Intelligent Computing in Smart Reality

2012. 05.

Sang-goo Lee
School of Computer Science & Engineering
SNU-SAMSUNG Joint Center for Intelligent Computing
Seoul National University, Seoul, Korea
sglee@snu.ac.kr

Gartner’s IT Predictions

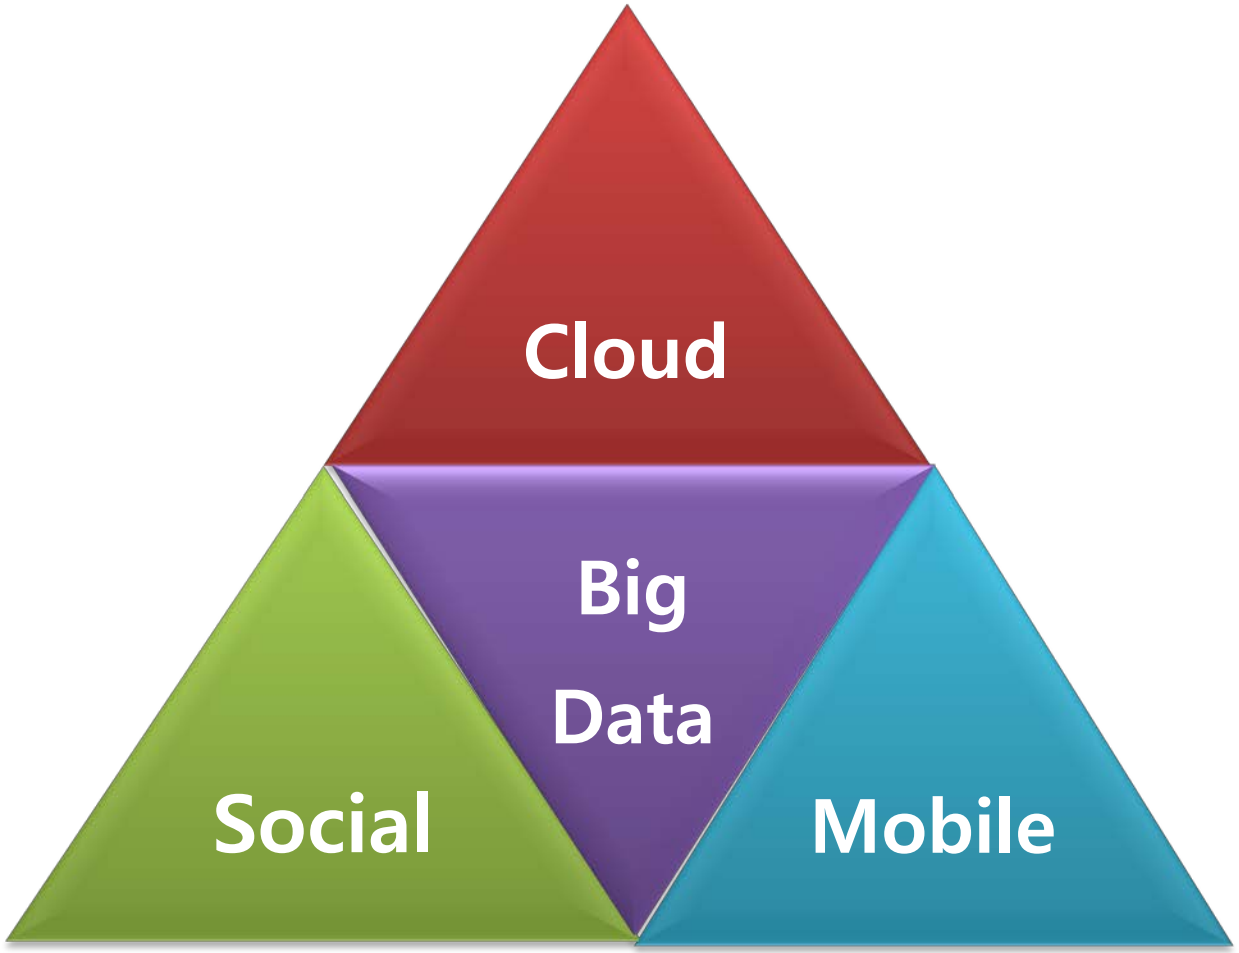

[Gartner IT Expo, 2011]

CENTER FOR INTELLIGENT COMPUTING

Cloud Computing

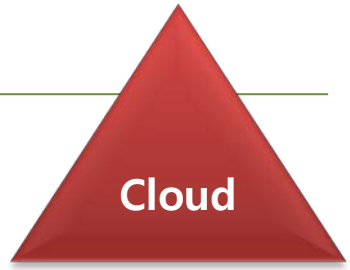

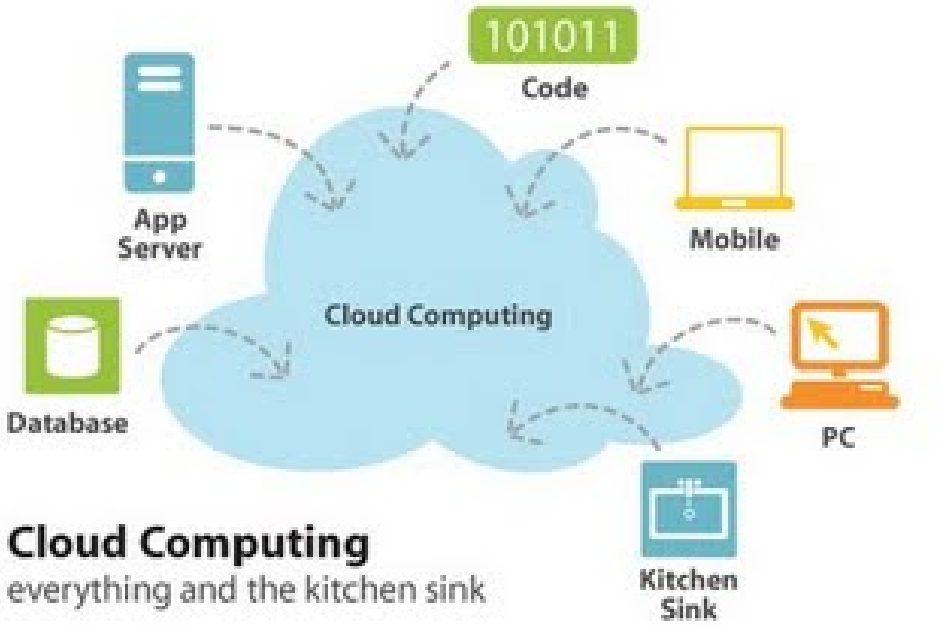

- Web as a Programming Platform
- Elastic Computing: SaaS, PaaS, IaaS, …
- **Supercomputing for everyone & everything!**

Smart Apps & Gadgets

Cloud

CENTER FOR INTELLIGENT COMPUTING

Data Explosion

Big Data

- 1,800,000,000,000,000,000,000 Bytes (1.8ZB)
  - Digital data produced in 2011 [EMC, 2012]
  - Doubles every 1.5 years => X 44 in 2020

The Digital Universe 2009-2020

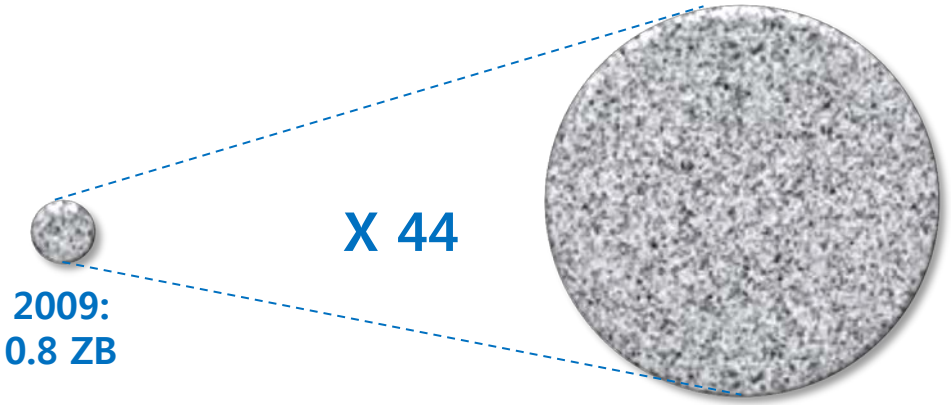

CENTER FOR INTELLIGENT COMPUTING

Data = Intelligence

Big Data

Large Scale Bio-Network Inference

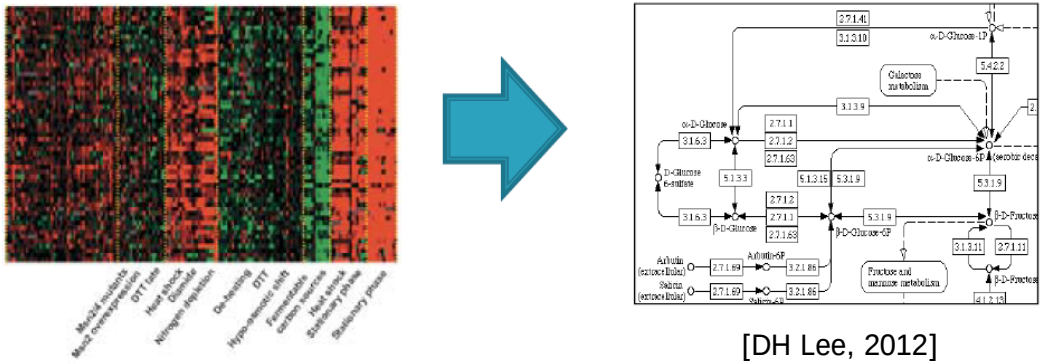

[DH Lee, 2012]

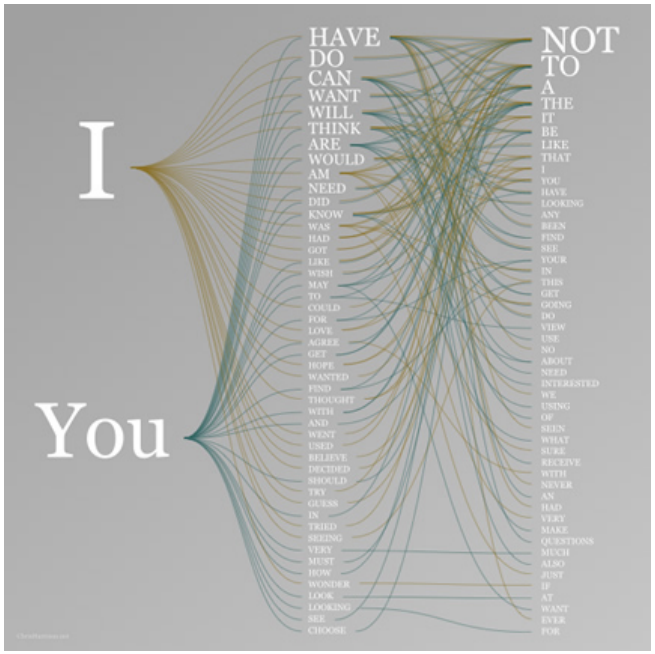

[Wikipedia]

sentence (문장):	95,119,665,584
unigrams (홑 단어):	13,588,391
bigrams (두 단어 조합):	314,843,401
trigrams (세 단어 조합):	977,069,902
4-grams (네 단어 조합)	1,313,818,354
5-grams: (다섯 단어 조합):	1,176,470,663

CENTER FOR INTELLIGENT COMPUTING

Robots => Cloud-bots

Social


Social


Understanding where and when an item was produced
Understanding where you are now
STORE
Automatic navigation in foreign languages for foreigners
Safe route navigation for the physically challenged

iPhone !!!

