SmallWorlds: Visualizing Social Recommendations

Brynjar Gretarsson, John O'Donovan, Svetlin Bostandjiev, Christopher Hall, Tobias Hollerer

Problema – Dataset

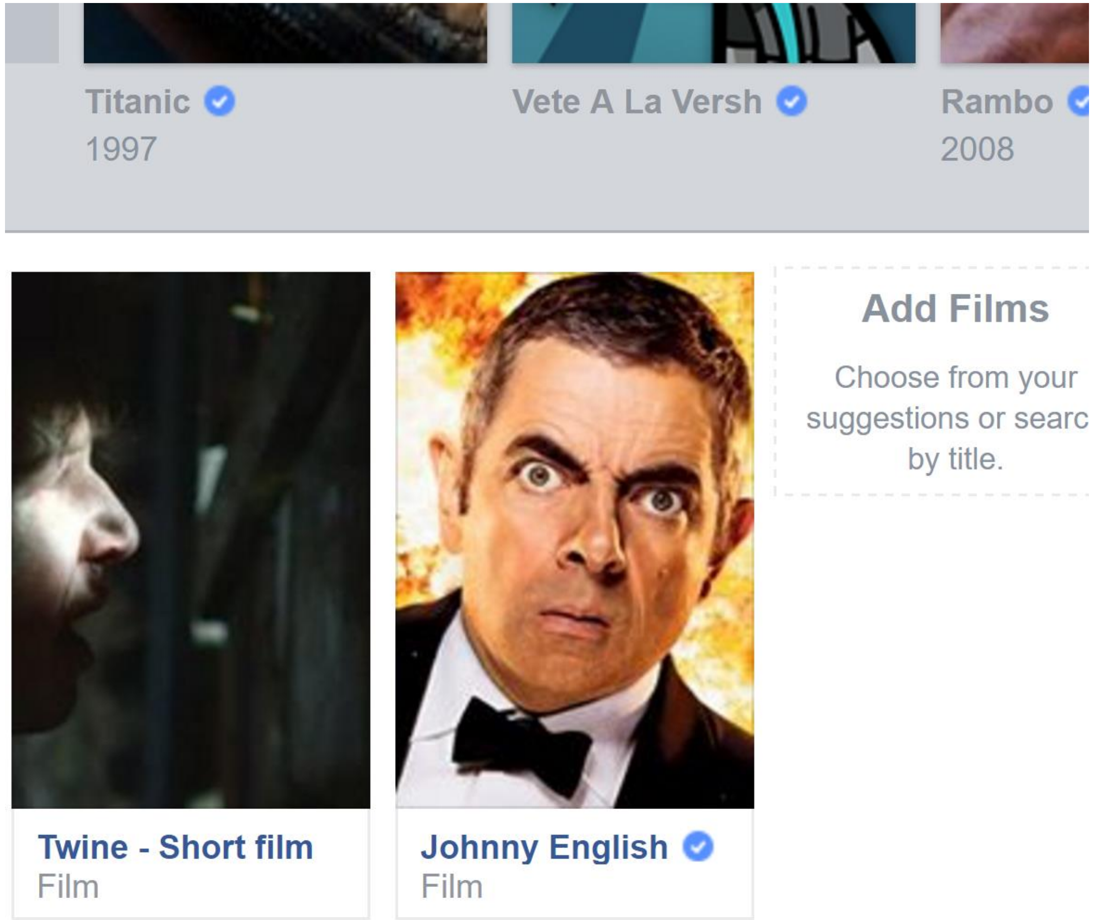

- Usuario activo – en su FB cuenta
- Items – Peliculas, musica
- Acceso restrigido
- Visualización y interacción

Human-computer Interaction (HCI)

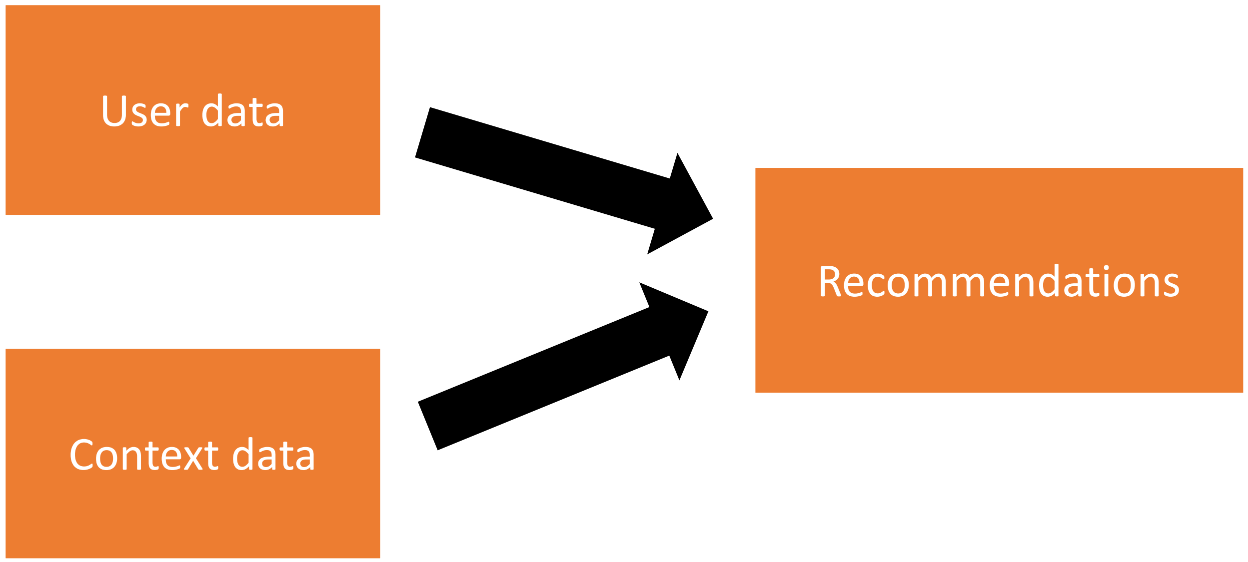

Human-computer Interaction (HCI)

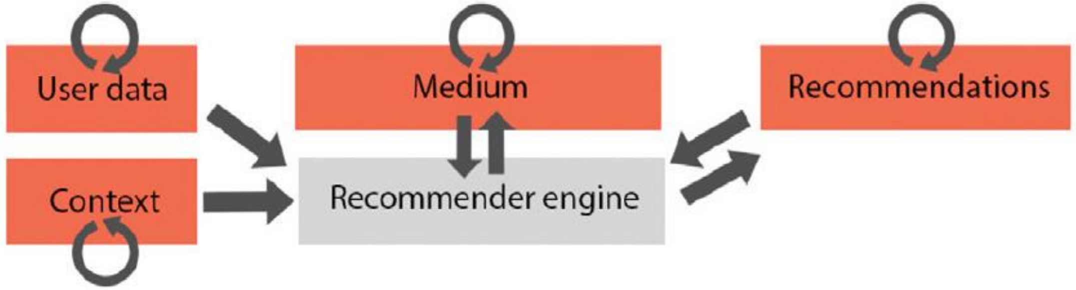

Fig. 1. Interactive recommender framework.

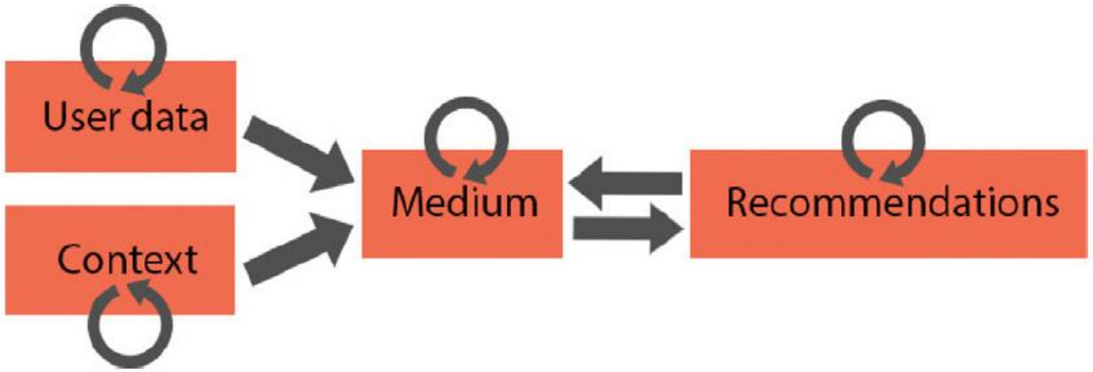

Fig. 2. User mental model of interactive recommender systems.

Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities

- Medium node as extra step
- Arrows = interactions

- User satisfaction
- Trust
- Transparency
- Sense of control

Preguntas de investigación

1. Visualización -> transparency + satisfaction
2. Interacción -> sense of control
3. Visualización -> “ambient information”
4. Social connections -> satisfaction + accuracy

Main topics

www.shutterstock.com · 255006670

Interfaz

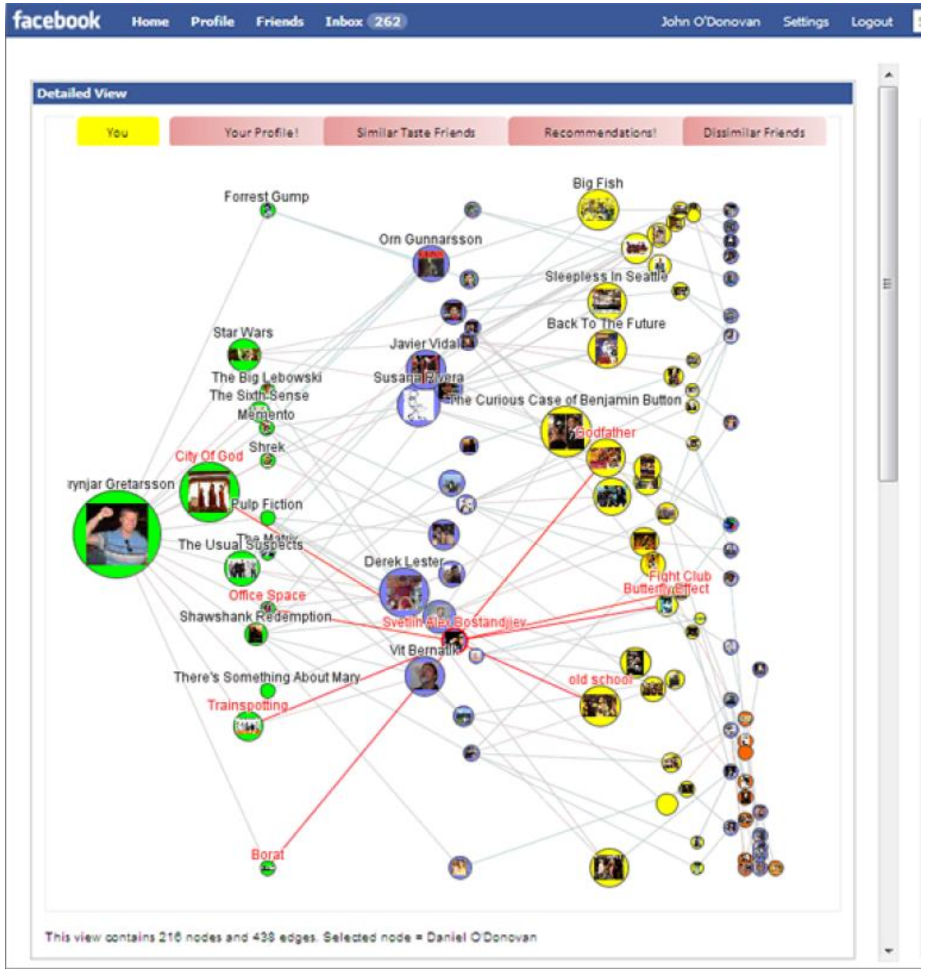
facebook
Home
Profile
Friends
Inbox 262
John O'Donovan
Settings
Logout
Detailed View
You
Your Profile!
Similar Taste Friends
Recommendations!
Dissimilar Friends
Forrest Gump
Big Fish
Orn Gunnarsson
Sleepless In Seattle
Star Wars
Back To The Future
Javier Vidal
The Big Lebowski
Susana Rivera
The Sixth Sense
The Curious Case of Benjamin Button
Memento
Godfather
City Of God
Shrek
Pulp Fiction
The Usual Suspects
Derek Lester
Fight Club
Butterfly Effect
Office Space
Shawshank Redemption
Svetlin Alex Bostandjiev
Vit Bernatik
There's Something About Mary
old school
Trainspotting
Borat
This view contains 216 nodes and 438 edges. Selected node = Daniel O'Donovan

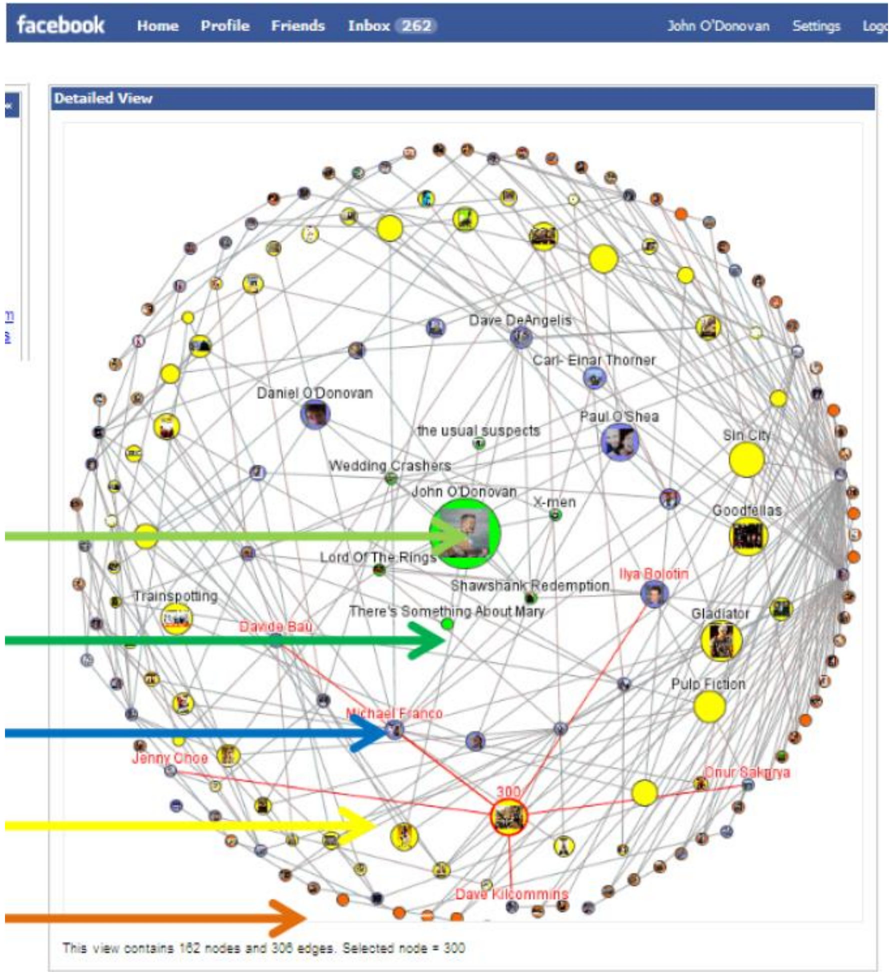
facebook
Home
Profile
Friends
Inbox 262
John O'Donovan
Settings
Detailed View
Dave DeAngelis
Carl- Einar Thorner
Daniel O'Donovan
Paul O'Shea
the usual suspects
Sin City
Wedding Crashers
John O'Donovan
X-men
Goodfellas
Lord Of The Rings
Shawshank Redemption
Ilya Bolotin
Trainspotting
There's Something About Mary
Gladiator
Davide Bau
Pulp Fiction
Michael Franco
Jenny Choe
Onur Sakarya
300
Dave Kilcommins
This view contains 162 nodes and 306 edges. Selected node = 300

Interfaz

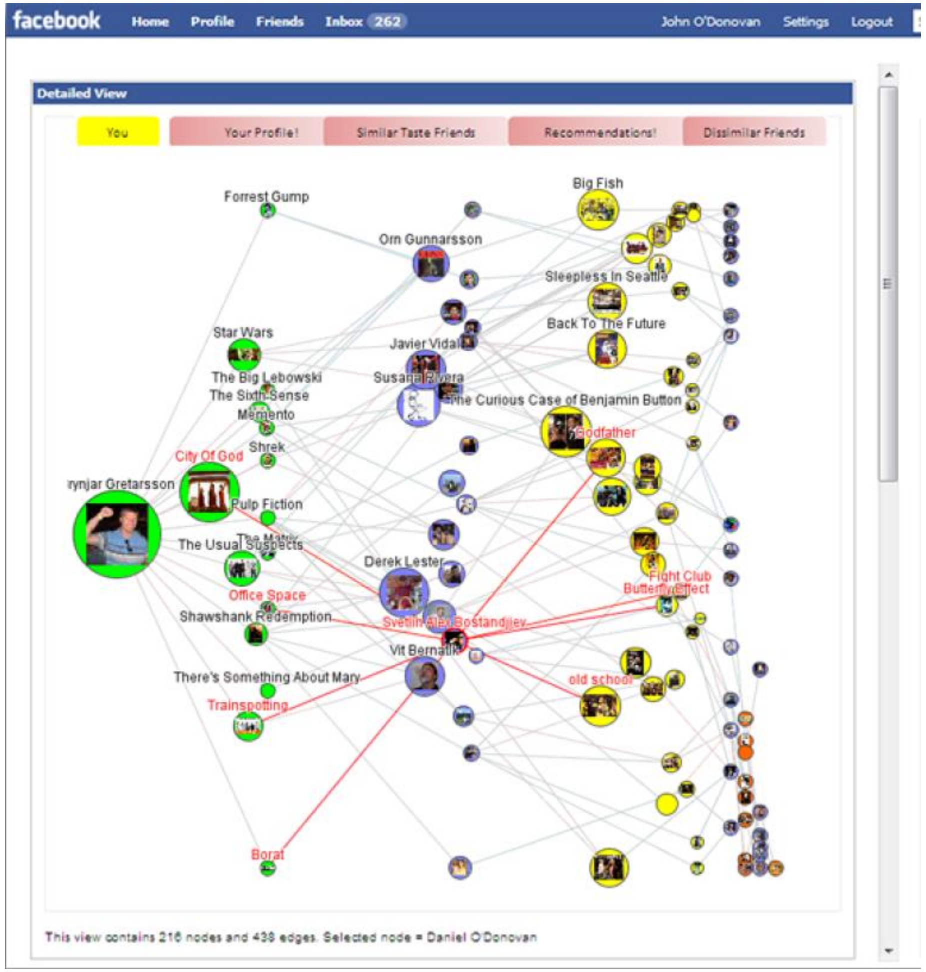

- Representacion
- Posición
- Visualización
- Interaccion
- Config. inicial
- Algoritmo

