, and Kristina Klinkner. 2010. *Beyond DCG: User Behavior as a Predictor of a Successful Search*. WSDM 2010.

Idea: Separate dialog modeling into two parts

1. Learn underlying process of goal/task/sequencing from web

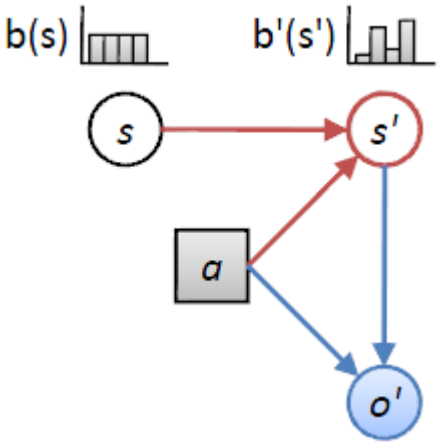

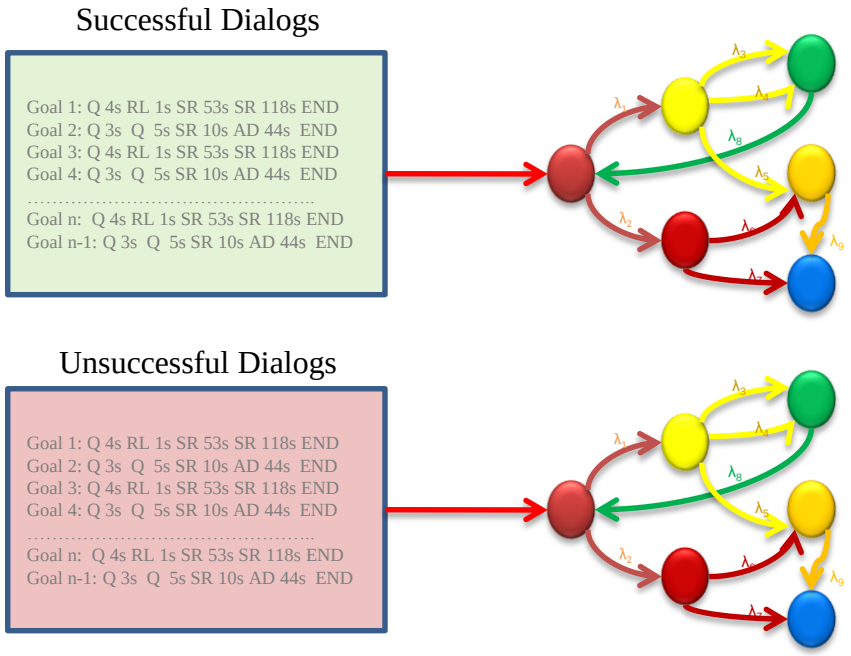

2. Learn **translation** of web queries/clicks into natural spoken conversations
   ➔ apply SMT methods developed for SLU training

Conversational Web

Major Research Themes in MSR

- Theme 1: *Dynamically adapt to **NUI Context** (visual content, personal, dialog)*
- Theme 2: *Scalable **Compositional** Models for Conversational Systems*
- Theme 3: *Exploiting Big Data to tap the **Web of Intents***
- Theme 4: *Exploiting* ***Web of Structured Knowledge***

Conversational Web

Theme 4: Exploiting ***Web of Structured Knowledge***

- Large community working on semantic graphs and semantic search
  - Explosion of structured/semi-structured knowledge on the web
  - Examples: Freebase , DBpedia, YAGO, etc.
  - Google, Bing are making rapid progress to leverage this structure
- Already defined semantic ontologies (www.schema.org June 2011)
- Why reinvent the wheel?

