Information Retrieval and Extraction

Berlin Chen

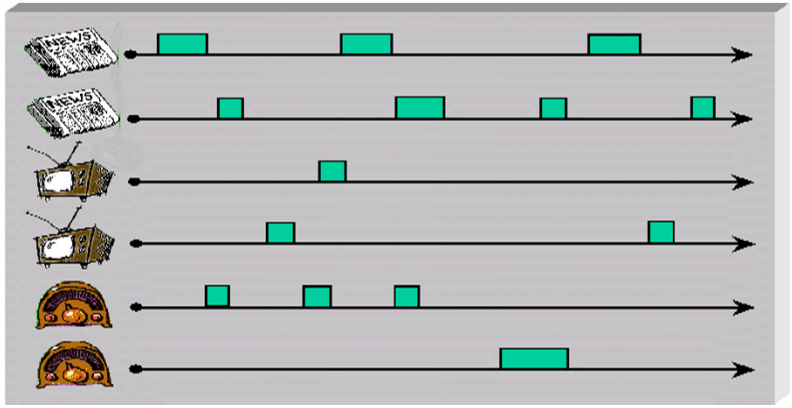

(Picture from the TREC web site)

Objectives of this Course

- Elaborate on the fundamentals of information retrieval (IR), a almost *fifty-year-old* discipline
  - Indexing, search, relevance, classification, organization, storage, browsing, visualization, etc.
- Focus on prominent *computer algorithms* and *techniques* used in IR systems from a computer scientist's perspective
  - How to provide users with easy assess to information of interest
  - Rather than from a "librarian" perspective that put great emphasis on *"human-centered"* studies (e.g., user behaviors, psychology, etc.)
- Practical Issues on the Web
  - Crawling, retrieval, and ranking of Web documents
  - Electronic commerce; security, privacy, copy rights and pattern rights; multimedia and cross-language retrieval; digital libraries

Textbook and References

- Textbooks
  - R. Baeza-Yates and B. Ribeiro-Neto. ***Modern Information Retrieval: The Concepts and Technology behind Search (2nd Edition)***, ACM Press, 2011
  - Christopher D. Manning, Prabhakar Raghavan and Hinrich Schütze, ***Introduction to Information Retrieval***, Cambridge University Press, 2008
  - W. Bruce Croft, Donald Metzler, and Trevor Strohman, ***Search Engines: Information Retrieval in Practice***, Addison Wesley, 2009

- References
  - C.X. Zhai, ***Statistical Language Models for Information Retrieval*** (Synthesis Lectures Series on Human Language Technologies),"Morgan & Claypool Publishers, 2008
  - W. B. Croft and J. Lafferty (Editors). ***Language Modeling for Information Retrieval***. Kluwer-Academic Publishers, July 2003
  - D. A. Grossman, O. Frieder, ***Information Retrieval: Algorithms and Heuristics***, Springer. 2004
  - I. H. Witten, A. Moffat, and T. C. Bell. ***Managing Gigabytes: Compressing and Indexing Documents and Images.*** Morgan Kaufmann Publishing, 1999
  - C. Manning and H. Schutze. ***Foundations of Statistical Natural Language Processing***. MIT Press, 1999

Motivation (1/2)

- Information Hierarchy
  - ***Data***
    - The raw material of information
  - ***Information***
    - Data organized and presented by someone
  - ***Knowledge***
    - Information read, heard or seen and understood
  - ***Wisdom***
    - Making appropriate use of distilled and integrated knowledge and understanding

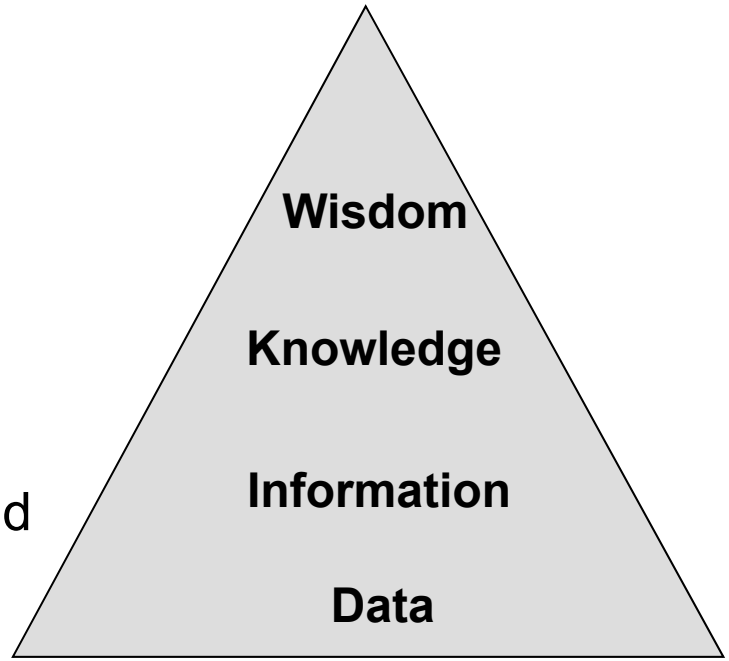

- *Search* and *communication* (of information) are by far the most popular uses of the computer

Motivation (2/2)

- User information need
  - Find all docs containing information on college tennis teams which:
    (1) are maintained by a USA university and
    (2) participate in the NCAA tournament
    (3) National ranking in last three years and contact information

Query

Emphasis is on the retrieval of information (not data)

Search engine/IR system

Information Retrieval

- Information retrieval (IR) is the field concerned with the structure, analysis, or organization, searching and retrieval of information
  - Defined by Gerard Salton, a pioneer and leading figure in IR
- Handle natural language text (or free text) which is not always well structured and could be semantically ambiguous
- Focus is on the user information need
  - Information about a subject or topic
  - Semantics is frequently loose
  - Small errors are tolerated

A user of an IR system is willing to accept documents that contain synonyms of the query terms in the result set, even when those documents do not contain any query terms.

Data Retrieval

- Determine which document of a collection contain the *keywords* in the user query
  - Such documents are regarded as database records, such as a bank account record or a flight reservation, consisting of structural elements such as fields or attributes (e.g., account number and current balance)
- Retrieve all objects (attributes) which satisfy clearly defined conditions in a regular expression or a relational algebra expression
  - Which documents contain a set of keywords (attributes) in some specific fields?
  - Well defined semantics & structures
  - A single erroneous object implies failure!

Data retrieval does not solve the problem of retrieving information about a **subject** or **topic**.

IR Systems: Schematic Depiction

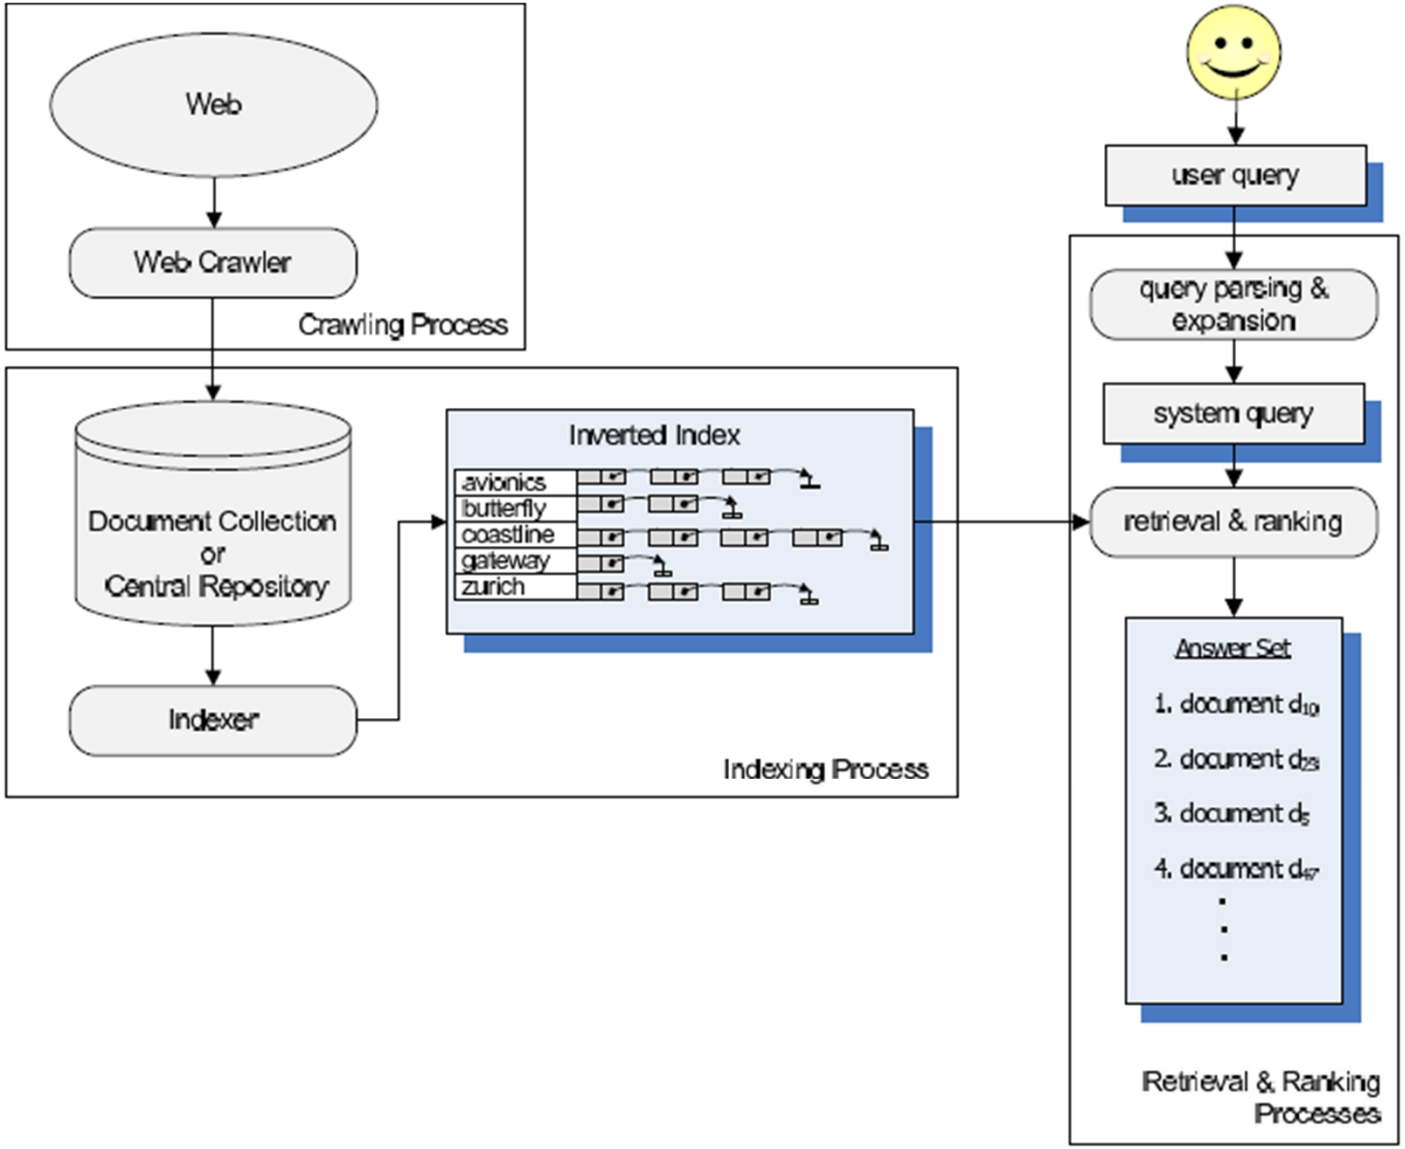

IR systems: Operations

- Indexing: assemble and interpret contents of information items (documents)
  - Most of the information in such documents is in the form of text which relatively unstructured
  - Efficient indexing is of much importance (inverted indexes)
- Retrieval process: generate a ranking that reflects relevance
  - A ranked list of documents returned according to a likelihood of relevance to the user
- Notion of *relevance* is most important
  - Relevance judgment

