Lost in Publications?

How to **Find Your Way** in **50 Million Scientific Documents**

Jaakko Peltonen
Aalto University and University of Tampere

based on papers by Tuukka Ruotsalo, Jaakko Peltonen, Manuel Eugster, Dorota Glowacka, Ksenia Konyushkova, Kumaripaba Athukorala, Ilkka Kosunen, Aki Reijonen, Petri Myllymäki, Giulio Jacucci, Samuel Kaski, Chirayu Wongchokprasitti, Payel Bandyopadhyay, and Peter Brusilovsky

Thanks to Revolution of Knowledge Work (Re:Know) project


Tampere
Helsinki

Aalto University

UNIVERSITY OF TAMPERE

How to find relevant data?

(when you don't yet know what you need)

Overview

Researchers must **navigate big data**. Current scientific knowledge includes 50 million published articles. How can a system **help a researcher find relevant documents** in her field?

Overview

Researchers must **navigate big data**. Current scientific knowledge includes 50 million published articles. How can a system **help a researcher find relevant documents** in her field?

We introduce **IntentRadar**, an interactive search user interface and search engine that **anticipates user's search intents** by estimating them from user's **interaction** with the interface. The estimated intents are **visualized** on a radial layout that organizes potential intents as directions in the information space.

Overview

The intent radar assists users to direct their search by allowing **feedback** to be targeted on **keywords** that represent the potential intents.

Users give feedback by dragging keywords

The system then learns and visualizes improved estimates and corresponding documents.

Overview

The intent radar assists users to direct their search by allowing **feedback** to be targeted on **keywords** that represent the potential intents.

Users give feedback by dragging keywords

The system then learns and visualizes improved estimates and corresponding documents.

IntentRadar **significantly improves users' task performance** and quality of retrieved information without compromising task execution time.

Scientific document search

Exploration and search in literature are main tasks of a researcher. Crucial for human analysis of big data.

Comprehensively following an interest is usually not feasible - too many potential sources of interesting information.

Need for search arises because:
- you have an interest in a new topic, but do not know where to find good information about it
- you were alerted that new information is available
- you forgot the location of information you have seen

Typical interfaces for scientific search

Typical interfaces for scientific search

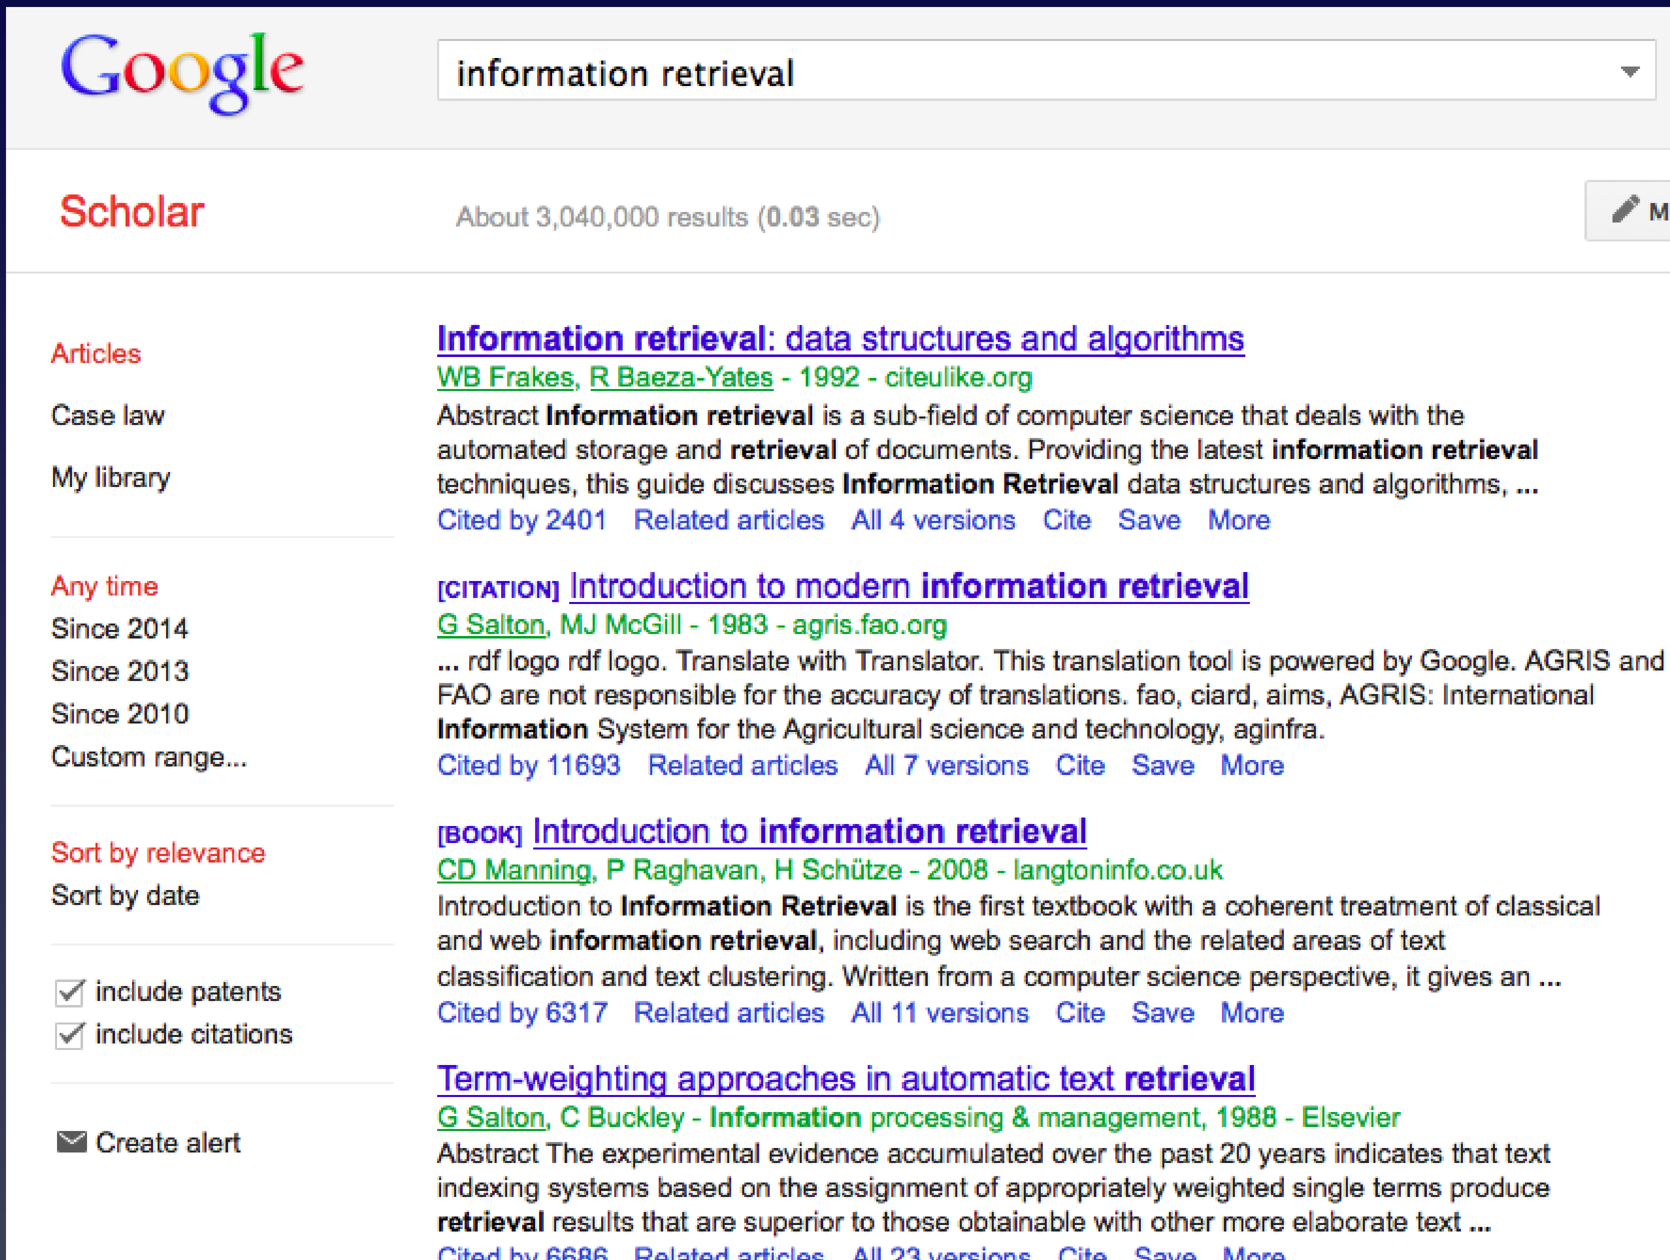

Typical interfaces for scientific search

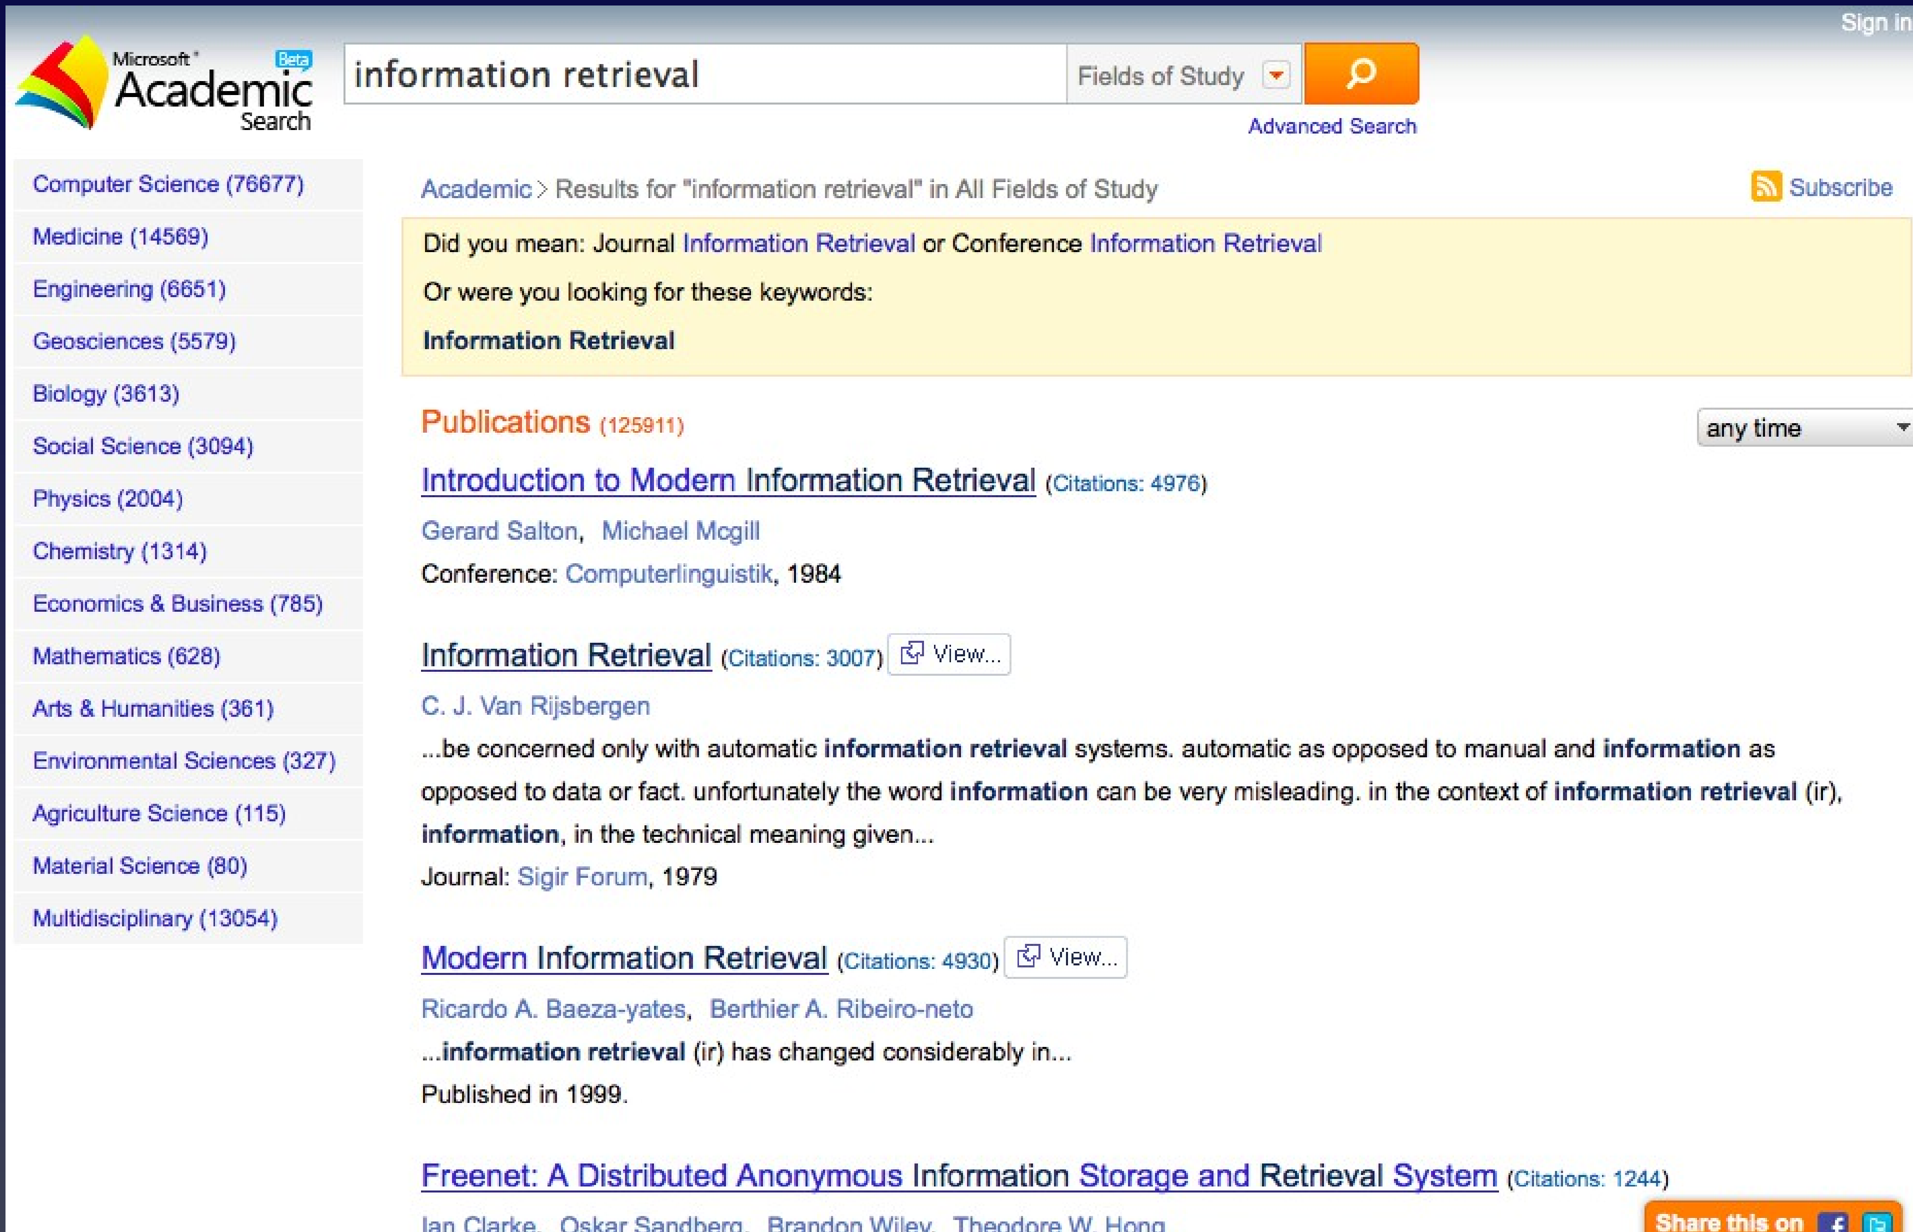

Typical interfaces for scientific search

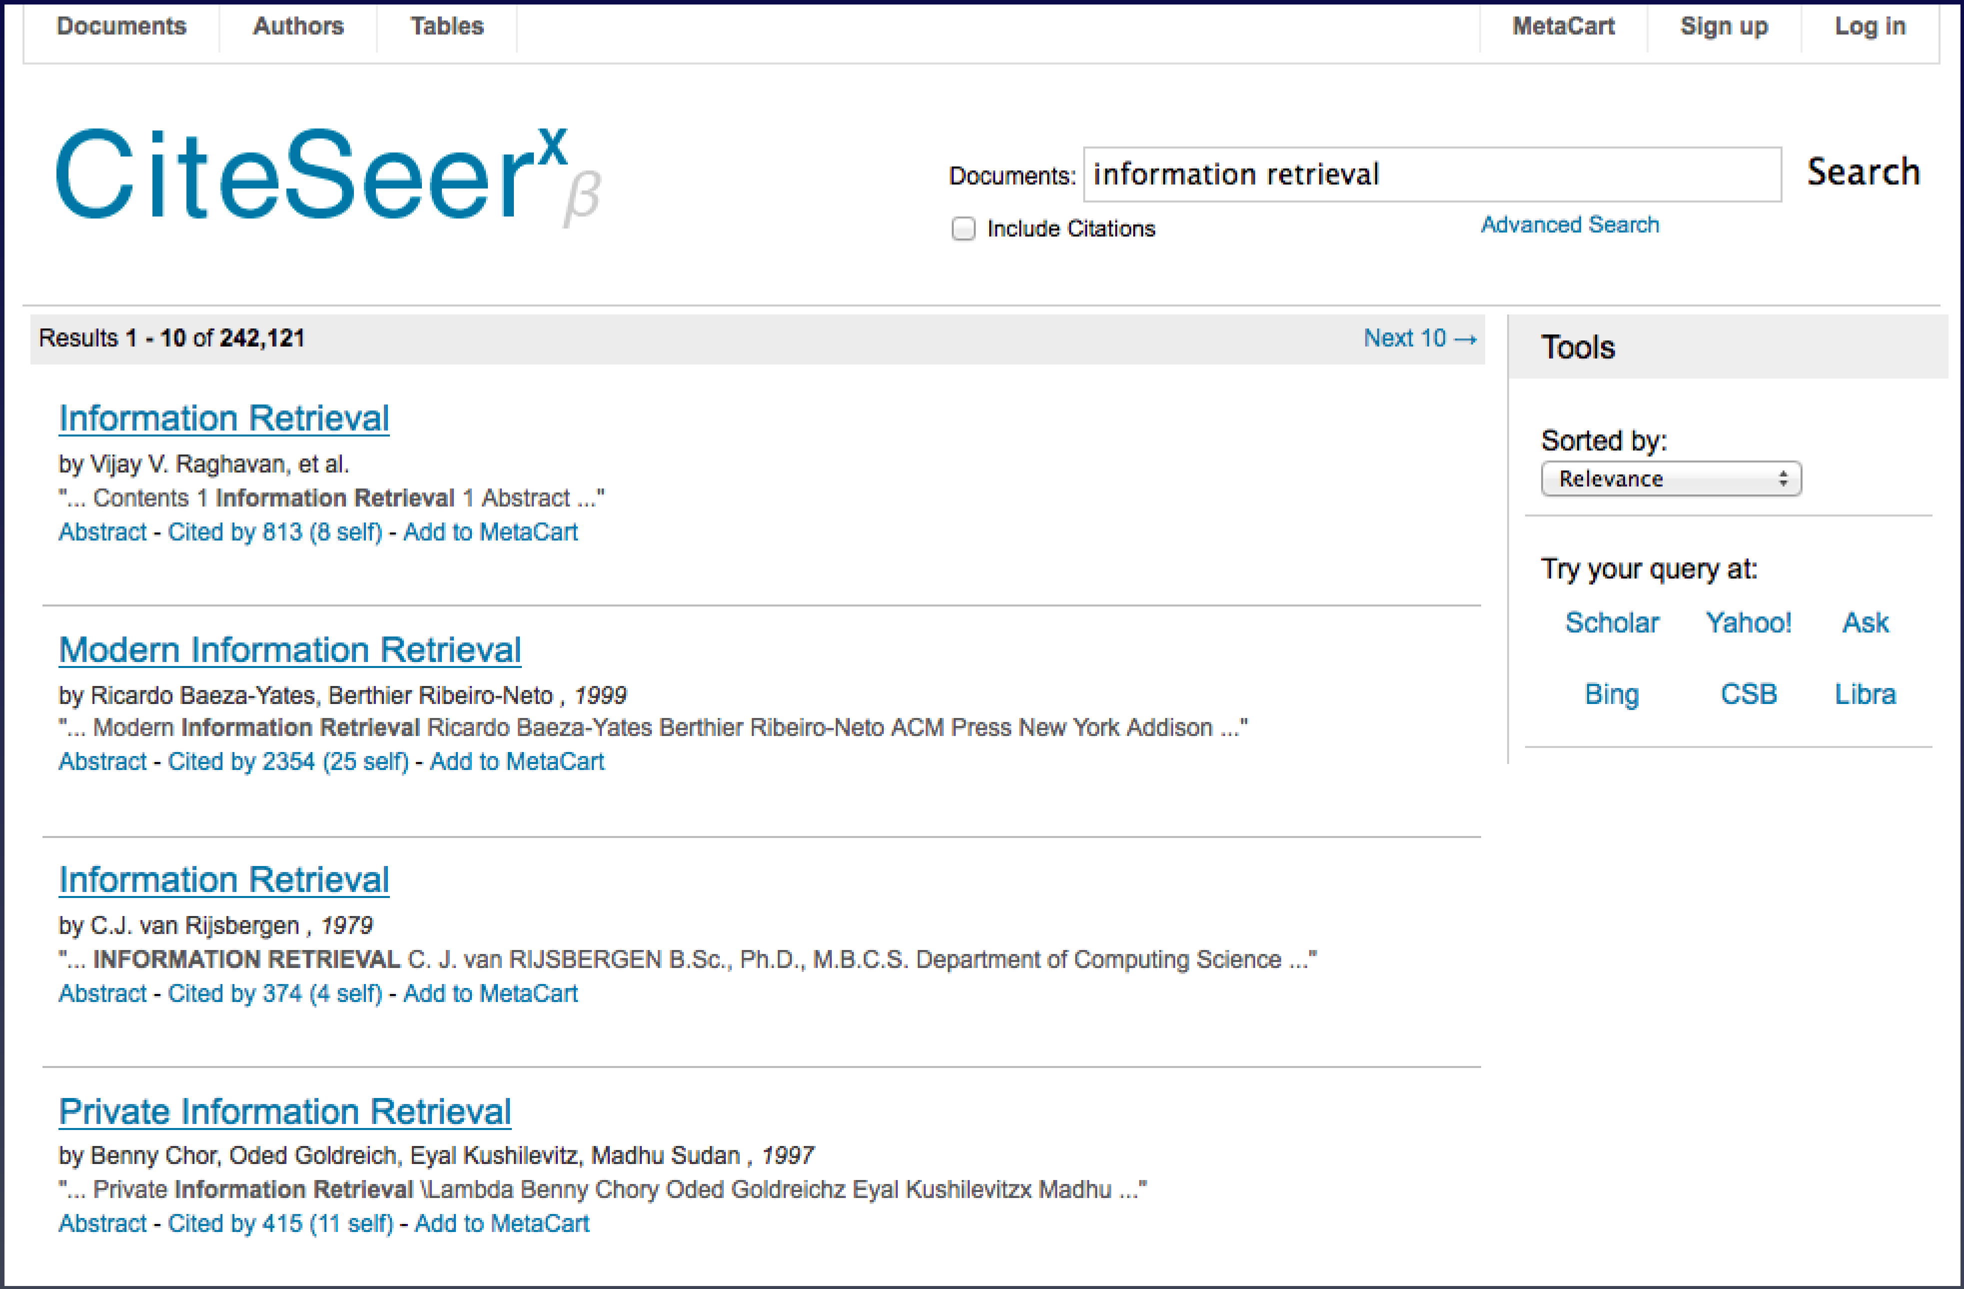

Different kinds of searches

Searching for a particular document
- e.g. someone told you there was an article about a new technology breakthrough, and you want to find that article

Searching for documents about a specific topic
- several documents complementing each other
- e.g. different commentaries on a politics event

Searching for documents about a general topic
- trying to understand/make sense of the topic
- diverse subtopics
- no single document may be enough
- no single search may be enough

Different kinds of searches

Studies have estimated that up to 50% of searching is informational and the corresponding search behavior is exploratory and spreads across individual queries and information needs

A main problem: hard to formulate queries precisely, information needs evolve.

A main problem: hard to formulate queries precisely, information needs evolve.

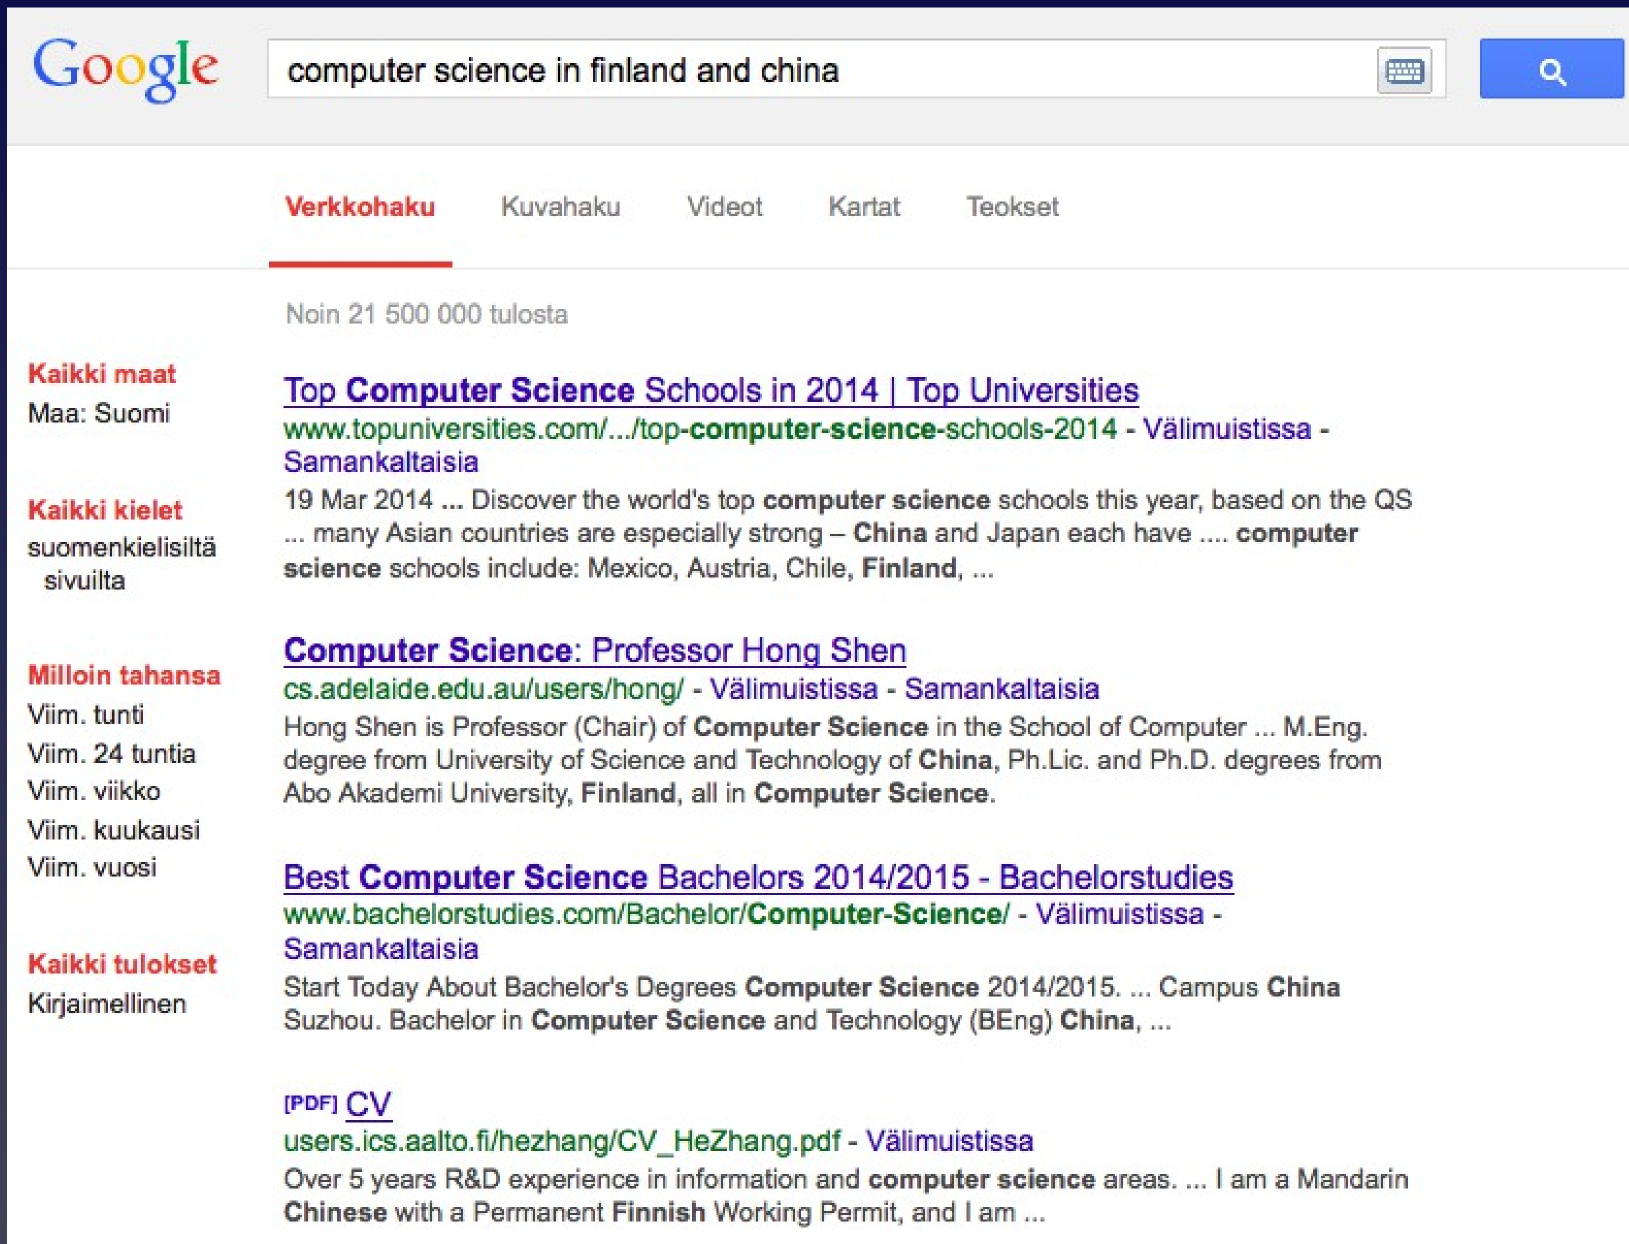

A main problem: hard to formulate queries precisely, information needs evolve.