Interfaz

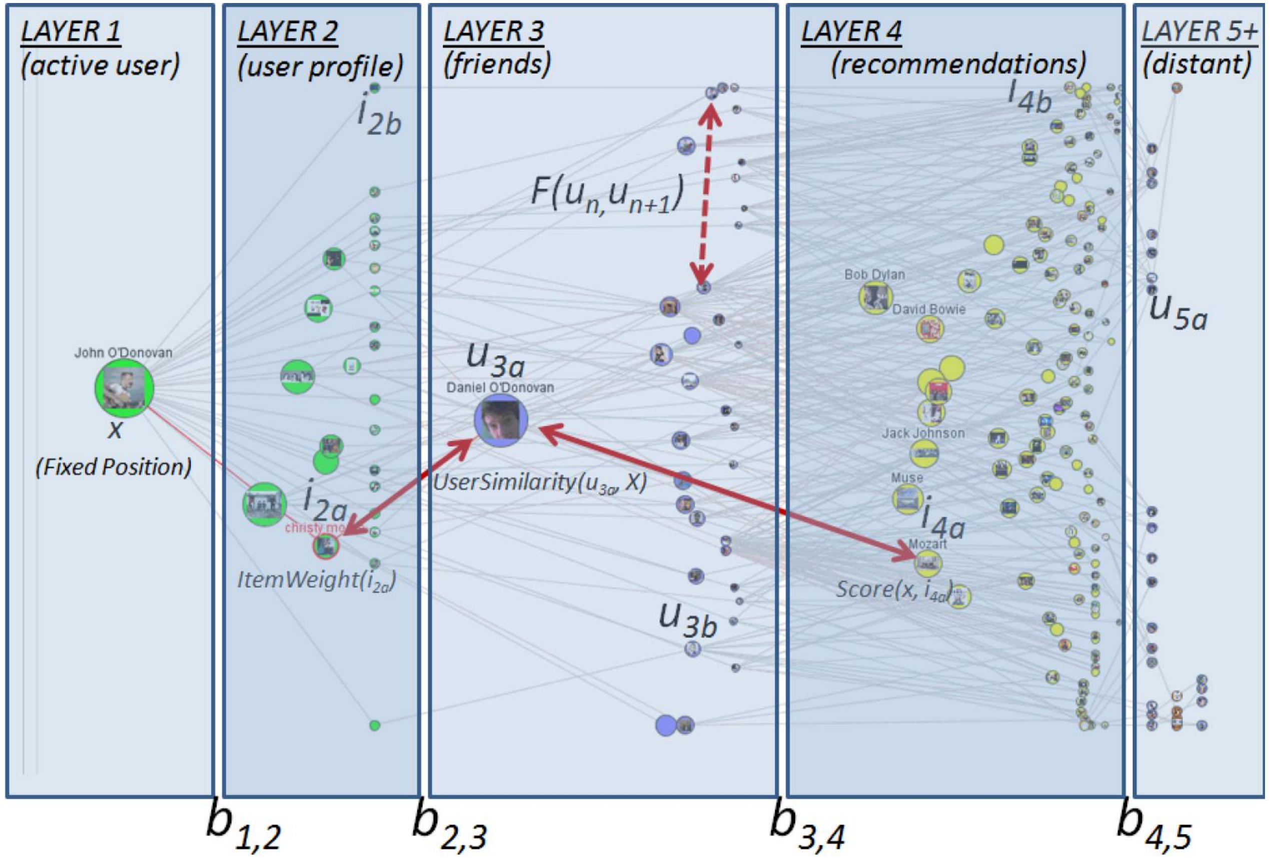

A

F

I_A

I_f

Capa 1

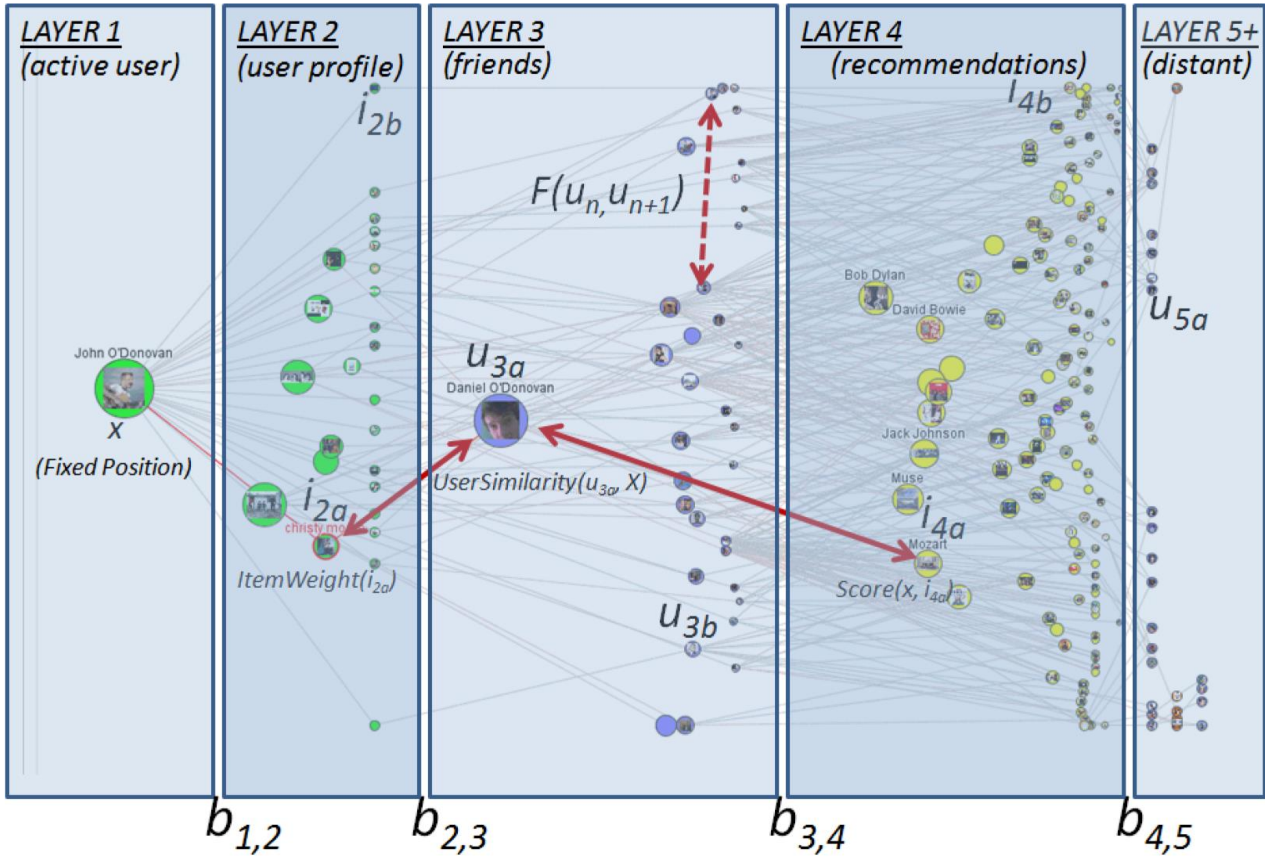
LAYER 1
(active user)
LAYER 2
(user profile)
LAYER 3
(friends)
LAYER 4
(recommendations)
LAYER 5+
(distant)
John O'Donovan
x
(Fixed Position)
i_{2b}
i_{2a}
ItemWeight(i_{2a})
$F(u_{n,}u_{n+1})$
u_{3a}
Daniel O'Donovan
UserSimilarity(u_{3a}, X)
u_{3b}
i_{4b}
Bob Dylan
David Bowie
Jack Johnson
Muse
i_{4a}
Mozart
Score(x, i_{4a})
u_{5a}
$b_{1,2}$
$b_{2,3}$
$b_{3,4}$
$b_{4,5}$

