MULTIMEDIA RECOMMENDER SYSTEMS: APPROACHES AND CHALLENGES

Prof. VINCENZO MOSCATO,
University of Naples Federico II,
Italy

RECOMMENDER SYSTEMS: FROM SEARCH TO DISCOVERY…

As early as 2006, in "The Long Tail", Chris Anderson states that we are leaving the age of *information* and entering the age of *recommendation*, and this statement seems to be a good summary of what is happening in the modern *Big Data and Artificial Intelligence* era…

- According to the "Magazine" writer Jeffrey M. O'Brien, in the last decade we have assisted the transition from the "season of search" to the "season of discovery"…
- Search is what you do when you are looking for something while Discovery is when something wonderful that you didn't know existed, or didn't know how to ask for, finds you.

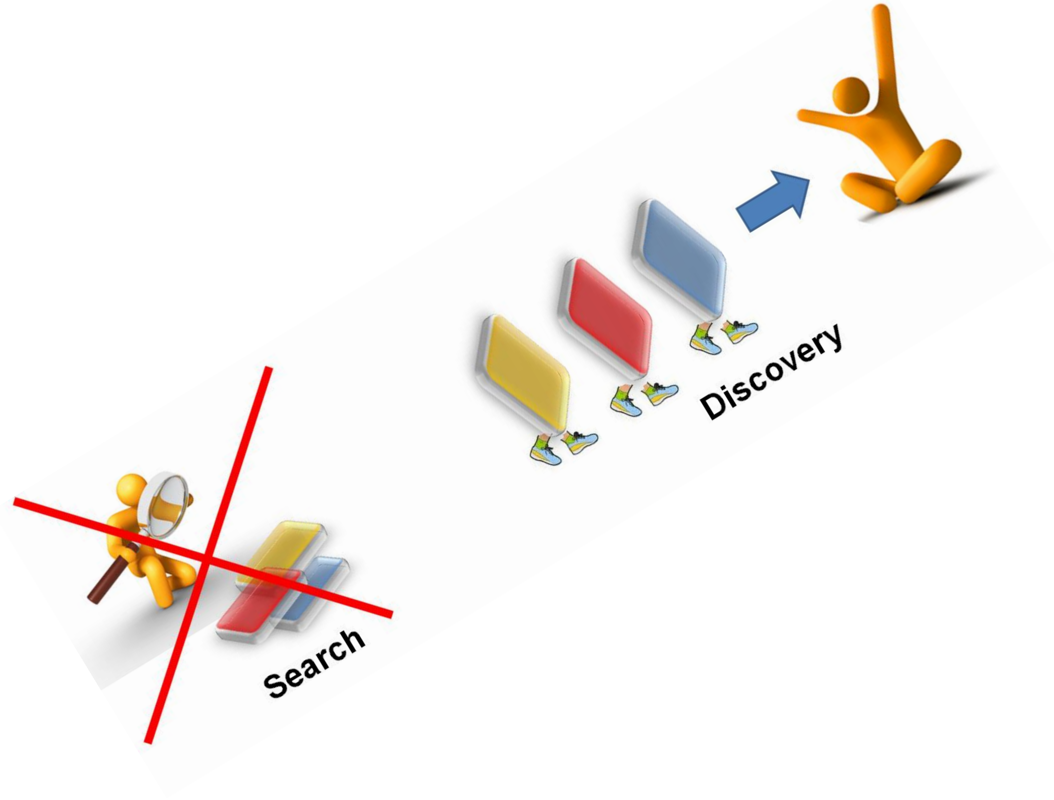

RECOMMENDER SYSTEMS: APPLICATIONS

Recommender systems help people in retrieving information that matches their preferences by recommending products or services from many candidates, and support people in making decisions in various contexts:

- what items to buy,
- which restaurant to book,
- which movie to watch,
- who they can invite to their social network,
- etc.

just to cite the first application contexts in which recommender have found their natural usage…

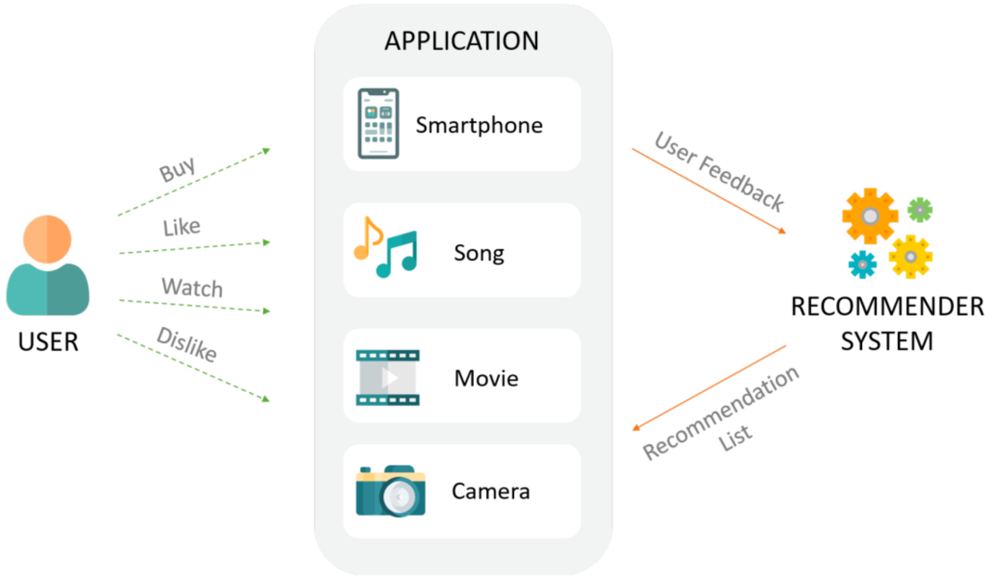

THE EXPLOSION OF MULTIMEDIA DATA…

Multimedia data is surely one of the most popular and pervasive information and communication media that accompanies us in almost every walks of life.

They allow fast and effective communication and sharing of information about peoples' lives, their behaviors, works, interests, and they are also the digital testimony of facts, objects, and locations and have become an essential component of Online Social Networks and Multimedia Streaming Platforms (e.g., Netflix, Spotify, etc.) …

From the other hand, the diffusion of such platforms has increasingly brought out the need for *multimedia recommendation algorithms*…

- They are known to suggest multimedia content of interest to users based on their preferences and needs, but also on the basis the behavior of similar users or user communities…

THE BORN OF MULTIMEDIA RECOMMENDER SYSTEMS

One of the most challenging research topics concerning multimedia data is surely to provide users with recommendation facilities that are able to suggest content of interest within very large collections of data.

To this aim, in the last decade, traditional recommendation strategies have been extended to handle multimedia data and the related features giving rise to the so-called *Multimedia Recommender Systems*.

One of the first work that introduces the term "Multimedia Recommender System" is that proposed by Moscato et al. in ACM Transaction on Internet Technologies (*A Multimedia Recommender System*, 2013), where multimedia features and similarity are for the first time fully considered and integrated in the recommendation process for the case related to the browsing of a virtual museum…

WHAT IS A RECOMMENDER SYSTEM?

Formally, a recommender system deals with a set of users ***U= {u1, u2,...ui,...um}*** and a set of objects ***O={o1,o2,...oj,...on}***.

For each pair ***(ui,oj)***, a recommender can compute a score ***ri,j*** that measures the expected interest of user ***ui*** in object ***oj*** (or the expected utility of object **oj** for user **ui**), using a knowledge base and a *scoring* (or *ranking*) algorithm that should take into account how users' preferences change with context.

In other terms, for each user ***u∈U***, the recommendation problem is to choose a set of items in **O** that maximize the user's utility, given the current context.

The interaction between users and object are usually represented a *Weighted Matrix*, where each element represents the interest of a given user for a given object. The challenge is to predict (using statistics, AI-data driven based, etc.) the missing values of such a matrix starting from a set of initial values.

TRADITIONAL RECOMMENDATION STRATEGIES (1/3)

Content-based

The utility ***r(u,o)*** of an item **o** is estimated using the utility ***r(u,oi)*** assigned by user ***u*** to items ***oi∈O*** that are in some way similar to item **o.**

For example, in a movie recommendation application, in order to recommend movies to user *u*, the recommender system tries to understand the commonalities among the movies that user *u* has rated highly in the past (actors, directors, genres, etc.).

Eventually, a subtle problem is that the system can only recommend items that are similar to those already rated by the user (*overspecialization*).