Search engines can mistake what you are looking for.

You may not know what precisely you are looking for, or may not be able to express it as a search phrase.

There is a disconnect between what the computer thinks you need, and what you actually need.

Traditional interfaces only allow you to try a search phrase, and try again if you don't like the results.

----> “guessing game”: what phrase (if any) will give the results I need

Our approach

- Current support methods (suggesting query terms, faceted browsing, result clustering) can trap user to initial context.

Existing techniques are effective for tasks where the user’s goal is well defined and success is measured based on system response to well formed queries.

But in exploratory search the user’s information needs evolve throughout the course of the search and her ability to direct the search to solve her task is critical.

- **Our system:** helps users explore effectively: rapid feedback loops

- Helps make sense of information around query context

Our approach: Radar layout

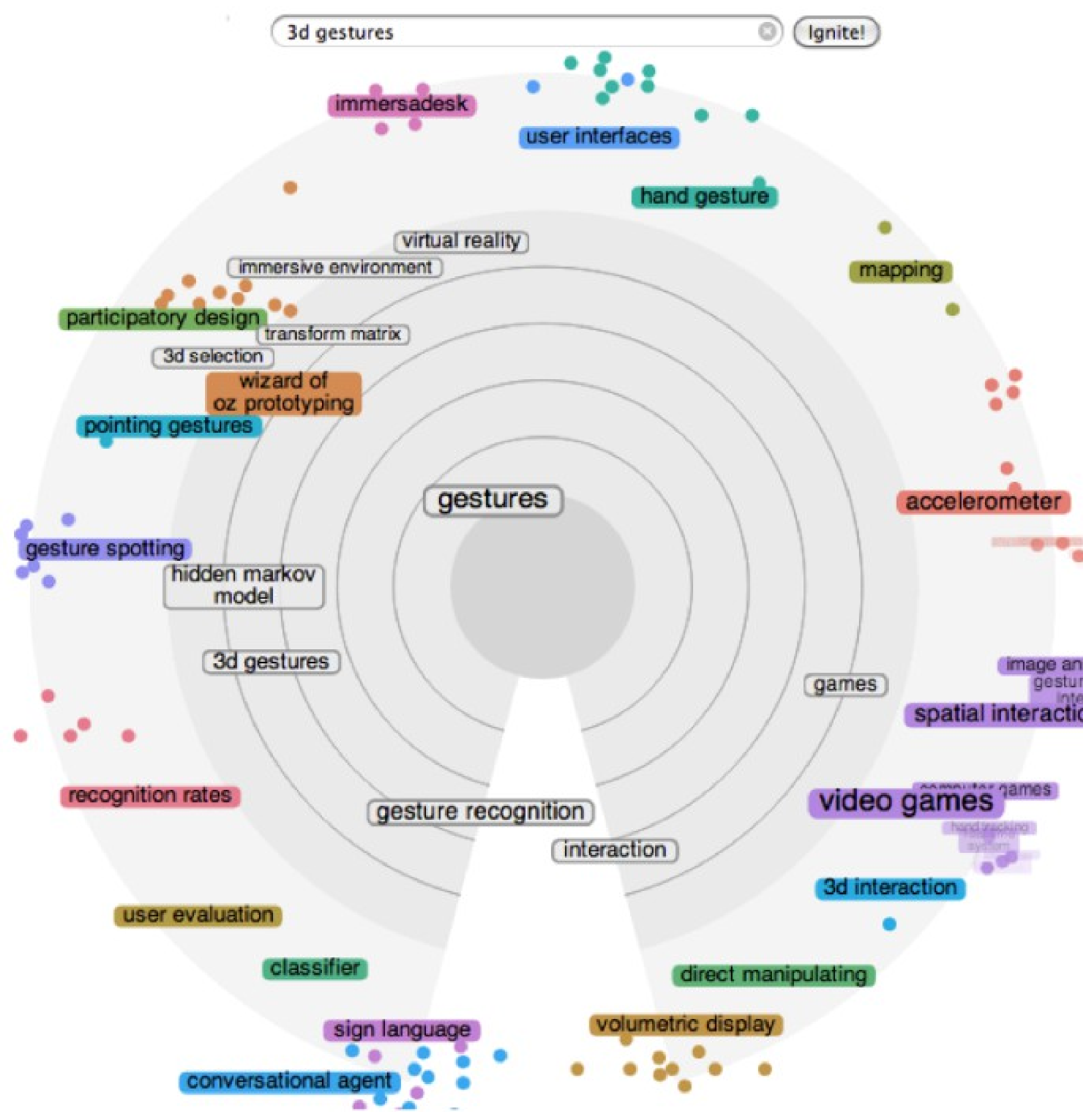

[Wiizards: 3D gesture recognition for game play input]
Louis Kratz, Matthew Smith, Frank J. Lee (Proceedings of the 2007 Conference on Future Play, 2007-01-01)
gestures | hidden markov model | interaction | 3d gestures | gesture recognition | accelerometer | games
`Gesture based input is an emerging techn...`

Feature Representations for the Recognition of 3D Emblematic Gestures
J Richarz, G A Fink (HUMAN BEHAVIOR UNDERSTANDING, 2010-01-01)
3d dynamic gesture recogmtion | human-machine interaction | smart rooms | time-senes analysis | trajectory | gestures | interaction
`In human-machine interaction, gestures p...`

A gesture recognition system using 3D data
S Malassiotis, N Aifanti, M G Strintzis (FIRST INTERNATIONAL SYMPOSIUM ON 3D DATA PROCESSING VISUALIZATION AND TRANSMISSION, 2002-01-01)
gestures | sign language | gesture recognition | classifier
`In this paper a gesture recognition syst...`

Gesture as an important factor in 3D kinematic assessment of the knee
R Lavoie, M Laplante, N Duval, S Dore, J A de (KNEE SURGERY SPORTS TRAUMATOLOGY ARTHROSCOPY, 2008-01-01)
knee | kinematics | gesture | knee-bend | variability | gestures
`Contradictions exist between studies of ...`

[Gameplay issues in the design of spatial 3D gestures for video games]
John Payne, Paul Keir, Jocelyn Elgoyhen, Mairghread McLundie, Martin Naef, Martyn Horner, Paul Anderson (Proceedings of ACM CHI 2006 Conference on Human Factors in Computing Systems, 2006-01-01)
spatial interaction | virtual reality | mapping | gestures | video games | interaction | 3d interaction | 3d gestures | games
`We describe preliminary tests that form ...`

The system uses a radar visualization metaphor.

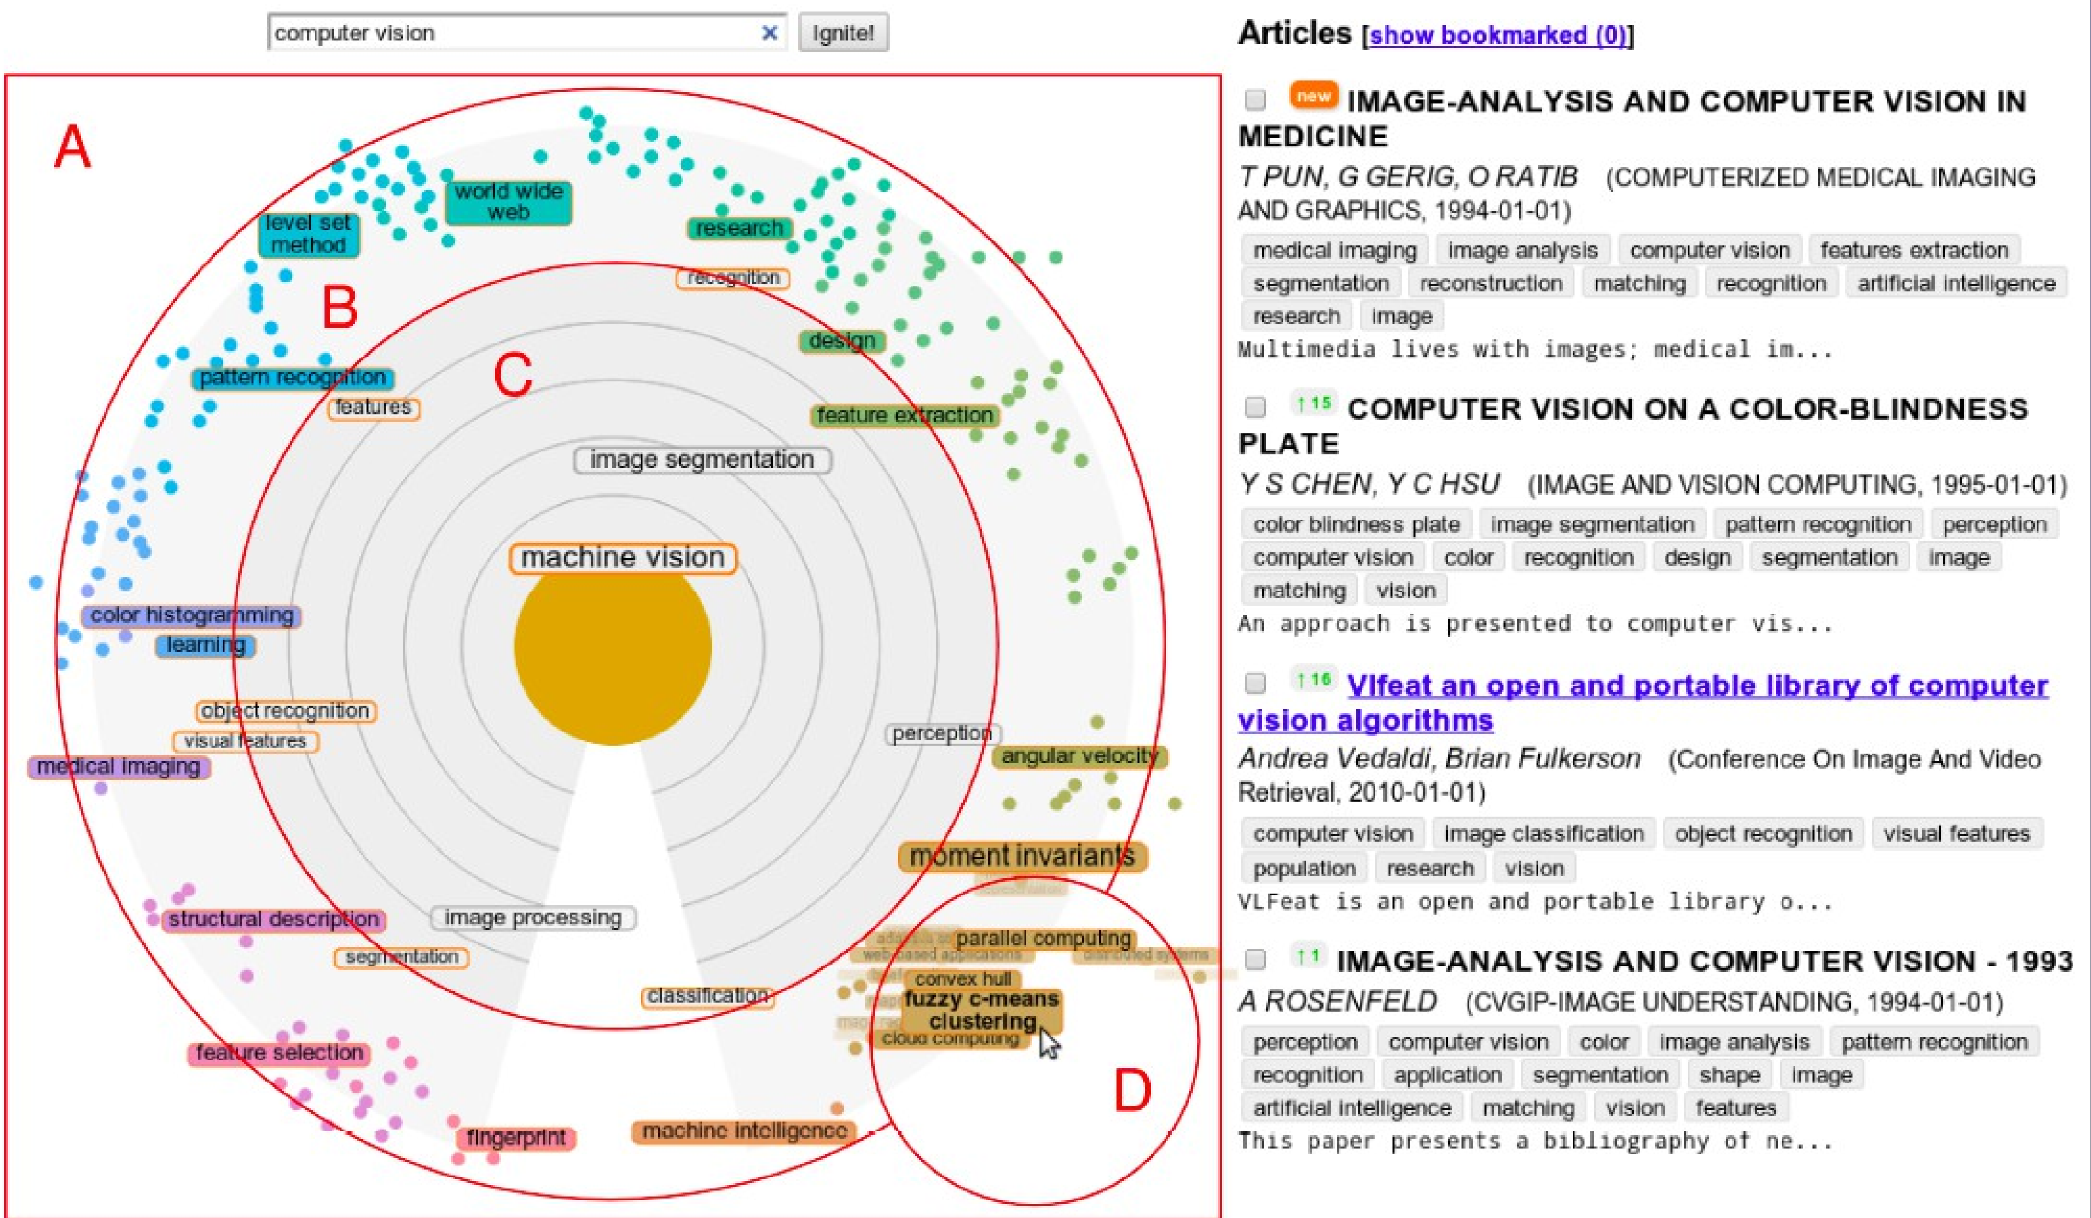

Article list

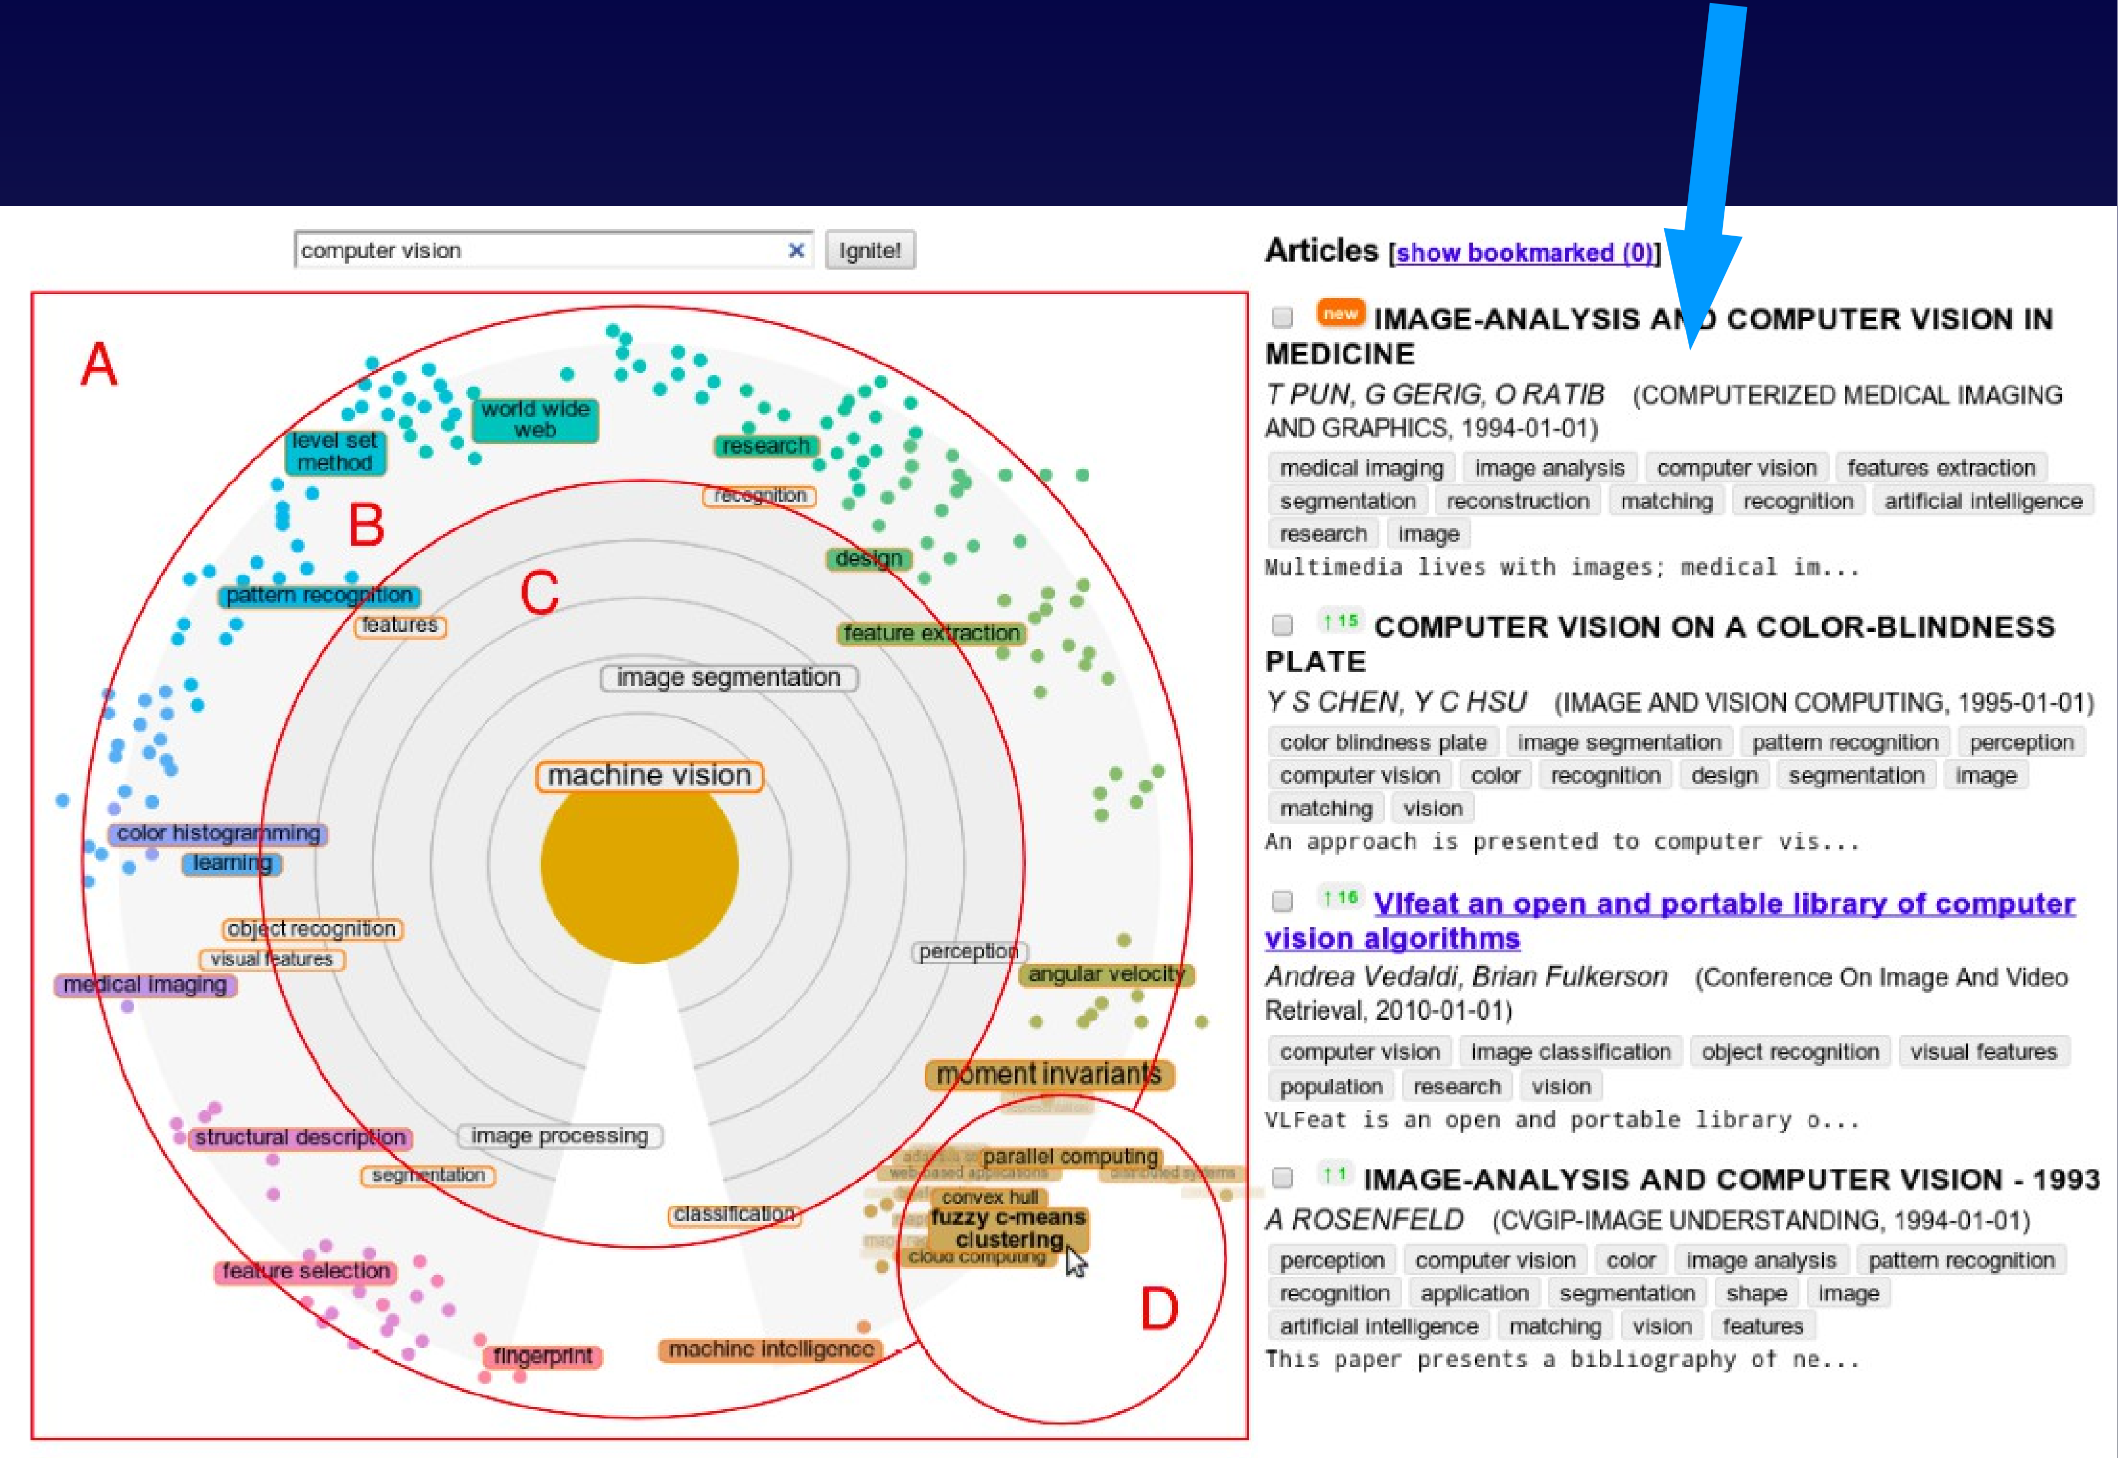

Radar screen

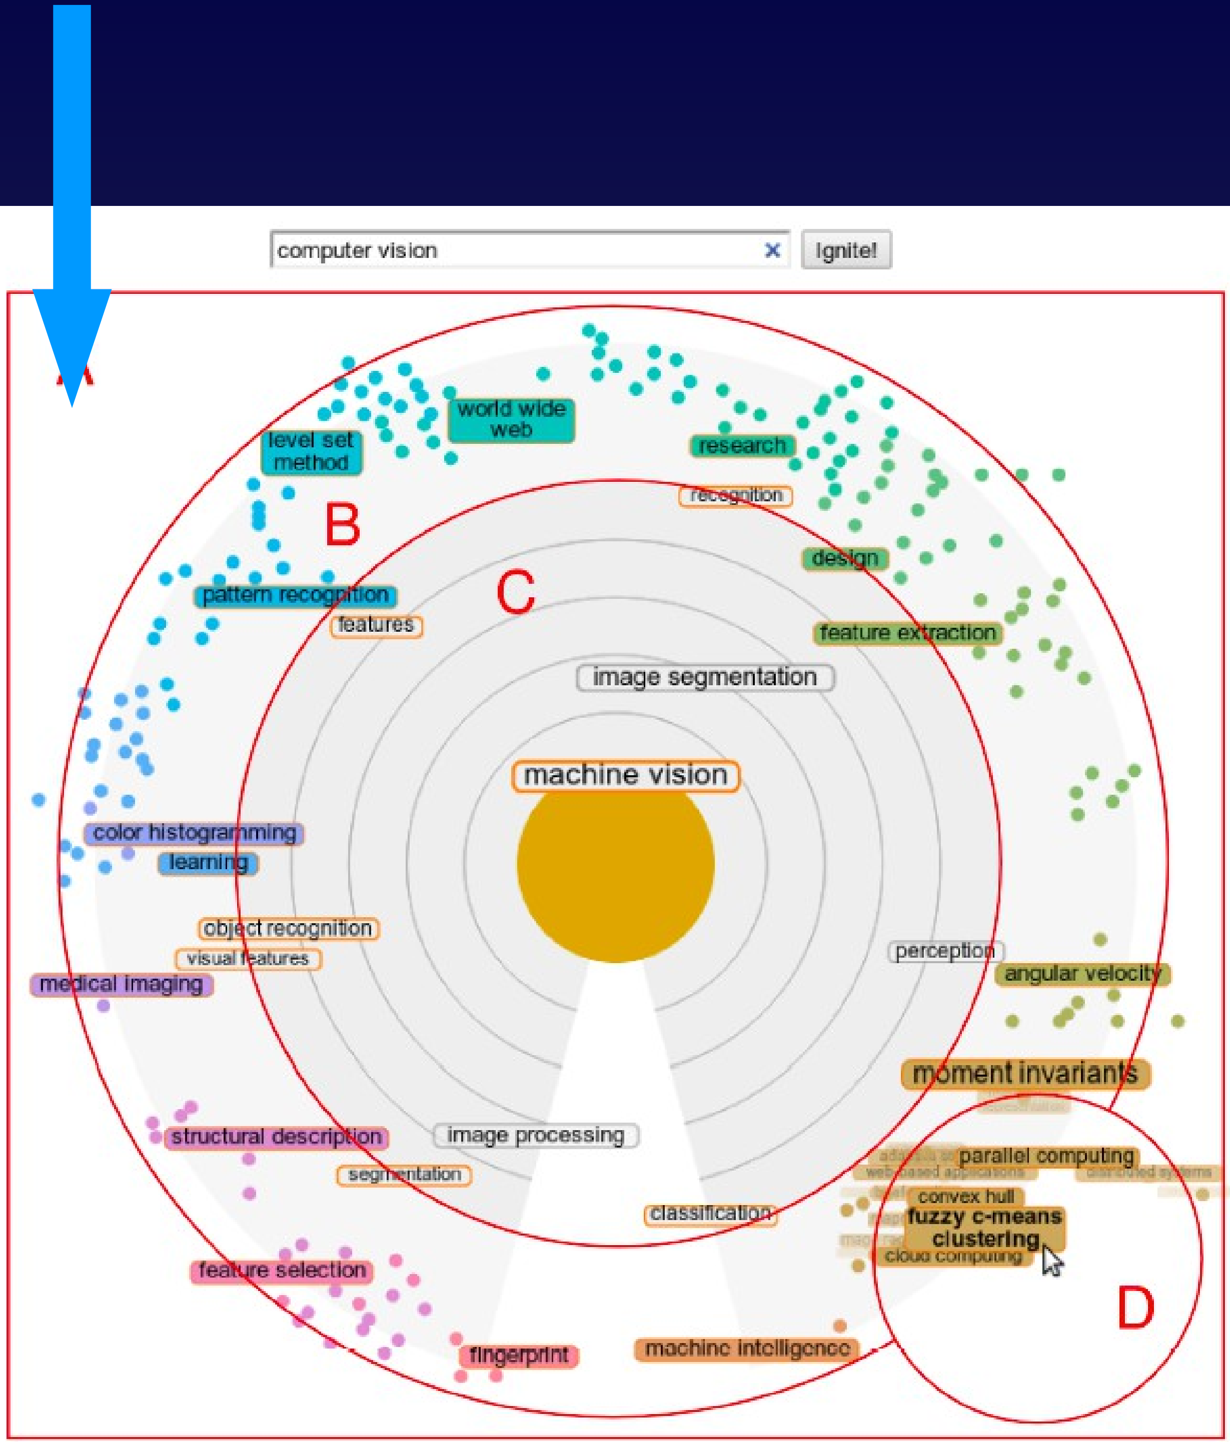

Articles [show bookmarked (0)]

new **IMAGE-ANALYSIS AND COMPUTER VISION IN MEDICINE**

T PUN, G GERIG, O RATIB (COMPUTERIZED MEDICAL IMAGING AND GRAPHICS, 1994-01-01)

medical imaging | image analysis | computer vision | features extraction | segmentation | reconstruction | matching | recognition | artificial intelligence | research | image