  http://schema.org/Movie

  http://schema.org/MusicAlbum

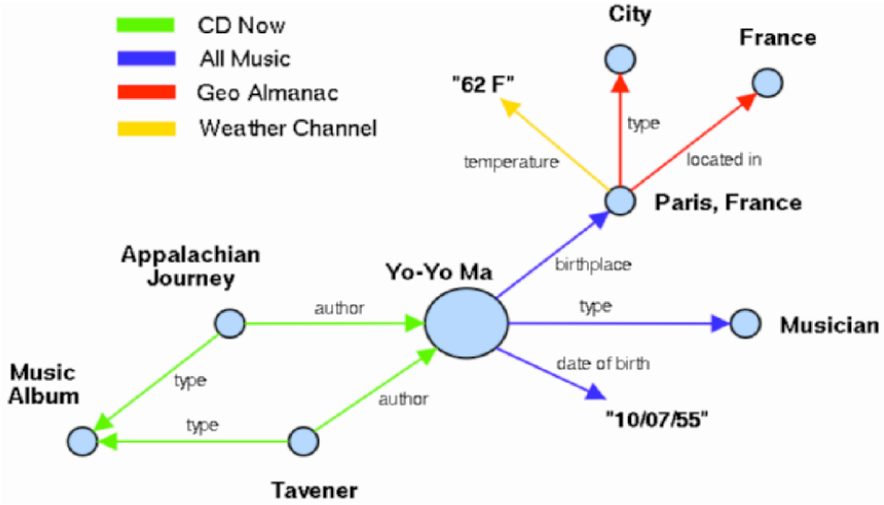

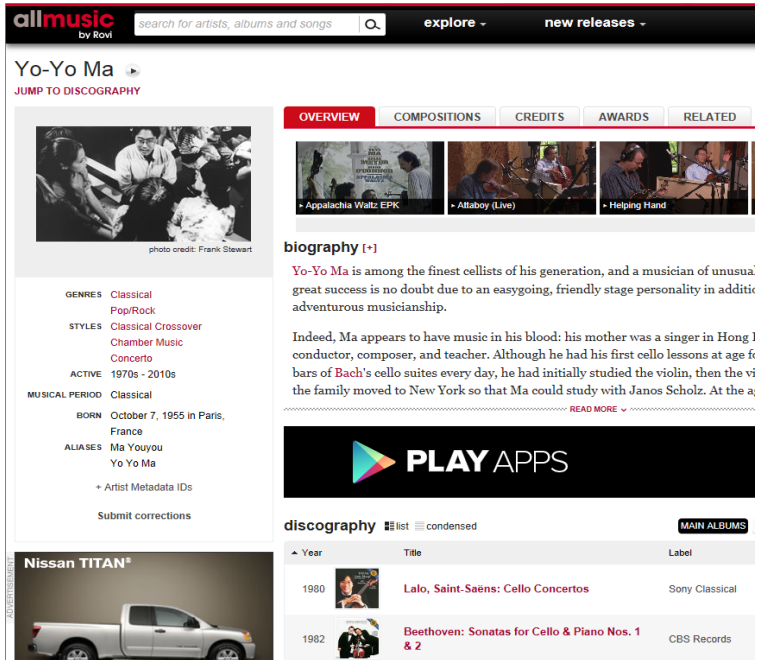

S. A. McIlraith, T. C. Sun, and H. Zeng, “Semantic web services,” IEEE Intelligent Systems, pp. 46–53, 2001.

Conversational Web

Major Research Themes in MSR

- Theme 1: *Dynamically adapt to **NUI Context** (visual content, personal, dialog)*
- Theme 2: *Scalable **Compositional** Models for Conversational Systems*
- Theme 3: *Exploiting Big Data to tap the **Web of Intents***
- Theme 4: *Exploiting* ***Web of Structured Knowledge***

  Unsupervised Learning of Spoken Language Understanding
  - *Approach #1: Enriching Knowledge-Bases with Patterns*
  - *Approach #2: Combining KBs with Search Query Logs*

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #1: Enriching Knowledge-Bases with Patterns

