LEVERAGING BIG DATA TO IMPROVE SEARCH PERFORMANCE

Clement Leung

Chinese University of Hong Kong, Shenzhen

(clementleung@cuhk.edu.cn)

- FULL PROFESSORSHIPS at
  - Chinese University of Hong Kong, Shenzhen, China
  - University of London, UK
  - Victoria University, Australia
  - Hong Kong Baptist University
  - National University of Singapore
- Two US patents & published five books and over 150 research articles
- Program Chair, Keynote Speaker, Panel Expert of major International Conferences
- Editorial Board of ten International Journals
- Listed in Who's Who in the World and Great Minds of the 21st Century
- Fellow of the British Computer Society, Fellow of the Royal Society of Arts, and Fellow of the International Academy, Research, and Industry Association

CONCEPT-BASED SEARCH

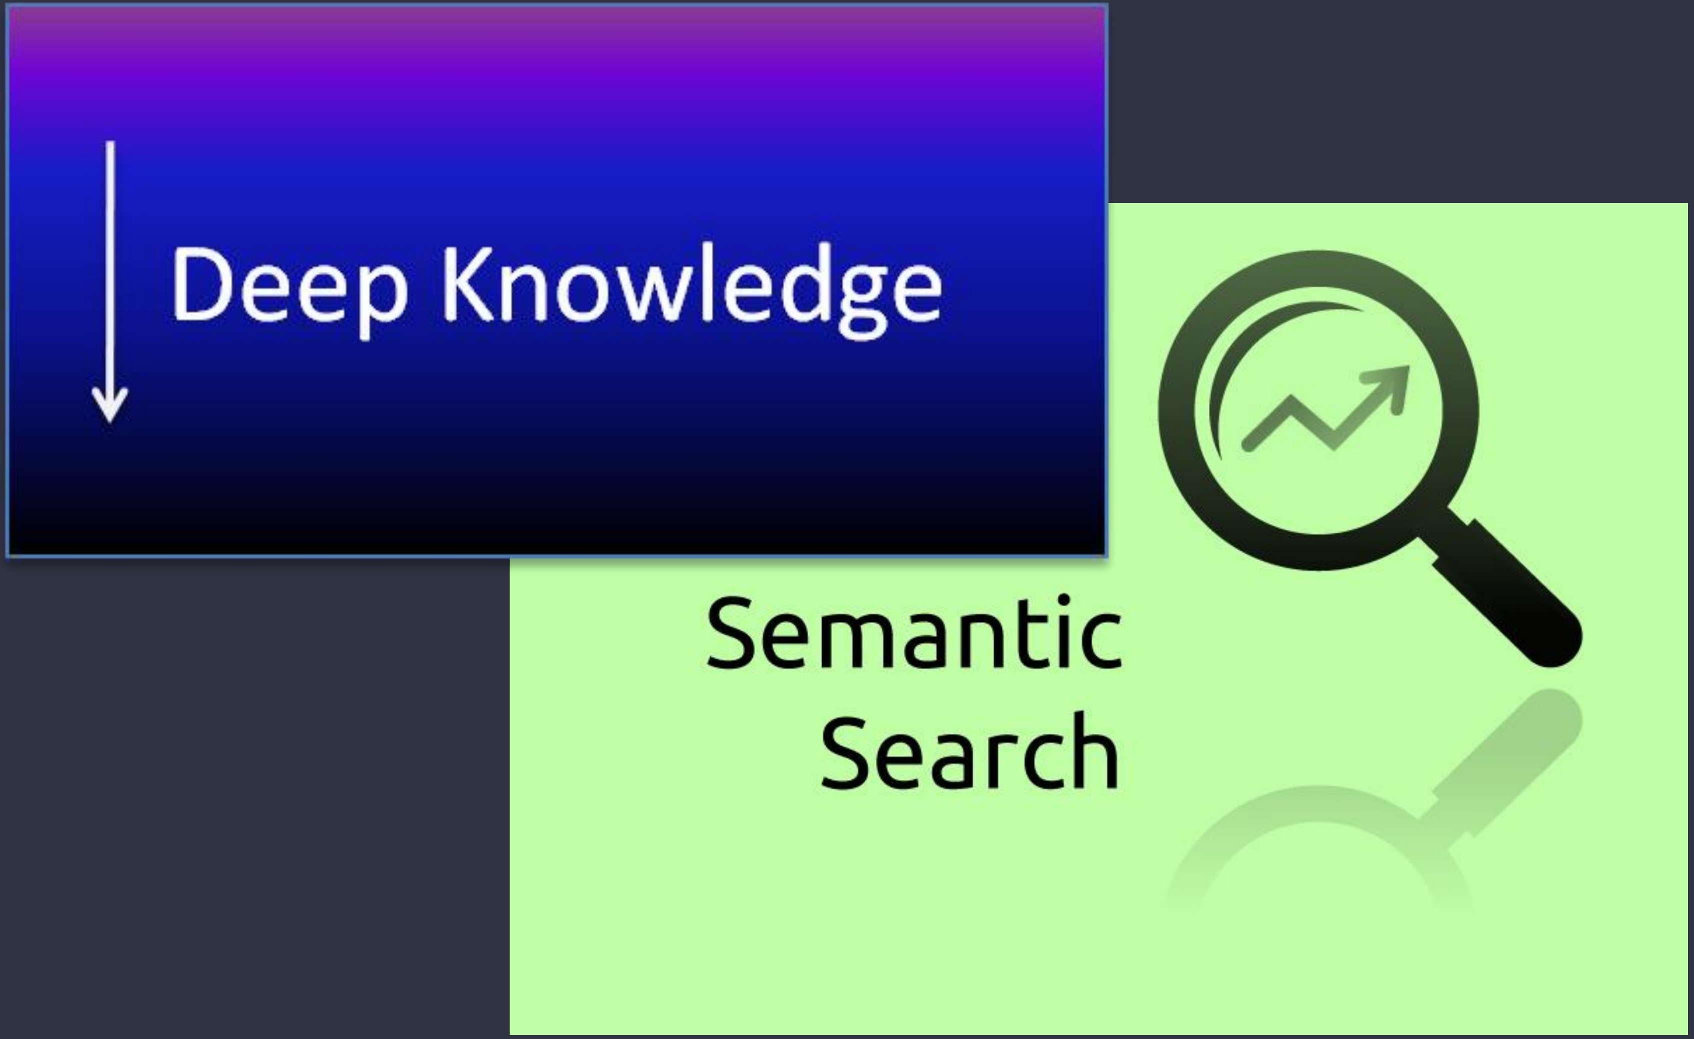

EXPONENTIAL GROWTH OF THE DIGITAL UNIVERSE

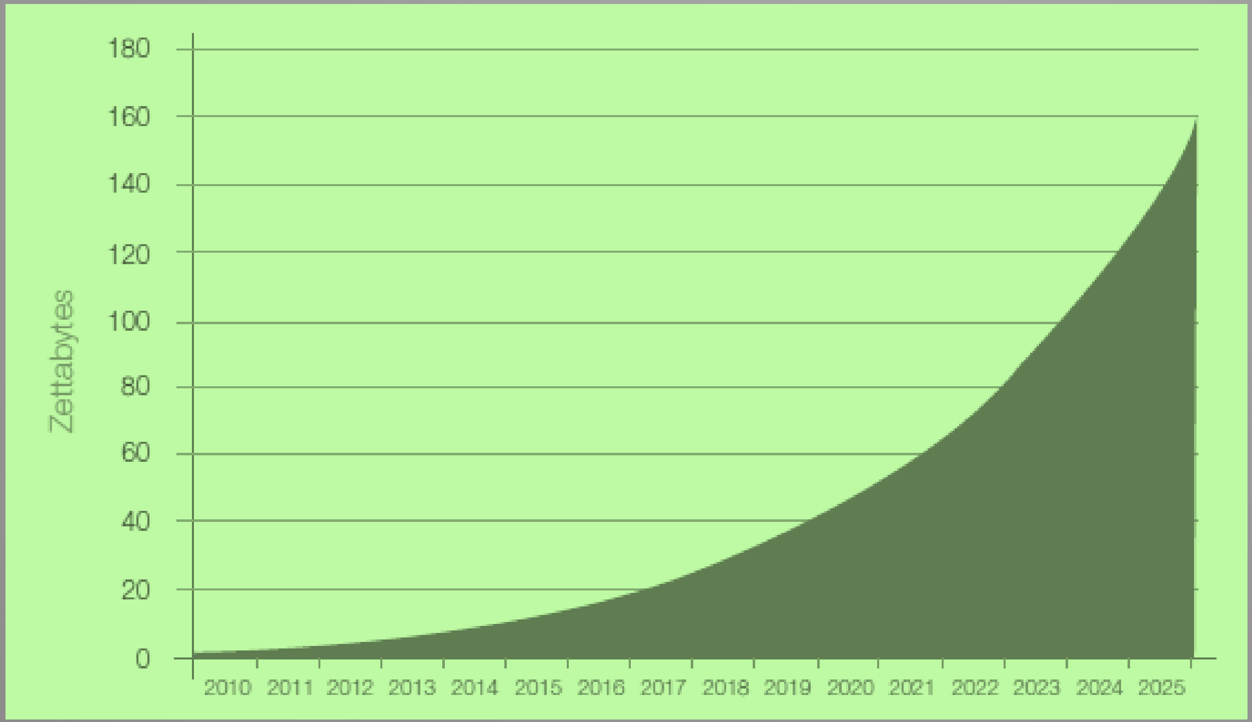

Source: IDC's Data Age 2025 study, sponsored by Seagate, April 2017

- By 2025, the amount of data ≈ 160 zettabytes (160×10^{21})
- 1 zetta seconds = 31.71 trillion years = 2300 x age of universe
- The volume of seawater in the Earth's oceans is approximately 1.37 zettalitres

TURING: “All Problems are Essentially Search Problems”

