Decision Theory and Network Science: Methods and Applications,

STOR-i Workshop at Lancaster University, Sep 18th, 2017

ACTIVE LEARNING ON NETWORKS AND ONLINE INFLUENCE MAXIMIZATION

Michal Valko, SequeL, Inria Lille - Nord Europe


2015-2016, AISTAST 2016




2016-2017, NIPS 2017

https://arxiv.org/abs/1605.06593

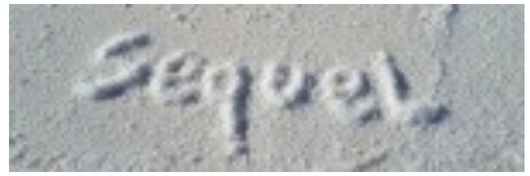

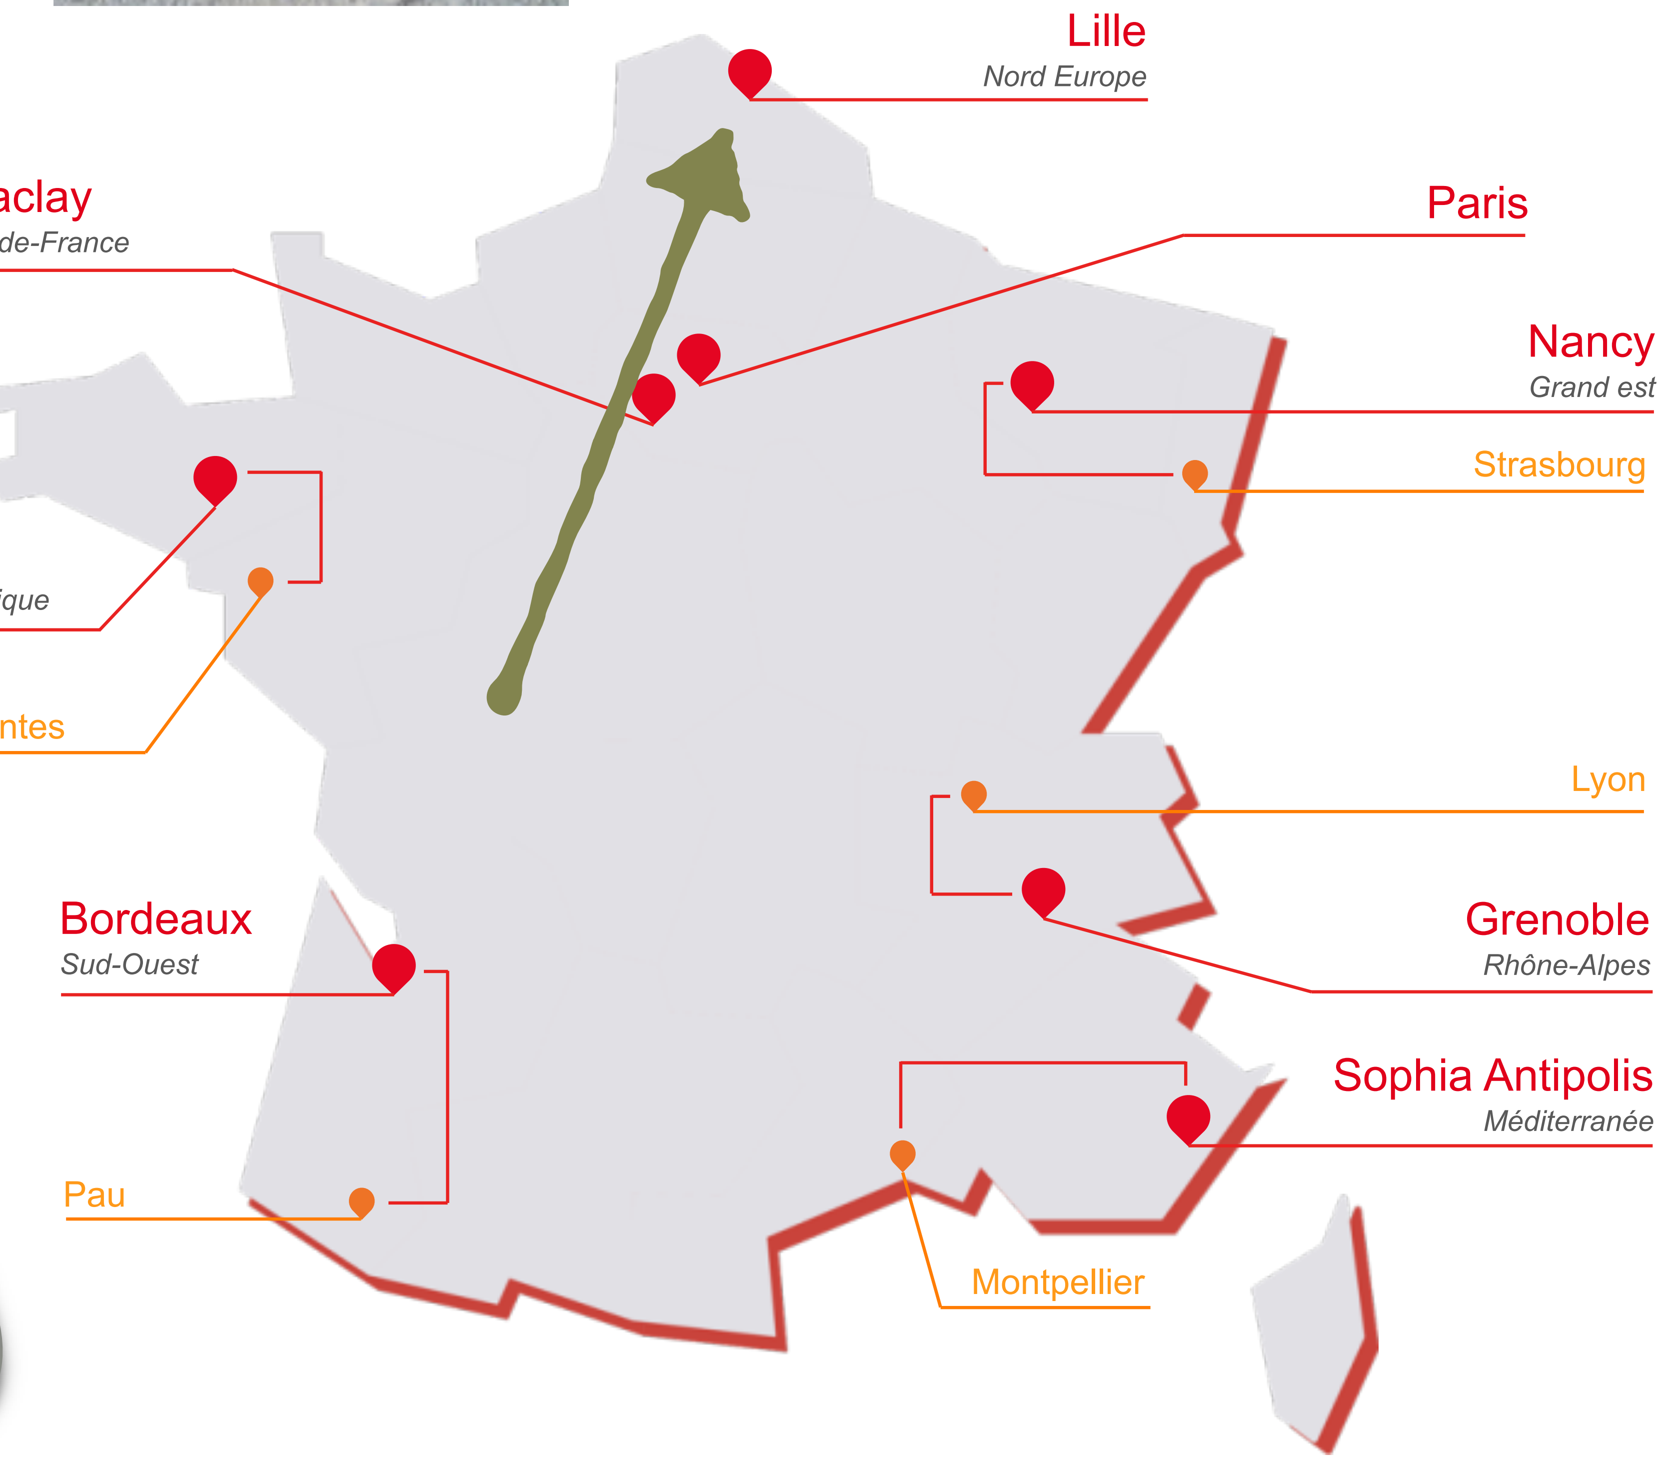

+

Philippe Preux
SequeL, Inria

Rémi Munos
Google DeepMind

10 YEARS

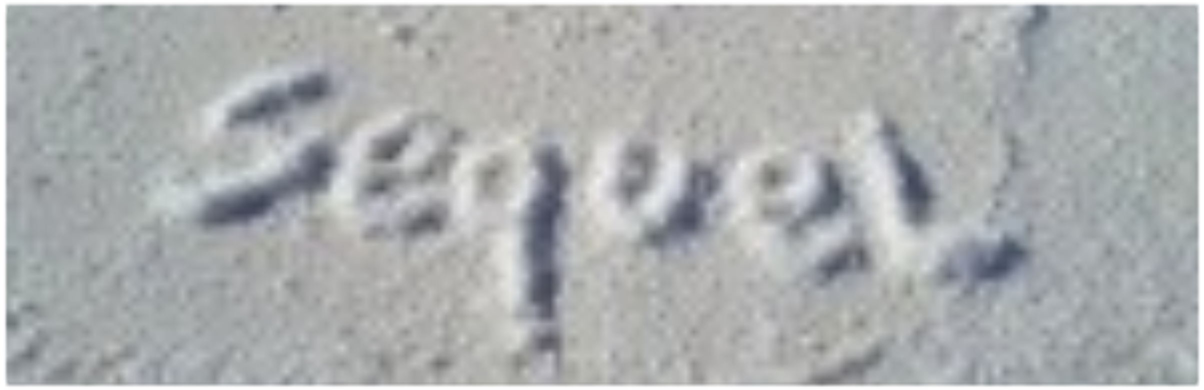

… LAST 10 YEARS AND INDUSTRY

Inria informatics mathematics

criteo

J. Mary

M. Davy

vekia THE HEART OF RETAIL

Adobe

M. Ghavamzadeh