Mobile

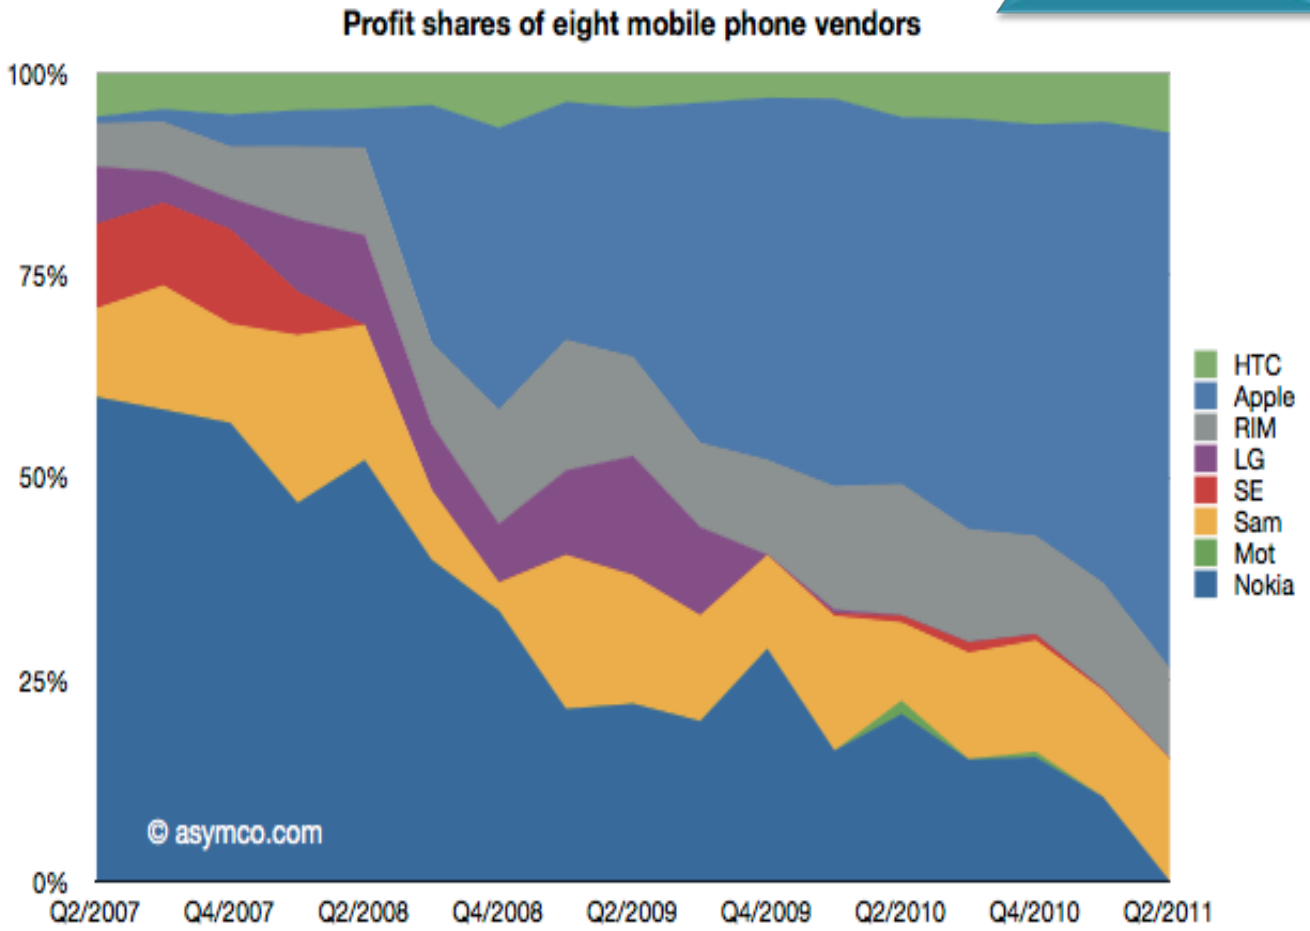

Smart Everything!

Mobile

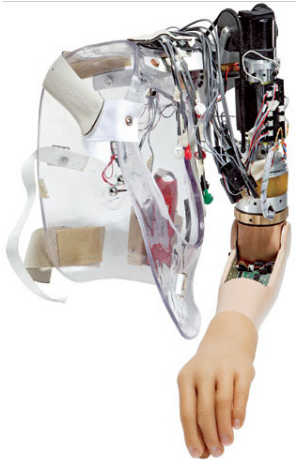

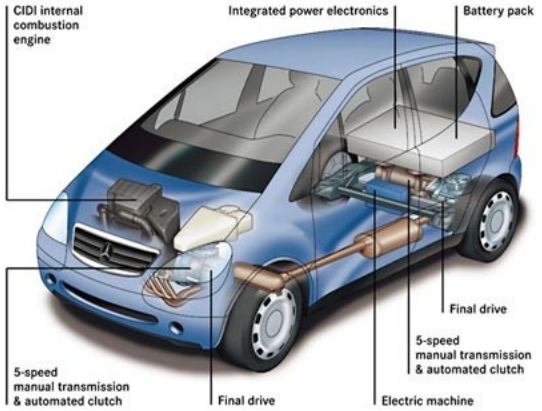

Sensors in Smartphones

Mobile

- Multi-touch
- Accelerometer (aka motion sensor)
  - Acceleration in 3-axis, gravity
- Proximity sensor (InfraRed)
- Ambient light sensor
- A-GPS
- Camera
- Microphone
  - E.g. stethoscope
- Digital/electronic Compass
  - Magnetic sensor
- Three-axis Gyroscope
- Water sensor

More sensors to come!

- Temperature, humidity
- Barometer
- Chemical sensors like smell
- Biomedical sensors
  - Apple has patents for an earbud
    - Blood oxygen, heat flux, body temperature, heart rate,…

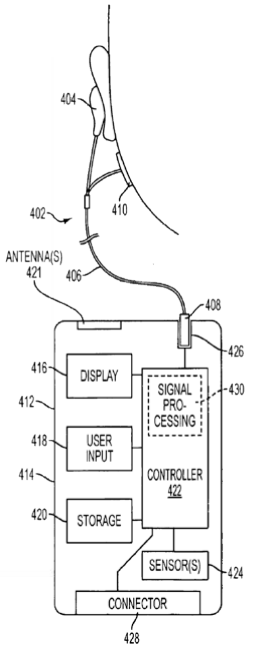

[T. Kwon, 2012]

Sensors & Actuators

Mobile

- Open-source physical computing platform
  - sense and control more of the physical world
  - board + SW

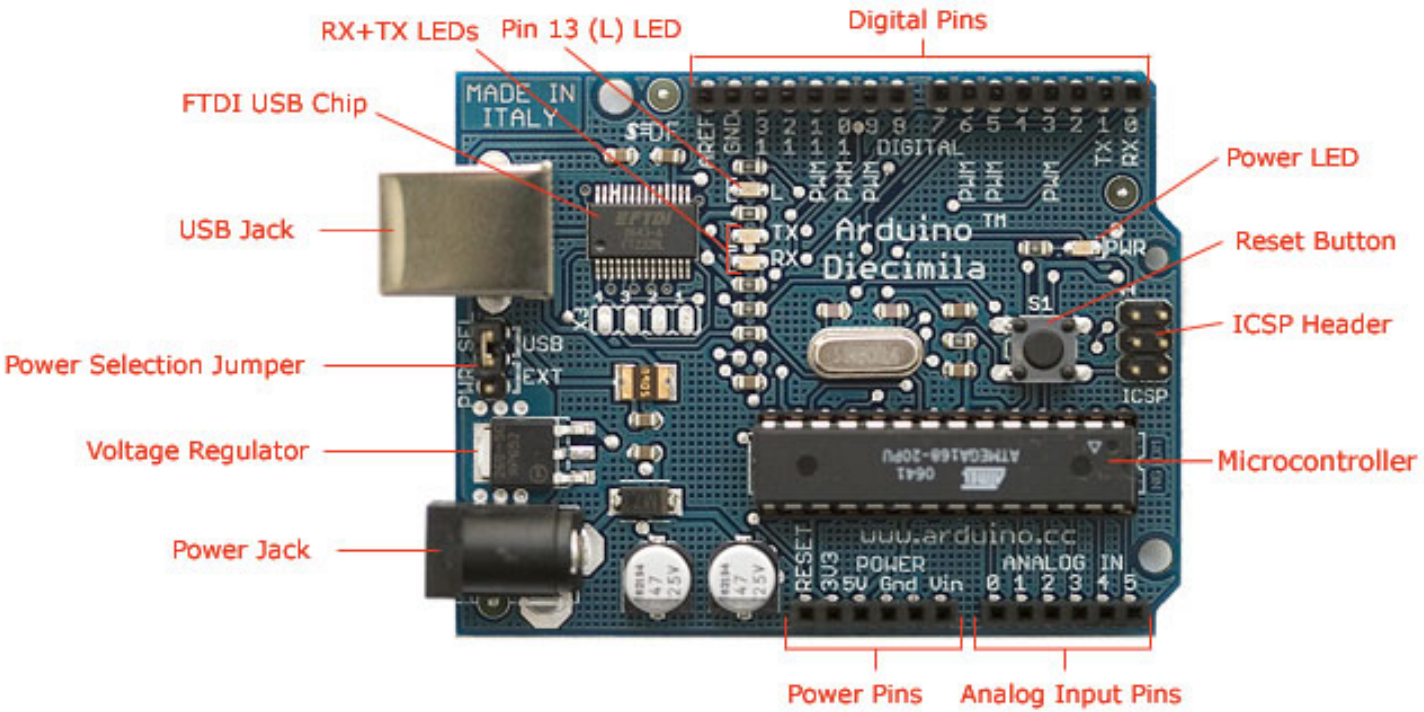

Photograph by SparkFun Electronics. Used under the Creative Commons Attribution

ARDUINO

~~15~~50,000,000,000 Smart Devices

Mobile

- *Programmable (프로그램 작동이 가능하고)*
- *Internet connected (인터넷 접속된)*

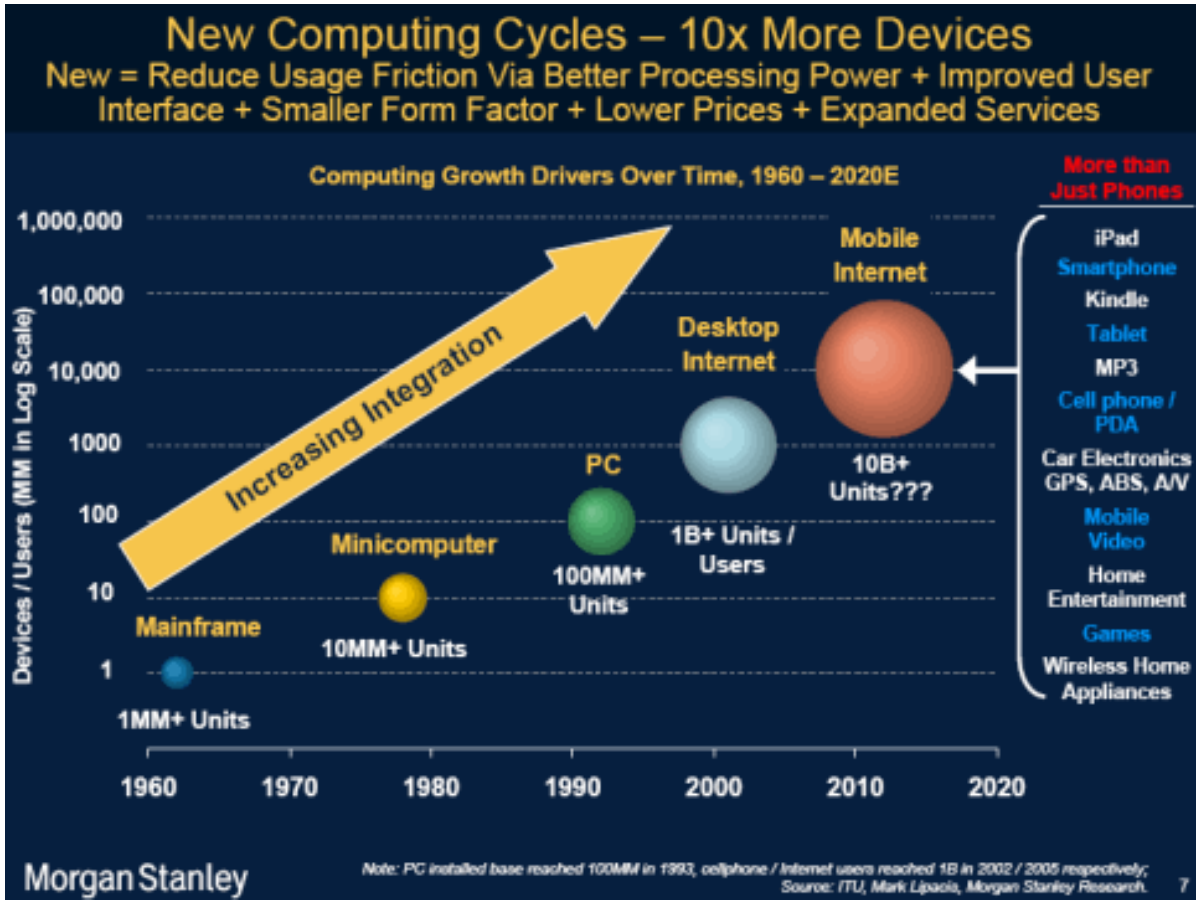

SMART REALITY

… a world where *reality and cyberspace are tightly integrated* by billions of *sensing and control devices* that are capable of autonomous computation and communication, thereby, ***empowering the reality with computational intelligence***.