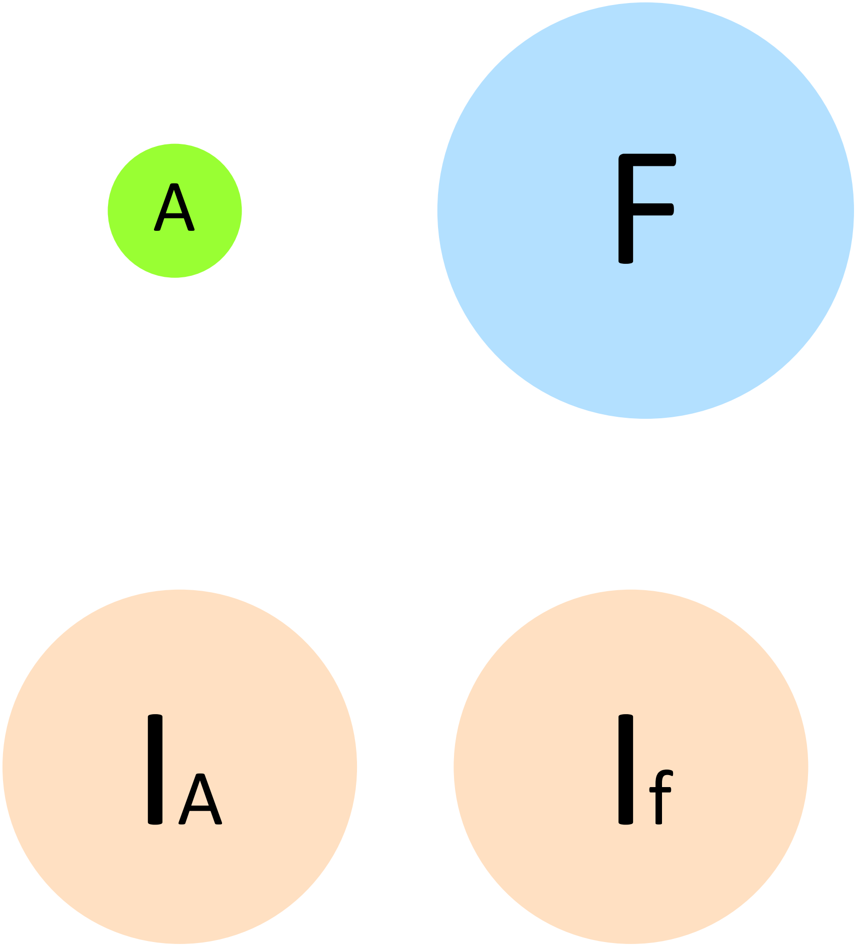
A
F
I_A
I_f

Capa 2

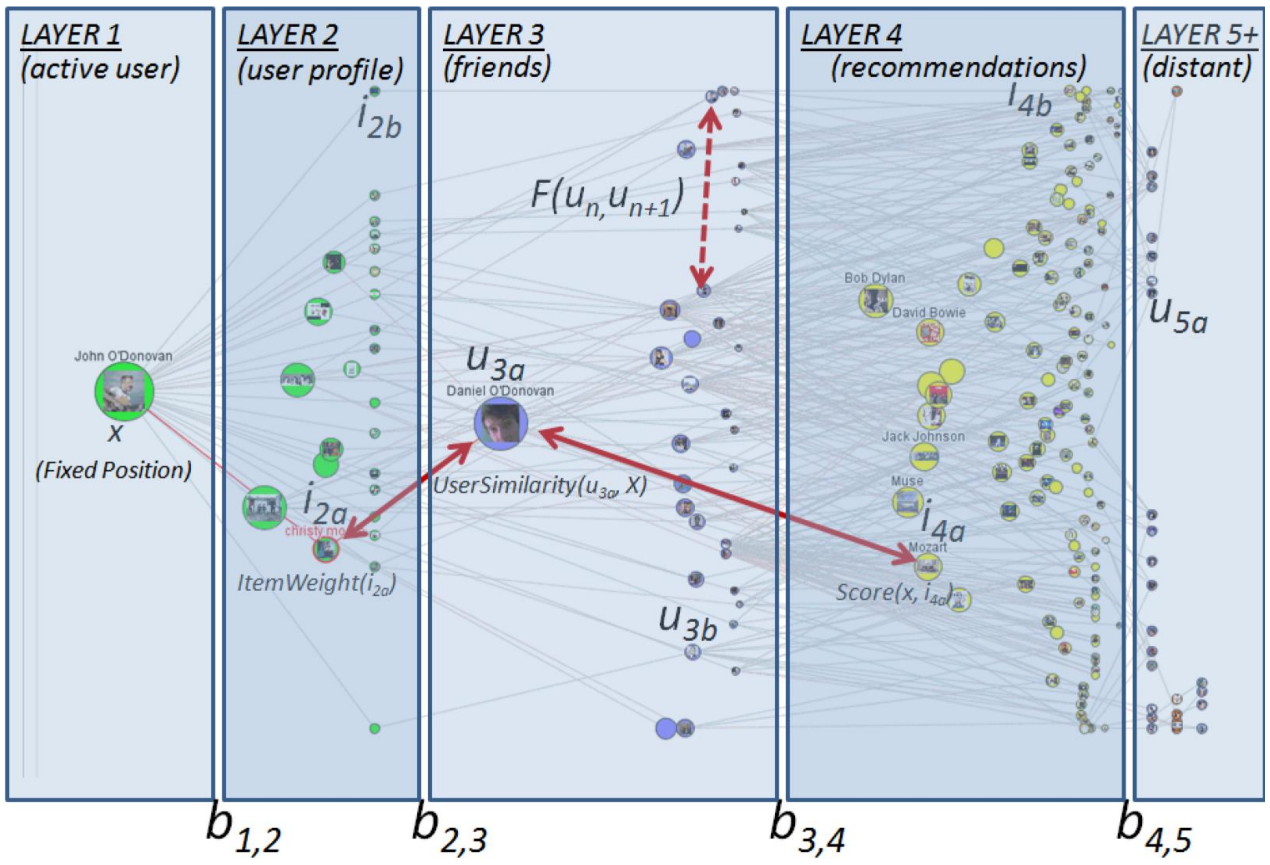

LAYER 1
(active user)
LAYER 2
(user profile)
LAYER 3
(friends)
LAYER 4
(recommendations)
LAYER 5+
(distant)
John O'Donovan
x
(Fixed Position)
i_{2b}
i_{2a}
ItemWeight(i_{2a})
$F(u_n, u_{n+1})$
u_{3a}
Daniel O'Donovan
UserSimilarity(u_{3a}, X)
u_{3b}
i_{4b}
Bob Dylan
David Bowie
Jack Johnson
Muse
i_{4a}
Mozart
Score(x, i_{4a})
u_{5a}
$b_{1,2}$
$b_{2,3}$
$b_{3,4}$
$b_{4,5}$

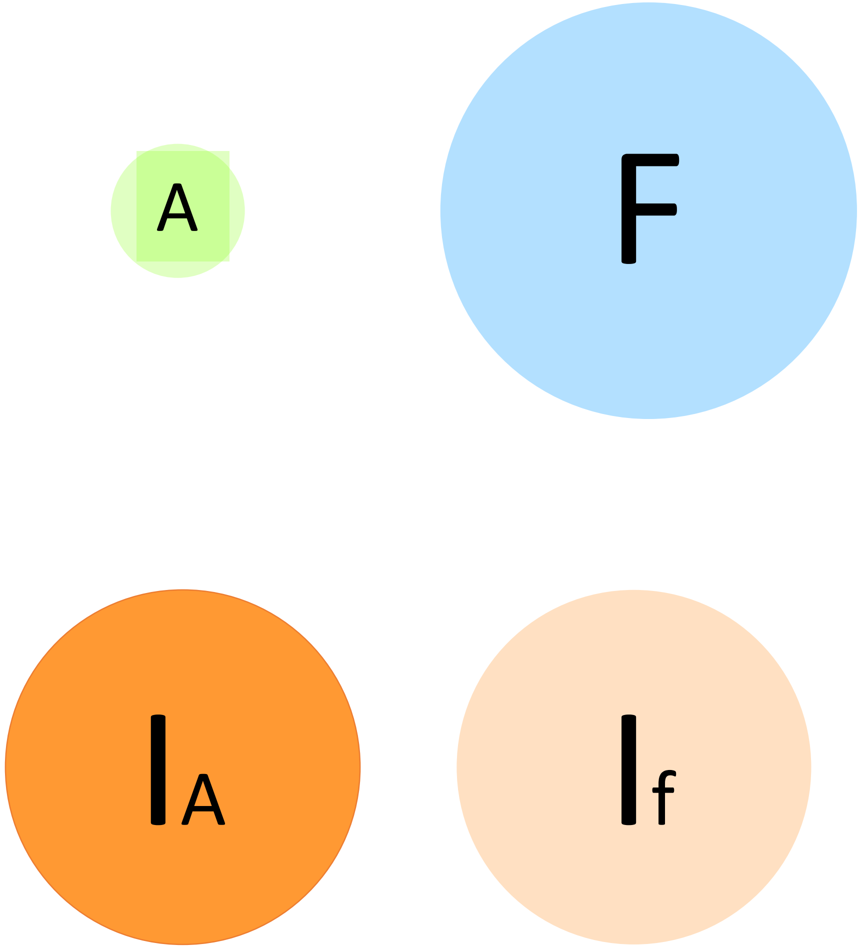

A
F
I_A
I_f

Capa 2

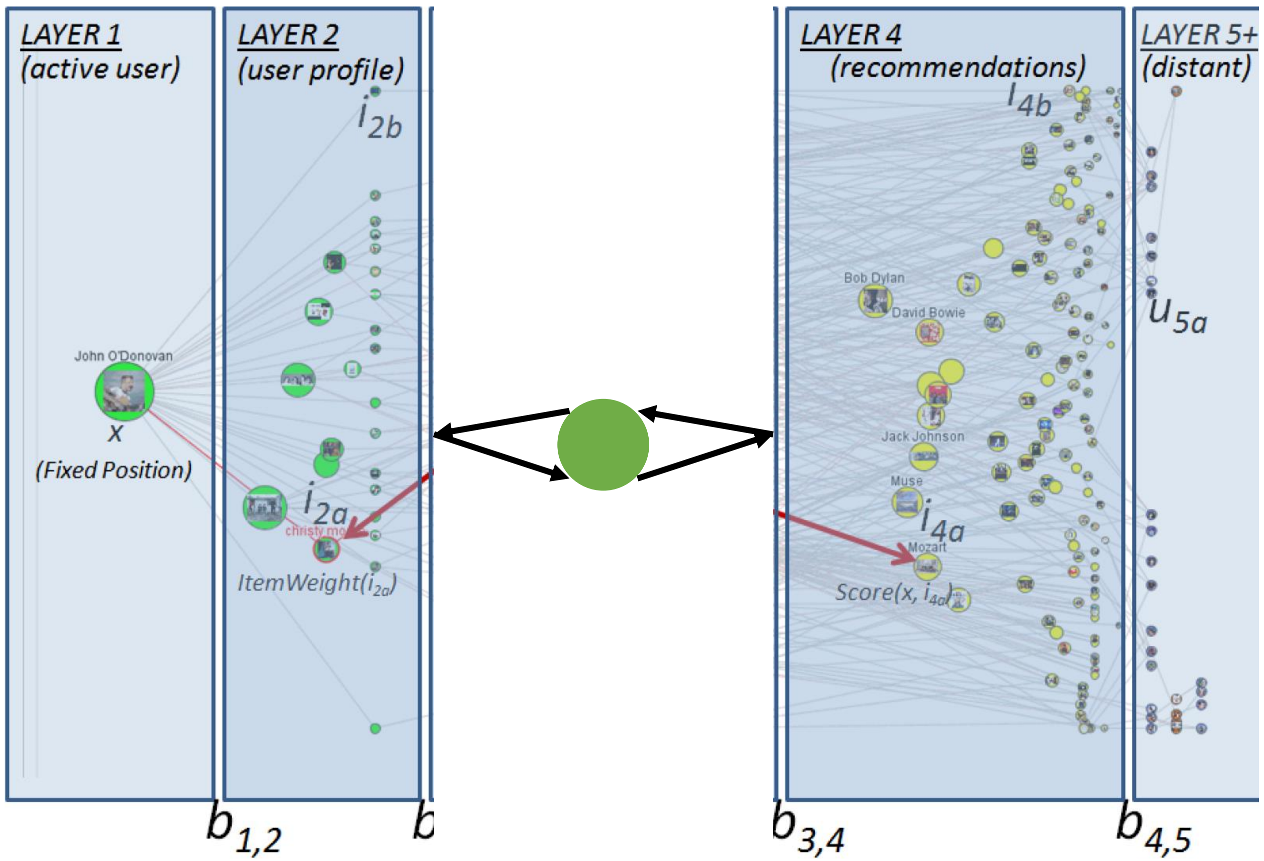

- Items de activo
- Weight – float
- Default: 1
- Cambia el dataset
- Dataset dinámica
- Tiene que ajustar todos los pesos