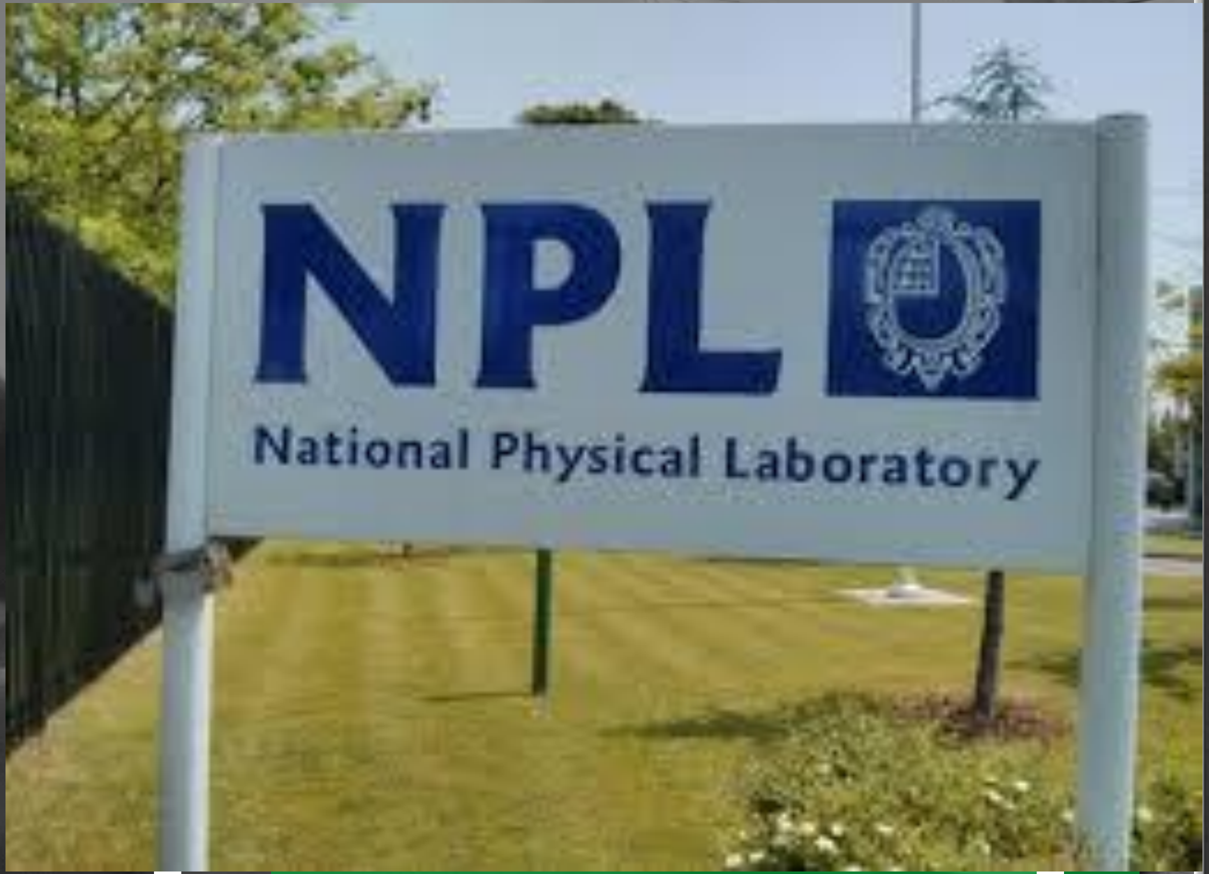

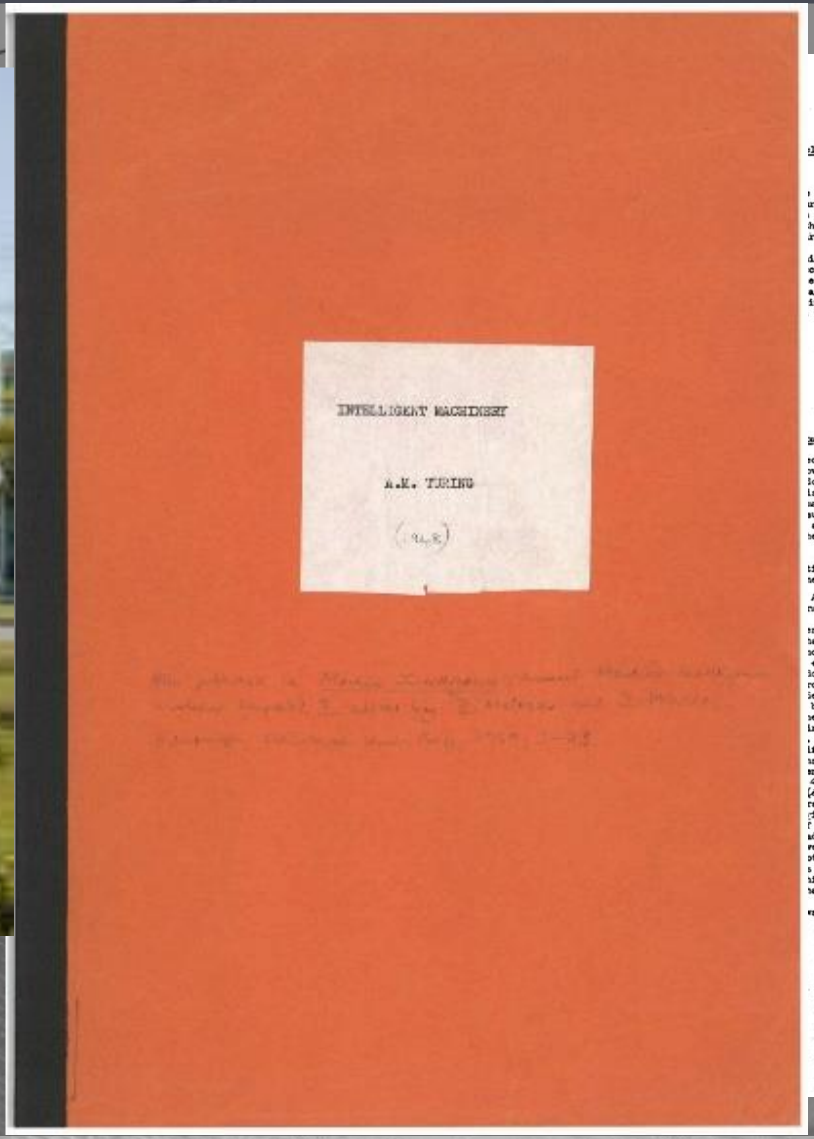

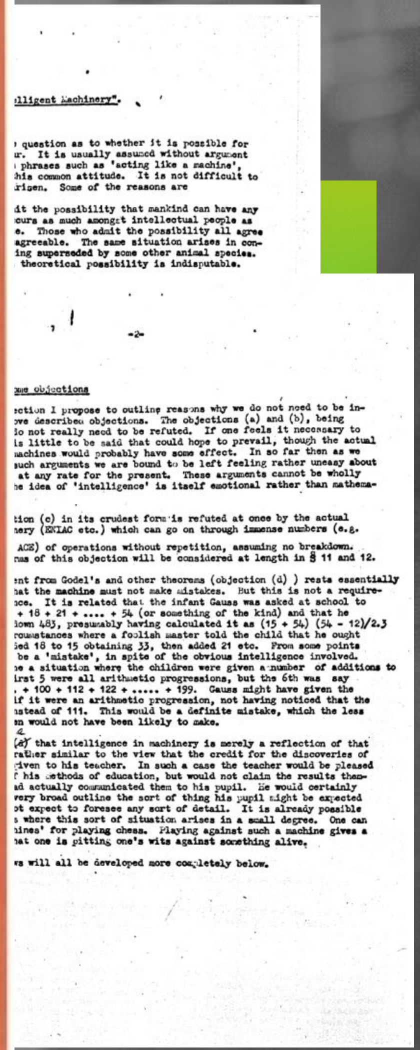

ALL PROBLEMS ARE ESSENTIALLY SEARCH PROBLEMS

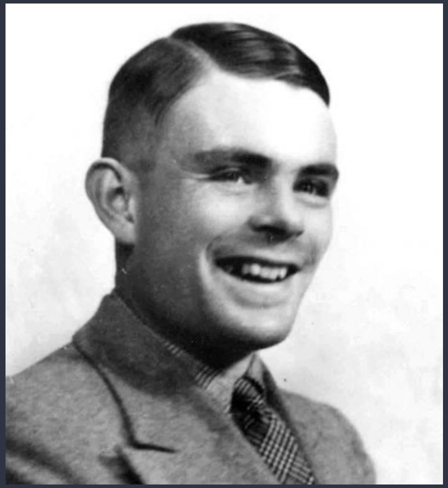

"…intellectual activity consists mainly of various kinds of search."

"…evolutionary search…"

"indexes of experiences"

"…cultural search…by the human community as a whole"

- In 2009, Google researchers publish a paper entitled “The Unreasonable Effectiveness of Data”
- The paper suggests that simple models with lots of data are better than elaborate models with less data
- No need to learn from algorithms how humans perform the tasks because the intelligence is in the data, not the algorithms


The Unreasonable Effectiveness of Data

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

No explanation, no justification?

DIGITAL MEDIA SEARCH

KEY ISSUES

- Distinct genres
  - Multimedia entities: Images, Videos, 3D Models, Stream Music…
- Semantic gap
  - Low-level features VS high-level semantics
  - Concept-Based VS Content-Based
- Lack of Efficient Methods
  - Trade-off between precision and complexity
  - Absence of a universally-accepted retrieval method

DEEPER SEMANTICS

Need to have significant background knowledge to retrieve this image (David and Goliath)

- Objective factual description
- Interpretation of the objects in the picture based on prior knowledge
- Familiarity with the subject matter

LEVELS OF SEMANTICS

Primary Level	• Objective description of content
Secondary Level	• Basic level of interpretation/summary of some/all entities
Tertiary Level	• Knowledge-based elements

Perception and interpretation of semantic content

- Depends on user knowledge & experience
- Degree of content richness can be distinguished into 3 levels
- Information encoded in a multimedia object can be potentially unlimited

Learning from Big Data

INDEX LEARNING AND EVOLUTION

ADOPTION OF REINFORCEMENT LEARNING ARCHITECTURE

INDEX BUILDING THROUGH EXPLOITING BIG DATA

- More than 2 billion people are connected to the Internet
- 20 sec of their time is more than 10 million man hours
- Empire State Building took 7 million man hours to build
- One Empire State Buildings built every 15 seconds

EXPLOITING BIG DATA

- System continuously learns and evolves
- Adapts its answer lists to queries
- Satisfies latest user preferences
- Incorporates deep knowledge from users
- Unsupervised Reinforcement Learning
- Semi-Markov Decision Process framework