Smart Reality Service Scenario

- Infinitely many choices of contents, devices, services
- Information overload, device overload, service overload

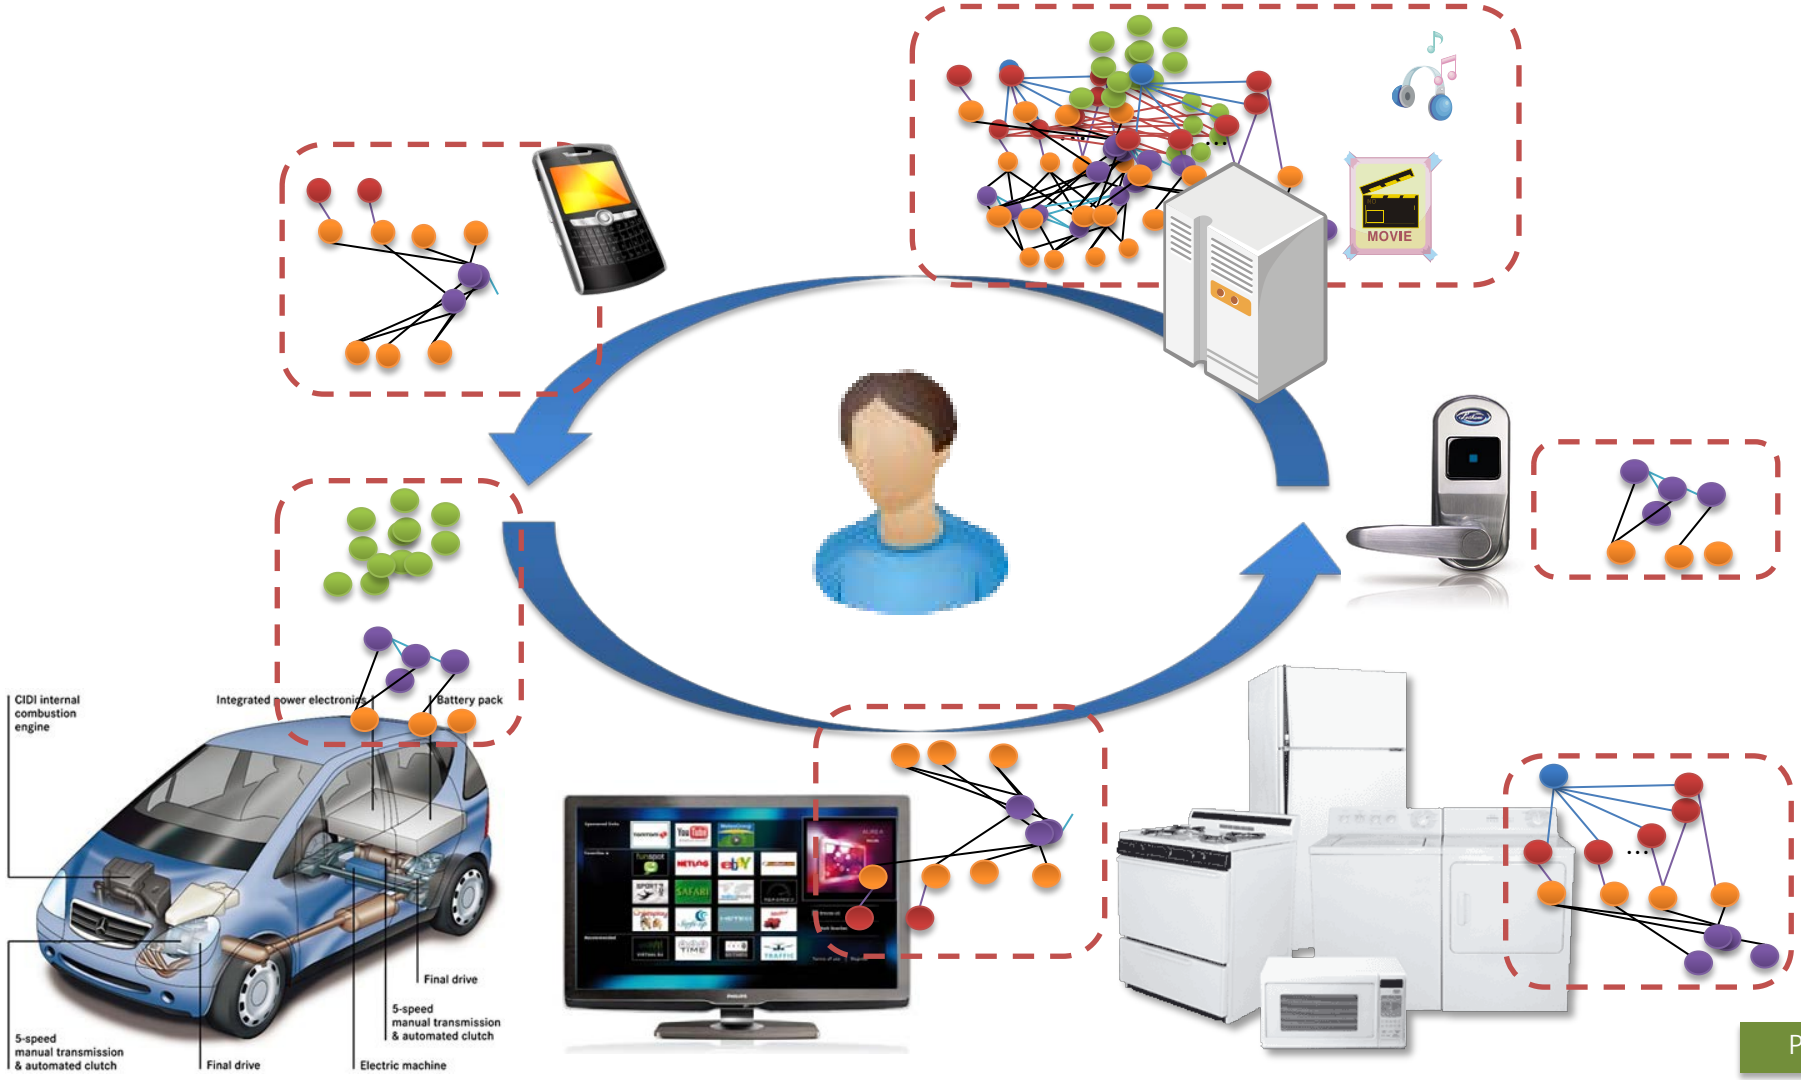

Context-Awareness in Smart Reality

- Ability to provide the right service through the right device at the right situation for the right person is essential in a successful implementation of smart reality.
- Challenges
  - Infinitely many contexts => a Long Tail problem
    - Must deal with *abnormal, unexpected, extraordinary* situations as well as the normal, expected, ordinary ones
  - Data integration and interoperability
  - Privacy!
  - New
    - No metrics, no testbeds, little or no logs available

CENTER FOR INTELLIGENT COMPUTING

What is *Context-Awareness*?

- Intelligent Office
  - Sense
  - Recognize
  - Decide
  - Action

- Augmented Reality
  - Sense
  - Recognize
  - Decide
  - Action

Recommender Systems

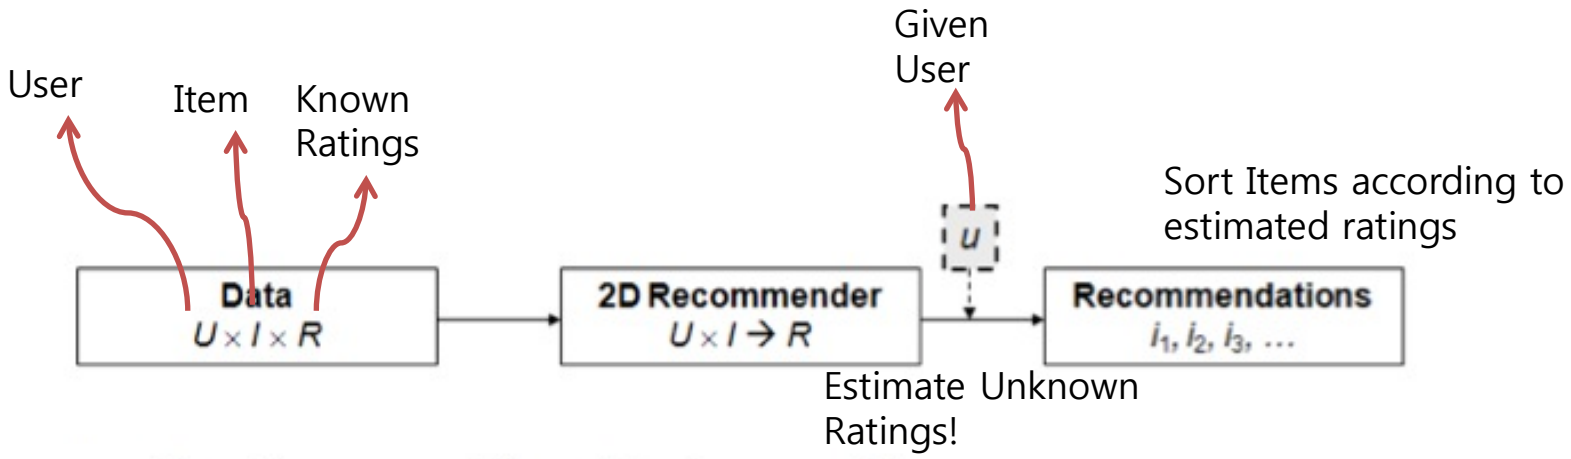

General components of the traditional recommendation process.