Collaborative filtering

The utility ***r(u,o)*** of an item ***o*** is estimated using the utility ***r(ui,o)*** assigned to ***o*** by other users similar to ***ui***, where ratings can take on a variety of forms: scalar ratings, ordinal ratings or unary rating.

The simplest method is *passive filtering*, which uses data aggregates to make predictions (such as the average rating for an item) and each user will be given the same predictions for a particular item. On the other side, *active filtering* uses patterns in a user's history to make user-specific and context-aware recommendations.

Collaborative systems have their own limitations, which mainly relate to *the cold start problem*, which describes situations in which a recommender is unable to make meaningful recommendations due to an initial lack of ratings. The most common solution is to provide rating incentives to a small bootstrap subset of items

TRADITIONAL RECOMMENDATION STRATEGIES (2/3)

COLLABORATIVE FILTERING

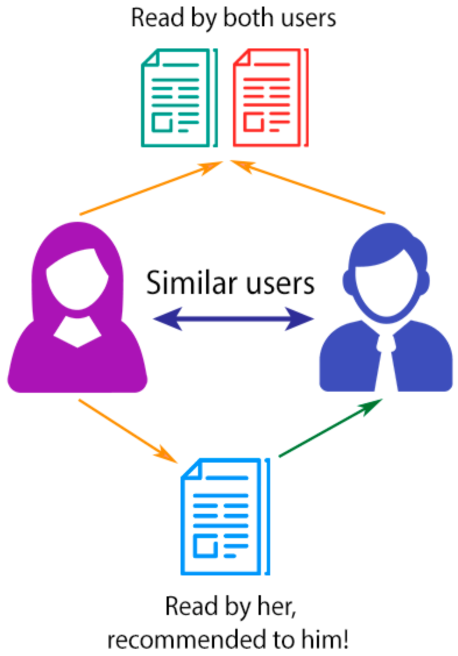

CONTENT-BASED FILTERING

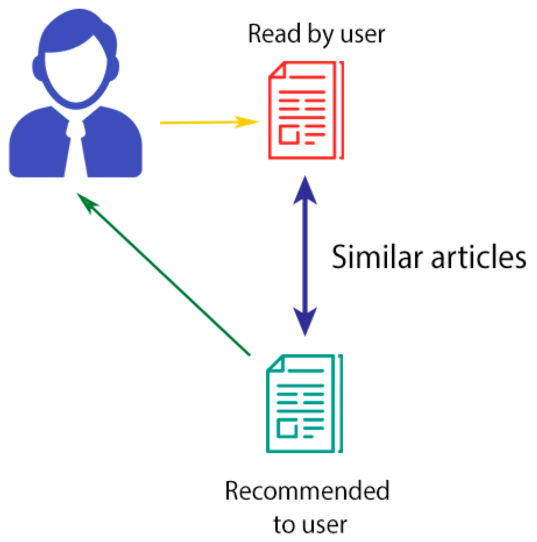

TRADITIONAL RECOMMENDATION STRATEGIES (3/3)

Hybrid

Different ways to combine collaborative and content-based methods:

- implementing collaborative and content based separately and combining their predictions;
- incorporating some content-based characteristics into a collaborative approach or viceversa;
- constructing a general unifying model.

The goal is to try to overcome the limitations of the two kinds of approaches.

Some Challenge

Importance of the last accessed items.

Sparsity and high-dimension of rating matrix.

Availability and quality of user profiles.

Proliferation of fake users that causes malicious ratings.

Different interactions (especially in social environments) between users and items can be considered as a sort of implicit rating.

Complex features of multimedia items.

Content-based similarity for multimedia items.

Difficulty in evaluating recommender systems.

Perfomances for very large data collections.

THE RECOMMENDATION PROCESS

Recommender Systems usually leverage several features to provide useful recommendations:

- user preferences and past behavior,
- preferences and past behavior of the user community,
- items' characteristics and the related similarity, and how they can match user preferences,
- user feedbacks,
- context information and how recommendations can change together with the context.

Thus, a Recommender Systems can exploit a pelthora of possibile features…

The recommendation process usually leverages one or more of the following steps:

- *prefiltering stage*: determines a set of useful candidate items for the recommendation, on the base of user actual needs and preferences;
- *ranking stage*: opportunely assigns to these items a rank, previously computed exploiting items' intrinsic features and users' past behaviors, and using as refinement, other social elements in the shape of users' opinions and feedbacks;
- *post-filtering stage*: dynamically, when a user "selects" as interesting one or more of the candidate items, determines the list of most suitable items, also considering other context information expressed by users in the shape of constraints on items' features;
- *presentation stage*; eventually, final recommended items can be arranged in specific "groups" and presented to user.

HOW TO TAKE INTO ACCOUNT MULTIMEDIA FEATURES AND INTERACTIONS BETWEEN USERS AND DATA?

Multimedia Data

- Multimedia items are carachterized by complex features:
  - *Low-level features* related to multimedia content (e.g., for images color, shape or texture of given objects)
  - *High-level features*, generally expressed as a set of metadata, an related to item semantics (e.g., keywords provding a description of objects within an image)
- It is possible to define a more fine-grained similarity between two items that leverages multimedia features by combining high and low level elements.
- Are these features really useful in the recommendation process? How it is possible to use multimedia similarity?
  - Items should be described also considering the related multimedia features.
  - In the definition of item similarity we have to consider multimedia similarity.

Interactions with multimedia data

- Interactions between users and items can be different and complex. Users can:
  - access a multimedia content multiple times (e.g., in the case of a favorite song this will surely be listened to multiple times by a user),
  - express a rating or feedback in different manners, even providing a comment,
  - interact with object along different time intervals and this time can be considered as a sort of rating (e.g., if a user looks at a picture or listens to a song for very few seconds, most likely he/she will not be interested in them),
  - Etc.
- Modeling interactions between users and media objects may therefore require data structures better suited to capture the multiplicity of relationships between users and objects w.r.t. simple weighted matrix, such as *graphs*.
- Also, presentation capabilities of Multimedia Recommender Systems have to handle different challenges as streaming…

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (1/7)

As already mentioned, the work by Moscato et al. was among the first in the literature to pose the problem of how to adapt the classical recommendation process to the nature of multimedia data…

The approach can be classified as a hybrid strategy that incorporates some content-based characteristics into a collaborative framework.

- It exploits system logs to implicitly derive information about individual users and the community of users as a whole, considering their past browsing sessions within an image repository as a sort of unary ratings.
- Similarly to collaborative filtering random walk techniques, it is a kind of active filtering strategy in which past browsing sessions, modeled as a directed graph, determine the most suitable items to be recommended.
- Similarly to rule mining approaches, transitive relationships among items are considered in computing the importance of an object.
- Similarly to content-based approaches, our approach gives high importance to the characteristics of the object a user is currently watching, in order to effectively compute the utility of other items.

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (2/7)

The proposed approach is based on an importance ranking method that strongly resembles the *Google PageRank* algorithm.

The idea is to model recommendation as a social choice problem and propose a method that computes customized recommendations by originally combining:

- *Single user implicit behaviour*, and in particular her/his browsing sessions, in the shape of directed graphs (from which we derive a *Local Browsing Matrix*), which nodes are objects and an edge between two nodes oi and oj represents the fact that in a particular browsing session of the user object ***oi*** has been accessed immediately after ***oj***. The edge is the labeled by the ratio of the number of times object oi has been accessed by the user immediately after ***oj*** to the number of times any object in **O** has been accessed by the user immediately after ***oj***.
- *Users Community behaviour,* modeled as a particular graph (from which we derive a *Global Browsing Matrix*) obtained by averaging the edge values of each local browsing graph.
- I*ntrinsec features* (low and high level) of multimedia objects by which a *Similariry Matrix* (containing for each element ***i,j*** the multimedia similarity between ***oi*** and ***oj***).