EXPLOITING BIG DATA

- Capture subtle nuances of human perceptions and deep knowledge
- Degree of reinforcement is suitably calibrated
- Indexes evolve naturally to accommodate dynamic user interests
- Continuously develop indexes
- Injecting progressive improvement in search performance

INDEX EVOLUTION

- Dynamically builds semantic indexes to associate query terms with multimedia objects
- Construction of new Indexes and Deconstruction of existing Indexes
- Advantages
  - Addresses newly-emerged concerns (e.g. privacy, legitimacy) in social networks
  - Automatic evolution without human invention

INDEX EVOLUTION

- A learning function $\mathcal{L}$ between the object space and the term space to measure the relevance liaison:

$$\boldsymbol{\mathcal{L}}: \boldsymbol{T} \times \boldsymbol{O} \rightarrow \mathbb{R}^{+} \cup \{\mathbf{0}\}$$

- For each link between an index term τ and a multimedia object o:
  - $\mathcal{R}$: relevance base; ω: normalized weight

$$\boldsymbol{\mathcal{L}}(\boldsymbol{\tau}, \boldsymbol{o}) \rightarrow \boldsymbol{\omega}\boldsymbol{\mathcal{R}}, \boldsymbol{\omega} \in (\mathbf{0}, \mathbf{1})$$

- Output – Relevance Index Value (RIV): non-negative real numbers that specify corresponding relevance with 0 indicating complete irrelevance

KNOWLEDGE EXTRACTION

- Clicking Model - implicit clicking information
- Click depends on pertinence of objects
- Reward Function
  - Scalar positive or negative reward
- Iterative learning process
  - Robustness: eliminate random errors and fault tolerant

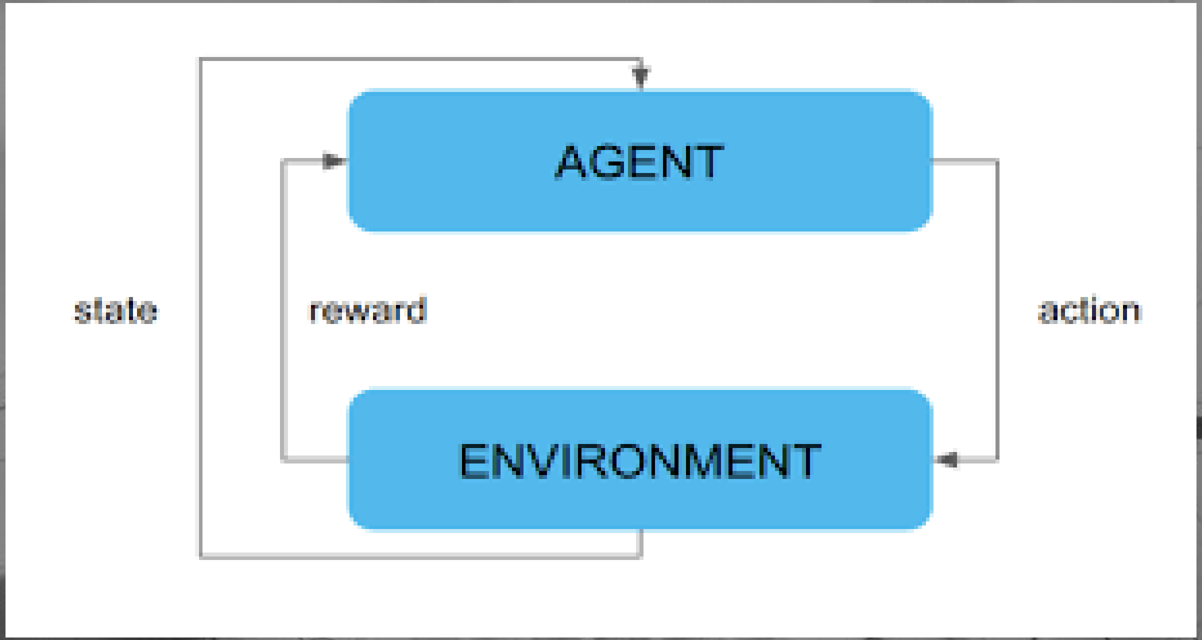

- Reward is observed via the clicking information
- Take actions based on current states
- Learn from useful evaluative information provided by users

REINFORCEMENT LEARNING

Semi-Markov Decision Process

Action Space *A*:

- Agent selects a set of objects and presents to user
- Intuitively huge due to a large number of combinatorial possibilities

State space S:

- A set of all indexes in dynamic indexing
- Include explored and unexplored Indexes

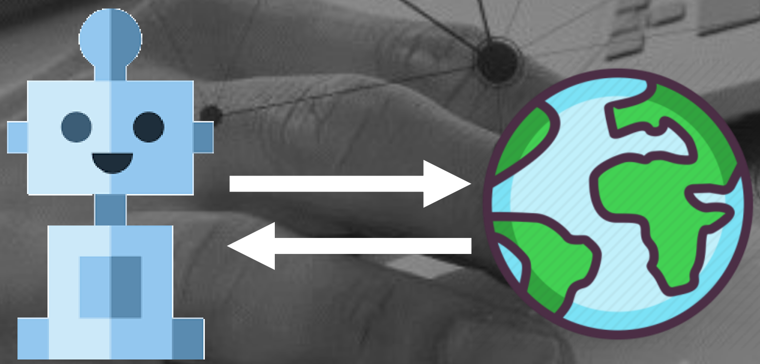

Reinforcement

- If a click results in a reinforcement made to a particular element x, then we have for its new value $x' = x + \Delta$, where Δ is a parameter calibrated to represent the extent of the reinforcement
  - A large value of Δ signifies strong reinforcement, while a small value of Δ signifies weak reinforcement