Capa 3

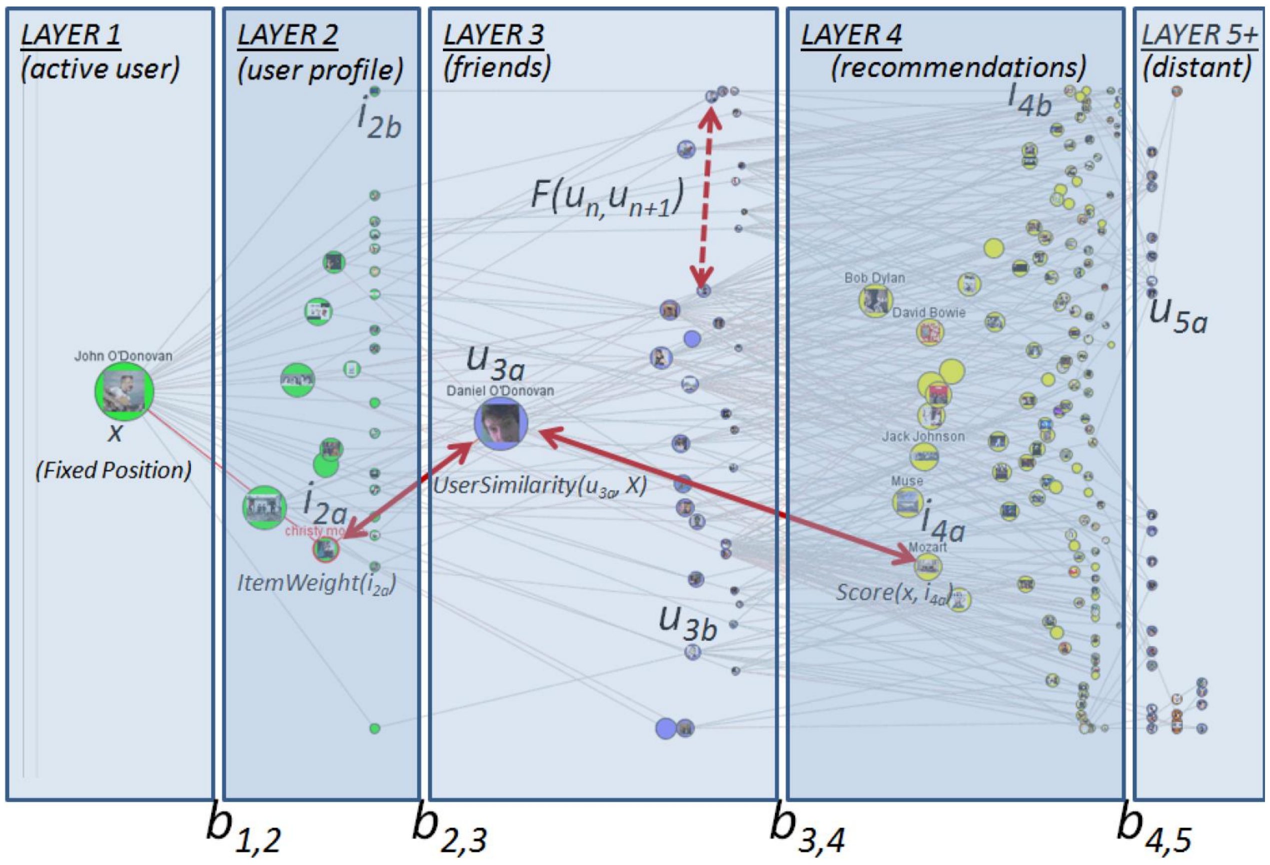
LAYER 1
(active user)
John O'Donovan
X
(Fixed Position)
LAYER 2
(user profile)
i_{2b}
i_{2a}
ItemWeight(i_{2a})
LAYER 3
(friends)
$F(u_n, u_{n+1})$
u_{3a}
Daniel O'Donovan
UserSimilarity(u_{3a}, X)
u_{3b}
LAYER 4
(recommendations)
i_{4b}
Bob Dylan
David Bowie
Jack Johnson
Muse
i_{4a}
Mozart
Score(x, i_{4a})
LAYER 5+
(distant)
u_{5a}
$b_{1,2}$
$b_{2,3}$
$b_{3,4}$
$b_{4,5}$

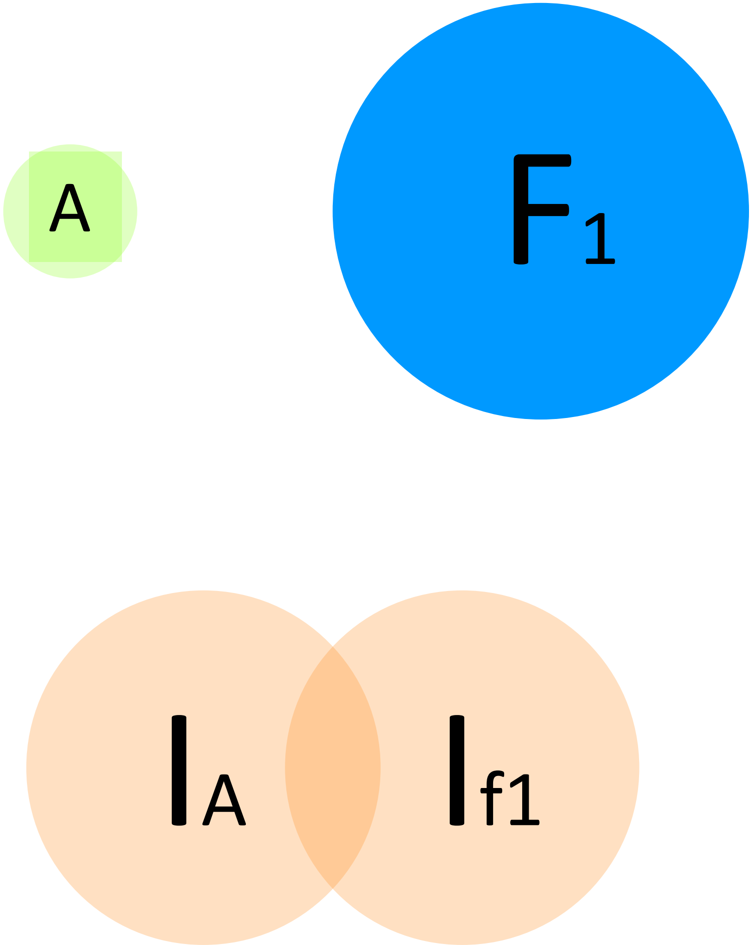
A
F_1
I_A
I_{f1}

Capa 3

$$UserSimilarity(x,u) = \frac{UserWeight(u) \cdot TotalWeightOfCommonItems(x,u)}{\sqrt{TotalWeightOfItems(x) \cdot TotalWeightOfItems(u)}}$$

- Basado en items comunes
- Basado en UserWeight

Capa 3

$$TotalWeightOfCommonItems(x,u) = \sum_{i \in I} (Likes(x,i) \cdot Likes(u,i) \cdot ItemWeight(i))$$

$$Likes(u,i) \begin{cases} 1 \\ 0 \end{cases}$$

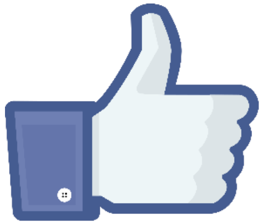

- Basado en items comunes

Capa 3

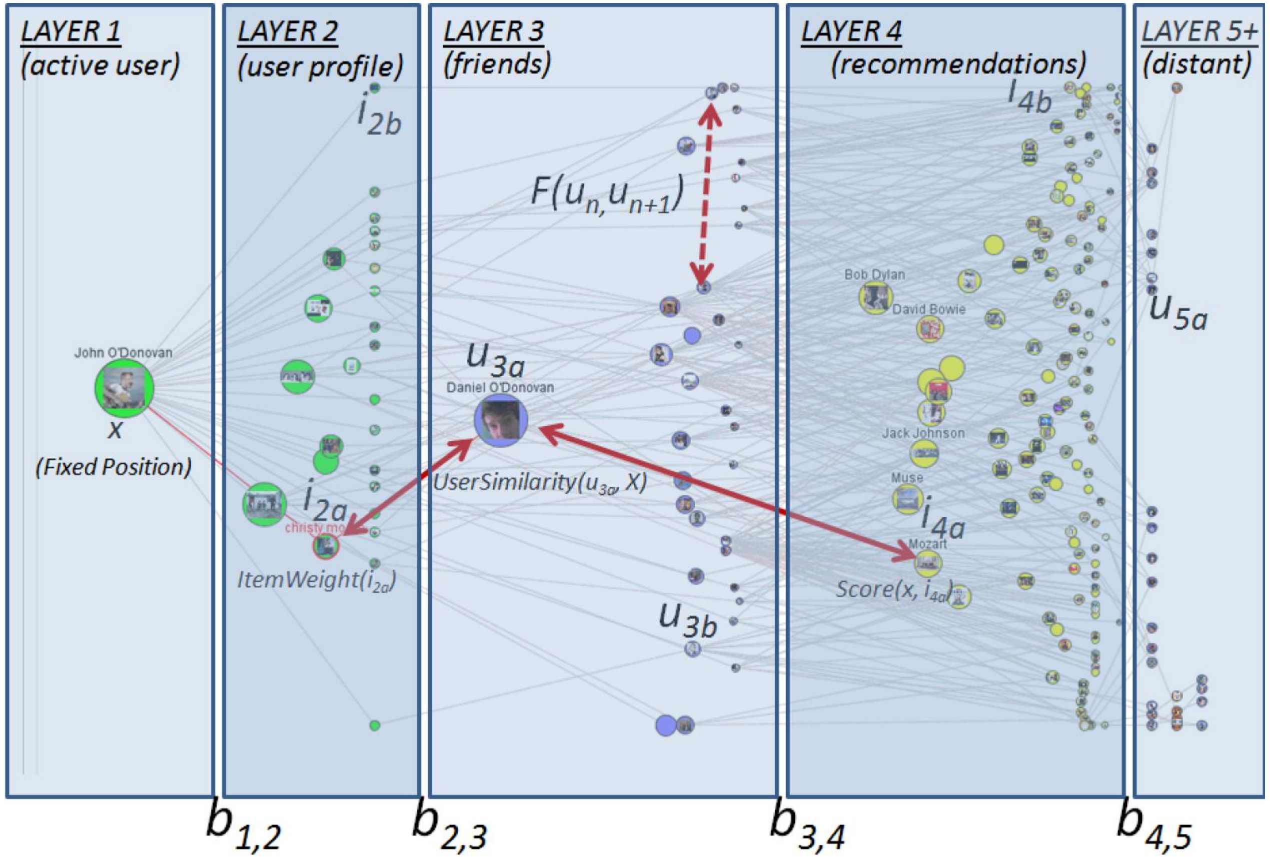

- Basado en UserWeights
- “tweaking” en tiempo real
- 0 <= UserWeight <= inf.
- Default: 1