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (3/7)

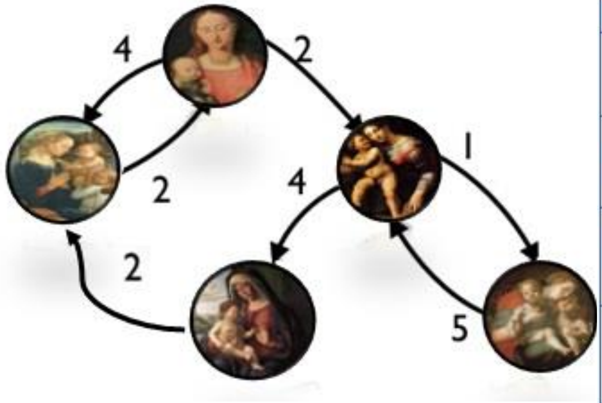

ID		19	154	204	228	391
19		0	2/2	0	0	0
154		4/6	0	0	0	2/2
204		2/6	0	0	5/5	0
228		0	0	1/5	0	0
391		0	0	4/5	0	0

Local and Global
Browsing Matrix ***Au, A***

LOW LEVEL

Color, shape, texture and position

SEMANTIC DESCRIPTION

Authors, artistic movements, subjects

Similarity Matrix **B**

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (4/7)

The proposed approach addresses two fundamental questions:

- How can we select a smaller set of objects from the collection that are good candidates for recommendation?
  - This set of candidates includes the objects that have been accessed by at least one user within ***k*** steps from ***oj***, with ***k*** between ***1*** and ***M***, and considering the objects that are most similar to ***oj***.
- How can we rank the set of candidates?
  - We use an importance ranking method that is very similar to the Google PageRank. We assume that the fact an object ***oi*** is chosen after an object ***oj*** in the same browsing session corresponds to ***oj*** voting for ***oi***.

$$\mathcal{C}_j = \bigcup_{k=1}^{M} \{o_i \in \mathcal{O} | A_{ij}^k > 0\} \cup \{o_i \in \mathcal{O} | B_{ij} > 0\}$$

ID		1	2	19	31	90	106	154	204	225	228
1		0	0.22	0.10	0	0	0.10	0.10	0	0	0.10
2		0.22	0	0.15	0.10	0.10	0	0.10	0.10	0.10	0.25
19		0.10	0.15	0	0.65	0.55	0	0.35	0.50	0.40	0.70
31		0	0.10	0.70	0	0.65	0	0.35	0.40	0.45	0.60
90		0	0.10	0.50	0.65	0	0.10	0.45	0.50	0.40	0.50
106		0.10	0	0	0	0.10	0	0.10	0	0	0
154		0.10	0.10	0.40	0.35	0.45	0.10	0	0.30	0.25	0.40
204		0	0.10	0.30	0.40	0.50	0	0.30	0	0.32	0.45
225		0	0.10	0.55	0.45	0.40	0	0.25	0.32	0	0.35
228		0.10	0.25	0.70	0.60	0.50	0	0.40	0.45	0.35	0

ID		1	2	19	31	90	106	154	204	225	228
1		0	0	0	3/4	0	0	0	0	0	0
2		0	0	0	0	2/3	0	0	0	0	3/8
19		0	0	0	0	0	0	0	0	0	0
31		0	0	0	0	0	4/5	0	0	0	4/5
90		0	0	0	0	0	0	0	0	0	0
106		0	0	0	0	0	7/8	0	0	0	0
154		1/5	0	0	1/5	0	0	2/5	0	1/6	0
204		0	0	0	0	0	0	0	0	0	0
225		0	0	3/4	0	0	0	0	0	0	2/3
228		0	0	0	0	0	0	0	0	0	0

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (5/7)

To better visualize the ideas which are behind the approach to recommendation in multimedia browsing system, a web application, which offers a virtual access to the multimedia collection of digital reproductions of Uffizi paintings, has been implemented.

Paintings set consists of 474 digital reproductions of Uffizi paintings, which belong to 144 artists such as Botticelli or Giotto etc, grouped by 16 artistic movements.

The recommender system is able to suggest paintings even if a user access to virtual museum for the first time. Thus, there is no cold start problem.

The recommender, for each image that a user is looking, suggests a list of images according to the proposed approach, but it offers also different search methods such as search by artist, search by genre and subject.

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (6/7)

During the tuning phase of the system:

- the most suitable multimedia features for images have been selected,
- the best matrix storage format has been chosen.

Concerning experiments, the following metrics have been evaluated:

- *User satisfaction* using i) empirical measurements of access complexity in terms of mouse clicks and time,ii) TLX (*NASA Task Load Index factor*);
- *Accuracy* of recommendations in terms of *Mean Absolute Error* (MAE) and the *Root Mean Square Error* (RMSE);
- *Efficiency* in terms of *computation times*.

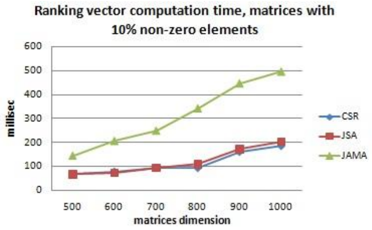

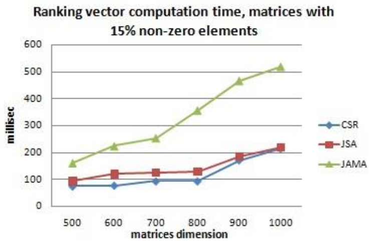

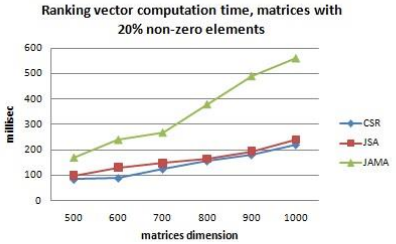

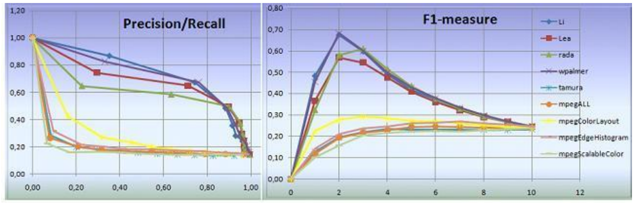

a)

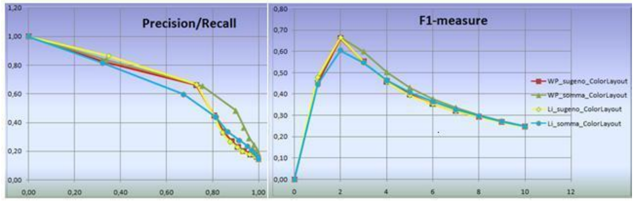

b)

AN EXAMPLE OF MULTIMEDIA RECOMMENDER SYSTEM (7/7)

TLX Factor	Expert users on art		Medium expert users on art		Not-expert users on art	
	Uffizi Gallery	Picasa	Uffizi Gallery	Picasa	Uffizi Gallery	Picasa
Mental demand	20.5	18.5	28.2	27	30	38
Physical demand	23	28	25.4	30	26.2	40
Temporal demand	26.2	31.1	27.6	32	27	31
Effort	23.3	28.2	28.7	37.2	32.3	46.4
Performance	70.1	66.3	70.6	69.3	60.2	70.1
Frustation	20.5	29	22.3	30.5	23.4	31.6