Semi-Markov Decision Process

Rewards

- If n elements from the results list $x_1, \ldots, x_n$ receive reinforcement, then the reward would be $n\Delta$.
- If no clicking response is received from the user on a given list, then all the relevant elements between the submitted query terms and retrieved objects will be decremented by a given amount Δ.
  - If there are q query terms, then for an object list of K objectS, the total decrement will be $qK\Delta$, which corresponds to a (negative) reward of $-qK\Delta$.

Semi-Markov Decision Process

- The entire process

$$\{\boldsymbol{Z}(t), t > 0\}$$

gives the state the process at time t and is a semi-Markov process, with the transition epochs being the *embedded Markov chain* of the process

- Here, the embedded Markov chain characteristics are such that a transition will trigger a move of the RIV of certain indexes an amount $\pm\Delta$

Semi-Markov Decision Process

Semi-Markov Decision Process

- Letting $\{\pi_i\}$ denote the stationary probabilities of the embedded Markov Chain, then it can be shown that the limiting state probabilities are given by

$$P_i = \frac{\pi_i \mu_i}{\sum_j \pi_j \mu_j}$$

where μ_i is the average time the process spends in state *i*

- The semi-Markov process takes into account the real time elapsed between state transitions, as simply sampling the system at transition points will not give a full picture
- The semi-Markov process is no longer memoryless because given the present state, the future evolution of the process is no longer independent of the past
  - In order to predict the future, it is not sufficient to just know the present state, but also the length of time that the process has been in the present state

No Longer Memoryless

LEARNING CONVERVENCE

The problem of *learning convergence* is concerned with whether unexplored indexes will eventually evolve to become explored indexes.
The index learning behaviour is considered to be convergent if the vast majority of unexplored indexes become explored indexes in the learning process

LEARNING CONVERVENCE

- **TOR-Tuple [τ_i, o_j, r_{ij}]**
  - represent the relationship between a query term τ_i and an object o_j, with r_{ij} as RIV
- **Index Criteria**
  - r_{ij} must attain or surpass the pre-defined threshold value h
- **Two categories of indexes**
  - Unexplored indexes: reside in the index generator
  - Explored indexes: promoted from index generator into the index pool

LEARNING CONVERVENCE

- Consider a query with multiple terms $Q(\tau_1,\ldots,\tau_t,\tau_{t+1},\ldots,\tau_n)$
- Explored indexes exist for $\{\tau_1,\ldots,\tau_t\}$ with a set of multimedia objects
- $\{\tau_{t+1},\ldots,\tau_n\}$ are entirely new terms
- *Returned Object List* is a union set
  - a k-object subset O_a: $\{o_1, o_2,\ldots,o_k\}$ that has the highest cumulative RIV scores
  - a subset O_b of random objects selected for exploration

- Naïve Strategy
  - Returns the best *K* result objects ordered by index scores in decreasing order (*k* = *K*)
    - Objects with high RIV will be shown repeatedly for user evaluation
    - The RIV of these objects tends to keep increasing even though they may not be the most relevant as these are selected as relevant (clicked) by the users
  - Would lead to ***local maxima problem***
    - Objects that have the highest RIV may not in fact be the most relevant

PURE EXPLOITATION

PURE EXPLORATION

- Randomized Strategy
  - Returns K result objects purely by random extractions (k = 0)
  - Discover ‘hidden’ objects by randomness
  - Exploration vs exploitation

Convex Combination

- We use a convex combination of exploration and exploitation
  - For $0 < u < 1$, we include a proportion of α randomly chosen objects, and v objects from exploitation, with $u+v=1$
    - $u = 0$ corresponds to pure exploitation
    - $u = 1$ corresponds to pure exploration
  - The β objects from exploitation can be chosen according to their ranks, or picked from the retrieved list using some randomized algorithm

State Classification

- ***Exposed***: If an unexplored index is elevated into the index pool, it can be regarded as exposed.
- ***Extant***: Conversely, it is regarded to be extant and still in the evolutionary process.

ANALYSIS OF LEARNING CONVERGENCE BEHAVIOUR

State Simplification

- The state of the process in this particular context can be simplified as the number of extant unexplored indexes.

