COSC572

GUEST LECTURE - PROF. GRACE HUI YANG
INTRODUCTION TO INFORMATION RETRIEVAL
NOV 2, 2016

TOPICS FOR TODAY

- **Modes of Search**
- **What is Information Retrieval**
- **Search vs. Evaluation**
- **Vector Space Model (VSM)**
- **Dynamic Search**

WHAT IS SEARCH?

LET'S PRACTICE SOMETHING YOU DO EVERY DAY...

EXERCISE

- **Find what city and state Dulles airport is in, what shuttles ride-sharing vans and taxi cabs connect the airport to other cities, what hotels are close to the airport, what are some cheap off-airport parking, and what are the metro stops close to the Dulles airport.**

HOW DO YOU SEARCH?

"(Search)/Information-seeking is a special case of problem solving. It includes recognizing and interpreting the information problem, establishing a plan of search, conducting the search, evaluating the results, and if necessary, iterating through the process again."

–GARY MARCHIONINI, 1989

SEARCH IS ALSO KNOWN AS INFORMATION SEEKING

TWO MODES OF SEARCH

1. SEARCH BY QUERY

SEARCH BY QUERY

- The user issues a textual query
  - types the query in a search box

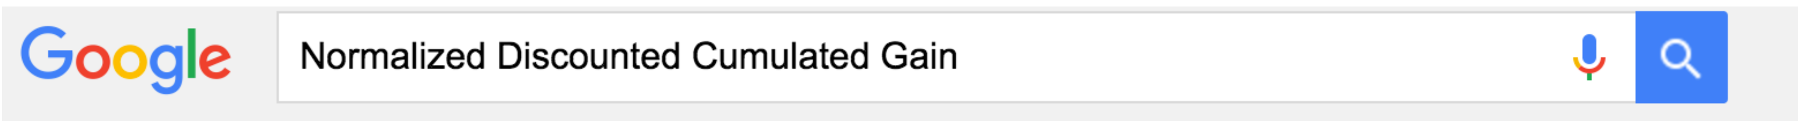

- After examining the search results, the user
  - either stops the search
  - or modifies the query to start another ad-hoc search

- **… it is like testing the water …**

E.g. Find what city and state Dulles airport is in, what shuttles ride-sharing vans and taxi cabs connect the airport to other cities, what hotels are close to the airport, what are some cheap off-airport parking, and what are the metro stops close to the Dulles airport.

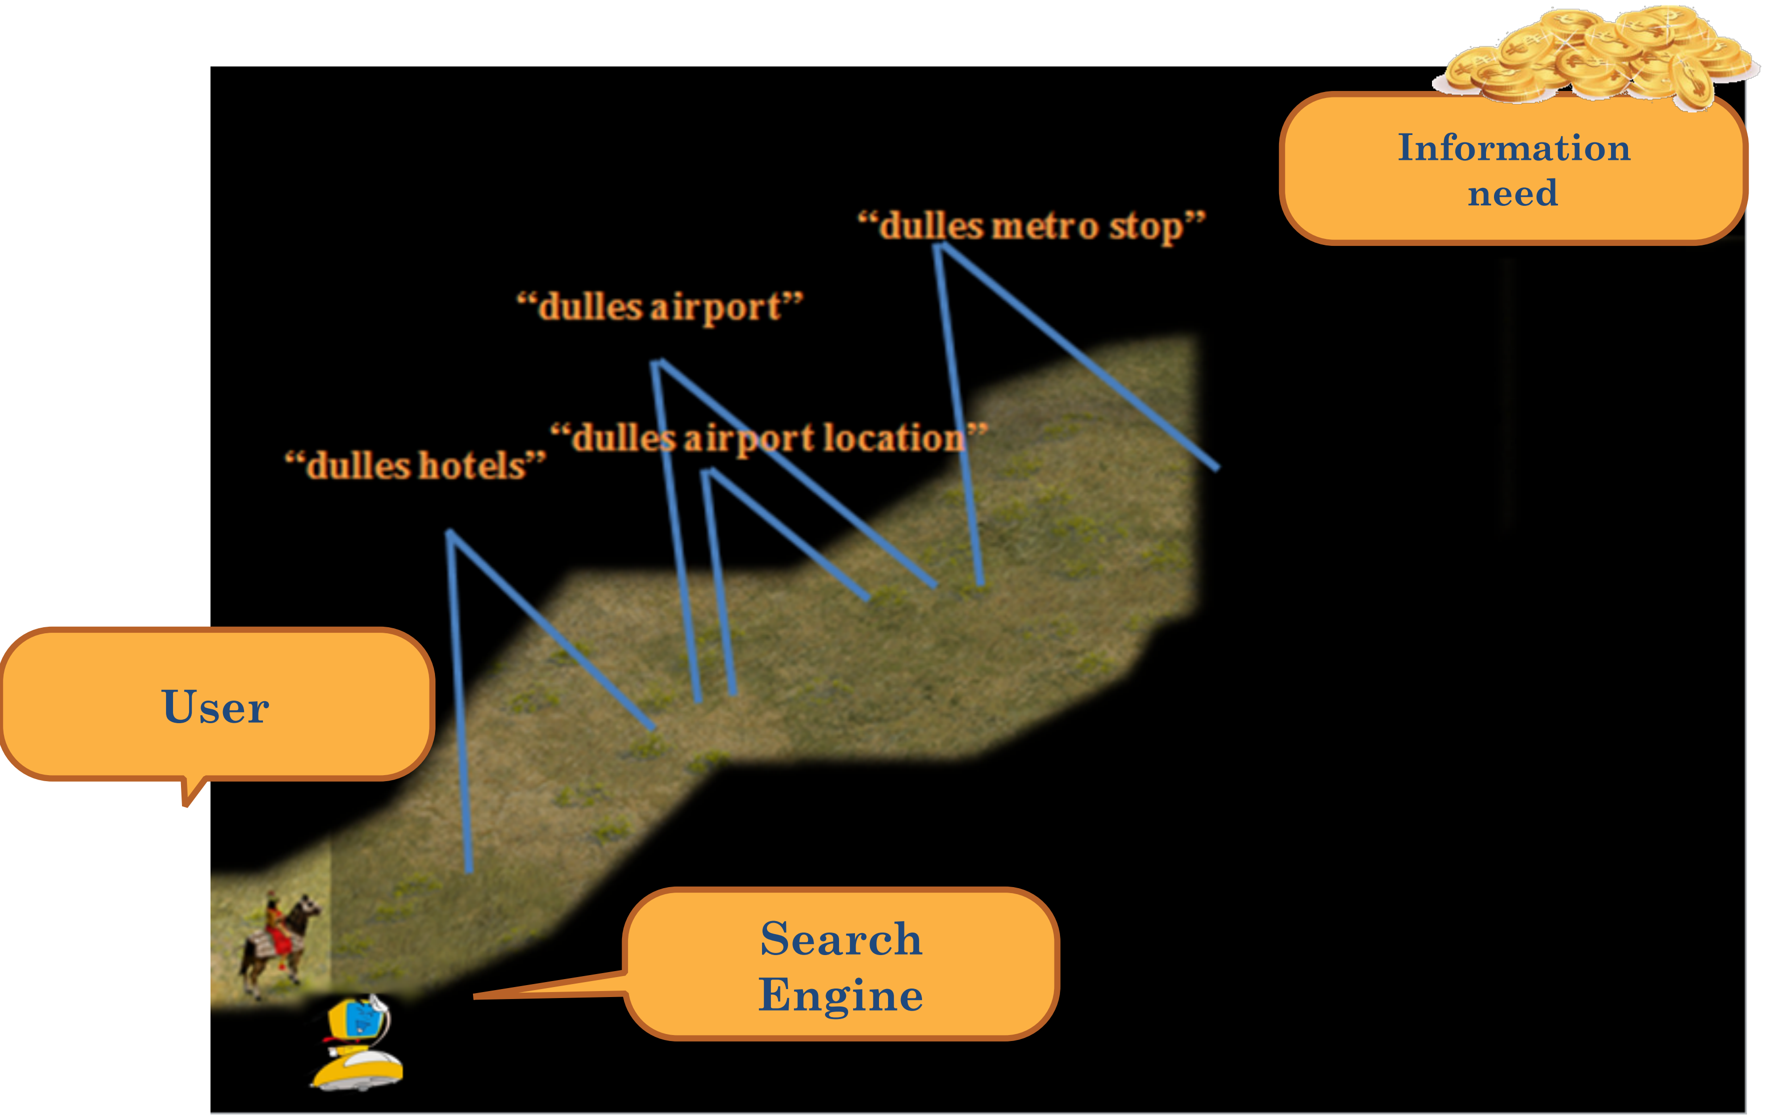


ENSEMBLE
STUDIOS
AGE of EMPIRES II

Age of Empire

2. SEARCH BY BROWSING

SEARCH BY BROWSING

- **The system can provide the searcher with structure that overviews/characterizes the available information**
- **The searcher (user) browses or navigate the organized information, by selecting links or categories that produce pre-defined groups of information items.**
- **Browsing involves following a chain of links, switching from one view to another, in a sequence of scan and select operations.**

YAHOO!
SMALL BUSINESS

Search Web Sign In Mail

Business to Business
Communications and Networking, Manufacturing, Computers...

U.S. States
California, Michigan, Virginia...

Shopping and Services
Apparel, Travel and Transportation, Communication and Information Management...

International
Fiji, Russia, Canada...

Employment and Work
Careers and Jobs,

Business and Economy
Directories, Organizations, Classifieds...

Entertainment
Television Shows, Consumer Electronics, Comedy...

Finance and Investment
Chats and Forums, Socially Responsible Investing, Reference and Guides...

Science
Engineering, Agriculture, Energy...

Arts
Design Arts, Visual Arts, Humanities...

Recreation
Outdoors, Sports, Games...

Society and Culture
Home and Garden, Weddings, Religion and Spirituality...

Social Science
Linguistics and Human Languages, Political Science, Psychology...

Government
U.S. Government, Law,

Education
Distance Learning, K-12, Career and Vocational...

News and Media
Business, Television, Newspapers...

Reference
Phone Numbers and Addresses, Calendars, Quotations...

Yahoo Directory Closes, Five Days Early

The core service that launched Yahoo as an internet company -- the Yahoo Directory -- finally closes 20 years after it began.

Danny Sullivan on December 27, 2014 at 10:33 pm

The Yahoo Directory, the core part of how Yahoo itself began in 1994, officially closed today, five days ahead of when Yahoo had said the end would come.

- **… but you are still using browsing …**

GEORGETOWN UNIVERSITY

GIVING CONNECT DIRECTORY

Academic Year

2015-2016 Fall and Spring

» *Course web sites* (Blackboard) (NetID login req
» *Schedule of classes*
» *Academic calendar*

Text search

Search course title and description:

Search

Browse by department/affiliation

- *Accounting*
- *African American Studies*
- *American Studies*
- *Anthropology*
- *Arab Studies*
- *Arabic*
- *Art*
- *Art and Museum Studies*
- *Art History*
- *Art, Music, and Theater*
- *Asian Studies*
- *Bachelor of Liberal Studies - Human Values*
- *Bachelor of Liberal Studies - Humanities & Social Sciences*
- *Biochemistry and Molecular & Cellular Biology*
- *Biology*
- *Biostatistics and Epidemiology*
- *Business Administration*
- *Catholic Studies*
- *Cell Biology*
- *Chemistry*
- *Chinese*
- *Classical Studies*
- *Classics*
- *Classics: Greek*
- *Classics: Latin*
- *Clinical/Translation Research*
- *Jewish Civilization*
- *Journalism*
- *Justice and Peace Studies*
- *Korean*
- *Latin American Studies*
- *Leisure and Recreation Education*
- *Liberal Studies - Human Values Evening Program*
- *Liberal Studies - Humanities and Social Sciences*
- *Linguistics*
- *Management*
- *Marketing*
- *Master of Professional Studies - Hospitality Management*
- *Master of Professional Studies - Human Resources*
- *Master of Professional Studies - Integrated Marketing Comm*
- *Master of Professional Studies - Journalism*
- *Master of Professional Studies - Public Relations*
- *Master of Professional Studies - Real Estate*
- *Master of Professional Studies - Sports Management*
- *Master of Professional Studies - Systems Engineering Management*
- *Master of Professional Studies - Technology Management*
- *Master of Professional Studies - Urban and Regional Planning*
- *Masters in Foreign Service*
- *Mathematics*
- *Medieval Studies*
- *Microbiology and Immunology*
- *Military Science*

QUERYING VS. BROWSING

- **Search queries tend to produce new, ad hoc collections of information that have not been gathered together before**
- **Browsing tend to be casual, mainly undirected exploration of navigation structures**

WHAT IS INFORMATION RETRIEVAL?

INFORMATION RETRIEVAL (IR)

- **A Computer Science Discipline studies search and provides solutions to it**
- **It involves a wide range of research areas and disciplines**
  - **psychology,**
  - **human-computer interactions,**
  - **math and statistics,**
  - **natural language understanding,**
  - **machine learning,**
  - **big data infrastructures,**
  - **distributed systems, …**

IR handles Unstructured data

- Typically refers to free text
- Allows
  - Keyword queries including operators
  - More sophisticated “concept” queries, e.g.,
    - find all web pages dealing with *drug abuse*

Databases handles structured data

Employee	Manager	Salary
Smith	Jones	50000
Chang	Smith	60000
Ivy	Smith	50000

THE MOST POPULAR IR TASK

Document Retrieval

- User issues a query
- Search engine returns a list of documents
  - in descending order of relevance
- It is to find relevant documents for the query
  - not to find answers - that is question answering
- Sometimes also known as “ad-hoc retrieval”

Google

infosense

Query

Web · Shopping · News · Videos

About 49,800 results (0.39 seconds)

[InfoSense: Sewer Line Cleaning, SL-RAT]
infosenseinc.com/
Sewer Line Rapid Assessment Tool (SL-RAT®) **InfoSense's** Revolutionary Acoustic Inspection Tool for Wastewater Collection System... Learn More ...

307 W Tremont Ave, Charlotte, NC 28203
(877) 747-3245
SL-RAT - Schedule a Demo

[About InfoSense, InfoSense Business Team | InfoSense]
infosenseinc.com/about-us/
InfoSense, Inc. is a technology-driven C-corporation spun out of the University of North Carolina at Charlotte to commercialize a promising acoustic inspection ...

[InfoSense Group]
infosense.cs.georgetown.edu/
Load More. This page is owned by the **InfoSense** group, Department of Computer Science, Georgetown University. Last Modified: 11/12/2014 13:33:58 PM.
You've visited this page many times. Last visit: 10/7/15

[Grace Hui Yang - InfoSense Group - Georgetown University]
infosense.cs.georgetown.edu/grace/
Dec 8, 2014 - I am leading the **InfoSense** group at Georgetown University. Check out more details from our group website. My research interests lie in ...
You've visited this page many times. Last visit: 10/17/15

[Infosense Solutions]
infosensesolutions.com/
Infosense Data Management in Data Warehouse, Data Integration and Analytics .

[Infosense: Understanding Information to Survive in the ...]
www.amazon.com › ... › Library Management · Amazon.com, Inc.

InfoSense, Inc.