`Multimedia lives with images; medical im...`

↑15 **COMPUTER VISION ON A COLOR-BLINDNESS PLATE**

Y S CHEN, Y C HSU (IMAGE AND VISION COMPUTING, 1995-01-01)

color blindness plate | image segmentation | pattern recognition | perception | computer vision | color | recognition | design | segmentation | image | matching | vision

`An approach is presented to computer vis...`

↑16 **Vlfeat an open and portable library of computer vision algorithms**

Andrea Vedaldi, Brian Fulkerson (Conference On Image And Video Retrieval, 2010-01-01)

computer vision | image classification | object recognition | visual features | population | research | vision

`VLFeat is an open and portable library o...`

↑1 **IMAGE-ANALYSIS AND COMPUTER VISION - 1993**

A ROSENFELD (CVGIP-IMAGE UNDERSTANDING, 1994-01-01)

perception | computer vision | color | image analysis | pattern recognition | recognition | application | segmentation | shape | image | artificial intelligence | matching | vision | features

`This paper presents a bibliography of ne...`

Current intent estimation for which results are retrieved.
Angular distance = similarity of intents, radius = relevance

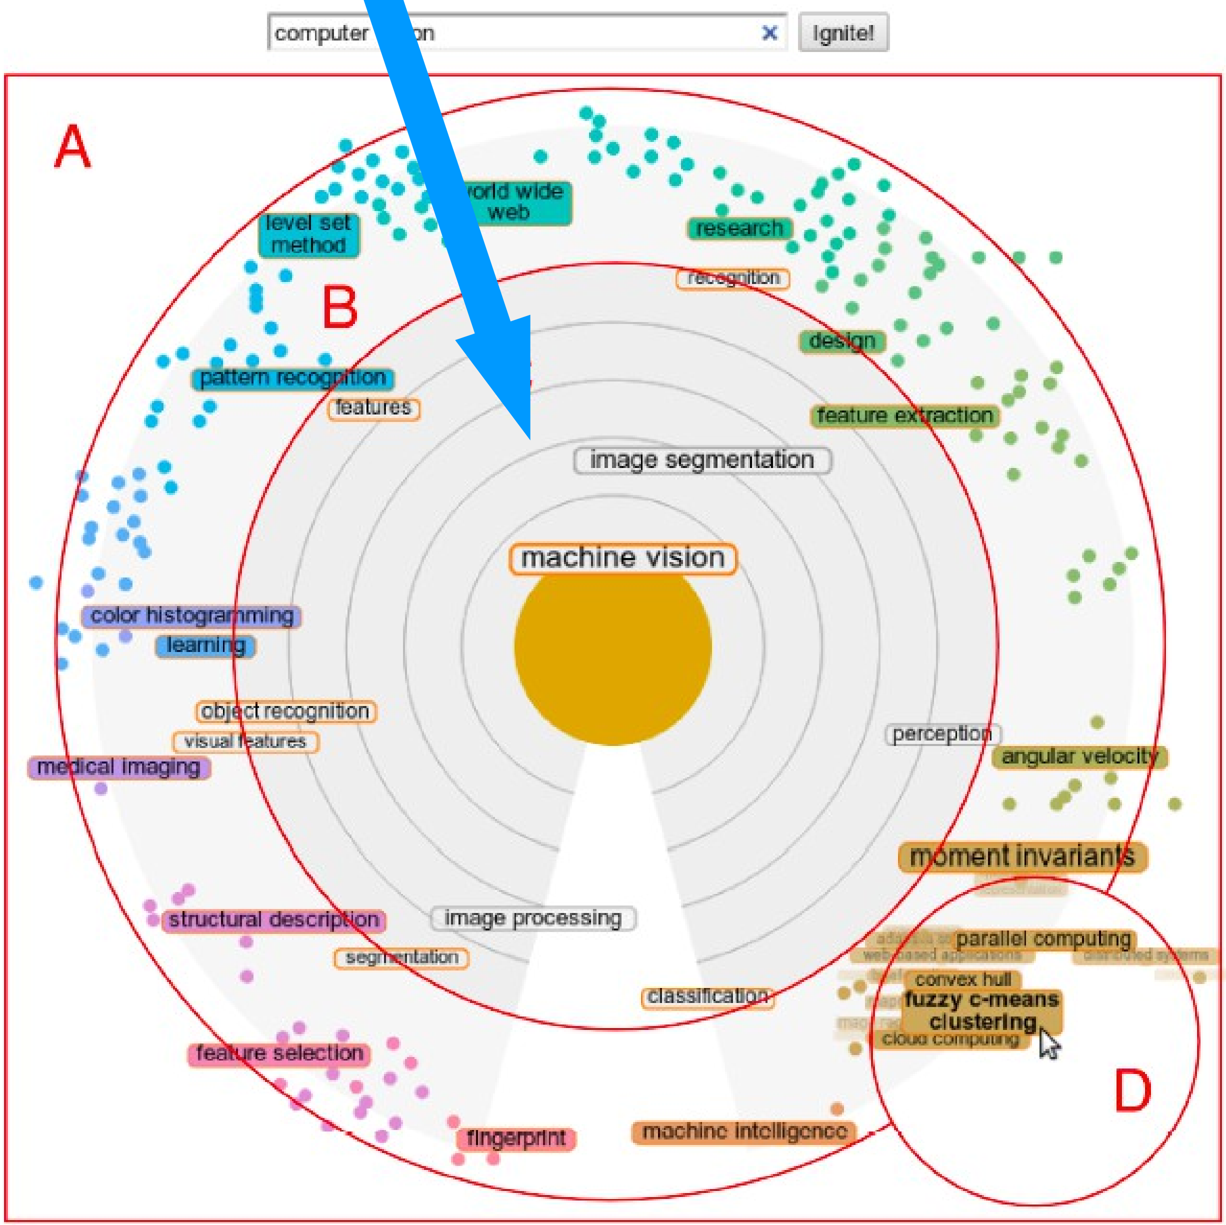

Articles [show bookmarked (0)]

new **IMAGE-ANALYSIS AND COMPUTER VISION IN MEDICINE**

T PUN, G GERIG, O RATIB (COMPUTERIZED MEDICAL IMAGING AND GRAPHICS, 1994-01-01)

medical imaging | image analysis | computer vision | features extraction | segmentation | reconstruction | matching | recognition | artificial intelligence | research | image

`Multimedia lives with images; medical im...`

↑15 **COMPUTER VISION ON A COLOR-BLINDNESS PLATE**

Y S CHEN, Y C HSU (IMAGE AND VISION COMPUTING, 1995-01-01)

color blindness plate | image segmentation | pattern recognition | perception | computer vision | color | recognition | design | segmentation | image | matching | vision

`An approach is presented to computer vis...`

↑16 **Vlfeat an open and portable library of computer vision algorithms**

Andrea Vedaldi, Brian Fulkerson (Conference On Image And Video Retrieval, 2010-01-01)

computer vision | image classification | object recognition | visual features | population | research | vision

`VLFeat is an open and portable library o...`

↑1 **IMAGE-ANALYSIS AND COMPUTER VISION - 1993**

A ROSENFELD (CVGIP-IMAGE UNDERSTANDING, 1994-01-01)

perception | computer vision | color | image analysis | pattern recognition | recognition | application | segmentation | shape | image | artificial intelligence | matching | vision | features

`This paper presents a bibliography of ne...`

Predicted intents (help users to find directions on the radar to move away from their currently estimated intent)

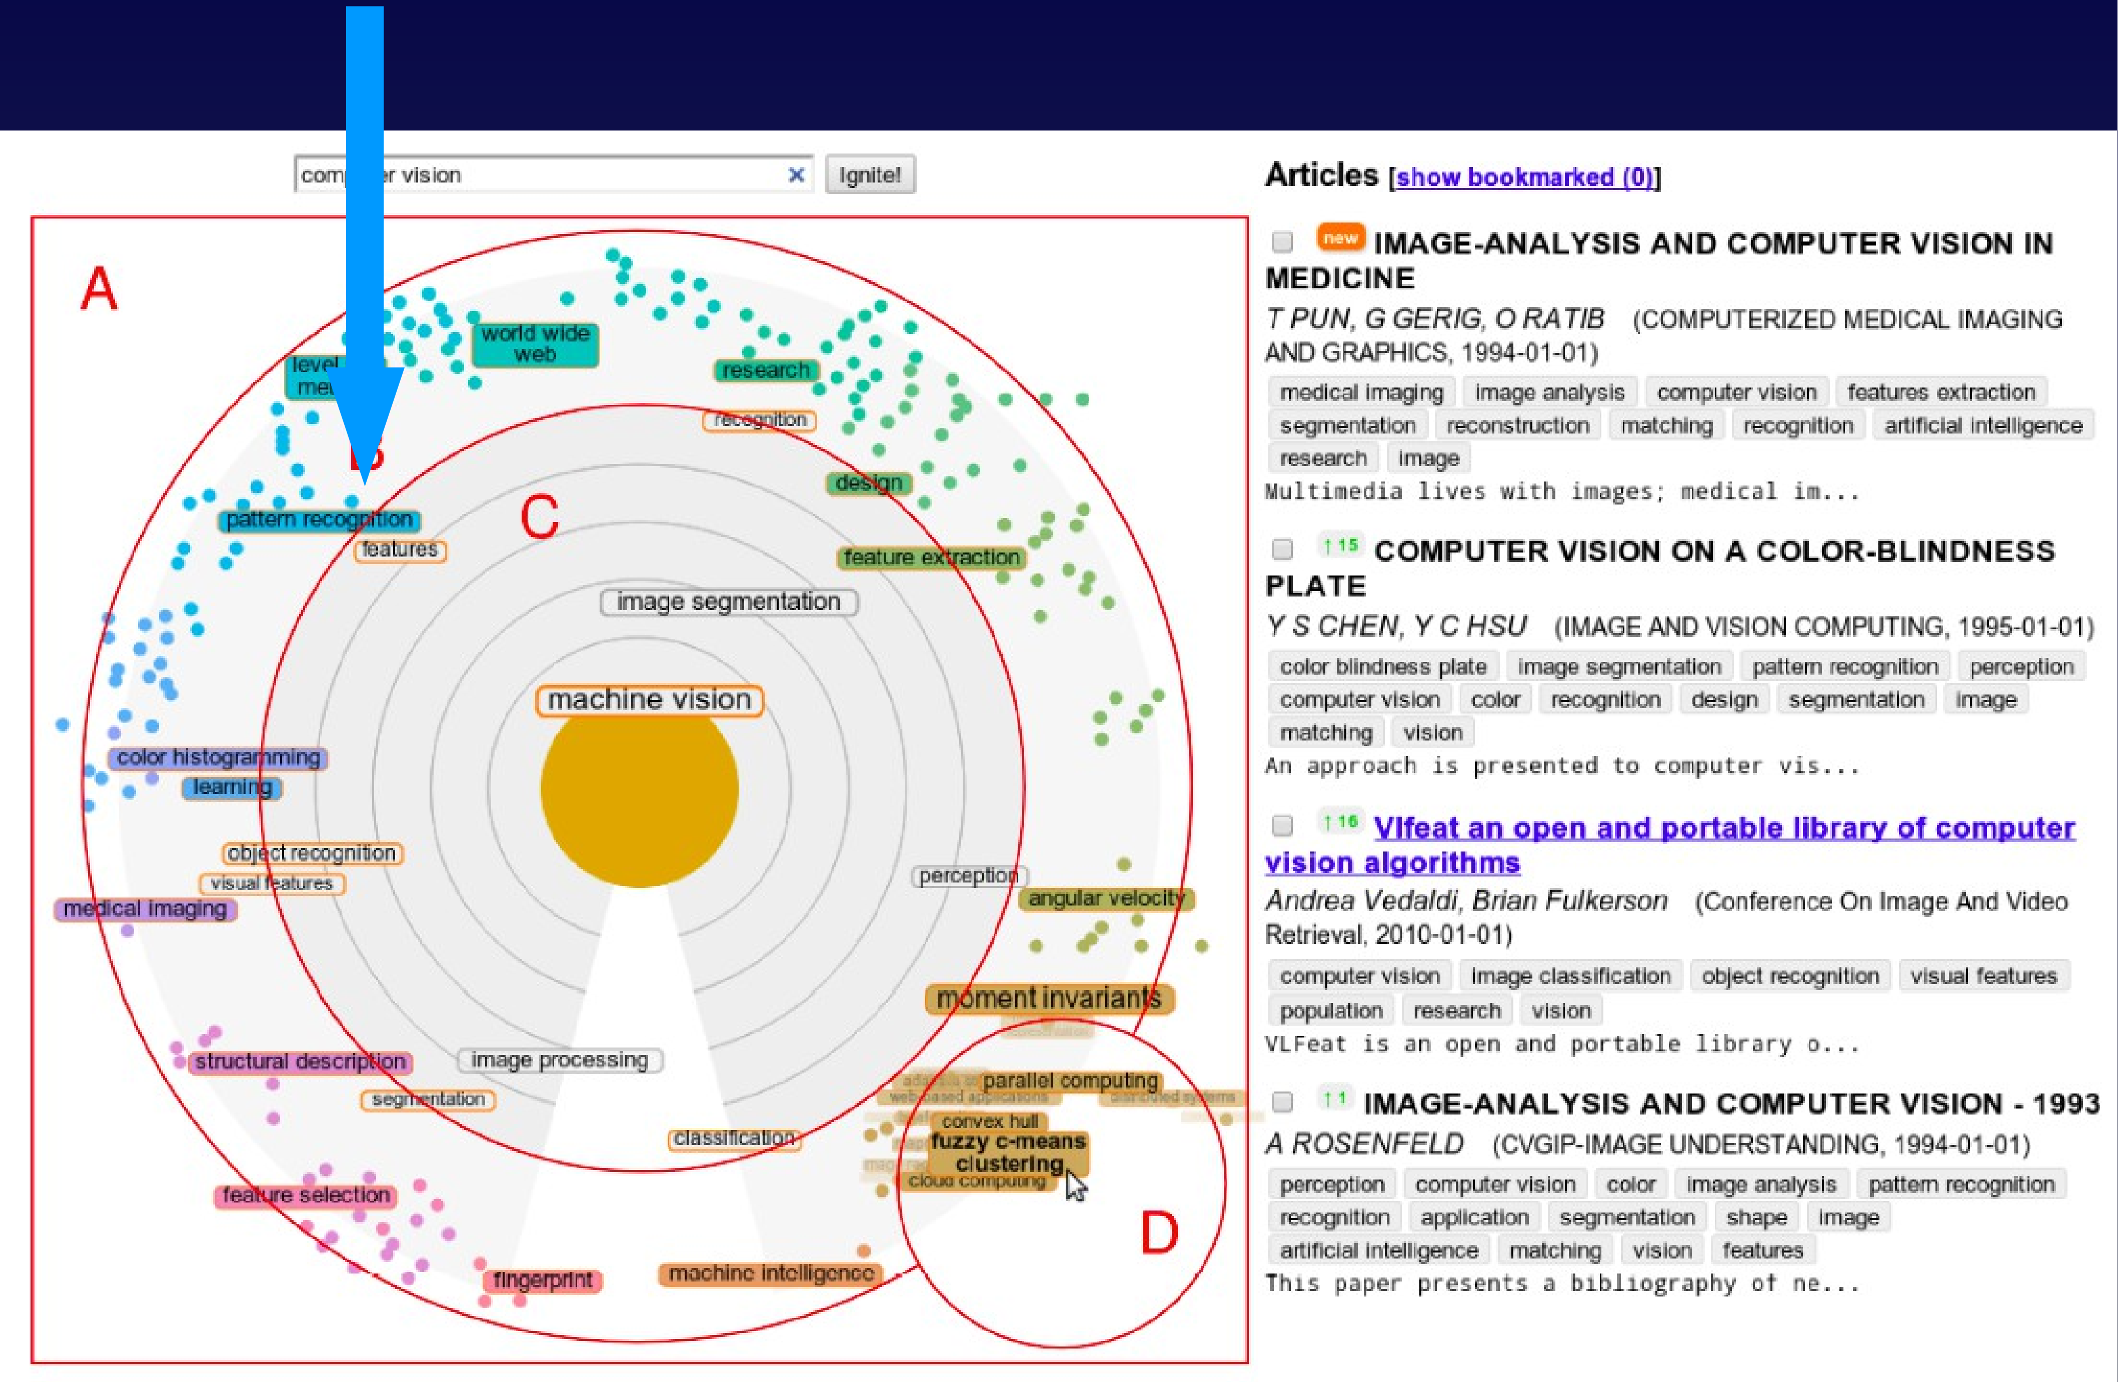

Predicted intents (help users to find directions on the radar to move away from their currently estimated intent). Can be inspected by moving the mouse as a fisheye lens.

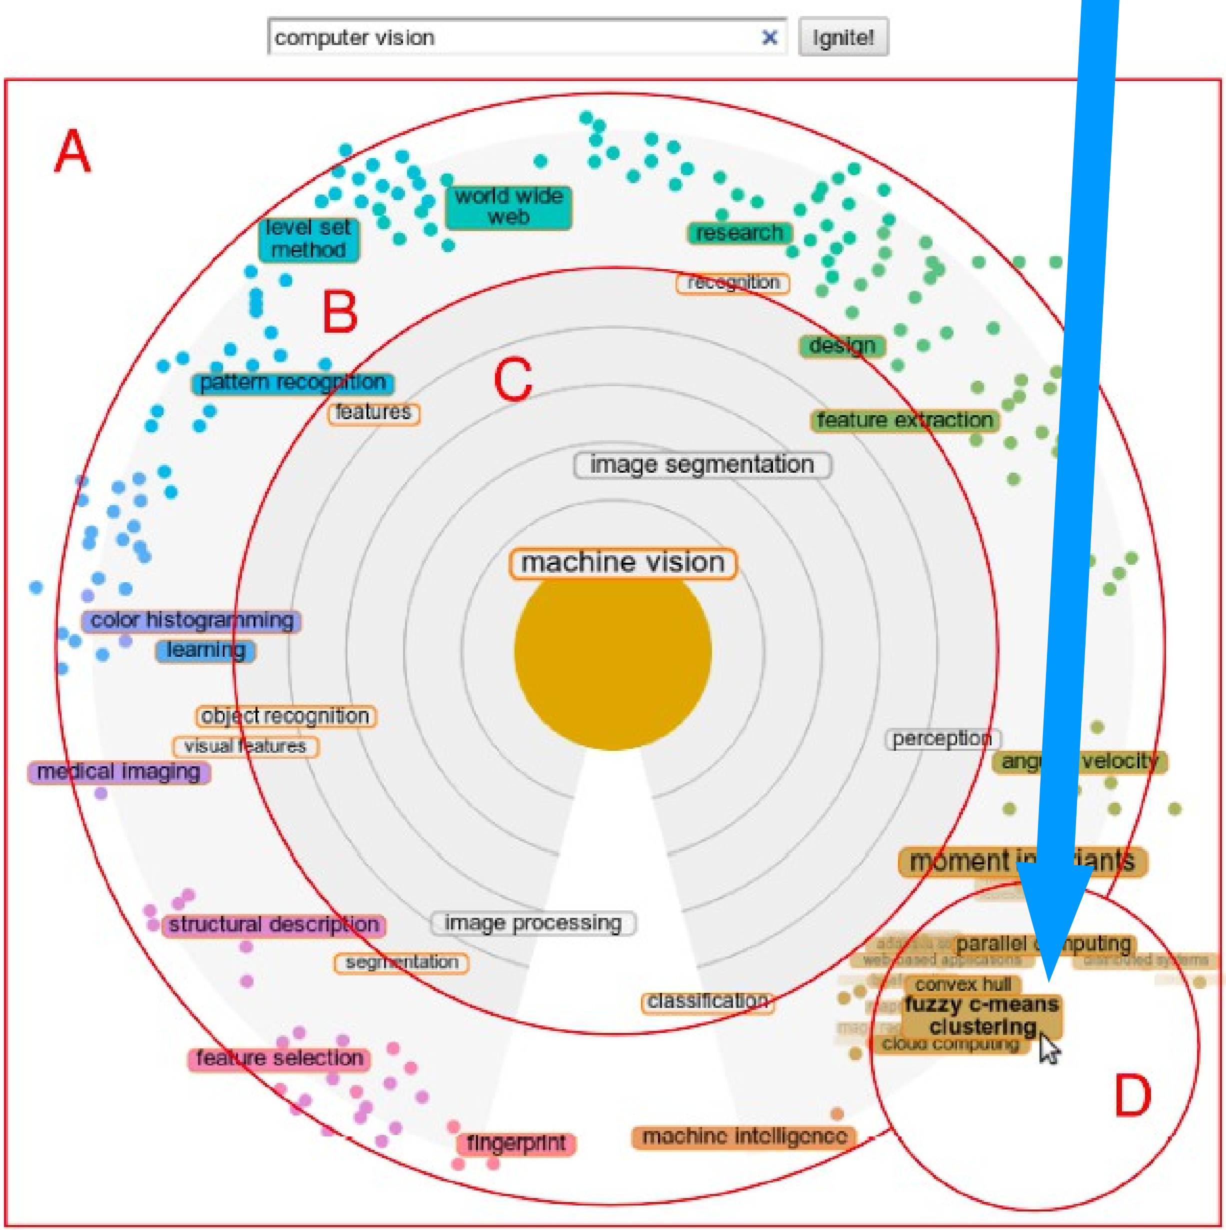

Articles [show bookmarked (0)]

new **IMAGE-ANALYSIS AND COMPUTER VISION IN MEDICINE**
T PUN, G GERIG, O RATIB (COMPUTERIZED MEDICAL IMAGING AND GRAPHICS, 1994-01-01)
medical imaging | image analysis | computer vision | features extraction | segmentation | reconstruction | matching | recognition | artificial intelligence | research | image
`Multimedia lives with images; medical im...`

↑15 **COMPUTER VISION ON A COLOR-BLINDNESS PLATE**
Y S CHEN, Y C HSU (IMAGE AND VISION COMPUTING, 1995-01-01)
color blindness plate | image segmentation | pattern recognition | perception | computer vision | color | recognition | design | segmentation | image | matching | vision
`An approach is presented to computer vis...`

↑16 **Vlfeat an open and portable library of computer vision algorithms**
Andrea Vedaldi, Brian Fulkerson (Conference On Image And Video Retrieval, 2010-01-01)
computer vision | image classification | object recognition | visual features | population | research | vision
`VLFeat is an open and portable library o...`

↑1 **IMAGE-ANALYSIS AND COMPUTER VISION - 1993**
A ROSENFELD (CVGIP-IMAGE UNDERSTANDING, 1994-01-01)
perception | computer vision | color | image analysis | pattern recognition | recognition | application | segmentation | shape | image | artificial intelligence | matching | vision | features
`This paper presents a bibliography of ne...`

The user can give **feedback** by dragging concepts towards the center. + traditional interactions: bookmarking documents, viewing abstracts/clicking links, or starting over by typing new search terms.)

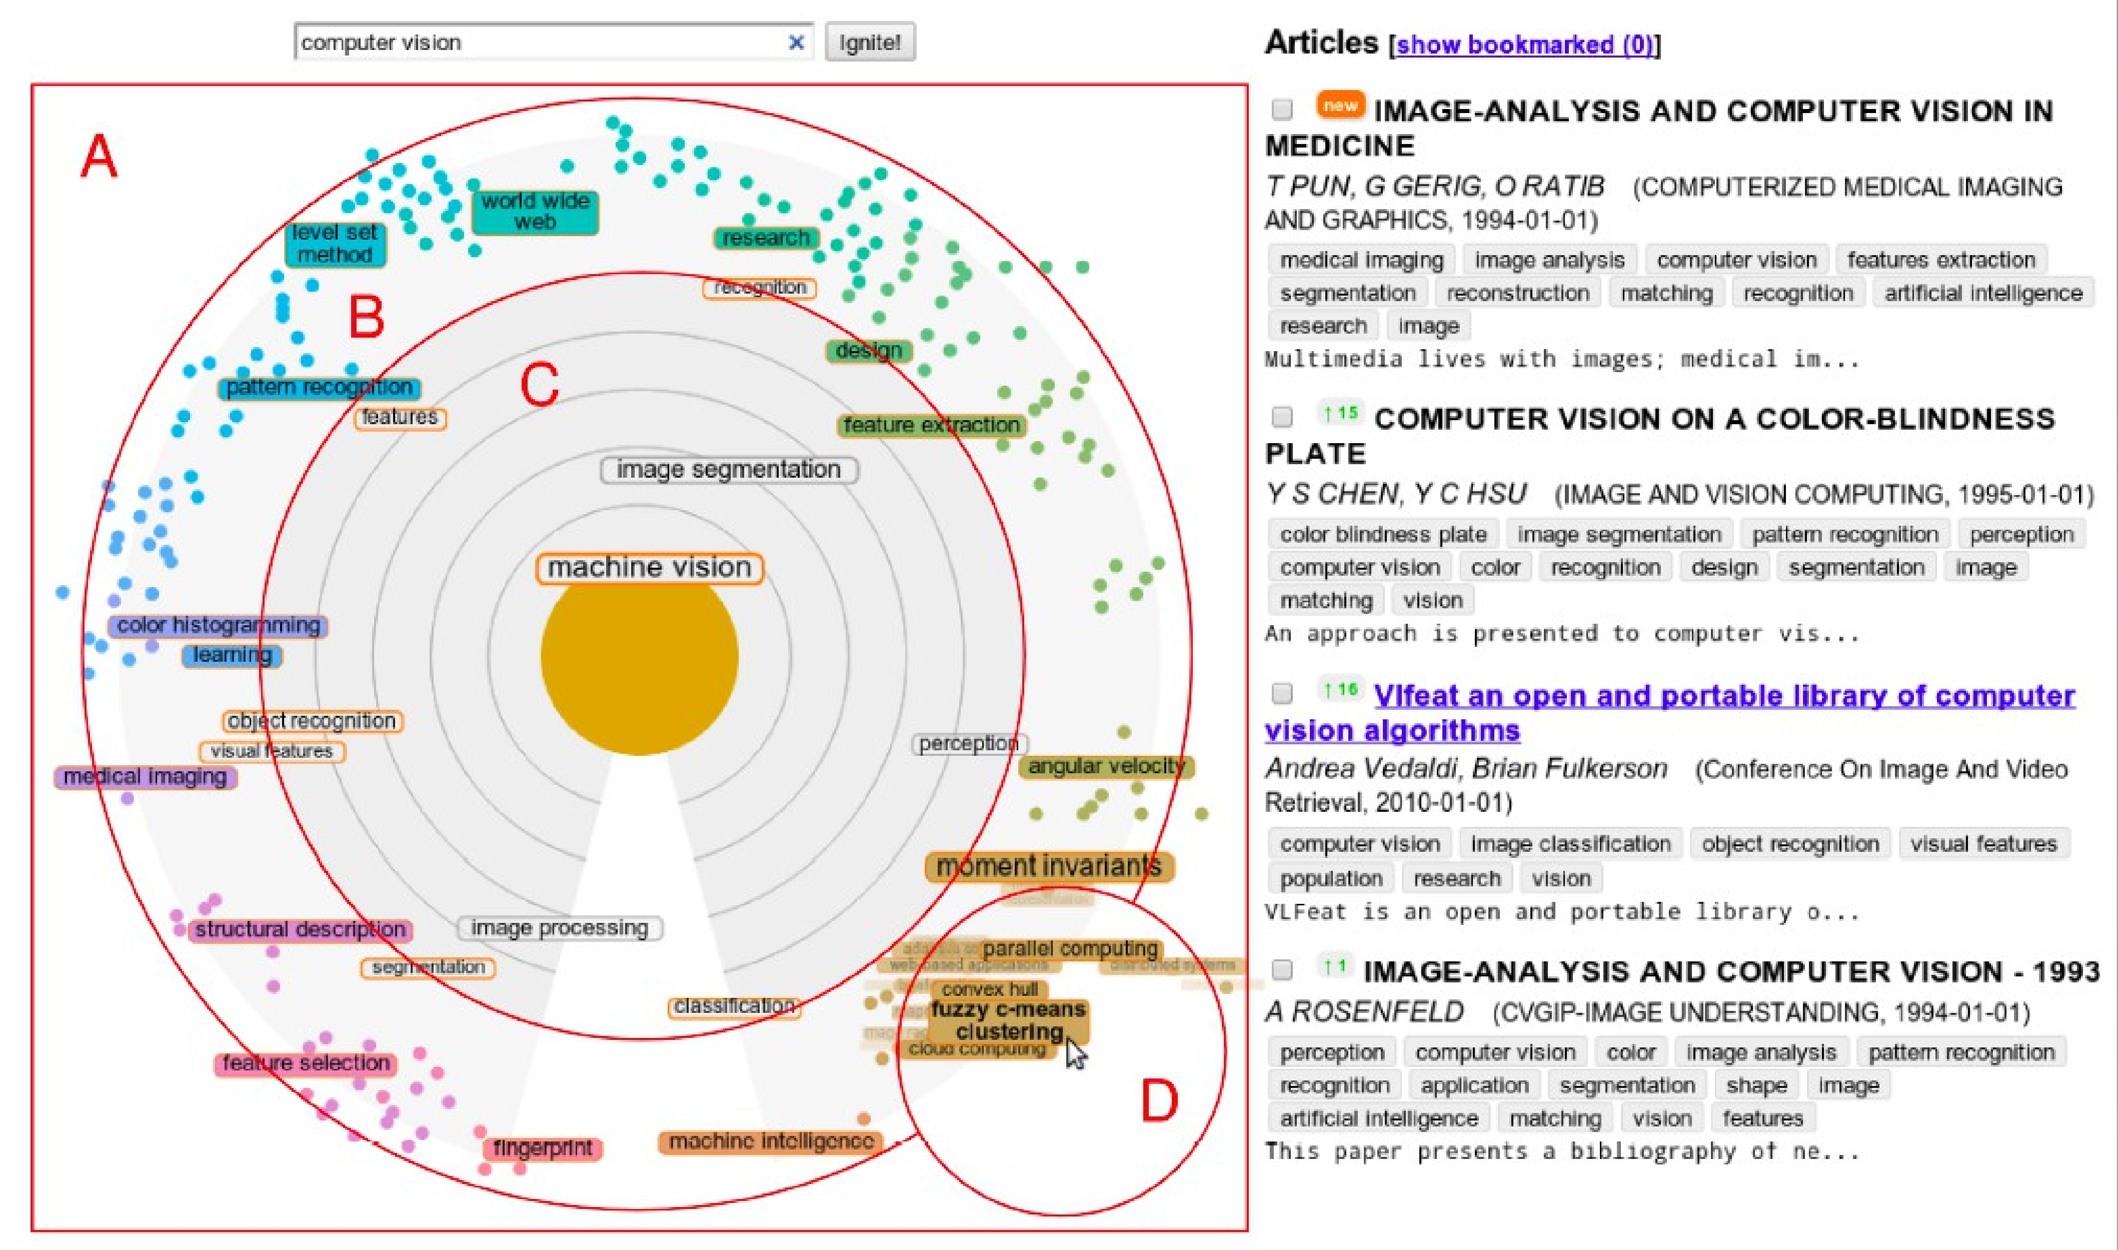

The user can give **feedback** by dragging concepts towards the center. + traditional interactions: bookmarking documents, viewing abstracts/clicking links, or starting over by typing new search terms.)