- Knowledge-bases (KBs) Freebase, YAGO, DBpedia include entities and relations:
  - Entities: e.g., movies, organizations, people
  - Relations: e.g., director, founder, release date

- KBs can be enriched with realizations of relations in natural language:

  COMPANY is founded by PERSON

  PERSON, founder of COMPANY

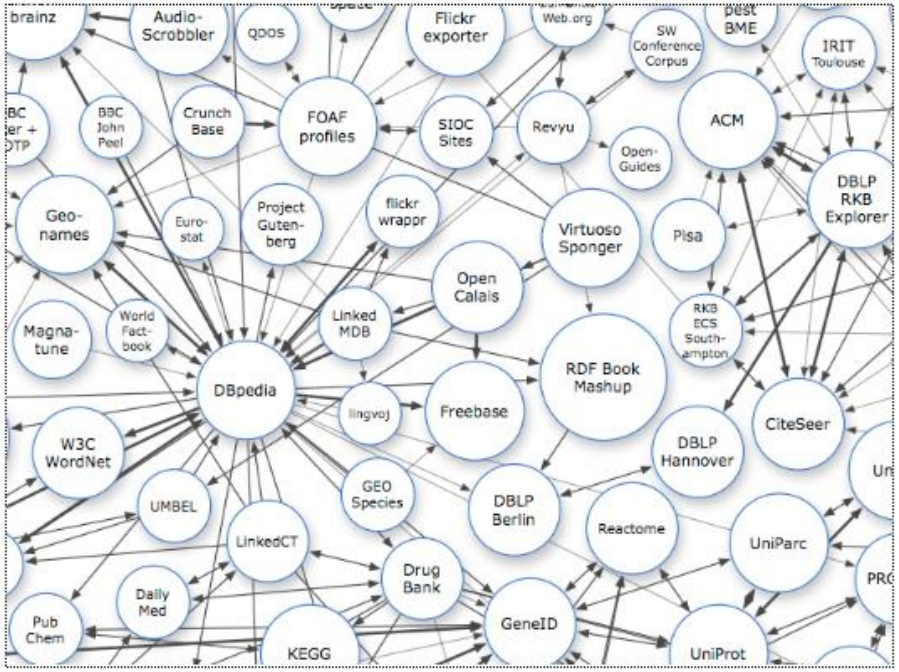

From www.linkeddata.org

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #1: Enriching Knowledge-Bases with Patterns