Content-based

- Pandora
  - Creat
    - Th
    - O
  - Music
    - m

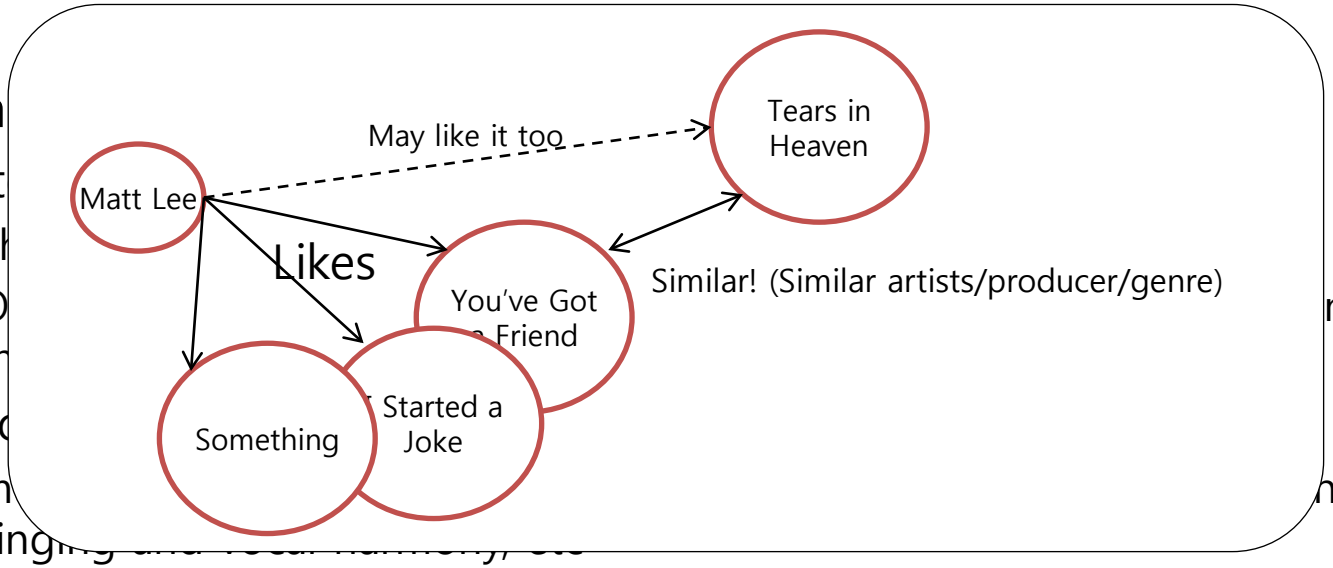

nalyzed

nt, lyrics, sing

Collaborative Filtering

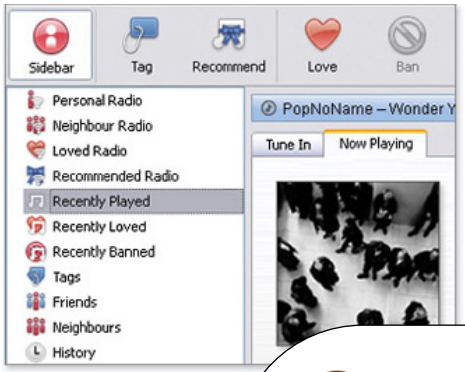

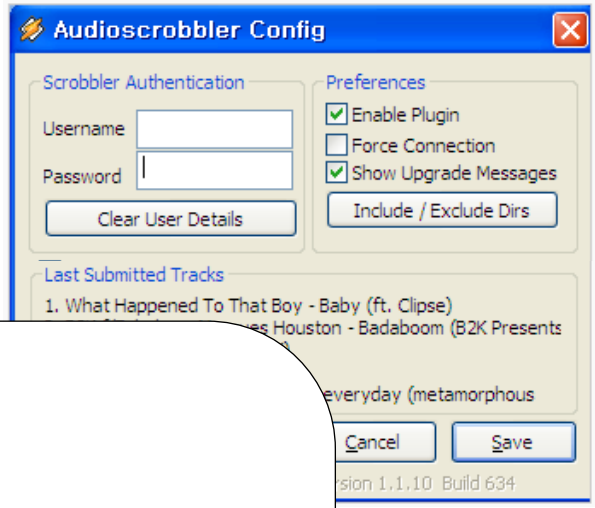

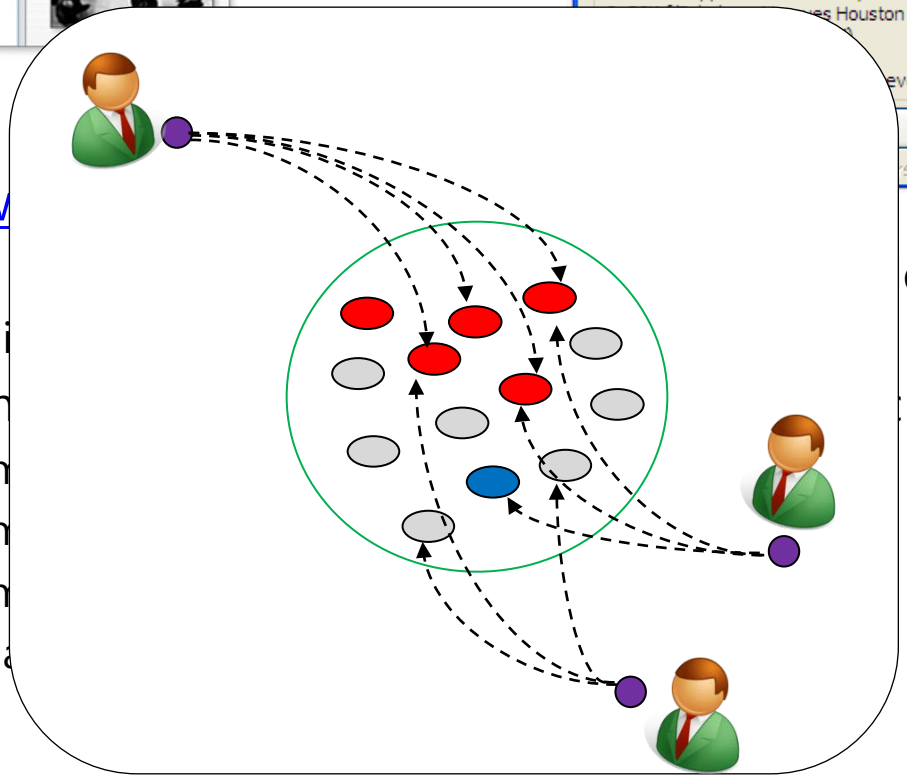

- Last.fm (www
  - Automatic
    - Scrobbli
  - Recommen
    - Recomm
    - Recomm
    - Recomm
    - Similar

e

to Last.fm

profiles

Modeling Context for Recommendation (Multidimensional)

$$R: User \times Season \times VacationLocation \rightarrow Rating$$

$$R: User \times Time \times Song \rightarrow Rating$$

$$\ldots$$

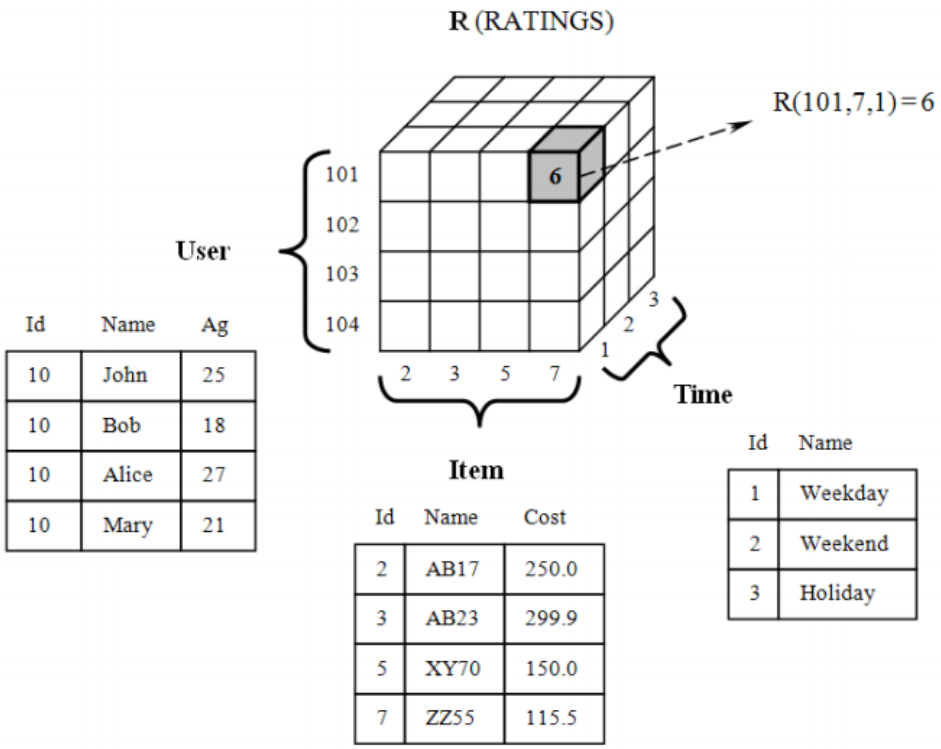

Id	Name	Ag
10	John	25
10	Bob	18
10	Alice	27
10	Mary	21

Id	Name	Cost
2	AB17	250.0
3	AB23	299.9
5	XY70	150.0
7	ZZ55	115.5

Id	Name
1	Weekday
2	Weekend
3	Holiday

Fig. 2 Multidimensional model for the *User* × *Item* × *Time* recommendation space.

Incorporating Context into Recommender Systems

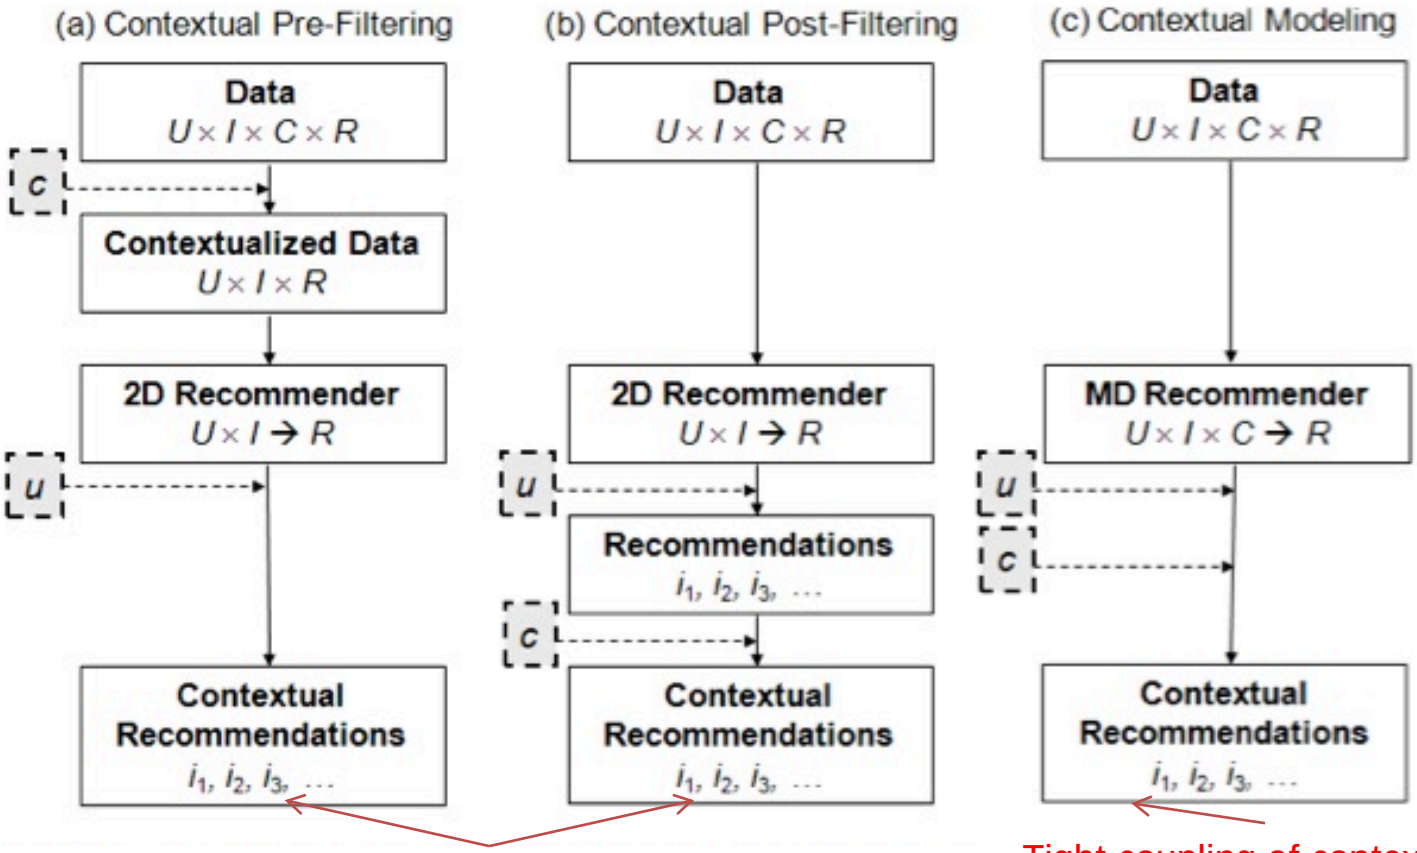

Loose coupling of context and user preferences

= can use traditional 2D recommender system

Tight coupling of context and user preferences

We need to do better!