Capa 3

- Normalizar

$$TotalWeightOfItems(u) = \sum_{i \in I} (Likes(u,i) \cdot ItemWeight(i))$$

$$UserSimilarity(x,u) = \frac{UserWeight(u) \cdot TotalWeightOfCommonItems(x,u)}{\sqrt{TotalWeightOfItems(x) \cdot TotalWeightOfItems(u)}}$$

Capa 3

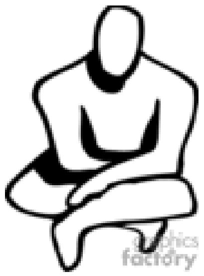

UserWeight	30
Common items	5
Total items (u)	5
Total items (x)	5
UserSimilarity	30

$$UserSimilarity(x,u) = \frac{UserWeight(u) \cdot TotalWeightOfCommonItems(x,u)}{\sqrt{TotalWeightOfItems(x) \cdot TotalWeightOfItems(u)}}$$

- UserWeight muy alto
- Afuera del rango [0, 1]
- Limite superior

$$BoundedUserSimilarity(x,u) = \min(1, UserSimilarity(x,u))$$

Capa 4

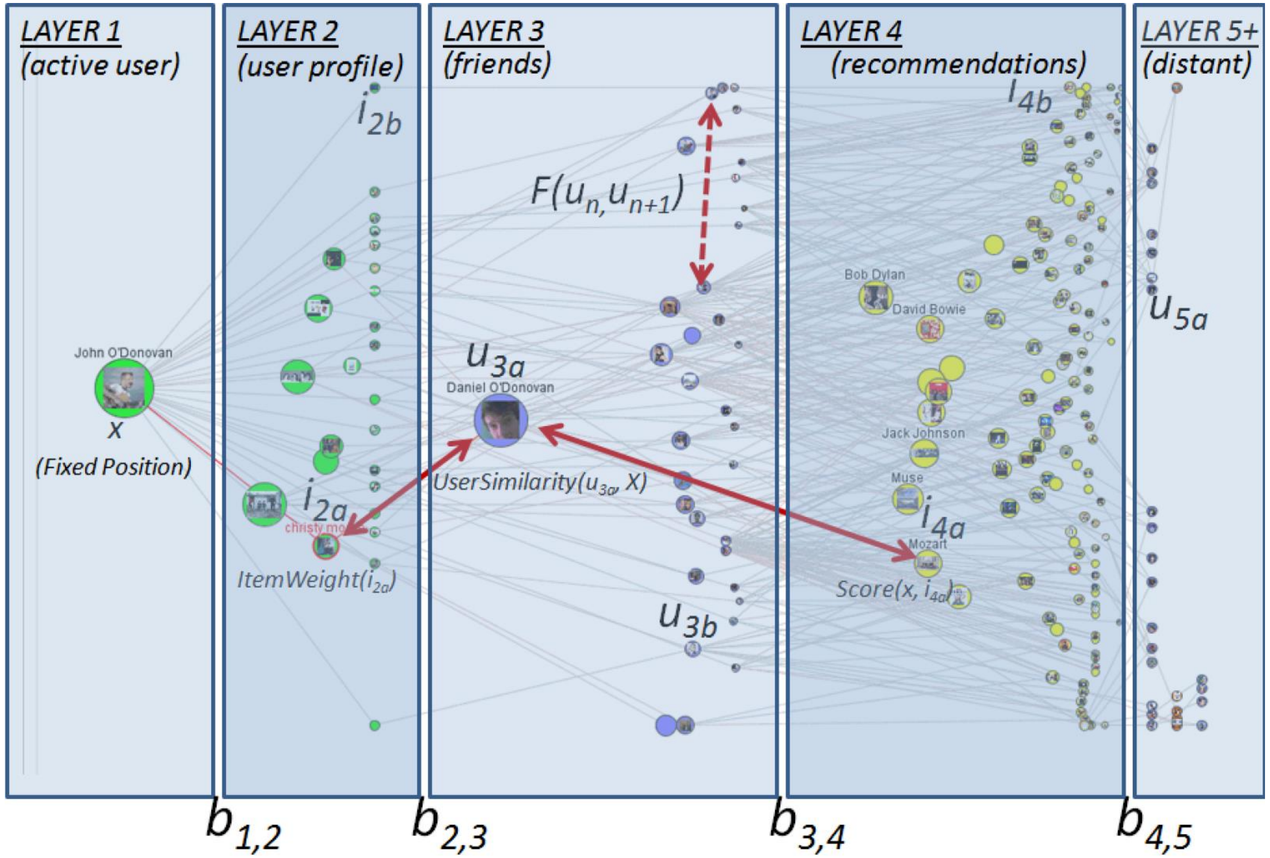
LAYER 1
(active user)
LAYER 2
(user profile)
LAYER 3
(friends)
LAYER 4
(recommendations)
LAYER 5+
(distant)
John O'Donovan
x
(Fixed Position)
i2b
i2a
ItemWeight(i2a)
F(un,un+1)
u3a
Daniel O'Donovan
UserSimilarity(u3a, X)
u3b
i4b
Bob Dylan
David Bowie
Jack Johnson
Muse
i4a
Mozart
Score(x, i4a)
u5a
b1,2
b2,3
b3,4
b4,5

