A Unified Game-Theoretic Framework for Recommender Systems and Search Engines

ChengXiang ("Cheng") Zhai
Department of Computer Science
(Carl R. Woese Institute for Genomic Biology
School of Information Sciences
Department of Statistics)
University of Illinois at Urbana-Champaign

czhai@illinois.edu http://czhai.cs.illinois.edu/

The Data and Information Systems Laboratories

KDD 2021 OARS Workshop, Aug. 15, 2021

Recommender Systems vs. Search Engines

- Commonality
  - Both involve users, a collection of items that are potentially interesting to the users
  - Same general goal: Connect users with the right items at the right time
  - Both benefit from using machine learning
- Difference
  - User taking initiative ("pull") vs. system taking initiative ("push")
  - Query-driven vs. context-driven
  - To what extent users know what they want
  - Expectation of a user

Research in Recommender Systems and Search Engines

Topics of interest for RecSys 2021

- **Algorithm** scalability, performance, and implementations
- **Bias, fairness, bubbles and ethics** of recommender systems
- Case studies of **real-world implementations**
- **Conversational** and **natural language** recommender systems
- **Cross-domain** recommendation
- **Economic models** and consequences of recommender systems
- **Interfaces** for recommender systems
- Novel approaches to recommendation, including **voice, VR/AR, etc**.
- **Preference elicitation**
- **Privacy** and **security**
- **Socially- and context-aware** recommender systems
- Systems challenges such as **scalability, data quality, and performance**
- **User studies**

Topics of interest for SIGIR 2021

Search and ranking. Core IR algorithmic topics, including IR at scale.

Foundations and theory of IR. Theoretical or empirical contributions on technical or social aspects of IR.

Domain-specific applications. Research focusing on domain-specific IR challenges.

Content recommendation, analysis and classification. Recommender systems, rich content representations and content analysis: Filtering and recommendation (e.g., content-based filtering, collaborative filtering, recommender systems, recommendation algorithms, zero-query and implicit search, personalized recommendation)

Artificial Intelligence, semantics, and dialog. Research bridging AI and IR, especially toward deep semantics and dialog with intelligent agents.

Human factors and interfaces. User-centric aspects of IR including user interfaces, behavior modeling, privacy, interactive systems.

Evaluation. Research that focuses on the measurement and evaluation of IR systems.

Can we study recommender systems and search engines in a unified framework?

A Cooperative Game-Playing Framework (CGF) for Search and Recommendation

- {Search, Recommendation} = cooperative game-playing
- **Players**: Player 1= system; Player 2= user
- **Rules** of game:
  - Players take turns to make "moves"
  - First move = "user entering the query" (in search) or "system recommending information" (in recommendation)
  - User makes the last move (usually)
  - For each move of the user, the system makes a response move (shows an interaction interface), and vice versa
- **Objective**: deliver relevant/useful information to the user with minimum user effort & minimum operating cost for system

Unification of search and recommendation

Search as cooperative game-playing

(Finish a user task
with minimum effort)

User

A_1 : Enter a query

Which items
to view?

A_2 : View item

No

View more?

A_3 : Scroll down or click on
"Back"/"Next" button

(Help user finish a task
with minimum effort, minimum system cost)

System

Which information items to present?
How to present them?

R_1: results (i=1, 2, 3, …)

Which aspects/parts of the item
to show? How?

R_2: Item summary/preview

Interface Card

Recommendation as cooperative game-playing

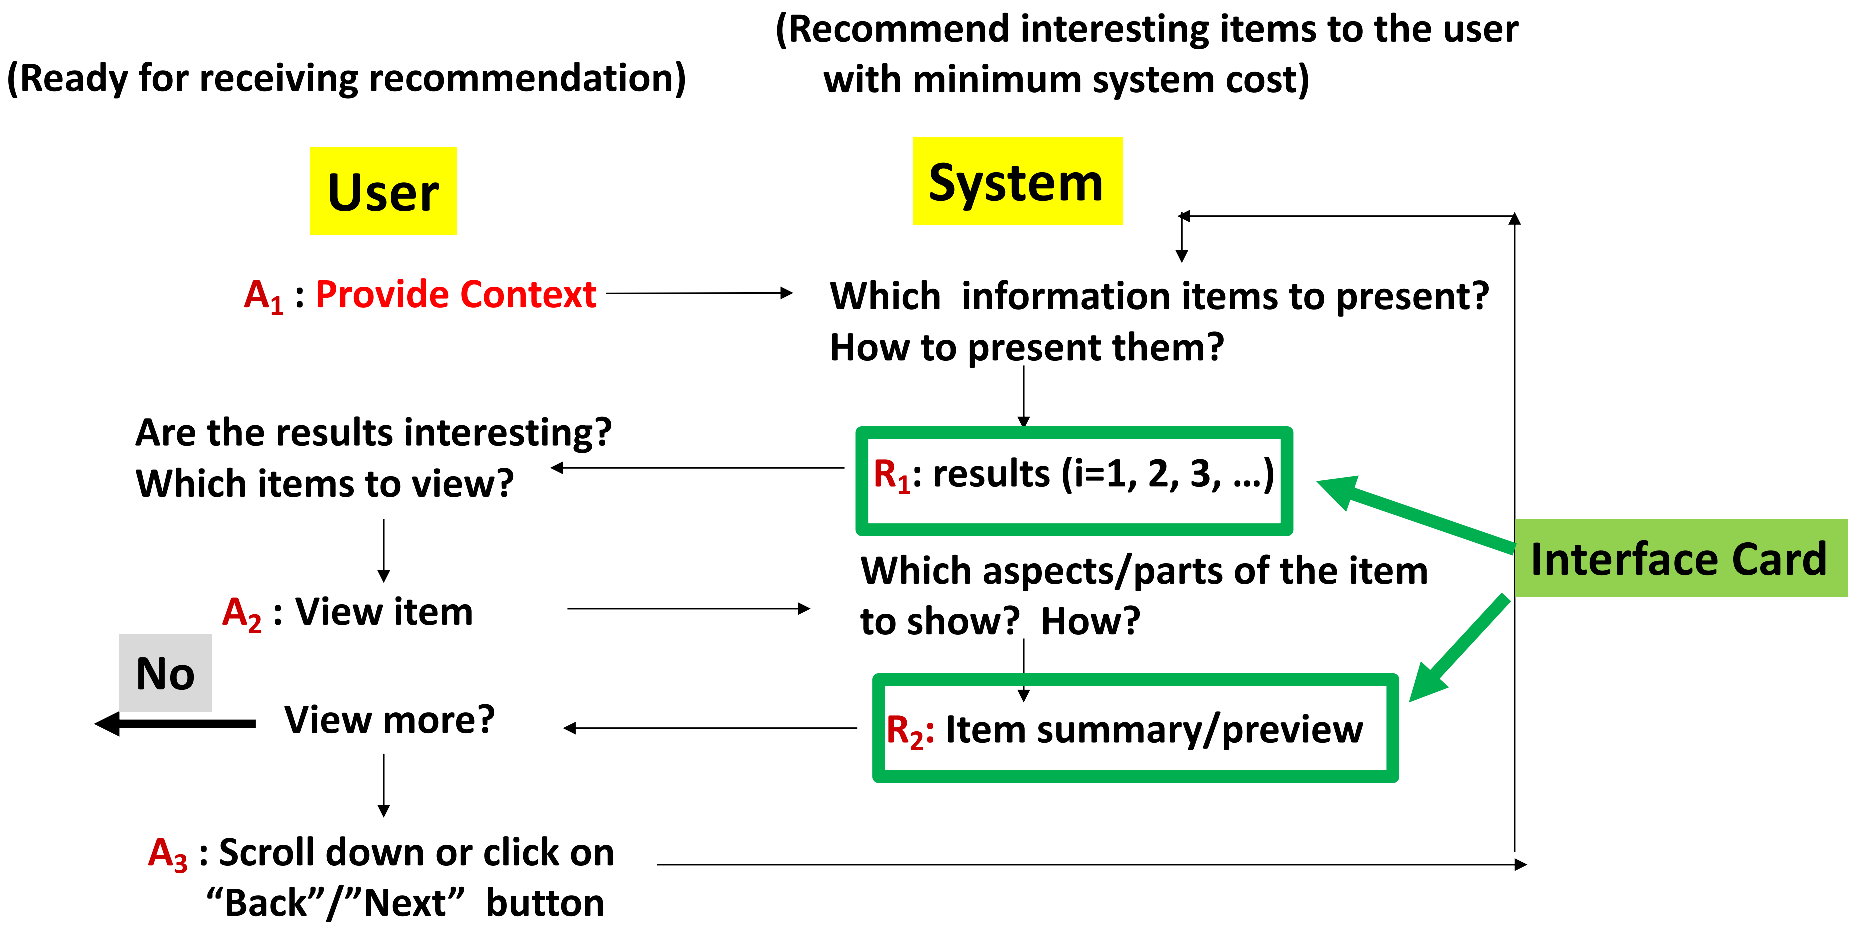

Major benefits of Cooperative Game-playing Framework (CGF)

- **General**
  - A formal framework to **integrate search and recommendation,** enabling study of user studies, evaluation, ranking models, and scalability in a unified framework
  - A general roadmap for identifying **unexplored important research topics in Interactive search & recommendation**
- **Specific**
  - Naturally optimize performance on an entire session instead of that on a single query or recommendation cycle (optimizing the chance of winning the <u>entire game)</u>
  - Optimize the collaboration of machines and users (maximizing <u>collective intelligence)</u> [Belkin 96]
  - Emphasize the two-way communications between a user and a system
  - …

4 Key Elements of the Game Framework (4 C's)

- **Collaboration**: Optimization of the collaboration (or **combined intelligence**, combined performance) of a user and a system
  - User knows well about what's useful, but doesn't know the whole information space
  - System "sees" the whole information space, but doesn't know which is most useful
- **Communication**: Optimization of the **two-way communications** between a user and a system
  - Communication of the shared goal and plan
  - Explanation of both user actions and system responses
- **Cognition:** Optimization of cognition for user (**bridge the cognition gap**) and system (**machine learning**)
  - Modeling of knowledge state and helping users learn during the interaction [Collins-Thompson et al. 17]
  - Helping system learn knowledge about a user's preferences and needs
- **Cost:** Optimization of system operation cost
  - Modeling operation cost and providing **cost-effective responses**

Formalization of the Cooperative Game

Situation **S**

History $\mathbf{H=\{(A_i,R_i)\}}$
$i=1, \ldots, t-1$

Given S, U, C, A_t , and H, choose the best R_t from all possible responses to A_t

User **U**: A_1 A_2 … … A_{t-1} $\mathbf{A_t}$

System: R_1 R_2 … … R_{t-1} $\mathbf{R_t}$ **=?**

Query="light laptop"

Click on "Next" button

Any action

The best ranking for the query

The best ranking of unseen items (recommendation)

Best interface!

C

Info Item Collection

$R_t \in r(A_t)$

All possible rankings of items in C

All possible rankings of unseen items

All possible interfaces

Bayesian Decision Theory for Interactive Recommendation & Search (IRS)

Observed

Situation:	S
User:	U
Interaction history:	H
Current user action:	A_t
Document collection:	C

User Model

$M=(K, \theta_U, B, T, \ldots)$

Knowledge State (seen items, readability level, …)

Information need

Browsing behavior

Task

All possible responses:

$r(A_t)=\{r_1, \ldots, r_n\}$

$L(r_i, M, S)$

Loss Function

Optimal response: r* (minimum loss)

$$R_t = \arg\min_{r \in r(A_t)} \int_M L(r, \underline{M}, S)\, p(\underline{M} \mid \underline{U, H, A_t, C, S})\, dM$$

Bayes risk **Inferred** **Observed**

ChengXiang Zhai. Towards a game-theoretic framework for text data retrieval, IEEE Data Eng. Bull. 39(3): 51-62 (2016).