Industrial Equipment Supplier

Address: 307 W Tremont Ave, Charlotte, NC 2

Phone: (877) 747-3245

Hours: Open today · 8:00 am – 5:00 pm

Reviews

Be the first to review

Search Results: a ranked list of documents

Document Retrieval Process

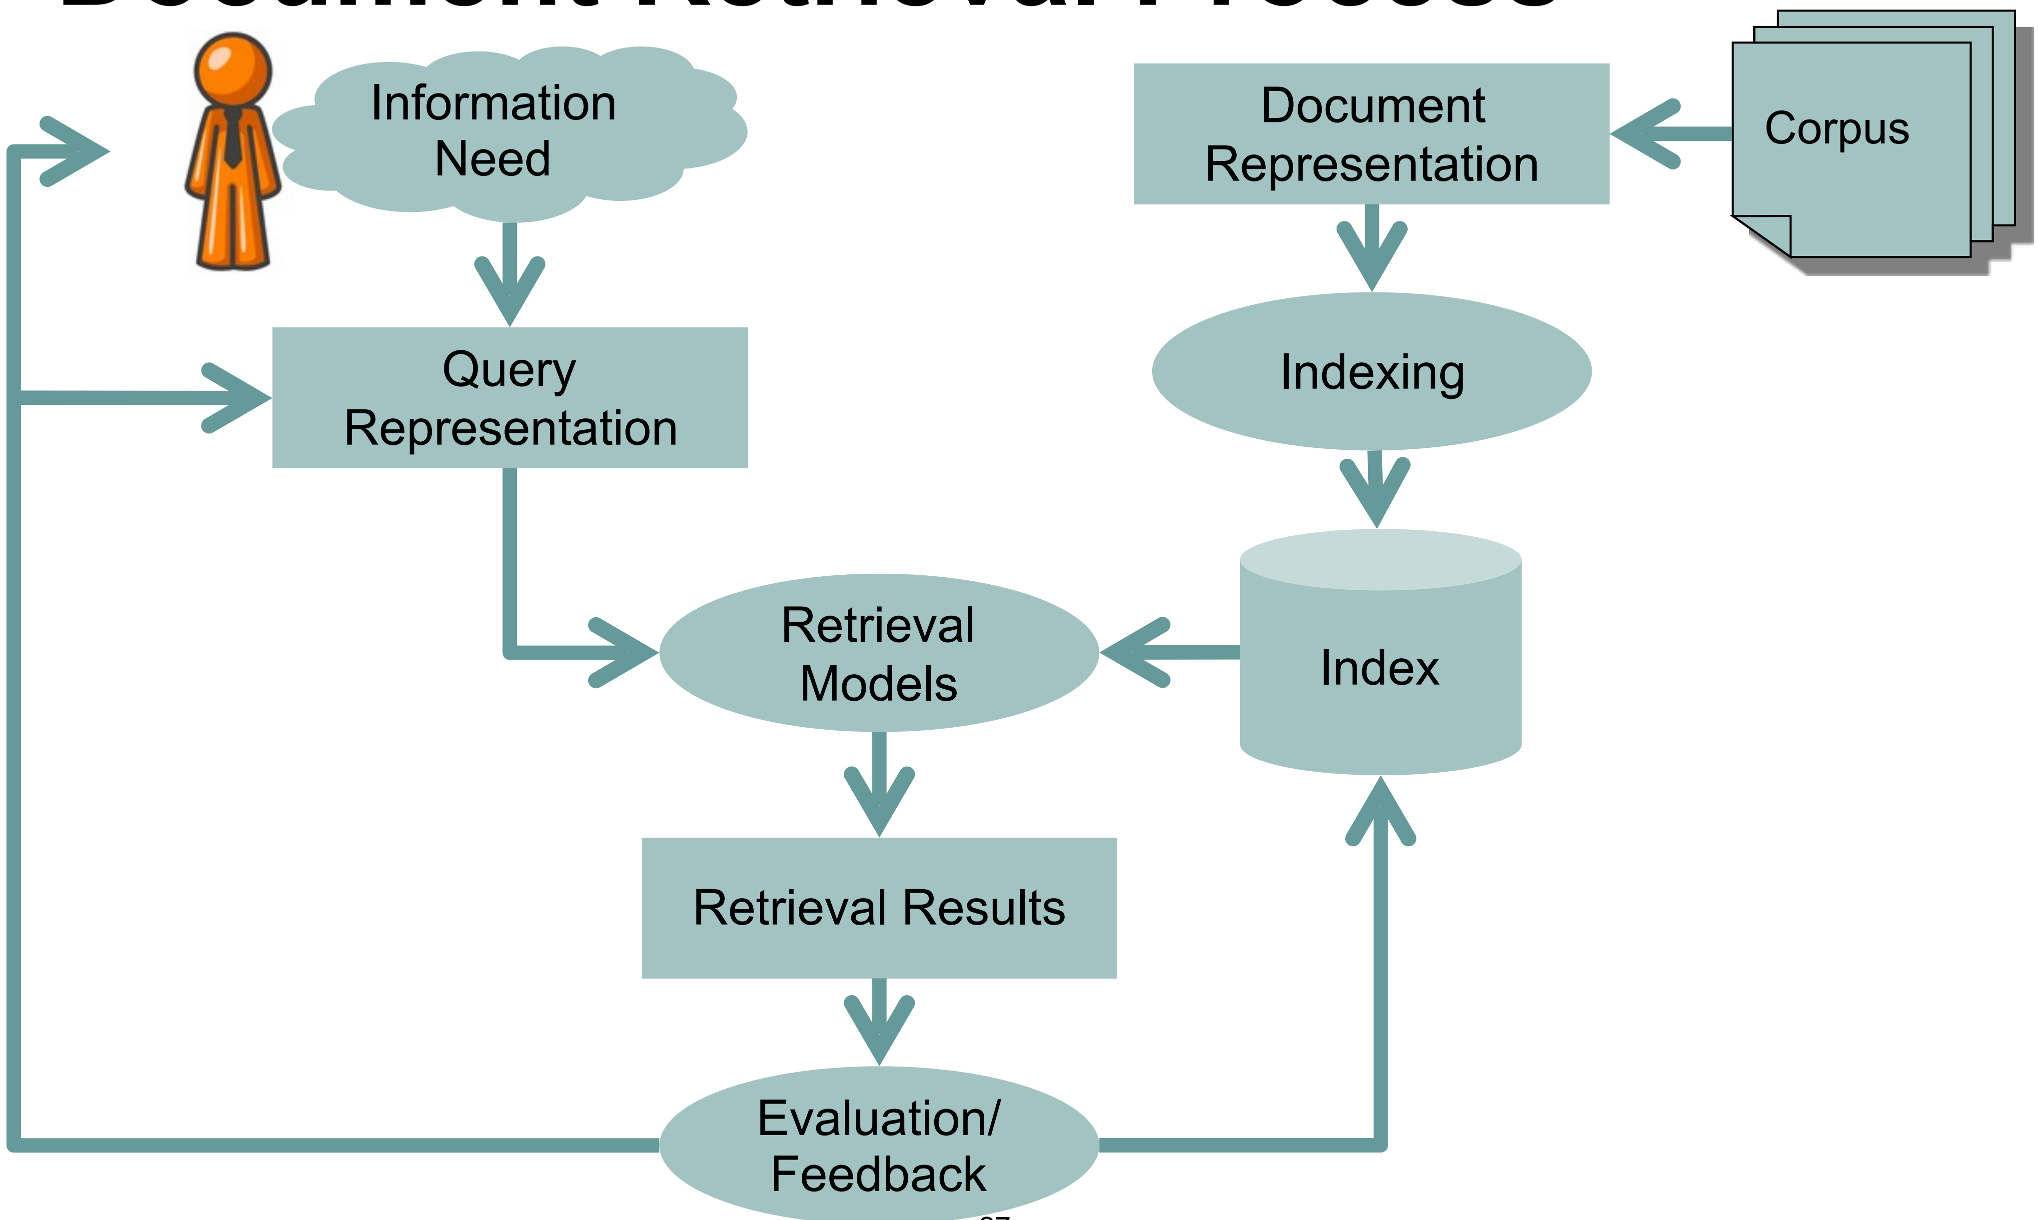

Terminology

- Query: text to represent an information need
- Document: a returned item in the index
- Term/token: a word, a phrase, an index unit
- Vocabulary: set of the unique tokens
- Corpus/Text collection
- Index/database: index built for a corpus
- Relevance feedback: judgment from human
- Evaluation Metrics: how good is a search system?
    - Precision, Recall, F1