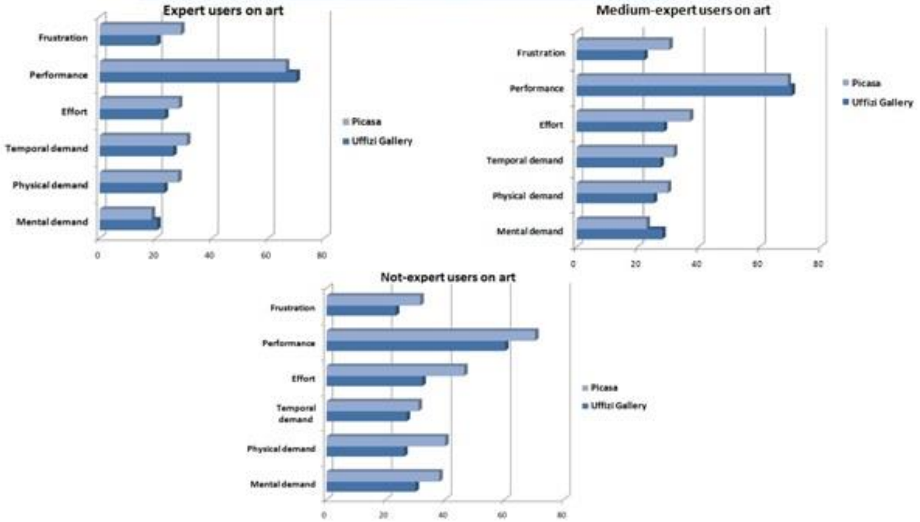

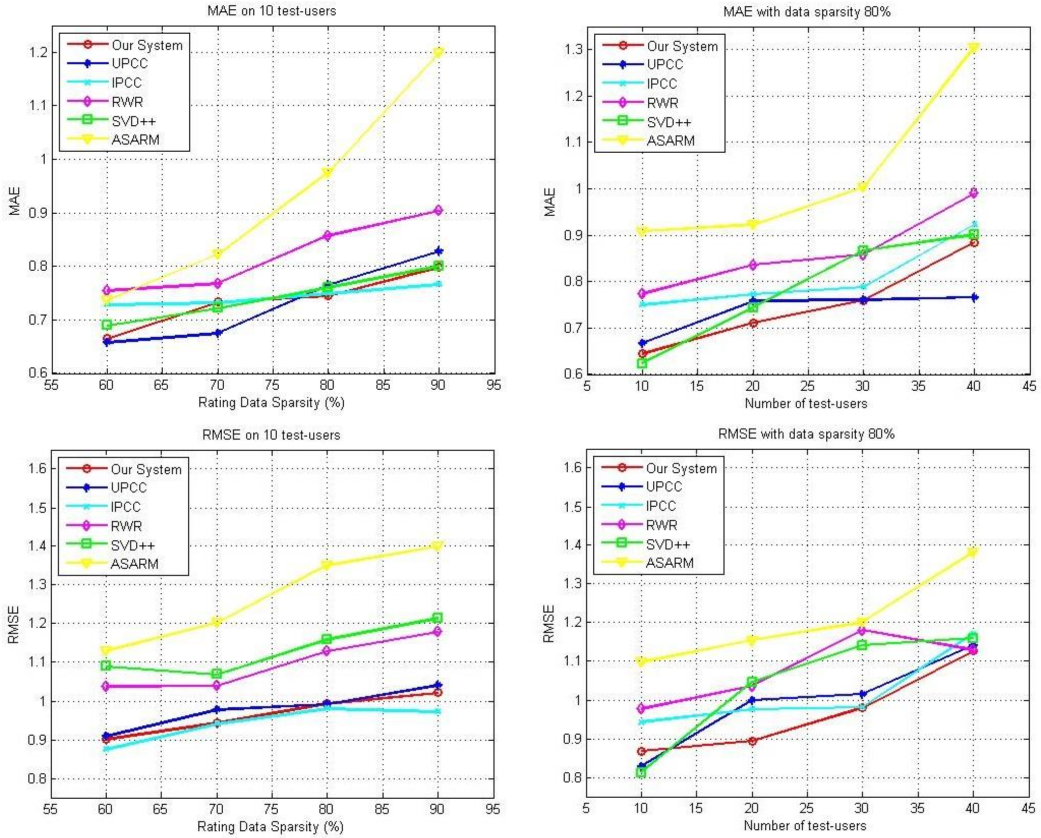

SEVERAL CHALLANGES

The proposed approach poses several reserach questions that have to be addressed:

- Is the methodology suitable of other kinds of multimedia data?
- Is the fact that a user has accessed an object really indicative that the unuser is interested in it?
- Should browsing time on a given object be taken into account?
- The browsing information can be combined with other forms of rating?
- What is the order of magnitude of data collections that begins to affect system performance? Can the multimedia similarity computation slow-down the recommendation process?
- What happen if multimedia data is not the object to be recommended but its component?
- Is it possible to group recommendations in proper groups?

THE FIRST WAVE OF APPROACHES FOR MULTIMEDIA RECOMMENDER SYSTEMS

Starting from the work by Moscato et al., several approches have been proosed in the literature for multimedia recommendation issue:

- by defining novel strategies that however take into account some of elements of the proposed strategy…
- by trying to improve in different way the basic idea behind the approach…

Some interesting recent surveys have schematized the research efforts of the first wave and the introduced novelties are basically the following:

- To extend the recommendation strategy to other kinds of multimedia data or heteregeous multimedia data
- To consider user interaction with multimedia in defining the related profile
- The majorty of them proposes content-based strategies customized for a single type of multimedia content or collaborative filtering methods that leverage visual and aural properties of items interacted with in modeling a user’s profile.
- They highlight that multimedia content can be also useful to recommend items that are not necessarily media types but may also be generic items (e.g., cultural heritage, tourism attractions, food, fashion,social media).
- Some of them poses the problem of arranging recommended items into groups (*Group Recommender Systems*)

THE SECOND WAVE OF APPROACHES FOR MULTIMEDIA RECOMMENDER SYSTEMS

In the second wave of multimedia recommender systems, the focus is on the user and in figuring out which media content can best meet their needs and preferences based on both the relevant features and the behavior of similar users, resulting in *user-centered* approaches.

In addition, the new techniques consider further issue to allow the usage of such system in real contexts:

- User profiles are very complex objects to describe, and they must consider in some way multimedia features.
  - User *mood* in another important feature to consider…
- Users can interact with multimedia content in different way and each one can contribute to express a sort of rating or to define user communities.
- Multimedia content of the is not the item to be suggested but only a part of the recommendable objects…
  - e.g., In the Cultural Heritage recommended object are cultural point of interest that can have also a multimedia description (images and texts)
- Recommendations can be affected by the context in which users use a Recommender Systems (*context-based* approches).
- Recommender System have to deal with very large set of objects and users (*Large Scale Recommender Systems*).

A MOTIVATING EXAMPLE (1/2)

A user who desires to have information about the coming soon movies…

- **user preferences** in terms of movies' metadata (e.g., favorite genre, director, stars, etc.),
- **item features** (i.e., movies' metadata) and their "similarity" also considering multimedia content (e.g., images of film poster),
- **user behavior** in terms of the sequence of items that in the past the community of users have observed and positively rated;
- **user feedbacks**,
- **user opinions** in terms of the average sentiment that items have aroused on the user community,
- **context information** (e.g. coming soon movies shown in theaters near the user or that have a good similarity with respect to the item that the user has recently selected).

A lot of aspects can influence and dtermine a good recommendation

A MOTIVATING EXAMPLE (2/2)

The user prefers the adventure and fantasy genres and has as favorite actors Ian McKellen and Hugh Jackman…

The system can initially suggest as first items to watch the X-Men saga movies…

The candidate items matching user preferences are initially ranked on the base of the related social popularity and similarity with user preferences…

Eventually, if the user chooses to limit the search to the coming soon movies and selects his/her position as context information, all the best movies matching user preferences that are showing in the next days in theaters near the user will be finally proposed…

Location-based theater booking services can be invoked…

A USER CENTERED APPROACH FOR MULTIMEDIA RECOMMENDER SYSTEMS

The Moscato et al. work has been extended by the authors trying to address the described issues…

A novel user-centered approach exploiting several aspects related to users:

- *preferences* (usually coded in the shape of items' metadata),
- *opinions* (textual comments to which it is possible to associate a particular sentiment),
- *behavior* (in most cases, logs of past items' observations/interactions made by users),
- *feedbacks* (usually expressed in the form of ratings)

All users' features are integrated together with:

- *items' features* (also multimedia descriptors) thar can be constituted by heterogeneous information…
- *context information* …

within a general recommendation framework that can support different applications using proper customizations…

- e.g. recommendation of news, photos, movies, travels, cultural objects, social element in OSNs, etc.

A USER CENTERED APPROACH:
PRE-FILTERING STAGE USING CO-CLUSTERING TECHNIQUES

We determine a set of useful items on the base of user actual needs and preferences…

- …for each user $\boldsymbol{u_h}$ we select a subset $\boldsymbol{O_h^c} \subset \boldsymbol{O}$ containing items that are good "candidates" to be recommended
- Items are represented in heterogeneous feature spaces $\mathbf{F} = \{\boldsymbol{F^1}, \dots, \boldsymbol{F^l}\}$ and are then simultaneously co-clustered in the various spaces…
- Each user is representd by a set of vectors in the same $\mathbf{l}$ feature spaces…
- We measure the *cosine distance* of the user vectors associated to the $\mathbf{k-th}$ space, with the centroids of each item clusters in the $\mathbf{k-th}$ space…
- For each space, the most similar items' cluster is chosen leading to l clusters $\{\boldsymbol{X_1^c}, \dots, \boldsymbol{X_l^c}\}$ of candidate items…
- List of candidate items is finally obtained using *set-union* or a *threshold* strategy…

A USER CENTERED APPROACH:

RANKING STAGE USING USER BEHAVIOR AND ITEMS SIMILARITY (1/2)

We automatically rank the set of items using the ranking technique that authors have proposed in previous works…

- combining low and high level features of items, past behavior of individual users and overall behavior of the whole user "community"…