An extension of risk minimization [Zhai & Lafferty 06, Shen et al. 05]

Simplification of Computation

- Approximate the Bayes risk (posterior mode)

$$R_t = \arg\min_{r \in r(A_t)} \boxed{\int_M} L(r, M, S) p(M \mid U, H, A_t, C, S) dM$$
$$\approx \arg\min_{r \in r(A_t)} L(r, M^*, S) \underline{p(M^* \mid U, H, A_t, C, S)}$$
$$= \arg\min_{r \in r(A_t)} L(r, M^*, S)$$
$$\text{where} \quad M^* = \arg\max_M p(M \mid U, H, A_t, C, S)$$

- Two-step procedure
  - Step 1: Compute an updated user model M* based on the currently available information
  - Step 2: Given M*, choose an optimal response to minimize the loss function

Optimal Interactive Recommendation & Search (IRS)

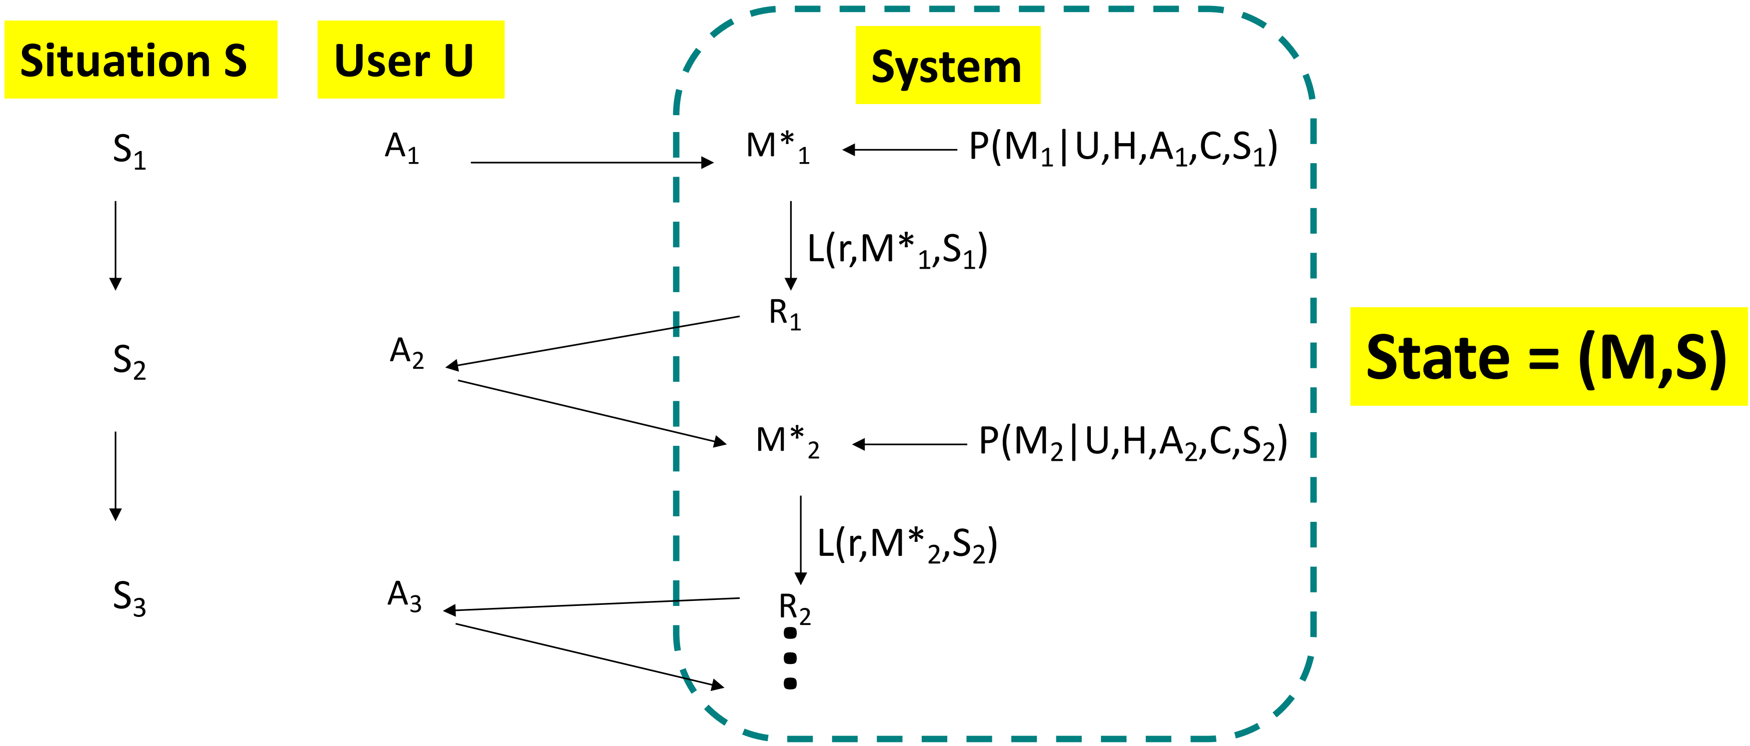

System's decision process can be modeled by a Partially Observable Markov Decision Process (POMDP) with (M, S) as State

Duality of User & System Decision Making

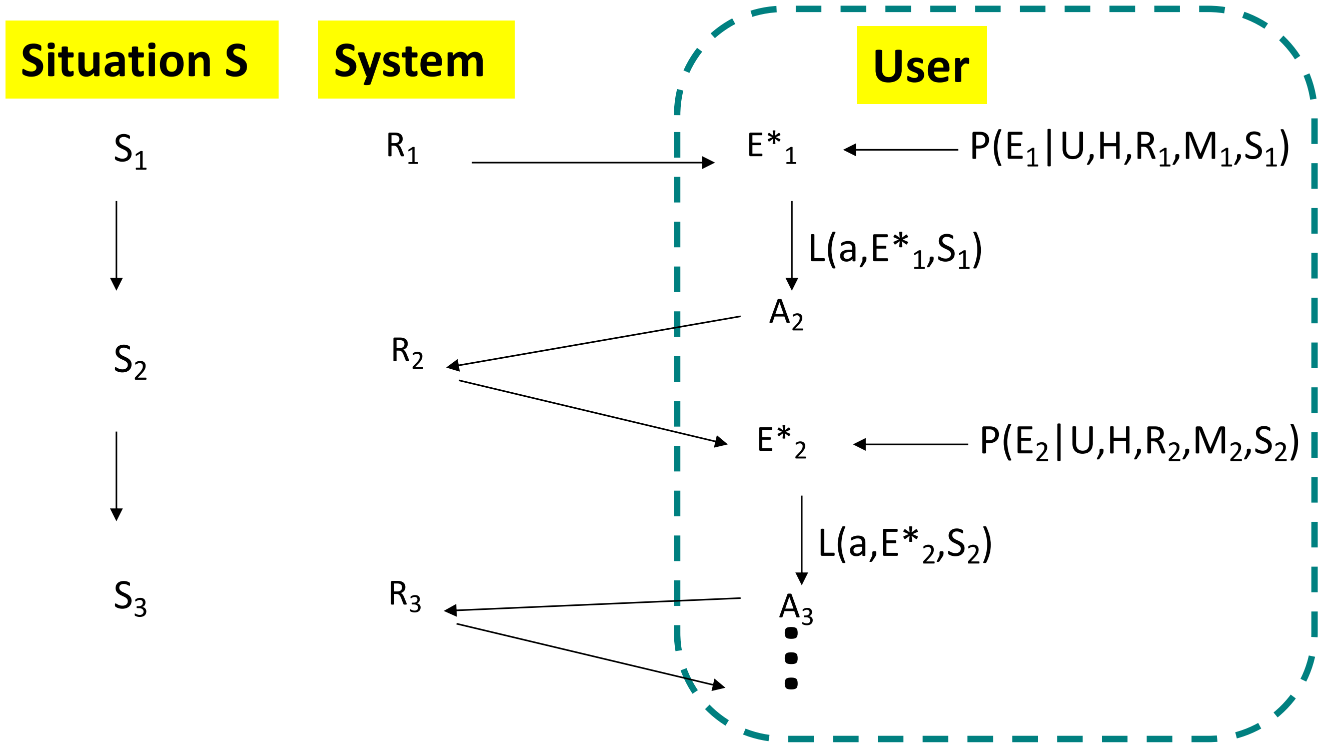

Interaction environment
E= (Collection, System)

State = (E,S)

User's decision process (behavior) can be modeled by a POMDP as well with (E,S) as State

Simulation of user agent for evaluating IRS [Zhang et al. 17]

Instantiation of the **Cooperative Game Framework(CGF)**

- **Situation S**: can include time, location, and other environmental factors that are relevant to a task
- **Information/Item Collection C:** naturally available in any application
- **User U**: can include any information we know about a user (or group)
- **User interaction history H:** naturally accumulated over time
- **User Actions** and **System Responses R(A)**: all interfaces (moves of the game)
- **Loss Function L(R,M,S):** captures the objective of the game
- **User Model M**: can include everything that we can infer about a user relevant to deciding how to respond to a user's action
- **Inference of User Model P(M|U, H, At, C,S):** capture system's belief about user model M

Instantiation of IR Game: Moves (Interface Design)

- User moves: Interactions can be modeled at different levels
  - Low level: keyboard input, mouse clicking & movement, eye-tracking
  - Medium level: query input, result examination, next page button
  - High level: each query session as one "move" of a user
- System moves: can be enriched via sophisticated interfaces, e.g.,
  - User action = "input one character" in the query: System response = query completion
  - User action = "scrolling down": System response = adaptive summary
  - User action = "entering a query": System response = recommending related queries
  - User action = "entering a query": System response = ask a clarification question

Example of new moves (new interface): Explanatory Feedback

- Optimize combined intelligence →
  - Leverage human intelligence to help search engines
- Add new “moves” to allow a user to help a search engine with minimum effort
- Explanatory feedback
  - I want documents similar to this one except for not matching “X” (user typing in “X”)
  - I want documents similar to this one, but also further matching “Y” (user typing in “Y”)
  - …

Instantiation of IR Game: User Model M

- **M = formal user model capturing essential knowledge about a user's state for optimizing system moves**
  - Essential component: θ_U = user's current information need
  - K = knowledge state (seen items)
  - Readability level
  - T= task
  - Patience-level
  - B= User behavior
  - Potentially include all findings from user studies!
- **An attempt to <u>formalize</u> existing models such as**
  - Anomalous State of Knowledge (ASK) [Belkin 80, Belkin et al. 82]
  - Cognitive IR Theory [Ingwersen 96, Ingwersen & Järvelin 06]

Instantiation of IR Game: Inference of User Model

- P(M|U, H, At, C,S) = system's current belief about user model M
  - Enables inference of the formal user model M based on everything the system has available so far about the user and his/her interactions
- Instantiation can be based on
  - Findings from user studies, and
  - Machine learning using user interaction log data for training
- Much work has been done on estimating/updating the information need θ_U and clicking behavior (e.g., implicit feedback [Joachims et al. 05, Shen et al. 05], intent understanding [Liu et al. 14], and many click models [Chuklin et al. 15, Liu et al. 17])
- Some work on inferring/updating other variables about the user, e.g.,
  - reading level [Collins-Thompson et al. 11]
  - modeling decision point [Thomas et al. 14]
- Similar work in the recommender system context

Instantiation of IR Game: Loss Function

- L(Rt ,M,S): loss function **combines** measures of
  - **Utility of Rt** for a user modeled as M to **finish the task** in situation S
  - **Effort of a user** modeled as M in situation S
  - **Cost of system** performing Rt (connected with **efficiency of IR systems** [Witten et al. 99])
- Tradeoff varies across users and situations
- **Utility of Rt** is a **sum** of
  - **ImmediateUtility(Rt)** and
  - **FutureUtilityFromInteraction(Rt)**, which depends on user's interaction behavior

Instantiation of IR Game: Loss Function (cont.)

- Formalization of utility depends on research on evaluation, task modeling, and user behavior modeling
- Traditional evaluation measures tend to use
  - Very simple user behavior model (sequential browsing)
  - Straightforward combination of effort and utility
- They need to be extended to incorporate **more sophisticated user behavior models** (e.g., [de Vries et al. 04] , [Smucker & Clarke 12], [Baskaya et al. 13])
- Much progress has been made recently on **incorporating click models** (simple user interaction models) into a loss function for **learning to rank or recommend** (e.g., online learning to rank [Hofmann et al. 11, Wang et al. 19] , dynamic IR [Yang et al. 06], recommendation [Zhao et al. 08], sequential browsing [Wei et al. 17])

Example of Instantiation:
Interface Card Model (ICM) [Zhang & Zhai 15, Zhang & Zhai 16]

How to optimize the interface design?

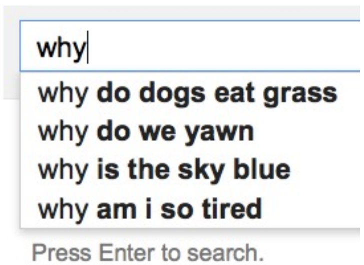

… or a combination of some of these?

How to allocate screen space among different blocks?

Yinan Zhang, ChengXiang Zhai, Information Retrieval as Card Playing: A Formal Model for Optimizing Interactive Retrieval Interface, *Proceedings of ACM SIGIR 2015.*
Yinan Zhang and Chengxiang Zhai. 2016. A Sequential Decision Formulation of the Interface Card Model for Interactive IR. In *Proceedings of ACM SIGIR 2016.*

Optimal User Interface = Optimal "Card Playing"

- In each interaction *lap*
- … facing an (evolving) interaction *context*
- … the system tries to play a *card*
- … that optimizes the user's *expected surplus*
- … based on the user's *action model* and *reward* / *cost* estimates
- … given all the *constraints* on card

Example of interface optimization

Context c^t

After a user clicks on "Colleges & Universities", which interface card q^t to show?

If showing card q^t

If user action a^{t+1}= view content

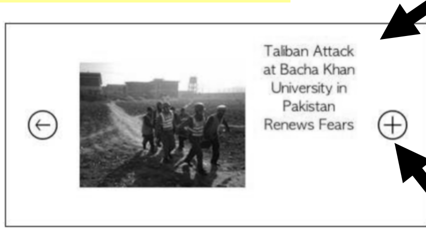

surplus for a^{t+1} :

$u(a^{t+1} \mid q^t, c^t)$= gain - cost

If user action a^{t+1}= "see more"?

A different card

If user action a^{t+1}= "navigate"?