  (using *clickthrough data*? how to interpret *clickthrough data* as an indicative of relevance.in an unsupervised manner?)
- The other important issues
  - Vocabulary mismatch problems
  - Evaluations of retrieval performance

IR systems: Distinctions

- IR systems can also be distinguished by the scale at which they operate
  - *Web search* (containing billions of documents)
  - *Enterprise, institutional, and domain-specific search*
  - *Personal (desktop) search*
  - *….*

IR at the Center of the Stage

- IR in the last 20 years:
  - Modelng, classification, clustering, filtering
  - User interfaces and visualization
  - Systems and languages

- WWW environment (90~)
  - Universal repository of knowledge and culture
  - Decentralized
  - Without frontiers: free universal access (*freedom to publish*)
  - Hypertext (HTTP protocal and browers by Tim Bermers-Lee)
  - Lack of well-defined data model

Restrictions imposed by mass communication media companies and by natural geographical barriers were almost entirely removed by the invention of the Web! (*e-Publishing Era*)

Web Changed Search!

- Characteristics of document collection
  - Distributed natural => *crawling*
- The size of document collection
  - ~20 billion pages=> *performance* and *scalability* are big issues
- Relevance judgment in the face of the vast size of document collections
  - Hyperlinks and user clicks in documents => *clickthrough data*
- Going beyond seeking text information
  - E.g., price of a book, phone number of a hotel
    => *effective answers* to various types of information needs
- Web advertising and economic incentives
  - E-commerce, advertising <=> *Web spam*

IR Main Issues

- The effective retrieval of relevant information affected by
  - The user task
    - Retrieval/searching and browsing
  - Logical view of the documents
    - Full-text/Keyword-based (text) operations; Indexing

The User Task

- Translate the information need into a query in the language provided by the system
  - A set of words conveying the semantics of the information need

- Browse the retrieved documents

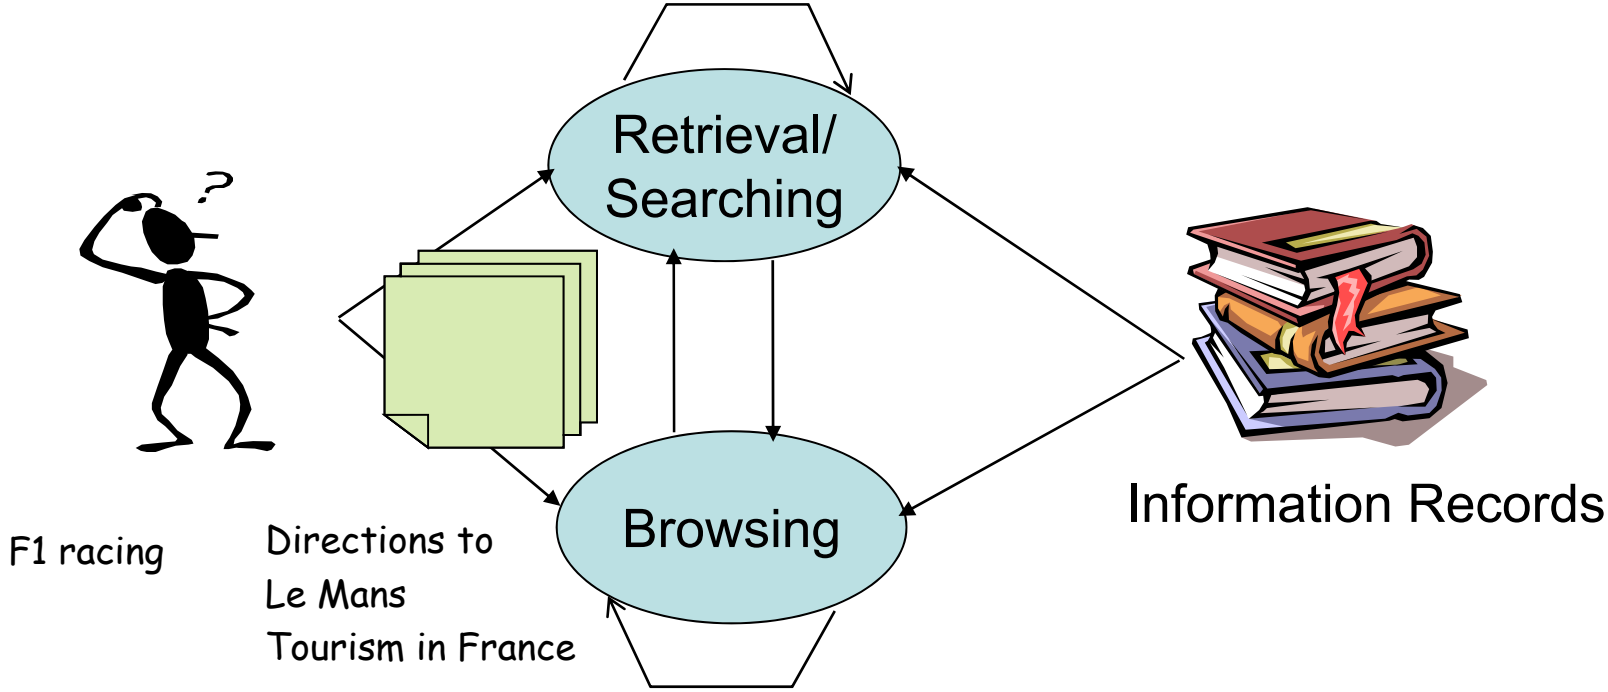

Logical View of the Documents (1/2)

- A full text view (representation)
  - Represent document by its whole set of words
    - Complete but higher computational cost
- A set of index terms by a human subject
  - Derived automatically or generated by a specialist
    - Concise but may poor
- An intermediate representation with feasible *text operations*

Logical View of the Documents (2/2)

- Text operations
  - Elimination of stop-words (e.g. articles, connectives, …)
  - The use of stemming (e.g. tense, …)
  - The identification of noun groups
  - Compression ….
- Text structure (chapters, sections, …)

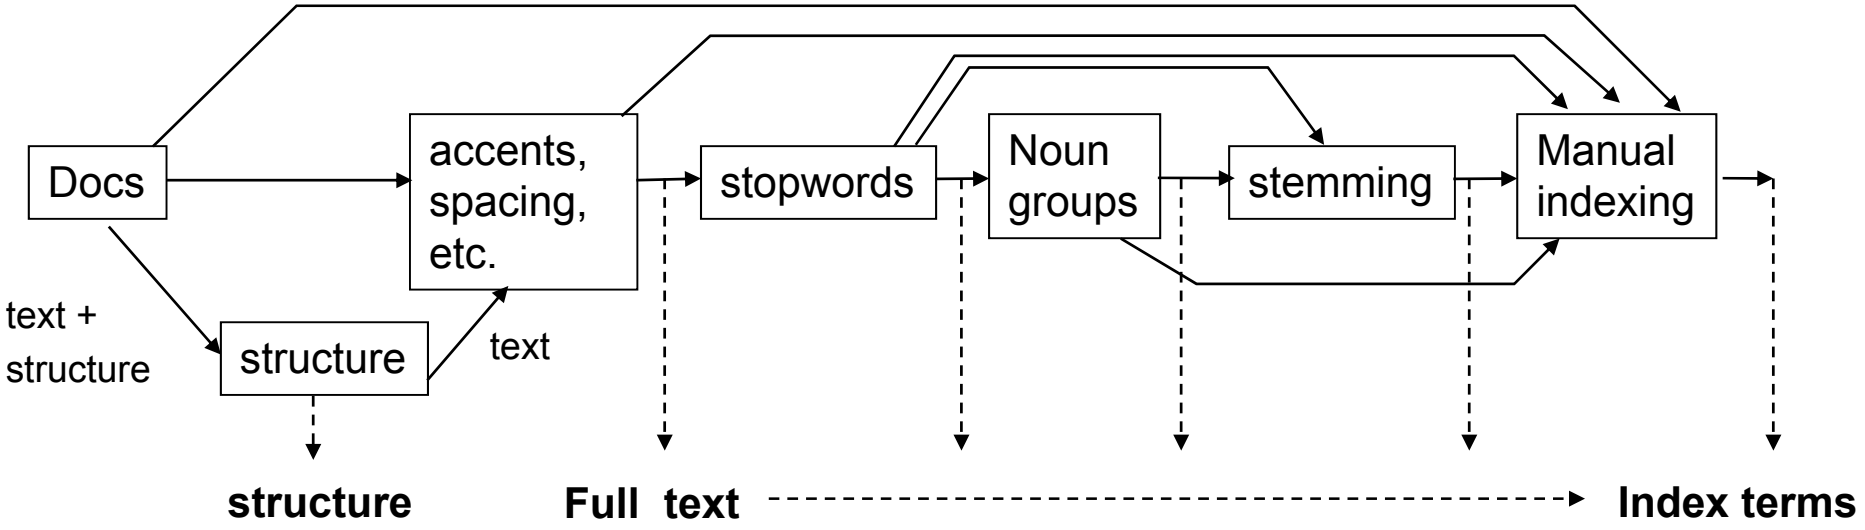

Different Views of the IR Problem

- Computer-centered (commercial perspective)
  - Efficient indexing approaches
  - High-performance matching ranking algorithms

- Human-centered (academic perceptive)
  - Studies of user behaviors
  - Understanding of user needs

Library science
psychology
....