ID		1	2	19	31	90	106	154	204	225	228
1		0	0.22	0.10	0	0	0.10	0.10	0	0	0.10
2		0.22	0	0.15	0.10	0.10	0	0.10	0.10	0.10	0.25
19		0.10	0.15	0	0.65	0.55	0	0.35	0.50	0.40	0.70
31		0	0.10	0.70	0	0.65	0	0.35	0.40	0.45	0.60
90		0	0.10	0.50	0.65	0	0.10	0.45	0.50	0.40	0.50
106		0.10	0	0	0	0.10	0	0.10	0	0	0
154		0.10	0.10	0.40	0.35	0.45	0.10	0	0.30	0.25	0.40
204		0	0.10	0.30	0.40	0.50	0	0.30	0	0.32	0.45
225		0	0.10	0.55	0.45	0.40	0	0.25	0.32	0	0.35
228		0.10	0.25	0.70	0.60	0.50	0	0.40	0.45	0.35	0

ID		1	2	19	31	90	106	154	204	225	228
1		0	0	0	3/4	0	0	0	0	0	0
2		0	0	0	0	2/3	0	0	0	0	3/8
19		0	0	0	0	0	0	0	0	0	0
31		0	0	0	0	0	4/5	0	0	0	4/5
90		0	0	0	0	0	0	0	0	0	0
106		0	0	0	0	0	7/8	0	0	0	0
154		1/5	0	0	1/5	0	0	2/5	0	1/6	0
204		0	0	0	0	0	0	0	0	0	0
225		0	0	3/4	0	0	0	0	0	0	2/3
228		0	0	0	0	0	0	0	0	0	0

- Our basic idea is to assume that:
  - when an object is chosen after another one in the same browsing session, this event means that first object "is voting" for the second one …
  - similarly, the fact that an object is very similar in terms of features to another one, it can also be interpreted as the first recommending the second one (and viceversa)…
- We model a browsing system for the set of candidate objects as a *graph* where each edge is labeled with:
  - a *pattern label* denoting the number of times an object was accessed immediately after another one…
  - a *similarity label* for denoting the similarity between the two objects…

A USER CENTERED APPROACH:
RANKING STAGE USING USER BEHAVIOR AND ITEMS SIMILARITY (2/2)

As in *PageRank* algorithm, *the item importance* is then measured by means the introduction of a ranking function $\boldsymbol{\rho(o)}$…

- The ranking vector $\boldsymbol{R} = [\boldsymbol{\rho}(\boldsymbol{o_1}), \ldots, \boldsymbol{\rho}(\boldsymbol{o_n})]$ of all the objects can be computed as the solution to the equation $\boldsymbol{R} = \boldsymbol{C} \cdot \boldsymbol{R}$
  - where $\boldsymbol{C}$ is an ad-hoc matrix that defines how the importance of each object is transferred to other objects…
  - Such a matrix can be seen as a linear combination of a *local browsing matrix*, a *global browsing matrix* and a *similarity matrix*…

In particular conditions, the equation can be solved using the *Power Method* algorithm…

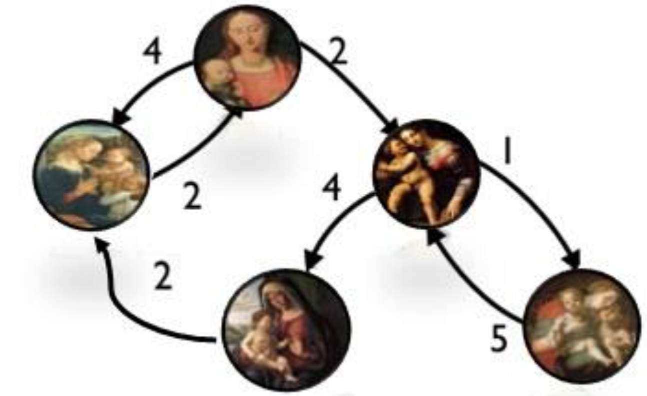

ID		19	154	204	228	391
19		0	2/2	0	0	0
154		4/6	0	0	0	2/2
204		2/6	0	0	5/5	0
228		0	0	1/5	0	0
391		0	0	4/5	0	0

A USER CENTERED APPROACH:

REFINING ITEMS RANKS USING USER SENTIMENTS AND FEEDBACKS

We used the sentiment extraction technique as an improvement of the approach presented by some of the authors in a previous work…

- …where *LDA* has been adopted for mining the sentiment inside documents…
- User comments are represented by a *mixed Graph of Terms* (mGT) that contains the most discriminative words and the probabilistic links between them…

For the ranking refinement, we introduce two probabilities $\boldsymbol{P}^{+}$ and $\boldsymbol{P}^{-}$ which express the probability that a sentiment, extracted from the set of comments related to a given item, is "positive" or "negative" (also feedbacks are considered)…

Such probabilities are then combined with the overall rank of an item by a proper function that

- increases the recommendation grade value if the sentiment within item's comments is positive, in the opposite decreases it in the case of negative mood…

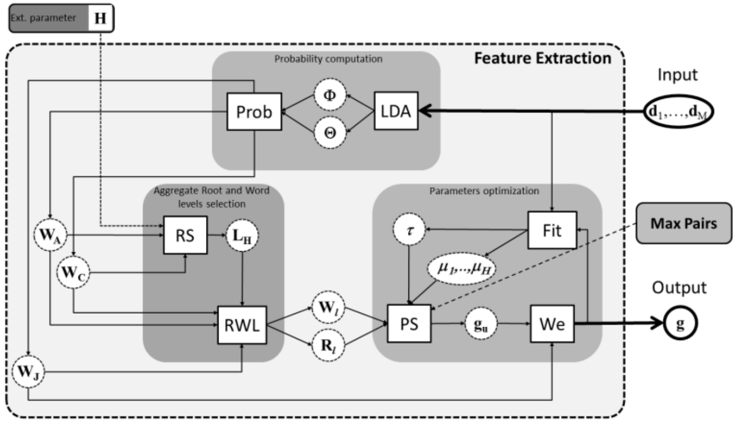

A USER CENTERED APPROACH:

POST-FILTERING USING CONTEXT INFORMATION

The context is represented by means of the well-known *key-value model* using as dimensions some of the different feature spaces related to items…

Context features can be expressed either directly

- using some target items (e.g. objects that have positively captured user attention),
- or specifying the related values in the shape of constraints that recommended items have to satisfy.

The set of final candidates includes the items that have been accessed by at least one user within k steps from the target object, and the items that are most similar to the target one according to the results of a *Nearest Neighbor Query*…

List of recommendations is then generated by ranking the items in for each item selected as interesting by a given user…

For each user all the items that do not respect possible *context constraints* are progressively removed from the final list…

A USER CENTERED APPROACH: EXPERIMENTAL RESULTS (IMDB + MOVIELENS DATASET)