Expected surplus of an interface card: $E(u^t|q^t,c^t)$

$E(u^t \mid q^t =$ 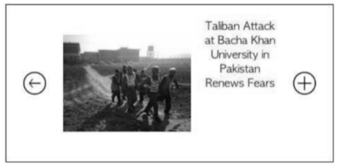 $, c^t)$

$= p($ a^t = "view content" $\mid c^t, q^t) \times u($  $\mid c^t, q^t)$

$u($  $\mid c^t, q^t) =$ Gain () – Cost(Viewing)

Gain ()=Relevance()

$+ p($ a^t = "see more" $\mid c^t, q^t) \times u($  $\mid c^t, q^t)$ + …

Depends on the next interface card q^{t+1}

Expected surplus of an interface card: $E(u^t|q^t,c^t)$

$E(u^t \mid q^t =$ 

$, c^t)$

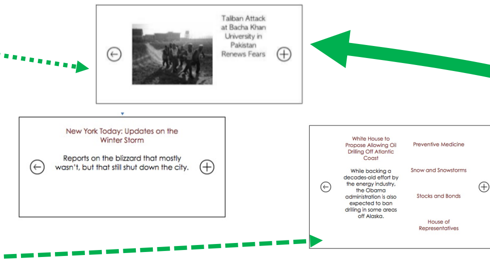

$= p(\, a^t = \text{"left-top tag"} \mid c^t, q^t) \times u($ 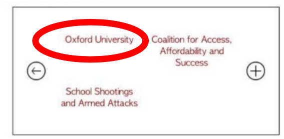

$\mid c^t, q^t)$

$+\; p(\, a^t = \text{"right-top tag"} \mid c^t, q^t) \times u($ 

$\mid c^t, q^t)$

$+\; p(\, a^t = \text{"left-bottom tag"} \mid c^t, q^t) \times u($ 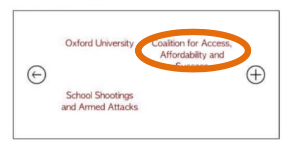

$\mid c^t, q^t)$

$+ \ldots$

ICM: Formal Definition

$$
\begin{aligned}
\underset{q^t}{\text{maximize}} \quad & E(u^t|c^t, q^t) \\
& = \sum_{a^{t+1} \in \mathcal{A}(q^t)} p(a^{t+1}|c^t, q^t)\, u(a^{t+1}|c^t, q^t) \\
\text{subject to} \quad & f_c^t(q^t) \leq 0
\end{aligned}
$$

The Data and Information Systems Laboratories at The University of Illinois at Urbana-Champaign
Large Scale Information Management and Mining

Interface card

$$\begin{aligned} \underset{\boxed{q^t}}{\text{maximize}} \quad & E(u^t|c^t,q^t) \\ &= \sum_{a^{t+1}\in\mathcal{A}(q^t)} p(a^{t+1}|c^t,q^t)\,u(a^{t+1}|c^t,q^t) \\ \text{subject to} \quad & f_c^t(q^t)\le 0 \end{aligned}$$

Context

$$
\begin{aligned}
\underset{q^t}{\text{maximize}} \quad & E(u^t \mid \boxed{c^t}, q^t) \\
& = \sum_{a^{t+1} \in \mathcal{A}(q^t)} p(a^{t+1} \mid \boxed{c^t}, q^t)\, u(a^{t+1} \mid \boxed{c^t}, q^t) \\
\text{subject to} \quad & f_c^t(q^t) \le 0
\end{aligned}
$$

$$\begin{aligned}
\underset{q^t}{\text{maximize}} \quad & E(u^t | c^t, q^t) \\
& = \sum_{a^{t+1} \in \boxed{\mathcal{A}(q^t)}} p(a^{t+1} | c^t, q^t)\, u(a^{t+1} | c^t, q^t) \\
\text{subject to} \quad & f_c^t(q^t) \leq 0
\end{aligned}$$

Action set

$$\begin{aligned} \underset{q^t}{\text{maximize}} \quad & E(u^t \mid c^t, q^t) \\ = & \sum_{a^{t+1} \in \mathcal{A}(q^t)} \boxed{p(a^{t+1} \mid c^t, q^t)}\, u(a^{t+1} \mid c^t, q^t) \\ \text{subject to} \quad & f_c^t(q^t) \leq 0 \end{aligned}$$

Action model

$$\overbrace{r(a^{t+1}|c^t,q^t)}^{\text{Reward}} - \overbrace{s(a^{t+1}|c^t,q^t)}^{\text{Cost}}$$

$$\begin{aligned} \underset{q^t}{\text{maximize}} \quad & E(u^t|c^t,q^t) \\ & = \sum_{a^{t+1}\in\mathcal{A}(q^t)} p(a^{t+1}|c^t,q^t)\,\boxed{u(a^{t+1}|c^t,q^t)} \\ \text{subject to} \quad & f_c^t(q^t) \le 0 \end{aligned}$$

Action surplus

Expected surplus

$$\underset{q^t}{\text{maximize}} \quad \boxed{E(u^t|c^t, q^t)}$$

$$= \sum_{a^{t+1} \in \mathcal{A}(q^t)} p(a^{t+1}|c^t, q^t)\, u(a^{t+1}|c^t, q^t)$$

$$\text{subject to} \quad f_c^t(q^t) \leq 0$$

The Data and Information Systems Laboratories at The University of Illinois at Urbana-Champaign
Large Scale Information Management and Mining

$$
\begin{aligned}
\underset{q^t}{\text{maximize}} \quad & E(u^t | c^t, q^t) \\
& = \sum_{a^{t+1} \in \mathcal{A}(q^t)} p(a^{t+1} | c^t, q^t)\, u(a^{t+1} | c^t, q^t) \\
\text{subject to} \quad & \boxed{f_c^t(q^t) \leq 0}
\end{aligned}
$$

Constraint(s)

Refinements/Instantiations of ICM

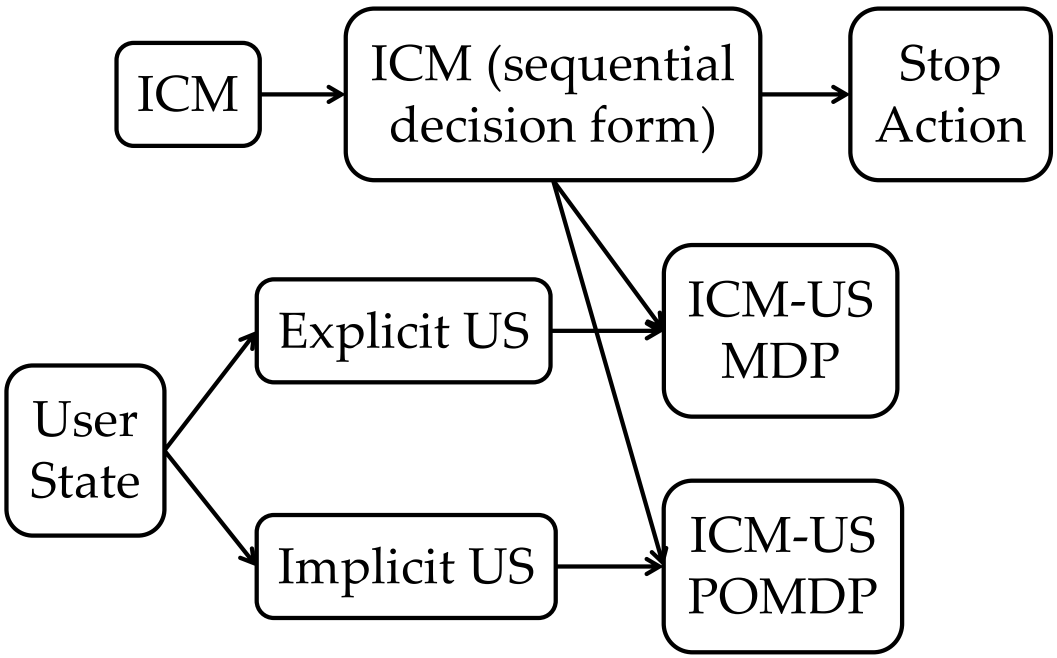

User study experiments

- Setting
  - Prototype interfaces for New York Times
  - Articles as items and keywords as tags
  - Two sizes: a medium sized one and a small one
- Comparison
  - # Interaction rounds to reach item of interest
  - We *automatically* optimize the interface layout
  - Compare with pre-designed static interfaces

Medium sized screen

Smaller screen

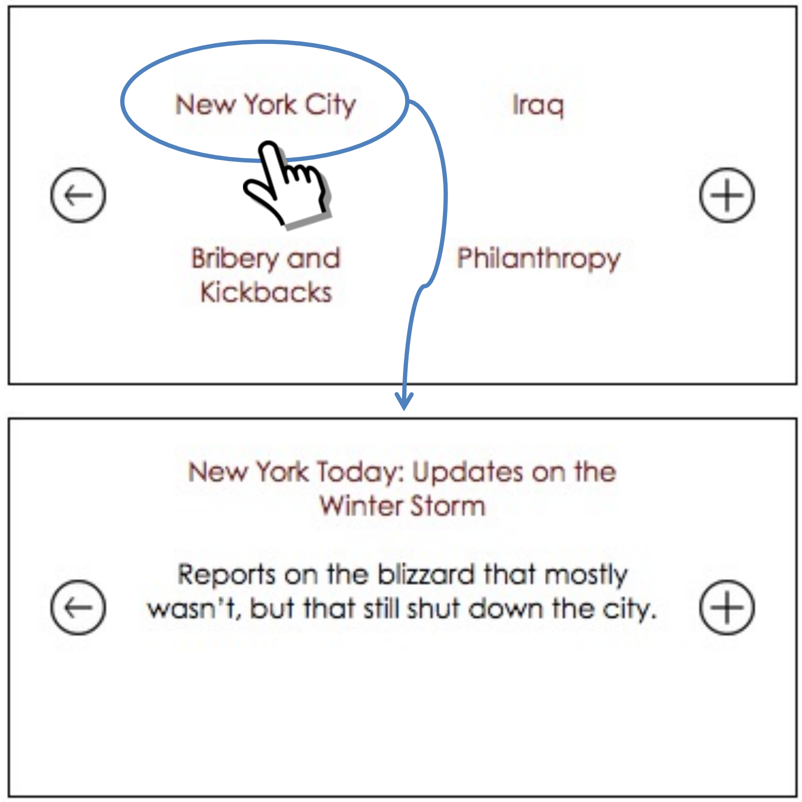

Interaction round comparison

More beneficial when screen is small and number of items large

Table 1: Significance levels of comparison results.

Card size	Item set size	Valid sample size	P-value
Small	20	19	**0.004753**
Small	50	23	**0.0003546**
Medium	20	18	0.09183
Medium	50	20	**0.01097**

CGF & Diversification:
3 Different Reasons for Diversification

1. Redundancy reduction → reduce user effort

2. Diverse information needs (e.g., overview, subtopic retrieval) → increase the immediate utility

3. Active relevance feedback → increase future utility

Capturing diversification with different loss functions

1. **Redundancy reduction**: Loss function includes a redundancy measure
   – Special case: list presentation + MMR [Zhai et al. 03]
2. **Diverse information needs**: loss function defined on latent topics
   – Special case: PLSA/LDA + topic retrieval [Zhai 02]
3. **Active relevance feedback**: loss function considers both relevance and benefit for feedback (online learning to rank, dynamic IR)
   – Special case: hard queries + feedback only [Shen & Zhai 05]

Whole Session/Page Optimization

- Special case of the Cooperative Game framework: Objective function includes expectation over future interactions
  – Whole session optimization: consider all future interactions with the user
  – Whole page optimization: consider all possible actions a user can take on the page
  – Both directly captured by the Interface Card Model
- Algorithms are generally based on multi-armed bandits and reinforcement learning and aim to optimize the tradeoff between exploitation (optimizing current benefit) and exploration (optimizing future benefit), leading to diversification of results
- The empirical benefit so far has been mostly optimizing the ranking of results, thus no "visible" impact on the interface design
- Exception: **Whole page optimization using ML** [Wang et al. 16]

Yue Wang, Dawei Yin, Luo Jie, Pengyuan Wang, Makoto Yamada, Yi Chang, and Qiaozhu Mei. 2016. Beyond Ranking: Optimizing Whole-Page Presentation. In *Proceedings of WSDM 2016.*

How to evaluate an Interactive Recommender & Search (IRS) system?