IR for Web and Digital Libraries

- Questions should be addressed
  - Still difficult to retrieve information relevant to user needs
  - Quick response is becoming more and more a pressing factor (Precision vs. Recall)
  - The user interaction with the system (HCI, Human Computer Interaction)
- Other concerns
  - Security and privacy
  - Copyright and patent

The Retrieval Process (1/2)

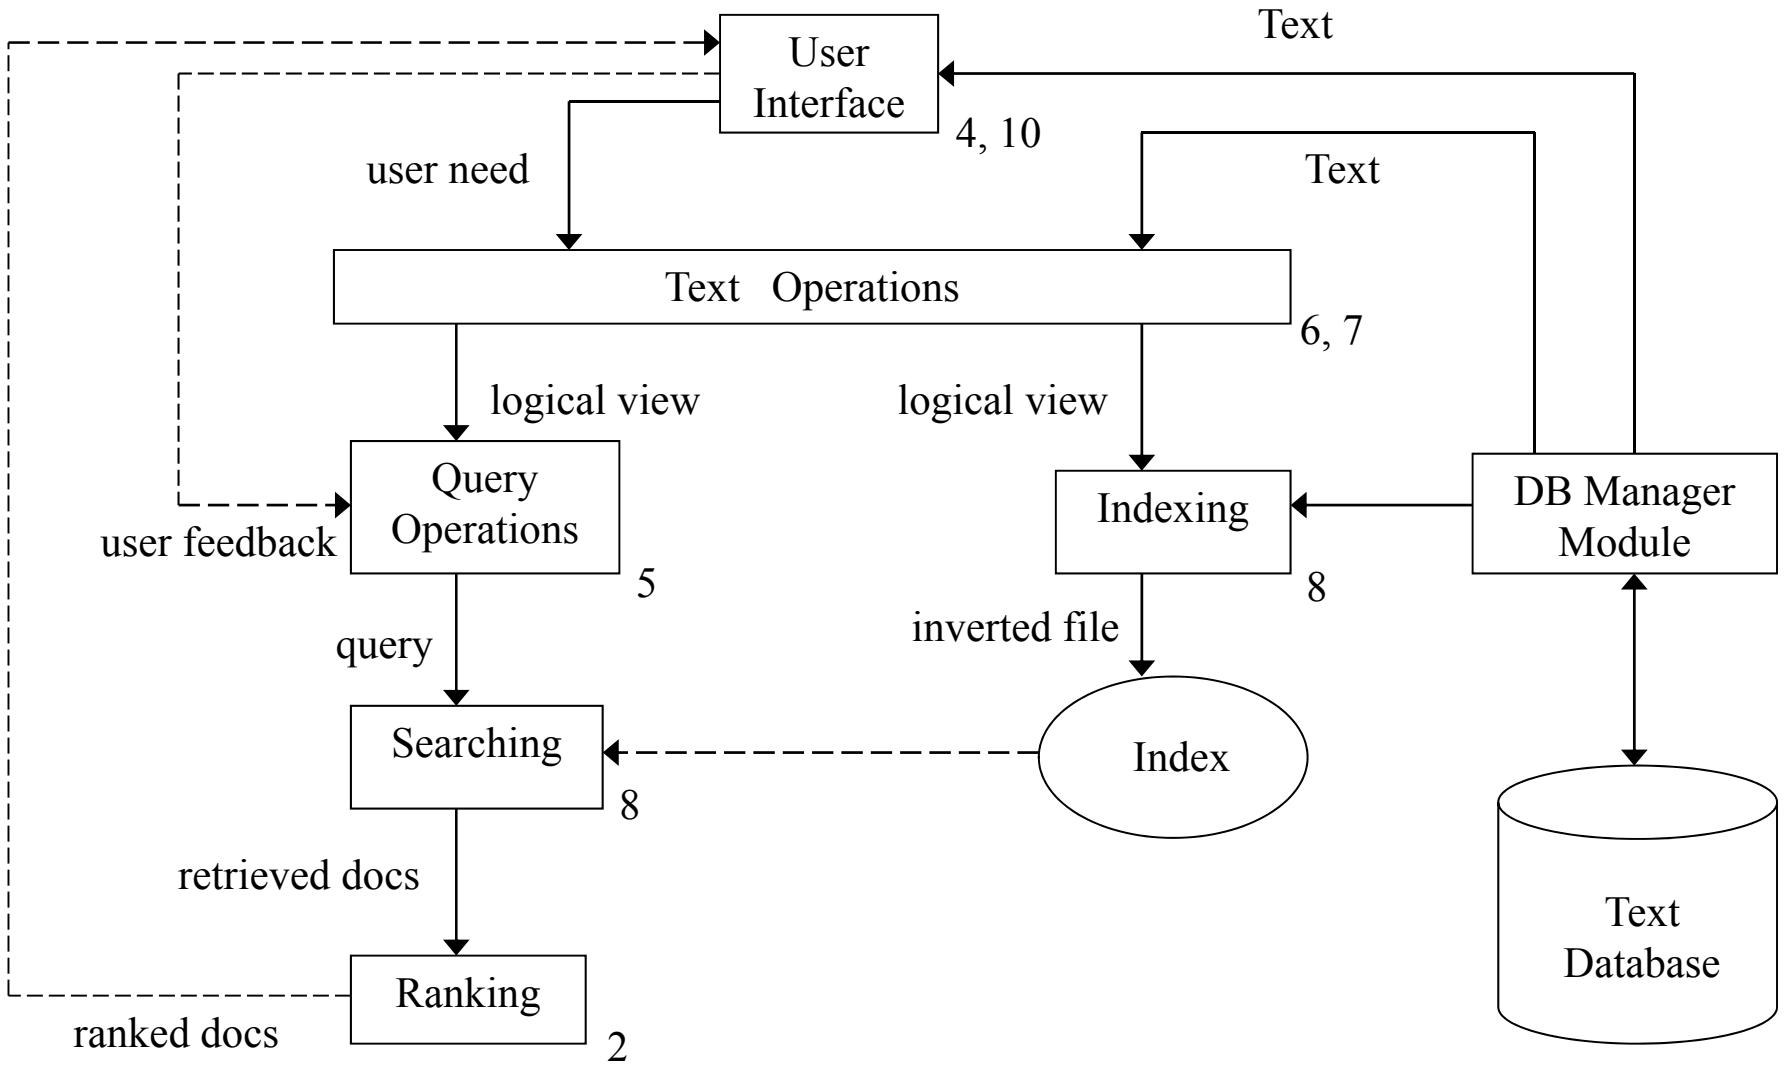

The Retrieval Process (2/2)

- In current retrieval systems
  - Users almost never declare his information need
    - Only a short queries composed few words (typically fewer than 4 words)
  - Users have no knowledge of the text or query operations

Poor formulated queries lead to poor retrieval !

Major Topics (1/2)

- Text IR
  - Retrieval models, evaluation methods, indexing
- Human-Computer Interaction (HCI)
  - Improved user interfaces and better data visualization tools
- Multimedia IR
  - Text, speech, audio and video contents
  - Multidisciplinary approaches
  - Can multimedia be treated in a unified manner?
- Applications
  - Web, bibliographic systems, digital libraries

Major Topics (2/2)

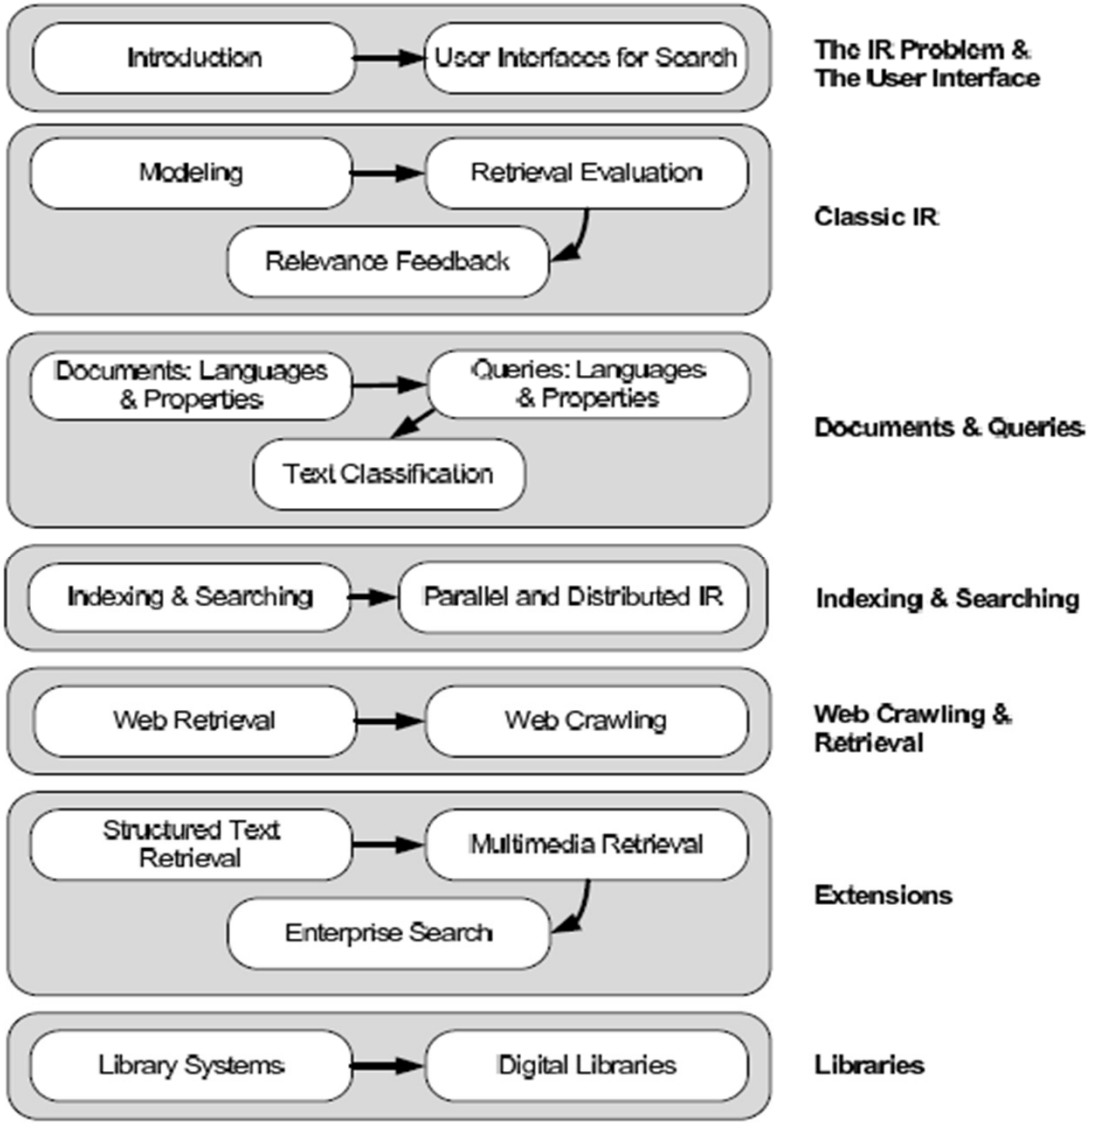

Some Directions of Information Retrieval

Example of Content	Example of Applications	Examples of Tasks
Text	Web search	Ad hoc search
Images	Vertical search	Filtering
Video	Enterprise search	Classification
Scanned documents	(Personal) Desktop search	Question answering
Audio (Speech)	Peer-to-peer search	
Music		