Document Retrieval Process

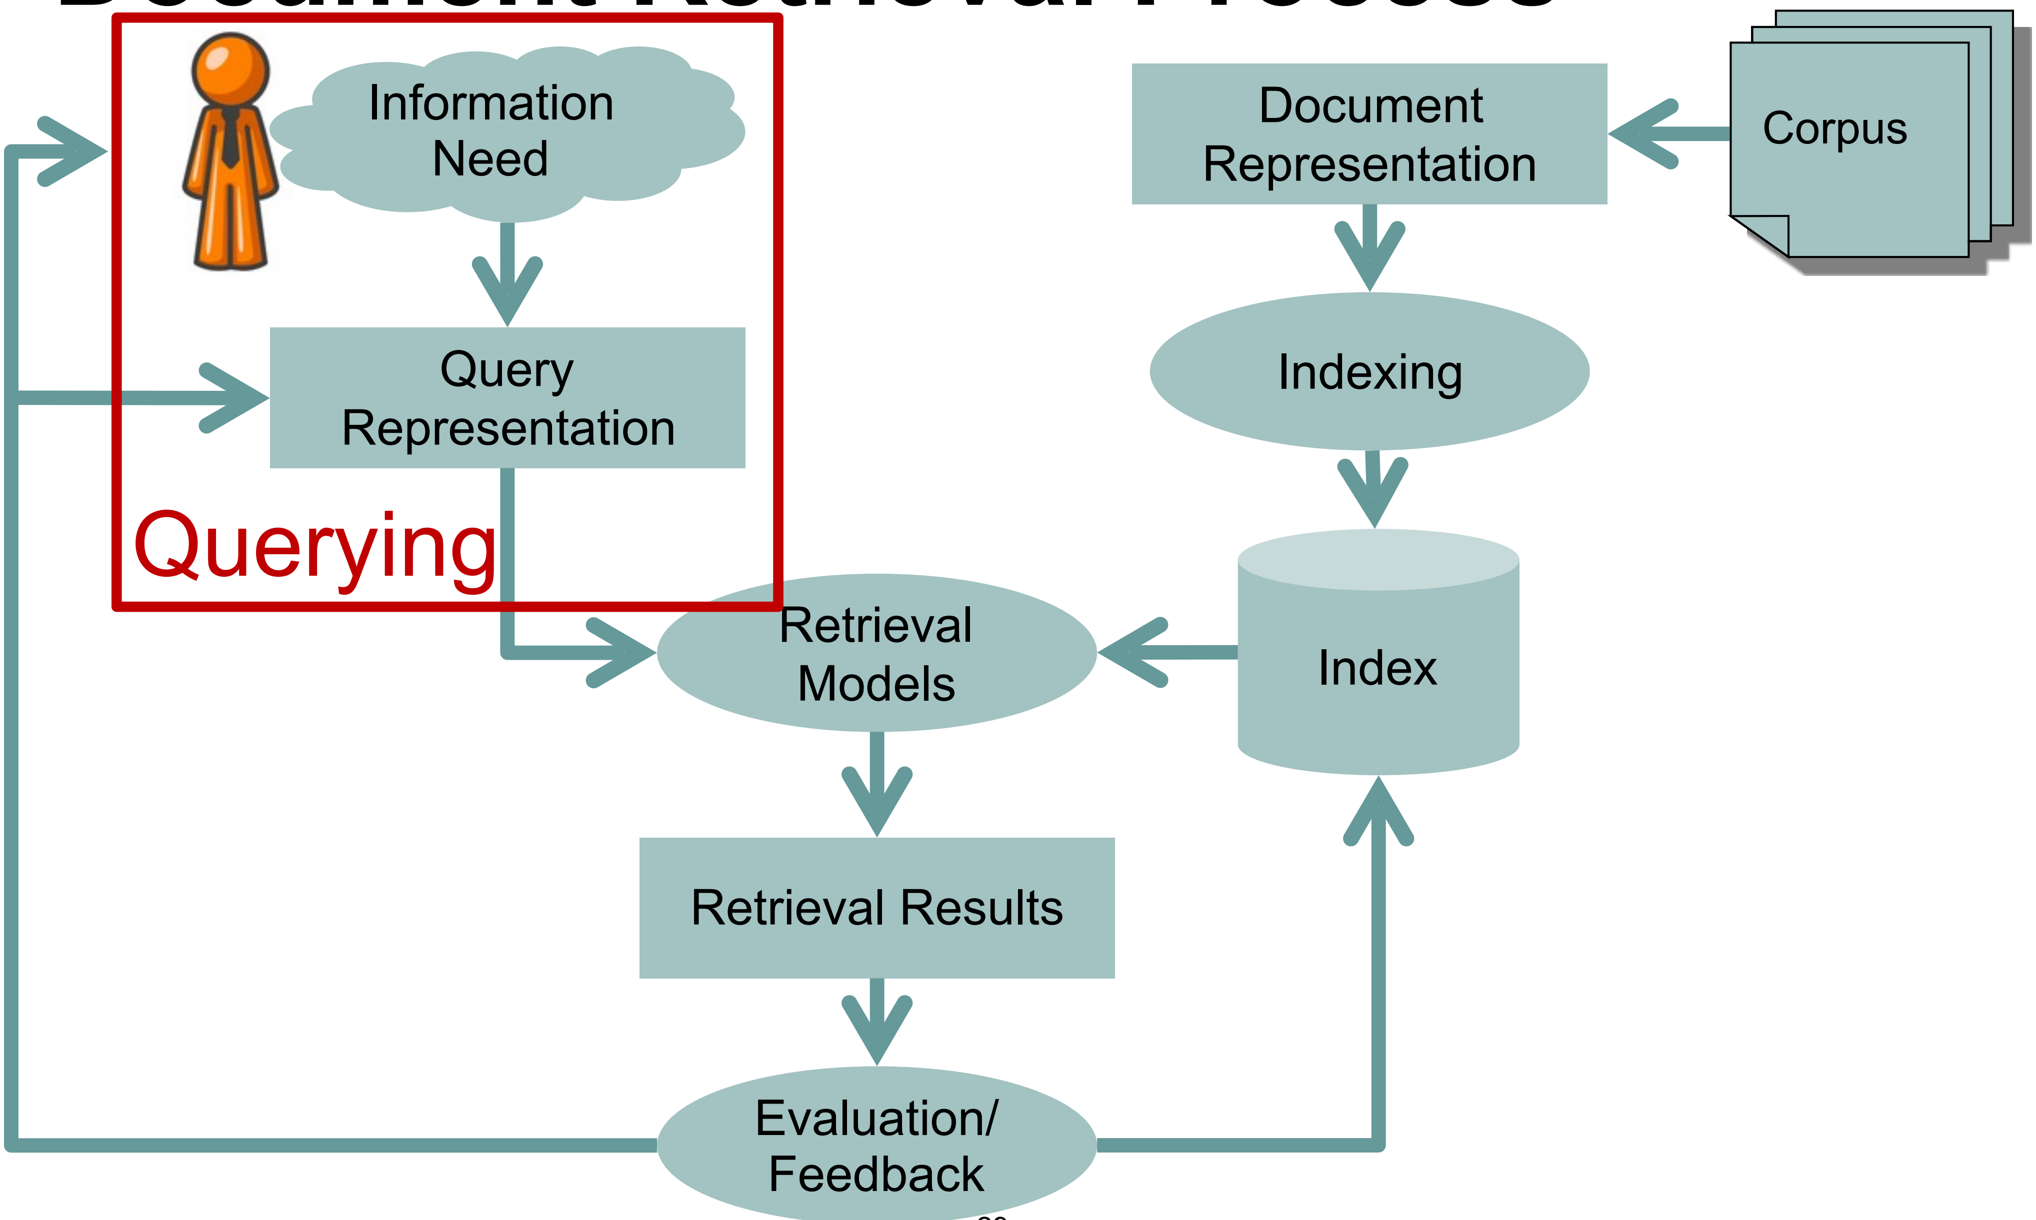

From Information Need to Query

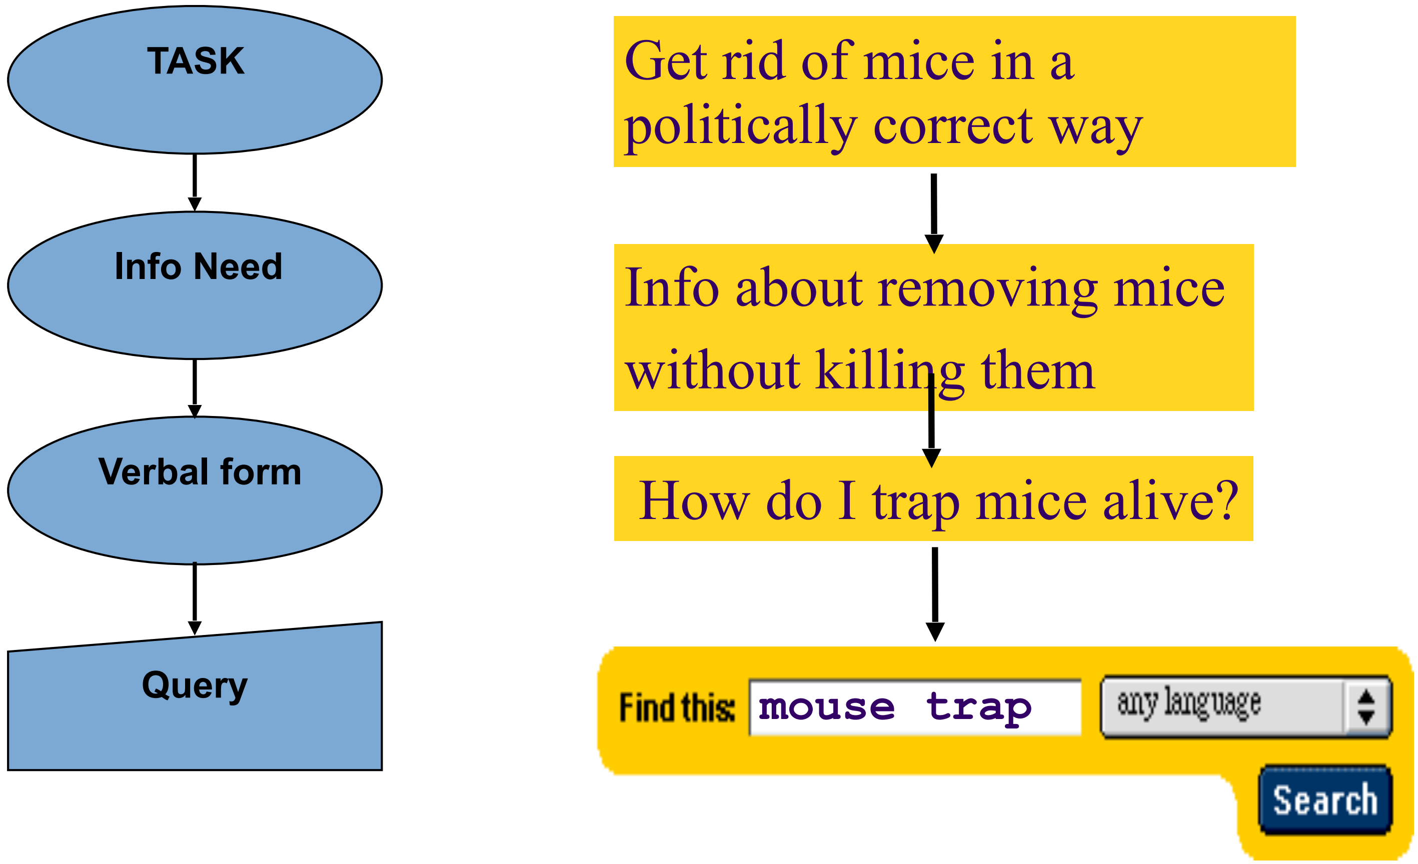

Document Retrieval Process

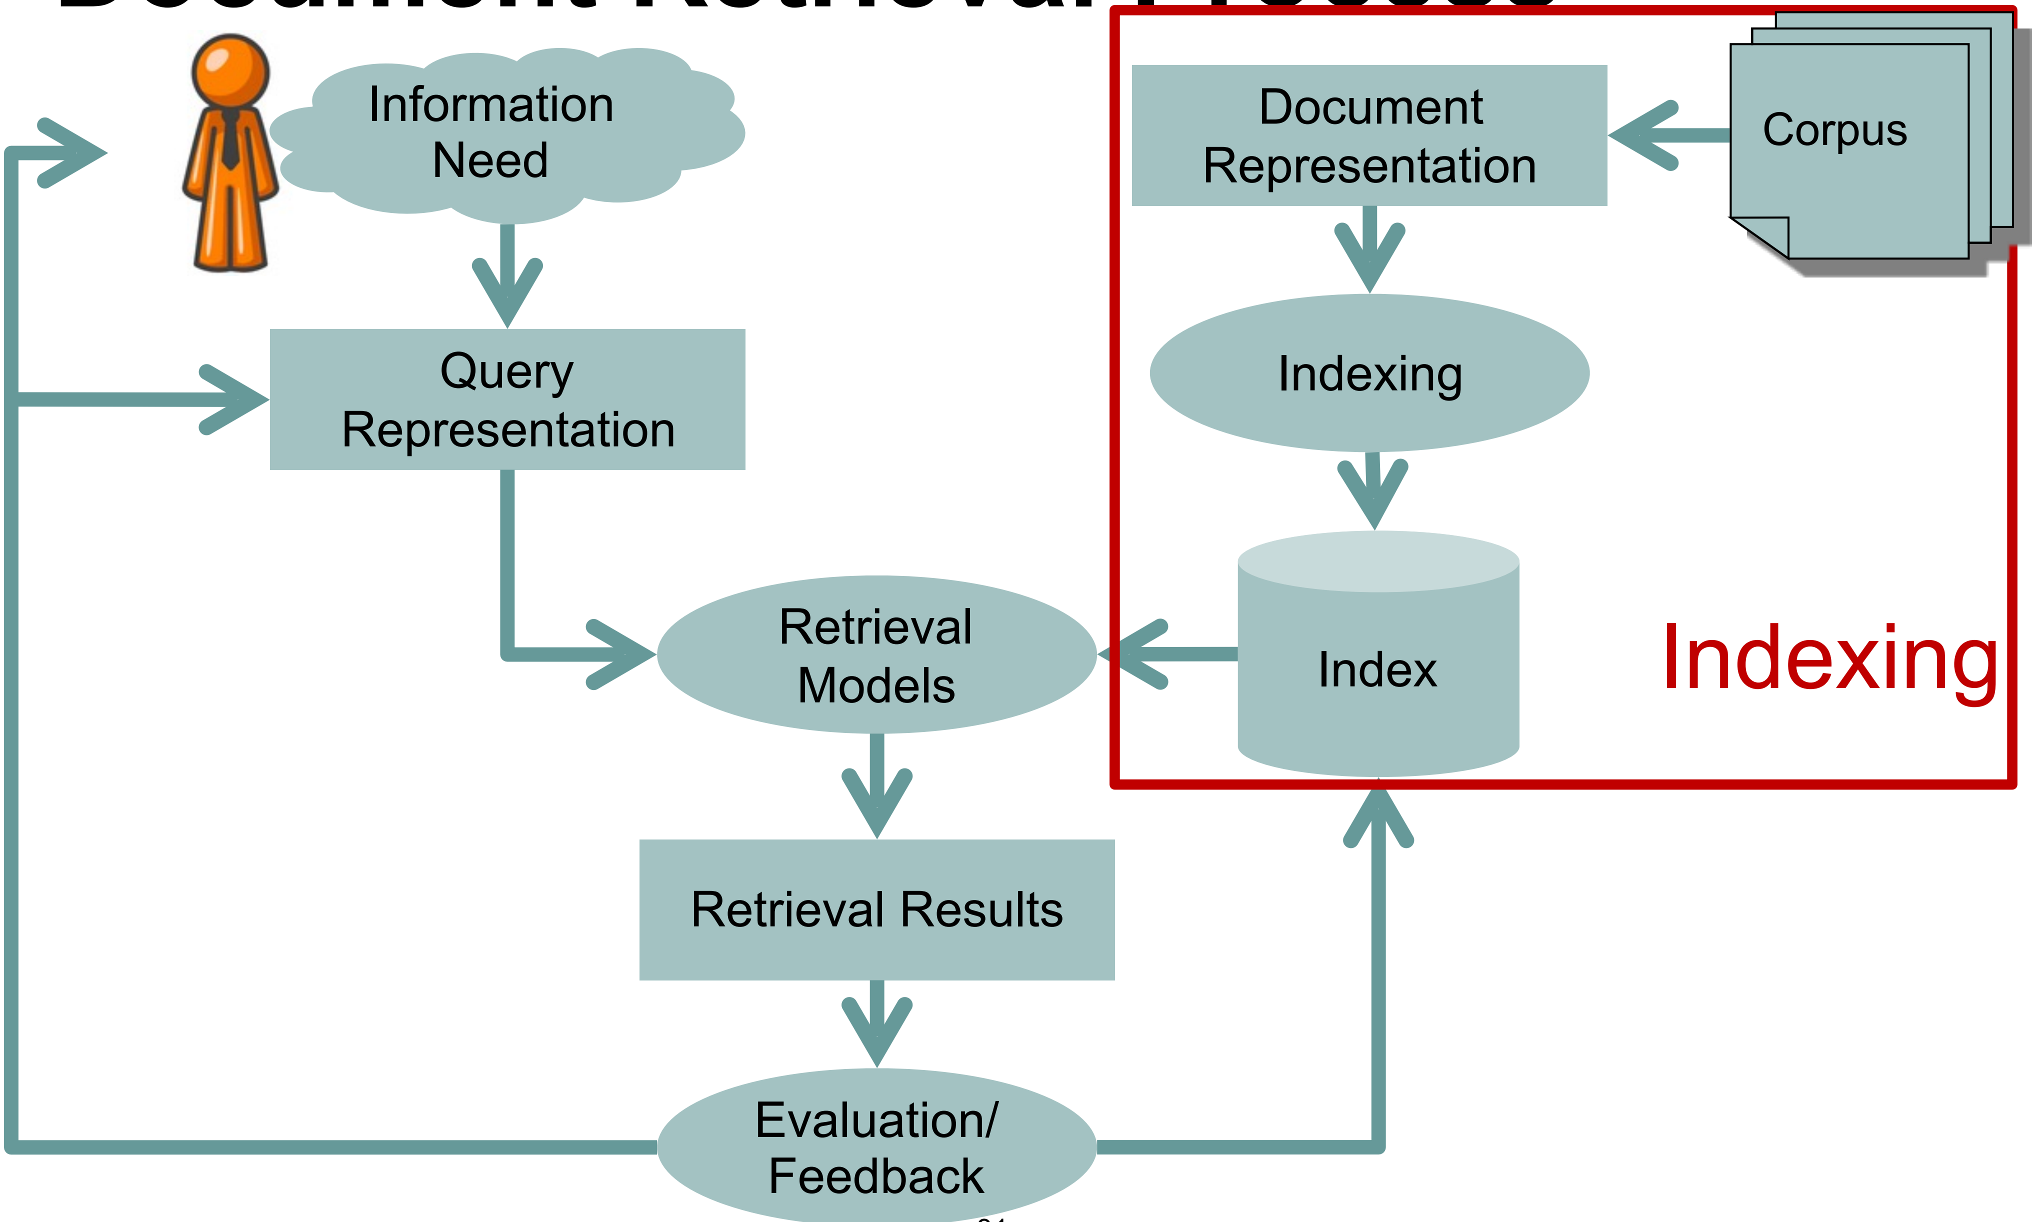

Term-Document Matrix

	Antony and Cleopatra	Julius Caesar	The Tempest	Hamlet	Othello	Macbeth
Antony	1	1	0	0	0	1
Brutus	1	1	0	1	0	0
Caesar	1	1	0	1	1	1
Calpurnia	0	1	0	0	0	0
Cleopatra	1	0	0	0	0	0
mercy	1	0	1	1	1	1
worser	1	0	1	1	1	0

1 if document contains term,
0 otherwise

Inverted index construction

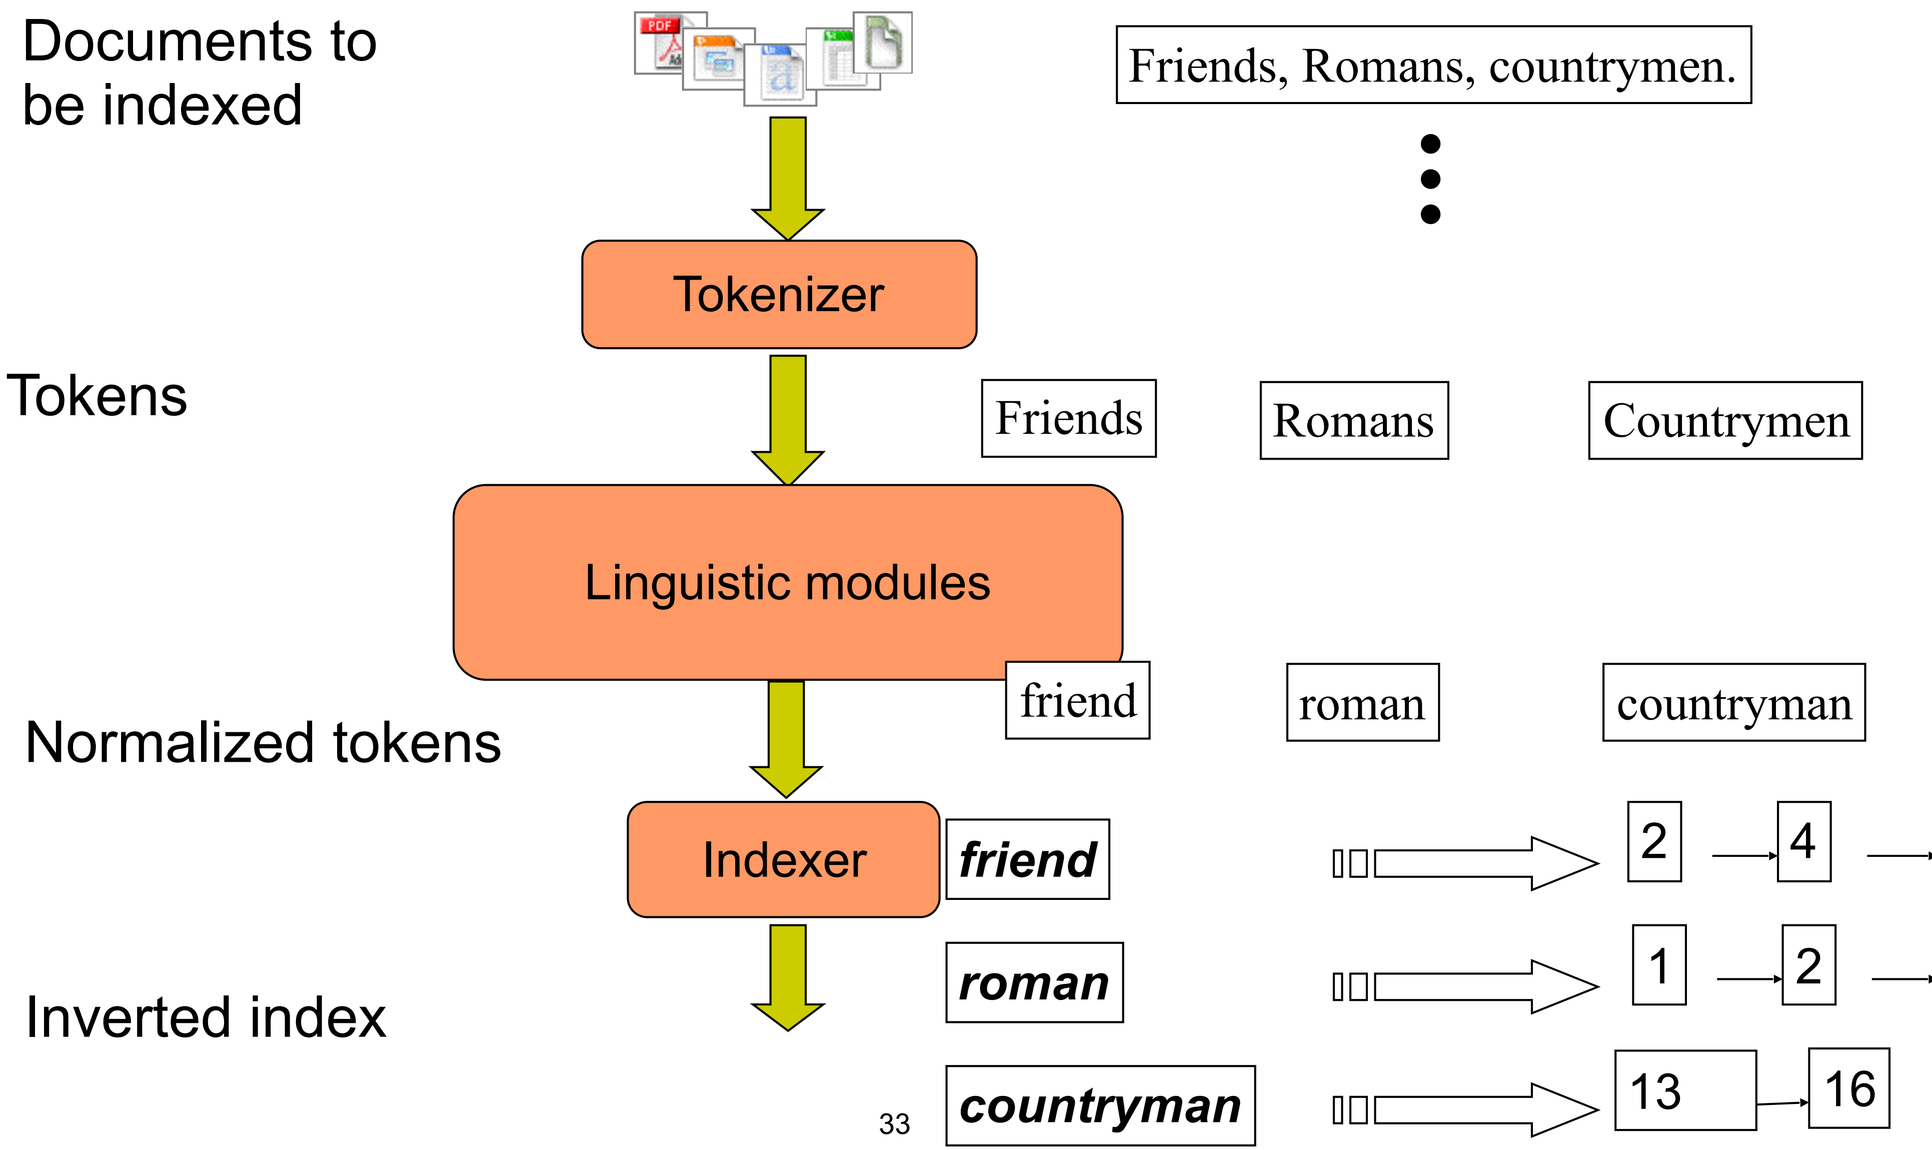

An Index

- Sequence of (Normalized token, Document ID) pairs.

Doc 1

I did enact Julius Caesar I was killed i' the Capitol; Brutus killed me.

Doc 2

So let it be with Caesar. The noble Brutus hath told you Caesar was ambitious

Term	docID
I	1
did	1
enact	1
julius	1
caesar	1
I	1
was	1
killed	1
i'	1
the	1
capitol	1
brutus	1
killed	1
me	1
so	2
let	2
it	2
be	2
with	2
caesar	2
the	2
noble	2
brutus	2
hath	2
told	2
you	2
caesar	2
was	2
ambitious	2

Document Retrieval Process

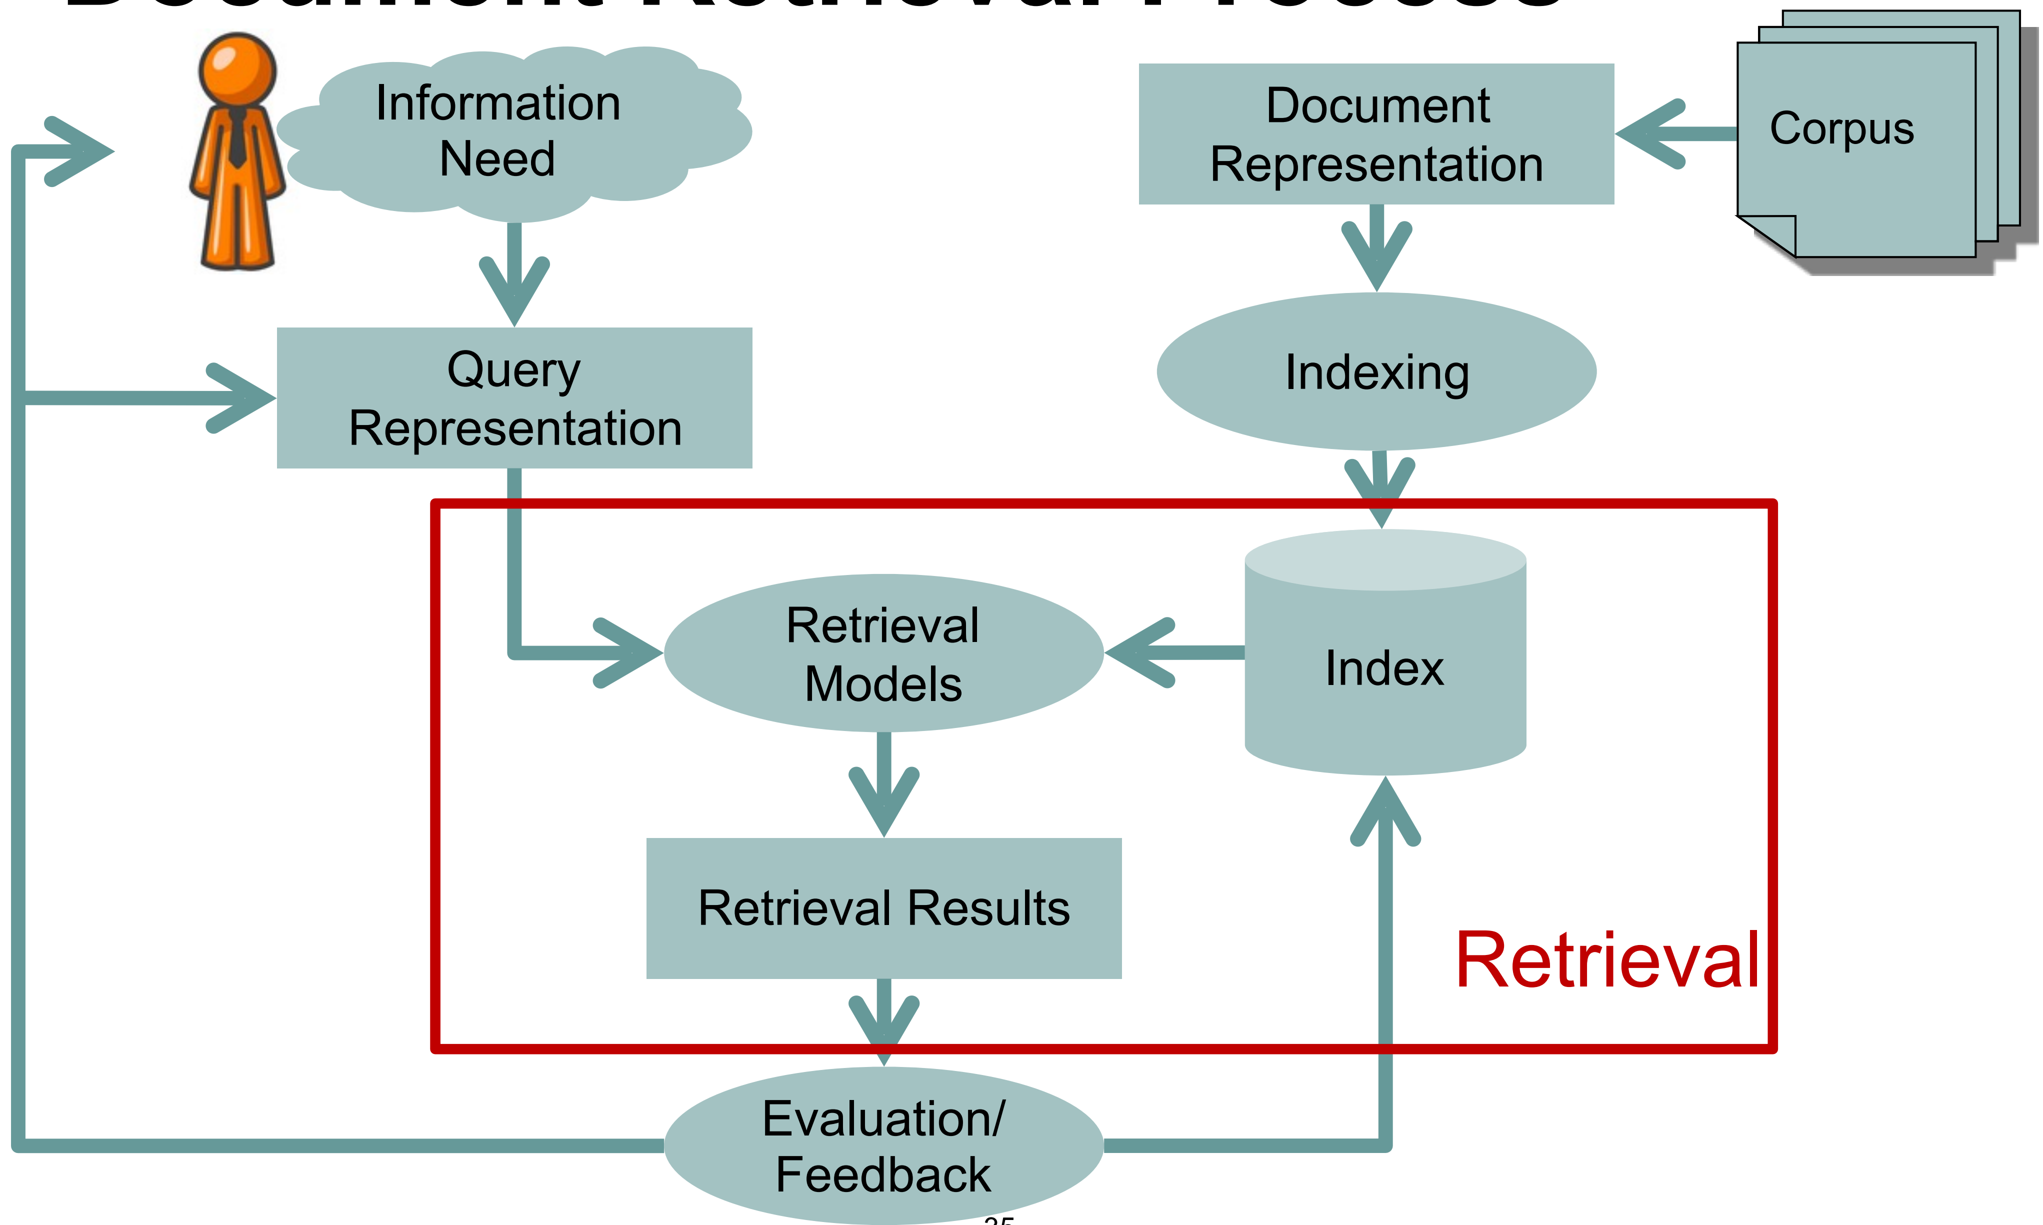

Boolean Retrieval

- The Boolean retrieval model is being able to ask a query that is a Boolean expression:
  - Boolean Queries use *AND, OR* and *NOT* to join query terms
    - Views each document as a set of words
    - Exact match: document matches condition or not.
  - Perhaps the simplest model to build an IR system on
- Primary commercial retrieval tool for 3 decades.
- Many search systems are still using Boolean
  - e.g. doctors and lawyers write very long and complex queries with boolean operators

Ranked Retrieval

- Boolean queries only give inclusion or exclusion of docs, not rankings
- Often we want to rank results
  - from the most relevant to the least relevant
  - Users are lazy
    - maybe only look at the first 10 results
- A good ranking is important

Google

infosense

Web | Shopping | News | Videos | Maps | More ▾ | Search tools

About 49,800 results (0.39 seconds)

InfoSense: Sewer Line Cleaning, SL-RAT
infosenseinc.com/ ▾
Sewer Line Rapid Assessment Tool (SL-RAT®) **InfoSense's** Revolutionary Acoustic Inspection Tool for Wastewater Collection System... Learn More ...