- **Deal with higher-dimensional information**
  - It is not trivial to extend existing recommendation model to support multidimensional information space

- **Need more than '*Item* to *User*' recommendation**
  - **Item** to **User** in the context of **Location** and **Mood**
  - **User** to **User** (Friend Recommendation)
  - **Item** to **Group** of Users (Group Recommendation)
  - **Item** to **Item** (Device Recommendation) in the context of **User**

- **Need to deal with sparsity problem**
  - As the number of dimensions increases, data becomes more sparse

CENTER FOR INTELLIGENT COMPUTING

Graph-based Data Model

- Graph is general and flexible
- Graph is good for dealing with heterogeneous information
- Exploiting indirect relationships among nodes is important
- Various graph ranking algorithms available

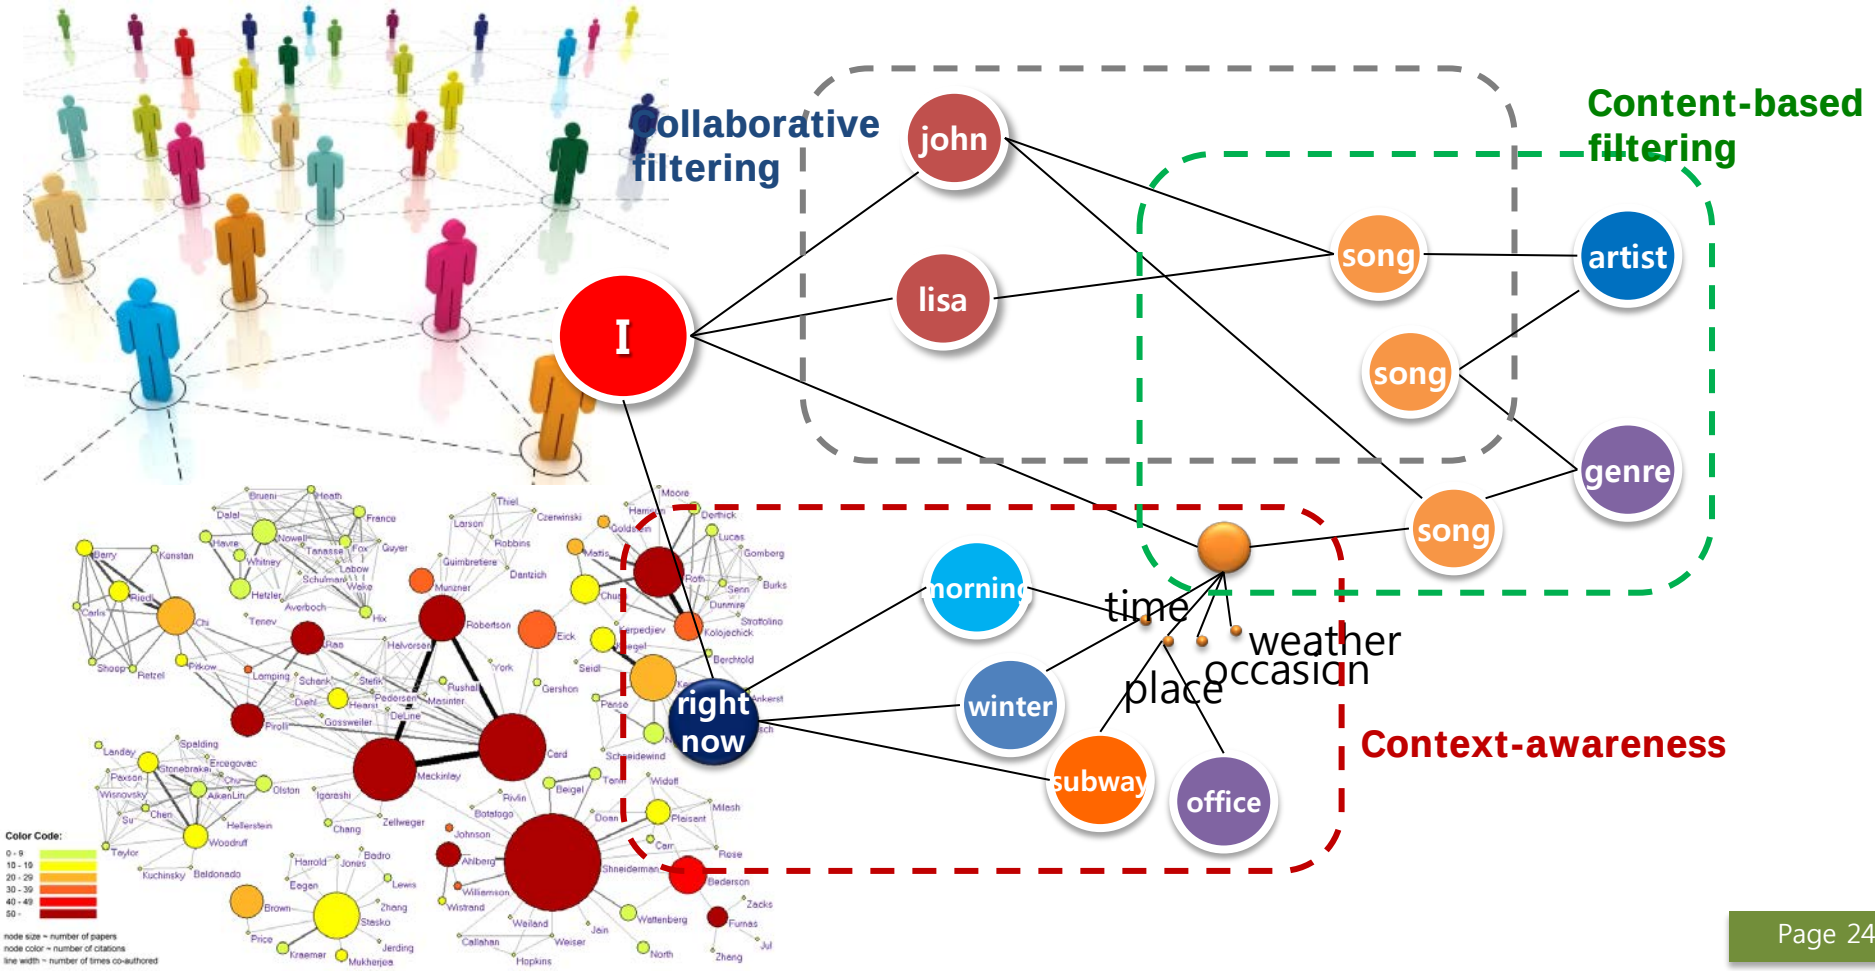

Goal: Ideal Situation

- **Task:** *Michael에게 저녁식사를 위해 가장 적합한 레스토랑을 추천하기?*

Query: 두 노드 'Person: Micheal', 'Context: Today'와 가장 근접한 노드는 (Constrant: Type=Restaurant) ?

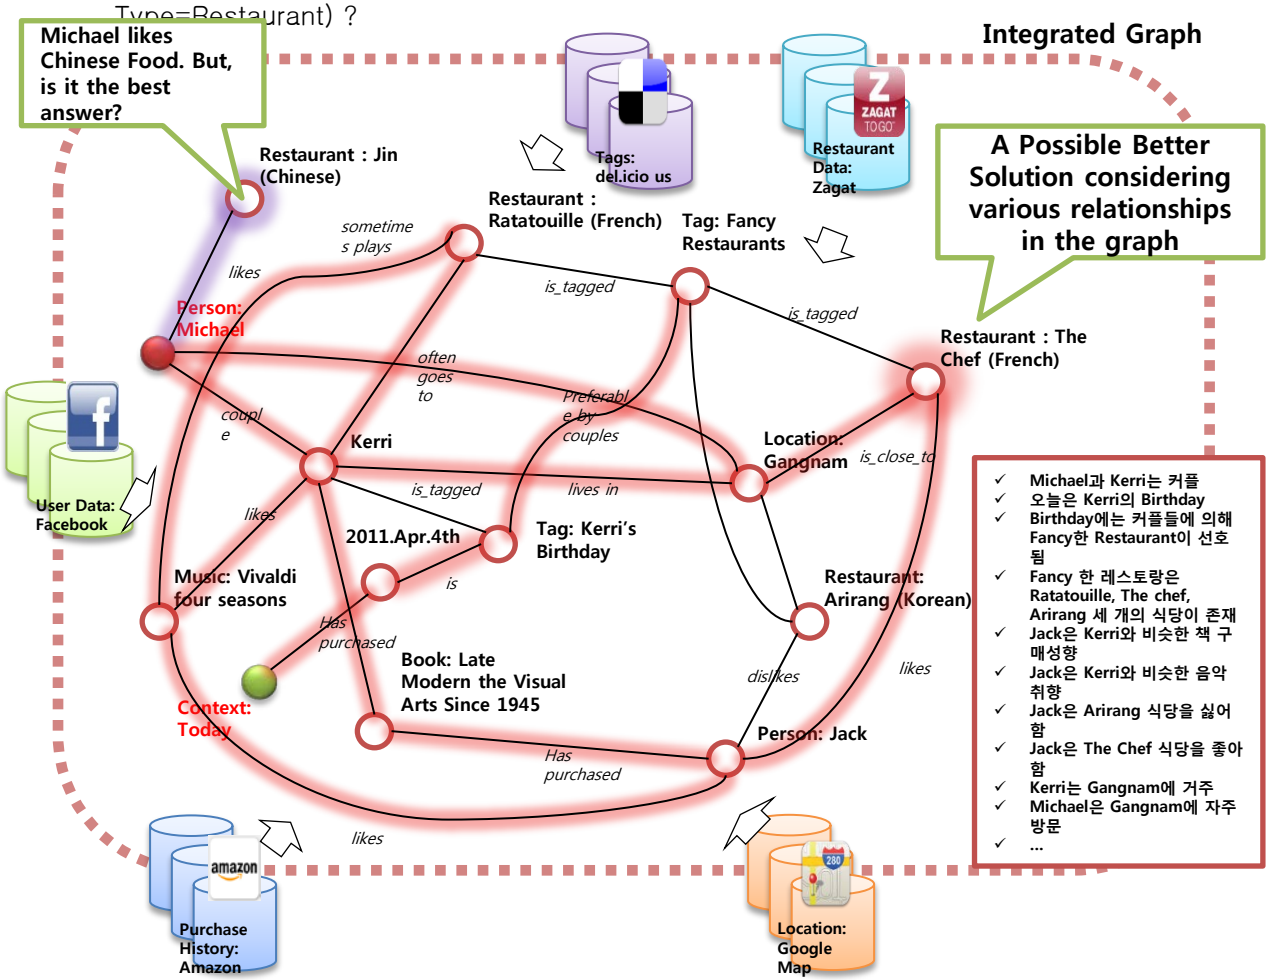

1. Put all information in the graph

2. Use a proximity measure that can fully convey the semantics that exist in the graph

3. For a given set of query nodes, produce the top-k answers depending on context/situation

4. Fully Scalable

Personalized PageRank

- Random walk based relevance measure
  - One of the most widely used methods in measuring relevance (distance) between two nodes
  - Represents the relative importance of the nodes with respect to a start distribution.