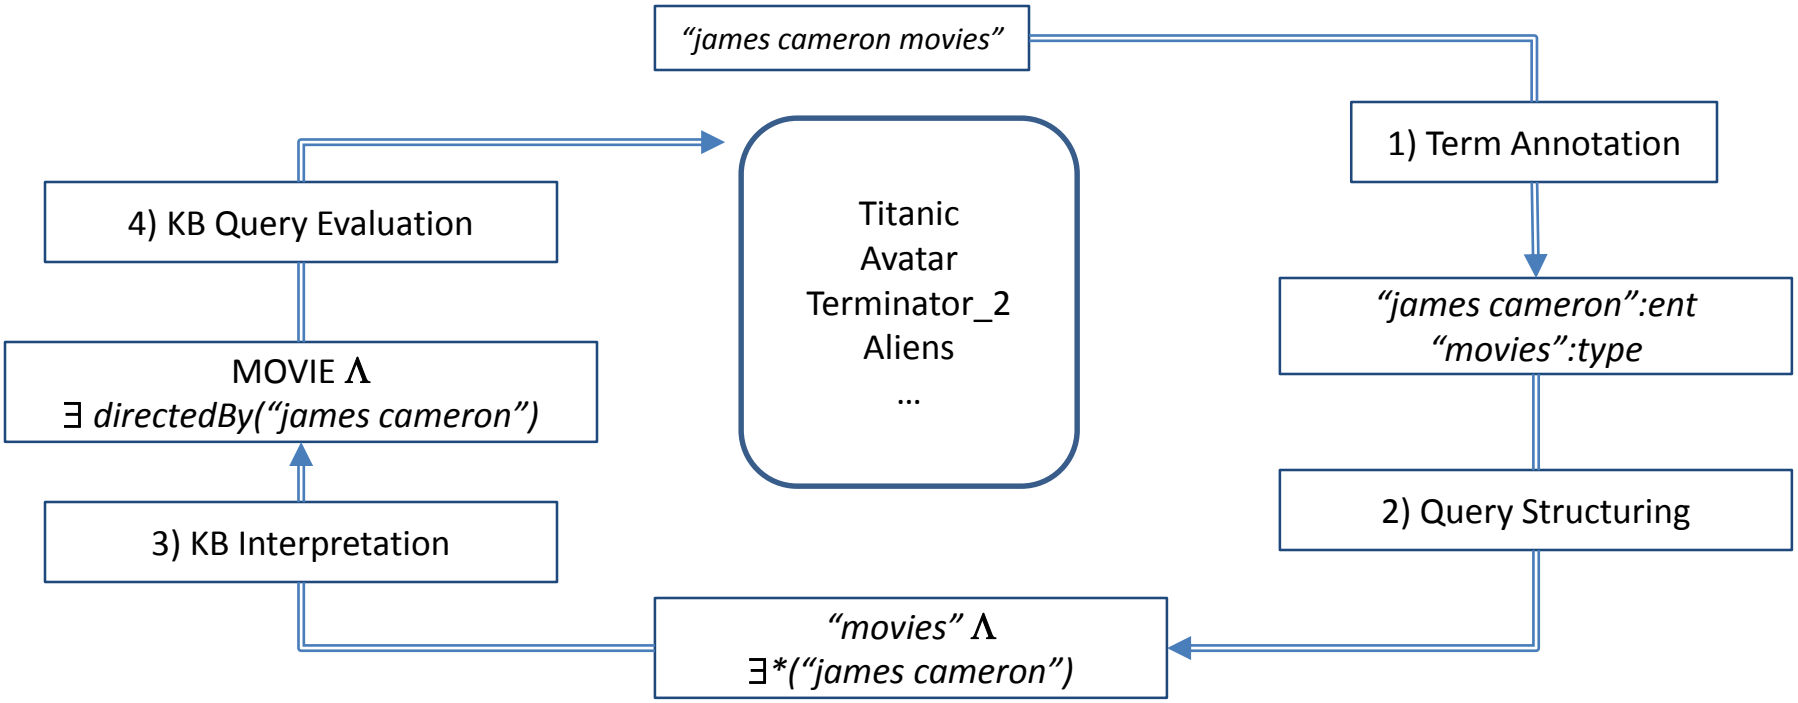

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #1: Enriching Knowledge-Bases with Patterns