307 W Tremont Ave, Charlotte, NC 28203
(877) 747-3245
SL-RAT - Schedule a Demo

About InfoSense, InfoSense Business Team | InfoSense
infosenseinc.com/about-us/ ▾
InfoSense, Inc. is a technology-driven C-corporation spun out of the University of North Carolina at Charlotte to commercialize a promising acoustic inspection ...

InfoSense Group
infosense.cs.georgetown.edu/ ▾
Load More. This page is owned by the **InfoSense** group, Department of Computer Science, Georgetown University. Last Modified: 11/12/2014 13:33:58 PM.
You've visited this page many times. Last visit: 10/7/15

Grace Hui Yang - InfoSense Group - Georgetown University
infosense.cs.georgetown.edu/grace/ ▾
Dec 8, 2014 - I am leading the **InfoSense** group at Georgetown University. Check out more details from our group website. My research interests lie in ...
You've visited this page many times. Last visit: 10/17/15

Infosense Solutions
infosensesolutions.com/ ▾
Infosense Data Management in Data Warehouse, Data Integration and Analytics .

Infosense: Understanding Information to Survive in the ...
www.amazon.com › ... › Library Management ▾ Amazon.com, Inc. ▾
Infosense: Understanding Information to Survive in the Knowledge Society [Keith, J

InfoSense, Inc.

Industrial Equipment Supplier

Address: 307 W Tremont Ave, Charlotte, NC 2

Phone: (877) 747-3245

Hours: Open today · 8:00 am – 5:00 pm ▾

Reviews

Be the first to review

Document Retrieval Process

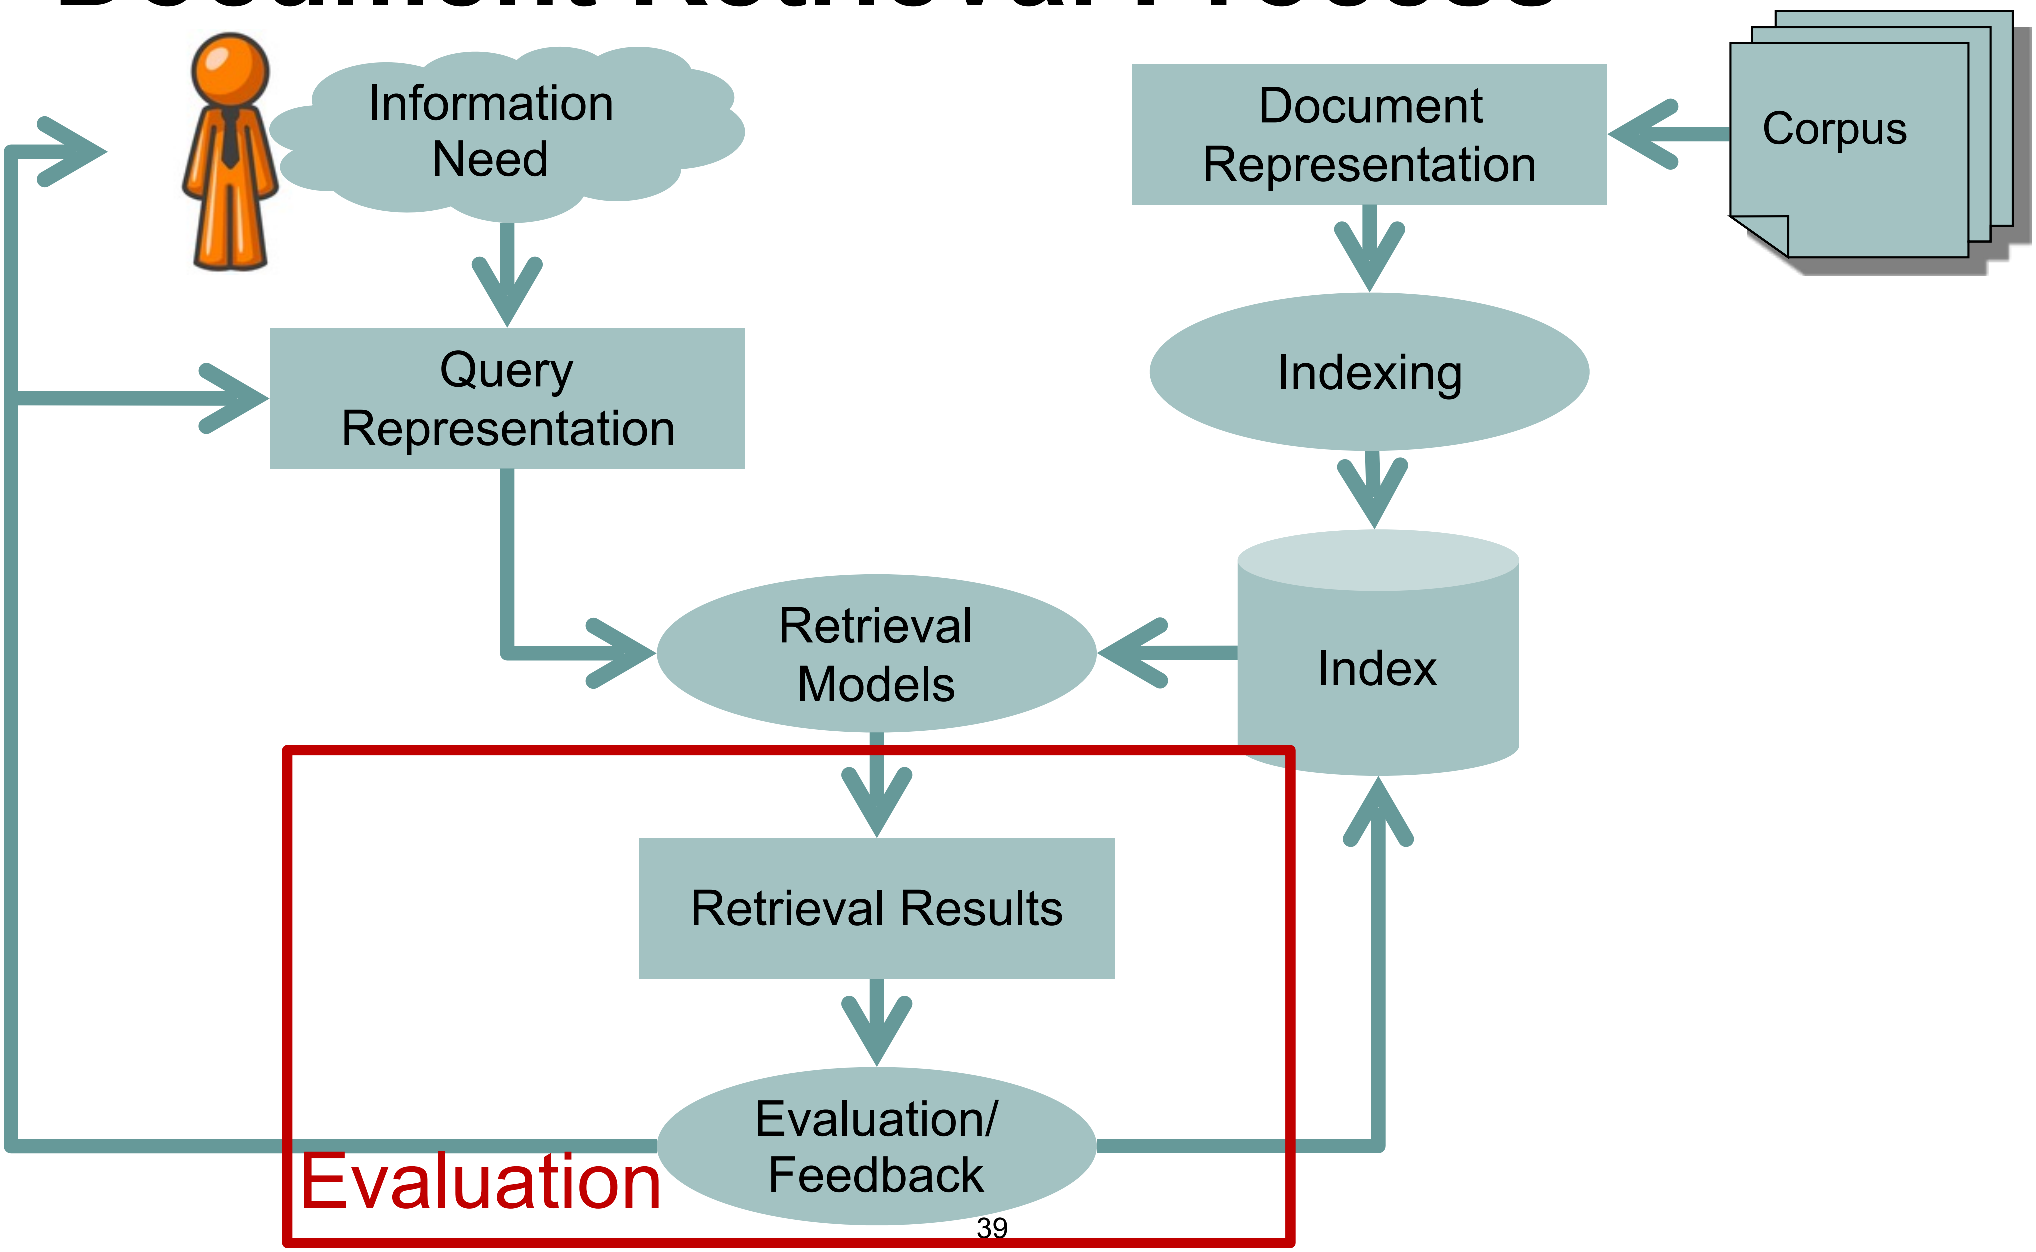

HOW TO JUDGE RELEVANCE

- **If a document (a result) can satisfy my information need**
- **If a document contains the answer to my query**

SEARCH VS. IR EVALUATION

- It is true that our understanding of what is a relevant document will help develop both search algorithms and evaluation methods
- However, they are two independent processes
- Search systems:
    - automated search engines analyze queries and documents then return the documents that search engines think that you might think they are relevant
- Evaluation:
    - manual methods to judge whether the results retuned by a search engine are relevant;
    - usually evaluated in batch

IR Evaluation metrics

- Precision
  - Correctly returned results / returned results
  - How good the returned ones are?
- Recall
  - Correctly returned results / all correct results
  - How many good ones can you find?
- NDCG
  - Normalized Discounted Cumulated Gain
  - Graded ratings to the results: Good, better, best
  - Popular for Web Search
- AB Test
  - side by side test to compare two lists of search results

Document Retrieval Process

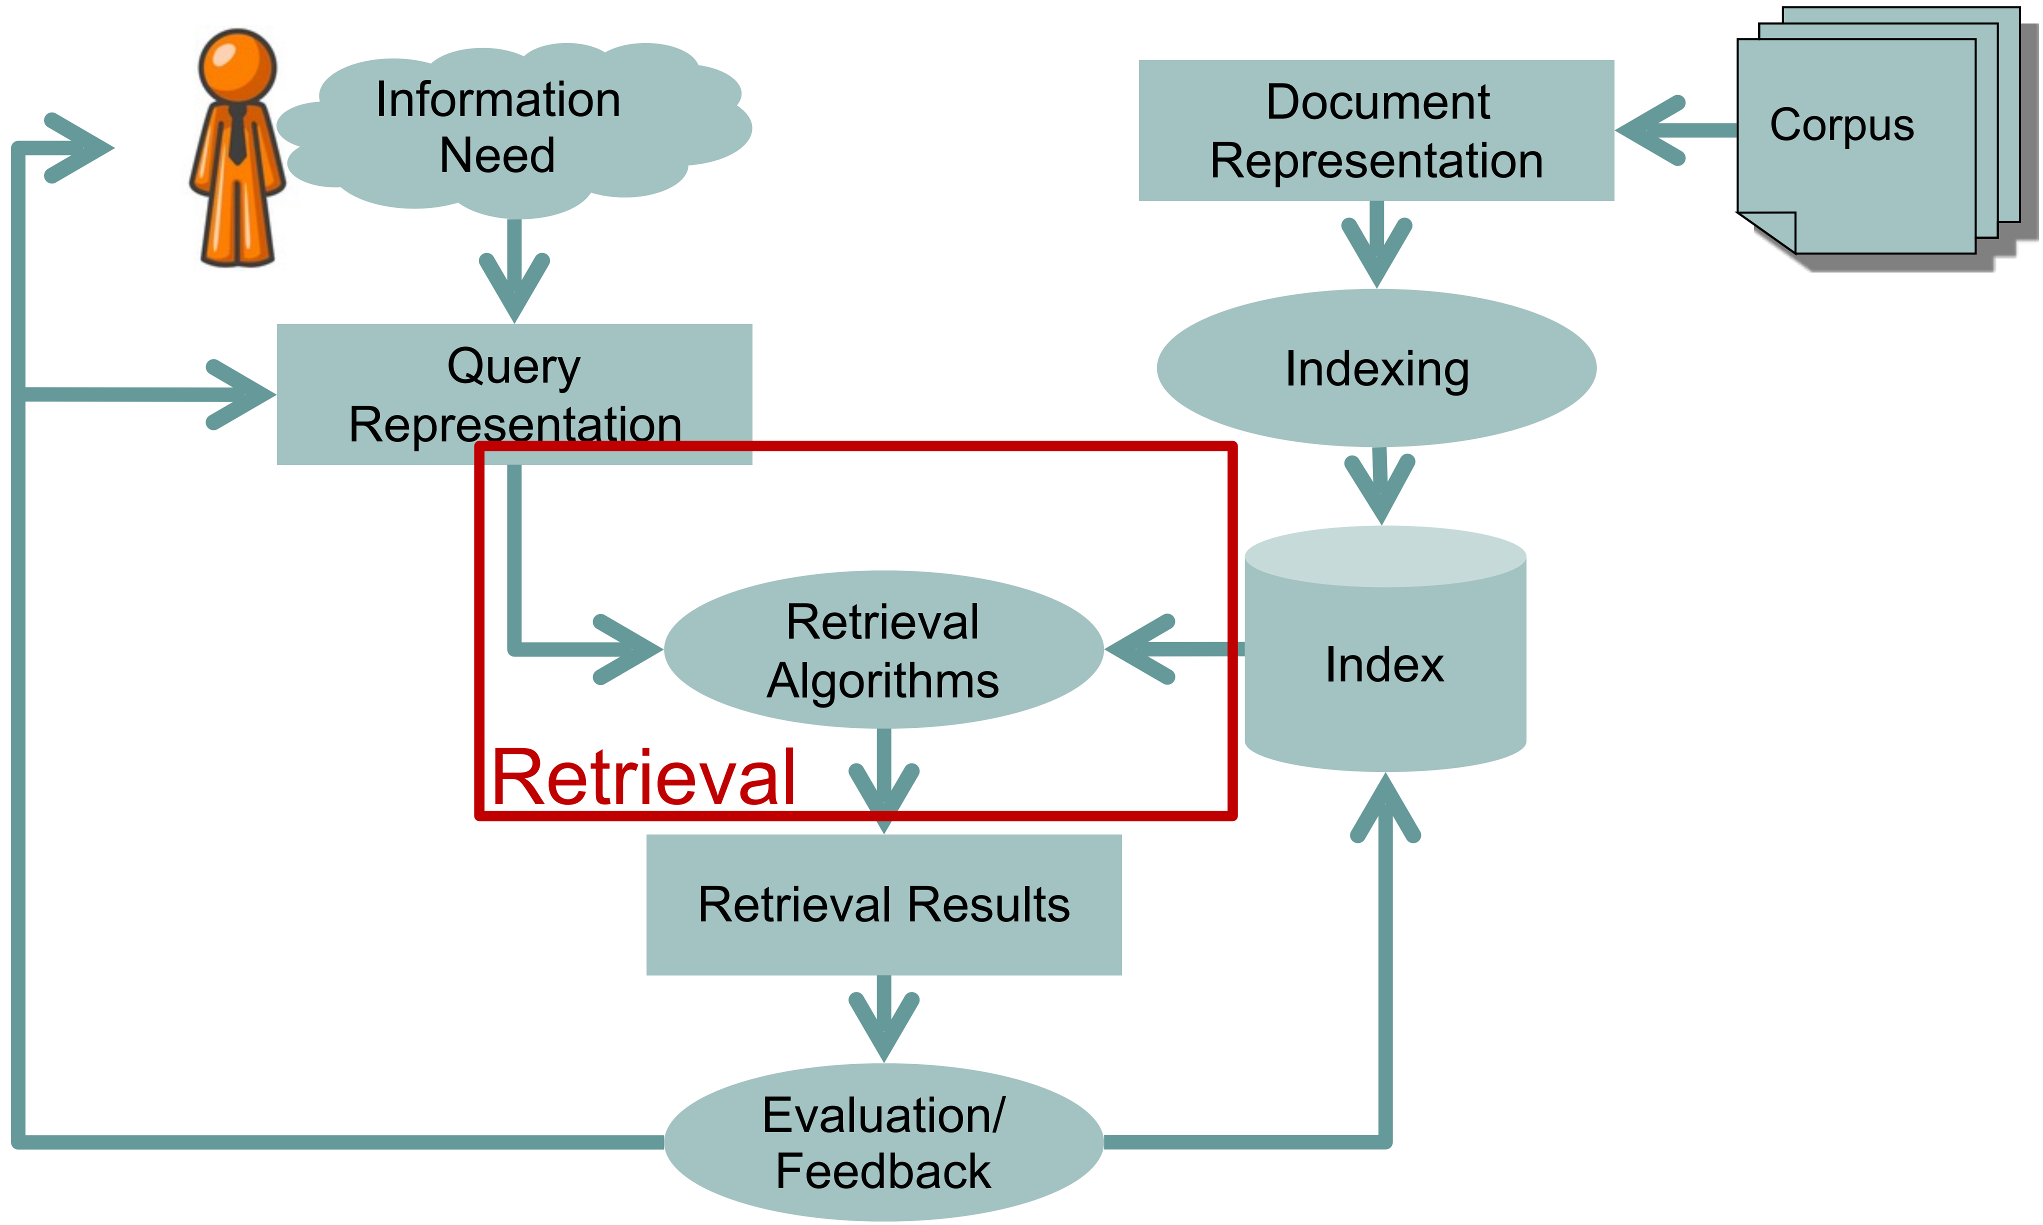

How to find the relevant documents for a query?