- Problem with using A/B test: Not reusable, not reproducible
- Cranfield evaluation methodology has the following benefit, but it cannot be used to evaluate IRS
  - Reusable test collection ➔ Can be reused and ensure fairness in comparison
  - Facilitate component testing
- How can we make a fair comparison of multiple IRS systems using reproducible experiments?
- Must control the users ➔ Using user simulators!
- SIGIR'21 has a workshop on user simulation for IR evaluation (https://sim4ir.org/)

IR evaluation as simulation [Zhang et al. 17, Pääkkönen et al. 17]

- Simulation provides a general way to evaluate IR systems
  - General formal framework [Zhang et al. 17]: Cranfield evaluation as a special instantiation case (simulating "naïve" users)
- Benefit
  - "Controlled" user study for reproducibility
  - "Generalized" Cranfield test for sophisticated IR interface
- Feasibility shown in some existing work (e.g., [Liu et al. 07], [Carterette et al. 15], [Zhang et al. 17], [Pääkkönen et al. 17])

Search simulation framework [Zhang et al. 17]

- Top level components
  - System: S
  - User / simulator: U
  - Task: T
  - Interaction sequence: I
- Metrics
  - Interaction reward and cost: R(I,T,U,S) and C(I,T,U,S)
  - Simulator reward and cost: R(T,U,S) and C(T,U,S)
    - Expectation w.r.t. p(I|T,U,S)

Classical IR simulator

- Task: find (all) relevant documents
- Interface card: document (snippet)
- User action: click / skip (and read next) / stop
  - User always clicks a relevant document
  - User may skip or stop at a non-relevant document
- Lap reward: 1 / 0 for relevant / non-relevant doc
  - Cumulative reward: # relevant docs
- Lap cost: 1 for each doc
  - Cumulative cost: # docs (the simulator scanned through)
- User state: cumulative reward and cost

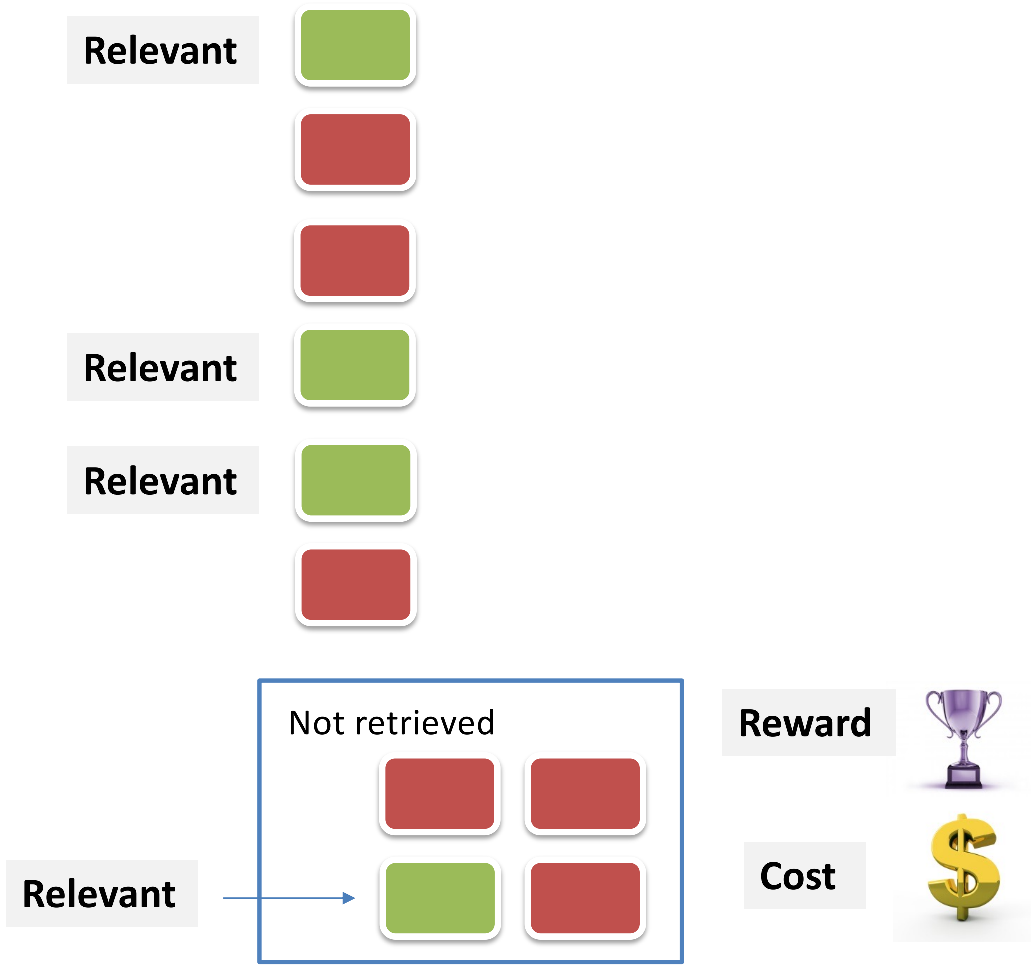
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Cost

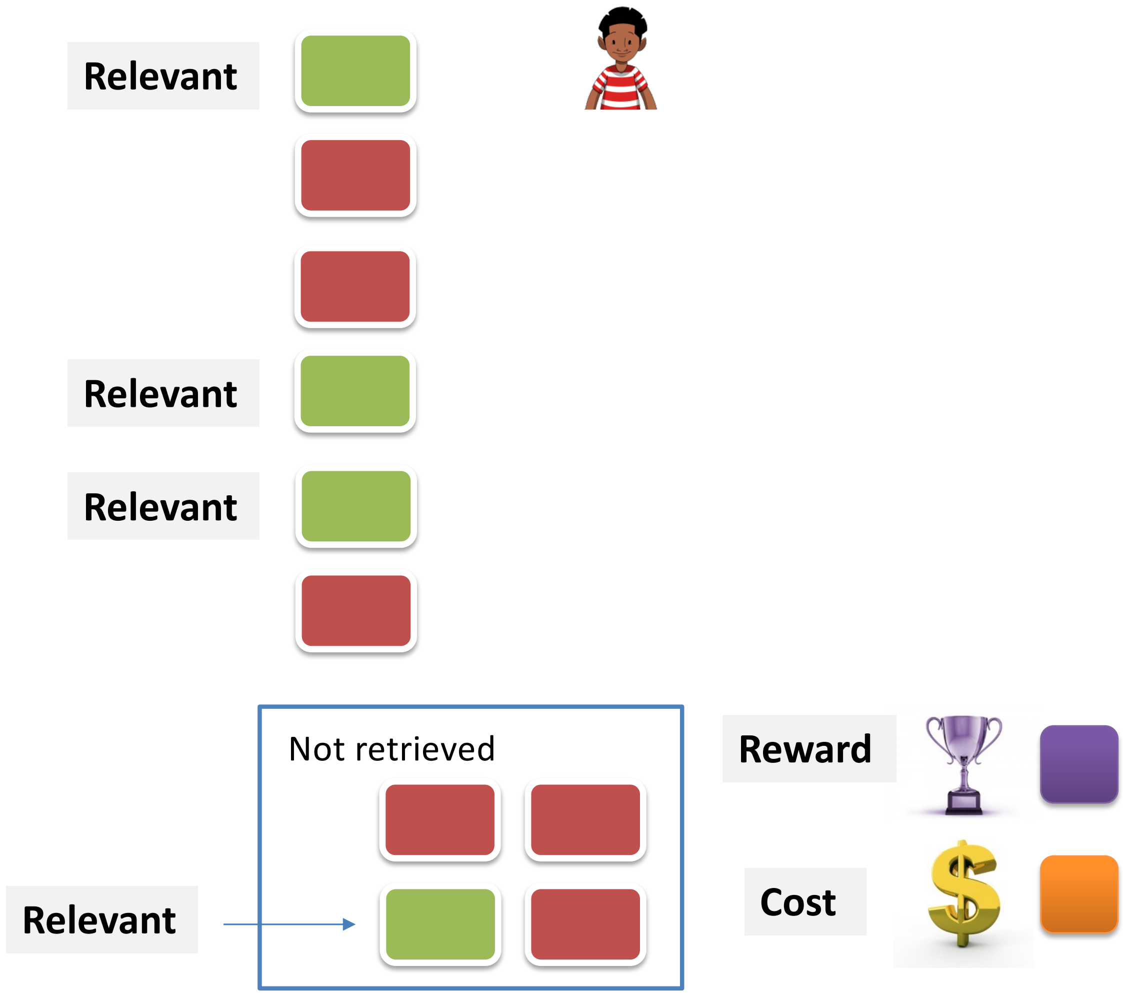
Relevant
Relevant
Relevant
Not retrieved
Reward
Cost
Relevant

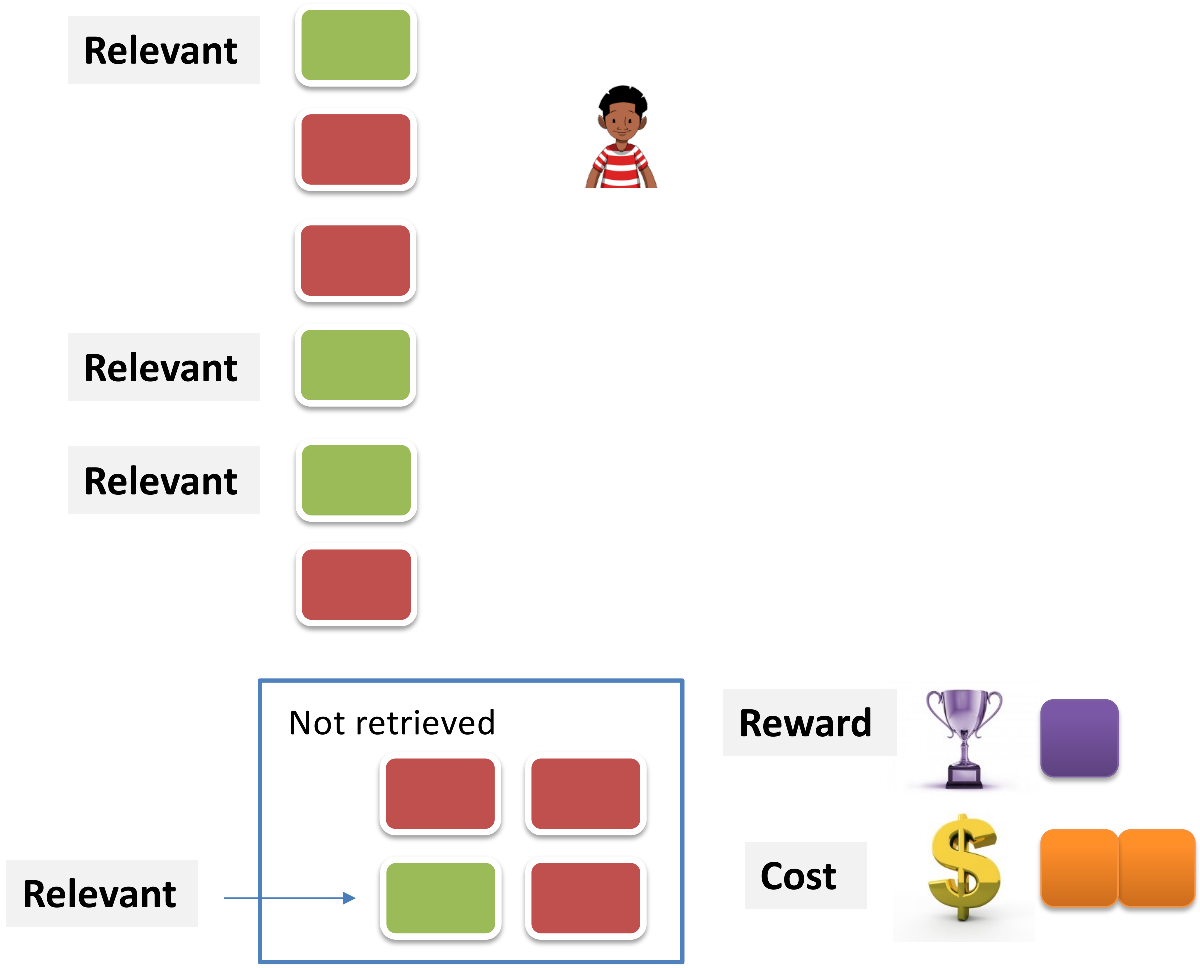
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Cost

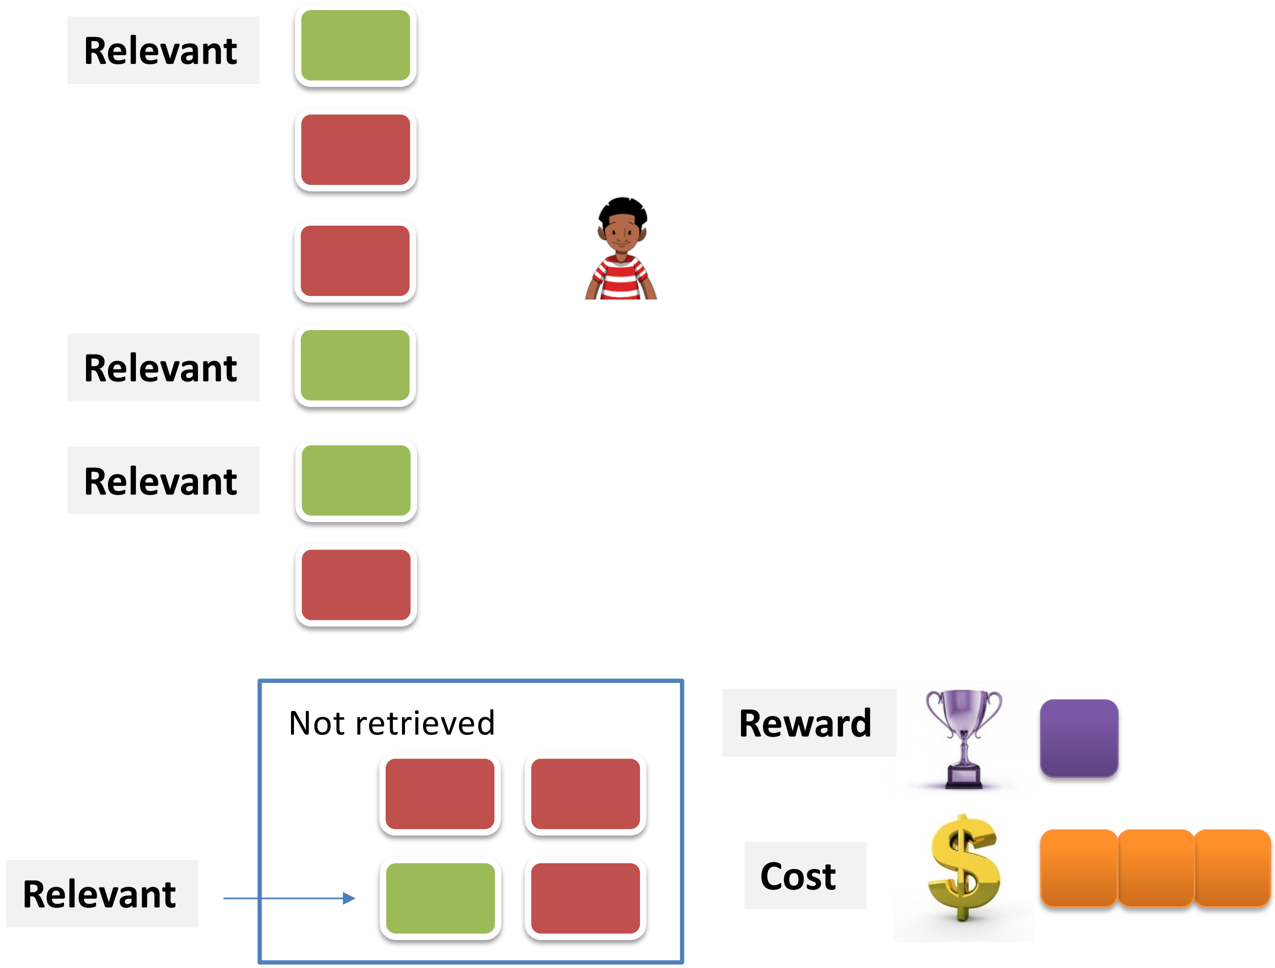
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Cost