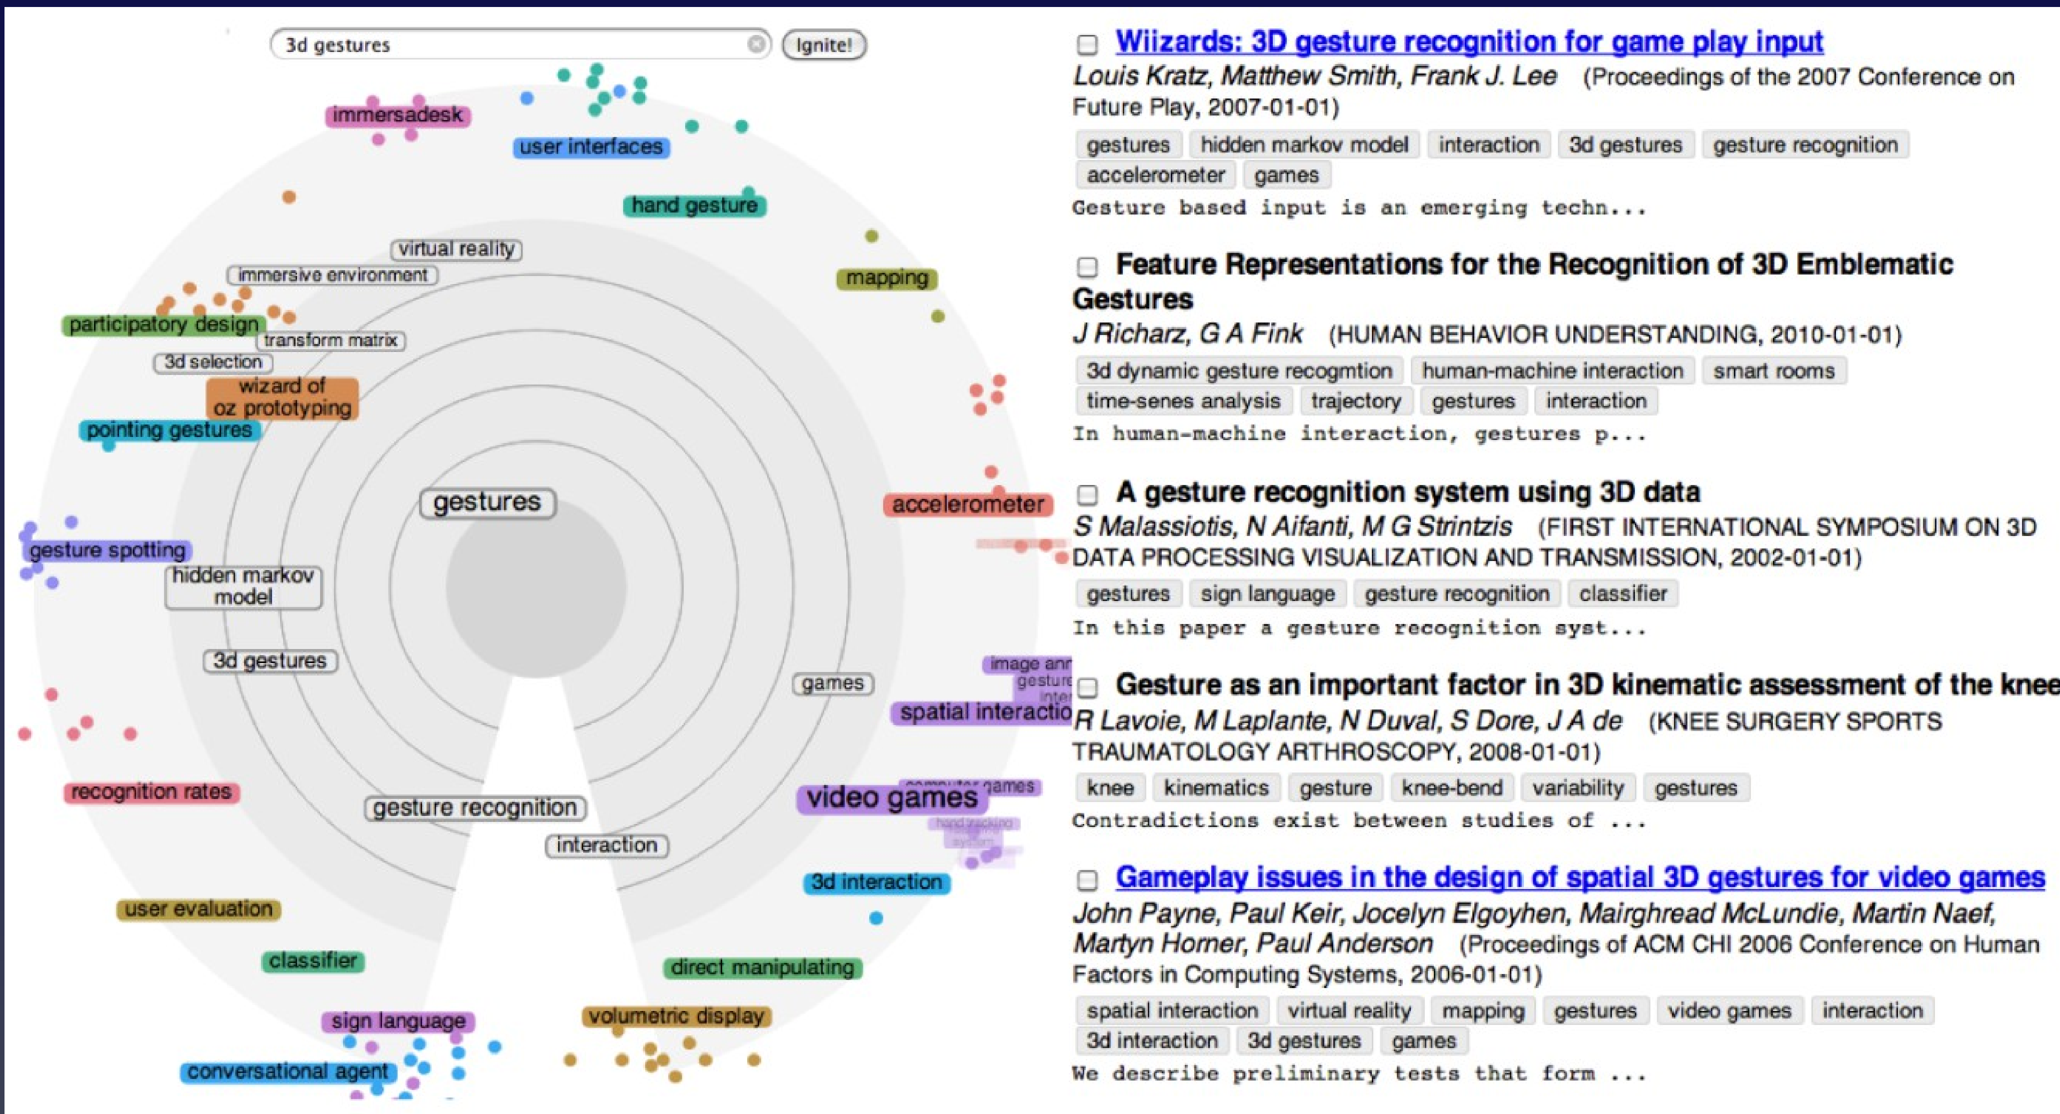

The user can give **feedback** by dragging concepts towards the center. + traditional interactions: bookmarking documents, viewing abstracts/clicking links, or starting over by typing new search terms.)

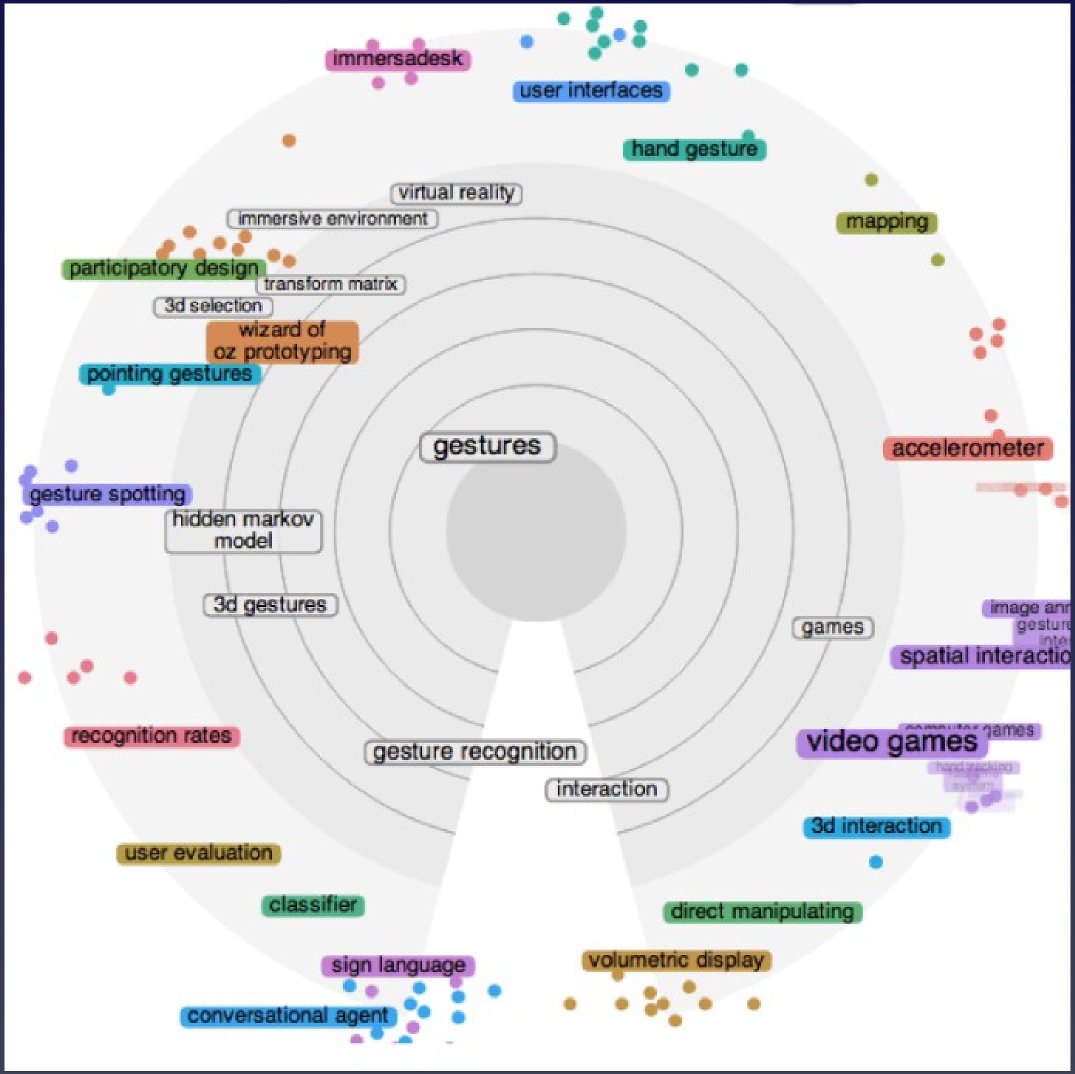

The user can give **feedback** by dragging concepts towards the center. + traditional interactions: bookmarking documents, viewing abstracts/clicking links, or starting over by typing new search terms.)

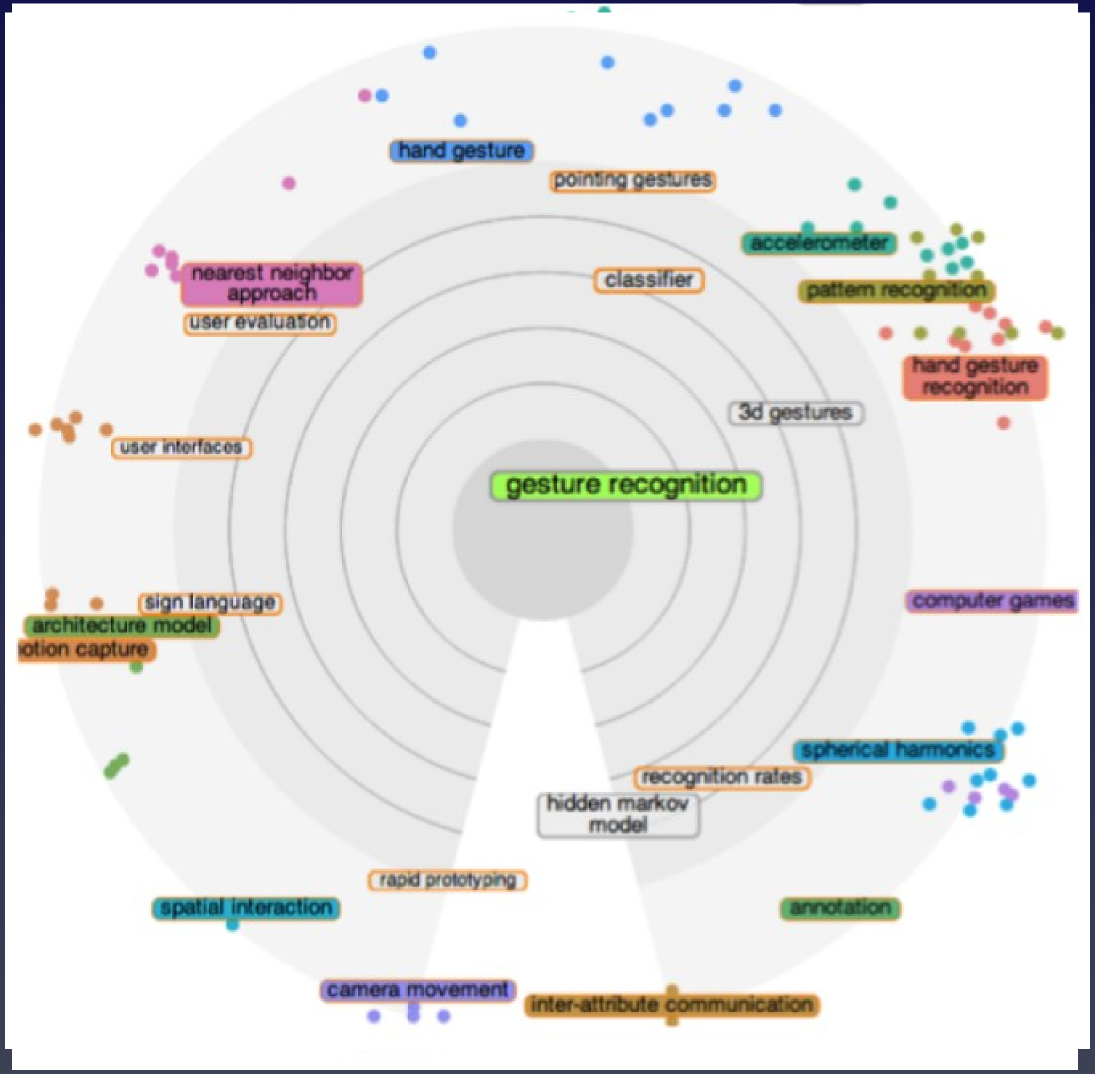

Behind the Scenes

Machine learning for the Intent Radar

- Learning of user's search intents during interactive search
- Based on a **retrieval model** and **intent model**; layout based on **neighbor embedding**

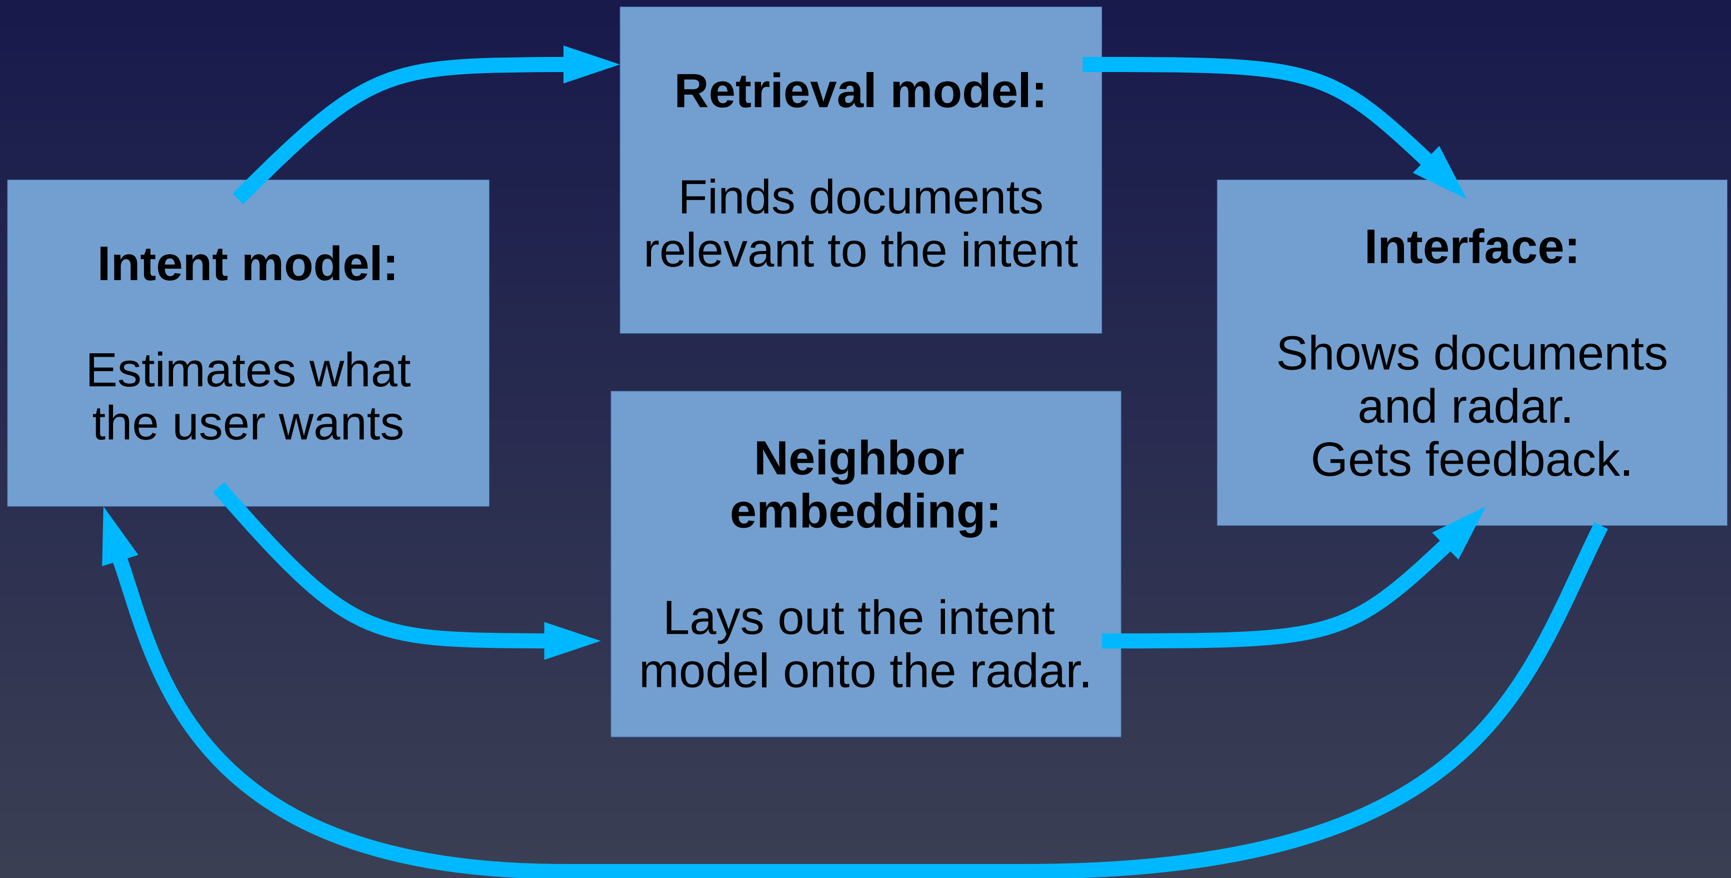

Retrieval Model

- Estimates probability of relevant documents based on estimates of the intent model

- We use the language modeling approach of information retrieval

- Unigram language model, Bayesian Dirichlet Smoothing

Retrieval Model

- Estimates probability of relevant documents based on estimates of the intent model

- We use the language modeling approach of information retrieval

- Unigram language model, Bayesian Dirichlet Smoothing

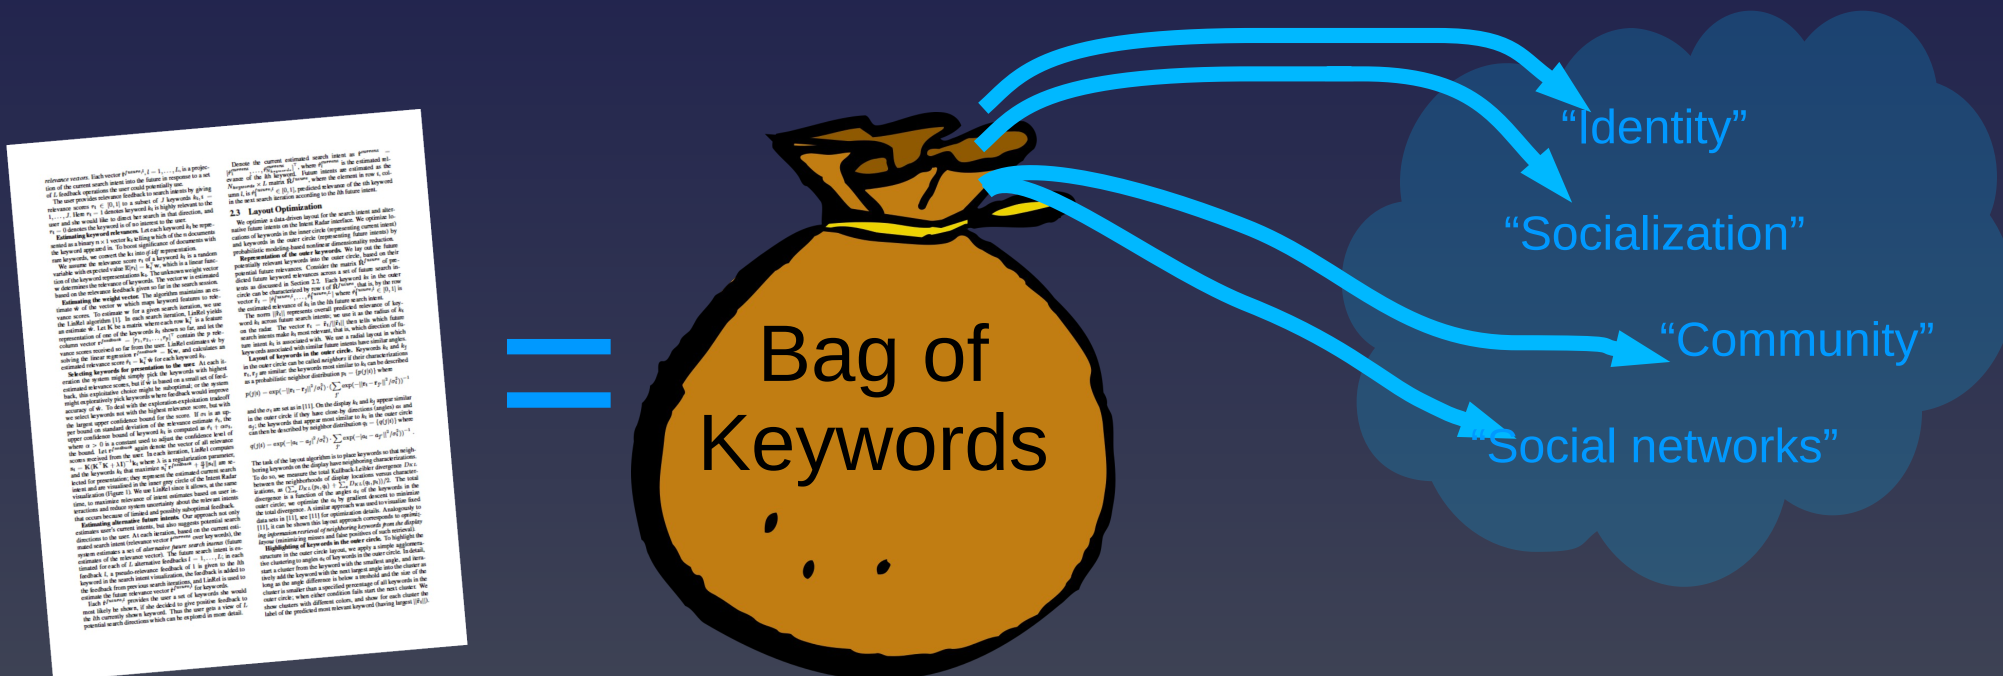

Retrieval Model

$$\hat{P}(\hat{\mathbf{v}}|M_{d_j}) = \prod_{i=1}^{|\hat{\mathbf{v}}|} \hat{P}_{mle}(k_i|M_{d_j})^{\hat{v}_i}$$

User model = sample of desired document

Count of keyword i in document j

Proportion of keyword i in the corpus

$$\hat{P}_{mle}(k_i|M_{d_j}) = \frac{c(k_i|d_j)+\mu p(k_i|C)}{\sum_k c(k|d_j)+\mu}$$

- Representation of a desired document: estimated by the intent model

- Rank documents by their probability to generate the desired document

- Expose user to more novel documents: sample documents from ranked list by Dirichlet Sampling: show documents with highest sampled values

$$f_j \sim Gamma(\alpha_j, 1) = f_j^{\alpha_j - 1} e^{-f_j} / \Gamma(\alpha_j)$$

Increase weight by 1 in each iteration for each document where at least one keyword got positive feedback

Intent Model

- Estimates **current search intent** and **alternative future intents** that could occur in response to user feedback

- We use the **LinRel algorithm**. Yields estimate of keyword weights in each iteration, based on interaction history.

- Observations = relevance scores given by user to keywords. Assumption: expected relevance = linear function of what documents the keyword appears in.

$$\mathbf{r}^{feedback} = [r_1, r_2, \ldots, r_p]^\top$$

Feedback scores in [0,1] given so far to a subset of keywords

$$\mathbf{r}^{feedback} = \mathbf{K}\mathbf{w}$$

Model feedback: regression based on what documents they appeared in (matrix **K**)

$$\hat{r}_i = \mathbf{k}_i^\top \hat{\mathbf{w}}$$

Use model to estimate relevance of the rest of the keywords

Intent Model

- **Choose keywords to show to the user:** the keywords represent the estimated current intent of the user. Choosing just keywords with highest estimated relevance would be pure exploitation, could trap users. Instead, control exploration-exploitation tradeoff. Show keywords with highest **upper confidence bound** of their relevance score

$$\mathbf{s}_i = \mathbf{K}(\mathbf{K}^\top\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{k}_i$$

Linear estimator of relevance of new keyword i based on previous feedback

$$\mathbf{s}_i^\top \mathbf{r}^{feedback} + \frac{\alpha}{2}||\mathbf{s}_i||$$

Upper confidence bound of the relevance, considering the previous feedback as independent random variables.