- By keyword matching
  - boolean models
- By similarity
  - vector space model
- By imaging how to write out a query
  - how likely a query is written with this document in mind
  - generate with some randomness
  - query generation language model
- By trusting how other people think about the documents /web pages
  - link-based methods, pagerank, hits

Vector Space Model

Vector Space Model

- Treat query as a tiny document
- Represent both query and documents as word vectors in a word space
- Rank documents according to their proximity to the query in the space of words

Represent Documents in a Space of Word Vectors

Suppose the corpus only has two words: 'Jealous' and 'Gossip'

They form a space of "Jealous" and "Gossip"

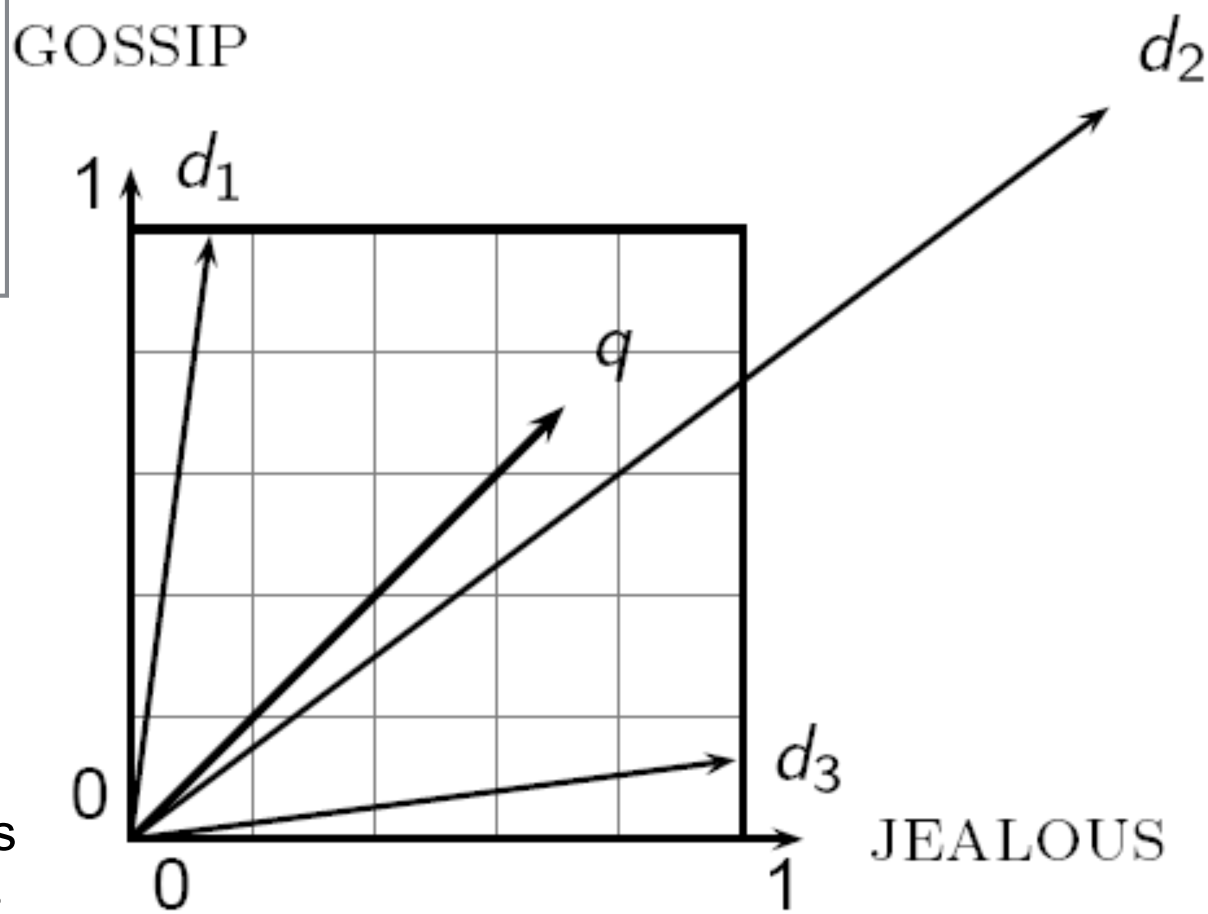

d1: gossip gossip jealous gossip gossip gossip gossip gossip gossip gossip gossip

d2: gossip gossip jealous gossip gossip gossip gossip gossip gossip gossip jealous jealous jealous jealous jealous jealous jealous gossip jealous

d3: jealous gossip jealous jealous jealous jealous jealous jealous jealous jealous jealous

q: gossip gossip jealous gossip gossip gossip gossip gossip jealous jealous jealous jealous

Calculate the Query-Document Similarity

Formalizing vector space proximity

- First cut: distance between the end points of the two vectors?

How to do it?

Euclidean Distance

- In mathematics, the Euclidean distance or Euclidean metric is the "ordinary" (i.e. straight-line) distance between two points in Euclidean space.

- If if p = (p1, p2,..., pn) and q = (q1, q2,..., qn) are two points in the Euclidean space, their Euclidean distance is

$$\mathrm{d}(\mathbf{p},\mathbf{q}) = \mathrm{d}(\mathbf{q},\mathbf{p}) = \sqrt{(q_1-p_1)^2 + (q_2-p_2)^2 + \cdots + (q_n-p_n)^2}$$

$$= \sqrt{\sum_{i=1}^{n}(q_i-p_i)^2}.$$

In a space of ‘Jealous’ and ‘Gossip’

Here, if you look at the content (or we say the word distributions) of each document, d2 is actually the most similar document to q

However, d2 produces a bigger distance score to q

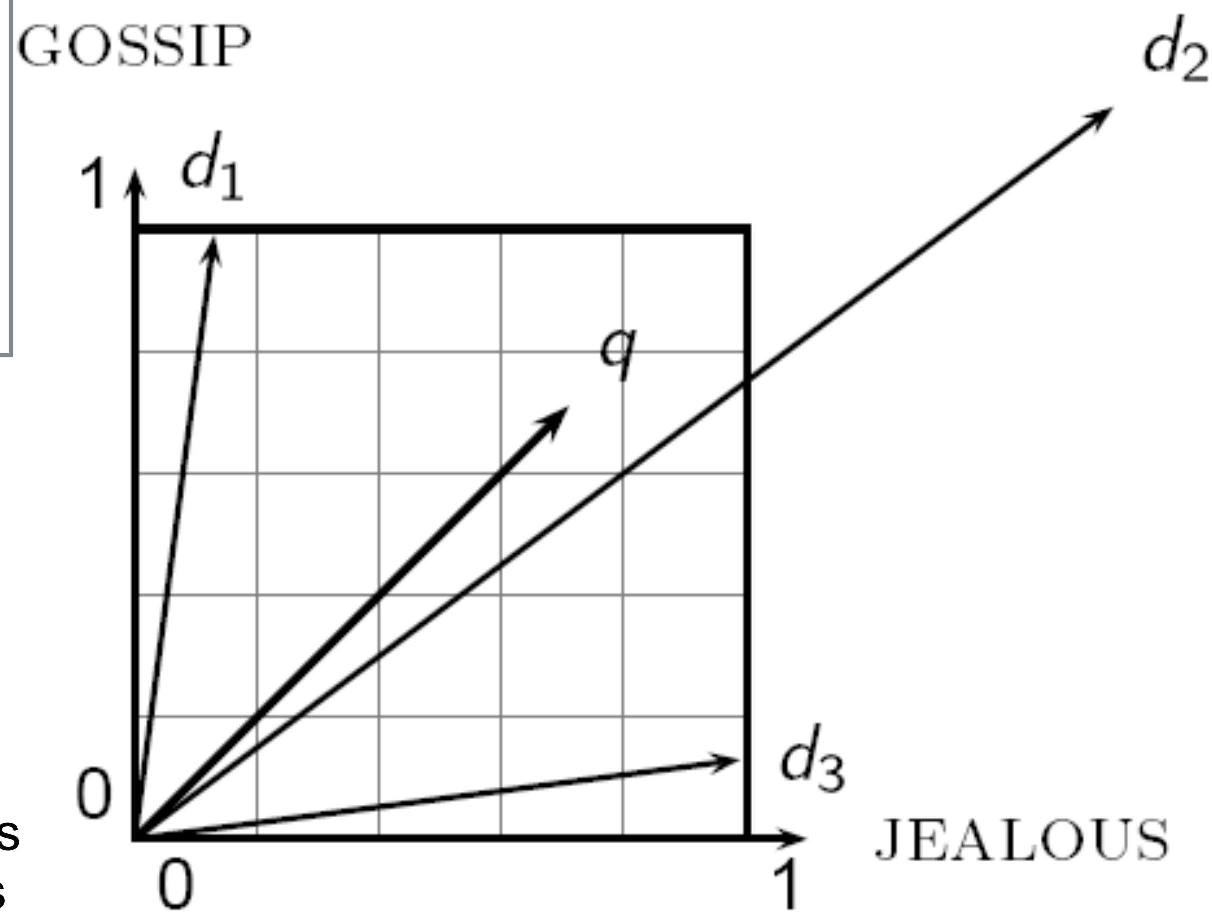

d1: gossip gossip jealous gossip gossip gossip gossip gossip gossip gossip gossip

d2: gossip gossip jealous gossip gossip gossip gossip gossip gossip gossip jealous jealous jealous jealous jealous jealous jealous gossip jealous

d3: jealous gossip jealous jealous jealous jealous jealous jealous jealous jealous jealous

q: gossip gossip jealous gossip gossip gossip gossip gossip jealous jealous jealous jealous

In a space of ‘Jealous’ and ‘Gossip’

The Euclidean distance between $\vec{q}$ and $\vec{d_2}$ is large even though the distribution of terms in the query $\vec{q}$ and the distribution of terms in the document $\vec{d_2}$ are very similar.

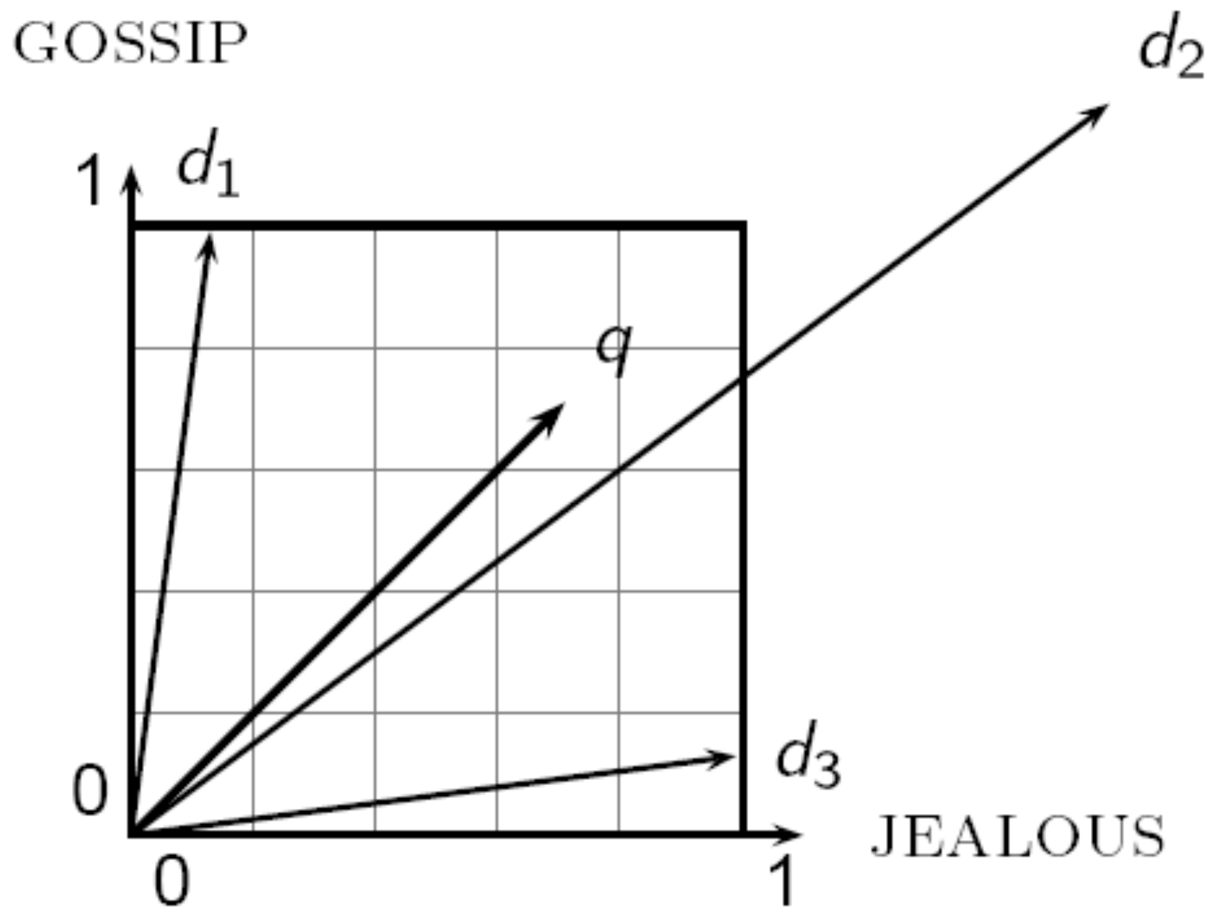

Why Euclidean Distance is A Bad Idea for Query-Document Similarity

- Because Euclidean distance is large for vectors with different lengths.
  - short query and long documents will be always have big Euclidean distance
  - we cannot rank them fairly, as compared with others
  - not possible to get a universal ranking

- How can we do better?

Use angle instead of distance

- Key idea: Rank documents according to angle with query

- The angle between similar vectors is small, between dissimilar vectors is large.

- This is exactly what we need to score a query-document pair.

- This is equivalent to perform a document length normalization

Cosine similarity illustrated

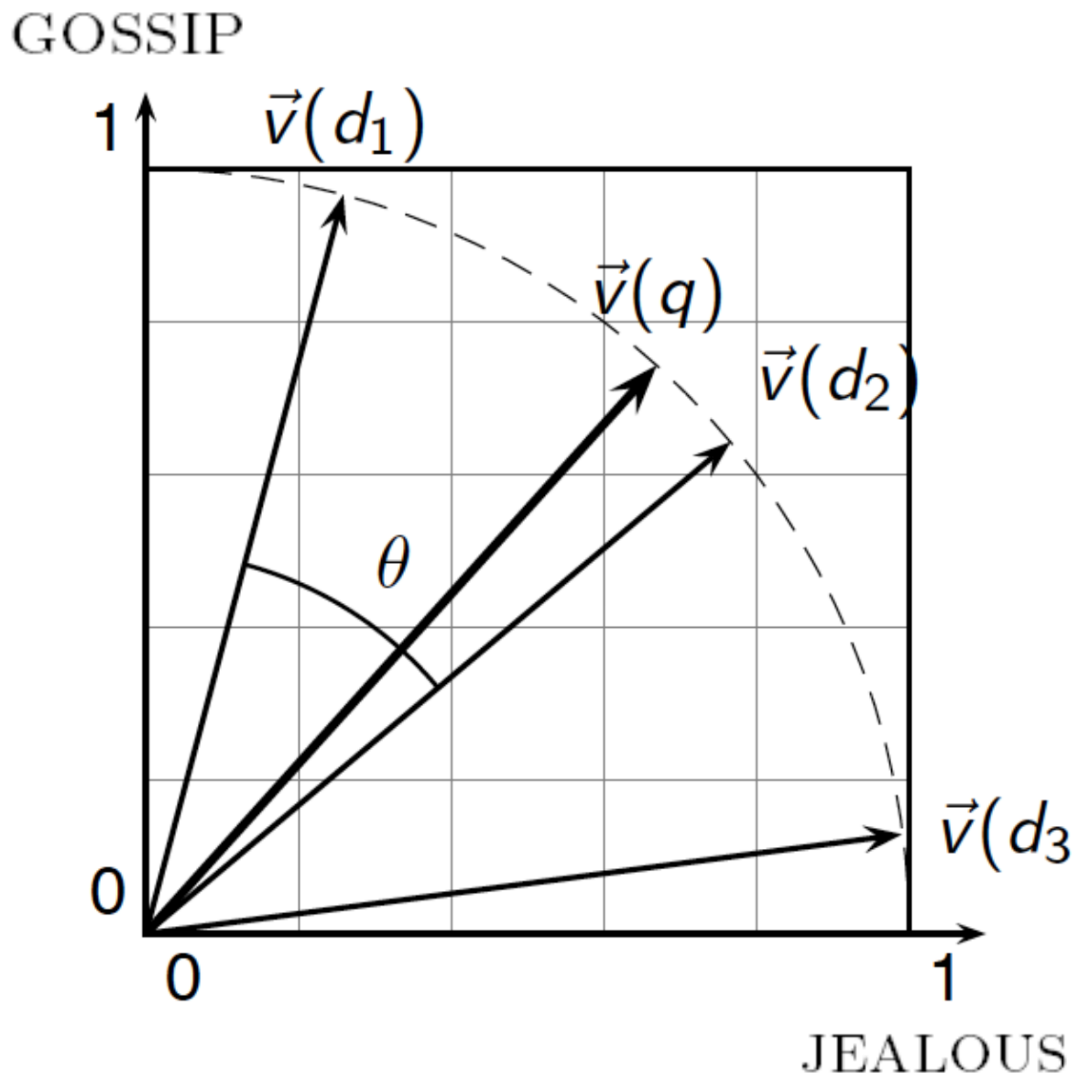

Cosine Similarity

$$\cos(\vec{q},\vec{d}) = \frac{\overbrace{\vec{q}\bullet\vec{d}}^{\text{Dot product}}}{|\vec{q}||\vec{d}|} = \overbrace{\frac{\vec{q}}{|\vec{q}|}\bullet\frac{\vec{d}}{|\vec{d}|}}^{\text{Unit vectors}} = \frac{\sum_{i=1}^{|V|} q_i d_i}{\sqrt{\sum_{i=1}^{|V|} q_i^2}\sqrt{\sum_{i=1}^{|V|} d_i^2}}$$

q_i is the tf-idf weight of term *i* in the query
d_i is the tf-idf weight of term *i* in the document

$\cos(\overrightarrow{q},\overrightarrow{d})$ is the cosine similarity of $\overrightarrow{q}$ and $\overrightarrow{d}$ … or, equivalently, the cosine of the angle between $\overrightarrow{q}$ and $\overrightarrow{d}$.

Exercise

Consider two documents D_1, D_2 and a query Q

$D_1 = (0.5, 0.8, 0.3)$, $D_2 = (0.9, 0.4, 0.2)$, $Q = (1.5, 1.0, 0)$

Results

Consider two documents D_1, D_2 and a query Q

D_1 = (0.5, 0.8, 0.3), D_2 = (0.9, 0.4, 0.2), Q = (1.5, 1.0, 0)

$$
\begin{aligned}
Cosine(D_1, Q) &= \frac{(0.5 \times 1.5) + (0.8 \times 1.0)}{\sqrt{(0.5^2 + 0.8^2 + 0.3^2)(1.5^2 + 1.0^2)}} \\
&= \frac{1.55}{\sqrt{(0.98 \times 3.25)}} = 0.87
\end{aligned}
$$

$$
\begin{aligned}
Cosine(D_2, Q) &= \frac{(0.9 \times 1.5) + (0.4 \times 1.0)}{\sqrt{(0.9^2 + 0.4^2 + 0.2^2)(1.5^2 + 1.0^2)}} \\
&= \frac{1.75}{\sqrt{(1.01 \times 3.25)}} = 0.97
\end{aligned}
$$

What are the numbers in a vector?