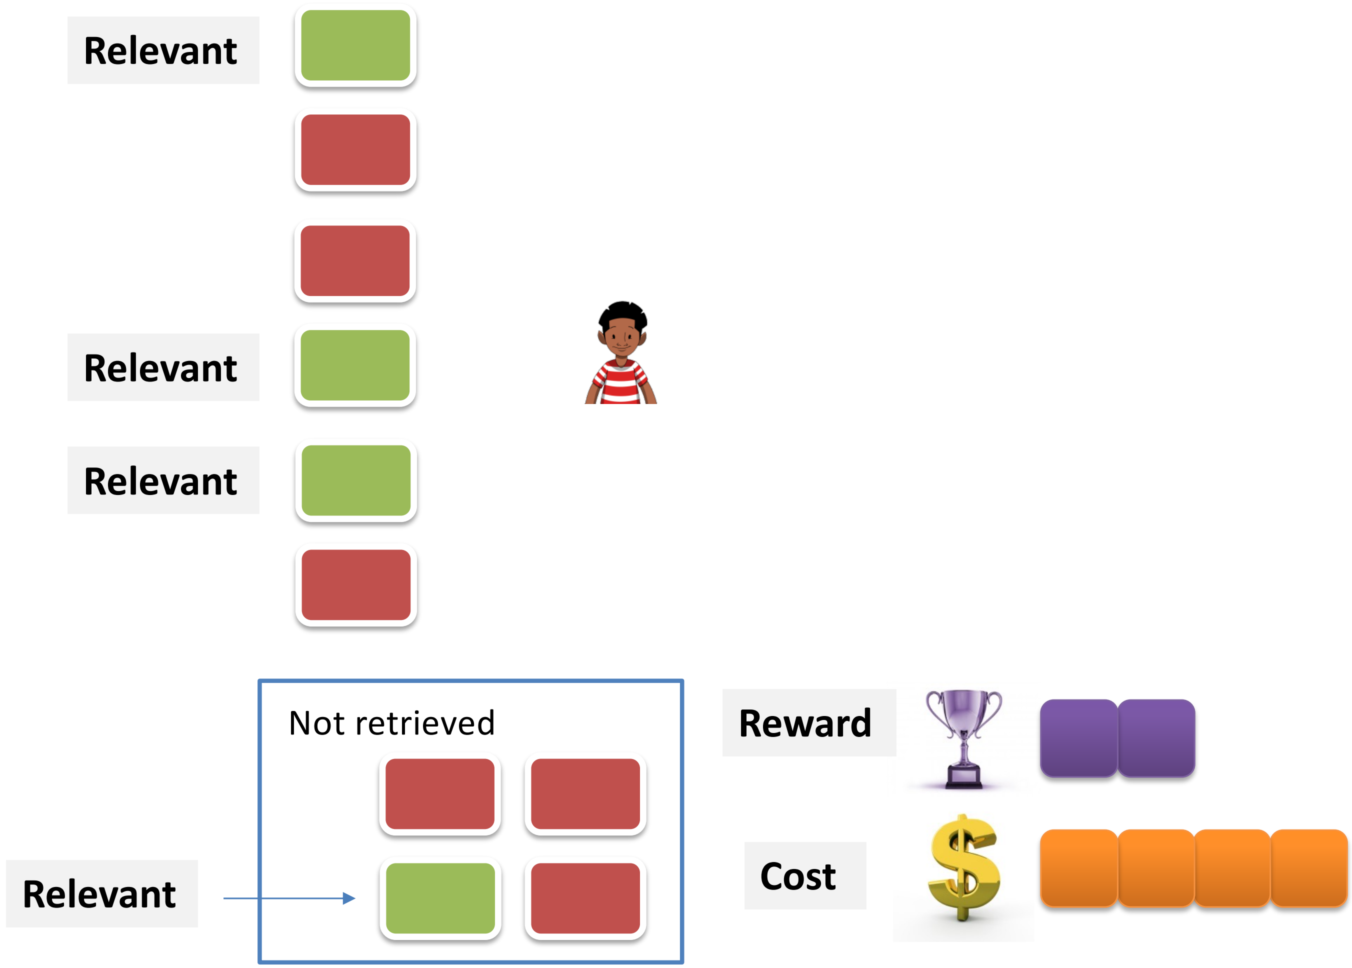
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Cost

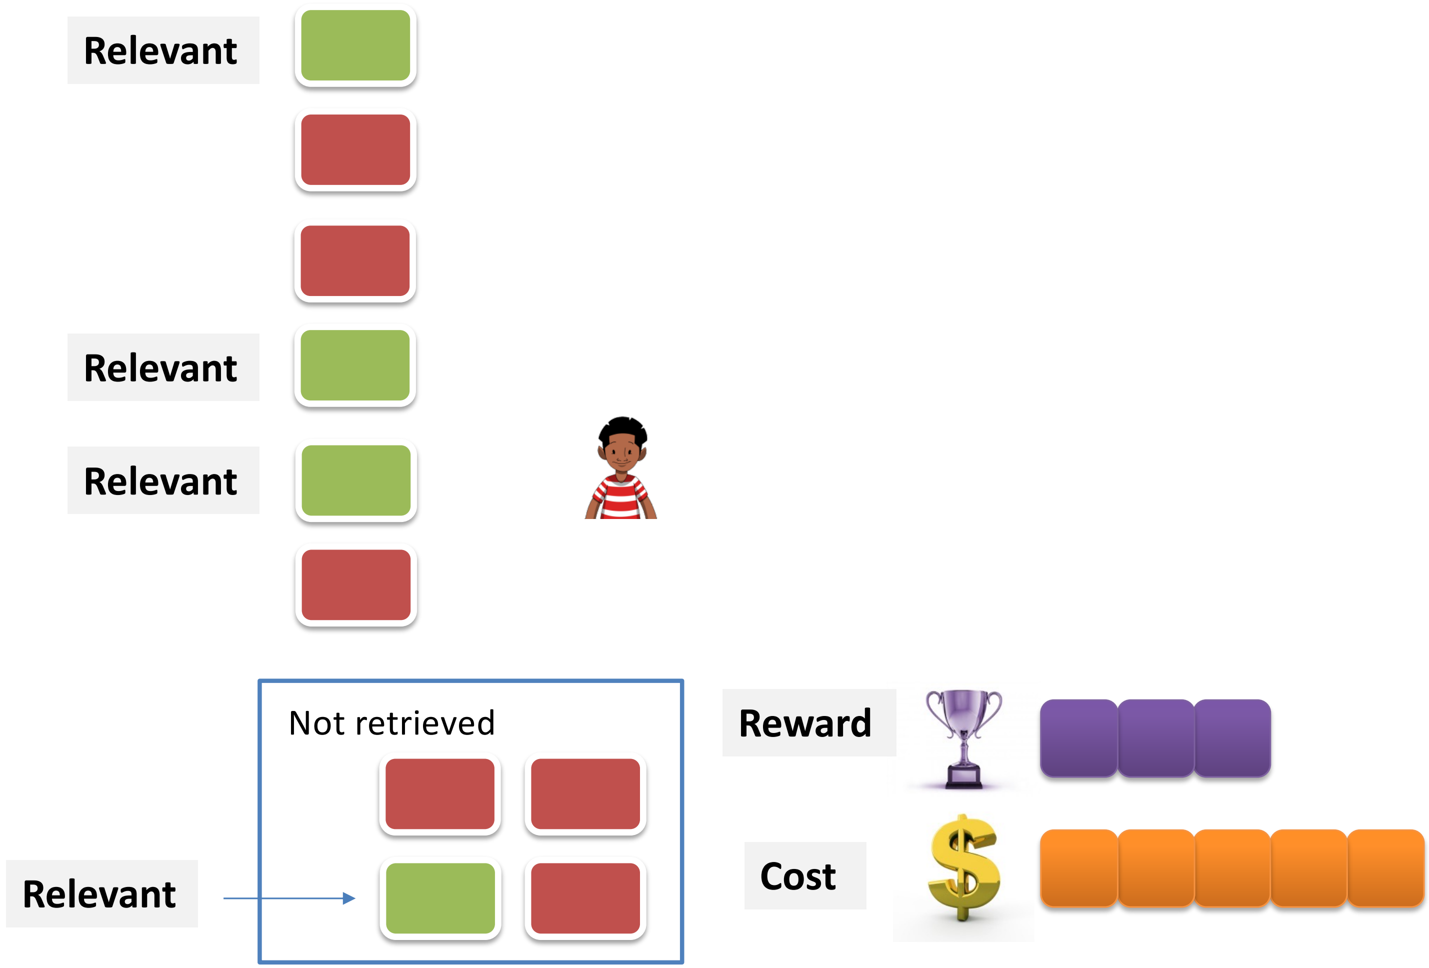
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Cost

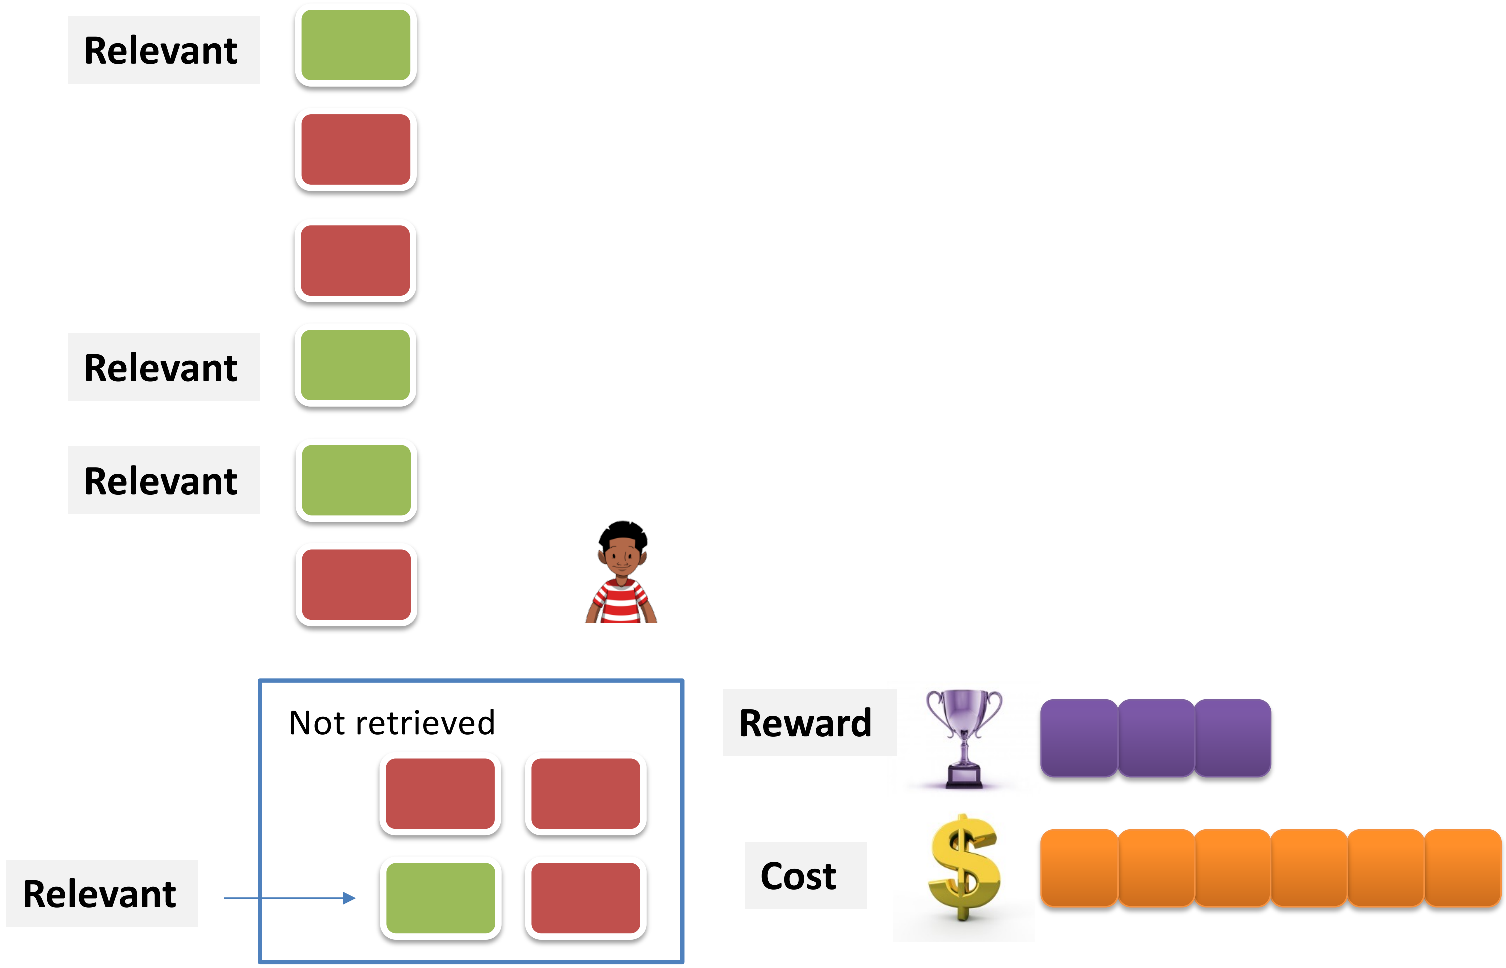
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Cost