PROBABILITY A GIVEN UNEXPLORED INDEX BE EVENTUALLY EXPOSED

$$\frac{1-[\frac{1-f}{f}]^k}{1-[\frac{1-f}{f}]^h}$$

where f (>1/2) is the positive feedback probability,

k (< h) is the current score of the index,

h is the threshold

TIME FOR A GIVEN UNEXPLORED INDEX TO BE EXPOSED

Average time:

$$1 + \frac{2q^{h-k+1} - (h-k+1)q^2 + h - k - 1}{(1-q)^2}$$

where $q = (1 - f)/f$

IMMIGRATION-DEATH PROCESS

Evolutionary Behaviour

- A total of C unexplored indexes initially
- Let S_t denote the reduced system state at time t, omitting the RIV scores: $S_0 = C$
- In the evolutionary stage, S_t decreases compared with S_0

IMMIGRATION -DEATH PROCESS

Evolutionary Behaviour

- Whenever an unexplored index is exposed, S_t is decreased by one.
- At a particular time t, the number of remaining unexplored indexes follows an immigration-death process with rate α :

$$P(S_t = k) = \binom{S_0}{k}(e^{-\alpha t})^k (1 - e^{-\alpha t})^{S_0 - k}$$

$$E(S_t) = S_0 e^{-\alpha t}$$

$$V(S_t) = S_0 e^{-\alpha t}(1 - e^{-\alpha t})$$

- $E(S_t)$ denotes the expected number of extant unexplored indexes after a time interval Δt
- $V(S_t)$ gives the corresponding variance

IMMIGRATION-DEATH PROCESS

Evolutionary Behaviour

- Let $t \to \infty$:

$$\lim_{t\to\infty} E(S_t) = 0, \quad \lim_{t\to\infty} V(S_t) = 0$$

- Indicates eventually the entire collection of unexplored indexes will be fully discovered
- The limit of $V(S_t)$ reveals that the effect of stochastic fluctuation tends to diminish in the course of the learning process, leading the system to evolve into a steady state
- The index learning behaviour acts like a deterministic evolution

THEORETICAL RESULTS

Evolutionary Behaviour

- Let T_s denote the expected time spent on indexing a proportion p of unexplored indexes
- A monotonically increasing function for T_s in regard to p can be obtained:

$$T_s = d \ln \frac{1}{(1-p)} \quad , \qquad d = \frac{1}{\alpha}$$

- The time unit of days is adopted for α, and 1/α can be seen as the average number of days spent to discover an unexplored index.
- Let $p \to 1$, the limit of T_s tends to be infinite, suggesting when most unexplored indexes are exposed, the indexing time for the remaining few ones tends to be large

THEORETICAL RESULTS

Evolutionary Behaviour in SLSE

- Taking the derivative with respect to *p*, we have the incremental indexing time I_t in discovering unexplored indexes:

$$I_t = \left(d \ln \frac{1}{(1-p)} \right)' = \frac{d}{(1-p)}$$

- The larger the values of *p, the higher* the incremental indexing time tends to be

THEORETICAL RESULTS

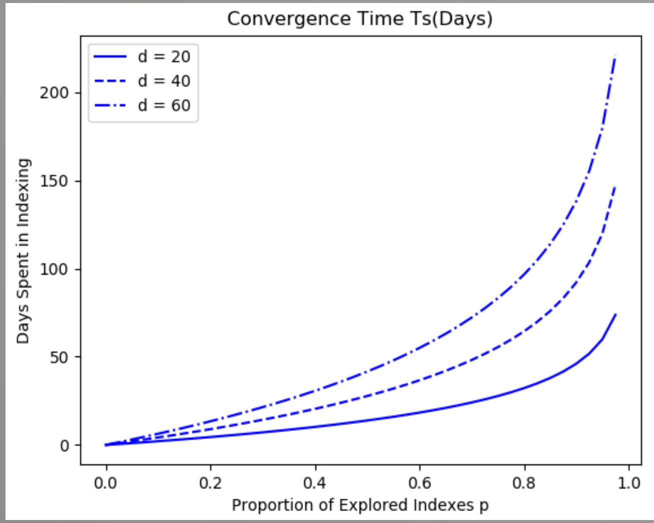

(a) Ts: d = 20, 40, 60

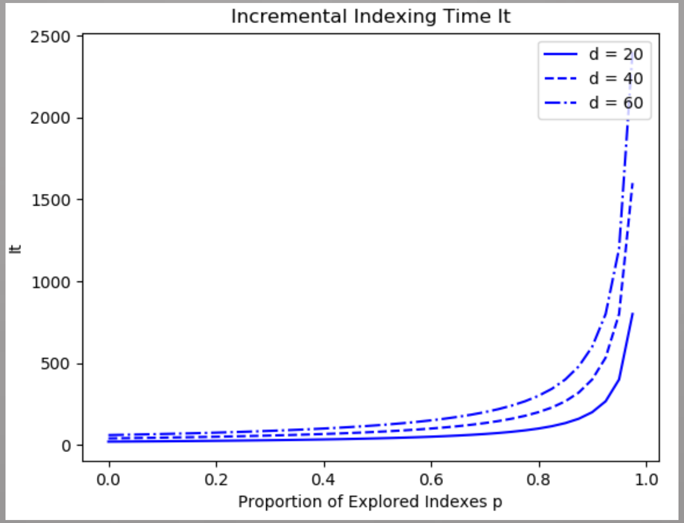

(b) It: d = 20, 40, 60

If 90% of the unexplored indexes needs to be indexed, approximately 2.3 times the amount of time d for discovering an individual index can be expected.