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

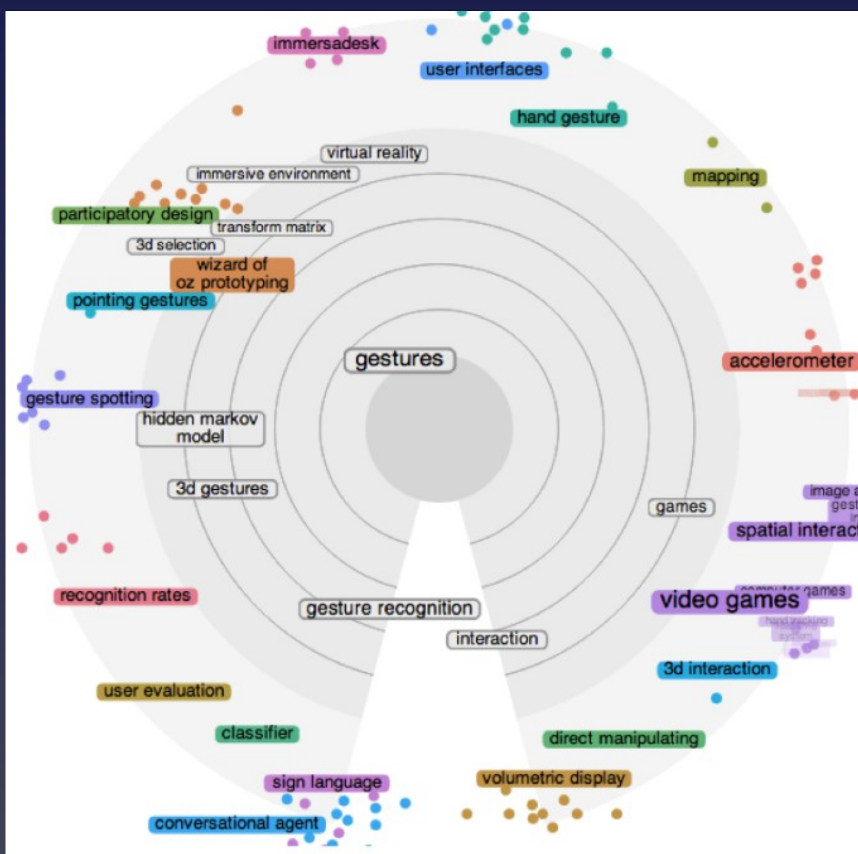

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

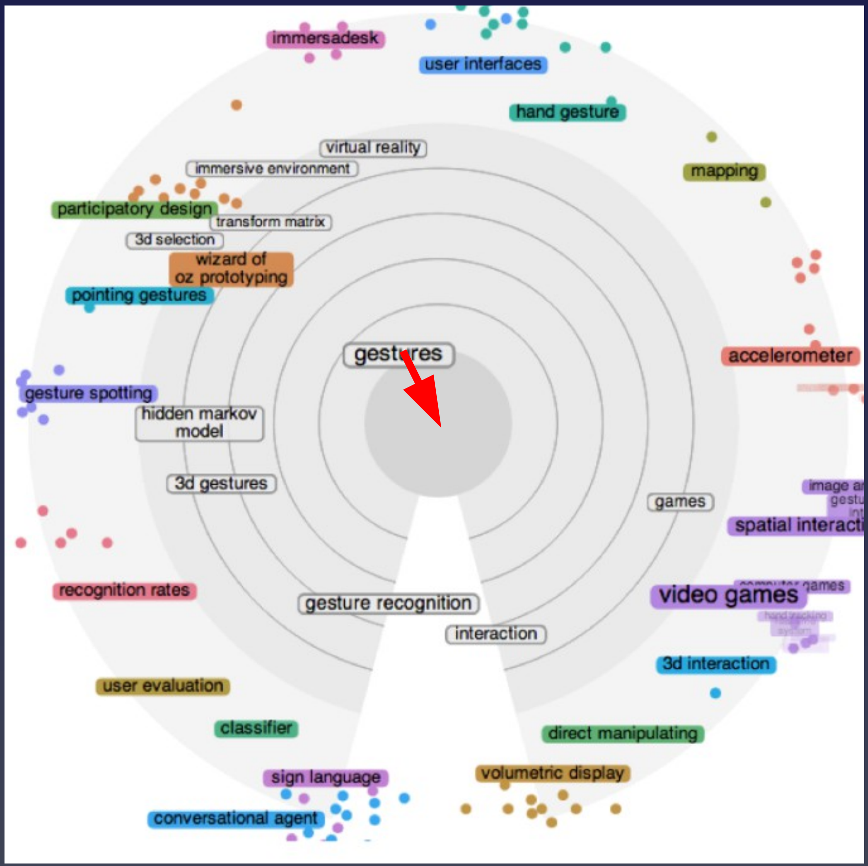

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

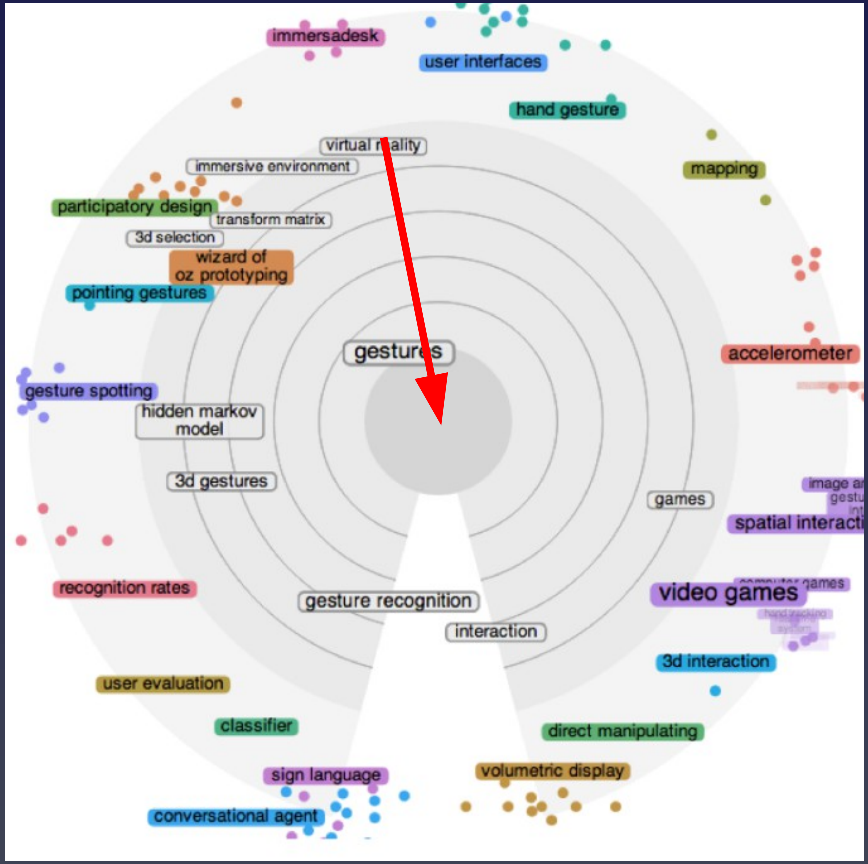

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

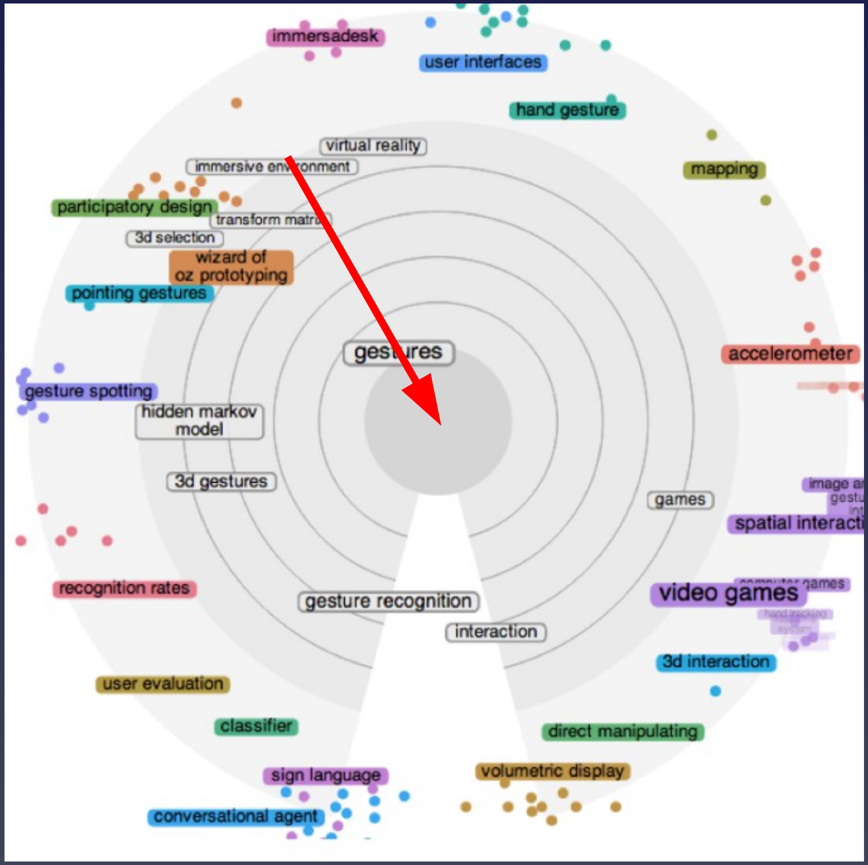

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

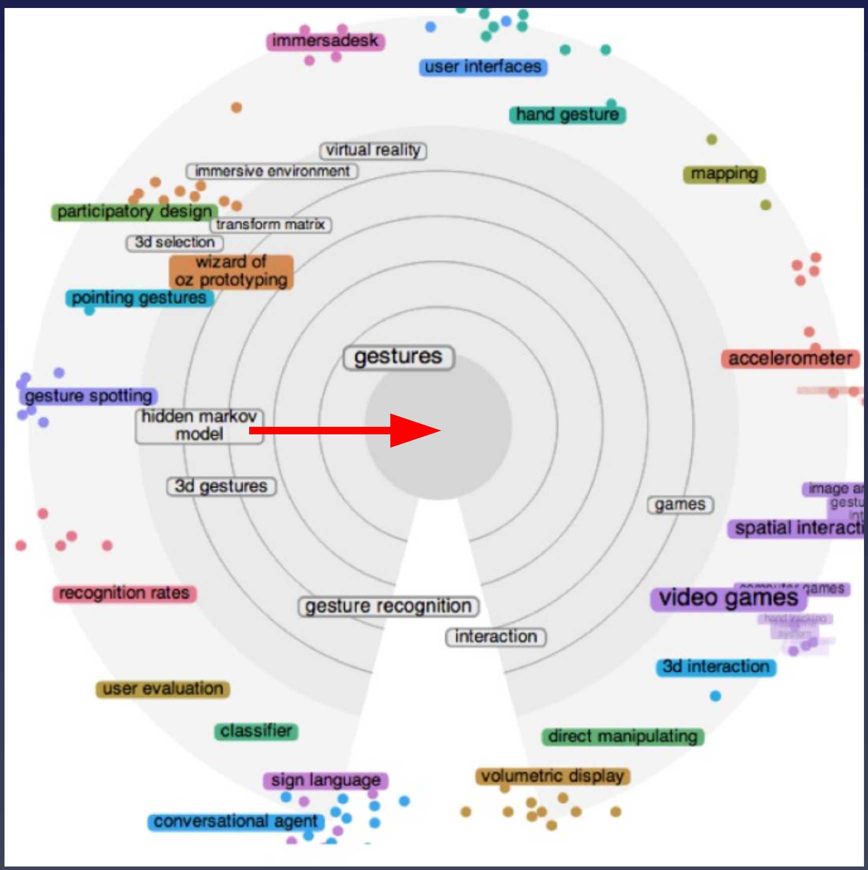

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

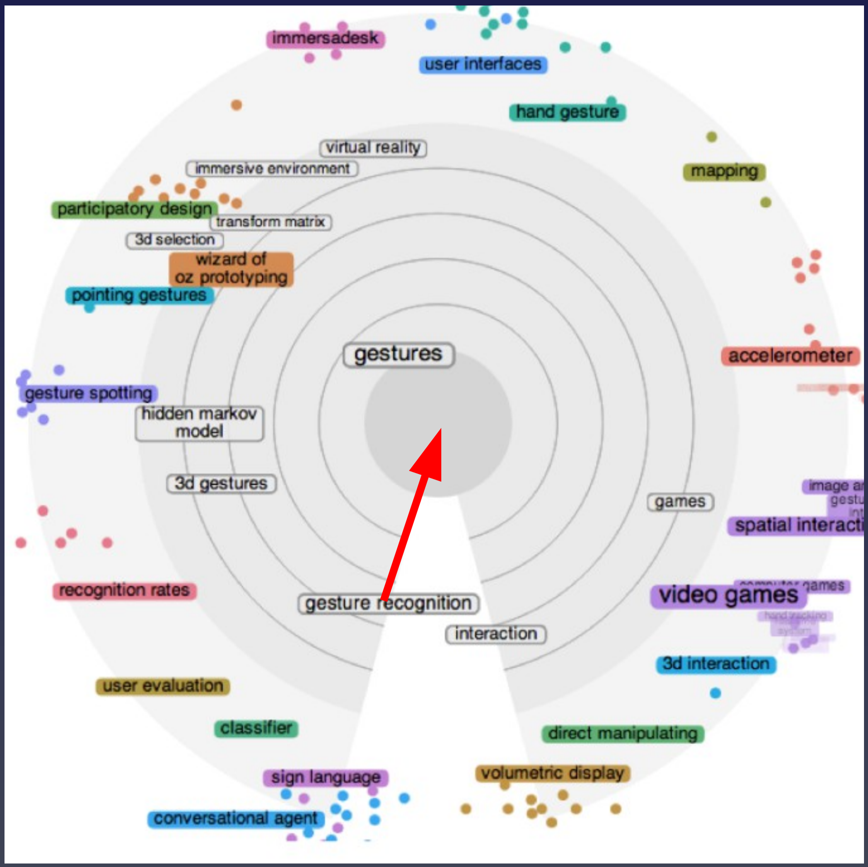

Intent Model

- **Estimate alternative future intents:** estimate future intent for L alternative feedbacks.

  - In each alternative, pseudo-relevance feedback 1 is given to the l:th keyword, adding to feedback from previous iterations.

  - For each alternative, LinRel is then used to estimate a new future relevance vector for all keywords = representation of an alternative future intent.

  - Collect future intents into a matrix:

$$N_{keywords} \times L \text{ matrix } \hat{\mathbf{R}}^{future}$$

Layout of Intents by Nonlinear Dimensionality Reduction

- Radial position of each keyword = current estimated relevance

- Angles are used to represent directions of future intent

- Each keyword is represented by its relevances in all future intents (high-dimensional representation):

$$\tilde{\mathbf{r}}_i = [\hat{r}_i^{future,l}, \ldots, \hat{r}_i^{future,L}]$$

Relevance of keyword i in future intent L

$$\bar{\mathbf{r}}_i = \tilde{\mathbf{r}}_i / ||\tilde{\mathbf{r}}_i||$$

Normalized vector, tells which future intents (or feedbacks) make keyword i most relevant

- Layout is **optimized for retrieval of keywords** with similar relevance in future intents, by **nonlinear dimensionality reduction.** We use a well-performing approach optimized for information retrieval, details in Wednesday's talk.

Layout of Intents

- Cluster keywords in outer circle, highlight with colors.

- Place inner circle keywords i at highest mode of angles of future keywords j, weighted by their relevance given feedback on i. Angle of inner keyword i tells which future keywords become relevant by interacting with i.

Experiments

User Experiment - Questions

Task-based user experiment to investigate effect of interactive intent modeling on exploratory search.

Research questions:

1. **User task performance** – does the interaction paradigm lead to better responses in user tasks?

2. **Quality of displayed information** – does the paradigm help users reach high quality information in response to interactions?

3. **Interaction support for directing exploration** – does the paradigm elicit more interaction from the user, is it targeted to relevant options? Does the paradigm let the user explore novel information more than a conventional system?

User Experiment - Setup

- Two search tasks: prepare materials to write essay on “semantic search” or “robotics”. Answer questions about the topic.

- Users: 30 graduate students: each had 30min to perform the task after 10-min demo

- Data: **50 million scientific documents** from Thomson Reuters, ACM, IEEE, Springer

- Comparison system: TypedQuery – traditional system, no keyword feedback

- Two versions of our system: IntentRadar (full), IntentList (lists estimated relevant keywords, no layout)

User Experiment - Evaluation

Ground truth from experts who evaluated all presented documents and keywords, and user answers: documents rated as relevant, novel, and obvious, keywords as relevant, general, and specific, user answers rated on 5-point Likert scale.

Evaluation measures:

- **User task performance** measured by average score of expert assessment of the written answers.We also measured the number of bookmarked relevant, obvious, and novel documents

- **Quality of displayed information** measured by precision, recall, and F-measure of shown articles and manipulated keywords, with respect to the ground truth categories novel, obvious, and relevant.

- **Interaction support for directing exploration** measured by number and type of interactions (typed query or interaction with intent model), and type of information (novel/obvious) received in response

User Experiment - Results

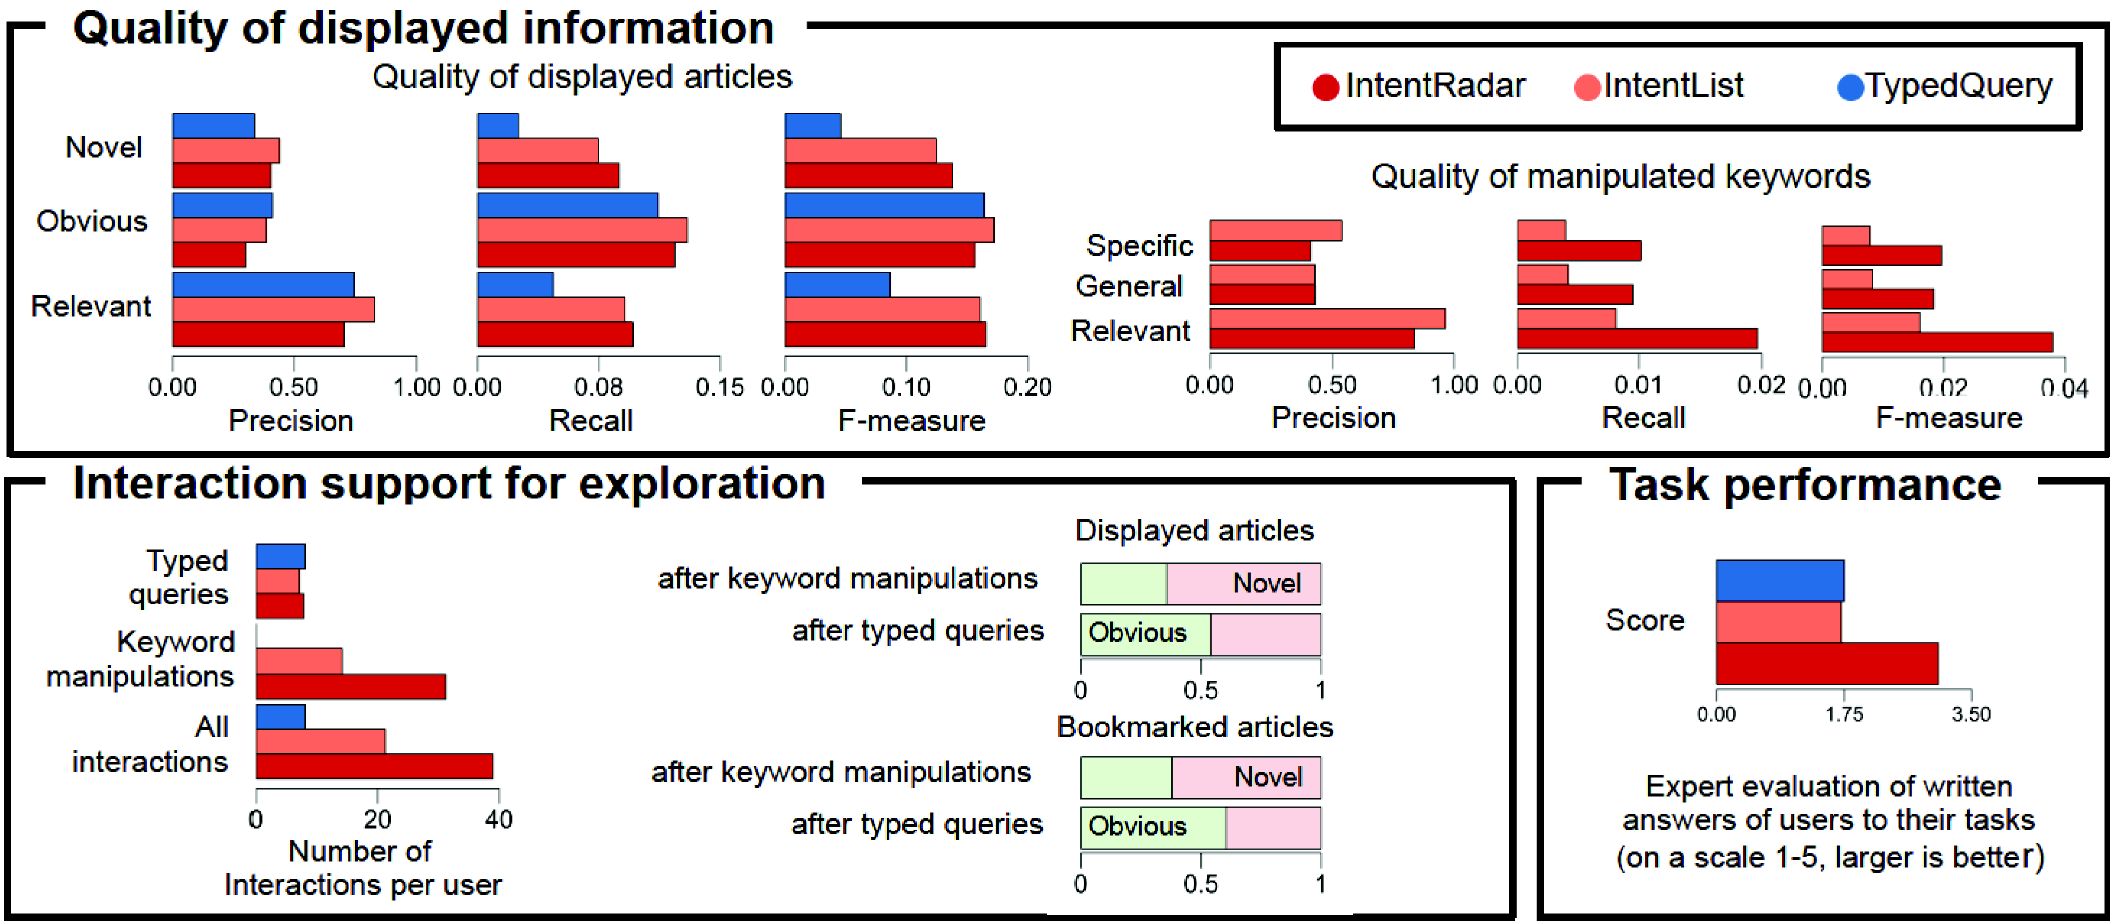

User Experiment - Results

- Users of Intent Radar get significantly better task performance than users of IntentList or TypedQuery

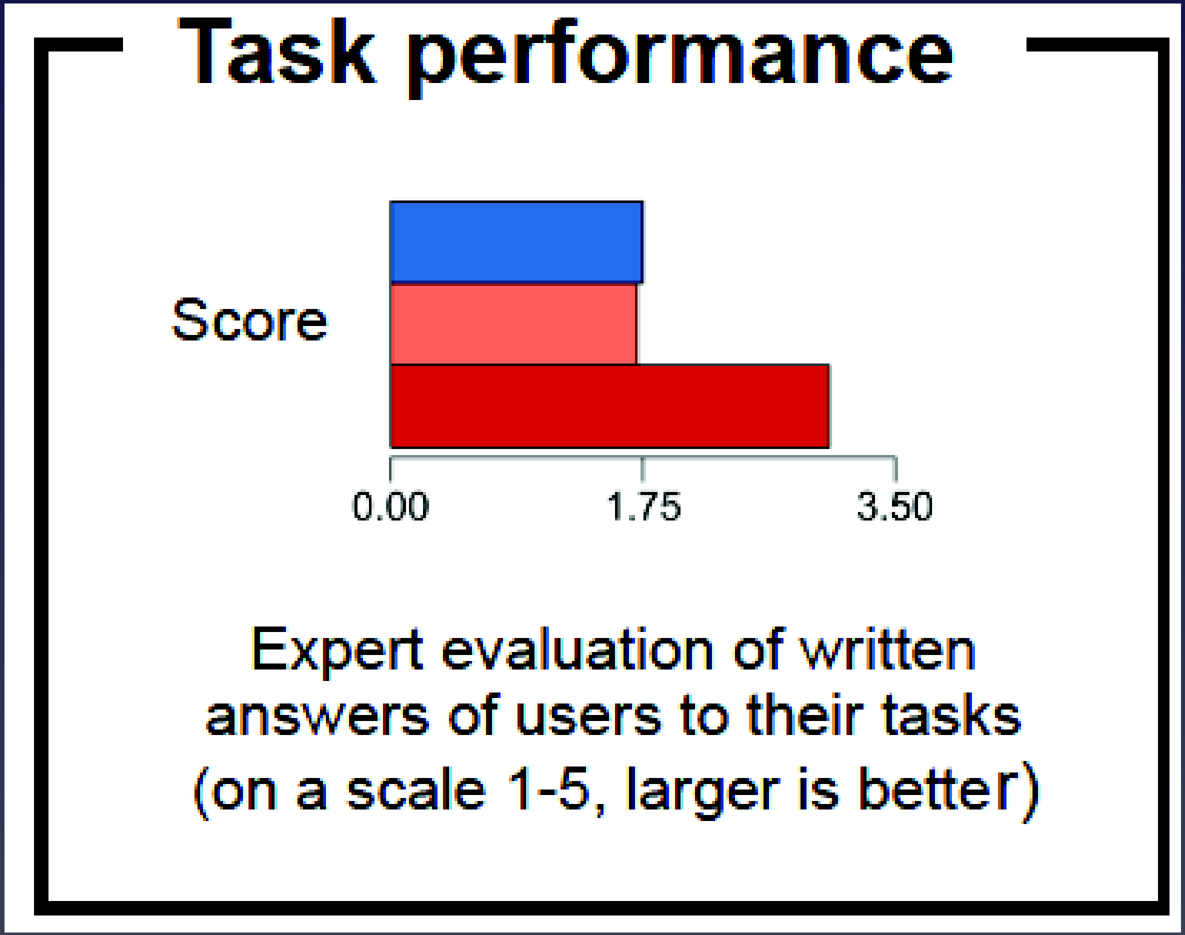

User Experiment - Results

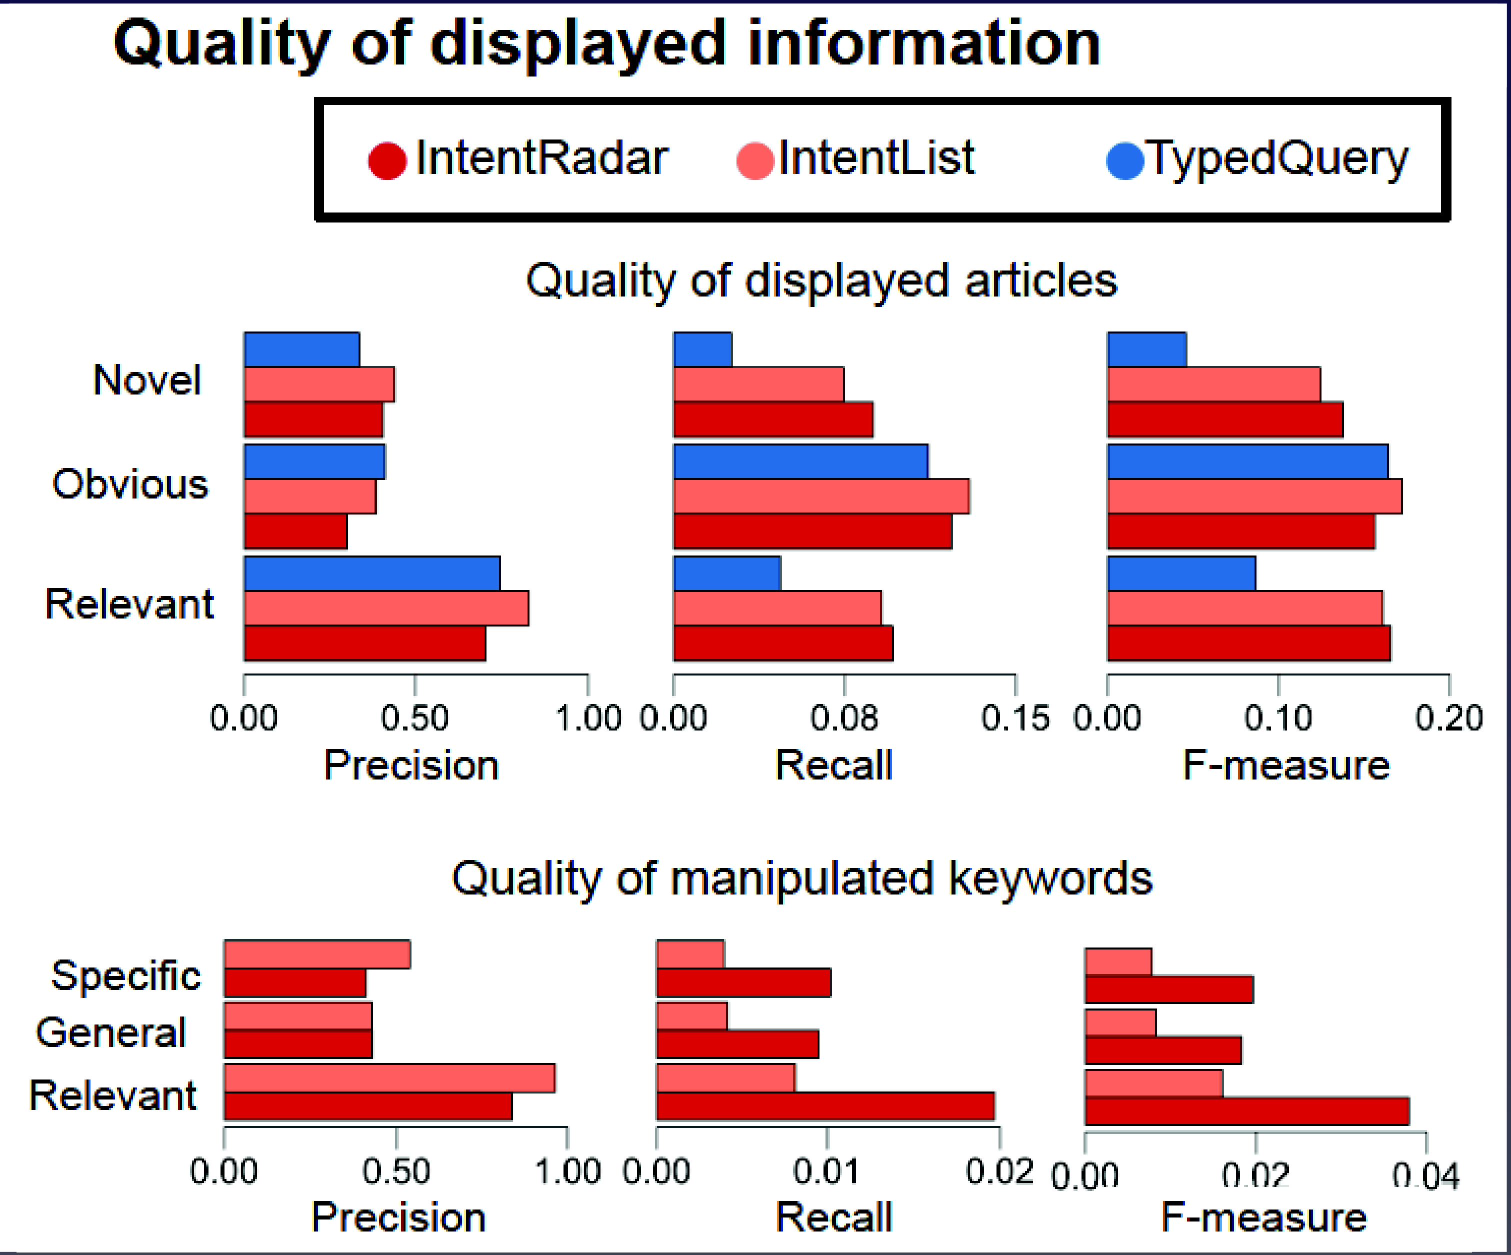

Interactive Intent modeling gets significantly better quality of displayed information than TypedQuery