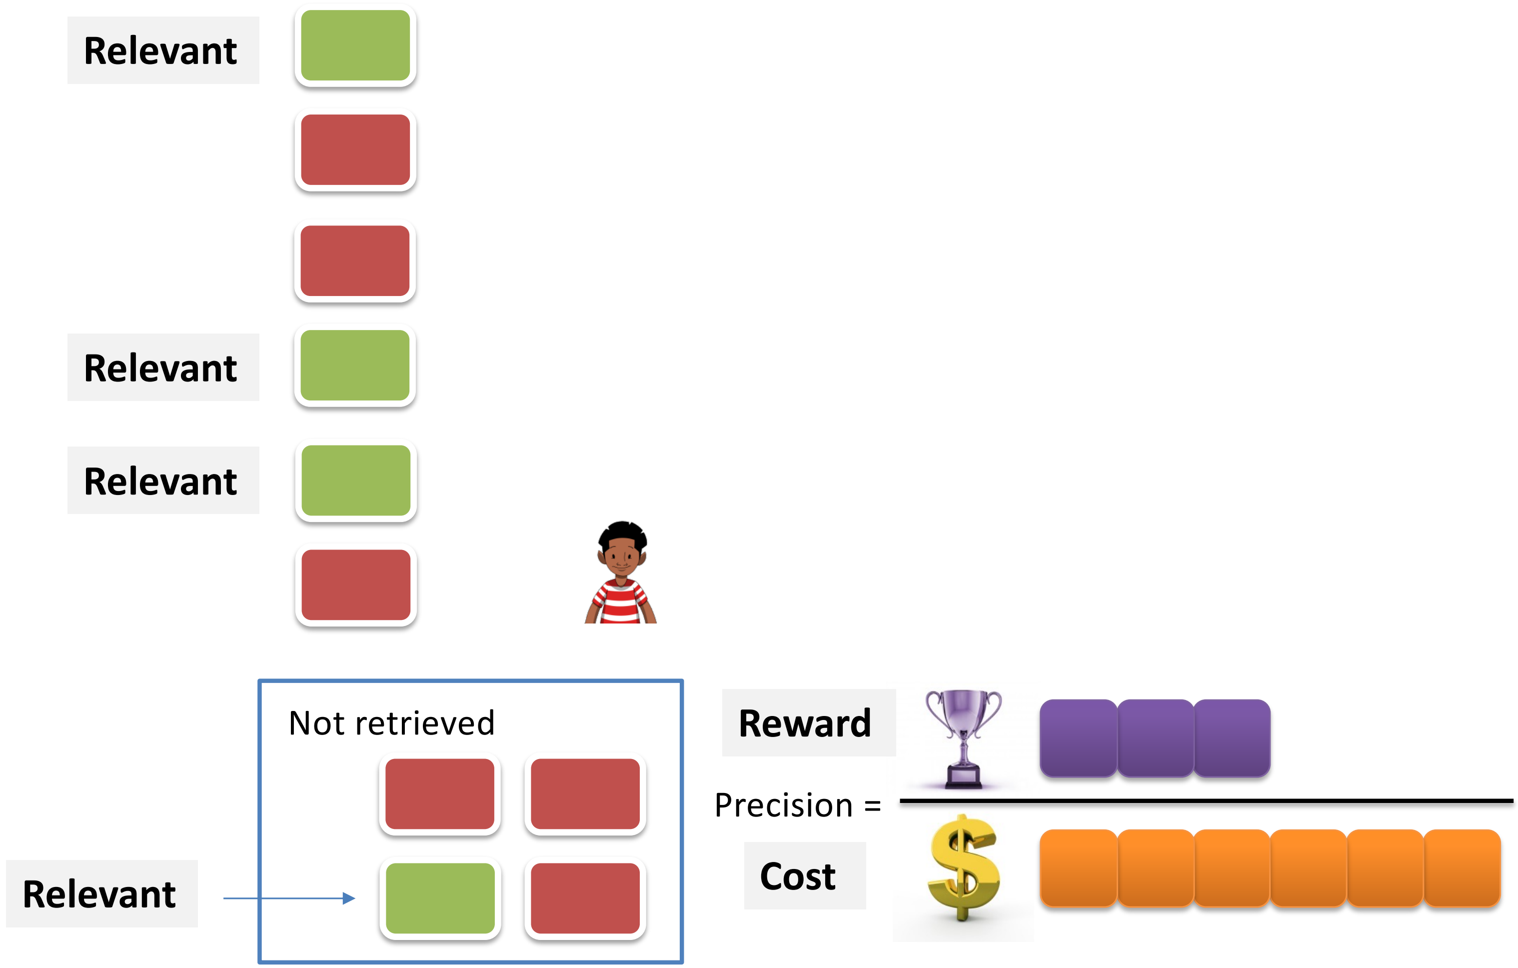
Relevant
Relevant
Relevant
Not retrieved
Relevant
Reward
Precision =
Cost

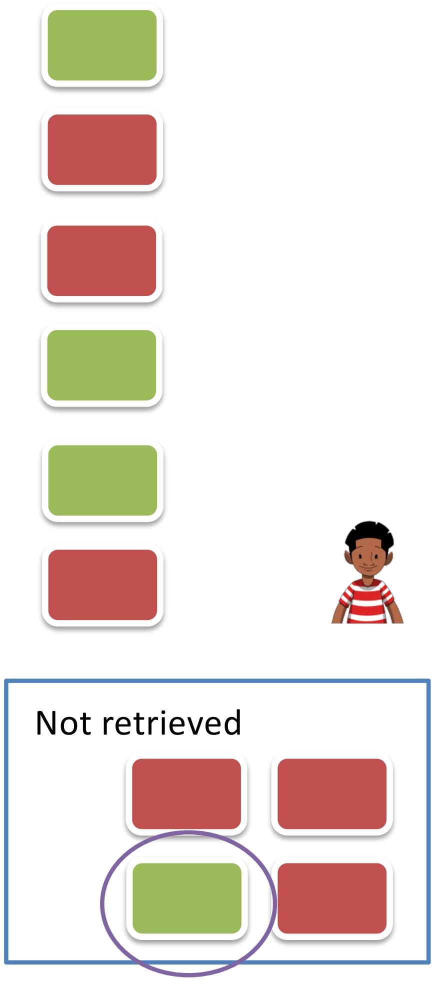
Not retrieved

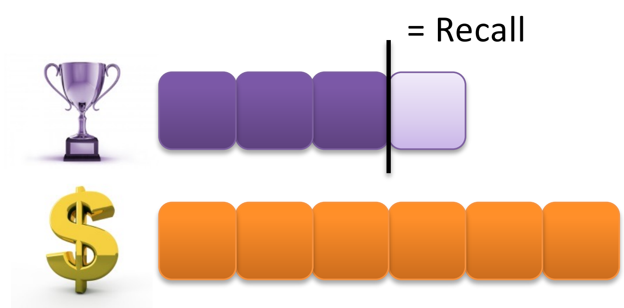
= Recall

Mean Average Precision (MAP)

- Variable-recall simulator
  - Classical IR simulator with task of finding N' relevant documents (N' between 1 and N)
  - Stops and only stops when the task is finished
- Average Precision (AP)
  - Average R(I,T,U,S) / C(I,T,U,S) across N variable-recall simulators with N' ranging from 1 to N respectively
  - AP@K: K = cost budget

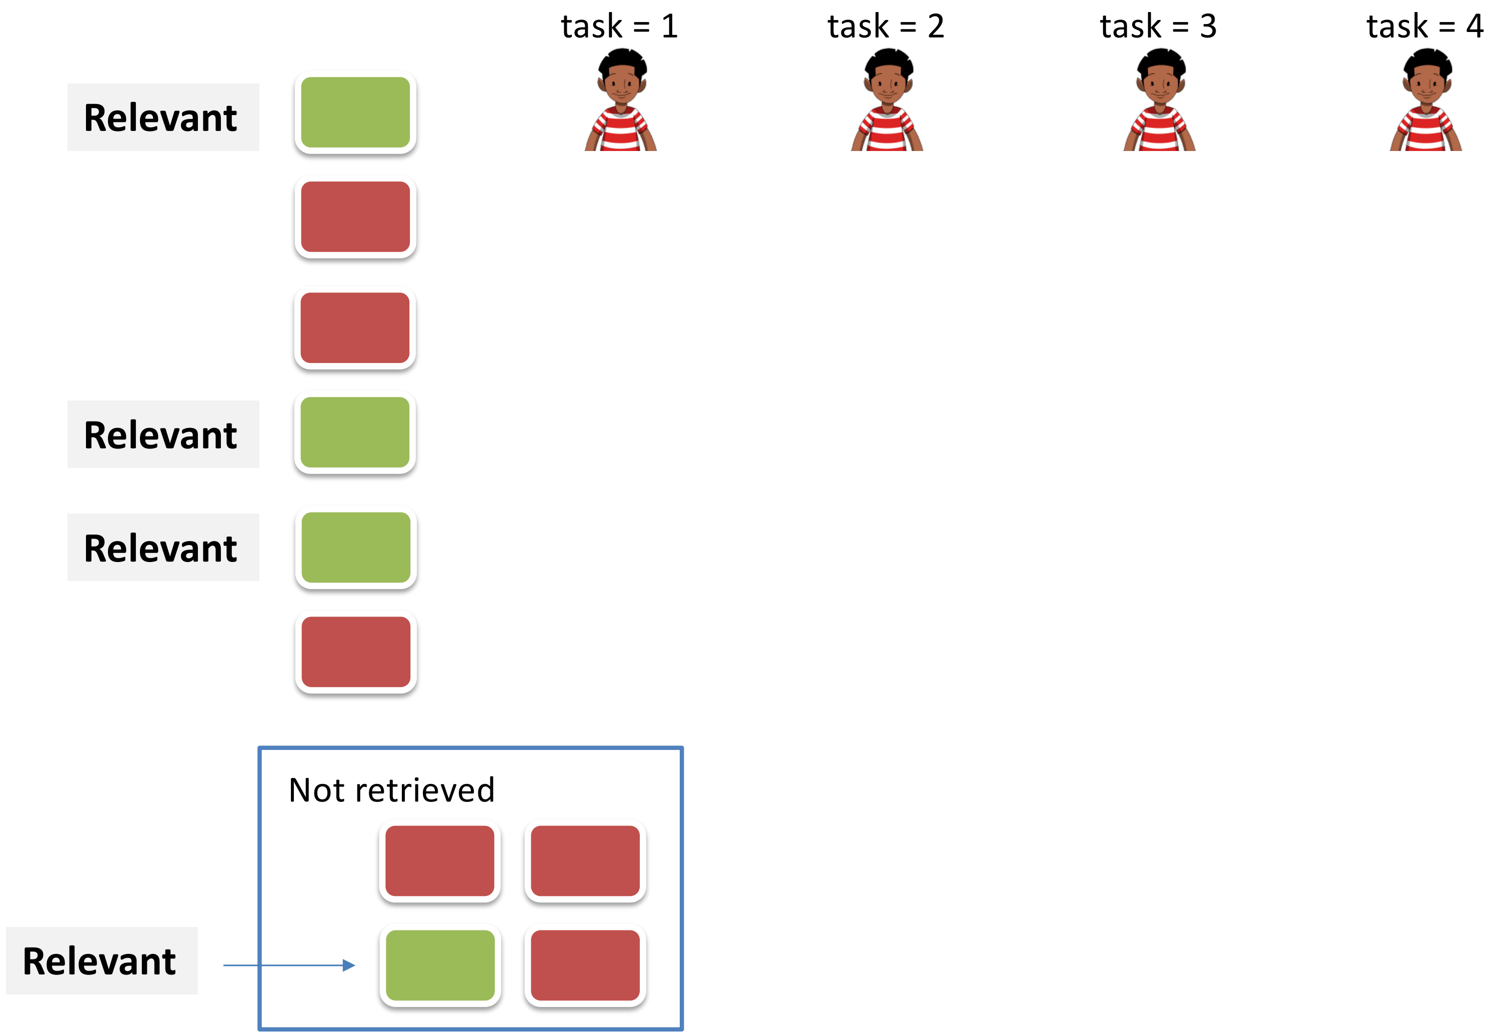
task = 1
task = 2
task = 3
task = 4
Relevant
Relevant
Relevant
Not retrieved
Relevant