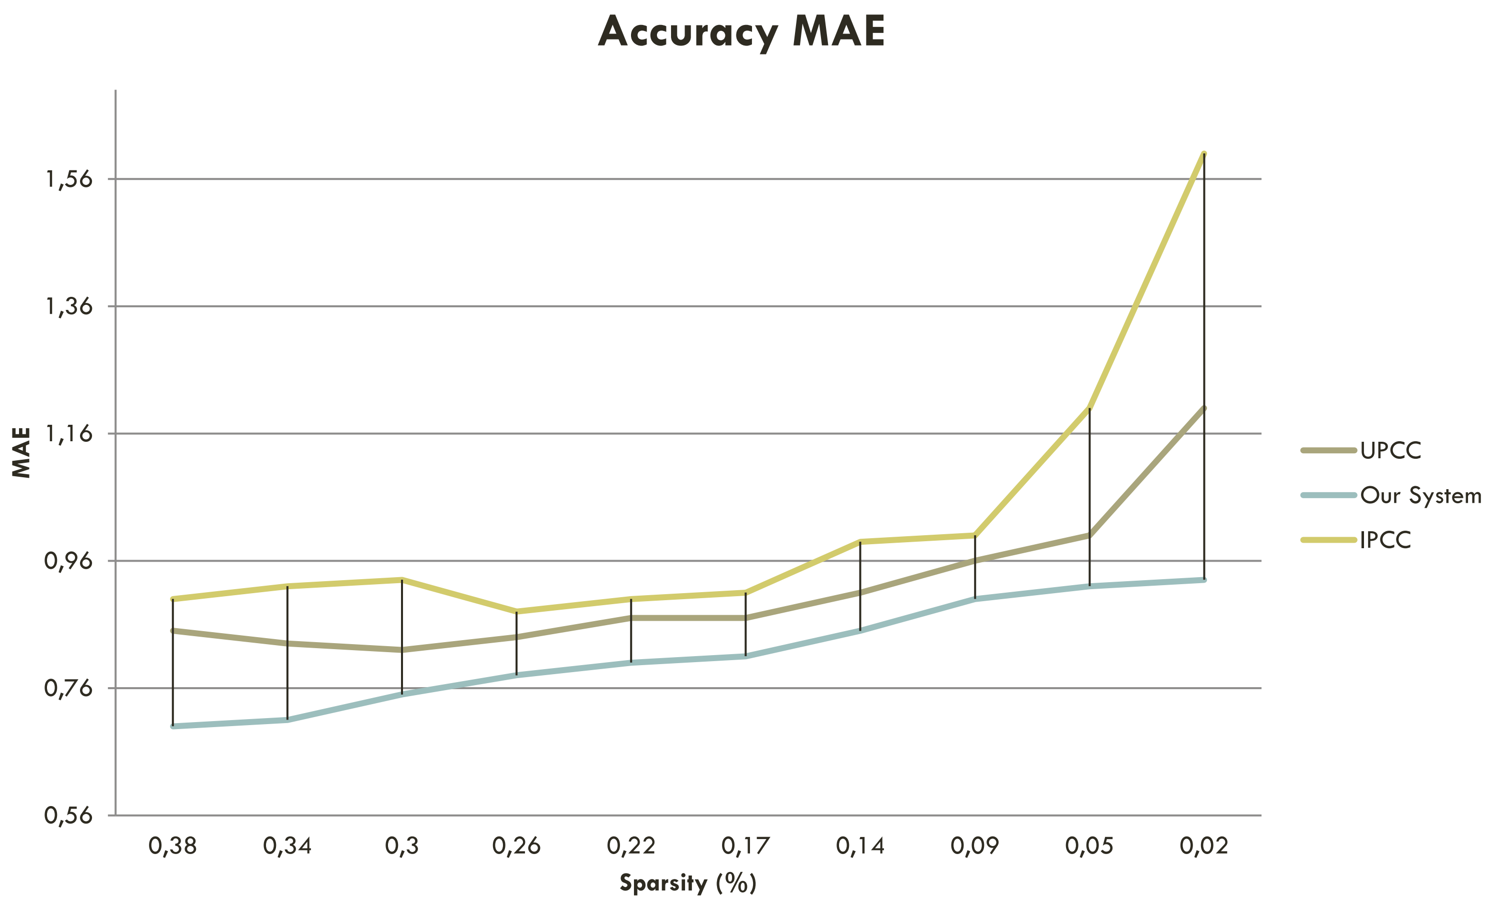

A USER CENTERED APPROACH:
EXPERIMENTAL RESULTS
(IMDB + MOVIELENS DATASET)

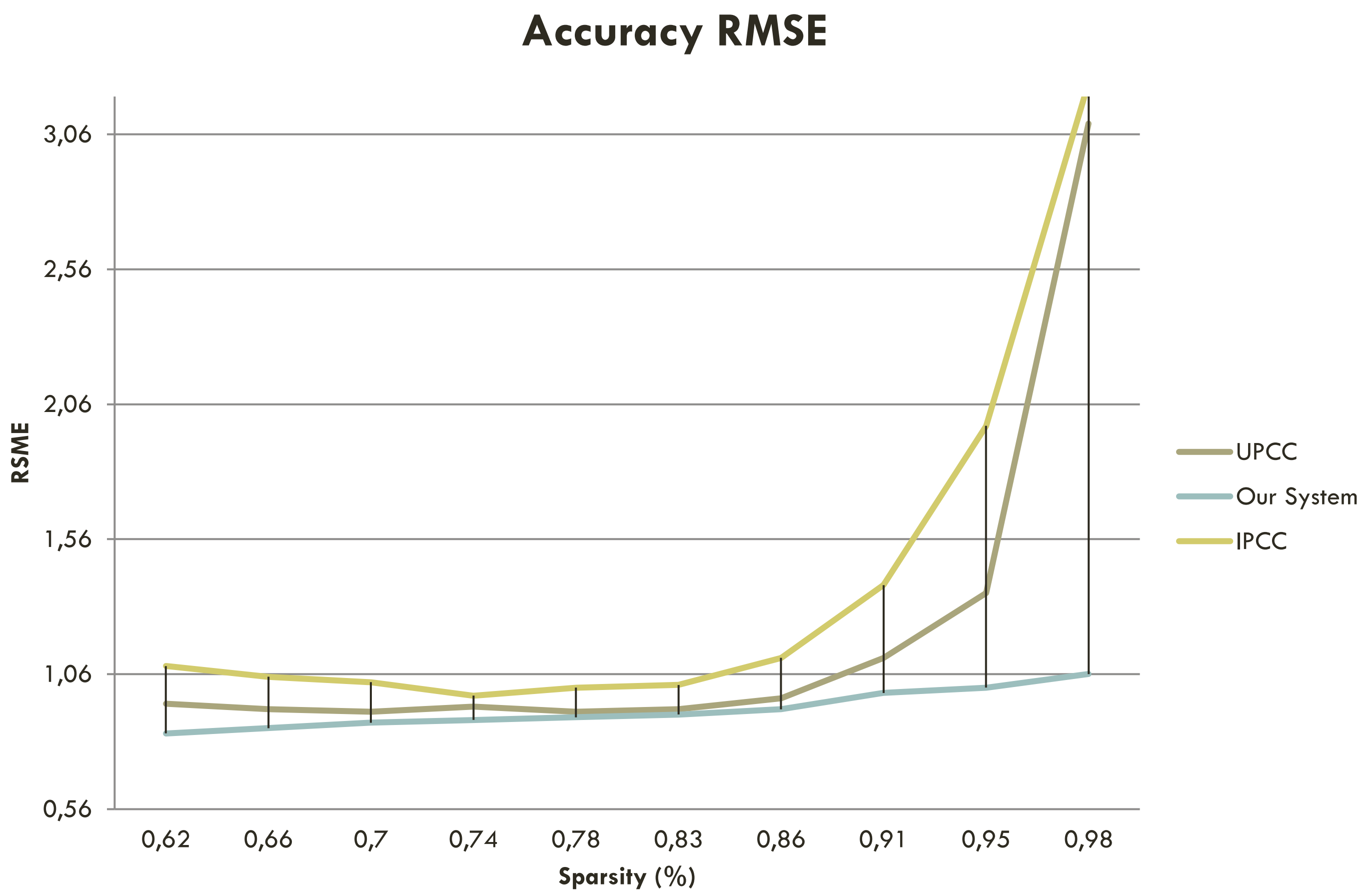

MANY YEARS LATER…THE NEFLIX ALGORITHM

Whenever you access the Netflix service, the recommendations system strives to help you find a show or movie to enjoy with minimal effort.

The likelihood that you will watch a particular title in the catalog is estimated on the basis of a number of factors including:

- interactions with service (such as your viewing history and how you rated other titles),
- other members with similar tastes and preferences on our service, and
- information about the titles, such as their genre, categories, actors, release year, etc.

In addition to knowing what users have watched on Netflix, to best personalize the recommendations we also look at things like:

- the time of day users watch,
- the devices users are watching Netflix on, and
- how long users watch.

All of these pieces of data are used as inputs that we process in the algorithm…

- A Large Scale Recommender System leveraging *RecSysOps* best practices
- Reinforcement Learning for Budget Constrained Recommendations
- Deep-learning and non-deep-learning approaches
- In-Session Adapted Recommendations

THE REMAINING CHALLANGES: THE NEW GENERATION OF MULTIMEDIA RECOMMENDER SYSTEMS

Even if different extensions to the initial Moscato et al. Recommendation framework have been performed, several challenges still remain for a possible application in real scenarios:

- Improving the efficiency performances addressing Large Scale Recommender Systems requirements and RecSysOPs best practices…
- Considering more complex user profile (*Demographic-based Recommender Systems*) and user mood in the different stages of recommendation process to obtain more personalized and dynamic recommendation services…
- Adopting with the increasing of items' features *deep-learning* recommendation strategies…
- Leveraging different data structure to capture more complex interactions between users and items…

All these issues should be addressed by the new generation of Multimedia Recommender Systems…

AN EMOTIONAL RECOMMENDER SYSTEM FOR MUSIC (1/2)

More recenlty, a Music Recommender System, which considers in the recommendation process user mood to provide more personalized recommendation, has been proposed by Moscato et al. In IEEE Intelligent Systems…

- Recent studies have demonstrated as user personality can effectively provide a more valuable information to significantly improve recommenders' performance, especially considering behavioural data captured from social network logs.
- The recommendation technique is based on the identification of personality traits, moods and emotions of a single user, starting from solid psychological observations recognized by the analysis of user behavior within a social environment.
- Users personality and mood have been embedded within a recommendation approach to obtain more accurate and dynamic results.