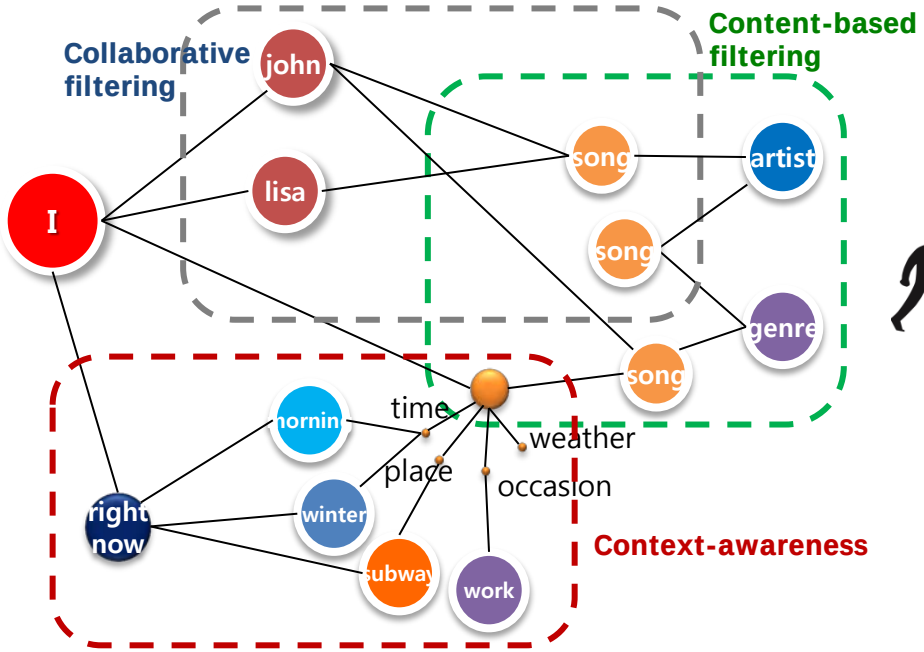

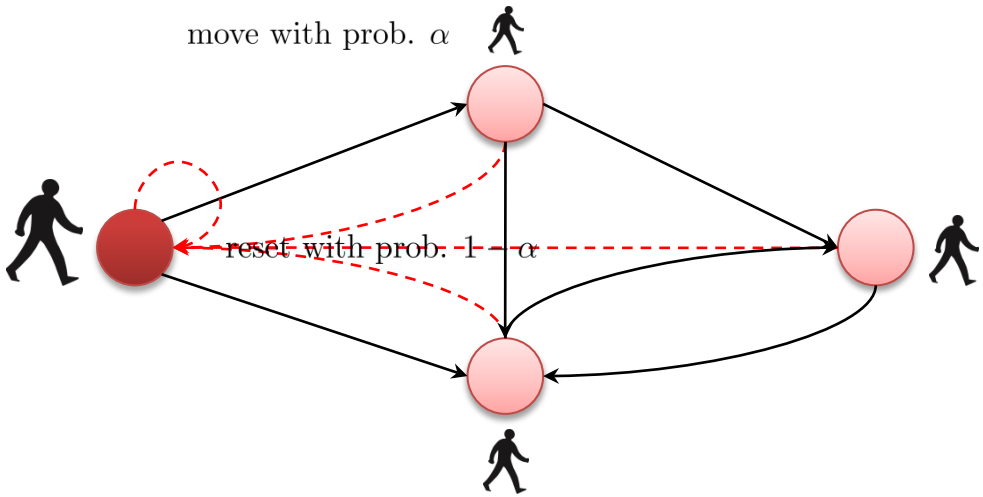

$$\overrightarrow{p_r} = \alpha W \overrightarrow{p_r} + (1-\alpha)\overrightarrow{r}$$

Random Walk based Entity Ranking

- Graph-based Approach that provides flexibility of recommendation
  - Step 1 : Transform Implicit Feedback (log) data into Bipartite Graph
  - Step 2 : Adapt Personalized PageRank and Rank Entities given a query

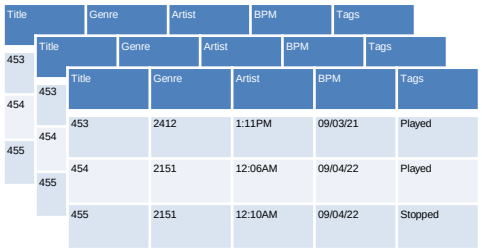

Implicit Feedback (log) data

$F_1 = \{\{USER\}\}, w_1 = 0.5$

$F_2 = \{\{USER, LOCATION\}\}, w_2 = 0.5$

Recommendation Factors

$f_1 = \{\{USER\}\}, w_1 = 0.5$

$f_2 = \{USER, LOCATION\}, w_2$=0.5

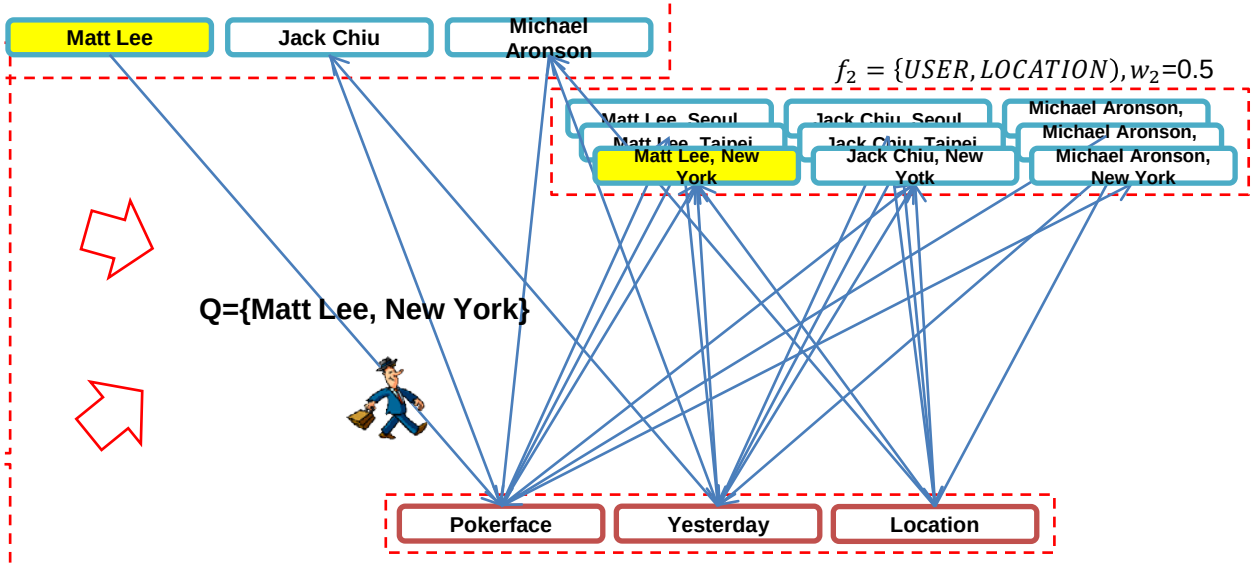

Run Personalized PageRank on the constructed *bipartite graph*

Sangkeun Lee, et al, Random Walk based Entity Ranking on Graph for Multidimensional Recommendation, *ACM RecSys 2011.*

Sangkeun Lee, et al, Flexible Recommendation using Random Walks on Implicit Feedback Graph, *ACM ICUIMC 2011.*

PathRank (Semantic PageRank)

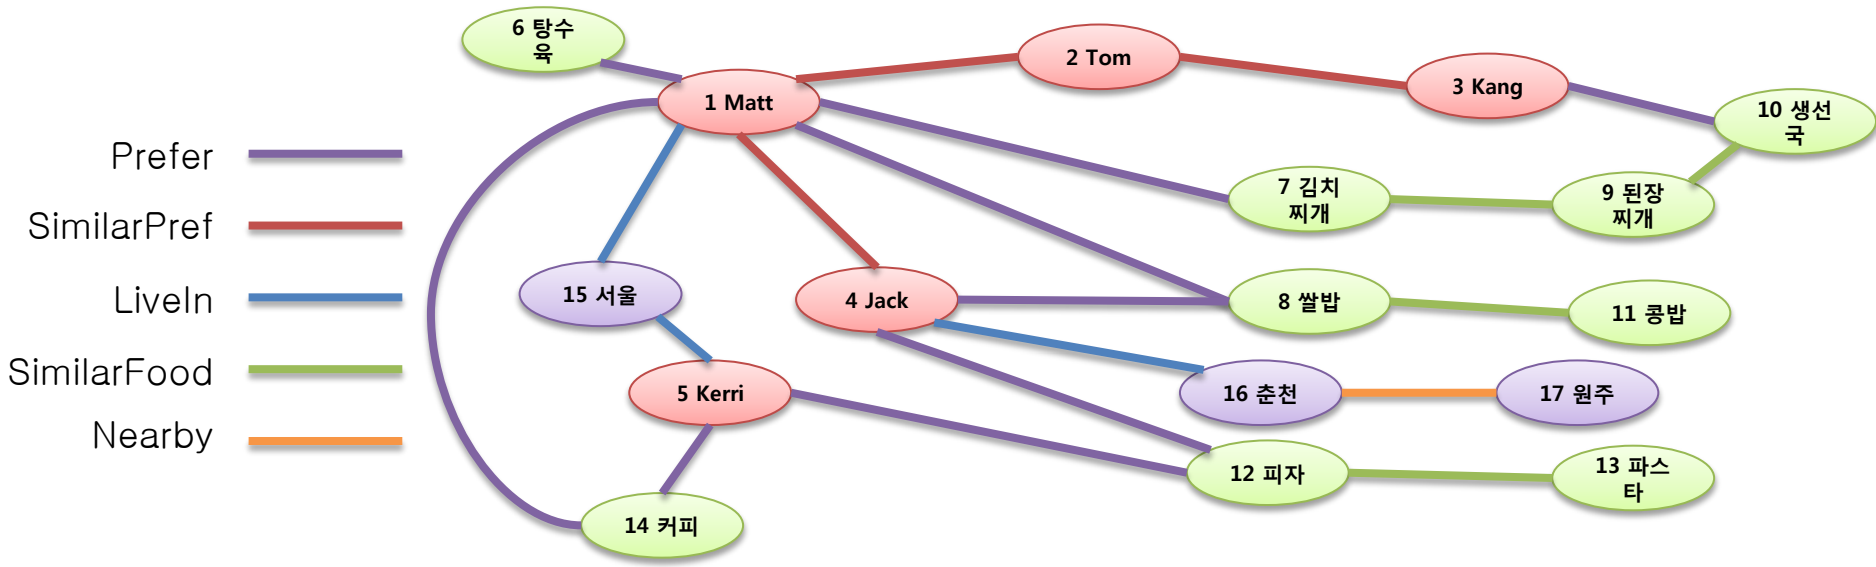

기존의 PageRank 식

$$\vec{r} = cM_G{}'\vec{r} + (1-c)\vec{q}'$$

- 새롭게 정의된 그래프 위에서의 PageRank 계산으로 이해할 수 있음
- 기존의 Pre-computation 기법 수용 가능

제시하는 알고리즘

$$\vec{r} = (w_t M_G{}'^T + w_1 G_1^T + w_2 G_2^T \ldots + w_n G_n^T)\vec{r} + w_r \vec{q}'$$
$$where\ w_t + w_1 + \cdots + w_n + w_r = 1,$$

- 노드의 중요도 (Importance, Authority)

Rule을 만족하는 노드들에 가중치

- 출발 노드와의 관련도 (Relatedness)

Efficient Processing

- Computation of Personalized PageRank is expensive
  - Pre-computation
    - High pre-computation storage costs for large graphs
    - Cannot cope with new data
- Matrix multiplication via multi-way join
  - over MapReduce
- Runtime retrieval of top-*k* elements
  - Utilize shortest path algorithm
- Random walk semantics on RDB

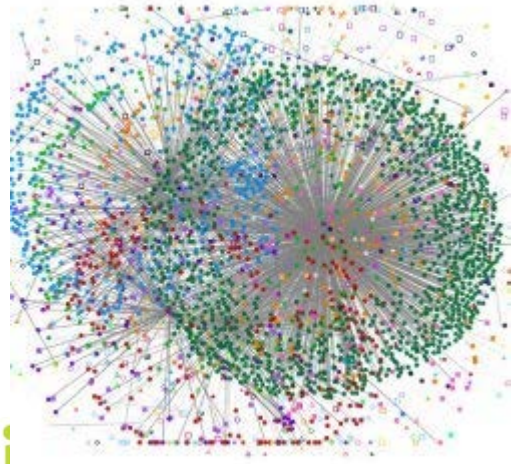

Pre-compute
Personalized PageRank

	n_1	n_2	n_3	$\cdots$	n_i
n_1		0.081	0.045	$\cdots$	0.002
n_2	0.099		0.125	$\cdots$	0.024
n_3	0.076	0.107		$\cdots$	0.009
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
n_i	0.015	0.019	0.038	$\cdots$	

Personalized PageRank Table

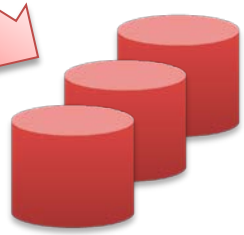

secondary storage

Random Walk Operations on RDB

Random walks on database schema graph

: much smaller than instance-level graph

Random walk on ***Data graph***

Start nodes:
{Tom Hanks, Steven Spielberg}

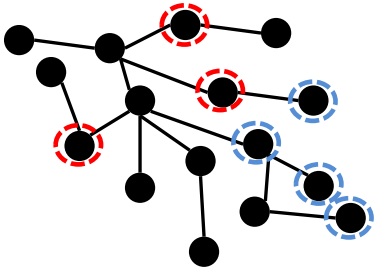

Result top-k nodes:
{Terminal, Catch me if you can, 111017 Tweets, …}

Data graph extraction

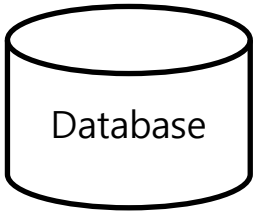

Limitations:

1. Hard to maintain the huge data graph in main memory
2. Much overhead to extract data graph
3. Missing semantics of original database design

Random walk on ***Schema graph***

Start nodes:
{Tom Hanks, Steven Spielberg}

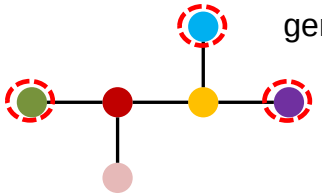

Structured query generation

Random walk results as queries

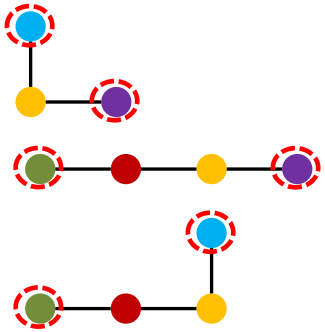

Schema graph extraction

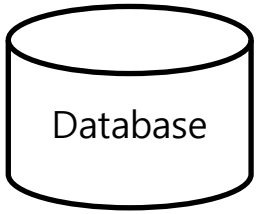

Query evaluation

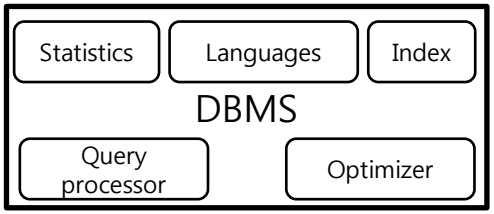

Random Walk Operations on RDB

- Two phase process
  1. Generate structured queries for the random walk
     - Interpret user's query semantics flexibly
     - Efficiently prune unpromising walk-paths
  2. Evaluate the random walk queries on DBMS
     - Utilizing the power of database systems, such as, query processing and optimization
     - No need to alter/transform existing databases to graphs

Processing performance

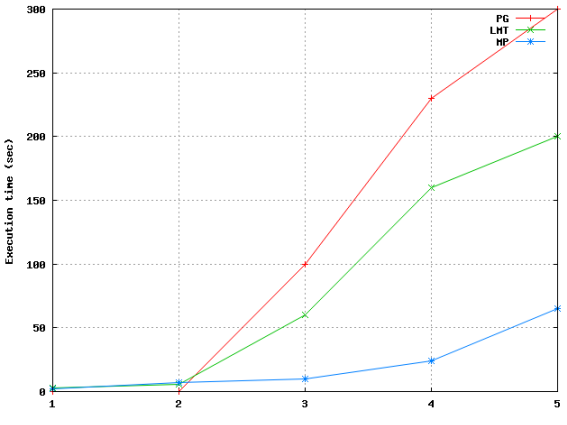

# of start nodes

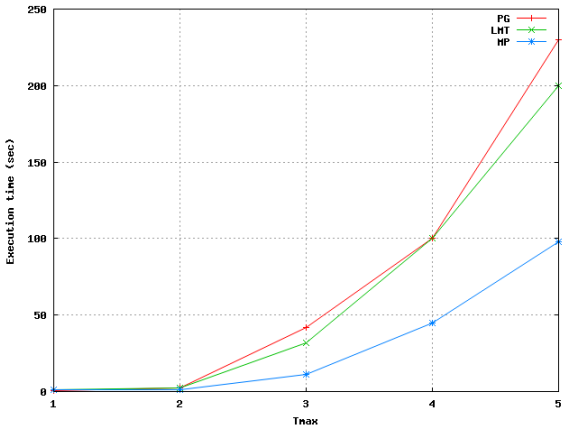

Max. size of selected paths

J. Park and S.-g. Lee. Exploiting correlation to rank database query results, *DASFAA 2011*.
J. Park and S.-g. Lee. Keyword search in relational databases. *Knowledge Information Systems*, 26(2), 2011.

Fast Retrieval of Top-*k* P-PageRank Elements

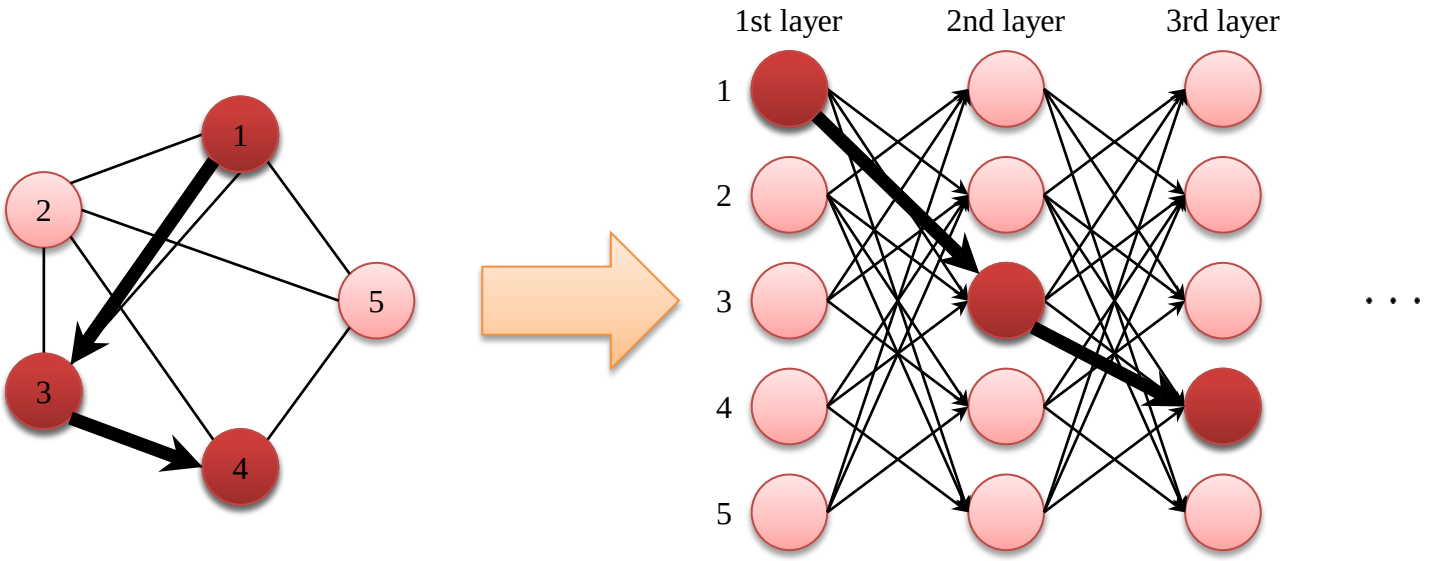

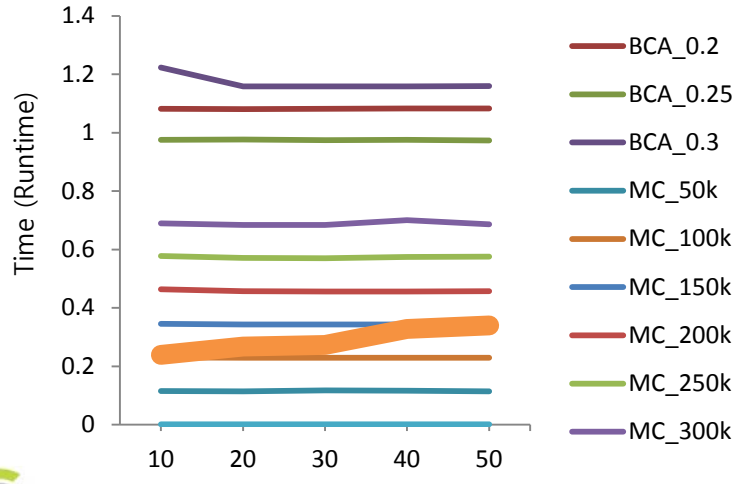

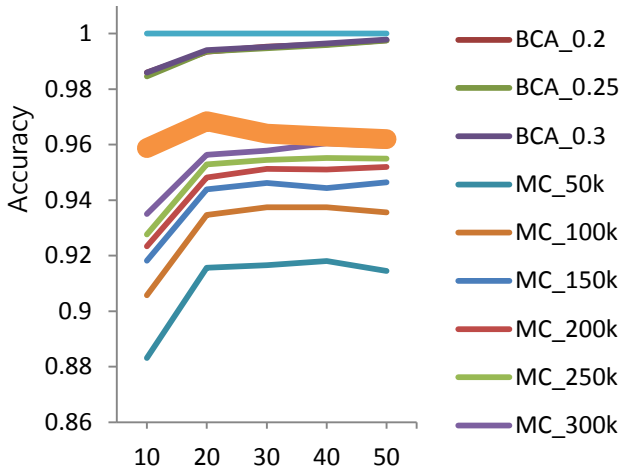

CIC & Testbed

SNU Smart Reality Testbed

- A testbed implementing an open platform for Smart Reality
- A sandbox for various devices, apps and services
  - Real users, real data and real feedback
  - Integrated space of multiple devices and databases
- A playing ground for Creative SW & Services Incubation (CSI) participants