E.g. when d = 40, roughly 92 days are required for convergence. This can be seen by examining the above equation, and letting p = 0.9:

$$T_s = d \ln \frac{1}{(1-0.9)} = 2.3d$$

SIMULATION EXPERIMENTS

- Consider the event that an end user submits a query and furnishes evaluative feedback by clicking interested returned objects.
- If such event streams are presented by a Poisson process with rate α, the inter-event time has the following exponential density function:

$$f(t) = \lambda e^{-\lambda t}$$

- Concerning the indexing behaviour, three different scenarios exist in regard to the submitted query $Q(\tau_1, \ldots, \tau_t, \tau_{t+1}, \ldots, \tau_n)$

SIMULATION EXPERIMENT

- Each time that an arrival of a user occurs, the user selection behaviour is enacted by randomly clicking on some objects in retrieved list
- Graphs of S_t at each arrival time t
  - The blue dot series represent experimental results
  - The black lines correspond to theoretical analysis
  - Sampling frequency is set at five-day interval

SIMULATION EXPERIMENT

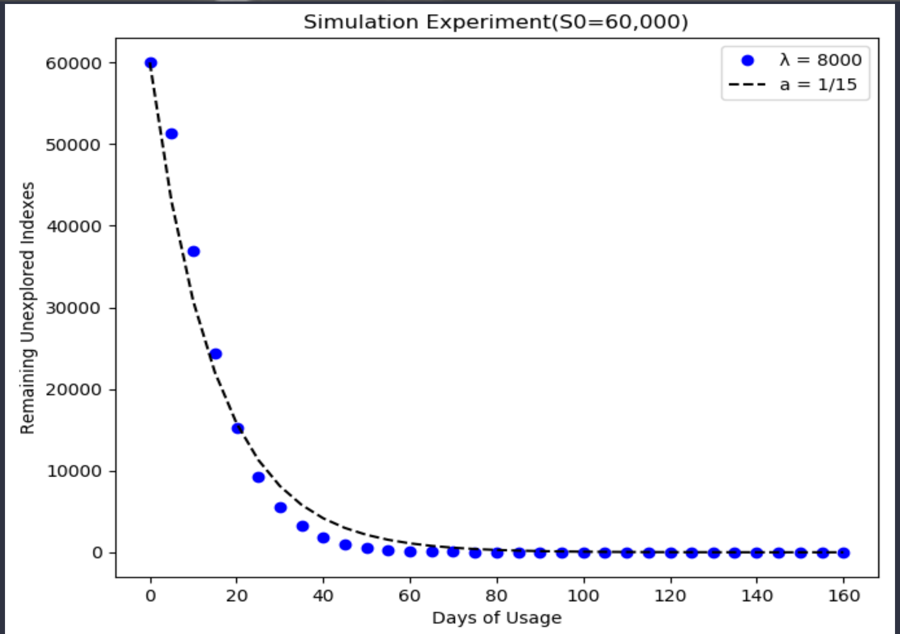

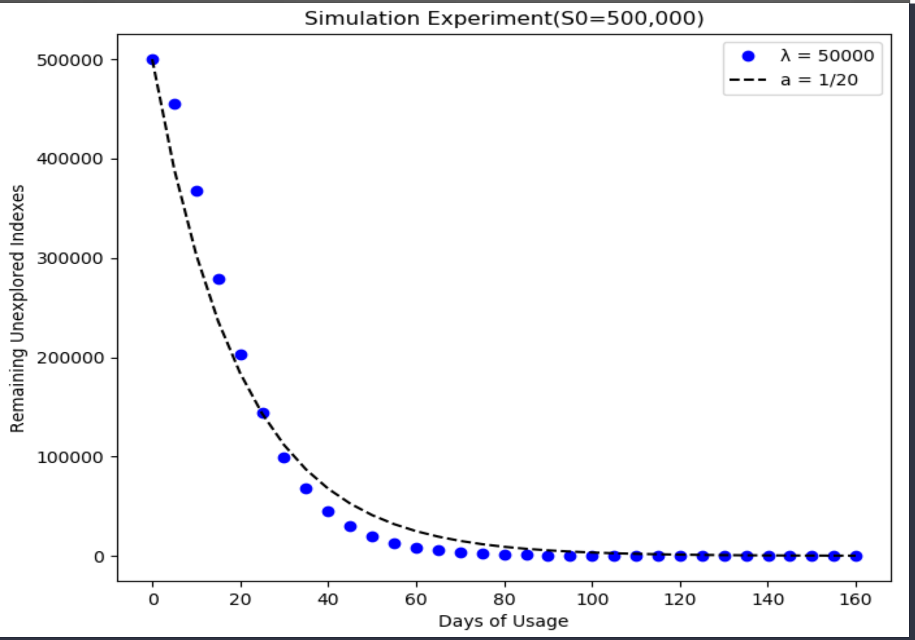

(a) St: S0 = 60,000; λ = 8,000, α = 1/15

(b) St: S0 = 500,000; λ = 50,000, α = 1/20

Comparison between Simulation and Theoretical Results

LEVERAGING BIG DATA

- Unlike text-based documents, the indexing of multimedia entities for deep knowledge retrieval is a learning and evolutionary process
- By applying reinforcement learning within a semi-Markov decision process framework, the subtle nuances of human perceptions and deep knowledge are captured and learnt for evolution
- Index evolution enables effective retrieval of multimedia objects
- Performance of the system automatically improves over time

THANK YOU!