D_1 = (0.5, 0.8, 0.3), D_2 = (0.9, 0.4, 0.2), Q = (1.5, 1.0, 0)

- They are **term weights**
- **to indicate the importance of a term in a document**

Summary: Vector Space Model

- Advantages
  - Simple computational framework for ranking documents given a query
  - Any similarity measure or term weighting scheme could be used
- Disadvantages
  - Assumption of term independence

Evolving IR

- Paradigm shifts in IR as new models emerge
- e.g. VSM → BM25 → Language Model
  - Different ways of defining relationship between query and document
- Static → Interactive → Dynamic
  - Evolution in **modeling** user interaction with search engine

Conceptual Model – Static IR

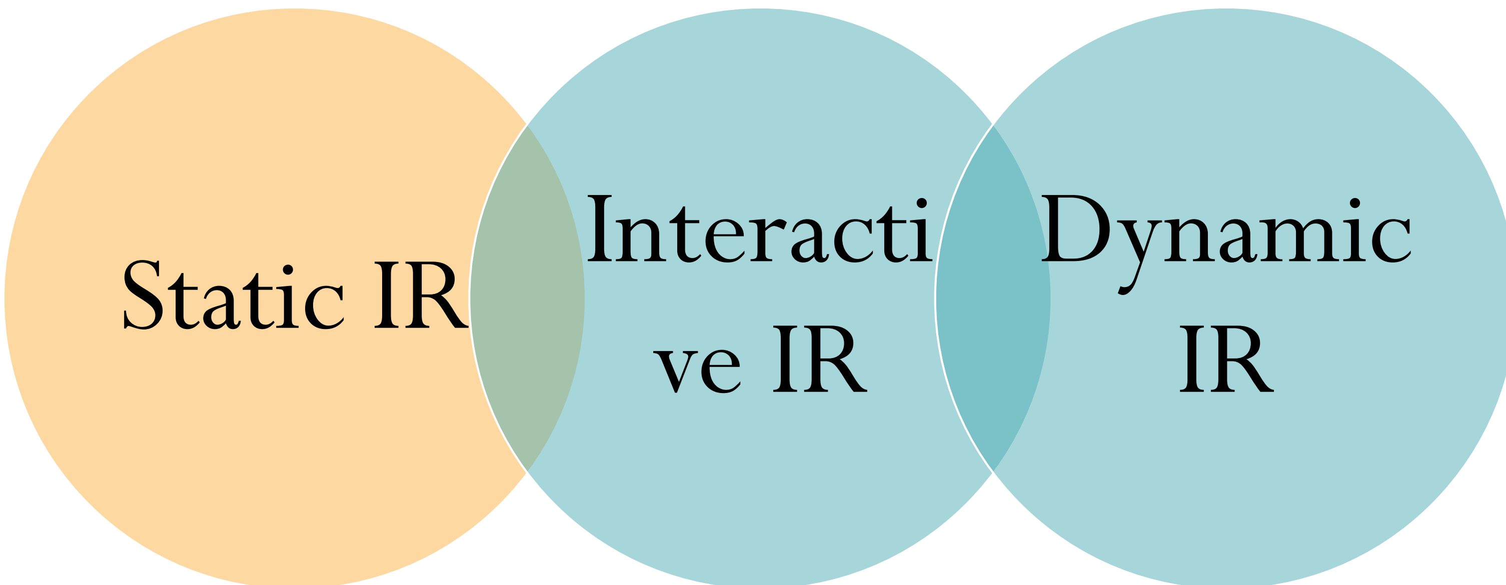

- No feedback

Conceptual Model – Interactive IR

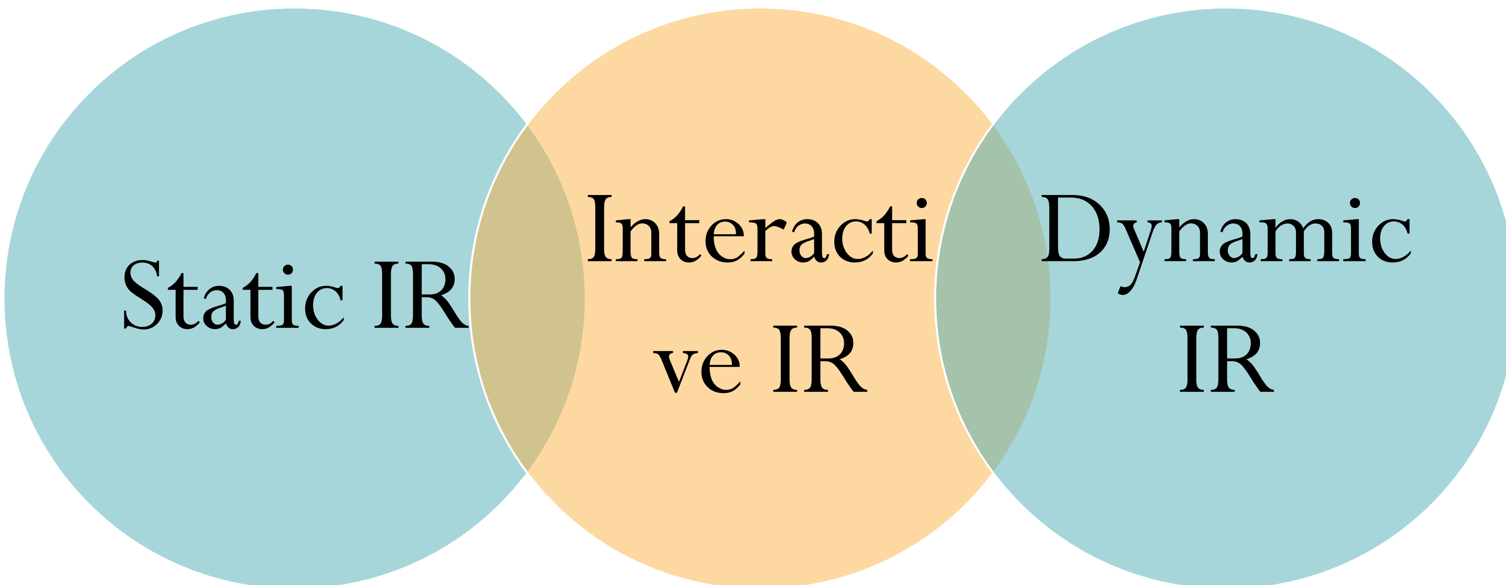

- Exploit Feedback

Conceptual Model – Dynamic IR

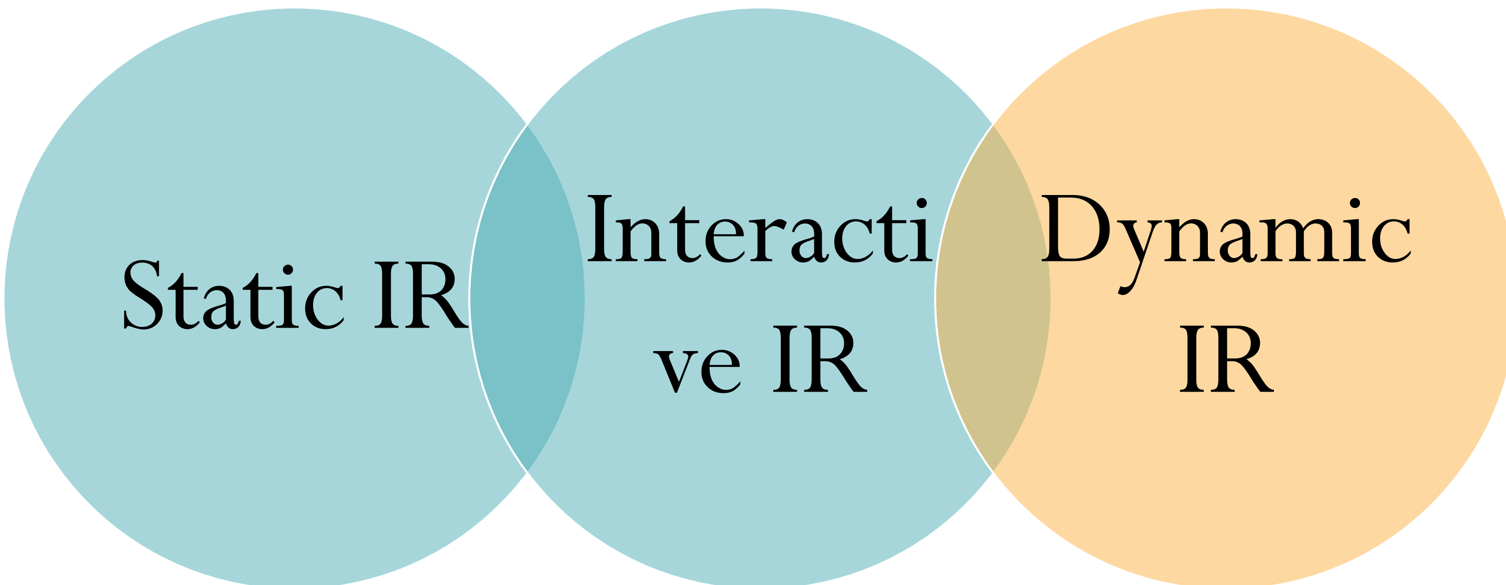

- Explore and exploit Feedback

Dynamic Information Retrieval

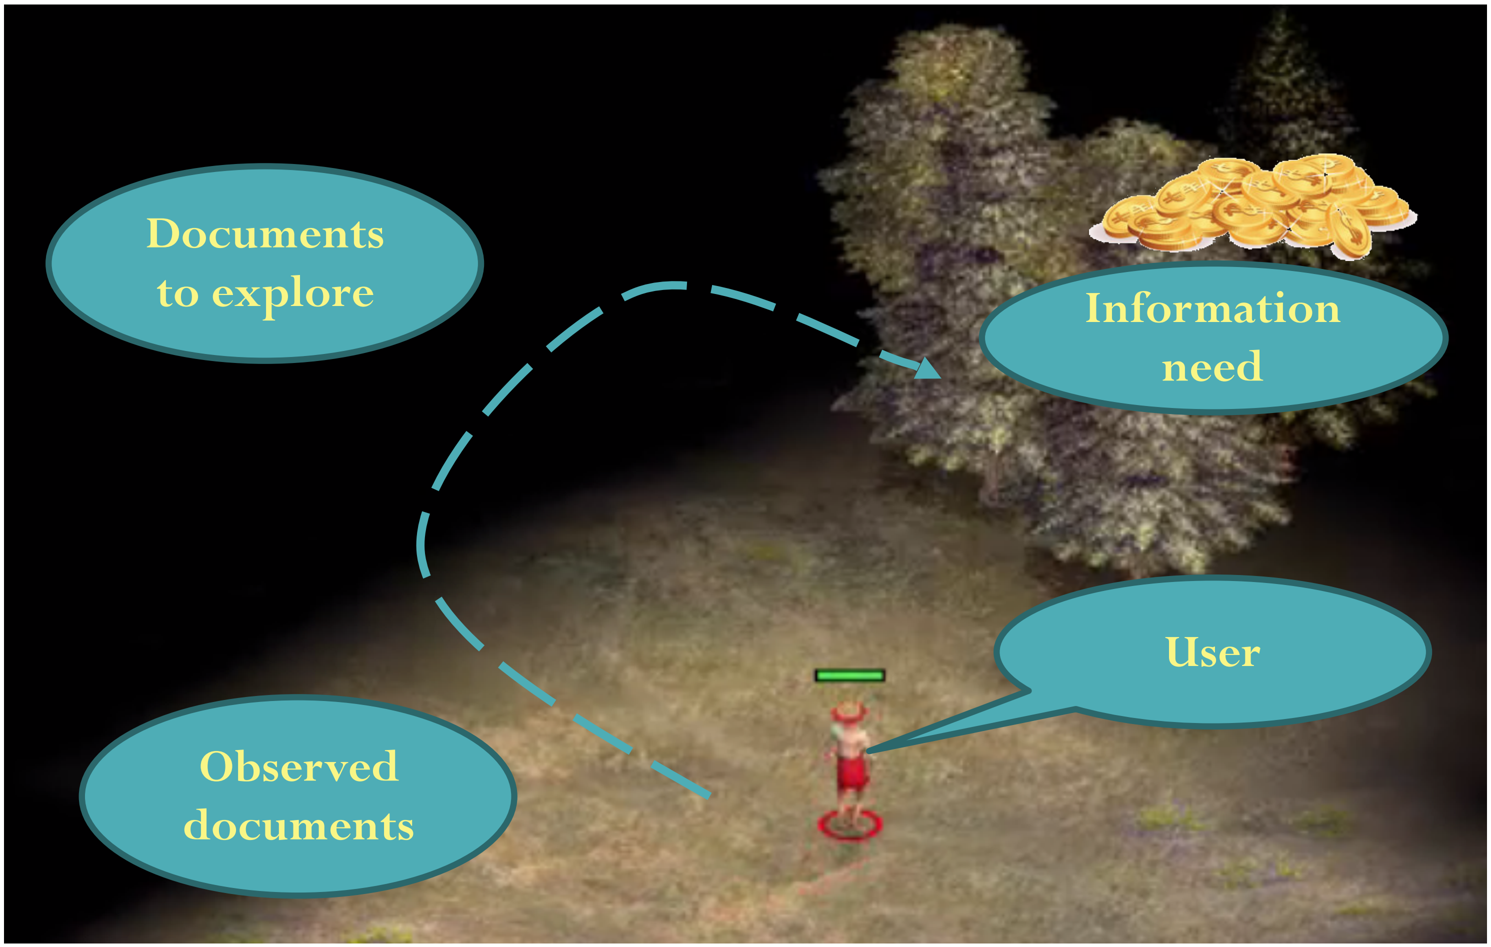

Devise a strategy for helping the user explore the information space in order to learn which documents are relevant and which aren't, and satisfy their information need.

Characteristics of Dynamic IR

- Rich interactions
  - Query formulation
  - Document clicks
  - Document examination
  - Eye movement
  - Mouse movements
  - etc.

[Luo et al., IRJ under revision 2014]

Characteristics of Dynamic IR

- Temporal dependency

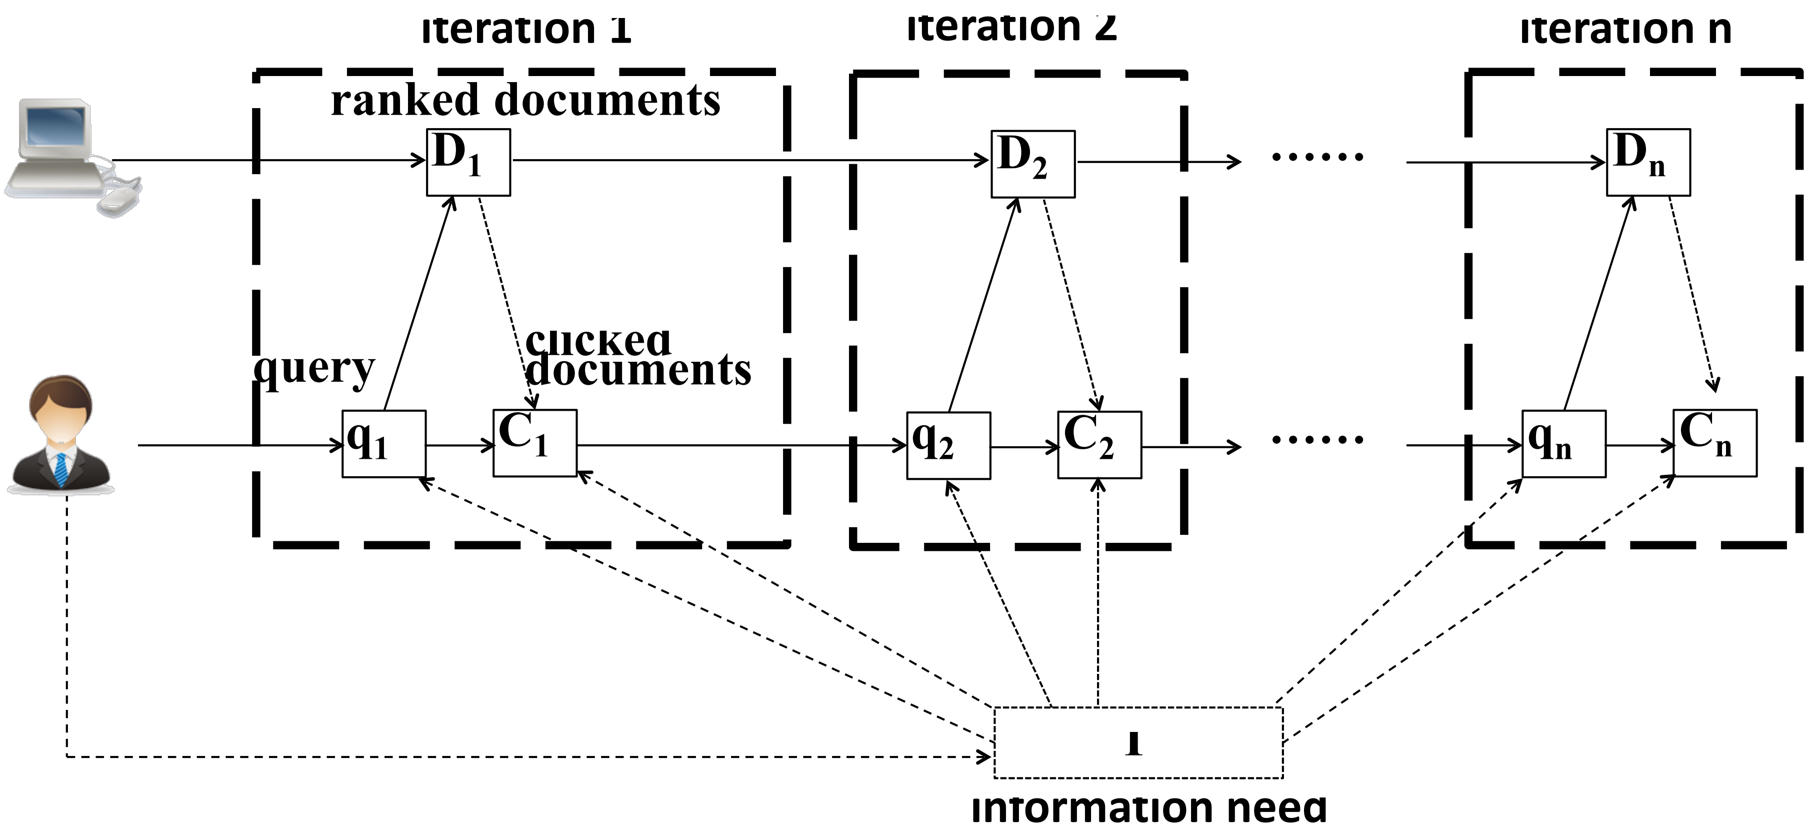

[Luo et al., IRJ under revision 2014]

Characteristics of Dynamic IR

- Overall goal
  - Optimize over all iterations for goal
  - IR metric or user satisfaction
  - Optimal policy

[Luo et al., IRJ under revision 2014]