AN EMOTIONAL RECOMMENDER SYSTEM FOR MUSIC (2/2)

The proposed recommendation process for audio contents' suggestion works following three different steps:

- *User personality recognition*: user personality is computed in terms of Big Five components (OCEAN), considering user's behavior in terms of user profile within social networks and proper classifiers. The obtained OCEAN user personality traits are subsequently mapped into *Mehrabian's Pleasure-Arousal- Dominance* (or PAD) emotional state space.
- *Content-based recommendation*: audio contents (mapped into the PAD space) are suggested to users on the basis of their content and the related similarity w.r.t. user profile.
- *Mood detection*: the last accessed objects are analyzed to discover current user mood that is then used to refine our recommending strategy.

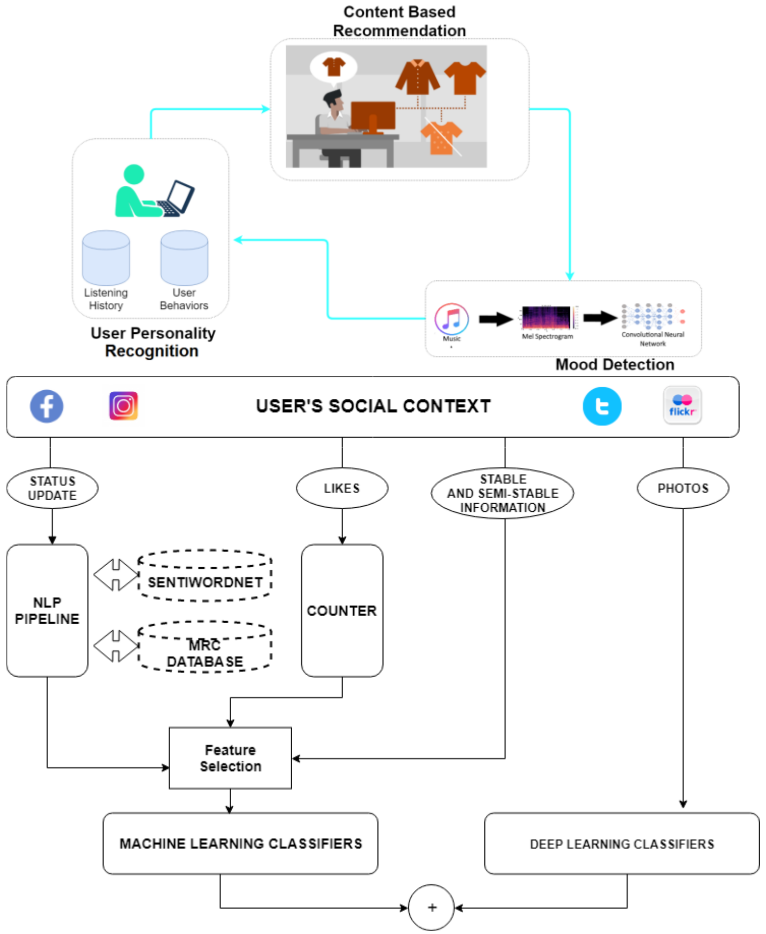

MUSIC RECOMMENDATION VIA HYPERGRAPH EMBEDDING (1/4)

Moscato et al. also propose a novel framework in IEEE Trans. On Neural Network and Learning Systems (2022) for song recommendation based on hypergraph embedding.

- The hypergraph data model allows to represent seamlessly all the possible and complex interactions between users and songs with the related characteristics;
- …meanwhile, embedding techniques provide a powerful way to infer the user–song similarities by vector mapping.
- A simple cosine-based distance is used to obtain the final recommendation.
- They experimented the effectiveness and efficiency of our approach with respect to the state-of-the-art most recent music recommender systems, exploiting the Million Song dataset. The results show that the approach significantly outperforms other state-of-the-art techniques, especially in scenarios where the cold-start problem arises.

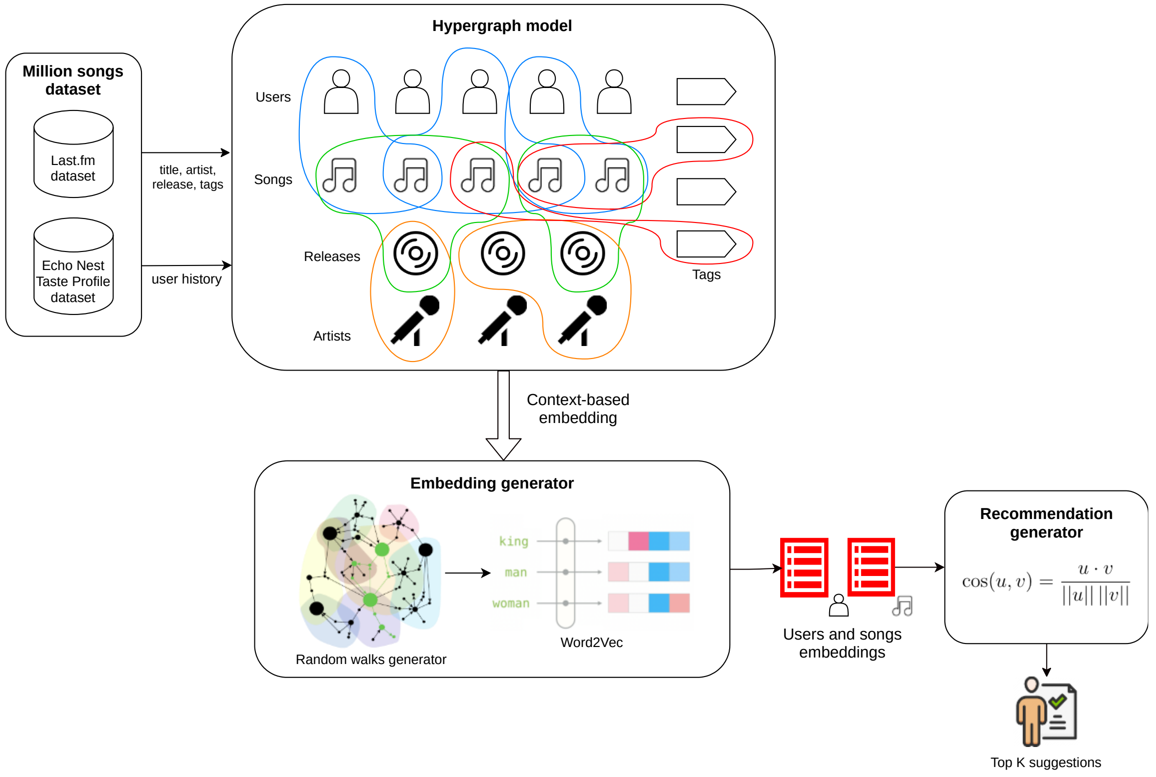

MUSIC RECOMMENDATION VIA HYPERGRAPH EMBEDDING (2/4)

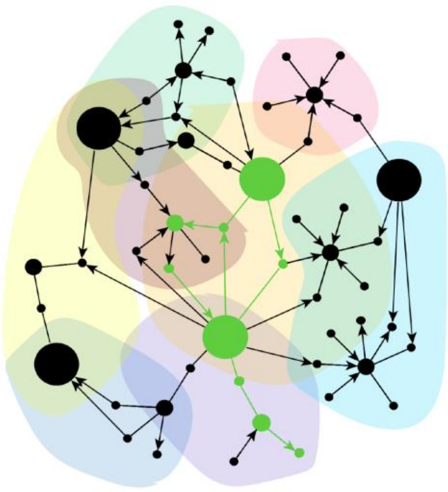

Random walks generator

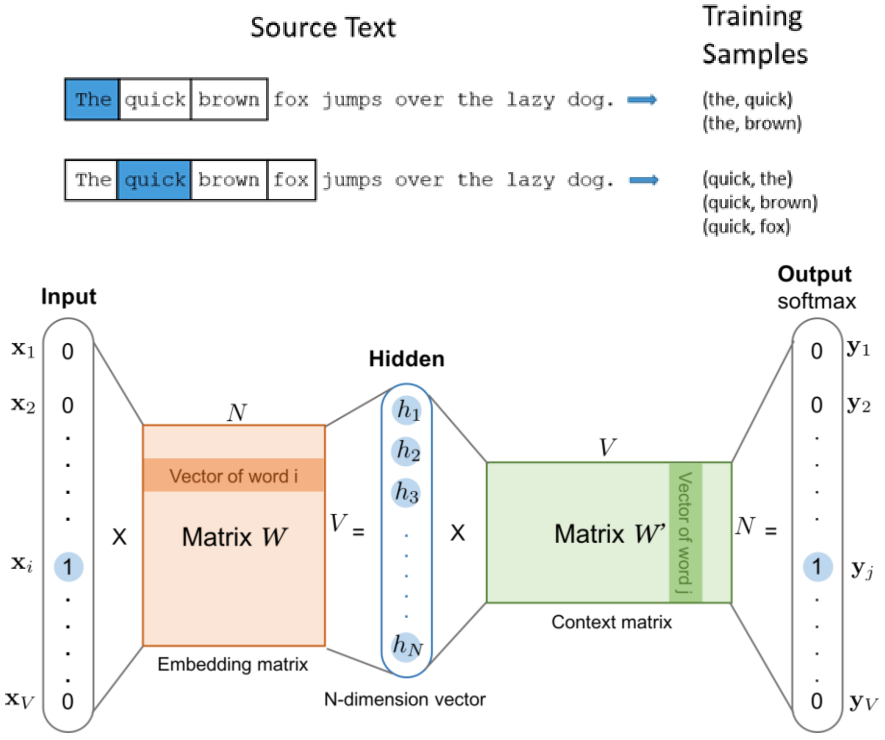

Word2vec (skip-gram)