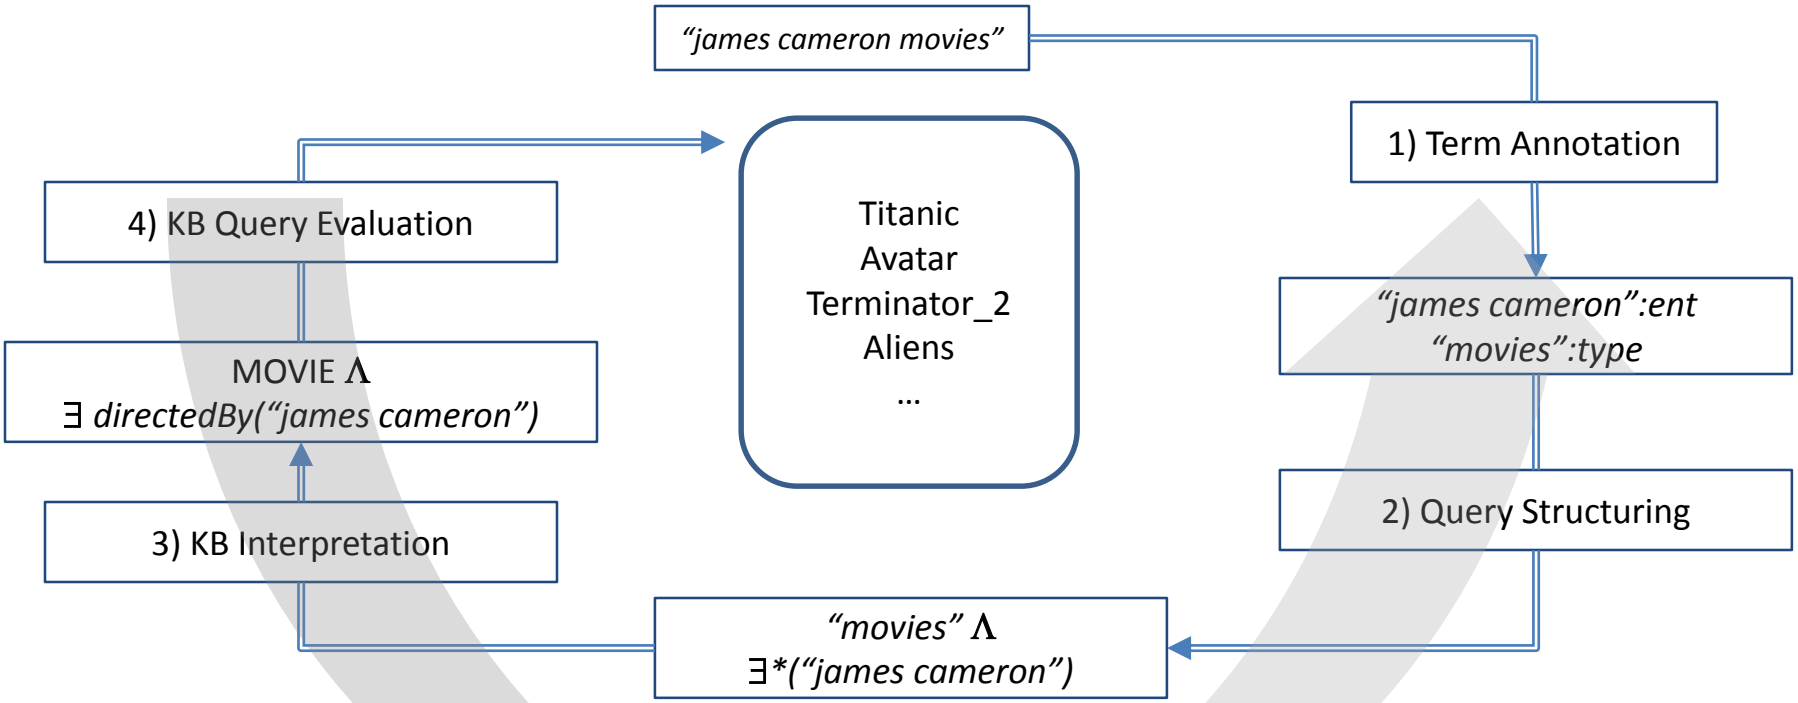

Unsupervised Training

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #1: Enriching Knowledge-Bases with Patterns