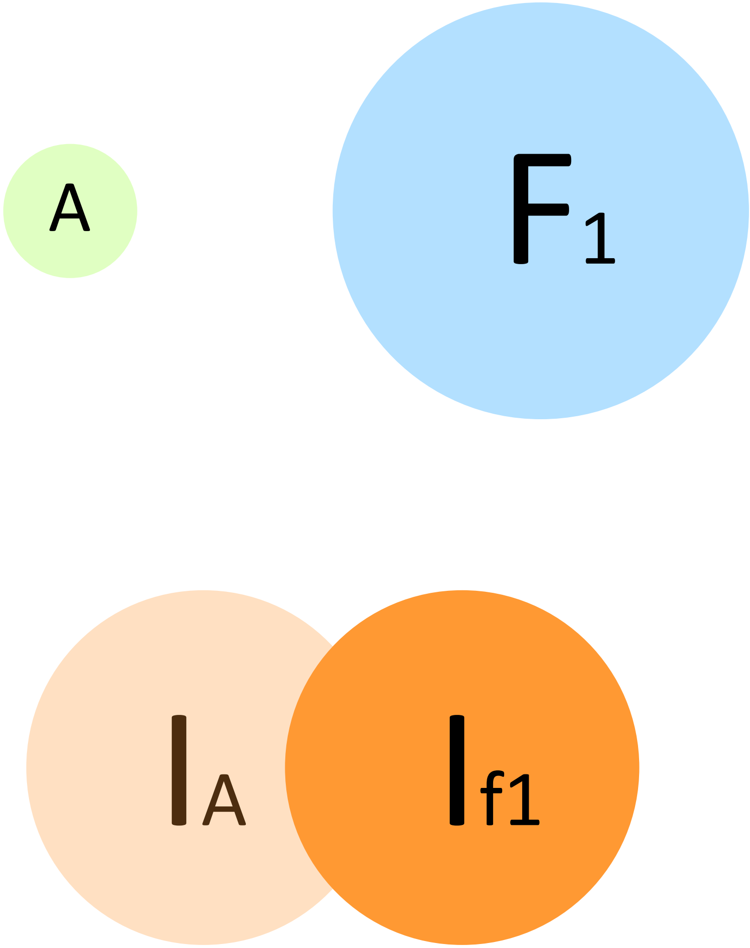
A
F1
IA
If1

Capa 4

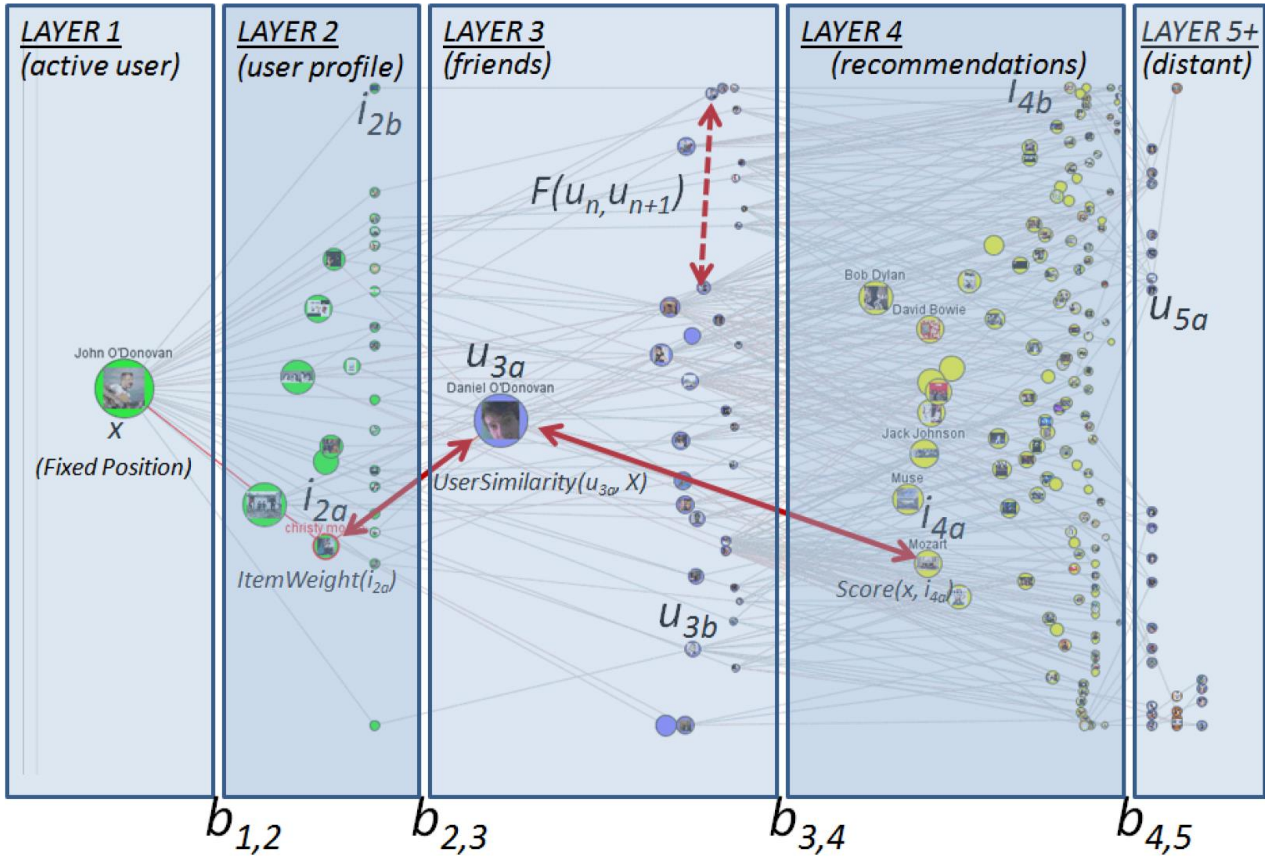

- Recomendaciones
- Conexiones sociales – filtrado
- Score
  - Aristas
  - BoundedUserSimilarity

Capa 4

- Aristas: Likes
- BoundedUserSimilarity

$$Score(x,i) =$$
$$\sum_{u \in U} (Likes(u,i) \cdot BoundedUserSimilarity(x,u))$$

Capa 4

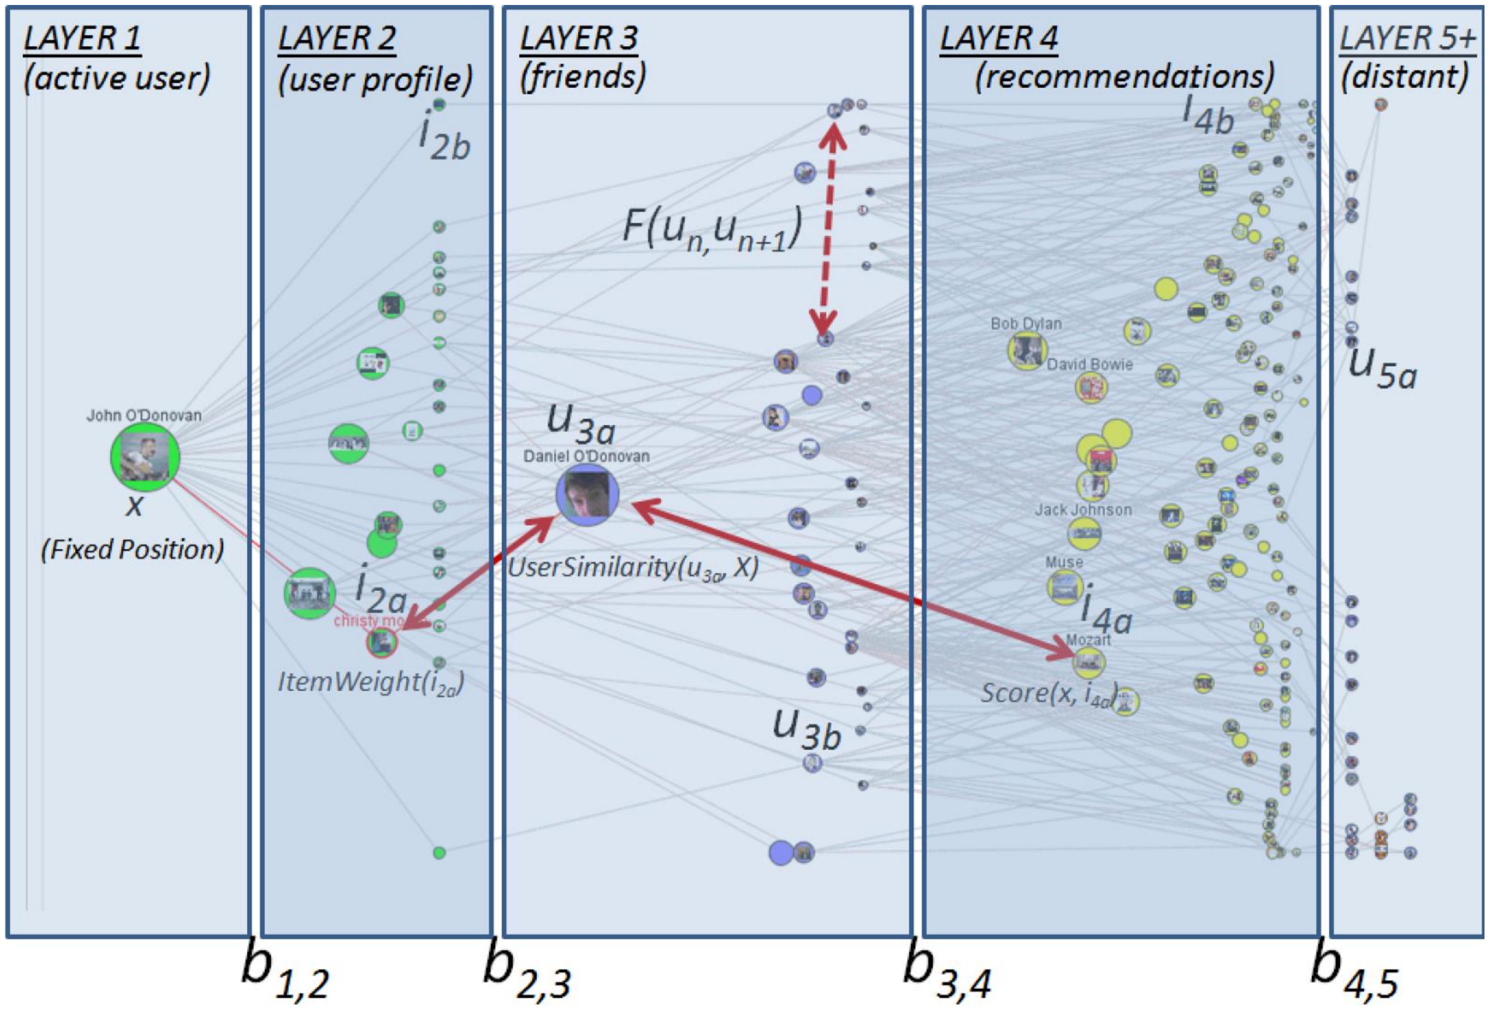

LAYER 1
(active user)
LAYER 2
(user profile)
LAYER 3
(friends)
LAYER 4
(recommendations)
LAYER 5+
(distant)
John O'Donovan
X
(Fixed Position)
i_{2b}
i_{2a}
ItemWeight(i_{2a})
$F(u_n, u_{n+1})$
u_{3a}
Daniel O'Donovan
UserSimilarity(u_{3a}, X)
u_{3b}
i_{4b}
Bob Dylan
David Bowie
Jack Johnson
Muse
i_{4a}
Mozart
Score(x, i_{4a})
u_{5a}
$b_{1,2}$
$b_{2,3}$
$b_{3,4}$
$b_{4,5}$

Evaluación

- User study – tasks
  - Facebook (Benchmark)
  - Interfaz arbol
  - Interfaz circulo
- User study – evaluar recomendaciones
  - SmallWorlds
  - MovieLens 100K dataset con CF (MovieLens) [MAL*03]
- Automated accuracy test – leave-one-out cross validation
  - SmallWorlds
  - MovieLens 1M (10K items) dataset con MovieLens QuickPick [Gro09]