- In the past, most technology for searching non-text document relies on the descriptions of their content rather than the contents themselves
  - The need of "*content-based*" image/audio/music retrieval !
- Peer-to-peer search involves finding information in networks of nodes or computers without any centralized control

IR and Search Engines

Information Retrieval

Relevance

-Effective ranking

Evaluation

-Testing and measuring

Information needs

-User interaction

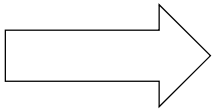

Search Engines

Performance

-Efficient search and indexing

Incorporating new data

-Coverage and freshness

Scalability

-Growing with data and users

Adaptability

-Tuning for applications

Specific problems

-e.g. Spam

Text Information Retrieval (1/4)

- Internet searching engine

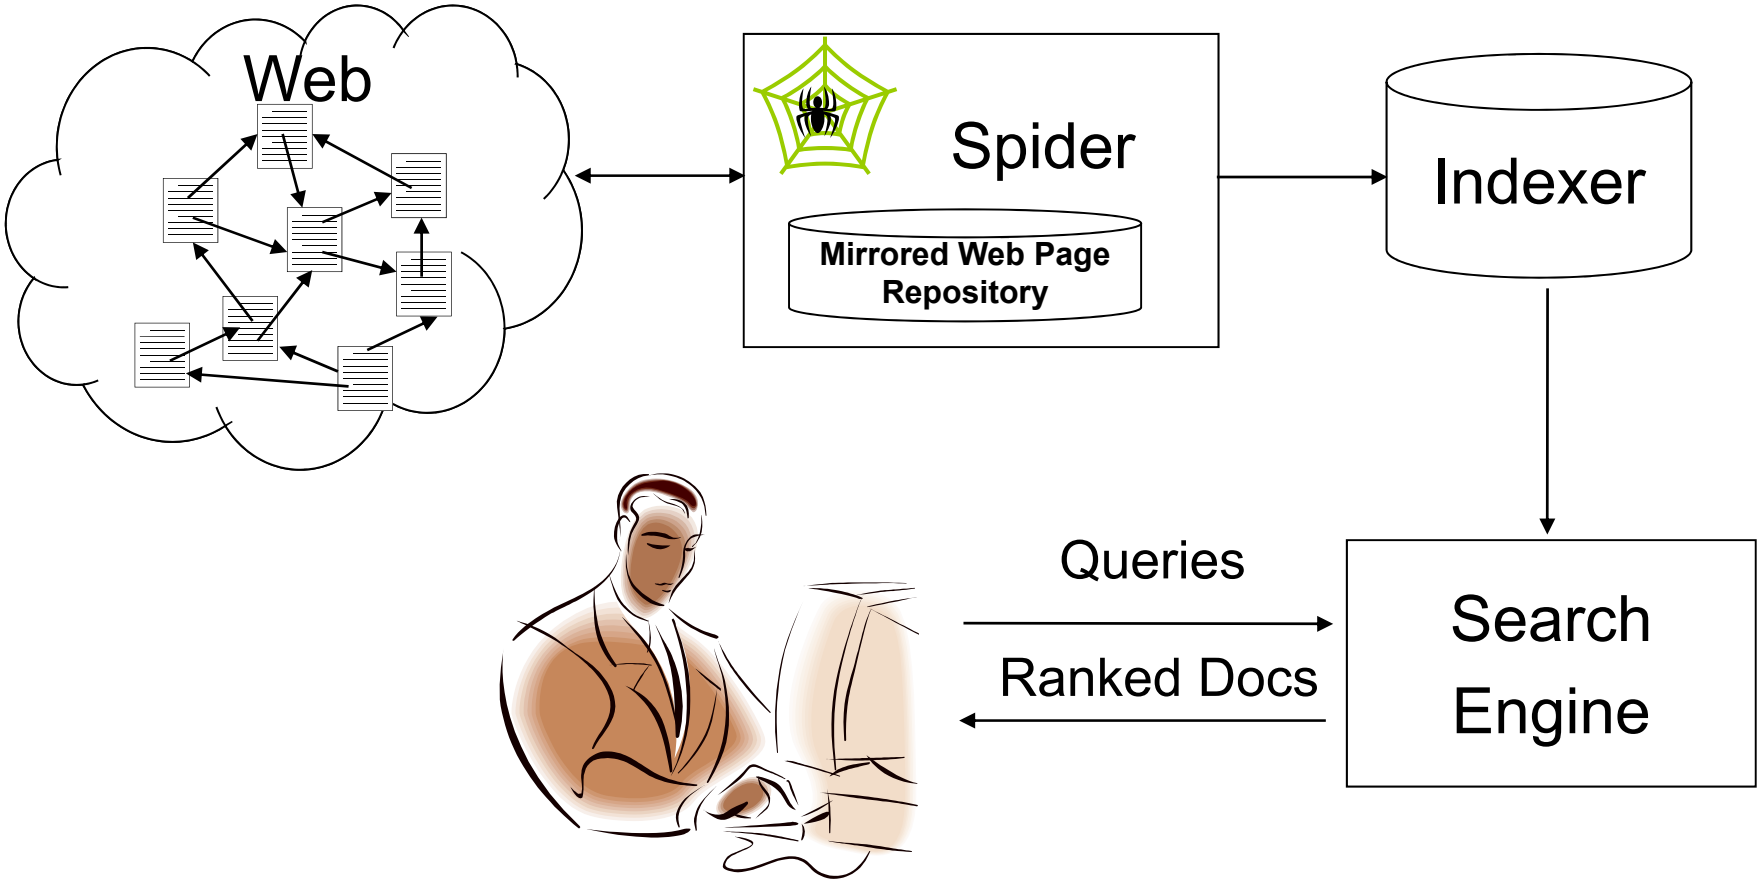

Text Information Retrieval (2/4)

- http://www.google.com

Text Information Retrieval (3/4)

- http://www.openfind.com.tw (Service is No Longer Available)

Text Information Retrieval (4/4)

- http://www.baidu.com

Speech Information Retrieval (1/4)