Dynamic Information Retrieval

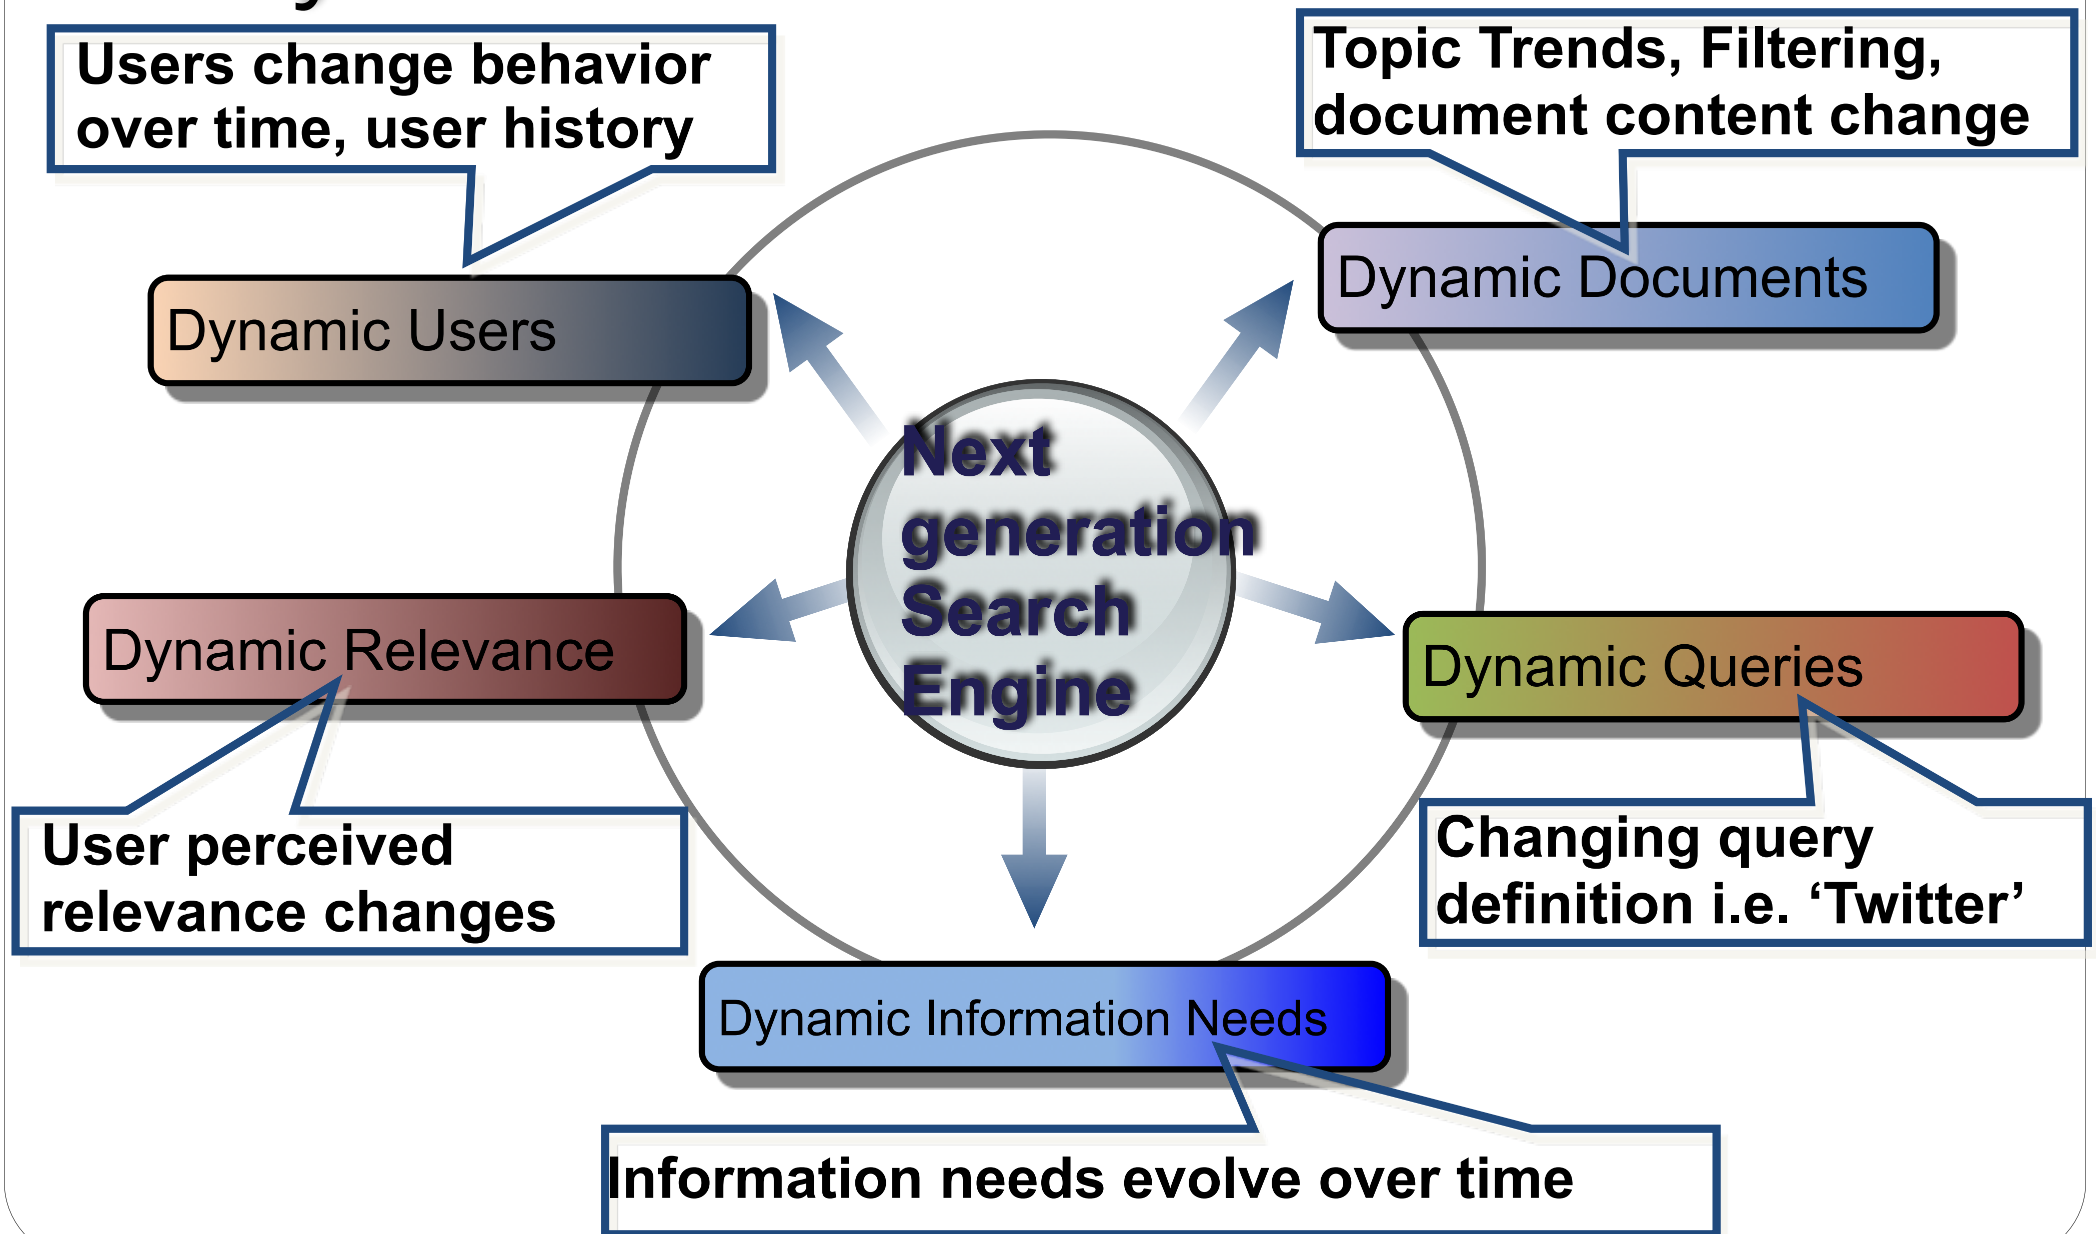

Trial and Error

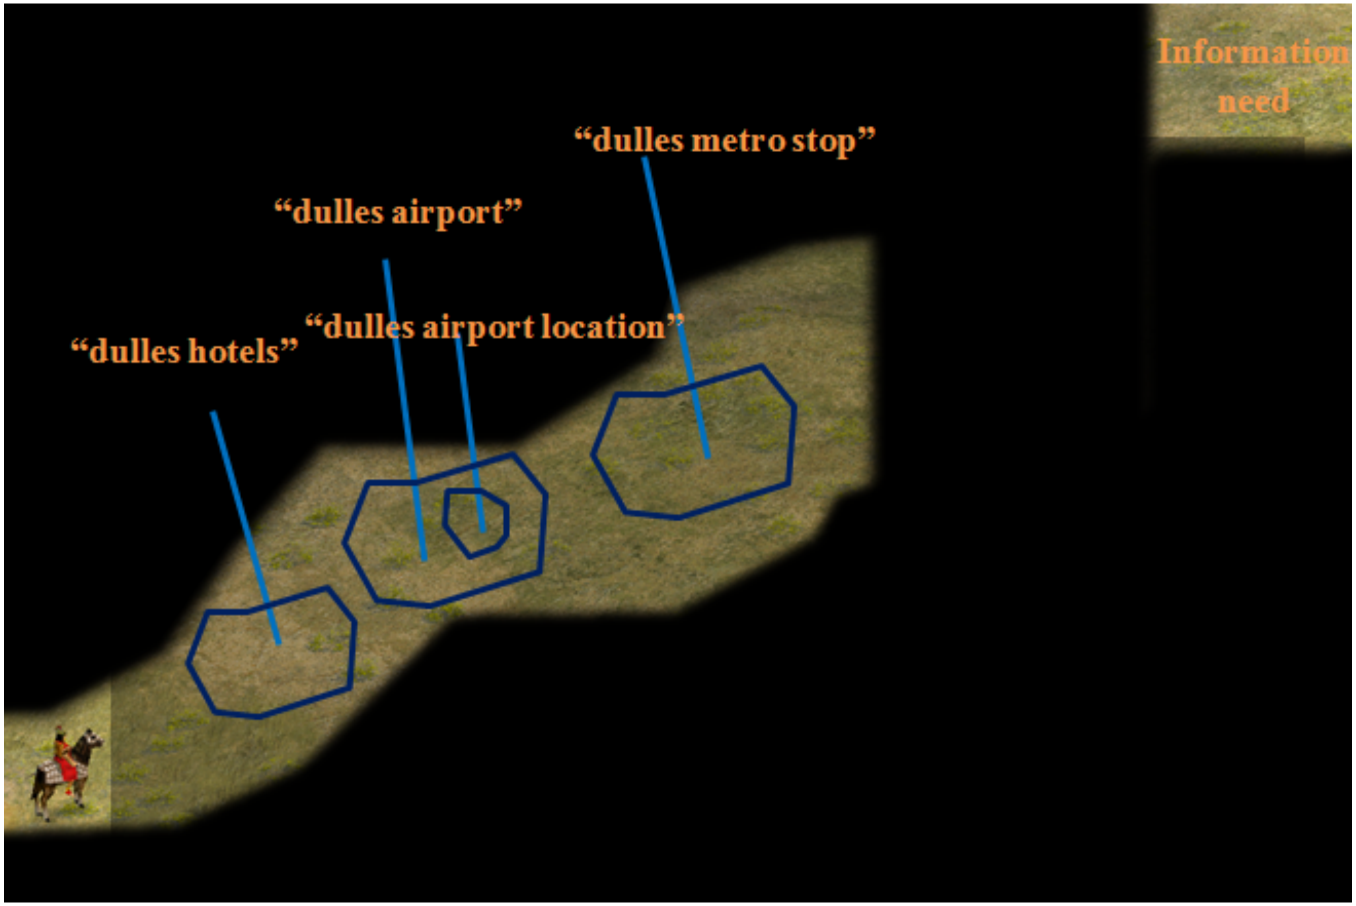

- q_1 – "dulles hotels"
- q_2 – "dulles airport"
- q_3 – "dulles airport location"
- q_4 – "dulles metrostop"

What is a Desirable Model for Dynamic IR

- Model **interactions**, which means it needs to have place holders for actions;
- Model **information need** hidden behind user queries and other interactions;
- Set up a **reward** mechanism to guide the entire search algorithm to adjust its retrieval strategies;

A model in Trial and Error setting will do!

- Represent Markov properties to handle the **temporal dependency.**

A Markov Model will do!

WE LOOK INTO THE FAMILY OF REINFORCEMENT LEARNING ALGORITHMS.

REINFORCEMENT LEARNING

- The 3rd type of Machine Learning Algorithms
- A computer program interacts with a dynamic environment in which it must perform a certain goal (such as driving a vehicle), without a teacher/the training data explicitly telling it whether it has come close to its goal.
- Many of them assume a Markov Process

  - Supervised machine learning
    - The computer program is presented with example inputs and their desired outputs, given by a “teacher" or the training data, and the goal is to learn a general rule that maps inputs to outputs.
  - Unsupervised machine learning
    - The computer program is given a bunch of data and must find patterns and relationships therein (such as clustering the data into groups)

Markov Process

- **Markov Property**[1] (the "memoryless" property)

  for a system, its next state depends on its current state.

$$\Pr(S_{i+1} \mid S_i, \ldots, S_0) = \Pr(S_{i+1} \mid S_i)$$

$$S_0 \rightarrow S_1 \cdots\cdots S_i \rightarrow S_{i+1} \cdots\cdots$$

e.g. 

- **Markov Process**

  a stochastic process with Markov property.

[1]A. A. Markov, '06

Family of Markov Models

- **Markov Chain**
- Hidden Markov Model
- Markov Decision Process
- Partially Observable Markov Decision Process
- Multi-armed Bandit

(S, M, A, R, γ)

Markov Decision Process

- MDP extends MC with actions and rewards[1]

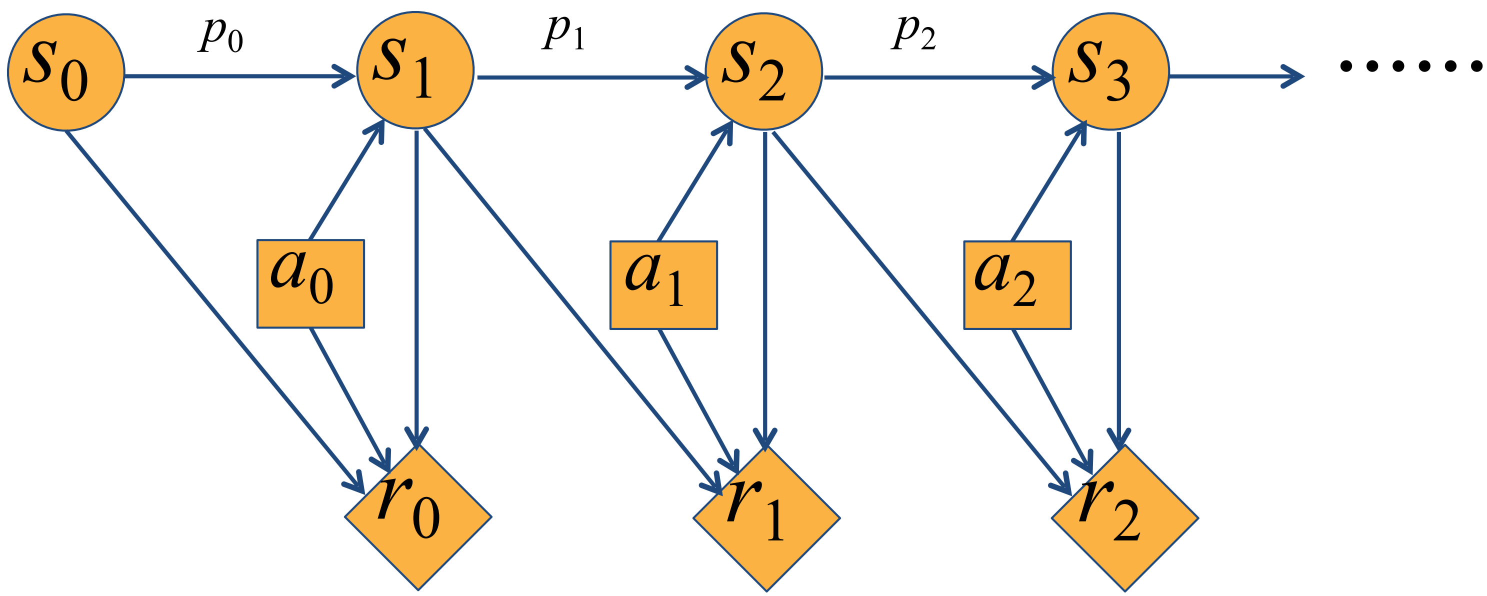

s_i – state a_i – action r_i – reward

p_i – transition probability

[1]R. Bellman, '57

Definition of MDP

- A tuple (S, M, A, R, γ)
  - S : state space
  - M: transition matrix

    $M_a(s, s') = P(s'|s, a)$
  - A: action space
  - R: reward function

    $R(s,a)$ = immediate reward taking action a at state s
  - γ: discount factor, $0 < \gamma \leq 1$
- policy π

  $\pi(s)$ = the action taken at state s
- Goal is to find an optimal policy π^* maximizing the expected total rewards.

Policy

Policy: $\pi(s) = \boldsymbol{a}$

According to which, select an action a at state s.