MUSIC RECOMMENDATION VIA HYPERGRAPH EMBEDDING (3/4)

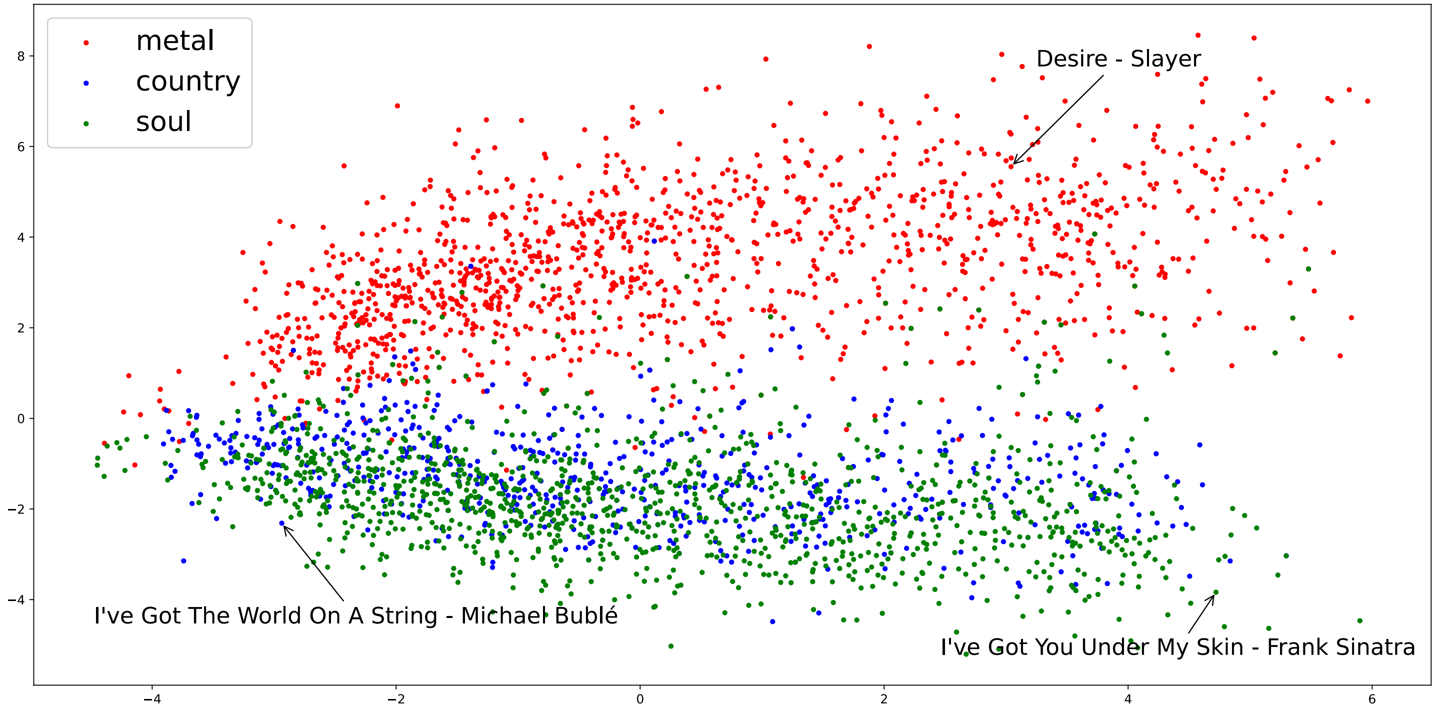

MUSIC RECOMMENDATION VIA HYPERGRAPH EMBEDDING (3/4)

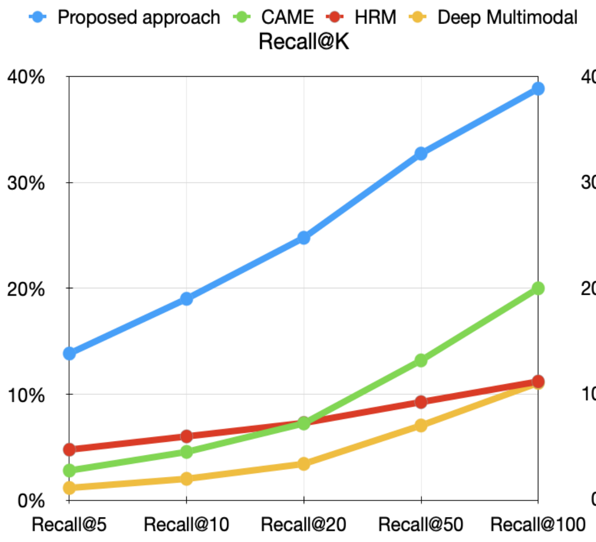

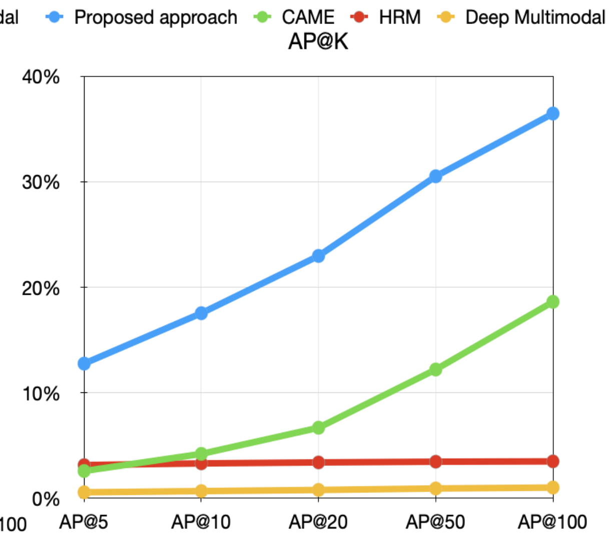

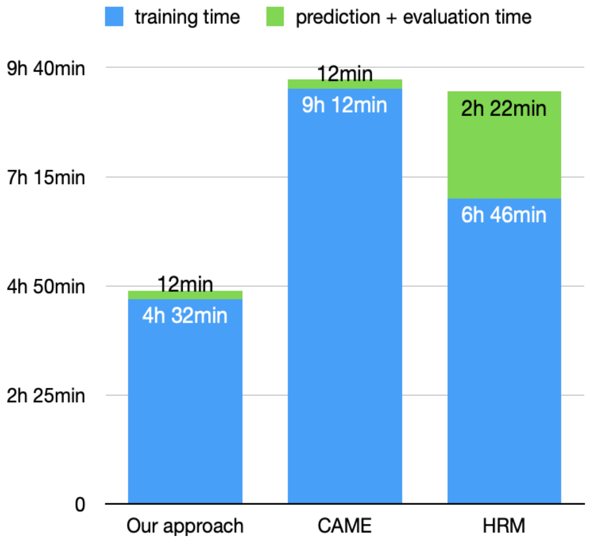

CONCLUSIONS

Despite the enormous research efforts made in the area of Multimedia Recommender Systems the open research challenges are still many, and as seen, they concern:

- Recommendation strategies able to address RecSysOps best practices and Big Data computation challenges;
- Recommendation models capable of effectively capture and handle all the possible interactions between users and multimedia content;
- Recommendation techniques able to leverage deep-learning models in order to improve significantly the system performances in some situations;
- Recommendation apporaches which fully consider complex user profile information;
- And so on…