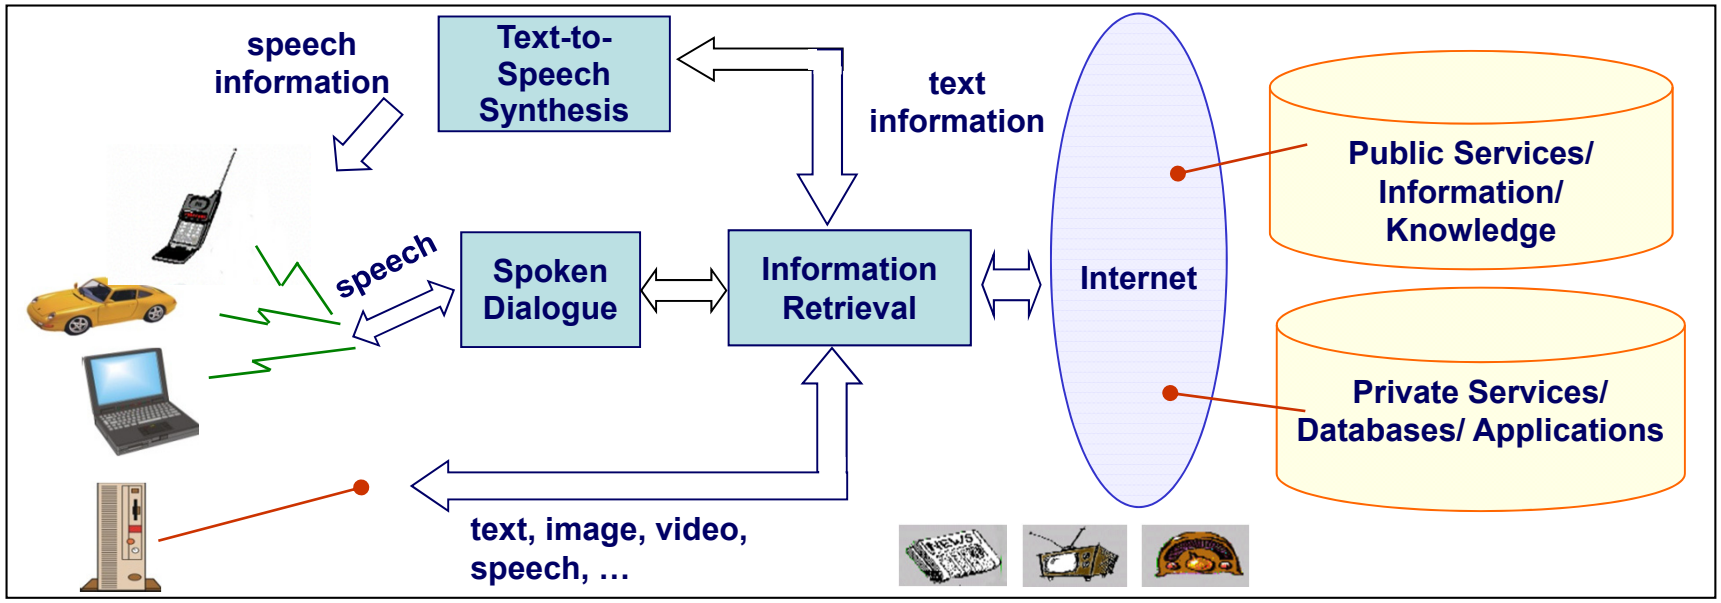

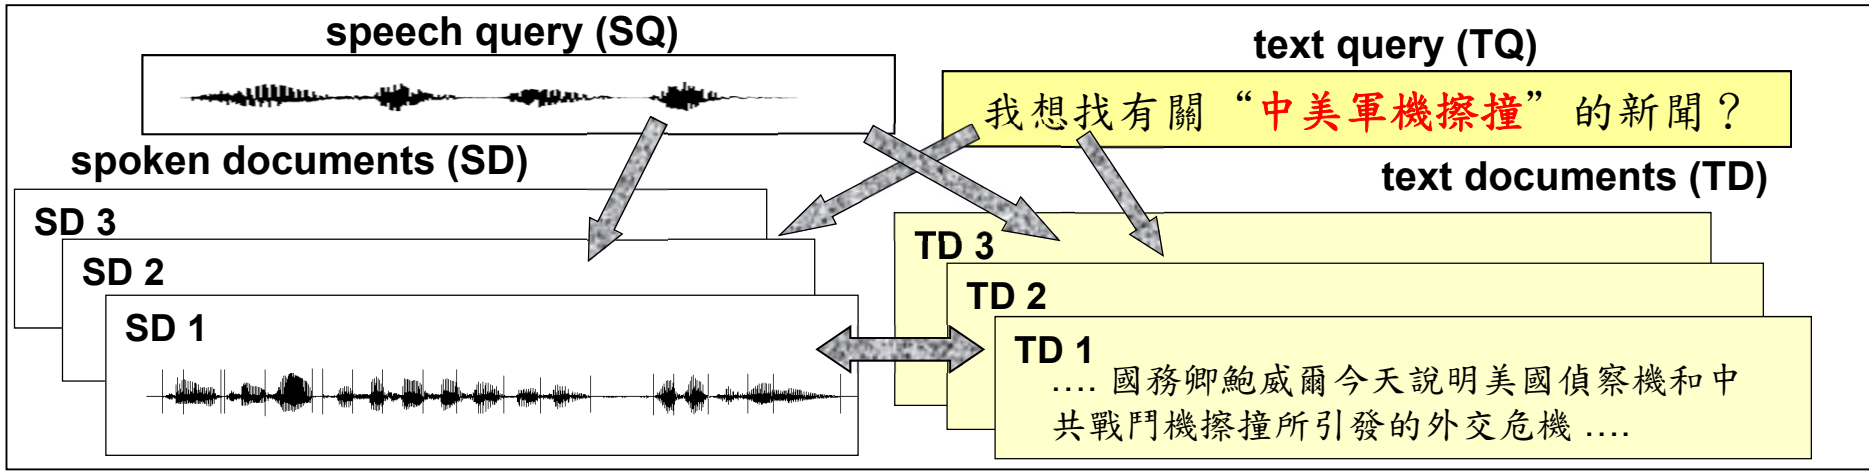

Speech Information Retrieval (2/4)

- HP Research Group – Speechbot System
  (Service is No Longer Available)
  - Broadcast news speech recognition, Information retrieval, and topic segmentation (SIGIR2001)
  - Currently indexes **14,791 hours of content** (2004/09/22, http://speechbot.research.compaq.com/)

Speech Information Retrieval (3/4)

- Speech Summarization and Retrieval

輸入聲音問句：“請幫我查總統府升旗典禮”

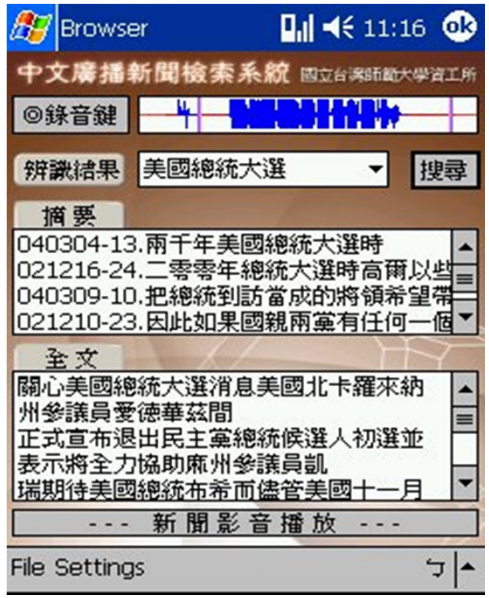

Speech Information Retrieval (4/4)

- Speech Organization

- L.-S. Lee and B. Chen, "Spoken Document Understanding and Organization," *IEEE Signal Processing Magazine* 22(5), pp. 42-60, Sept. 2005

Visual Information Retrieval (1/4)

- Content-based approach

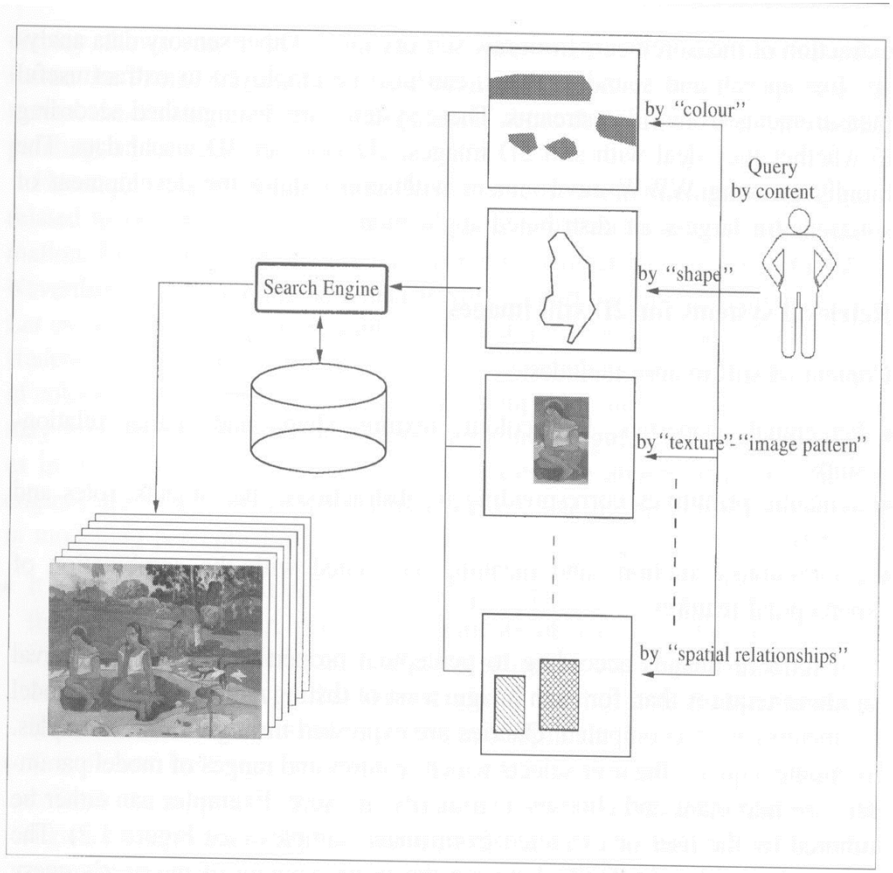

Figure 1.2 Different types of query by example.

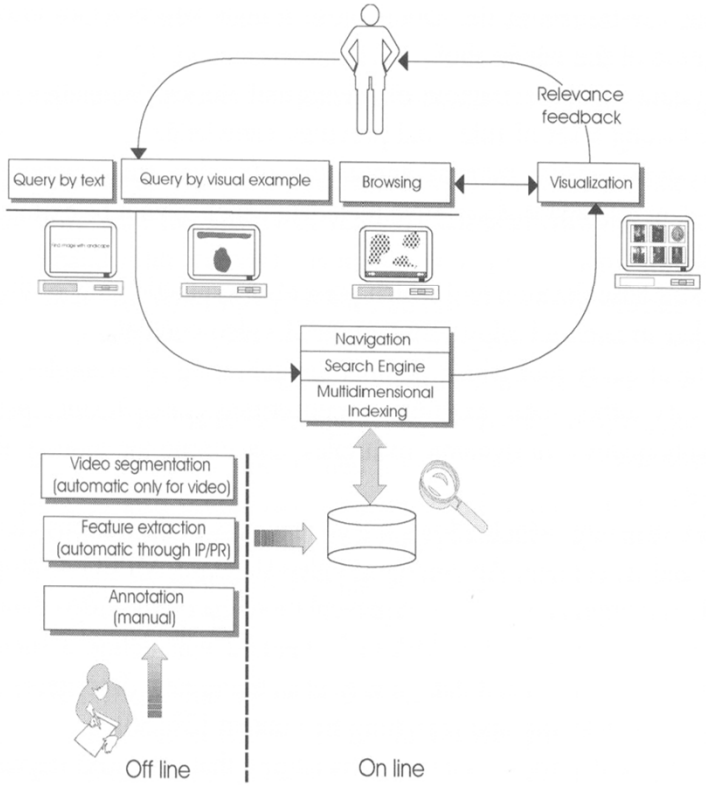

Figure 1.5 Sketch of a new-generation visual information retrieval system for video.

Visual Information Retrieval (2/4)

- Images with Texts (Metadata)

Visual Information Retrieval (3/4)

- Content-based Image Retrieval

Visual Information Retrieval (4/4)