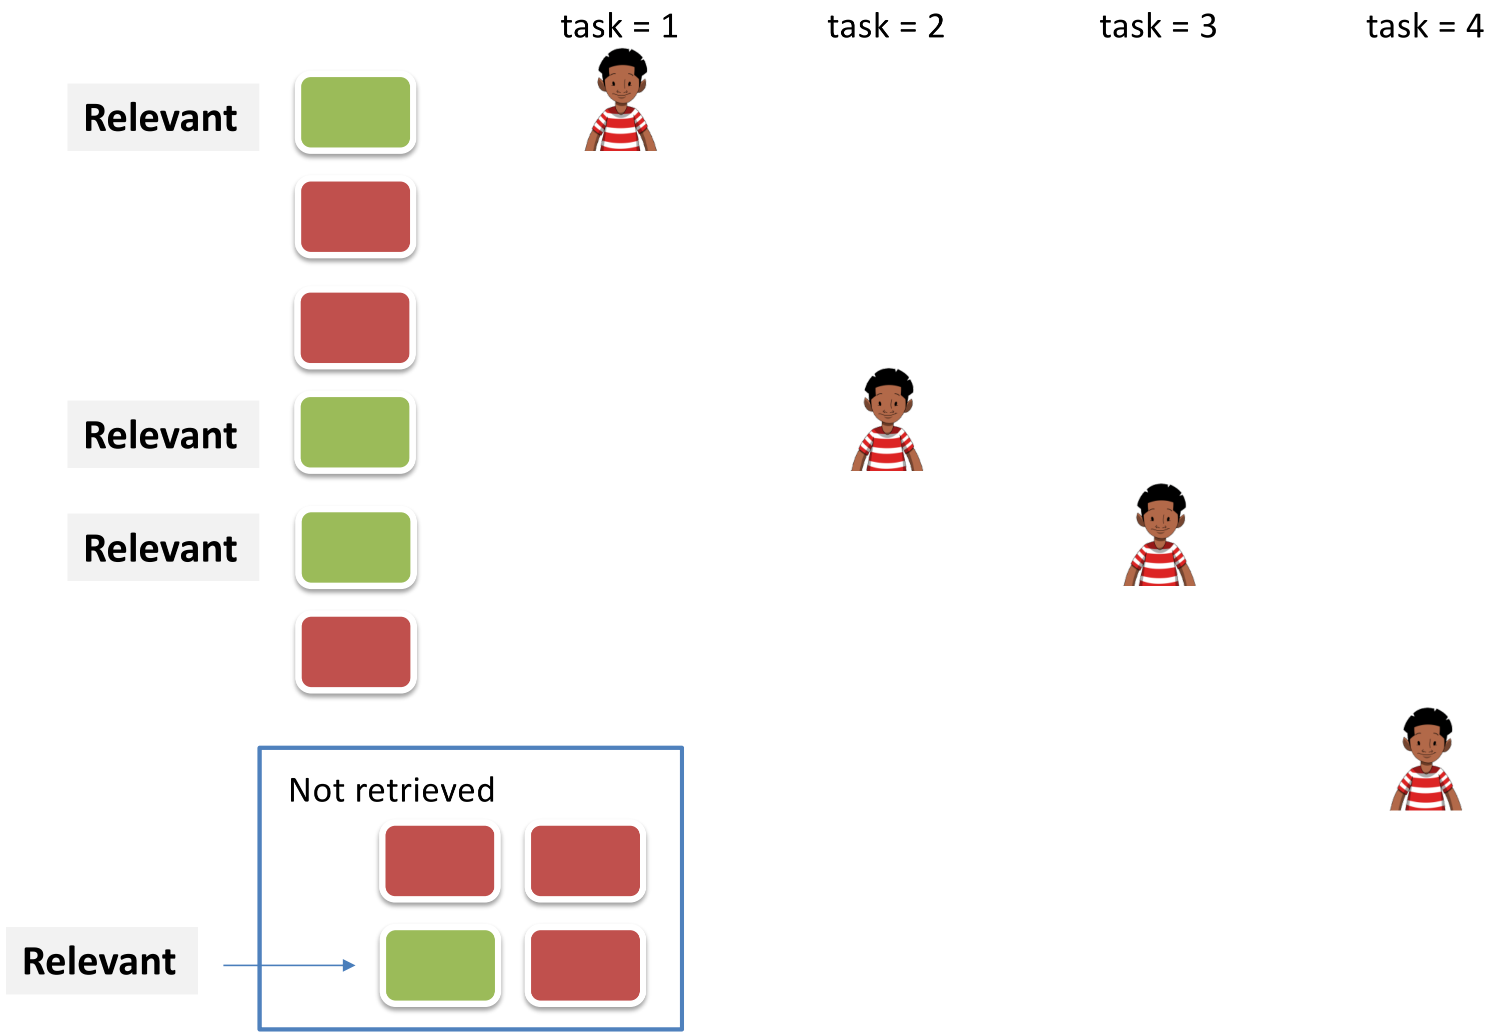

task = 1
task = 2
task = 3
task = 4
Relevant
Relevant
Relevant
Not retrieved
Relevant

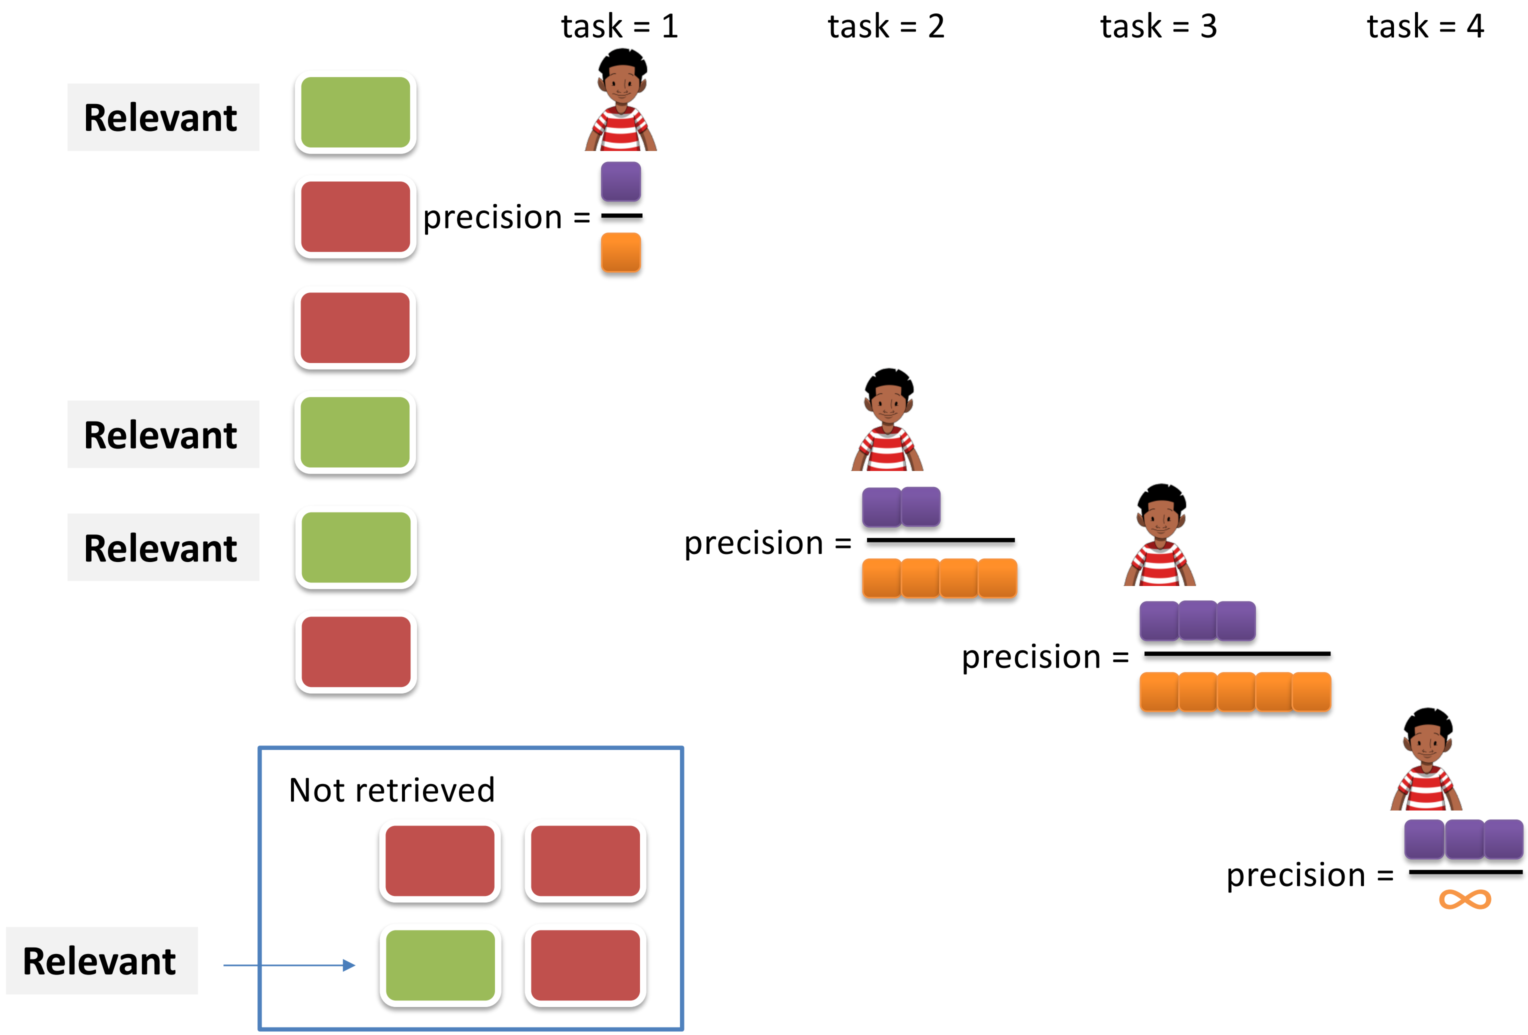
task = 1
task = 2
task = 3
task = 4
Relevant
precision =
Relevant
Relevant
precision =
precision =
Not retrieved
precision =
∞
Relevant