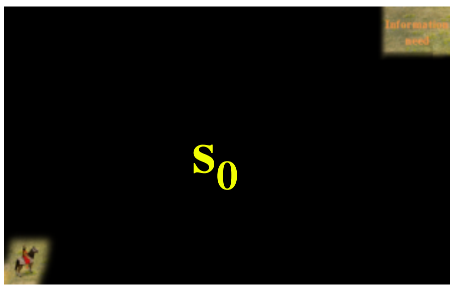

$\pi(s_0)$ =move right and up

$\pi(s_1)$ =move right and up

$\pi(s_2)$ = move right

[Slide altered from Carlos Guestrin's ML lecture]

Value of Policy

Value: $V^{\pi}(s)$ → Expected long-term reward starting from s

$$V^{\pi}(\mathbf{s_0}) = \mathbf{E}[R(\mathbf{s}_0) + \gamma\, R(\mathbf{s}_1) + \gamma^2\, R(\mathbf{s}_2) + \gamma^3\, R(\mathbf{s}_3) + \gamma^4\, R(\mathbf{s}_4) + \cdots]$$

Future rewards discounted by $\gamma \in [0,1)$

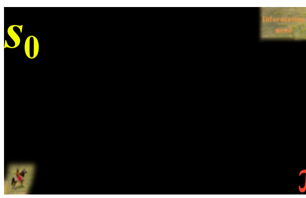

$\pi(s_0)$

$R(s_0)$

Start from s_0

[Slide altered from Carlos Guestrin's ML lecture]

Value of Policy

Value: $V^{\pi}(s)$

Expected long-term reward starting from s

$$V^{\pi}(\mathbf{s}_0) = \mathbf{E}[R(\mathbf{s}_0) + \gamma\, R(\mathbf{s}_1) + \gamma^2\, R(\mathbf{s}_2) + \gamma^3\, R(\mathbf{s}_3) + \gamma^4\, R(\mathbf{s}_4) + \cdots]$$

Future rewards discounted by $\gamma \in [0,1)$

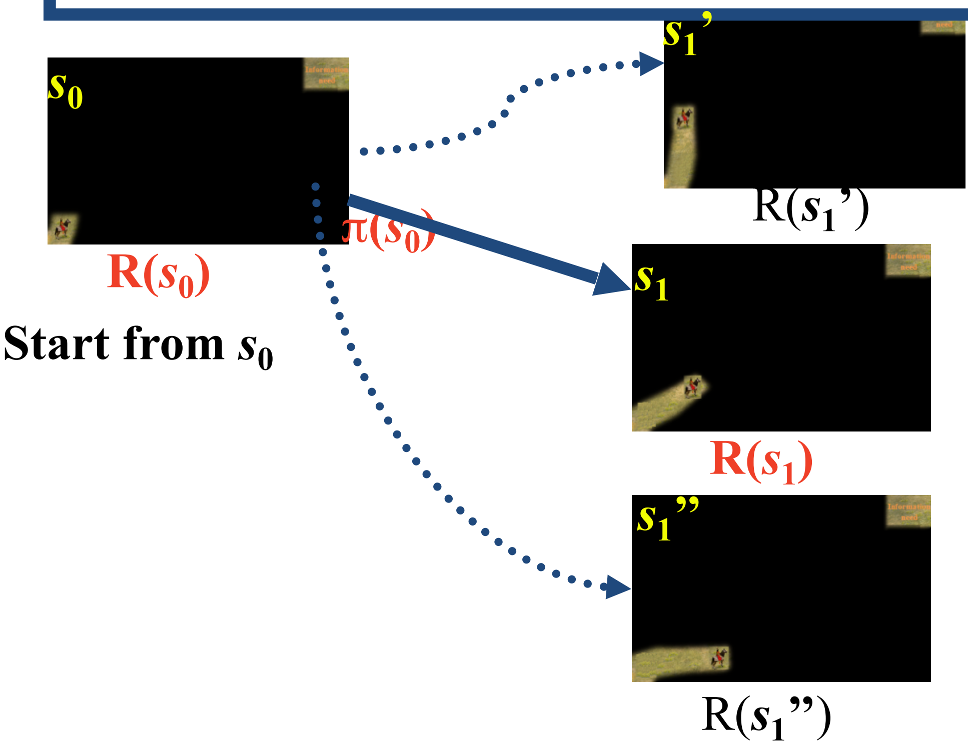

[Slide altered from Carlos Guestrin's ML lecture]

Value of Policy

Value: $V^{\pi}(s)$

Expected long-term reward starting from s

$$V^{\pi}(\mathbf{s}_0) = \mathbf{E}[R(\mathbf{s}_0) + \gamma\, R(\mathbf{s}_1) + \gamma^2\, R(\mathbf{s}_2) + \gamma^3\, R(\mathbf{s}_3) + \gamma^4\, R(\mathbf{s}_4) + \cdots]$$

Future rewards discounted by $\gamma \in [0,1)$

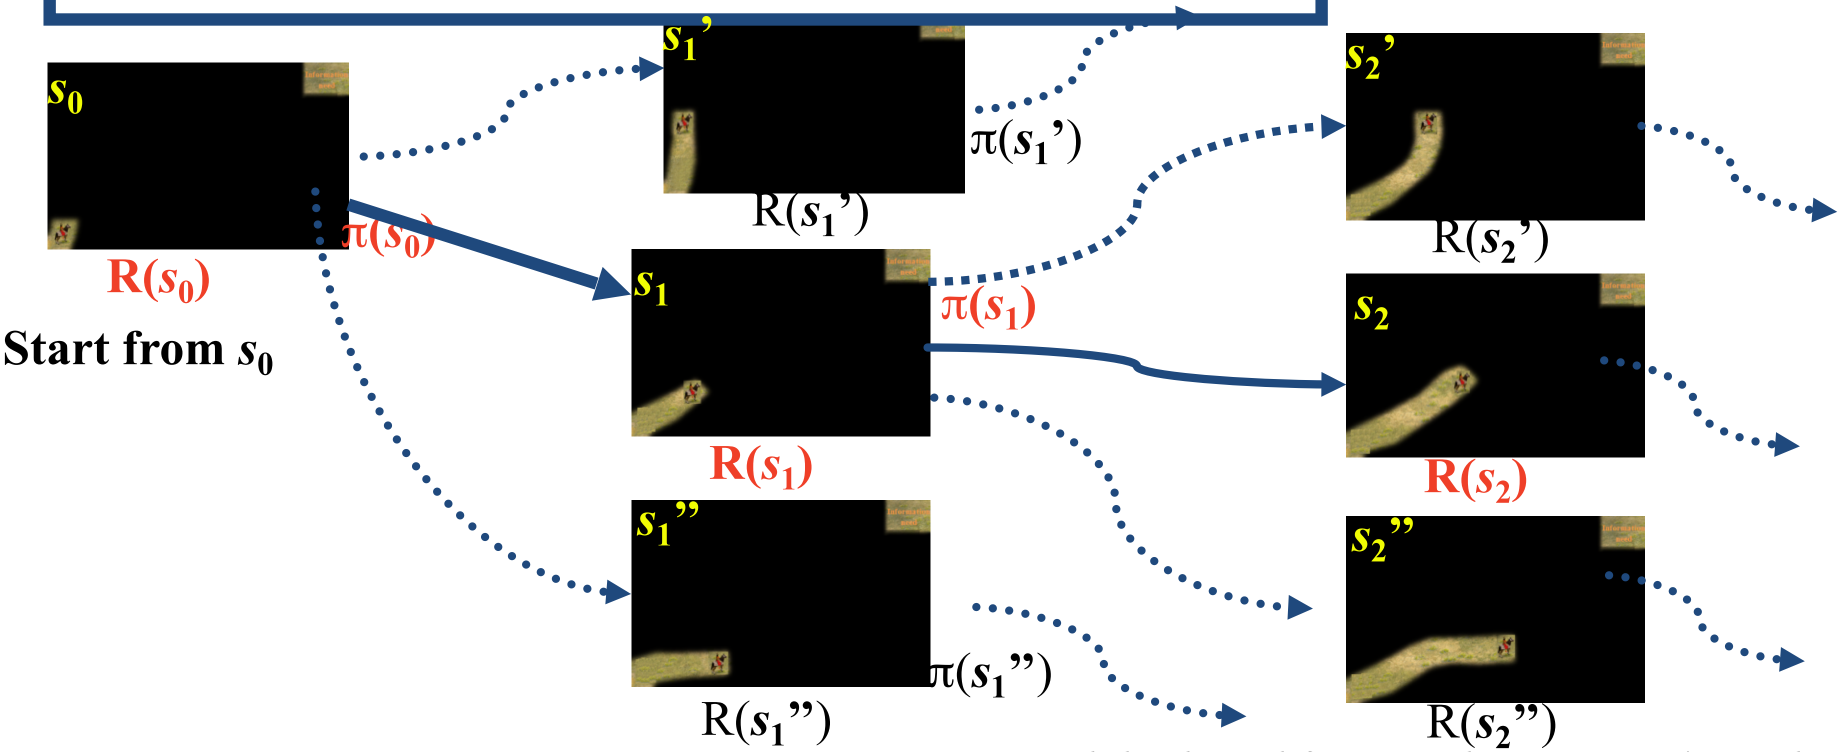

[Slide altered from Carlos Guestrin's ML lecture]

Computing the value of a policy

Value function

$$
\begin{aligned}
V^{\pi}(\mathbf{s_0}) &= E^{\pi}[R(s_0, a) + \gamma R(s_1, a) + \gamma^2 R(s_2, a) + \gamma^3 R(s_3, a) + \cdots] \\
&= E^{\pi}[R(s_0, a) + \gamma \textstyle\sum_{t=1}^{\infty} \gamma^{t-1} R(s_t, a)] \\
&= R(s_0, a) + \gamma E^{\pi}[\textstyle\sum_{t=1}^{\infty} \gamma^{t-1} R(s_t, a)] \\
&= R(s_0, a) + \gamma \textstyle\sum_{s'} M_{\pi(s)}(s, s')\, V^{\pi}(s')
\end{aligned}
$$

The current state

A possible next state

Optimality — Bellman Equation

- The Bellman equation[1] to MDP is a recursive definition of the optimal value function $V^*(.)$

state-value function $$V^*(s) = \max_a \left[R(s,a) + \gamma \sum_{s'} M_a(s,s') V^*(s') \right]$$

- Optimal Policy

$$\pi^*(s) = \arg\max_a \left[R(s,a) + \gamma \sum_{s'} M_a(s,s') V^*(s') \right]$$

[1]R. Bellman, '57

Optimality — Bellman Equation

- The Bellman equation can be rewritten as

$$V^*(s) = \max_{\mathrm{a}}[Q(s,a)]$$

Relationship between V and Q

action-value function

$$Q(s,a) = R(s,a) + \gamma \sum_{s'} M_a(s,s')V^*(s')$$

- Optimal Policy

$$\pi^*(\mathrm{s}) = \arg\max_{a} Q(s,a)$$

MDP algorithms

Solve Bellman equation → Optimal value $V^*(\mathbf{s})$ → Optimal policy $\pi^*(\mathbf{s})$

- Value Iteration
- Policy Iteration
- Modified Policy Iteration
- Prioritized Sweeping
- Temporal Difference (TD) Learning
- Q-Learning

Model-based approaches

Model free approaches

[Bellman, '57, Howard, '60, Puterman and Shin, '78, Singh & Sutton, '96, Sutton & Barto, '98, Richard Sutton, '88, Watkins, '92]

[Slide altered from Carlos Guestrin's ML lecture]

REINFORCEMENT LEARNING FOR DYNAMIC IR

Apply Reinforcement Learning to the Search Problem

- We can model IR systems using a Markov Decision Process
- Is there a temporal component?
- States – What changes with each time step?
- Actions – How does your system change the state?
- Rewards – How do you measure feedback or effectiveness in your problem at each time step?
- Transition Probability – Can you determine this?
  - If not, then model free approach is more suitable

Session Search

TREC 2012 Session 6

1. pocono mountains pennsylvania
2. pocono mountains pennsylvania hotels
3. pocono mountains pennsylvania things to do
4. pocono mountains pennsylvania hotels
5. pocono mountains camelbeach
6. pocono mountains camelbeach hotel
7. pocono mountains chateau resort
8. pocono mountains chateau resort attractions
9. pocono mountains chateau resort getting to
10. chateau resort getting to
11. pocono mountains chateau resort directions

Information needs:

You are planning a winter vacation to the Pocono Mountains region in Pennsylvania in the US. Where will you stay? What will you do while there? How will you get there?

In a session, queries change constantly

Apply MDP to Session Search

- States – user's relevance judgement
- Action – new query, new returned document lists
- Reward – relevant information gained

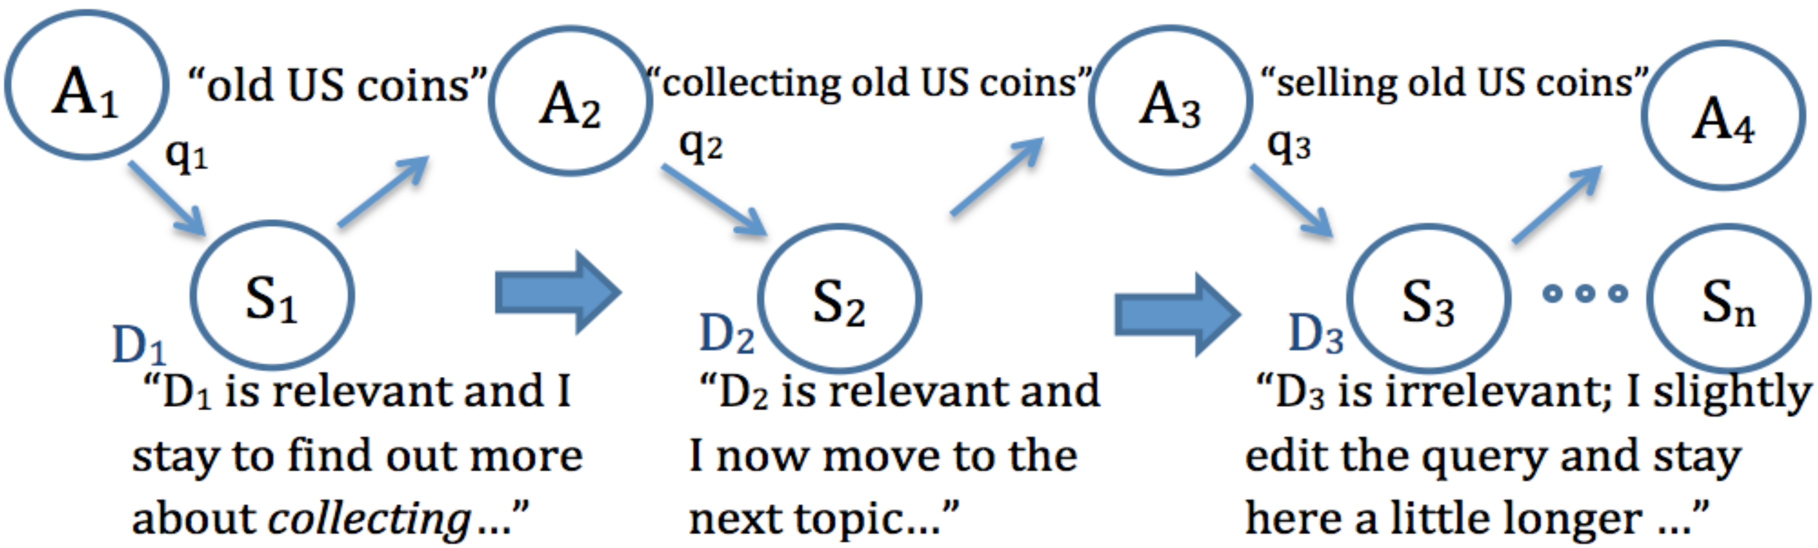

Settings of the Session MDP

[Guan, Zhang and Yang SIGIR 2013]

- States: Queries
- Environments: Search results
- Actions:
  - User actions:
    - Add/remove/ unchange the query terms
      - Nicely correspond to our definition of query change
  - Search Engine actions:
    - Increase/ decrease /remain term weights

Search Engine Agent's Actions

	$\in D_{i-1}$	action	Example
q_{theme}	Y	increase	"pocono mountain" in s6
	N	increase	"france world cup 98 reaction" in s28, france world cup 98 reaction stock market→ france world cup 98 reaction
$+\Delta q$	Y	decrease	'policy' in s37, Merck lobbyists → Merck lobbyists US policy
	N	increase	'US' in s37, Merck lobbyists → Merck lobbyists US policy
$-\Delta q$	Y	decrease	'reaction' in s28, france world cup 98 reaction → france world cup 98
	N	No change	'legislation' in s32, bollywood legislation →bollywood law

[Guan, Zhang and Yang SIGIR 2013]

Query Change retrieval Model (QCM) [Guan, Zhang and Yang SIGIR 2013]

- Bellman Equation gives the optimal value for an MDP:

$$V^*(s) = \max_a R(s,a) + \gamma \sum_{s'} P(s' \mid s,a) V^*(s')$$

- The reward function is used as the document relevance score function and is tweaked backwards from Bellman equation:

$$\text{Score}(q_i, d) = P(q_i \mid d) + \gamma \sum_a P(q_i \mid q_{i-1}, D_{i-1}, a) \max_{D_{i-1}} P(q_{i-1} \mid D_{i-1})$$

Document relevant score

Current reward/ relevance score

Query Transition model

Maximum past relevance

Calculating the Transition Model

- According to Query Change and Search Engine Actions

[Guan, Zhang and Yang SIGIR 2013]

Current reward/ relevance score

Increase weights for theme terms

$$
\begin{aligned}
\mathrm{Score}(q_i, d) = {} & \log P(q_i \mid d) + \alpha \sum_{t \in qtheme} [1 - P(t \mid d^*_{i-1})] \log P(t \mid d) \\
& - \beta \sum_{\substack{t \in +\Delta q \\ t \in d^*_{i-1}}} P(t \mid d^*_{i-1}) \log P(t \mid d) + \varepsilon \sum_{\substack{t \in +\Delta q \\ t \notin d^*_{i-1}}} idf(t) \log P(t \mid d) \\
& - \delta \sum_{t \in -\Delta q} P(t \mid d^*_{i-1}) \log P(t \mid d)
\end{aligned}
$$

Increase weights for novel added terms

Decrease weights for old added terms

Decrease weights for removed terms

Maximizing the Reward Function

- $$\max_{D_{i-1}} P(q_{i-1} \mid D_{i-1})$$
  - Generate a maximum rewarded document denoted as d^*_{i-1}, from D_{i-1}
    - That is the document(s) most relevant to q_{i-1}
  - The relevance score can be calculated as

$$P(q_{i-1}|d_{i-1}) = 1 - \prod_{t \in q_{i-1}} \{1 - P(t|d_{i-1})\}$$

$$P(t|d_{i-1}) = \frac{\#(t, d_{i-1})}{|d_{i-1}|}$$

  - From several options, we choose to only use the document with top relevance

[Guan, Zhang and Yang SIGIR 2013]

Scoring the Entire Session

- The overall relevance score for a session of queries is aggregated recursively :

$$\mathrm{Score}_{\mathrm{session}}(q_n, d) = \mathrm{Score}(q_n, d) + \gamma\,\mathrm{Score}_{\mathrm{session}}(q_{n-1}, d)$$

$$= \mathrm{Score}(q_n, d) + \gamma[\mathrm{Score}(q_{n-1}, d) + \gamma\,\mathrm{Score}_{\mathrm{session}}(q_{n-2}, d)]$$

$$= \sum_{i=1}^{n} \gamma^{n-i}\,\mathrm{Score}(q_i, d)$$

[Guan, Zhang and Yang SIGIR 2013]

Experiments

- TREC 2011-2012 query sets, datasets

	2011	2012		2011	2012
#topics	62	48	#queries/session	3.68	3.03
#sessions	76	98	#sessions/topic	1.23	2.04
#queries	280	297	#pages judged	19,413	17,861
#dups	16	5	#sessions w/o rel. docs	2	4

- ClubWeb09 Category B

Search Accuracy (TREC 2012)

- nDCG@10 (official metric used in TREC)

Approach	nDCG@10	%chg	MAP	%chg
Lemur	0.2474	-21.54%	0.1274	-18.28%
TREC'12 median	0.2608	-17.29%	0.1440	-7.63%
Our TREC'12 submission	0.3021	–4.19%	0.1490	-4.43%
TREC'12 best	0.3221	0.00%	**0.1559**	**0.00%**
QCM	0.3353	4.10%†	0.1529	-1.92%
QCM+Dup	**0.3368**	**4.56%†**	0.1537	-1.41%

Search Accuracy for Different Session Types

- TREC 2012 Sessions are classified into:
  - Product: Factual / Intellectual
  - Goal quality: Specific / Amorphous

	Intellectual	%chg	Amorphous	%chg	Specific	%chg	Factual	%chg
TREC best	**0.3369**	0.00%	**0.3495**	0.00%	0.3007	0.00%	0.3138	0.00%
Nugget	**0.3305**	-1.90%	**0.3397**	-2.80%	0.2736	-9.01%	0.2871	-8.51%
QCM	**0.3870**	14.87%	**0.3689**	5.55%	0.3091	2.79%	0.3066	-2.29%
QCM+DUP	**0.3900**	15.76%	**0.3692**	5.64%	0.3114	3.56%	0.3072	-2.10%

- **Better handle sessions that demonstrate evolution and exploration Because QCM treats a session as a continuous process by studying changes among query transitions and modeling the dynamics**

SUMMARY

- **IR Essentials**
- **Vector Space Model**
- **Dynamic Search**

Grace Hui Yang
huiyang@cs.georgetown.edu