Video Analysis and Content Extraction

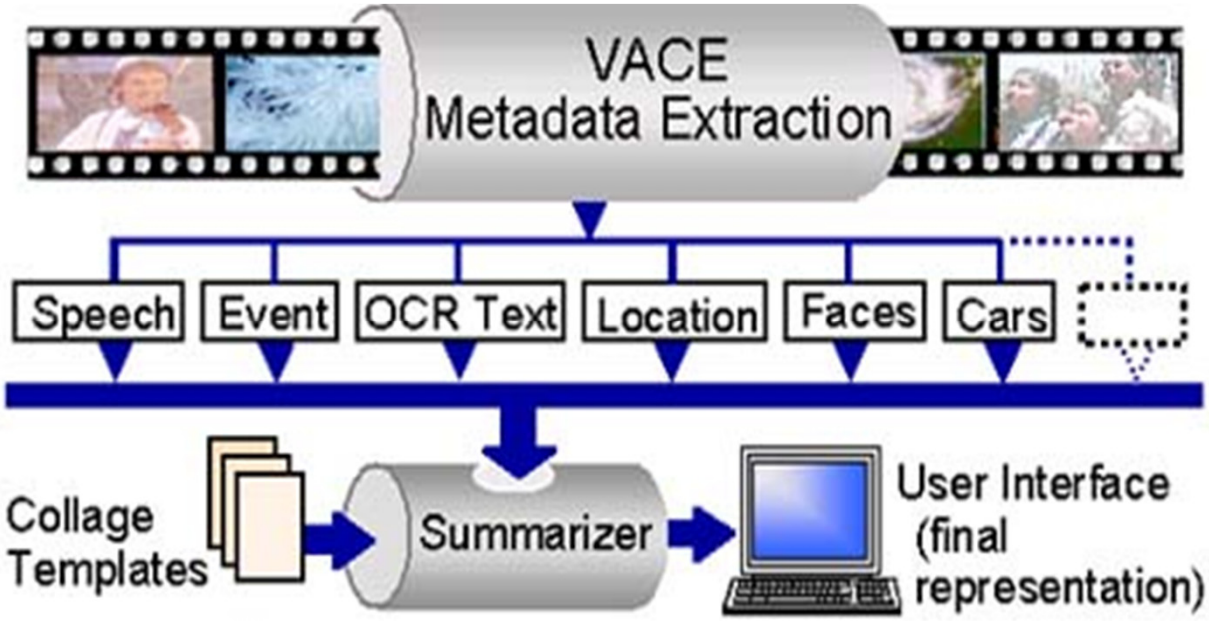

Scenario for Multimedia information access

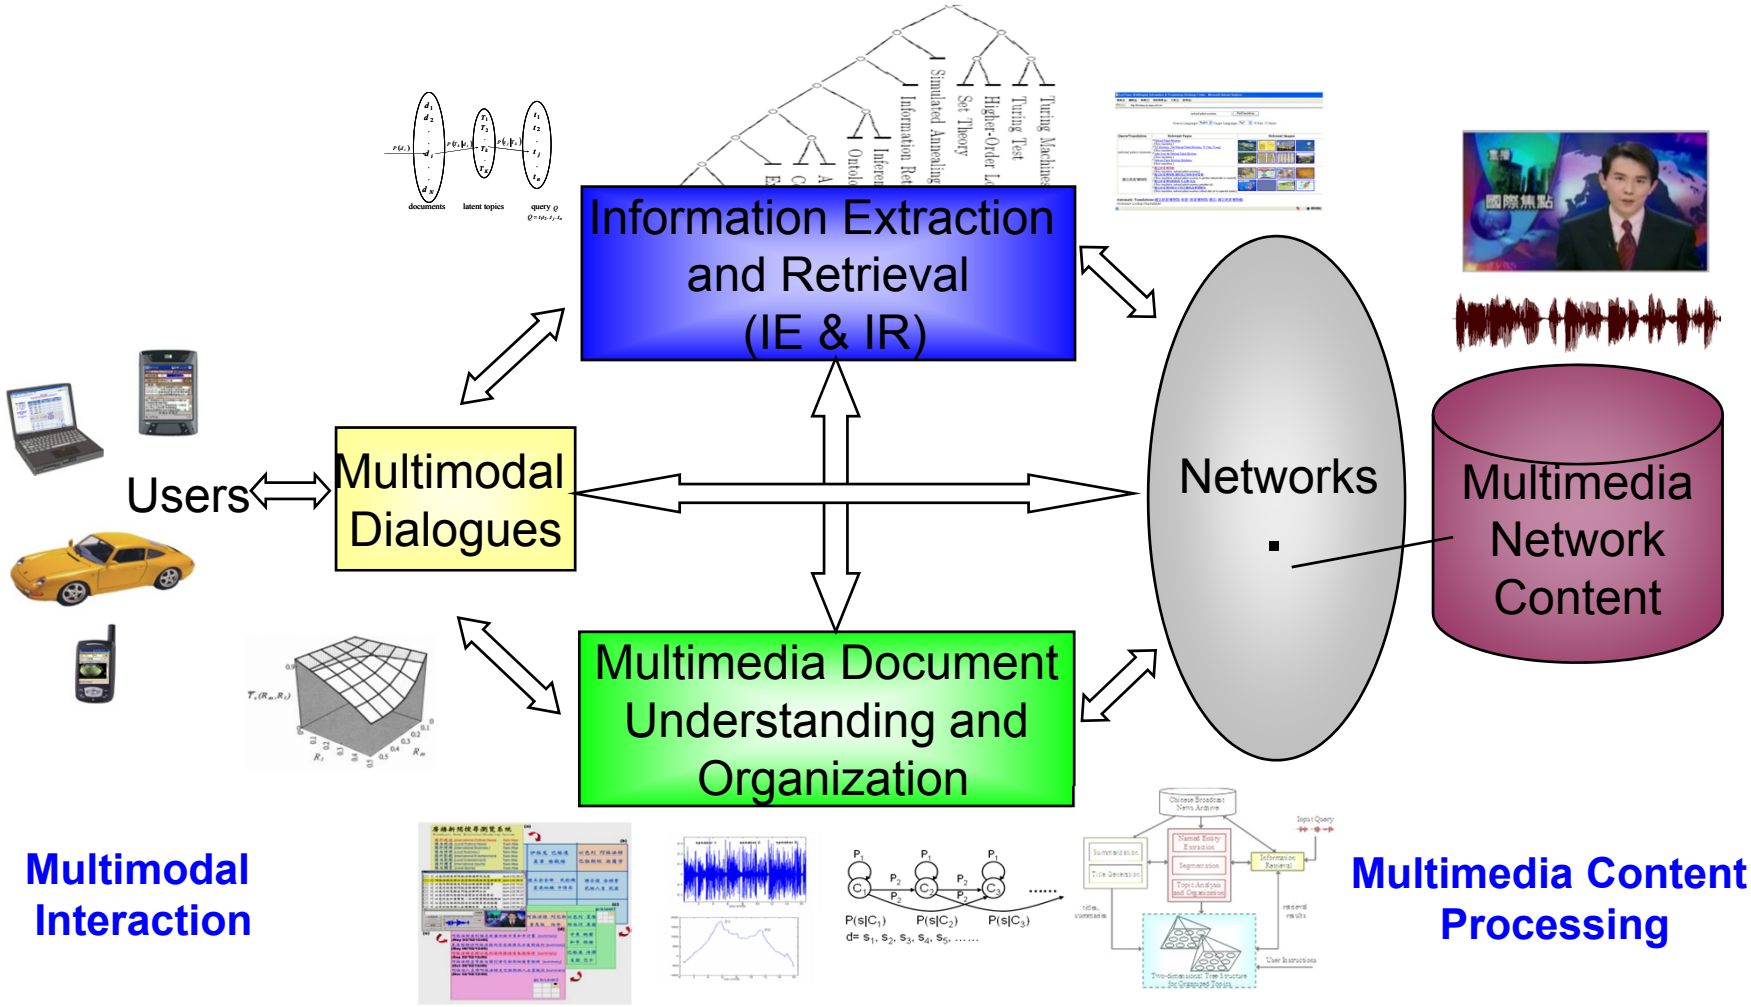

Other IR-Related Tasks

- Information filtering and routing
- **Term/Document categorization**
- **Term/Document clustering**
- **Document summarization**
- **Information extraction**
- Question answering
  - "*What is the height of Mt. Everest?*"
- Crosslingual information retrieval
- …..

Document Summarization

- Audience
  - Generic summarization
  - User-focused summarization
    - Query-focused summarization
    - Topic-focused summarization
- Function
  - Indicative summarization
  - Informative summarization
- Extracts vs. abstracts
  - Extract: consists wholly of portions from the source
  - Abstract: contains material which is not present in the source
- Output modality
  - Speech-to-text summarization
  - Speech-to-speech summarization
- Single vs. multiple documents

Information Extraction

- E.g., Named-Entity Extraction
  - NE has it origin from the Message Understanding Conferences (MUC) sponsored by U.S. DARPA program
    - Began in the 1990's
    - Aimed at extraction of information from text documents
    - Extended to many other languages and spoken documents (mainly broadcast news)
  - Common approaches to NE
    - Rule-based approach
    - Model-based approach
    - Combined approach

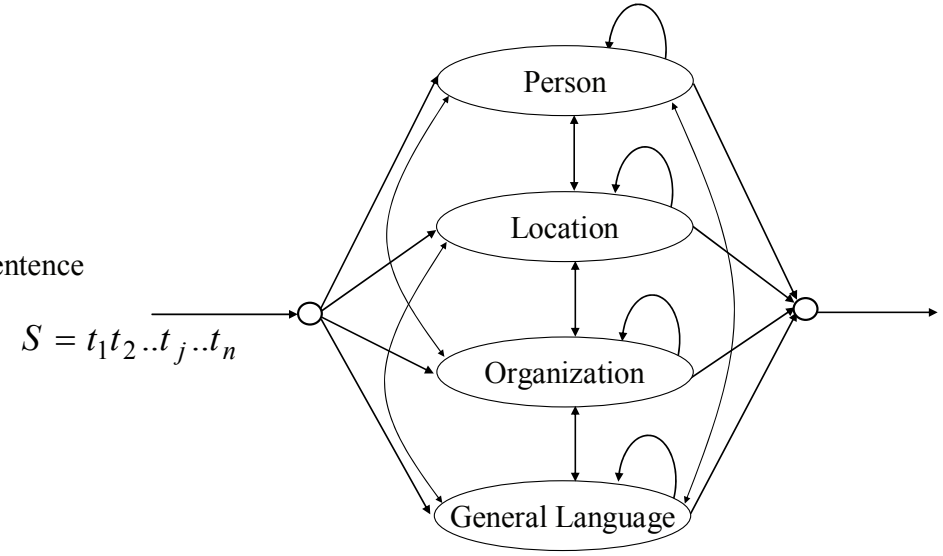

Cross-lingual Information Retrieval

- E.g., Automatic Term Translation
  - Discovering translations of unknown query terms in different languages
  - E.g., The Live Query Term Translation System (LiveTrans) developed at Academia Sinica/by Dr. Chien Lee-Feng