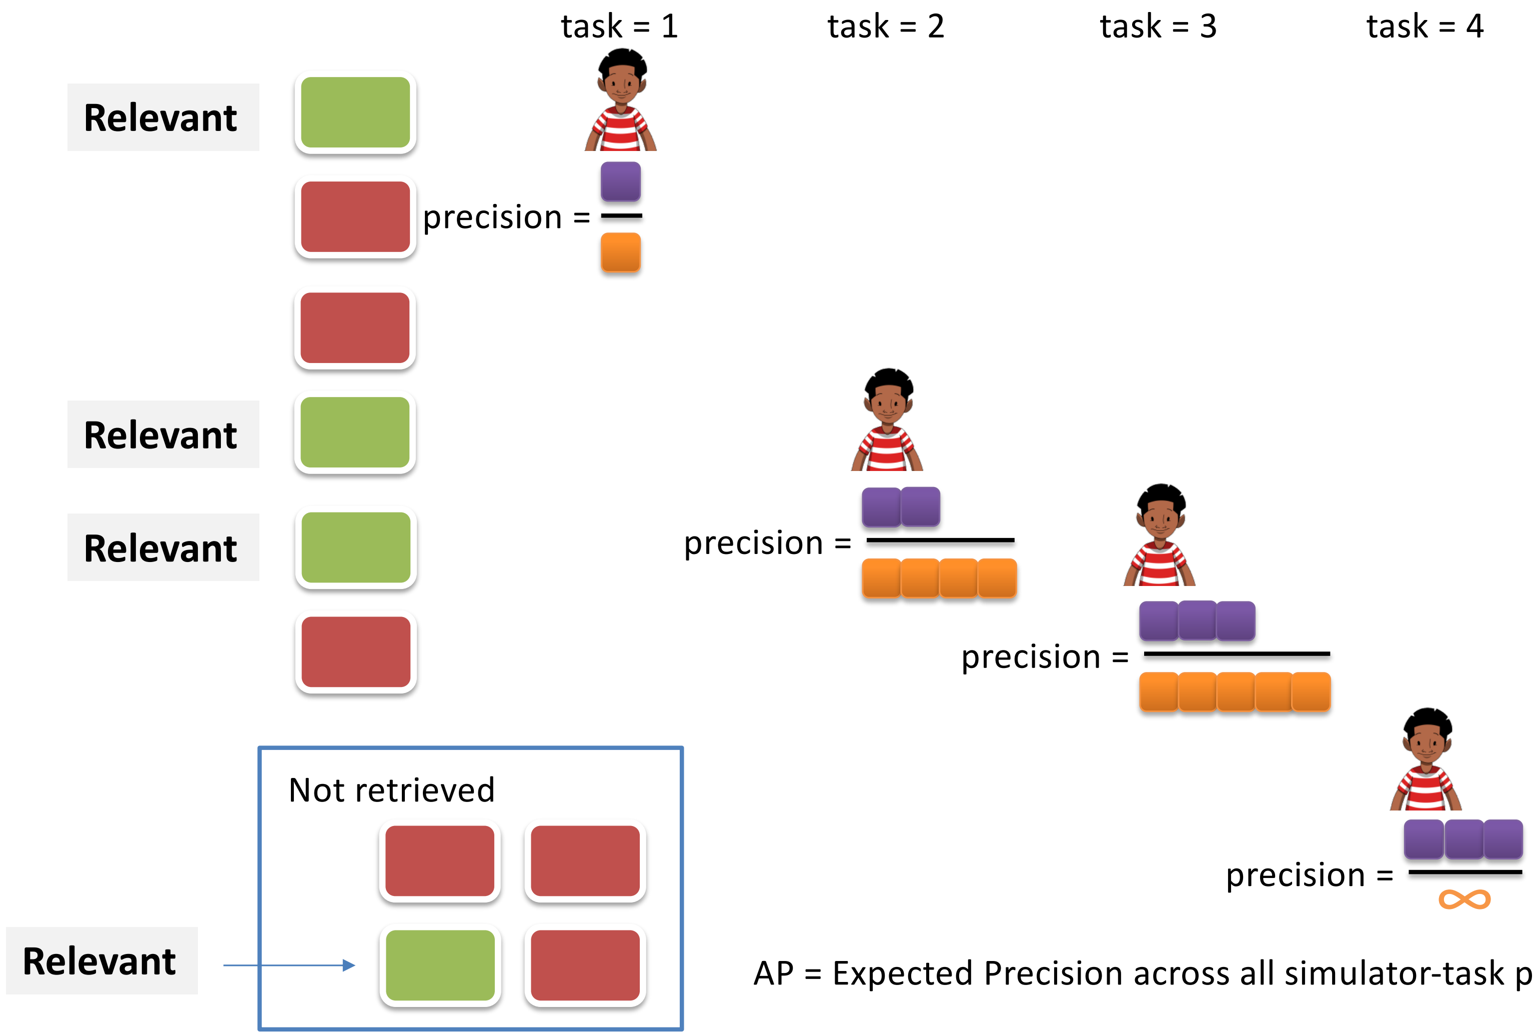

AP = Expected Precision across all simulator-task pairs

Future Work …

Major Challenges for Future Research in Interactive Recommendation and Search (IRS)

1. **How to evaluate an IRS system (with controlled experiments)?**
   - How to build realistic user simulators? User search logs? User study designed specifically for eliciting user behavior? How to evaluate simulators [Labhishetty & Zhai 21]?
   - How to measure task performance and measure user effort?
   - How to incorporate situation/context into an evaluation framework?
2. **How to formally (mathematically) represent and model a user?**
   - How to leverage theory from Psychology to design a formal user model?
   - How to represent a user's state of knowledge?
   - How to model many other aspects of a user (e.g., potential needs, browsing behavior, situational constraints, cognitive state, …)
   - How to model shared characteristics of users? Structure on users?

Major Challenges for Future Research (cont.)

3. How to infer and update a user model over time?

- Given all the observed data about a user, how can we infer knowledge about the user and update the user model over time?
- How can we recognize and correct errors in a user model (misunderstanding of users)?

4. How to model and infer a user's task?

- What is exactly a user task?
- How do we assess whether a user task has been completed? Assess progress toward task completion?
- How do we go beyond supporting query formulation to task specification?

Major Challenges for Future Research (cont.)

5. How do we design an "IRS game" with richer user actions and system responses?

- How can we systematically enumerate the possibilities of "interface cards"? Are there a finite number of basic interface elements that would be sufficient when combined in a flexible way?
- How can we design interfaces to encourage/optimize user-system collaboration? (**Interface = Language for communication between users and system**)
- How do we design interfaces to enable multi-mode interactions (e.g., speech + touch screen)?
- How can we design interfaces to enable a system to explain its responses to users?
- How can we parameterize an interface to enable automated optimization of interface using an algorithm?

6. How should we formalize the optimization problem of an "IRS game"?

- How do we formally define the multiple objectives (task performance, user effort, system cost, …)?
- How do we set up the optimization problem so as to make it feasible to solve it?

Major Challenges for Future Research (cont.)

7. How can we efficiently solve the optimization problem of IRS?

– POMDP and reinforcement learning are generally complex to compute. How can we simplify the objective function and make approximations?
– How can we leverage advances in machine learning to improve modeling and algorithms for IRS?
– How can we engage users to help simplify the optimization problem (resolve uncertainties)? How to simplify the exploration-exploitation tradeoff?

8. How can the system dynamically adapt the interface to each individual user in a context-sensitive and task-sensitive way?

– Novice vs. expert users?
– User sitting in a train vs. being at home?
– Medical diagnosis task vs. solve a homework problem?
– How can the system adapt the interface while minimizing the cognitive load on users? How can the system "train" a user to recognize changes in the interface?

Major Challenges for Future Research (cont.)

9. How can the system provide help for users all the time?

- Many “help me to do X” buttons and many explanations
- “Reporting problem” button on every interface page?
- How to maximize the flexibility for a user to dynamically reconfigure an interaction interface (let the user “program” the interface)?
- How to sense a user’s emotion during IRS?

10. How to support multi-mode interactions and engage a user to go beyond search or recommendation to support user tasks?

- An IRS system can evolve into a personalized intelligent task support agent

Summary

- **Recommendation** and **search** are **complementary** ways to serve users with useful information and can be studied in the **same unified cooperative game framework (CGF)**
- The unified problem can be called
  – Interactive/Intelligent Recommendation & Search (**IRS**), or
  – Interactive/Intelligent Search & Recommendation (**ISR**)
- Key challenges for future research:
  – Mathematical **modeling of users** (build **user simulators**)
  – **Continuous updating** of user model & **adaptive context-sensitive service** to each individual user
  – Collaboration with users (**learning to collaborate**)
  – Optimization of multiple objectives (**learning to make adaptive tradeoff**)
  – **Evaluation** of IRS/ISR (particularly using user simulators)
  – Optimization of system operation (minimization of **operation cost** and **energy consumption**)

Thank You!

Questions/Comments?

czhai@illinois.edu

http://czhai.cs.illinois.edu/

Looking forward to opportunities for collaboration!

References

- [Baskaya et al. 13] Feza Baskaya, Heikki Keskustalo, and Kalervo Järvelin. 2013. Modeling behavioral factors ininteractive information retrieval. In *Proceedings of ACM CIKM 2013,* 2297-2302.
- [Belkin 80] Belkin, N.J. "Anomalous states of knowledge as a basis for information retrieval". The Canadian Journal of Information Science, 5, 1980, pages 133-143.
- [Belkin et al. 82] Belkin, N.J., Oddy, R.N., Brooks, H.M. "ASK for information retrieval: Part I. Background and theory". The Journal of Documentation, 38(2), 1982, pages 61-71.
- [Belkin 96] Belkin, N. J. (1996). Intelligent information retrieval: Whose intelligence? *Proceedings of the Fifth International Symposium for Information Science*, Konstanz: Universitätsverlag Konstanz, 25-31.
- [Carterette et al. 15] Carterette, Ben, Ashraf Bah, and Mustafa Zengin. "Dynamic test collections for retrieval evaluation." *Proceedings of the 2015 international conference on the theory of information retrieval*. ACM, 2015.
- [Chuklin et al. 15] Chuklin, Aleksandr, Ilya Markov, and Maarten de Rijke. "Click models for web search." *Synthesis Lectures on Information Concepts, Retrieval, and Services* 7.3 (2015): 1-115.
- [Collins-Thompson et al. 11] Kevyn Collins-Thompson, Paul N. Bennett, Ryen W. White, Sebastian de la Chica, and David Sontag. 2011. Personalizing web search results by reading level. In *Proceedings of ACM CIKM 2011*, 403-412.
- [Collins-Thompson et al. 17] Collins-Thompson, Kevyn, Preben Hansen, and Claudia Hauff. "Search as learning (dagstuhl seminar 17092)." In *Dagstuhl reports*, vol. 7, no. 2. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

References (cont.)

- [de Vries et al. 04] A. P. de Vries, G. Kazai, and M. Lalmas. Tolerance to irrelevance: A user-effort oriented evaluation of retrieval systems without predefined retrieval unit. In Proc. RIAO, pages 463–473, 2004.
- [Hofmann et al. 11] Katja Hofmann, Shimon Whiteson, Maarten de Rijke: Balancing Exploration and Exploitation in Learning to Rank Online. ECIR 2011: 251-263
- [Ingwersen 96] Peter Ingwersen, Cognitive Perspectives of Information Retrieval Interaction: Elements of a Cognitive IR Theory. Journal of Documentation, v52 n1 p3-50 Mar 1996
- [Ingwersen & Järvelin 06] Ingwersen, Peter, and Kalervo Järvelin. *The turn: Integration of information seeking and retrieval in context*. Vol. 18. Springer Science & Business Media, 2006.
- [Joachims et al. 05] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In Proceedings of ACM SIGIR 2005, pp. 154-161. DOI=http://dx.doi.org/10.1145/1076034.1076063
- [Liu et al. 07] Yiqun Liu, Yupeng Fu, Min Zhang, Shaoping Ma, and Liyun Ru. 2007. Automatic search engine performance evaluation with click-through data analysis. In *Proceedings of the 16th international conference on World Wide Web* (WWW '07). ACM, New York, NY, USA, 1133-1134. DOI: https://doi.org/10.1145/1242572.1242731
- [Liu et al. 14] Liu, Yiqun, et al. "Overview of the NTCIR-11 IMine Task." *NTCIR*. 2014.
- [Liu et al. 17] Y. Liu, X Xie, C Wang, JY Nie, M Zhang, S Ma, Time-aware click model, ACM Transactions on Information Systems (TOIS) 35 (3), 2017.
- [Pääkkönen et al. 17] Pääkkönen, Teemu, Jaana Kekäläinen, Heikki Keskustalo, Leif Azzopardi, David Maxwell, and Kalervo Järvelin. "Validating simulated interaction for retrieval evaluation." *Information Retrieval Journal* 20, no. 4 (2017): 338-362.

References (cont.)

- [Shen et al. 05] Xuehua Shen, Bin Tan, and ChengXiang Zhai, Implicit User Modeling for Personalized Search , In *Proceedings of the 14th ACM International Conference on Information and Knowledge Management* (CIKM'05), pages 824-831.
- [Shen & Zhai 05] Xuehua Shen, ChengXiang Zhai, Active Feedback in Ad Hoc Information Retrieval, *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (SIGIR'05), 59-66, 2005.
- [Smucker & Clarke 12] Mark D. Smucker and Charles L.A. Clarke. 2012. Time-based calibration of effectiveness measures. In *Proceedings of ACM SIGIR 2012;* 95-104.
- [Thomas et al. 14] Paul Thomas, Alistair Moffat, Peter Bailey, and Falk Scholer. 2014. Modeling decision points in user search behavior. In *Proceedings of the 5th Information Interaction in Context Symposium* (IIiX '14). 239-242.
- [Wang et al. 16] Yue Wang, Dawei Yin, Luo Jie, Pengyuan Wang, Makoto Yamada, Yi Chang, and Qiaozhu Mei. 2016. Beyond Ranking: Optimizing Whole-Page Presentation. In Proceedings of WSDM 2016.
- [Wang et al. 19] Huazheng Wang, Sonwoo Kim, Eric McCord-Snook, Qingyun Wu, and Hongning Wang. 2019. Variance Reduction in Gradient Exploration for Online Learning to Rank. In Proceedings of ACM SIGIR 2019.
- [Wei et al. 17] Wei, Zeng, Jun Xu, Yanyan Lan, Jiafeng Guo, and Xueqi Cheng. "Reinforcement learning to rank with Markov decision process." In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 945-948. ACM, 2017.
- [Witten et al. 99] Ian H. Witten, Alistair Moffat, and Timothy C. Bell. 1999. *Managing Gigabytes (2nd Ed.): Compressing and Indexing Documents and Images*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

References (cont.)

- [Yang et al. 16] Grace Hui Yang, Marc Sloan, and Jun Wang. 2016. *Dynamic Information Retrieval Modeling*. Morgan & Claypool Publishers
- [Zhai 02] ChengXiang Zhai, Risk Minimization and Language Modeling in Information Retrieval, Ph.D. thesis, Carnegie Mellon University, 2002.
- [Zhai 16] ChengXiang Zhai. Towards a game-theoretic framework for text data retrieval, IEEE Data Eng. Bull. 39(3): 51-62 (2016).
- [Zhai et al. 03] ChengXiang Zhai, William W. Cohen, and John Lafferty, Beyond Independent Relevance: Methods and Evaluation Metrics for Subtopic Retrieval , *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (SIGIR'03), pages 10-17, 2003.
- [Zhai & Lafferty 06] ChengXiang Zhai, John D. Lafferty: A risk minimization framework for information retrieval. Inf. Process. Manage. 42(1): 31-55 (2006)
- [Zhang et al. 2017] Yinan Zhang, Xueqing Liu, ChengXiang Zhai: Information Retrieval Evaluation as Search Simulation: A General Formal Framework for IR Evaluation. ICTIR 2017: 193-200
- [Zhang & Zhai 15] Yinan Zhang, ChengXiang Zhai, Information Retrieval as Card Playing: A Formal Model for Optimizing Interactive Retrieval Interface. In *Proceedings of ACM SIGIR 2015, pp.* 685-694.
- [Zhao et al. 18] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep reinforcement learning for page-wise recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems* (RecSys '18). ACM, New York, NY, USA, 95-103. DOI: https://doi.org/10.1145/3240323.3240374