Smart Reality Open Platform

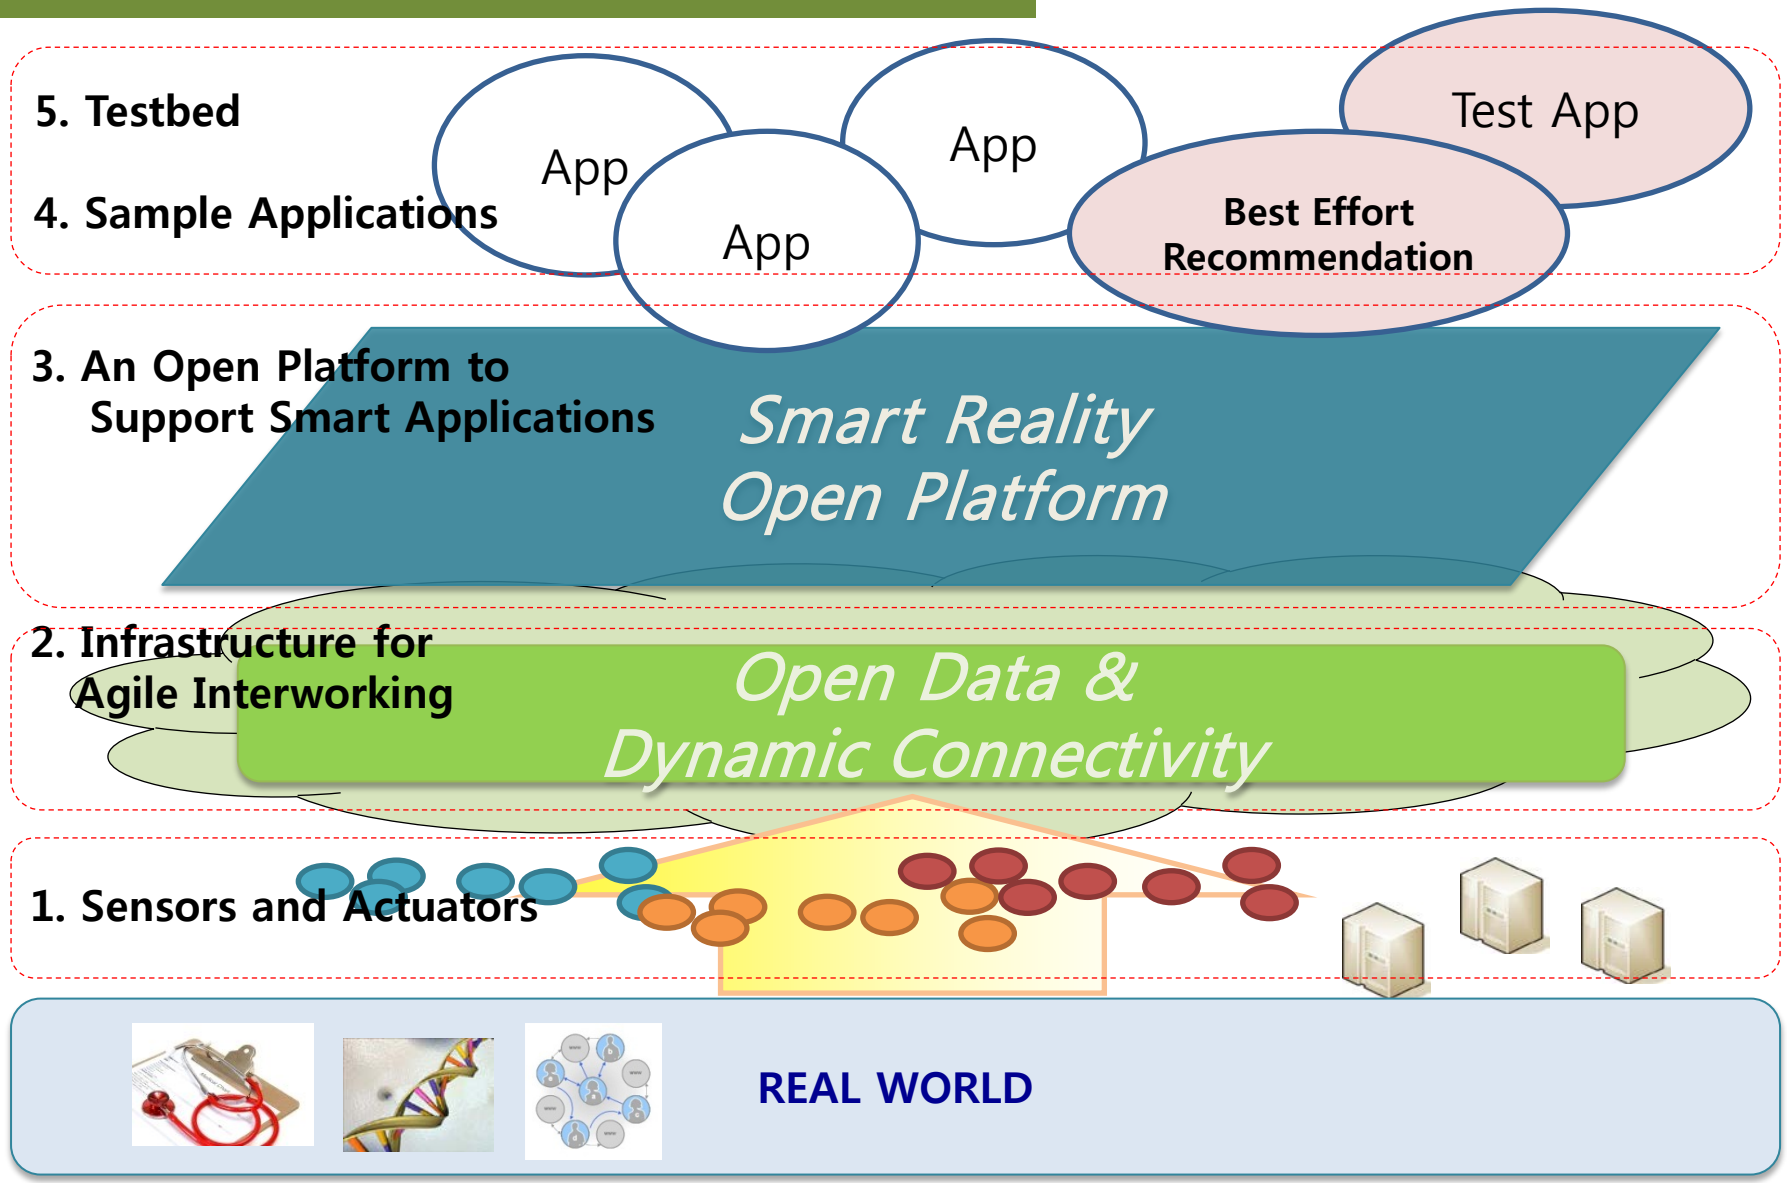

기대효과

- Low-Risk & Low-Cost Test Environment
  - Controlled environment
  - Subscribed users
  - Readily available infrastructure
- Data-Driven Intelligence
  - Smartphone log data
    - Call: 0.25%
    - Text message: 0.44% => abundance of data to mine on
    - eMail: 0.75%
    - App usage: 13.15%
    - GPS & Other activities: 85.46%
- Network Effects of Smart Objects
  - More devices talking to one another => $Value \propto N^2$
- Crowd-Sourcing
  - Open platform for users/developers/vendors to freely add new services
- Standards
  - 다양한 스마트 서비스를 지원하기 위한 개발자 API
  - 다양한 스마트 서비스 사이의 데이터 전달을 위한 네트워크 프로토콜
  - 다양한 스마트 서비스에서 발생하는 데이터 / 로그 통합 모델

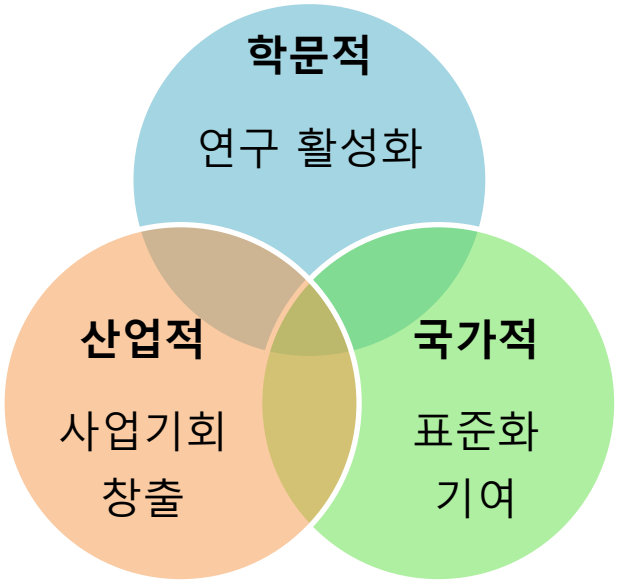

CENTER FOR INTELLIGENT COMPUTING

SNU Smart Reality Testbed

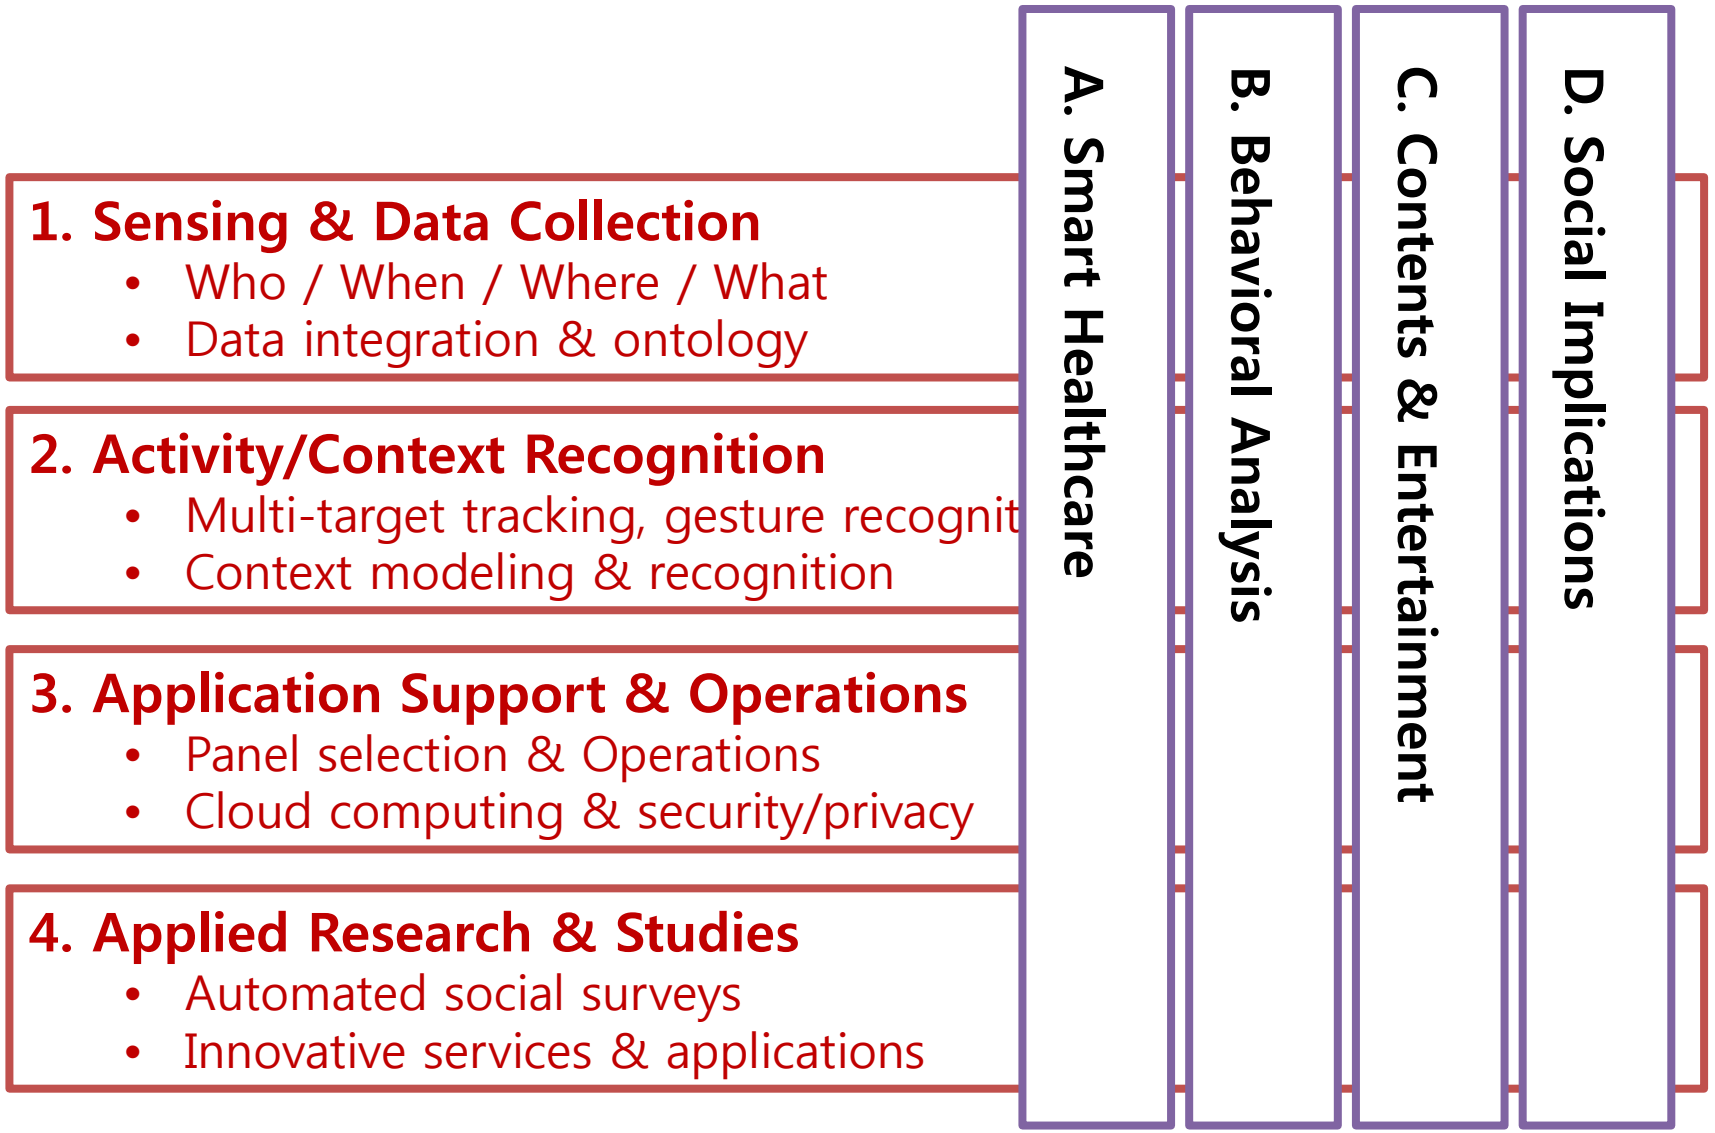

Intelligent Computing in Smart Reality

CENTER FOR INTELLIGENT COMPUTING
서울대-삼성전자 SW 공동연구센터