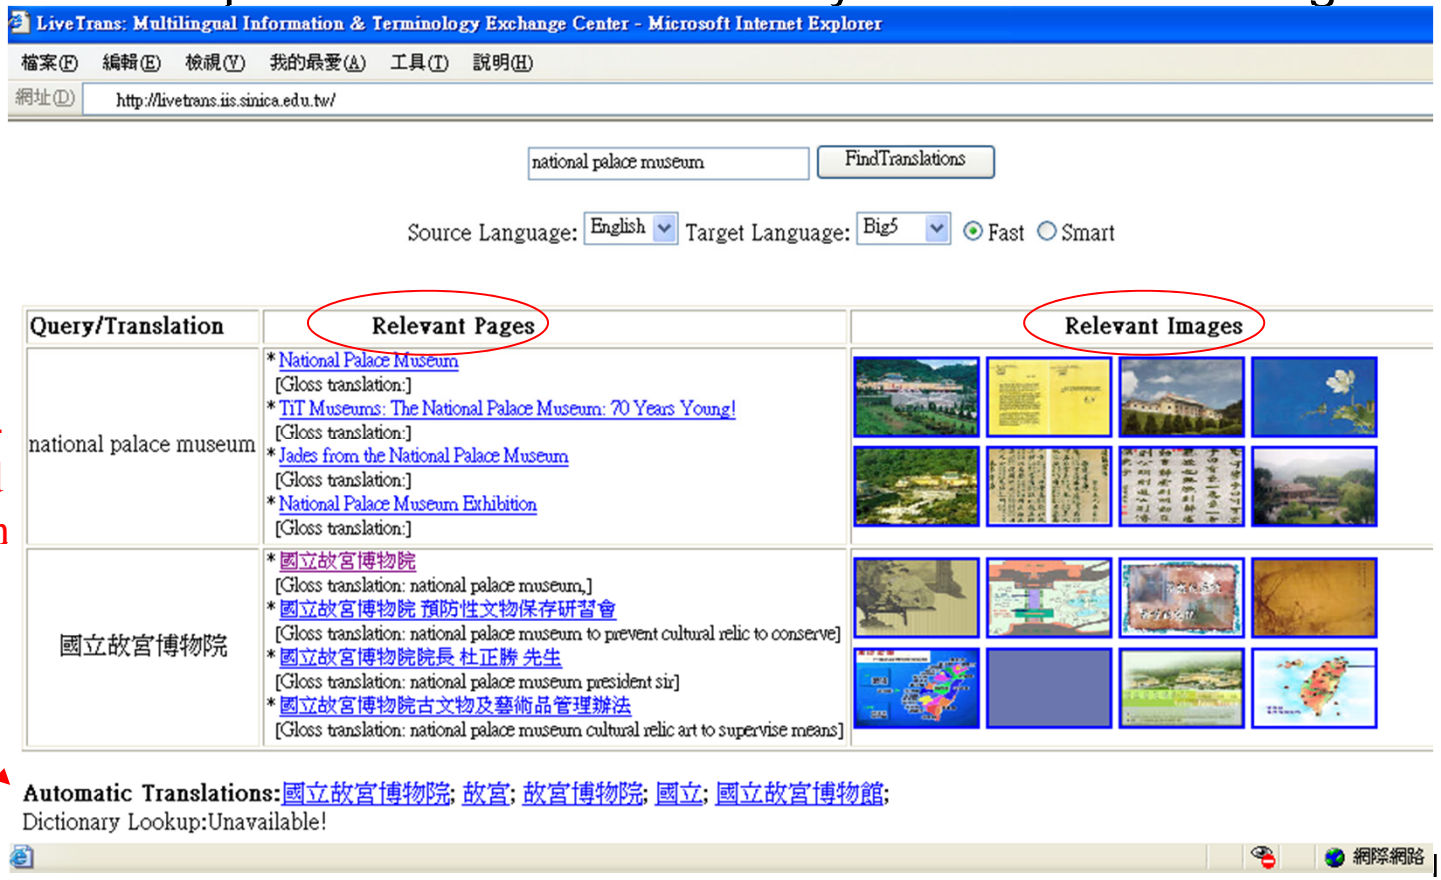

Machine-Extracted Translation

Multidisciplinary Approaches

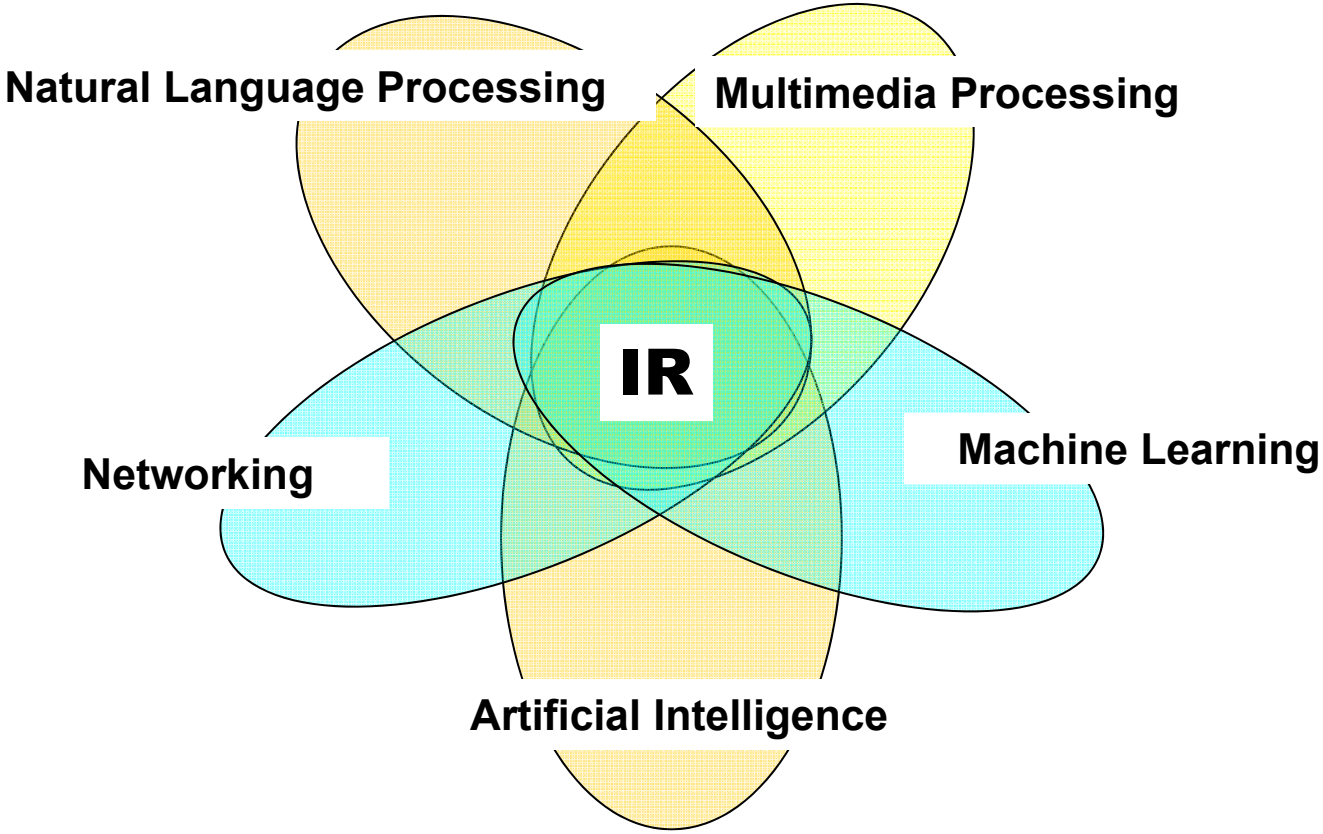

Resources

- Corpora (Speech/Language resources)
  - Refer speech waveforms, machine-readable text, dictionaries, thesauri as well as tools for processing them
    - LDC - Linguistic Data Consortium

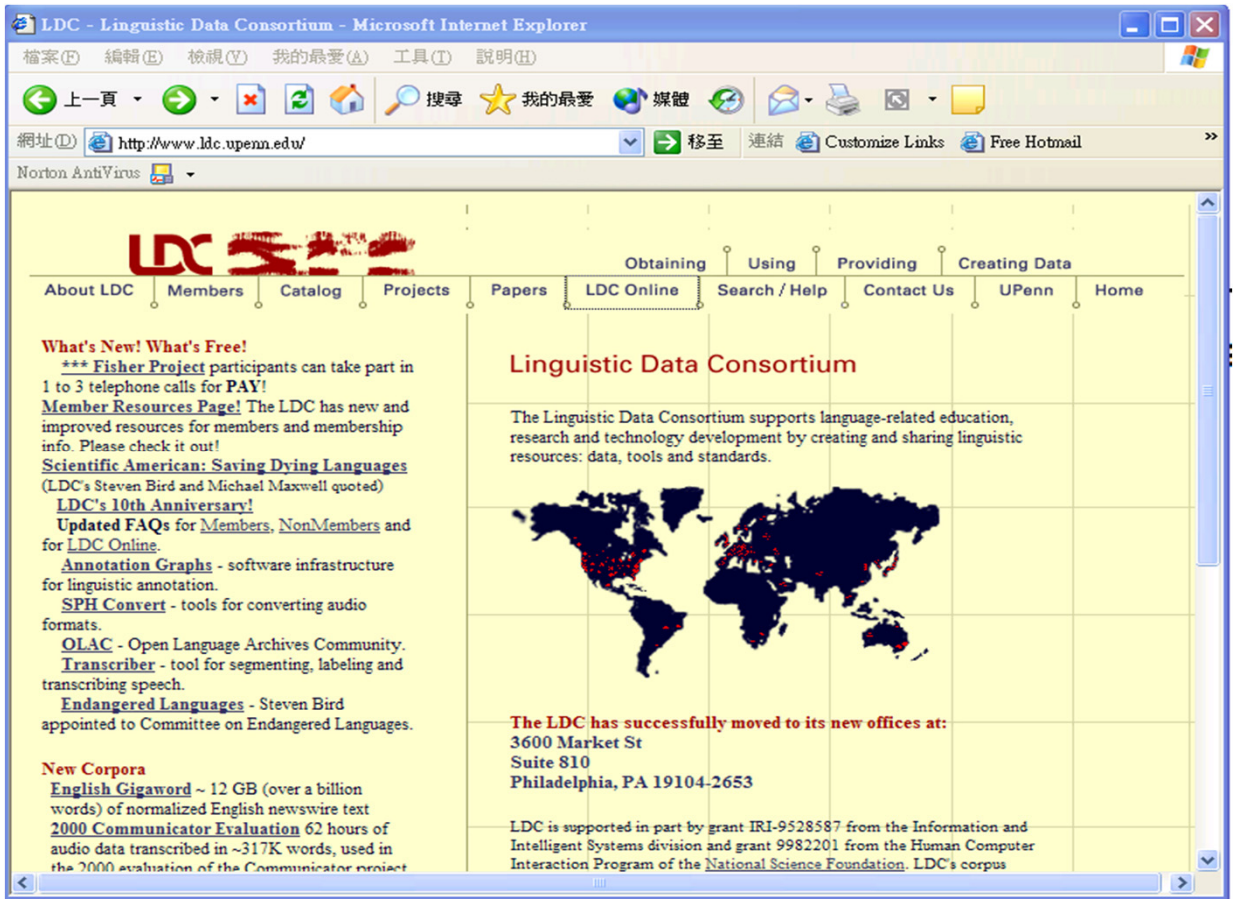

Contests (1/2)

- Text REtrieval Conference (TREC)