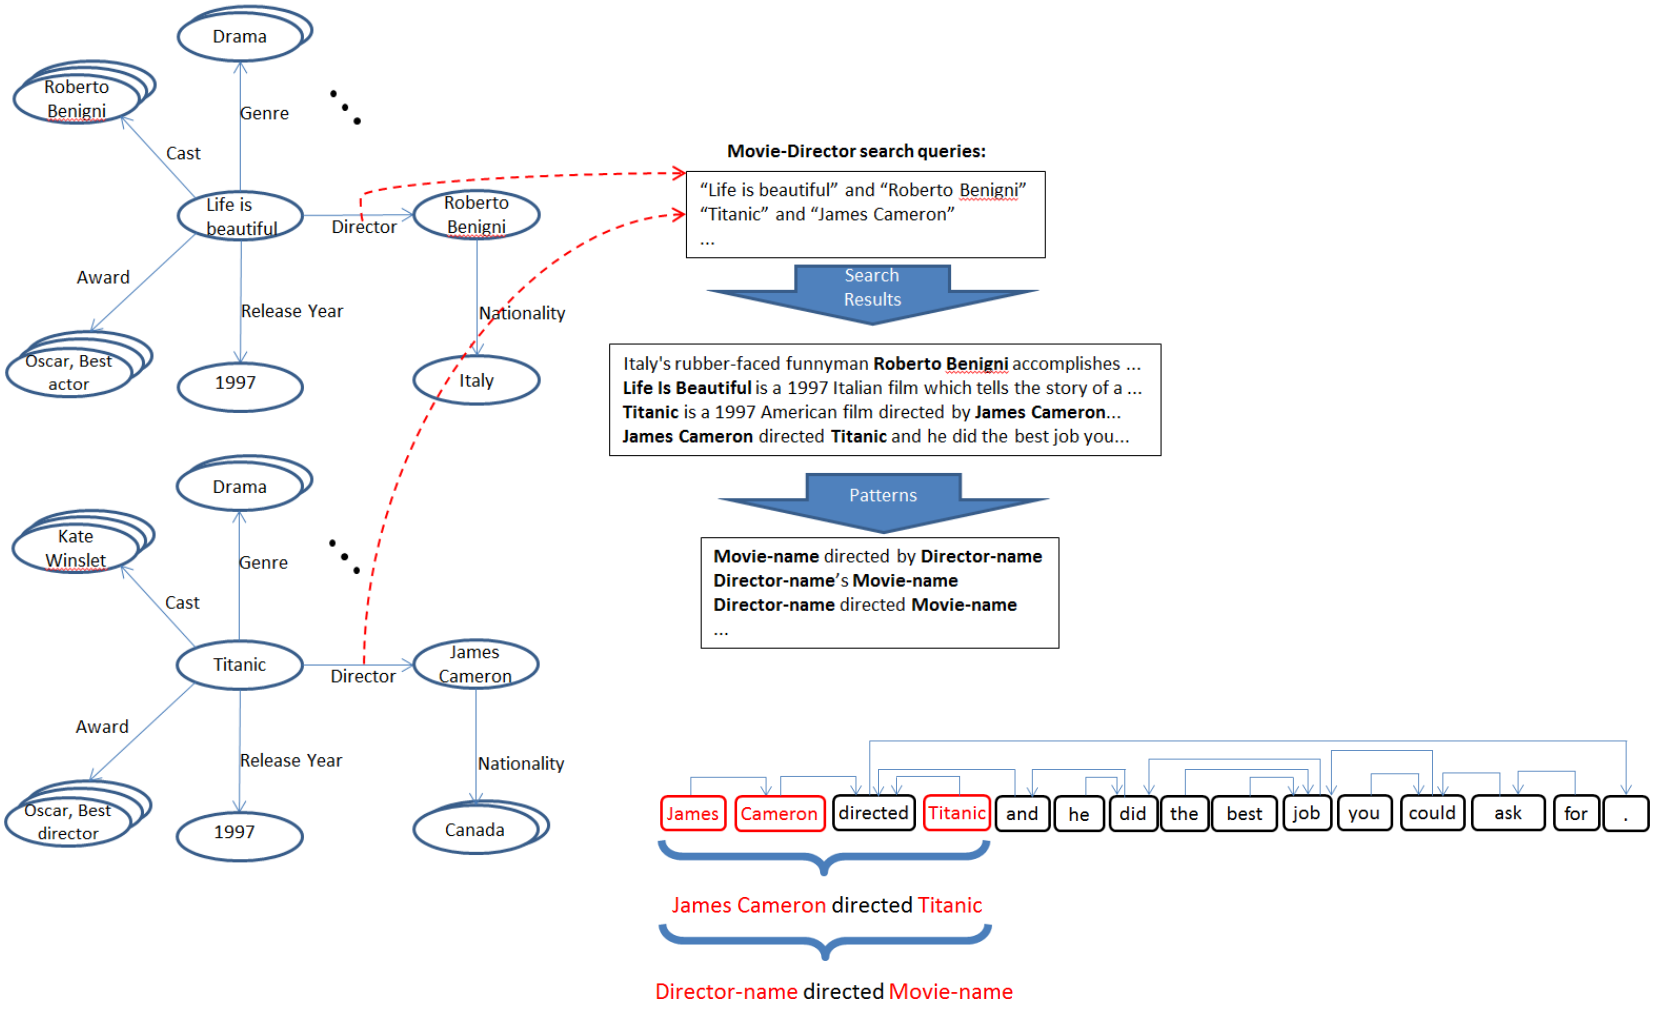

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #1: Enriching Knowledge-Bases with Patterns

Query	Intent
"create a list of the top ten banks by employees"	Find Company (General)
"what is the price of your common preferred and adjustable stock"	Find Stock Information
"show how much money was spent by Microsoft on advertising"	Find Finances
"what are analysts saying about investing in the Coca Cola company stock"	Find News
"can you tell me about the sales revenue from the last quarter"	Find Revenue
"which cell phone model had the largest number of complaints in 2011"	Find Products
"show the highest paid tech CEO and his salary versus company revenue"	Find Leadership
"chart Apple's sales for last year"	Find Annual Sales
"show me any history and info on the treasurer of Dell"	Find People
"find me all the overseas offices for Apple and rank them by highest market cap then by liabilities"	Find Location

	EER	P_{miss}@Pfa=10%
Baseline IDU (no BGM)	36.5%	73.4%
(no BGM, Named Entity Recognition)	35.1%	72.2%
IDU (w/ Supervised BGM)	26.4%	54.3%
IDU (w/ Unsupervised BGM)	**27.0%**	**53.5%**

Larry Heck and Dilek Hakkani-Tur, Exploiting the Semantic Web for Unsupervised Spoken Language Understanding, IEEE Workshop on Spoken Language Technologies (SLT), December 2012.