User Experiment - Results

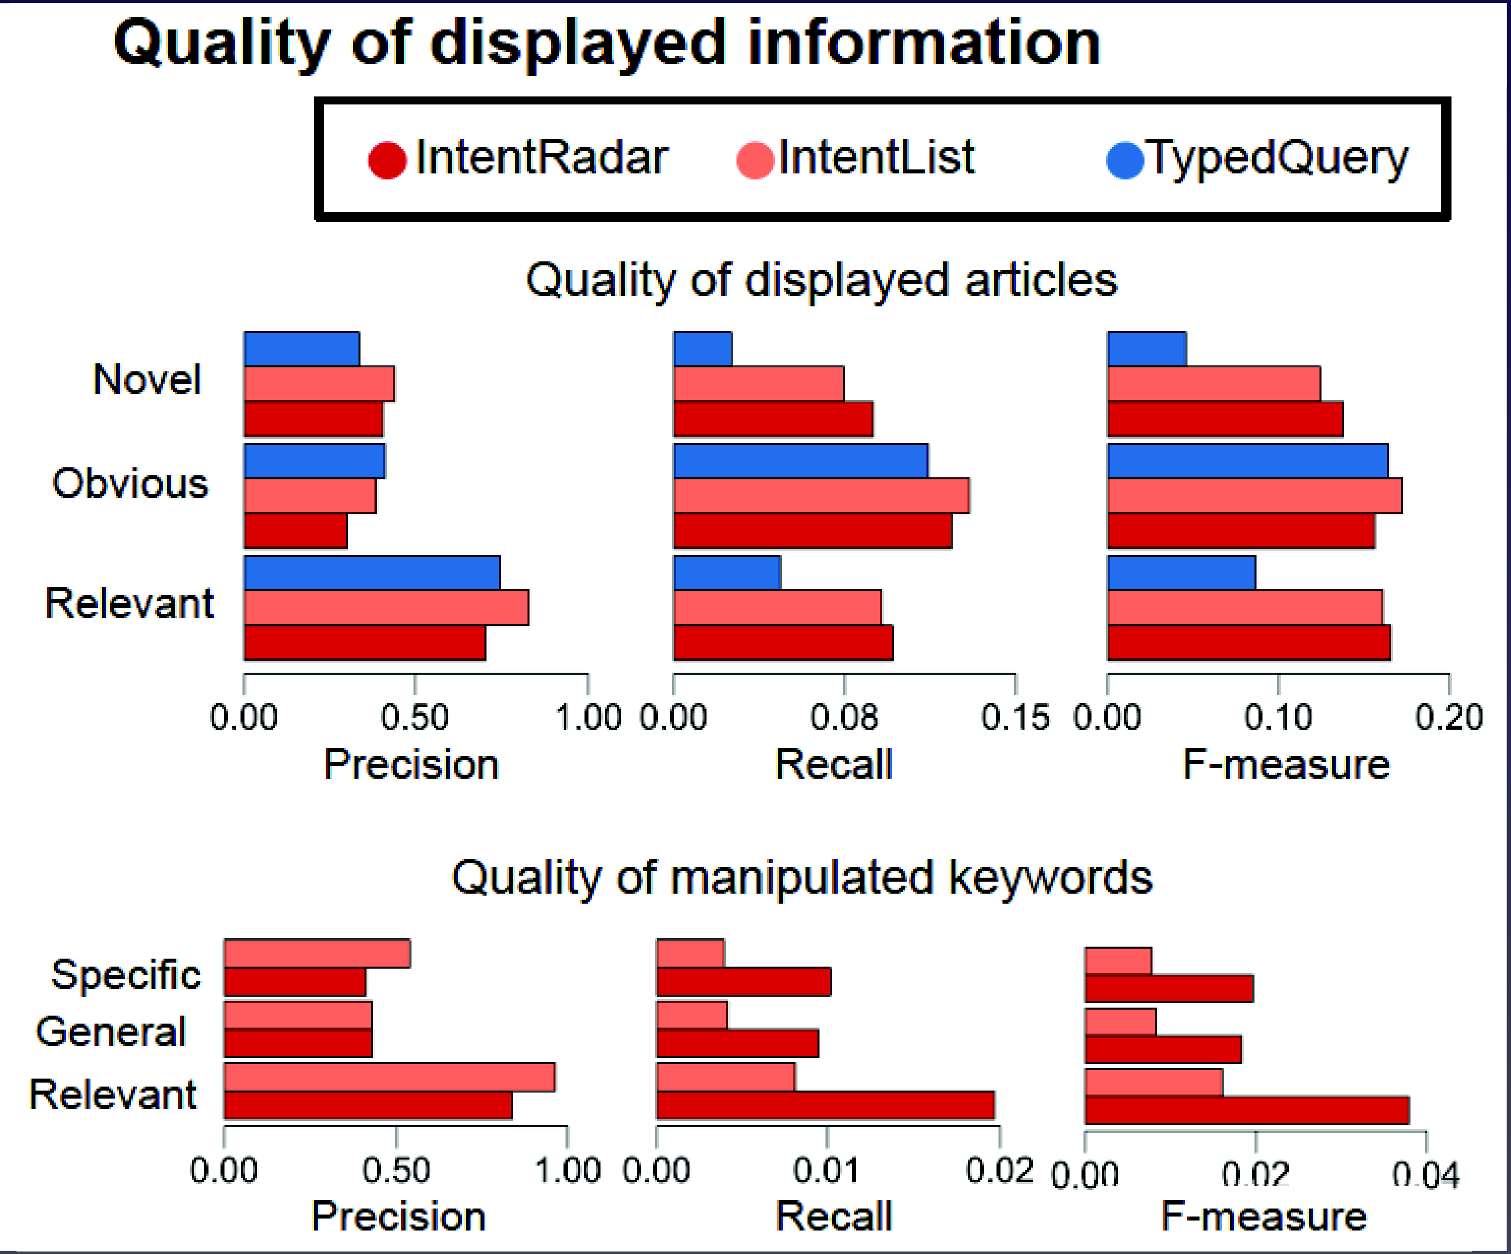

IntentList is slightly better for obvious documents – harder to move from initial context?

User Experiment - Results

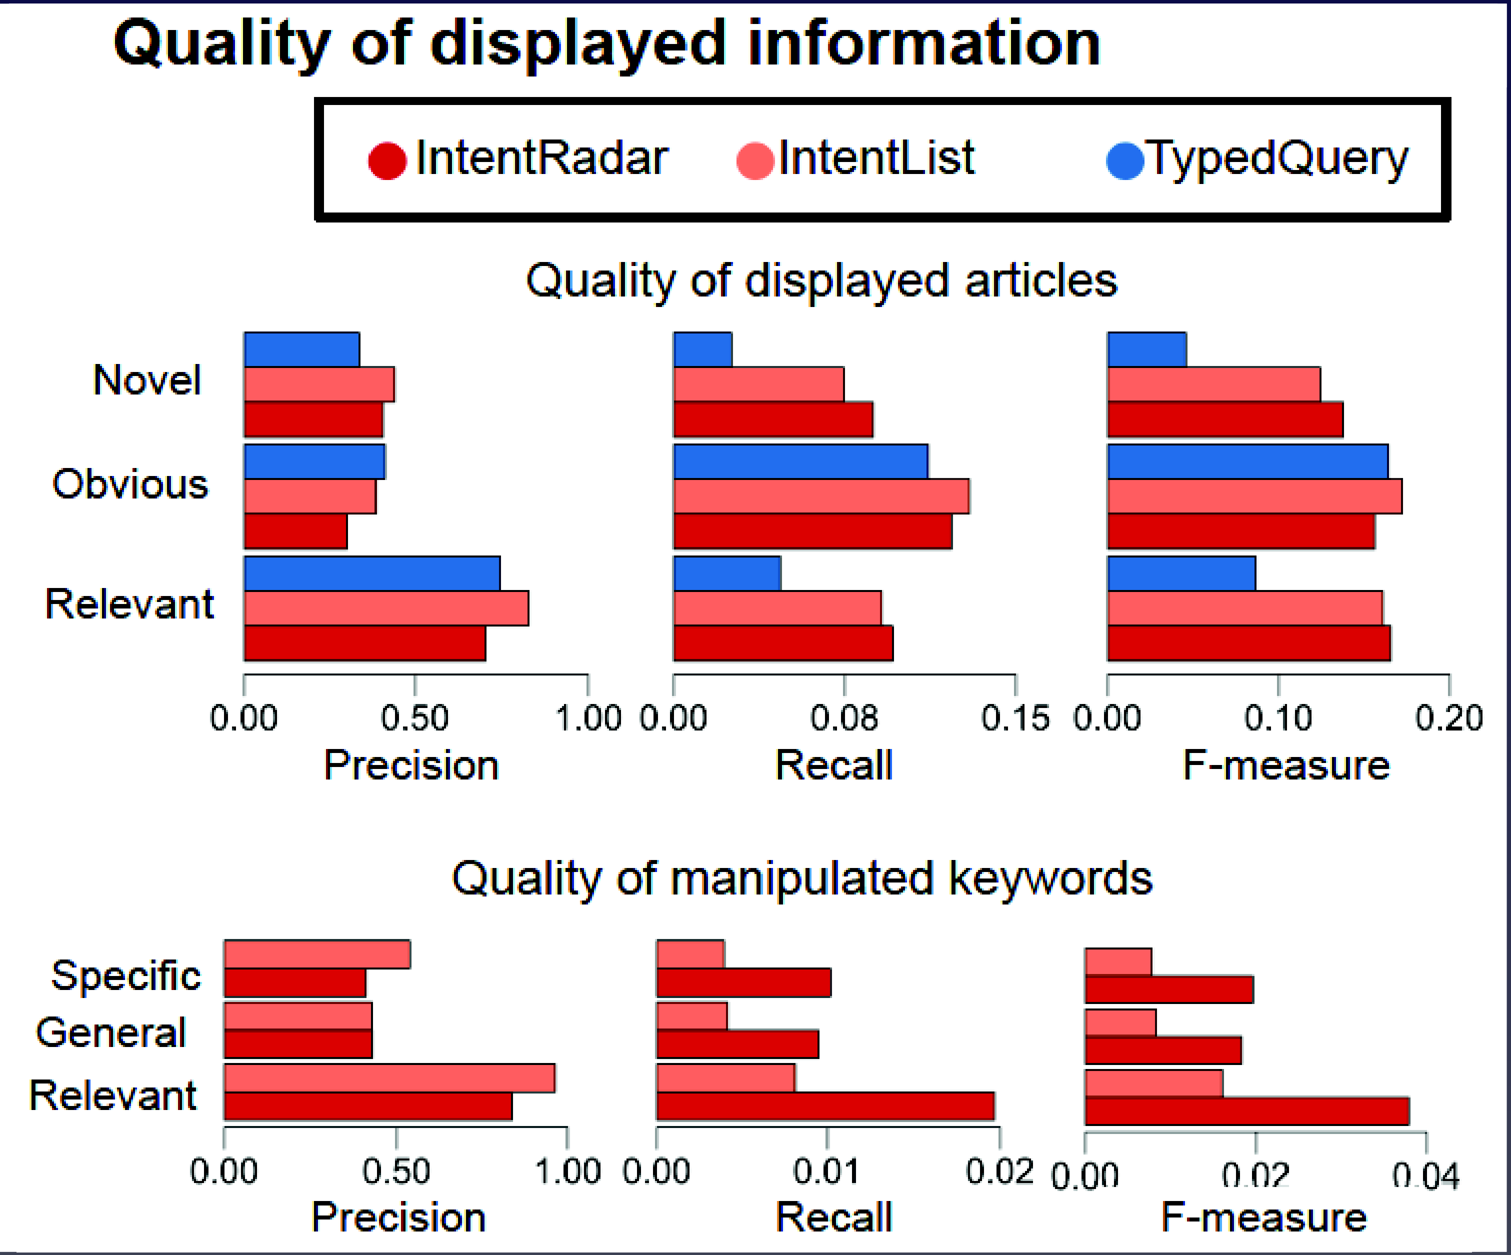

IntentRadar has significantly higher keyword quality than IntentList – made targeting interactions to relevant keywords easier

User Experiment - Results

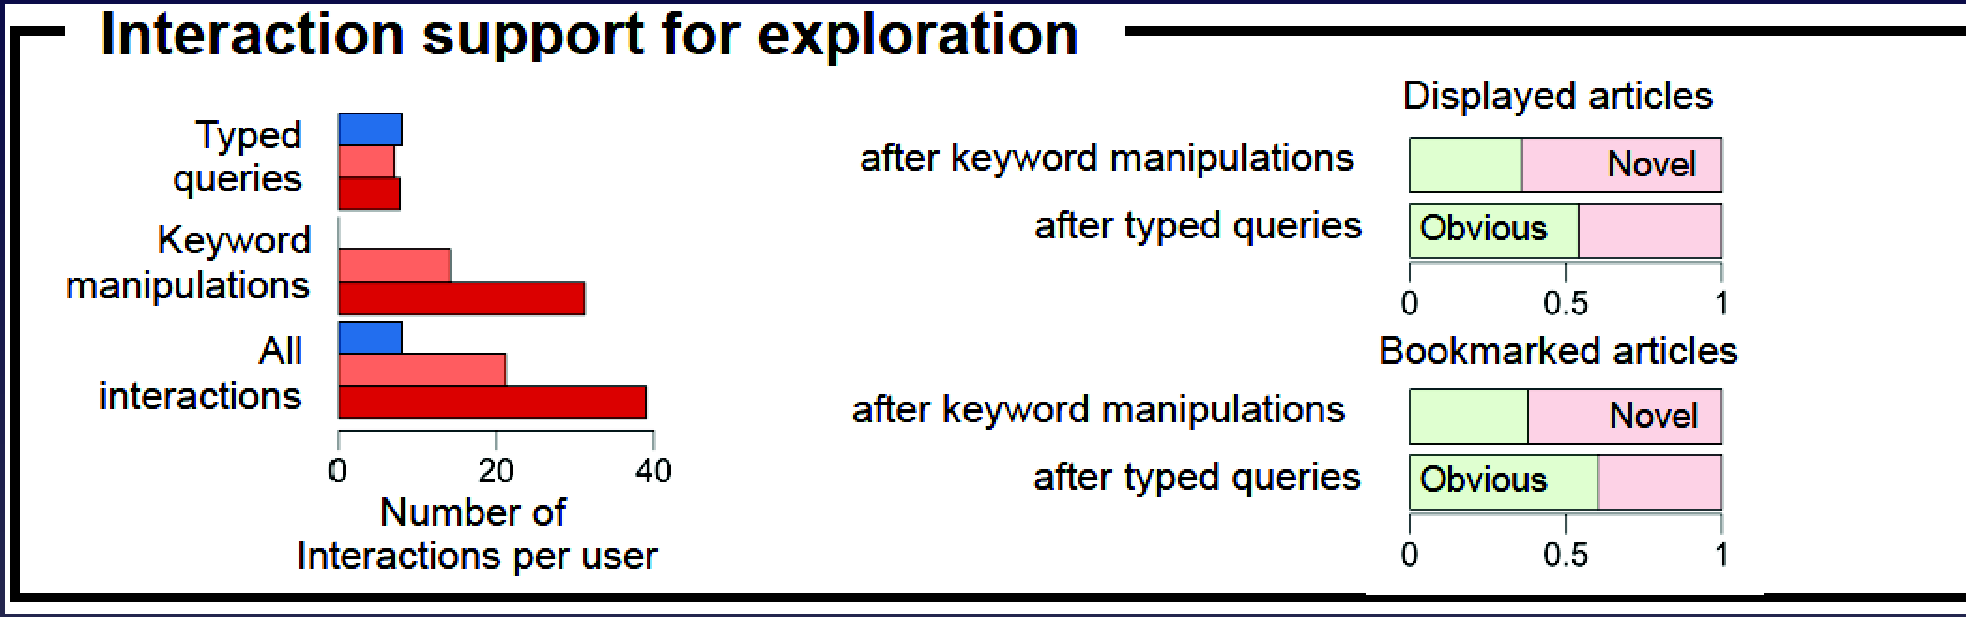

- Users interact with IntentRadar 2x as much as IntentList, nearly 4x as much as TypedQuery
- Intent models do not replace typed queries, they are used to direct search from initial imprecise query
- TypedQuery users had trouble reaching novel information
- Directing search with interactive intent modeling was successful: users got significantly more novel documents after interaction than after typed queries (same for bookmarked documents)

Conclusions

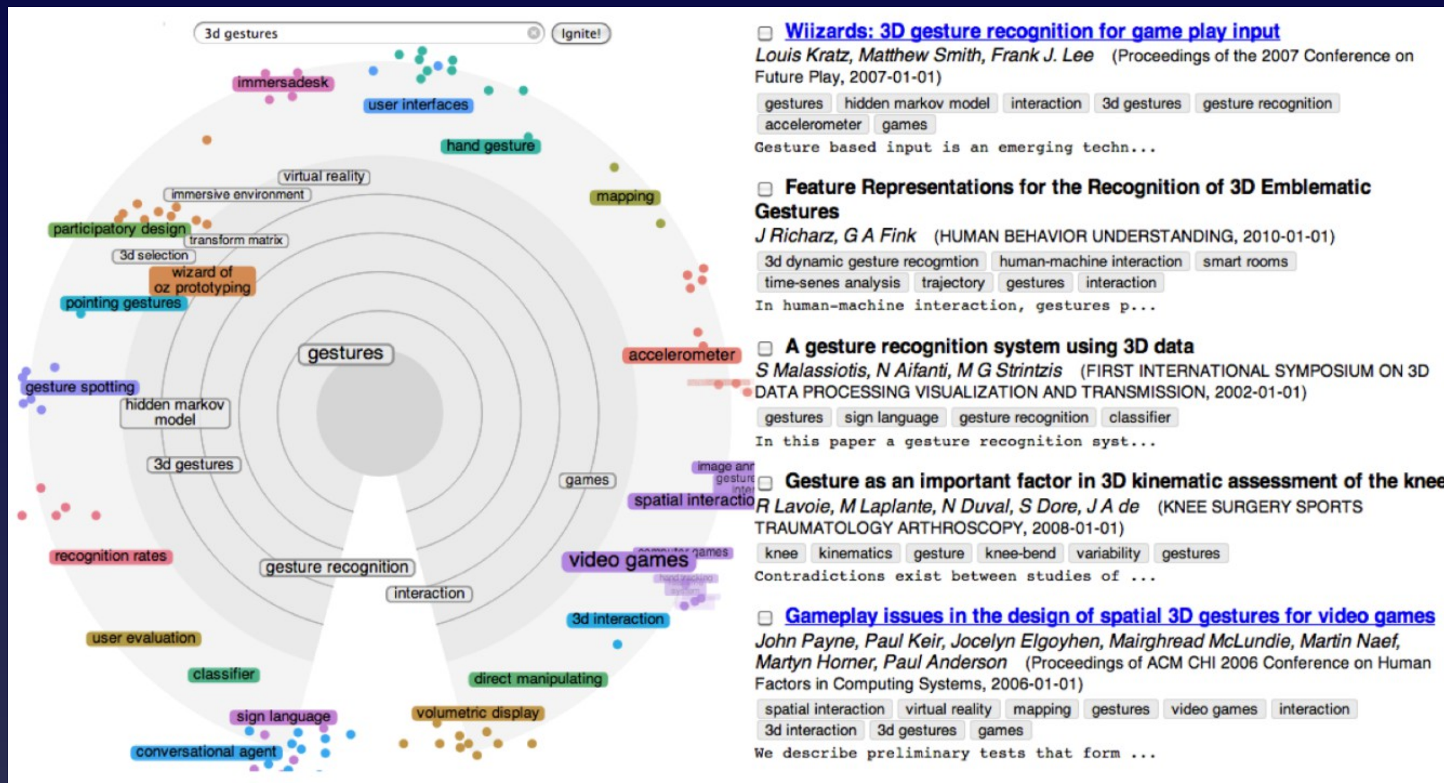

We introduced **interactive intent modeling for directing exploratory search**. It significantly improves users' performance in exploratory search tasks. Improvements can be attributed to better quality of displayed information in response to interactions, better targeted interaction, and better support for directing search to achieve novel information.

Conclusions

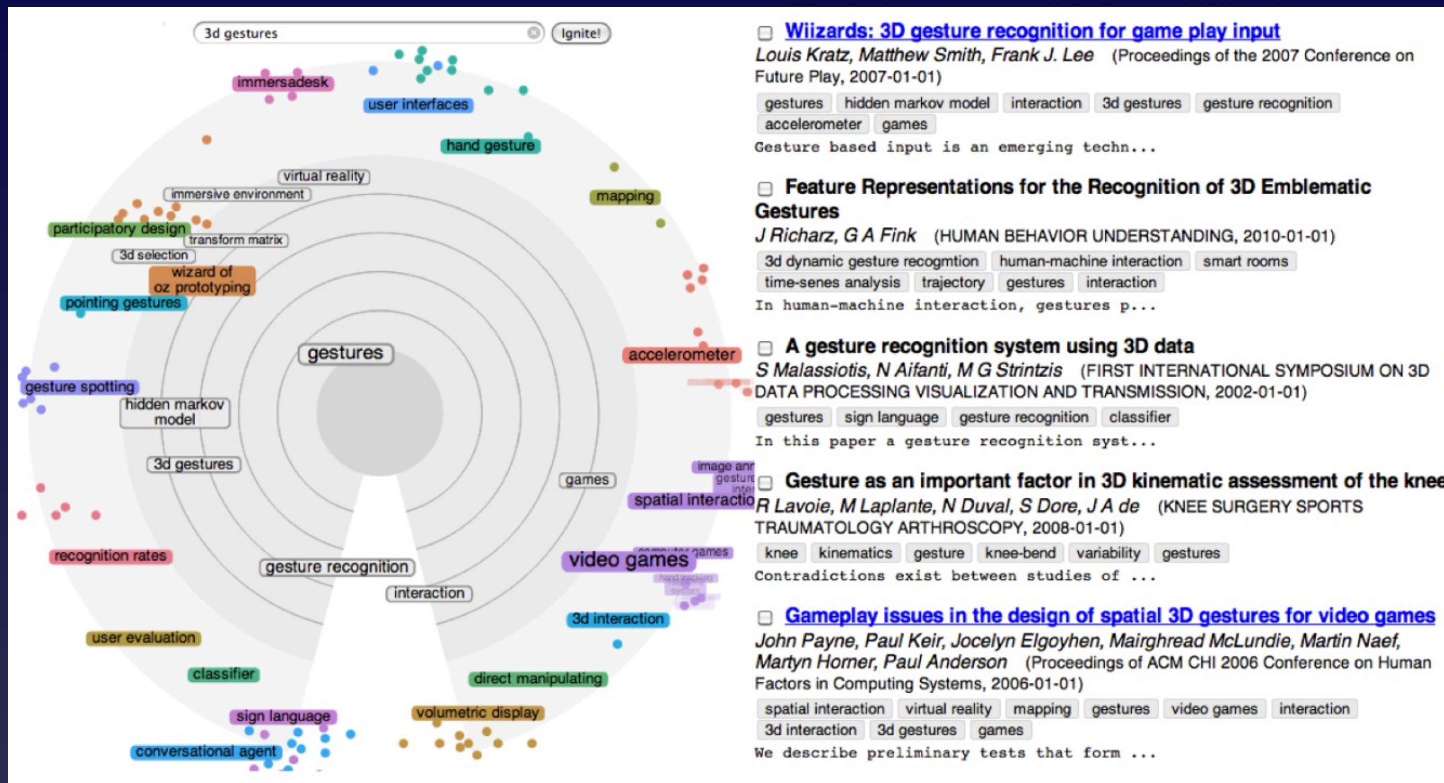

Reference: Tuukka Ruotsalo*, Jaakko Peltonen*, Manuel J. A. Eugster, Dorota Głowacka, Ksenia Konyushkova, Kumaripaba Athukorala, Ilkka Kosunen, Aki Reijonen, Petri Myllymäki, Giulio Jacucci, Samuel Kaski. **Directing Exploratory Search with Interactive Intent Modeling**. In Proceedings of CIKM 2013, ACM Conference on Computational Intelligence and Knowledge Management, 2013. (* equal contributions)

Can SciNet help talk recommendation?

Can SciNet help talk recommendation?

In Brief:

- We compare methods of **cross-system user model transfer**

In Brief:

- We compare methods of **cross-system user model transfer**

- Two **large real-life systems:**

SciNet for scientific document search

CoMeT for managing scientific talks

In Brief:

- We compare methods of **cross-system user model transfer**

- Two **large real-life systems:**

SciNet for scientific document search

CoMeT for managing scientific talks

- transfer of novel **explicit open user models** (curated by user during information seeking) strongly improves **cold-start talk recommendation**

Recommender systems face a **cold-start problem:** recommendations are needed for users who have rated few or no items

Recommender systems face a **cold-start problem:** recommendations are needed for users who have rated few or no items

We investigate **user model transfer** to enable warm start: establish in **source** system, use in **target** system

Cross-system/domain recommendation has grown in popularity, but still few studies exploring real information transfer (lack of paired users across systems).

Major focus has been on approaches not assuming common users. Major approaches: collaborative filtering or content-based.

Results mixed, especially content-based has been hard. Focus has been on settings having shared semantic features (social tags, Wikipedia).

We expand earlier research by exploring transferability of **open user models** across related but different domains.

Users of the source system can **explore** and **curate** their model by **visual interaction**.

→ better quality user models, valuable for cross-system transfer

1st work exploring transferability of open user models.

Contributions:

1) cross-system transfer of open user models greatly improves cold-start recommendation

2) we investigate **ways of transferring** open user models, as well as transfer of more traditional implicit and explicit document information.

Open user models bring greatest benefit. We explain it by analysis of cross-system similarities of the different information types.

Academic Information Setting

- Academic users attend **research talks**.
- **A talk management system** can recommend interesting talks given the user's preference.
- Relatively many talks but few bookmarks and ratings (Farzan et al., 2008)
- New users face the cold start problem
- Academic users also **search for scientific documents** in a scientific search system. Can its user model help talk recommendation?

Target system: **CoMeT** system for talk management and recommendation

Target system: **CoMeT** system for talk management and recommendation

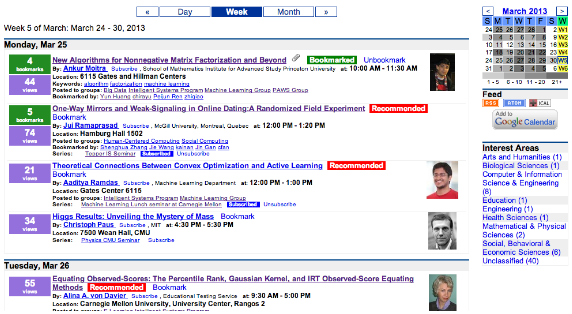

System for sharing information about research talks at Carnegie Mellon University and University of Pittsburgh.

- Available online: halley.exp.sis.pitt.edu/comet/
- Collaborative tagging system: anyone can announce, find, bookmark, and tag talks.
- Has **content-based recommender** - builds interest profile of individual users, recommends new talks to users immediately when posted.

Academic Information Setting

- We use CoMeT as the **target system**
- Academic users also **search for scientific documents** in a scientific search system. Can its user model help talk recommendation?

- Unlike traditional search systems (e.g. Google Scholar, Microsoft Academic Search, Citeseer), as the **source system** we use a recent search system having an **open user model: SciNet**

Source system: **SciNet** system for interactive exploratory search of scientific documents

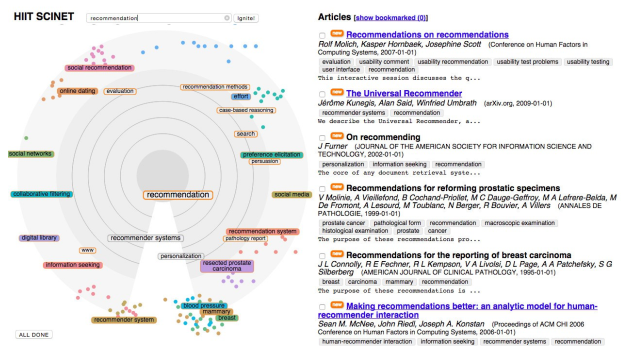

Exploratory search system. Indexes over 50m scientific documents from Thomson Reuters, ACM, IEEE & Springer

- Goes beyond text-based queries.
- SciNet **opens its user model:** users can interact with a **visualization** of the model, and curate the model by feedback.
- User can direct exploratory search by interacting with the open user model.
- Significantly improves information seeking task performance and quality of retrieved information.
- Open user models are promising for cross-system transfer.