User study - results

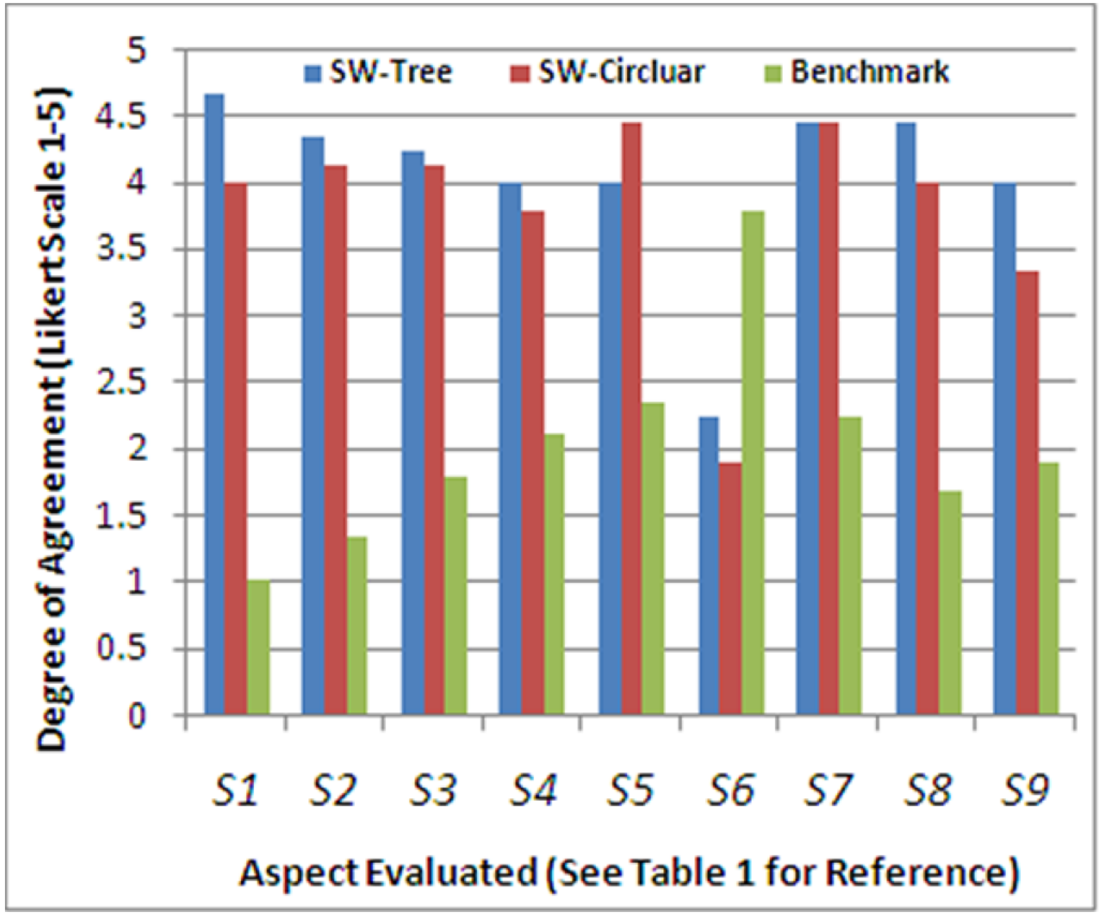

#	*QuestionDescription*
S1	effective for finding commonalities in taste
S2	item popularity easily discoverable
S3	interesting items easily discoverable
S4	was easy to use
S5	was intuitive overall
S6	was clumsy overall
S7	was informative overall
S8	helped you to explore the given topic
S9	helped you to build your movie profile

- S5
- Sense of control ✓

User study – other results

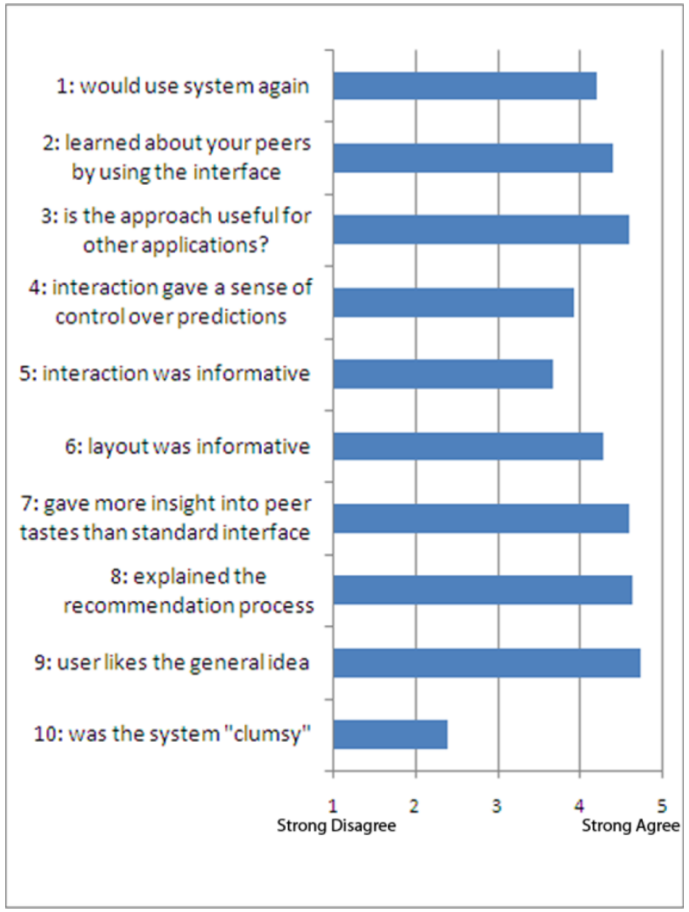

Method	MovieLens	SW-Tree (interactive)	SW-Tree
Satisfaction	4.25	4.19	3.78

Table 2: Satisfaction ratings of item predictions for Movielens and for SmallWorlds with and without user interactions.

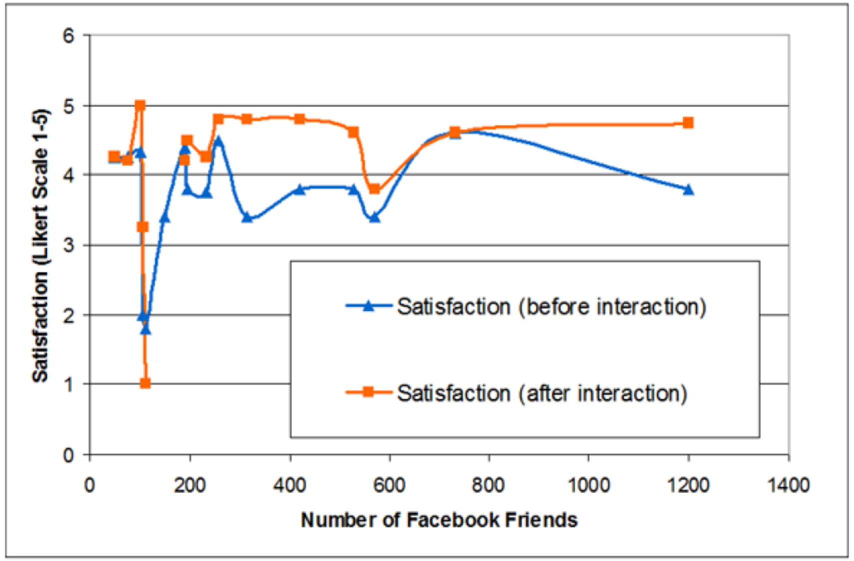

Figure 6: Relation between satisfaction with recommendations and number of Facebook friends.

Conclusion & future work

	Approaches					Metrics							Data collection methods						Results						
	Comparison with baseline without recommendations	Comparison with baseline without user control or visual explanation	Comparing different visualizations	Comparing different recommender algorithms	Asking users to explore freely	Effectiveness	Efficiency	Engagement	Satisfaction	Trust	Usability	Usefulness	Recommendation accuracy test	Task performance analysis	User behavior	Questionnaire	Interview	Think-aloud	Increase of acceptance	Better task performance	Increase of efficiency	Increase of engagement	Increase of satisfaction	Increase of trust	Positive usefulness / usability feedback
SmallWorlds (Figure 4)	+		+			+			+		+	+	+			+			+				+		+

- User satisfaction ✓
- Transparency ✓
- Sense of control ✓

Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities

Referencias

- Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C., & Höllerer, T. (2010, June). Smallworlds: visualizing social recommendations. In *Computer Graphics Forum* (Vol. 29, No. 3, pp. 833-842). Blackwell Publishing Ltd.
- [MAL*03] Miller, B. N., Albert, I., Lam, S. K., Konstan, J. A., & Riedl, J. (2003, January). MovieLens unplugged: experiences with an occasionally connected recommender system. In *Proceedings of the 8th international conference on Intelligent user interfaces* (pp. 263-266). ACM.
- [Gro09] GROUPLENS: Movielens quickpick recommender system. http://www.movielens.org/quickpick, 2009.
- He, C., Parra, D., & Verbert, K. (2016). Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications*, *56*, 9-27.

Demo video

- Por el cambio de Facebook API, ya no funciona SmallWorlds
- https://vimeo.com/21060974

Web based architecture

- Procesamiento en el servidor
- Captura los movimientos del raton

User study – 7 tasks

1. *Task 1:* Familiarization (5 mins, supervised)
2. *Task 2:* Find popular items in your peer-group.
3. *Task 3:* Find your 3 most similar peers
4. *Task 4:* Find your 3 least similar peers
5. *Task 5:* Get recommendations through layout only
6. *Task 6:* Get recommendations through layout and interaction
7. *Task 7:* Get recommendations through layout and interaction, with layer 4 (candidate-set) items hidden.

- 17 participantes
- 50 – 1200 amigos (mediano 215)

Task 1, Facebook	Task 1, Tree-like	Task 1, Circular
Task i, Facebook	Task i, Tree-like	Task i, Circular

User satisfaction evaluation

1. For each participant, a list of recommendations are generated by SW or MovieLens.
2. The participants then rate these items on a 5 point rating scale.
3. The ratings are subtracted from 5 (assumed ground truth) as the MAE.

Automated accuracy test

Leave one out analysis

1. Train using n-1 items in the profile (Layer 2)
2. Aim to have the left out item in the top 12 recommendations of SW and MovieLens
3. Record for each user the number of times the left out item is included in the list (iteration depends on n), e.g. User X has 10 items, left out item is recommended 2 times in the 10 iterations.
4. Compute the average per user per system: 1.94 items (SW) 0.82 items (MovieLens)
5. Top 5: 1.00 (SW) 0.65 (MovieLens)

Winner-loser analysis

1. The same as above but compare between the two, the number of times each system had the removed item ranked higher than the other.
2. 1.88 items SW ranked higher than MovieLens, 0.65 items MovieLens ranked higher.

Note: MovieLens have much more data.