Thomas Lin, Patrick Pantel, Michael Gamon, Anitha Kannan, Ariel Fuxman. 2012. Active Objects: Actions for Entity-Centric Search. In Proceedings of World Wide Web Conference (WWW-12). pp. 589-598. Lyon, France. [pdf]

Conversational Web

Major Research Themes in MSR

- Theme 1: *Dynamically adapt to **NUI Context** (visual content, personal, dialog)*
- Theme 2: *Scalable **Compositional** Models for Conversational Systems*
- Theme 3: *Exploiting Big Data to tap the **Web of Intents***
- Theme 4: *Exploiting* ***Web of Structured Knowledge***

  Unsupervised Learning of Spoken Language Understanding
  - *Approach #1: Enriching Knowledge-Bases with Patterns*
  - ***Approach #2: Combining KBs with Search Query Logs***

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #2: Combining KBs with Search Query Logs

Semantic Parsing of Structured Web Pages

structured document

Id	Name	City	State
r_1_3671	wild ginger	seattle	washington

+

query log

wild ginger restaurant in seattle wa

AutoLabel

automatically annotated query

```
<restaurant parse="wild ginger restaurant in
seattle wa">
   <name start="0" end="2">wild ginger</name>
   <city start="4" end="5">seattle</city>
   <state start="5" end="6">wa</state>
</restaurant>
```

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #2: Combining KBs with Search Query Logs