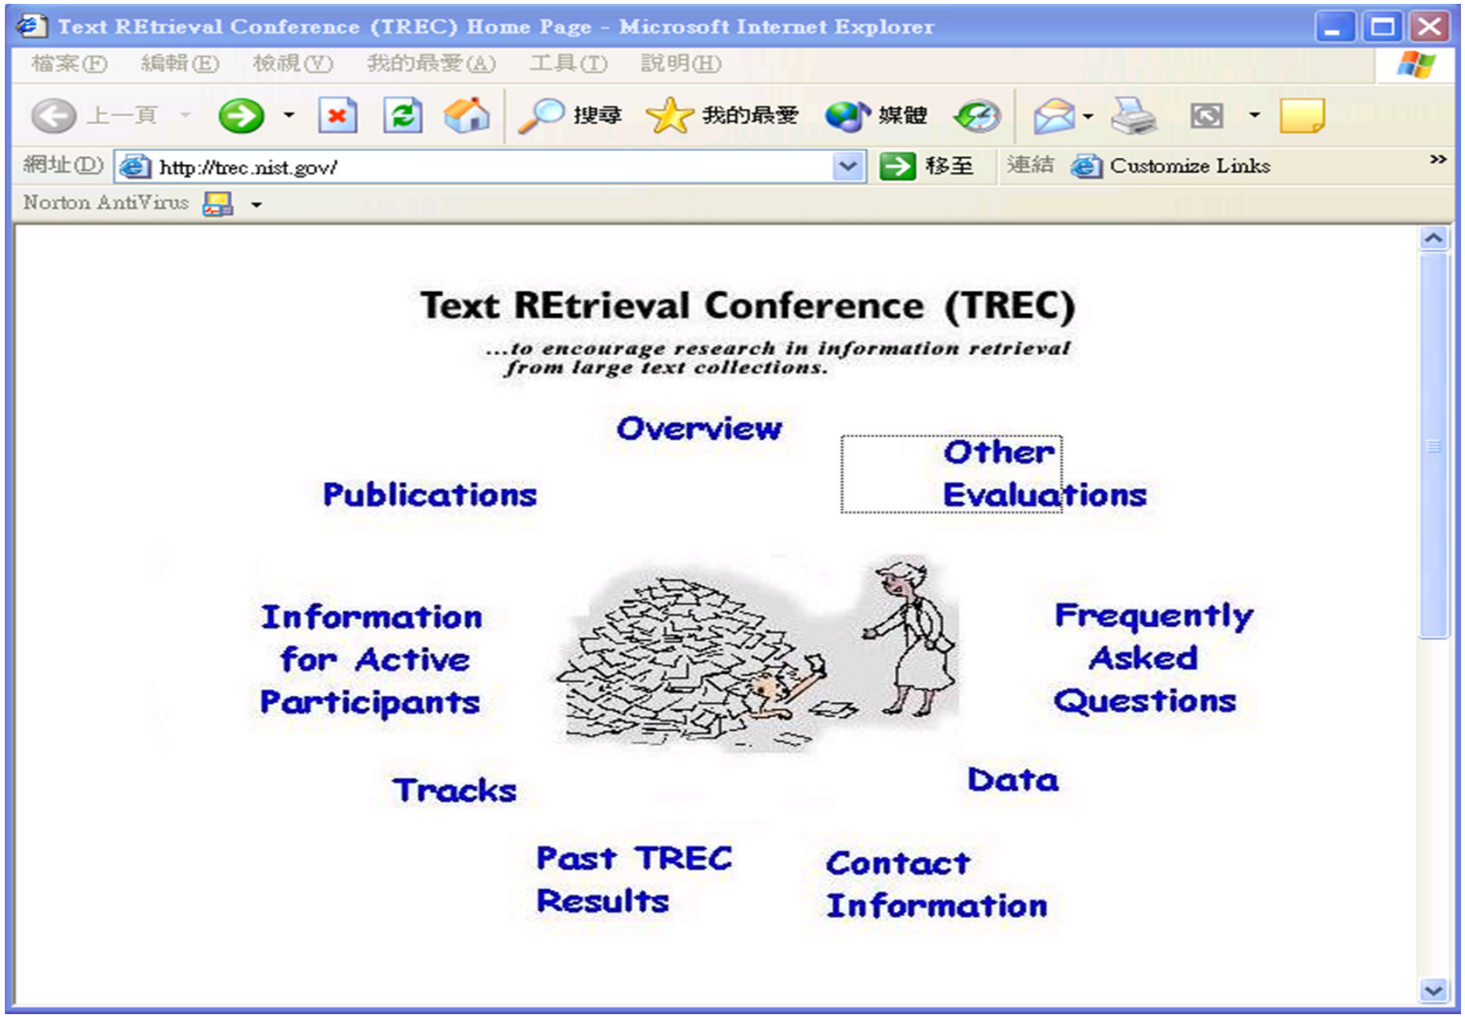

Contests (2/2)

- US National Institute of Standards and Technology

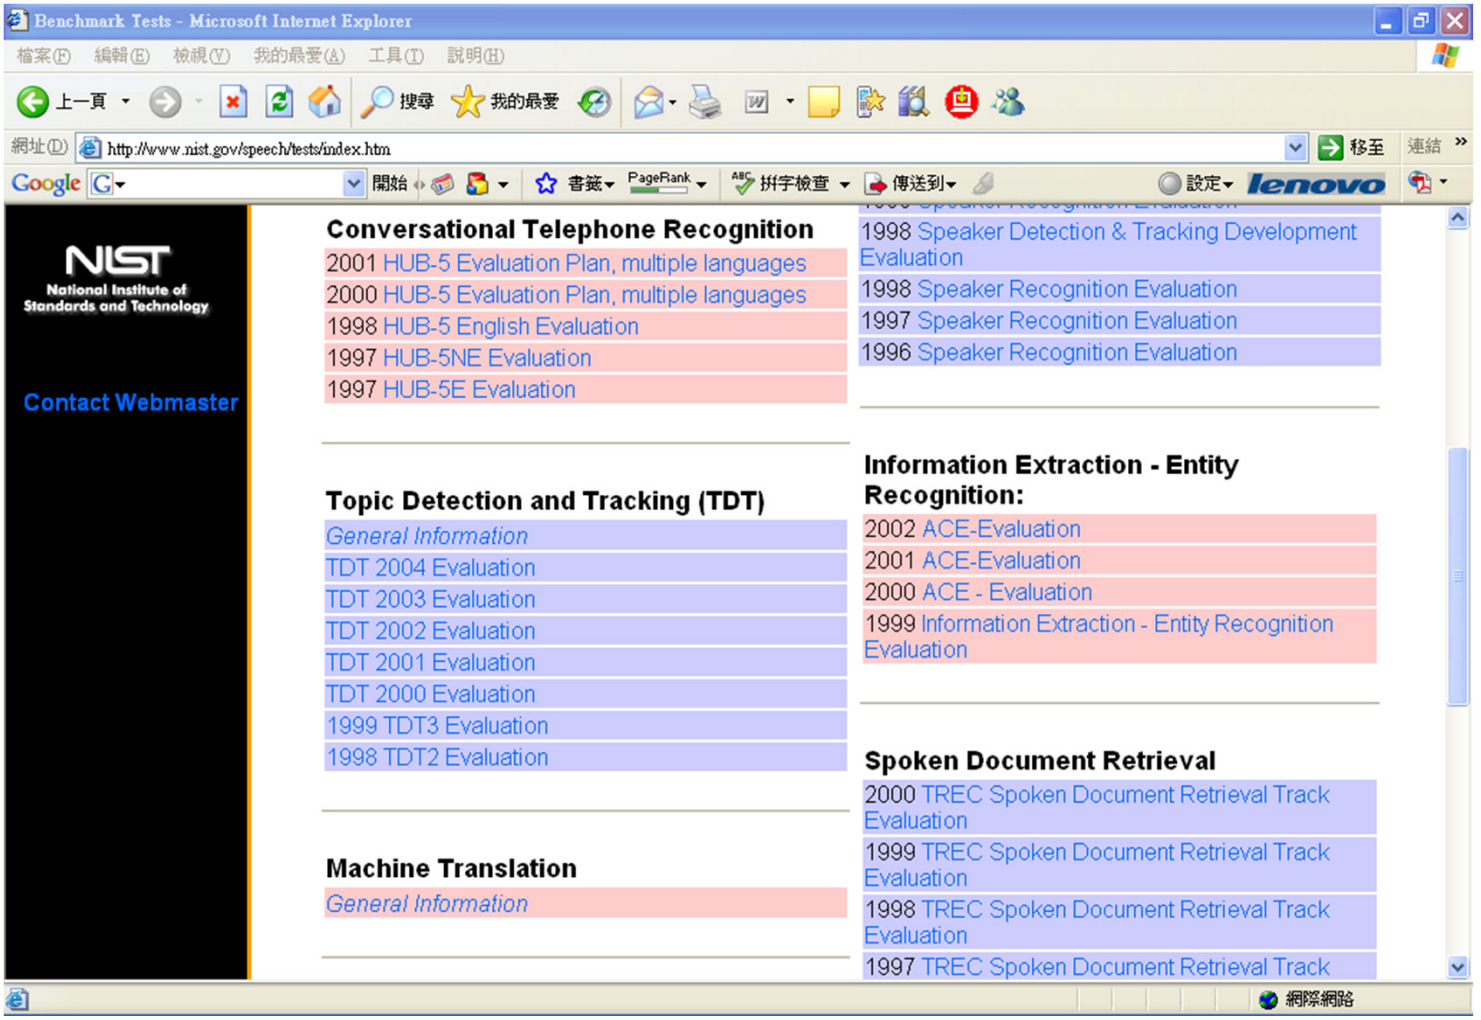

Conferences/Journals

- Conferences
  - ACM Annual International Conference on Research and Development in Information Retrieval (SIGIR)
  - ACM Conference on Information Knowledge Management (CIKM)
  - …
- Journals
  - Journal of the American Society for Information Science (JASIS)
  - ACM Transactions on Information Systems (TOIS)
  - Information Processing and Management (IP&M)
  - ACM Transactions on Asian Language Information Processing (TALIP)
  - …

Tentative Topic List

Course Overview & Introduction
Retrieval Models (I) - Classic Retrieval Models (Boolean, Vector Space and Probabilistic Models)
Retrieval Performance Evaluation - Measures
Retrieval Performance Evaluation - Collections
Retrieval Models (II) - Improved Approaches (Fuzzy Set, Extended Boolean, Generalized Vector Space Models)
Query Operations (Query Expansion and Term Re-weighting)
Retrieval Models (III) - Latent Semantic Analysis (LSA)
Retrieval Models (IV) - Language Models
Retrieval Models (V) - Learning to Rank
Clustering for Information Retrieval
Classification for Information Retrieval
Efficient Indexing and Searching
Web Search Basics
Cross-lingual Information Retrieval
Spoken Document Recognition, Retrieval and Summarization

Grading (Tentative)

- Midterm (or Final): 45%
- Homework/Projects: 30%
- Presentation: 15%
- Attendance/Other: 10%