Ways of Transferring a User Model

Our interest is to use

1. the **whole content of the open user model**

2. its **curated subset** (the keywords the user moved in the process of curation).

As a baseline, we also explore transfer of:

3. the **relevant documents** selected by the user during search (could be considered a hidden, implicit user model)

4. a broader set of all **documents retrieved** in response to user queries (weaker reflection of user interests)

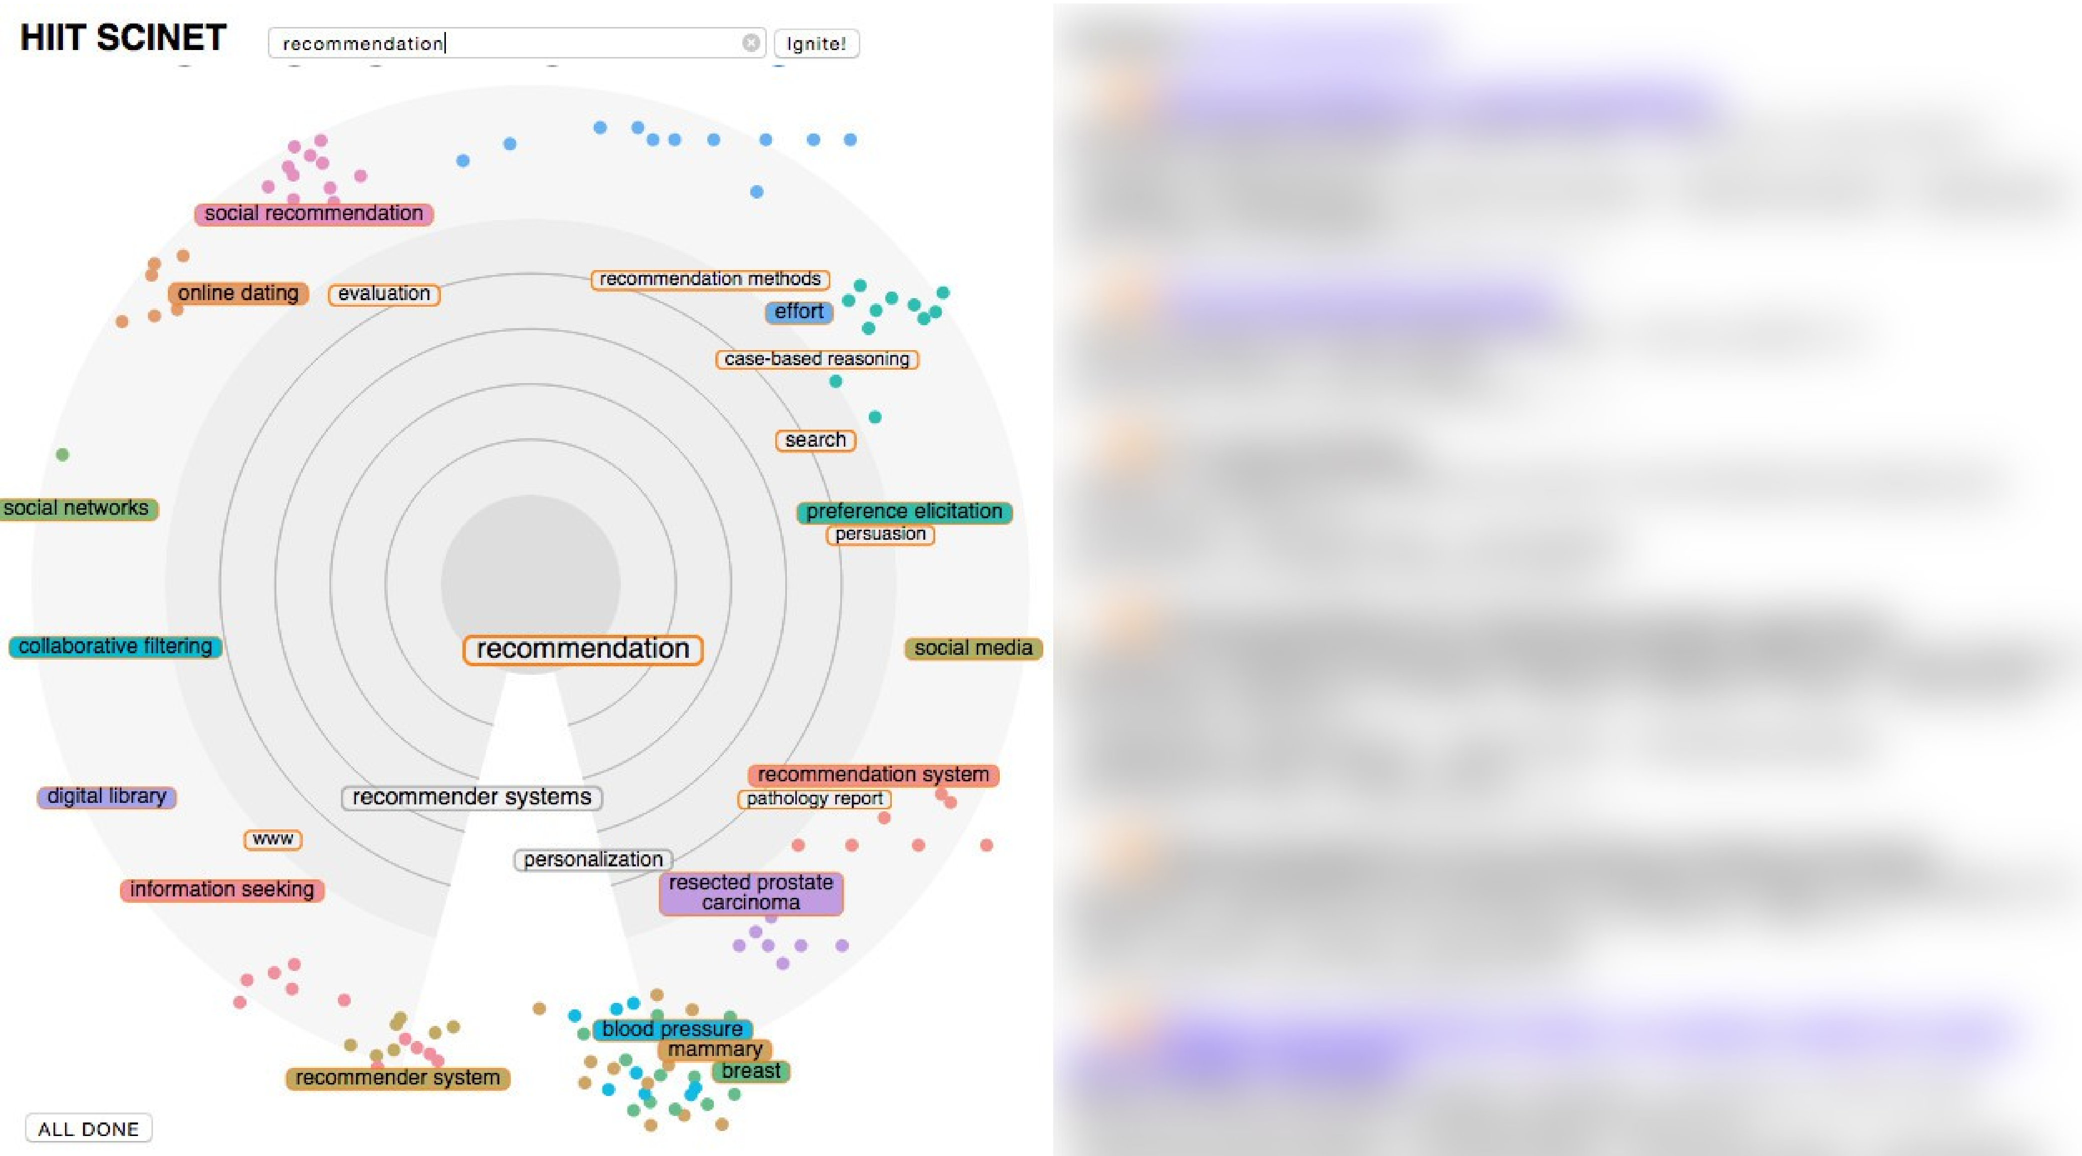
HIIT SCINET
recommendation
Ignite!
social recommendation
online dating
evaluation
recommendation methods
effort
case-based reasoning
search
social networks
preference elicitation
persuasion
collaborative filtering
recommendation
social media
recommendation system
pathology report
digital library
recommender systems
www
personalization
information seeking
resected prostate carcinoma
blood pressure
mammary
breast
recommender system
ALL DONE

Explicit model 1: manipulated keywords

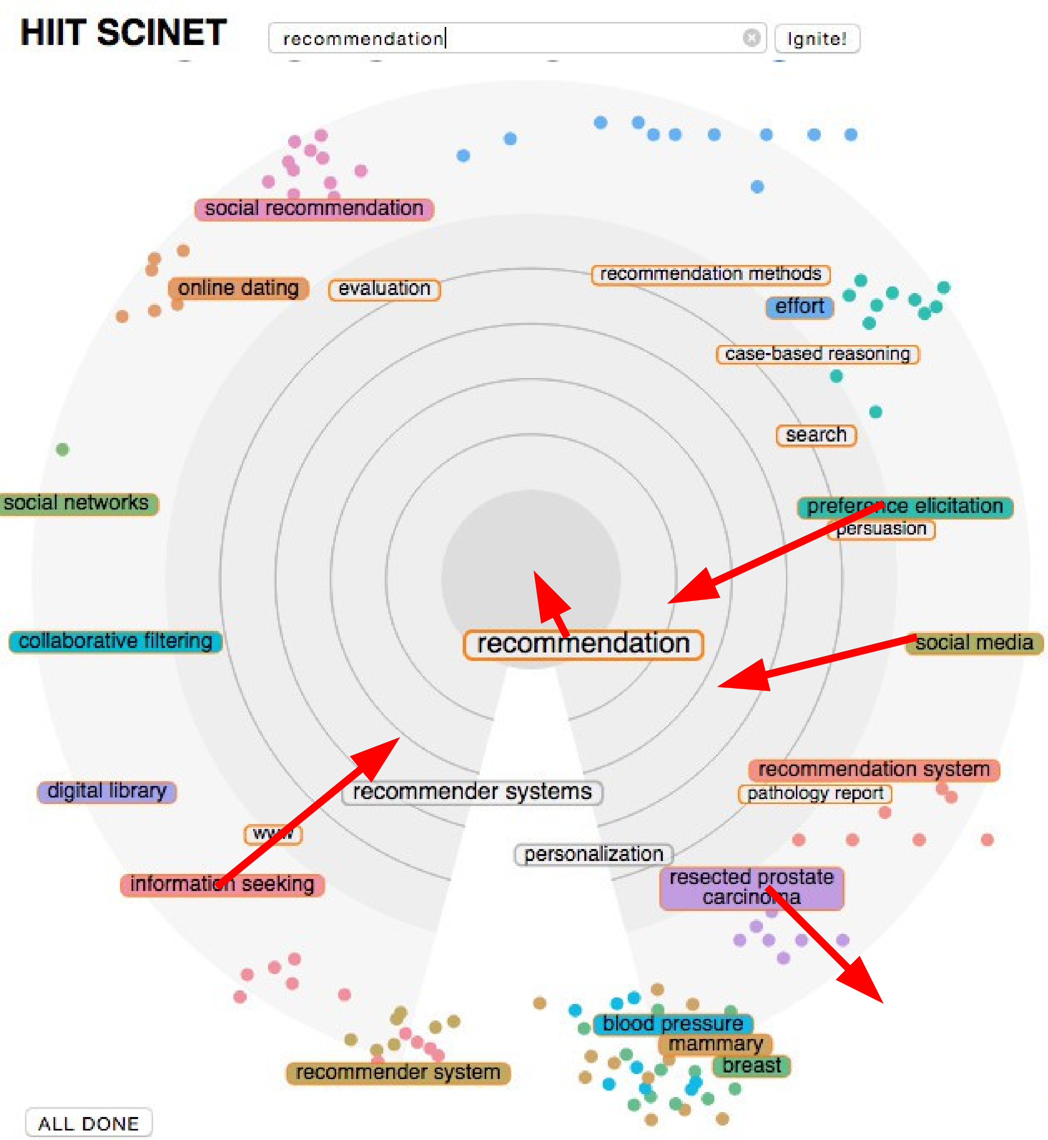

Take keywords **manipulated** by users, with their **assigned** relevances, convert to unigrams, and form a pseudo-document (bookmarked talk abstract) from them.

(discard unigrams not occurring in target system)

Explicit model 2: shown keywords

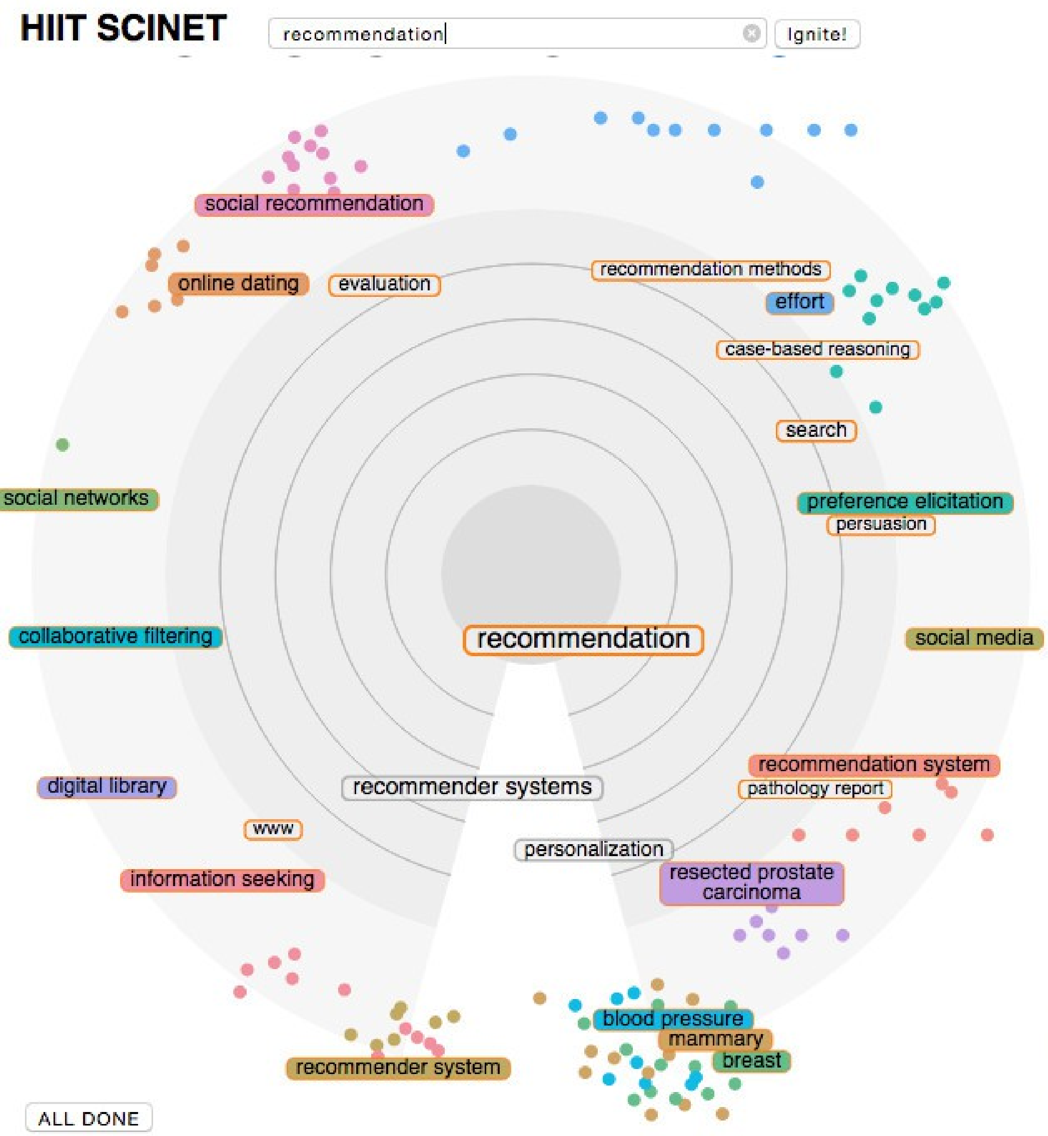

At each iteration,
Take keywords
seen by users,
with their **predicted**
relevances, convert to
unigrams, and form a
pseudo-document
(bookmarked talk abstract)
from them.

(discard unigrams not occurring in target system)

Articles [show bookmarked (0)]

new **Recommendations on recommendations**
Rolf Molich, Kasper Hornbaek, Josephine Scott (Conference on Human Factors in Computing Systems, 2007-01-01)
evaluation | usability comment | usability recommendation | usability test problems | usability testing | user interface | recommendation
This interactive session discusses the q...

new **The Universal Recommender**
Jérôme Kunegis, Alan Said, Winfried Umbrath (arXiv.org, 2009-01-01)
recommender systems | recommendation
We describe the Universal Recommender, a...

new **On recommending**
J Furner (JOURNAL OF THE AMERICAN SOCIETY FOR INFORMATION SCIENCE AND TECHNOLOGY, 2002-01-01)
personalization | information seeking | recommendation
The core of any document retrieval syste...

new **Recommendations for reforming prostatic specimens**
V Molinie, A Vieillefond, B Cochand-Priollet, M C Dauge-Geffroy, M A Lefrere-Belda, M De Fromont, A Lesourd, M Toublanc, N Berger, R Bouvier, A Villers (ANNALES DE PATHOLOGIE, 1999-01-01)
prostate cancer | pathological form | recommendation | macroscopic examination | histological examination | prostate | cancer
The purpose of these recommendations pro...

new **Recommendations for the reporting of breast carcinoma**
J L Connolly, R E Fechner, R L Kempson, V A Livolsi, D L Page, A A Patchefsky, S G Silberberg (AMERICAN JOURNAL OF CLINICAL PATHOLOGY, 1995-01-01)
breast | carcinoma | mammary | recommendation
The purpose of these recommendations is ...

new **Making recommendations better: an analytic model for human-recommender interaction**
Sean M. McNee, John Riedl, Joseph A. Konstan (Proceedings of ACM CHI 2006 Conference on Human Factors in Computing Systems, 2006-01-01)
human-recommender interaction | information seeking | recommender systems | recommendation

Implicit model 1: bookmarked documents

Scientific documents **bookmarked** by the user during the search session are implicit information about user interests.

Convert each into unigrams, add into CoMeT as a bookmarked talk.

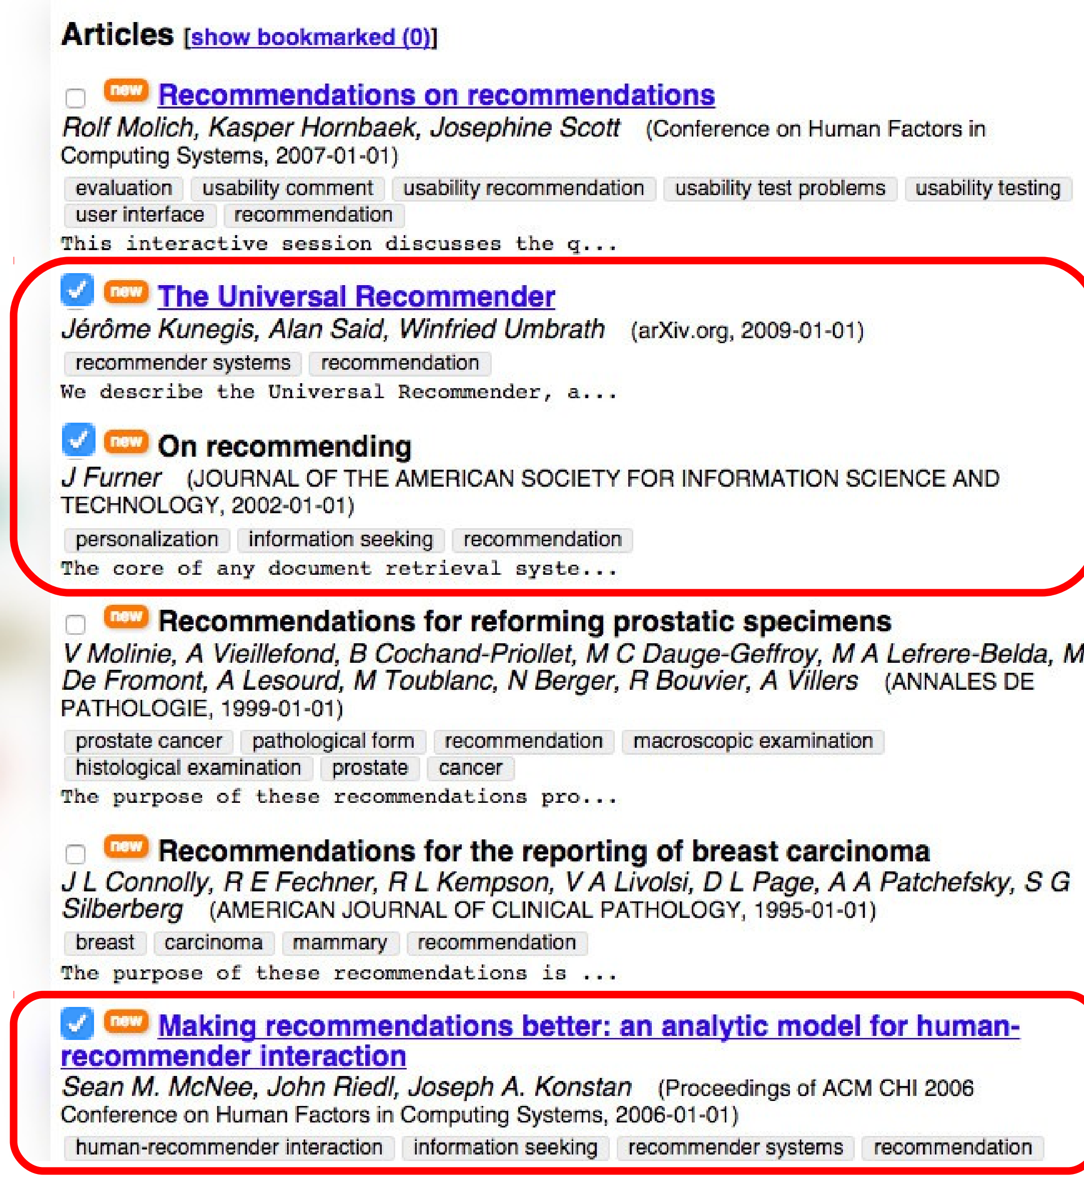

Implicit model 2: seen documents

Scientific documents **seen** by the user during the search session are implicit information about user interests (momentary responses to user search).

Convert each into unigrams, add into CoMeT as a bookmarked talk.

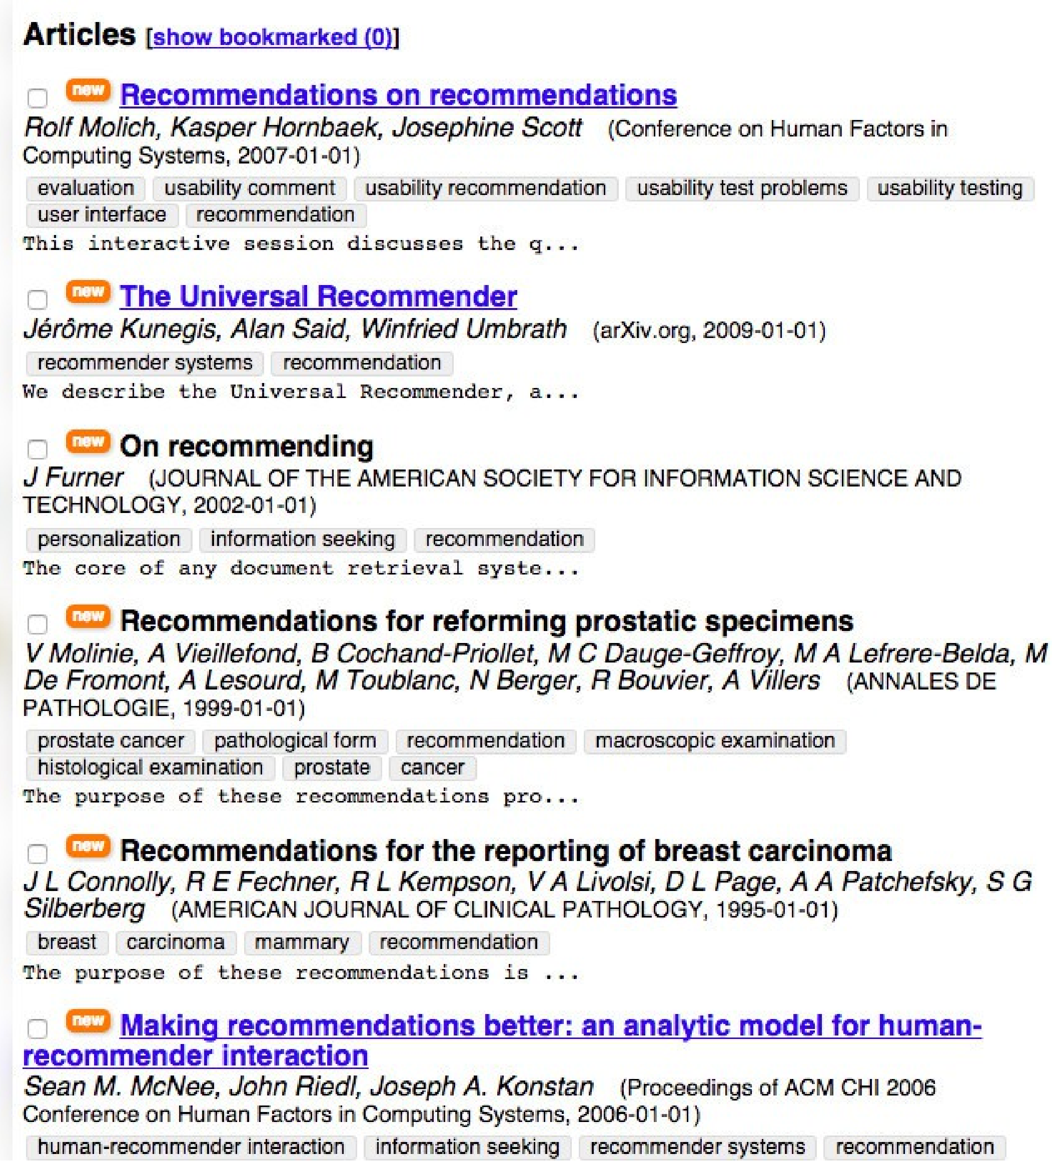

Articles [show bookmarked (0)]

new **Recommendations on recommendations**
Rolf Molich, Kasper Hornbaek, Josephine Scott (Conference on Human Factors in Computing Systems, 2007-01-01)
evaluation | usability comment | usability recommendation | usability test problems | usability testing | user interface | recommendation
This interactive session discusses the q...

new **The Universal Recommender**
Jérôme Kunegis, Alan Said, Winfried Umbrath (arXiv.org, 2009-01-01)
recommender systems | recommendation
We describe the Universal Recommender, a...

new **On recommending**
J Furner (JOURNAL OF THE AMERICAN SOCIETY FOR INFORMATION SCIENCE AND TECHNOLOGY, 2002-01-01)
personalization | information seeking | recommendation
The core of any document retrieval syste...

new **Recommendations for reforming prostatic specimens**
V Molinie, A Vieillefond, B Cochand-Priollet, M C Dauge-Geffroy, M A Lefrere-Belda, M De Fromont, A Lesourd, M Toublanc, N Berger, R Bouvier, A Villers (ANNALES DE PATHOLOGIE, 1999-01-01)
prostate cancer | pathological form | recommendation | macroscopic examination | histological examination | prostate | cancer
The purpose of these recommendations pro...

new **Recommendations for the reporting of breast carcinoma**
J L Connolly, R E Fechner, R L Kempson, V A Livolsi, D L Page, A A Patchefsky, S G Silberberg (AMERICAN JOURNAL OF CLINICAL PATHOLOGY, 1995-01-01)
breast | carcinoma | mammary | recommendation
The purpose of these recommendations is ...

new **Making recommendations better: an analytic model for human-recommender interaction**
Sean M. McNee, John Riedl, Joseph A. Konstan (Proceedings of ACM CHI 2006 Conference on Human Factors in Computing Systems, 2006-01-01)
human-recommender interaction | information seeking | recommender systems | recommendation

Experiments

- 20 researchers from University of Helsinki:
  14 male, 6 female; 10 PhD researchers and 10 research assistants
- SciNet: Search relevant papers to their interest
  - "Write down 3 areas of your research interests. Imagine you are preparing for a course/seminar for each interest. Search scientific documents you find useful for preparing for the courses/seminars."
  - Bookmark at least 5 documents for each interest.
  - 7min demonstration, 30min for task
  - Complex enough: users must interact with the system to get needed information. Broad enough to reveal research interests.
- CoMeT: Rate 500 talks (Jan 1 to May 17, 2013)
  - Consider attending (Yes/No)? If yes, rate willingness 1 – 5
  - 7min demonstration, 75min for task
- All interactions logged (shown/manipulated keywords shown/bookmarked documents, queries, read abstracts...)