Semantic Parsing of Structured Web Pages

structured document

Id	Movie Name	Director	Genre
1392170	the hunger games	gary ross	action; drama; sci-fi; thriller

+

query log

full cast of hunger games with jennifer lawrence

AutoLabel

automatically annotated query

```
<movies parse="full cast of hunger games with jennifer lawrence">
        <name>hunger games</name>
        <actor>jennifer lawrence</actor>
</movies>
```

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #2: Combining KBs with Search Query Logs

- How to define a semantic template the target domain
  - Just use ontologies used by semantic web (like schema.org)
- How to mine relevant data
  - Use NL-like queries hitting to target webpages
- How to annotate that data for modeling
  - Automatically annotate the NL-like queries using the semantic web parses

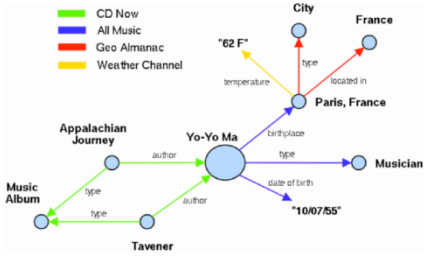

parse

annotate

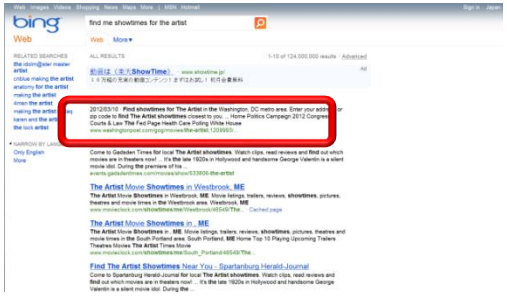

schema

structured web pages

search query click logs

Theme 4: Exploiting **Web of Structured Knowledge**

Approach #2: Combining KBs with Search Query Logs

Unsupervised Slot Filling Results with Unlabeled NL Sentences

	Movie Name	Actor Name	All Slots
Supervised (Upper Bound)	55.22%	81.25%	64.26%
NL-Like	47.94%	84.26%	57.73%
NL-Like + Unlabeled Set	50.21%	85.47%	**60.03%**[1]

NO ANNOTATED DATA NECESSARY!

[1] Gokhan Tur, Minwoo Jeong, Ye-Yi Wang, Dilek Hakkani-Tur, Larry Heck,
Exploiting the Semantic Web for Unsupervised Natural Language Semantic Parsing, Interspeech 2012.

Conversational Web

Summary and Conclusions

- Convergence of Conversational Systems, Web Search, and Web Application

- Theme 1: *Adapting to **Context** (visual content, personal, dialog)*
  - Provides not only better user experience, but much better accuracy/scalability
- Theme 2: *Leveraging the **Taxonomy of the Conversational Web***
  - Multi-modal NUI adds new dimensions to the "taxonomy of the web"
  - Touch/Gesture combined with speech are powerful constraints
- Theme 3: *Learning/Training with Big Data - the **Web of Intents***
  - Massive source of unsupervised/semi-supervised learning from the web
  - Described SLU methods to add training data /generate features : very promising results!
  - Proposed new idea to leverage Web of Intents for training statistical dialog managers
- Theme 4: *Exploiting the **Web of Structured Knowledge***
  - Bridging the research communities of semantic knowledge graphs and NLU
  - Unsupervised Training: using no manual annotations nearly matching accuracy of supervised
  - Consistent with the semantic web representation of the target domain, no interpretation issues

Conversational Web: SLT can benefit greatly from leveraging the web

Microsoft Research

Thank You!

Larry Heck

Contributors:
Dilek Hakkani-Tür, Gokhan Tur, Andreas Stolcke, Ashley Fidler, Lisa Stifelman, Liz Shriberg, Madhu Chinthakunta, Malcolm Slaney, Patrick Pantel, Benoit Dumoulin

Microsoft