Non-cold-start Setting

- We first evaluated a traditional non-cold-start learning setting
- 10-fold cross-validation setup, in each fold rank the held-out CoMeT talks by 3 predictors
- **Centroid:** rank test talks by cosine similarity to centroid of bookmarked talks
- **k-Nearest-Neighbor:** find nearest training neighbors for each test talk, rank by $s_{pos}-s_{neg}$ (sum of cosine similarities to positive nearest neighbors - sum of cosine similarities to negative neighbors)
- **positive-only kNN:** find nearest positive-rated talks, rank by sum of cosine similarity to them

Non-cold-start Setting

- Results evaluated by Mean Average Precision of ranked test talks
  (mean of precision values at positive test talks in the ranking, averaged over users and folds)
- no significant improvement from transfer compared to baseline in non-cold-start setting
  from traditional or open-user-model approach
- User profiles in CoMeT had enough data to work well on their own

Mean Average Precision		Centroid	k-NN				k-NN.PO			
			5nn	10nn	20nn	30nn	5nn.po	10nn.po	20nn.po	30nn.po
baseline		0.47	0.45	0.47	0.48	0.46	0.48	0.50	0.50	0.50
Implicit User Model	ex.papers	0.42	0.44	0.45	0.45	0.44	0.43	0.44	0.44	0.44
	im.papers	0.36	0.36	0.36	0.35	0.35	0.36	0.37	0.36	0.36
Explicit Open User Model	ex.keywords	0.48	0.46	0.48	0.48	0.47	0.49	0.51	0.51	0.50
	im.keywords	0.47	0.46	0.48	0.49	0.49	0.48	0.49	0.49	0.48

Cold-start Setting

- In each cross-validation fold we subsample a small pool of cold-start talks (0-20 positive talks, proportionally same amont of negative talks)
- Cold-start talks used to predict test talk ranking, evaluate by mean average precision
- We report average results over 10 subsamplings
- Same predictors as before (Centroid, k-Nearest-Neighbor, positive-only kNN)

Cold-Start Impact

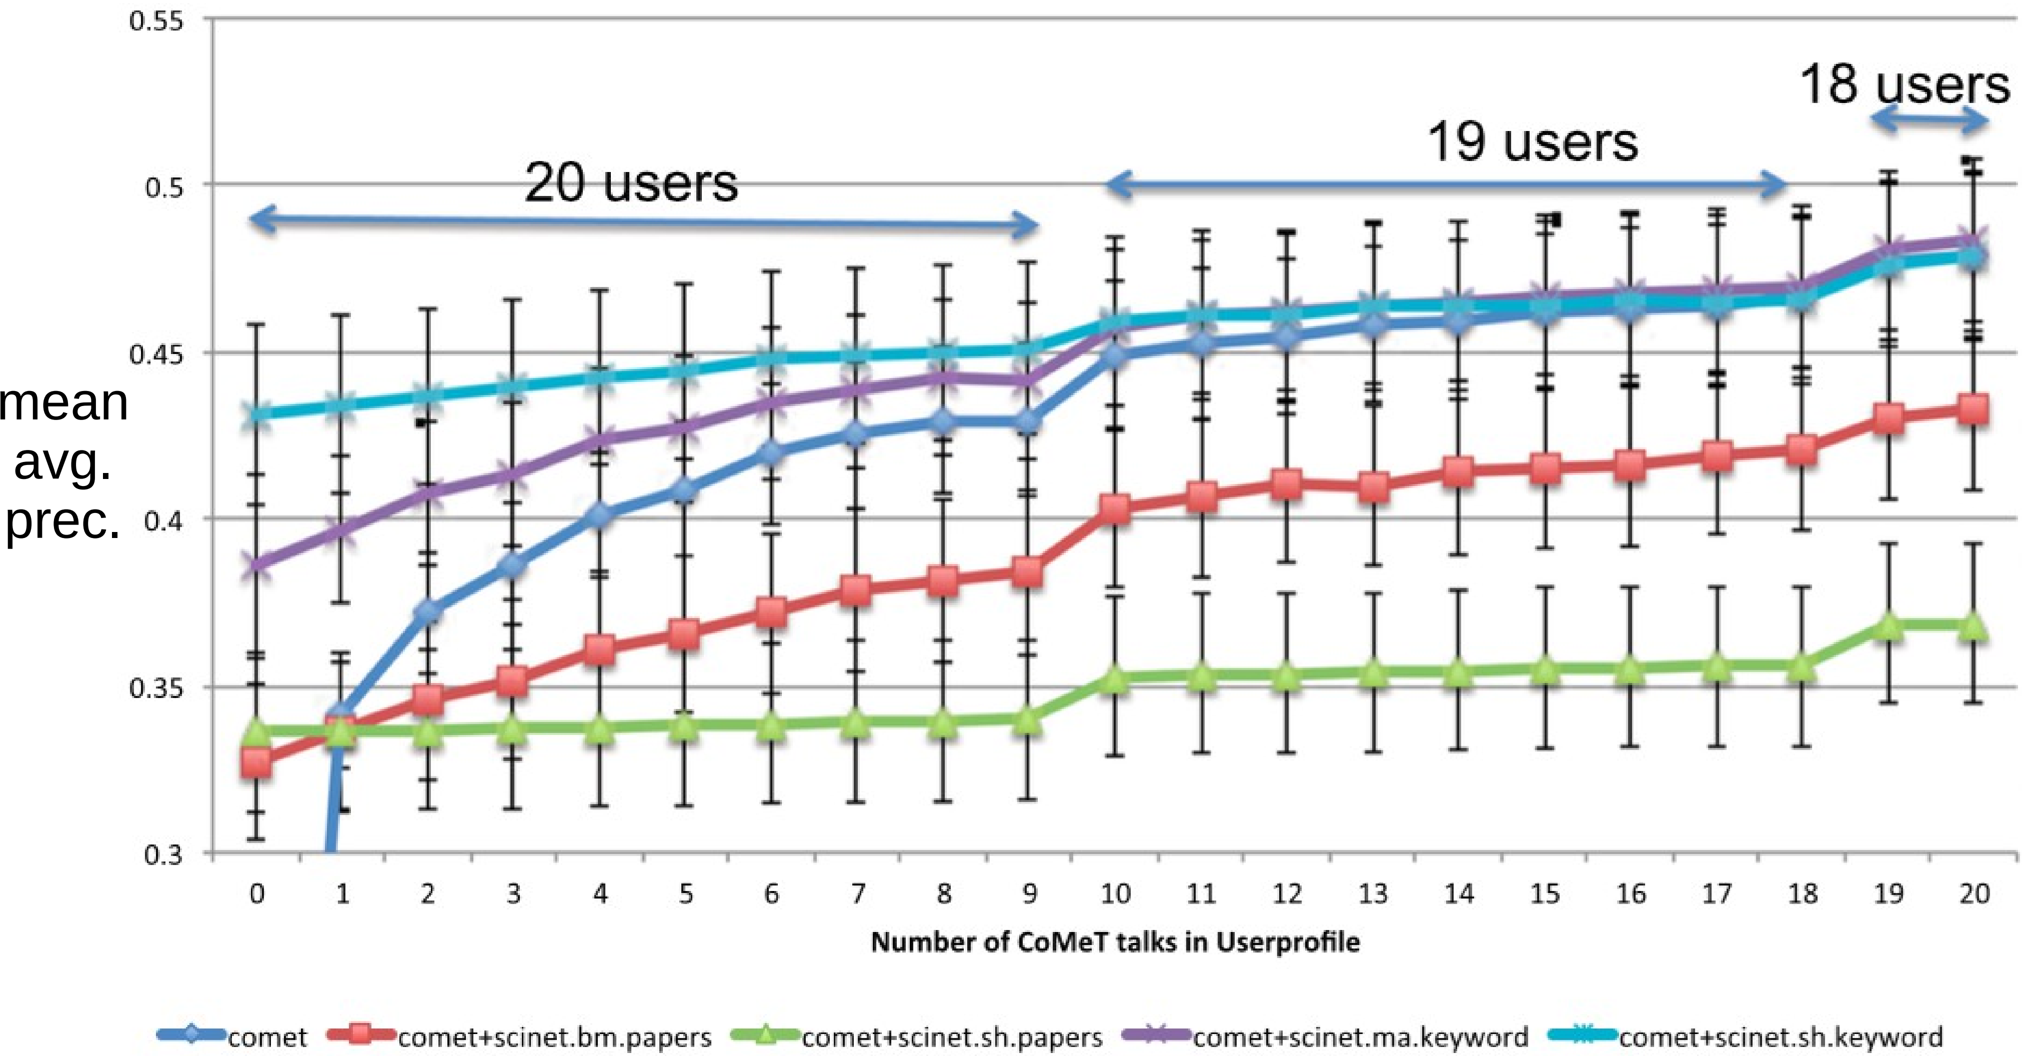

Results with centroid predictor shown: positive-only kNN performs essentially the same, and outperforms kNN

Cold-Start Impact

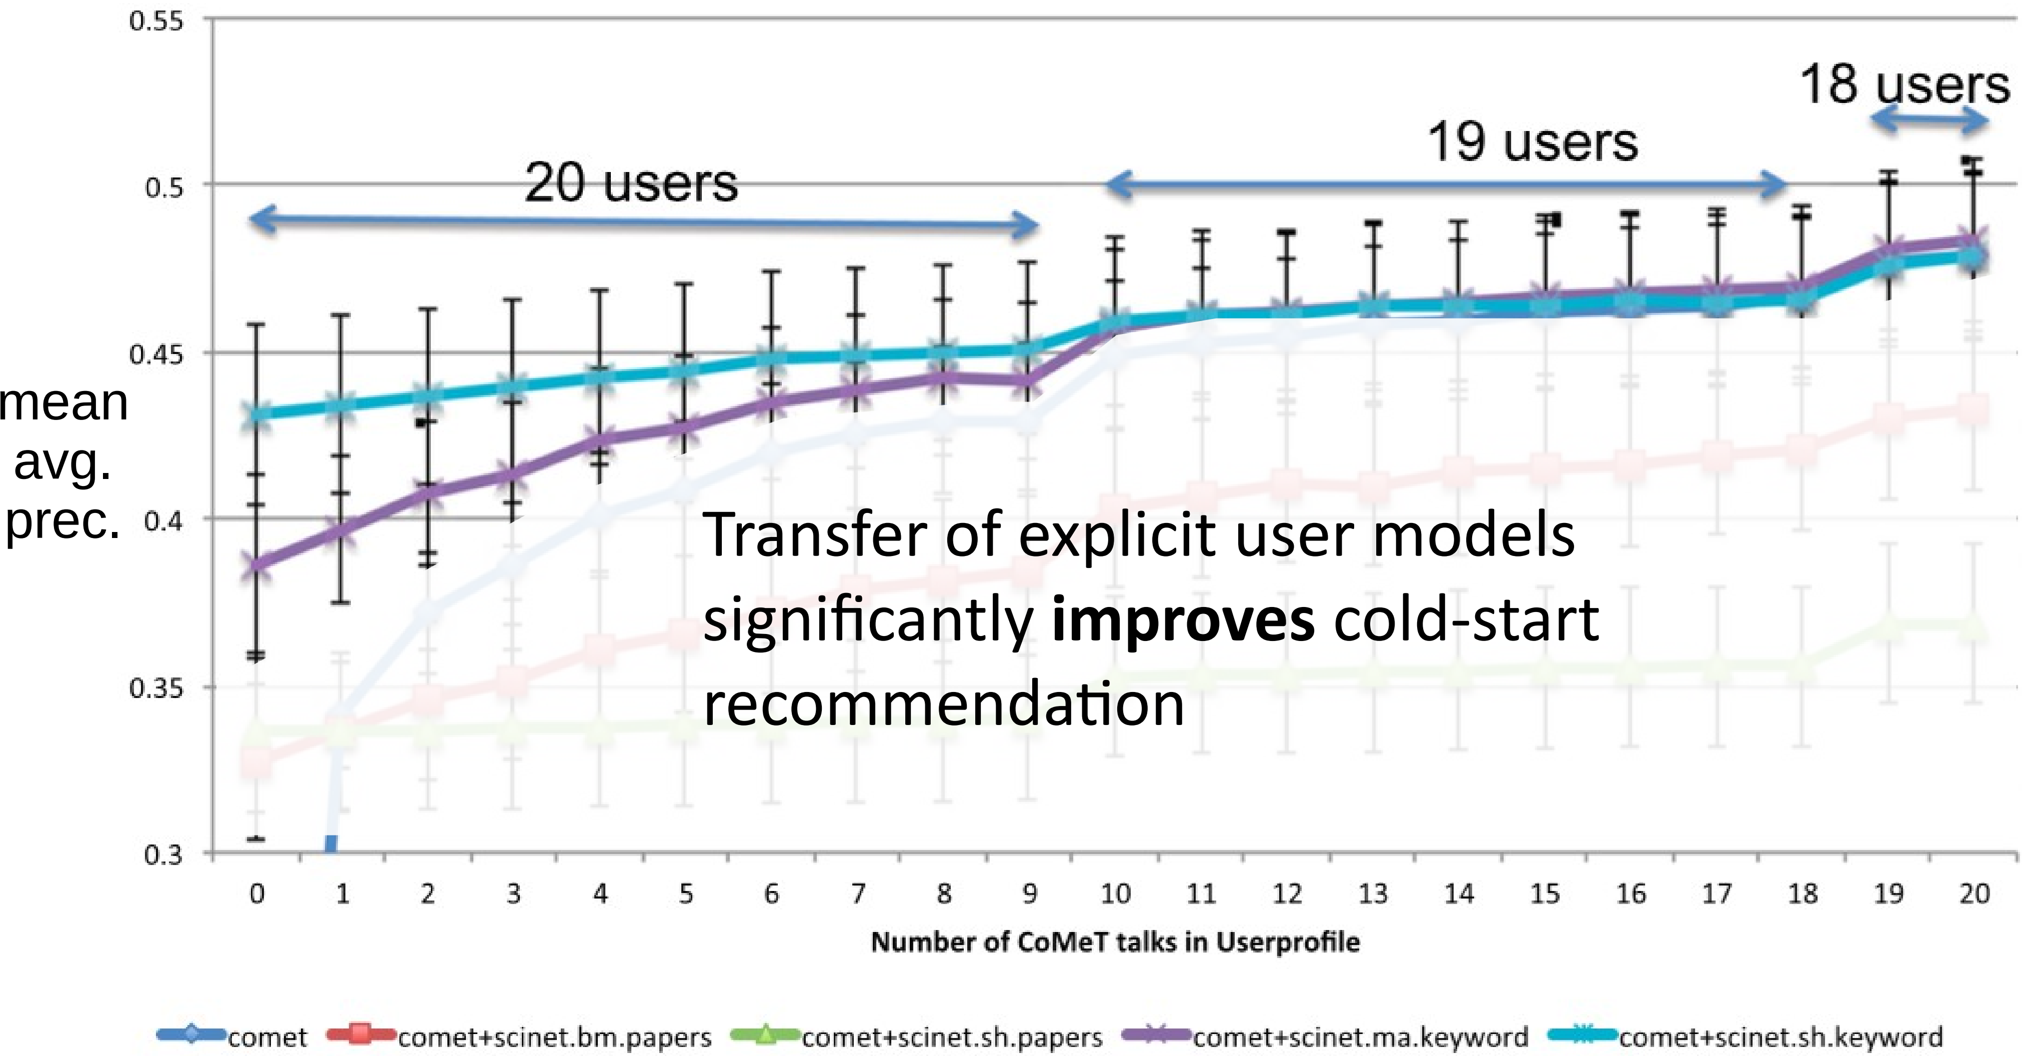

Transfer of explicit user models significantly **improves** cold-start recommendation

Results with centroid predictor shown: positive-only kNN performs essentially the same, and outperforms kNN

Cold-Start Impact

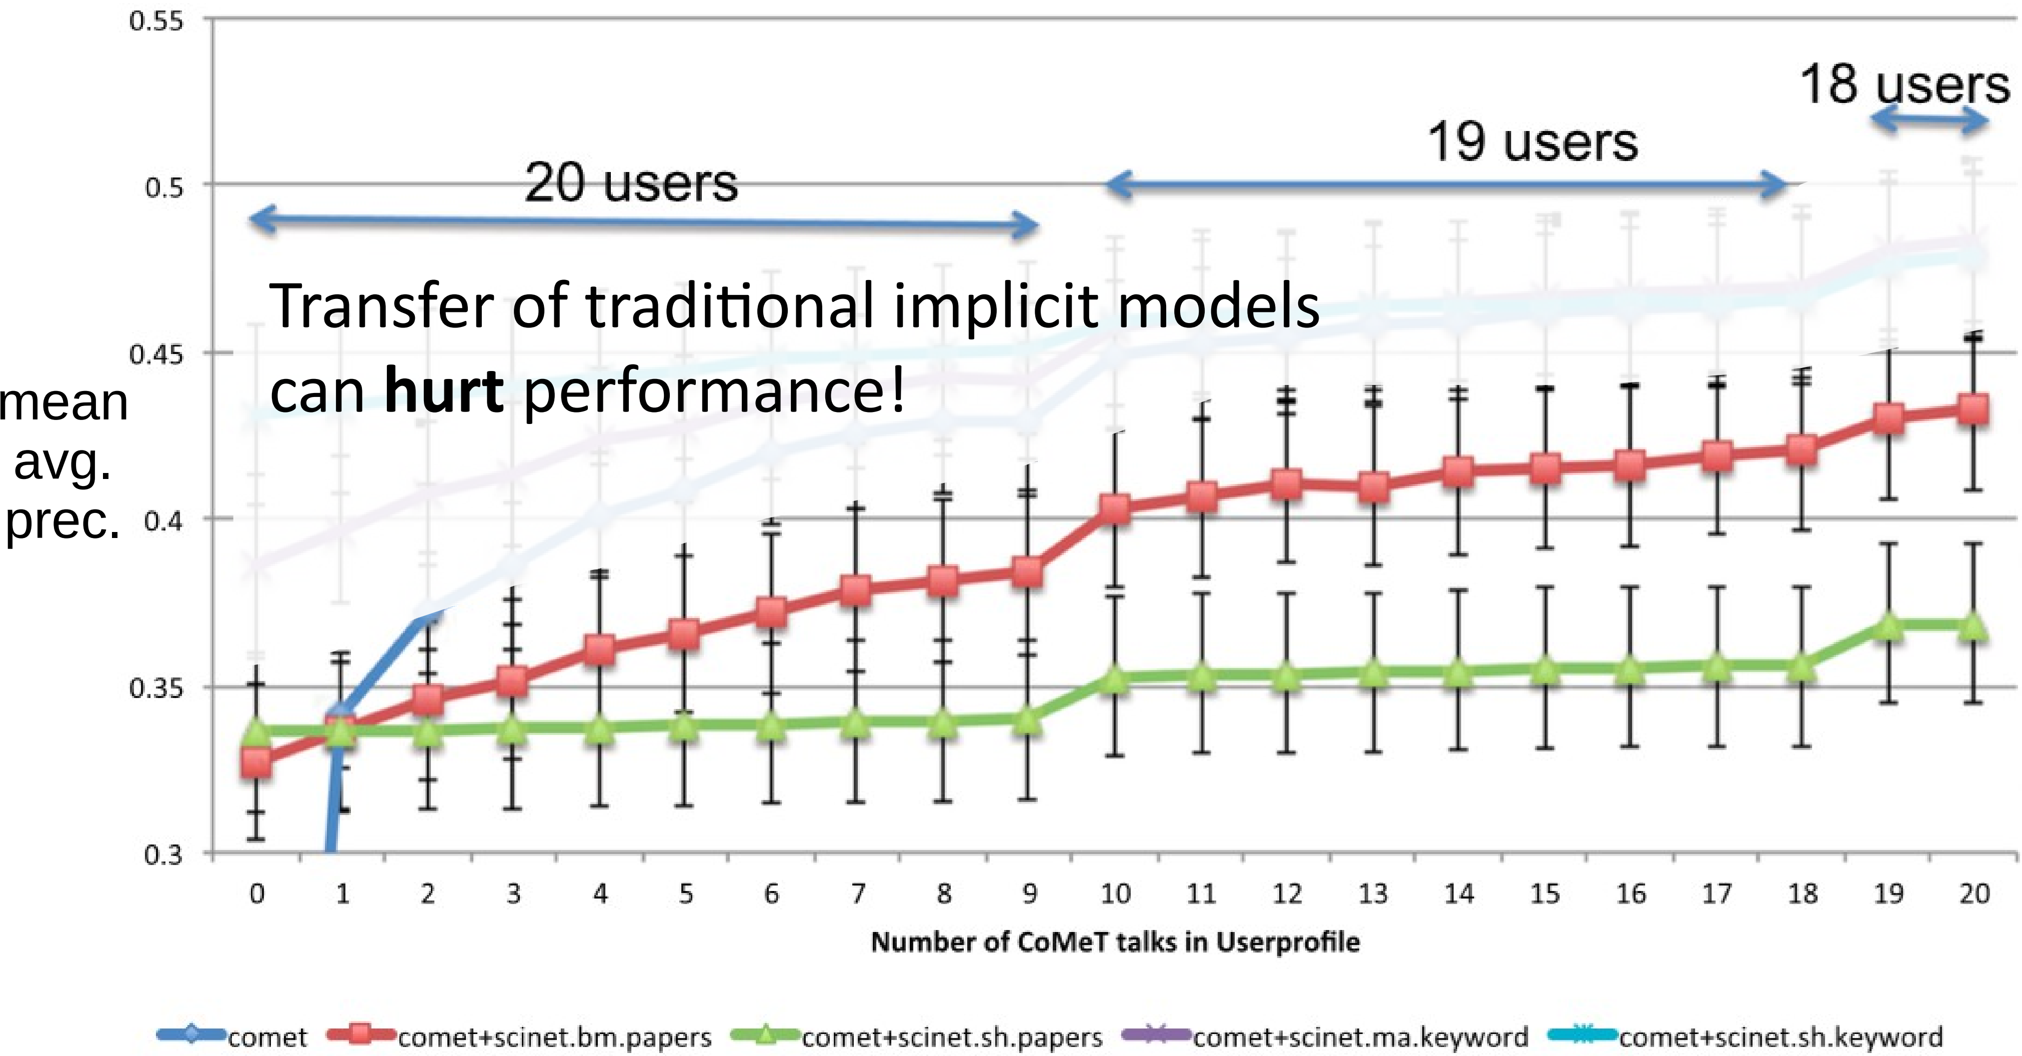

Results with centroid predictor shown: positive-only kNN performs essentially the same, and outperforms kNN

Analysis

- Cosine similarities between different information types

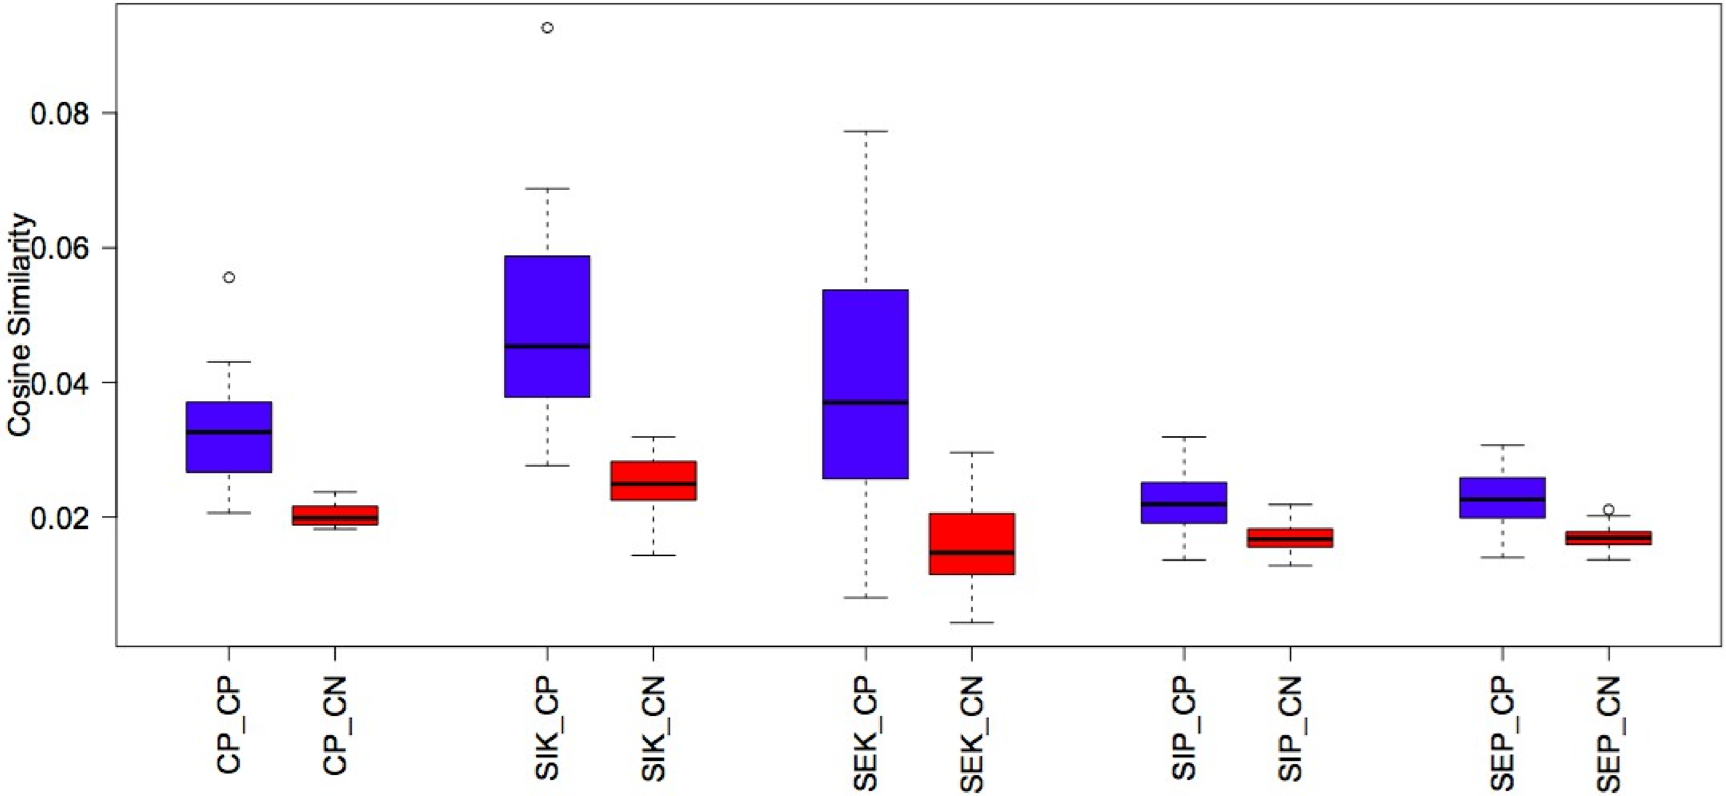

Explicit open user models have good similarity to positive-rated talks, well separated from uninteresting talks

Implicit models from papers are far from bookmarked talks. They do not separate positive-rated from uninteresting talks.
→ add more noise than value

Summary

- Cross-system personalization by transferring an **explicit, open, and editable user model.**
- Transfer from a literature search system to a talk recommendation system.
- Cross-system model transfer is challenging: no impact in general case
- However, **significant impact in cold-start case!**
- **Use of open, explicitly curated user models is critical** for the success of user model transfer
- Transferring implicit models (here through shown or bookmarked documents) can damage performance

Overall Conclusions

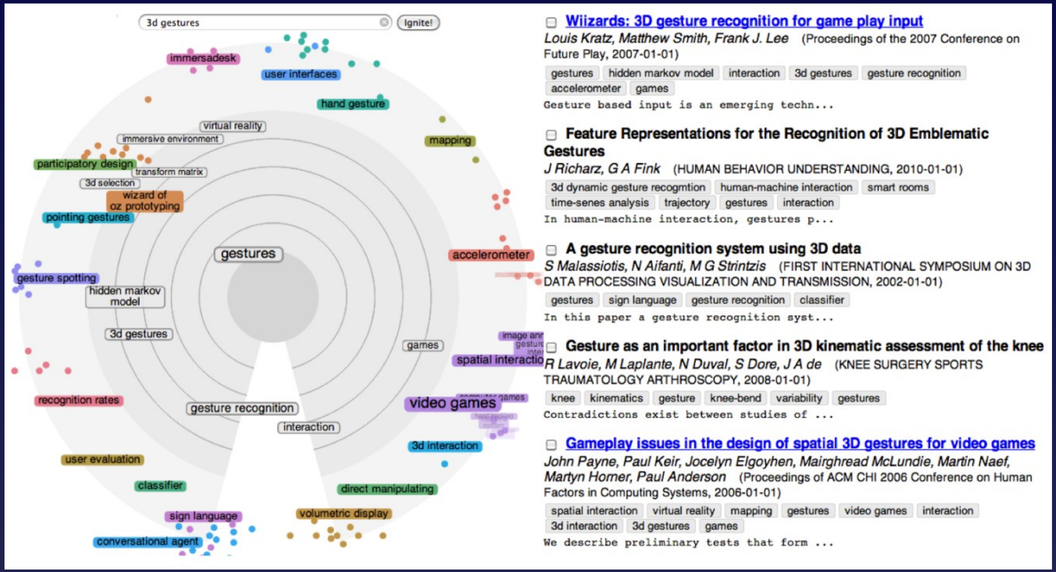

Novel systems to search for relevant data with open user models

Transfer of open user models helps in cold-start recommendation

References

Tuukka Ruotsalo*, Jaakko Peltonen*, Manuel Eugster*, Dorota Glowacka, Ksenia Konyushkova, Kumaripaba Athukorala, Ilkka Kosunen, Aki Reijonen, Petri Myllymäki, Giulio Jacucci, and Samuel Kaski. **Directing Exploratory Search with Interactive Intent Modeling.** In Proceedings of CIKM 2013, ACM Conference on Information and Knowledge Management, pages 1759-1764. ACM, 2013. (* equal contributions)

Chirayu Wongchockprasitti, Jaakko Peltonen, Tuukka Ruotsalo, Payel Bandyopadhyay, Giulio Jacucci and Peter Brusilovsky. **User Model In a Box: Cross-System User Model Transfer for Resolving Cold Start Problems.** In Proceedings of UMAP'15, The 23rd Conference on User Modelling, Adaptation and Personalization, pages 289-301, Springer, 2015.