Sequel

R.Munos
O. Pietquin
B. Piot
G. Dulac-Arnold
A. Huang
M. Azar
JB. Grill

A. Lazaric

R. Coulom
CrazyStone

Facebook AI Research PARIS

UNBALANCE

DeepMind

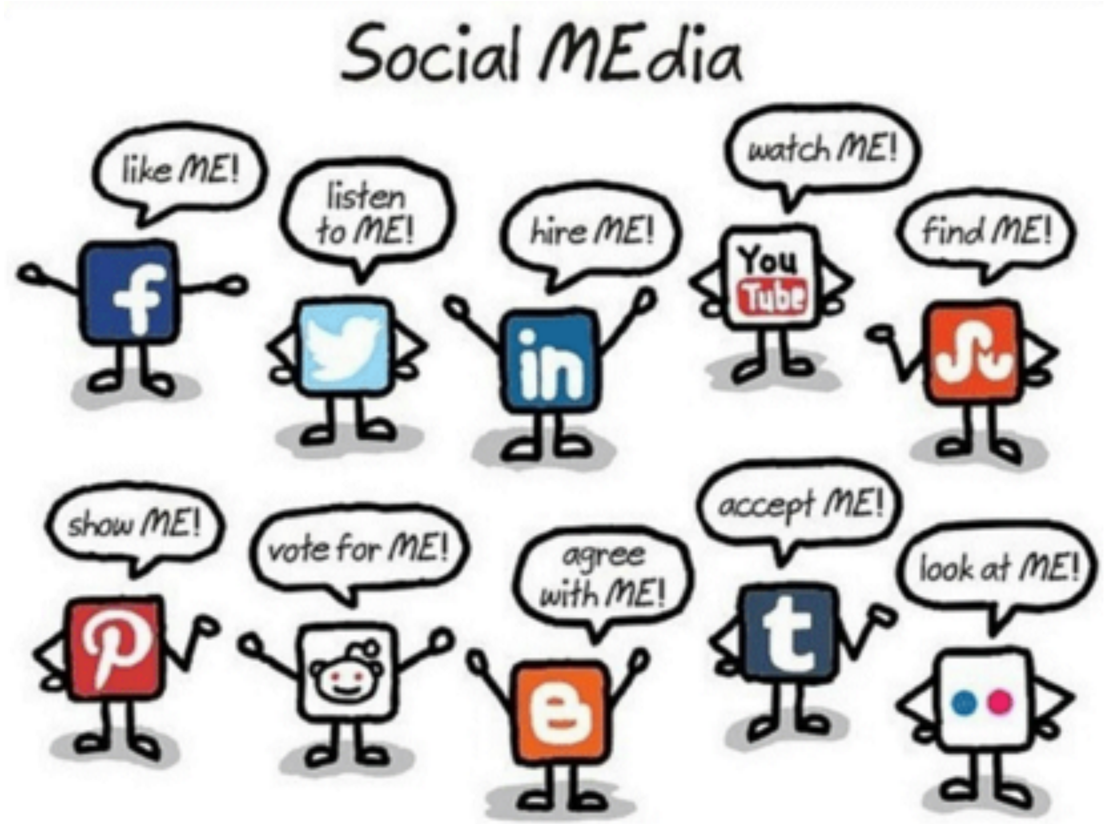

online social networks

Berkeley's floating sensor network

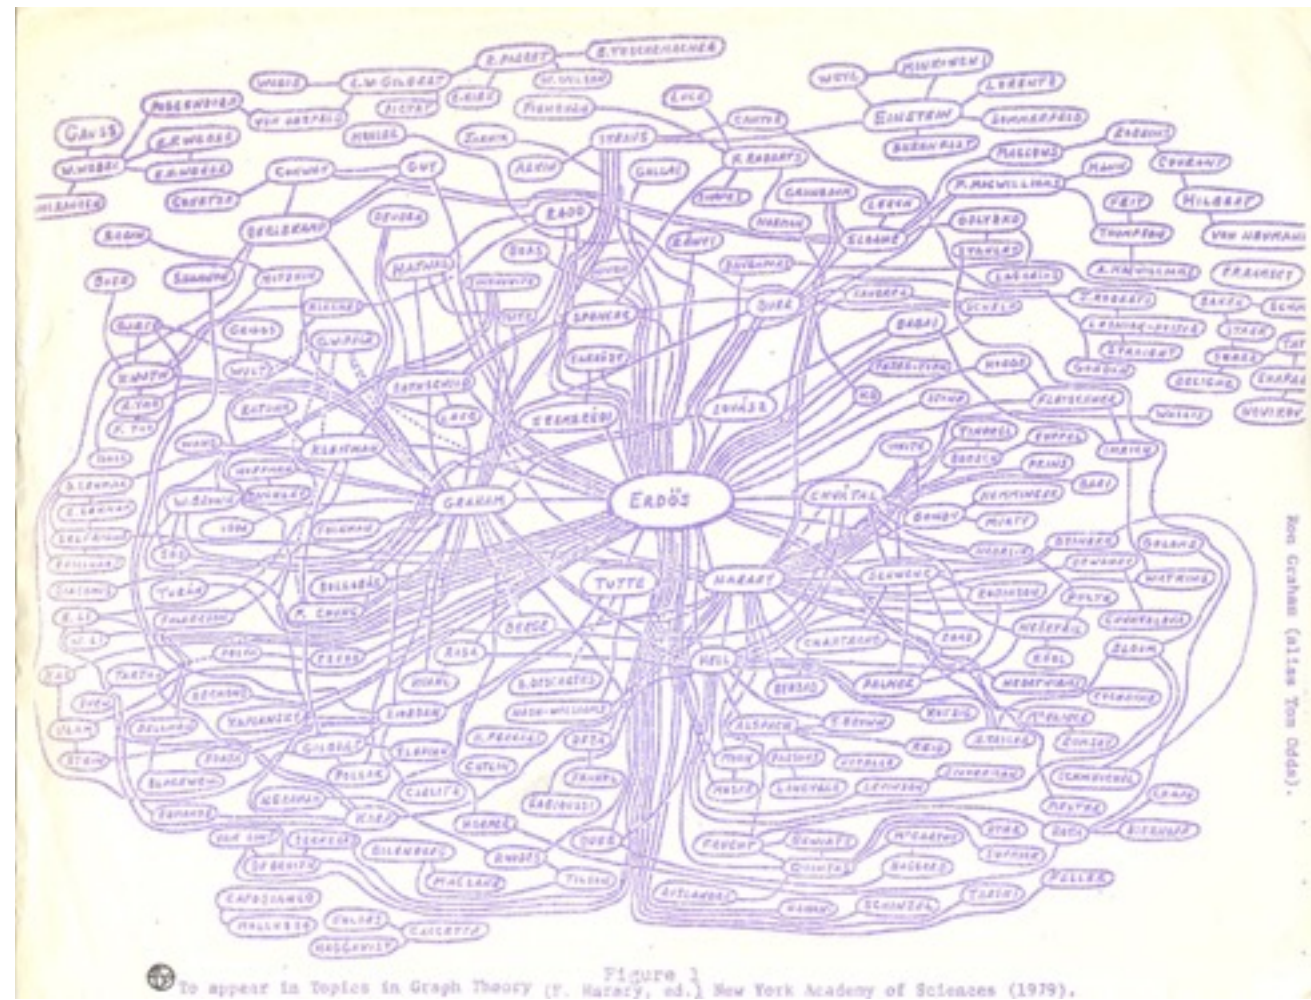

Erdös number project


KUNG FU PANDA
Harry Potter
THOR
QUENTIN TARANTINO
DJANGO
From the Creators of TOY STORY
MONSTERS, INC.
In Cinemas Soon
FROZEN
LORD OF THE RINGS
SIN CITY
MUMMY
NICOLAS CAGE
MEG RYAN
CITY OF ANGELS
NICHOLSON
PFEIFFER
WOLF
BRIDGET JONES'S DIARY
ALL MEN
MUST DIE
GOT
the BIG BANG theory
CBS
the uninvited
HEROES
HANNIBAL
RISING

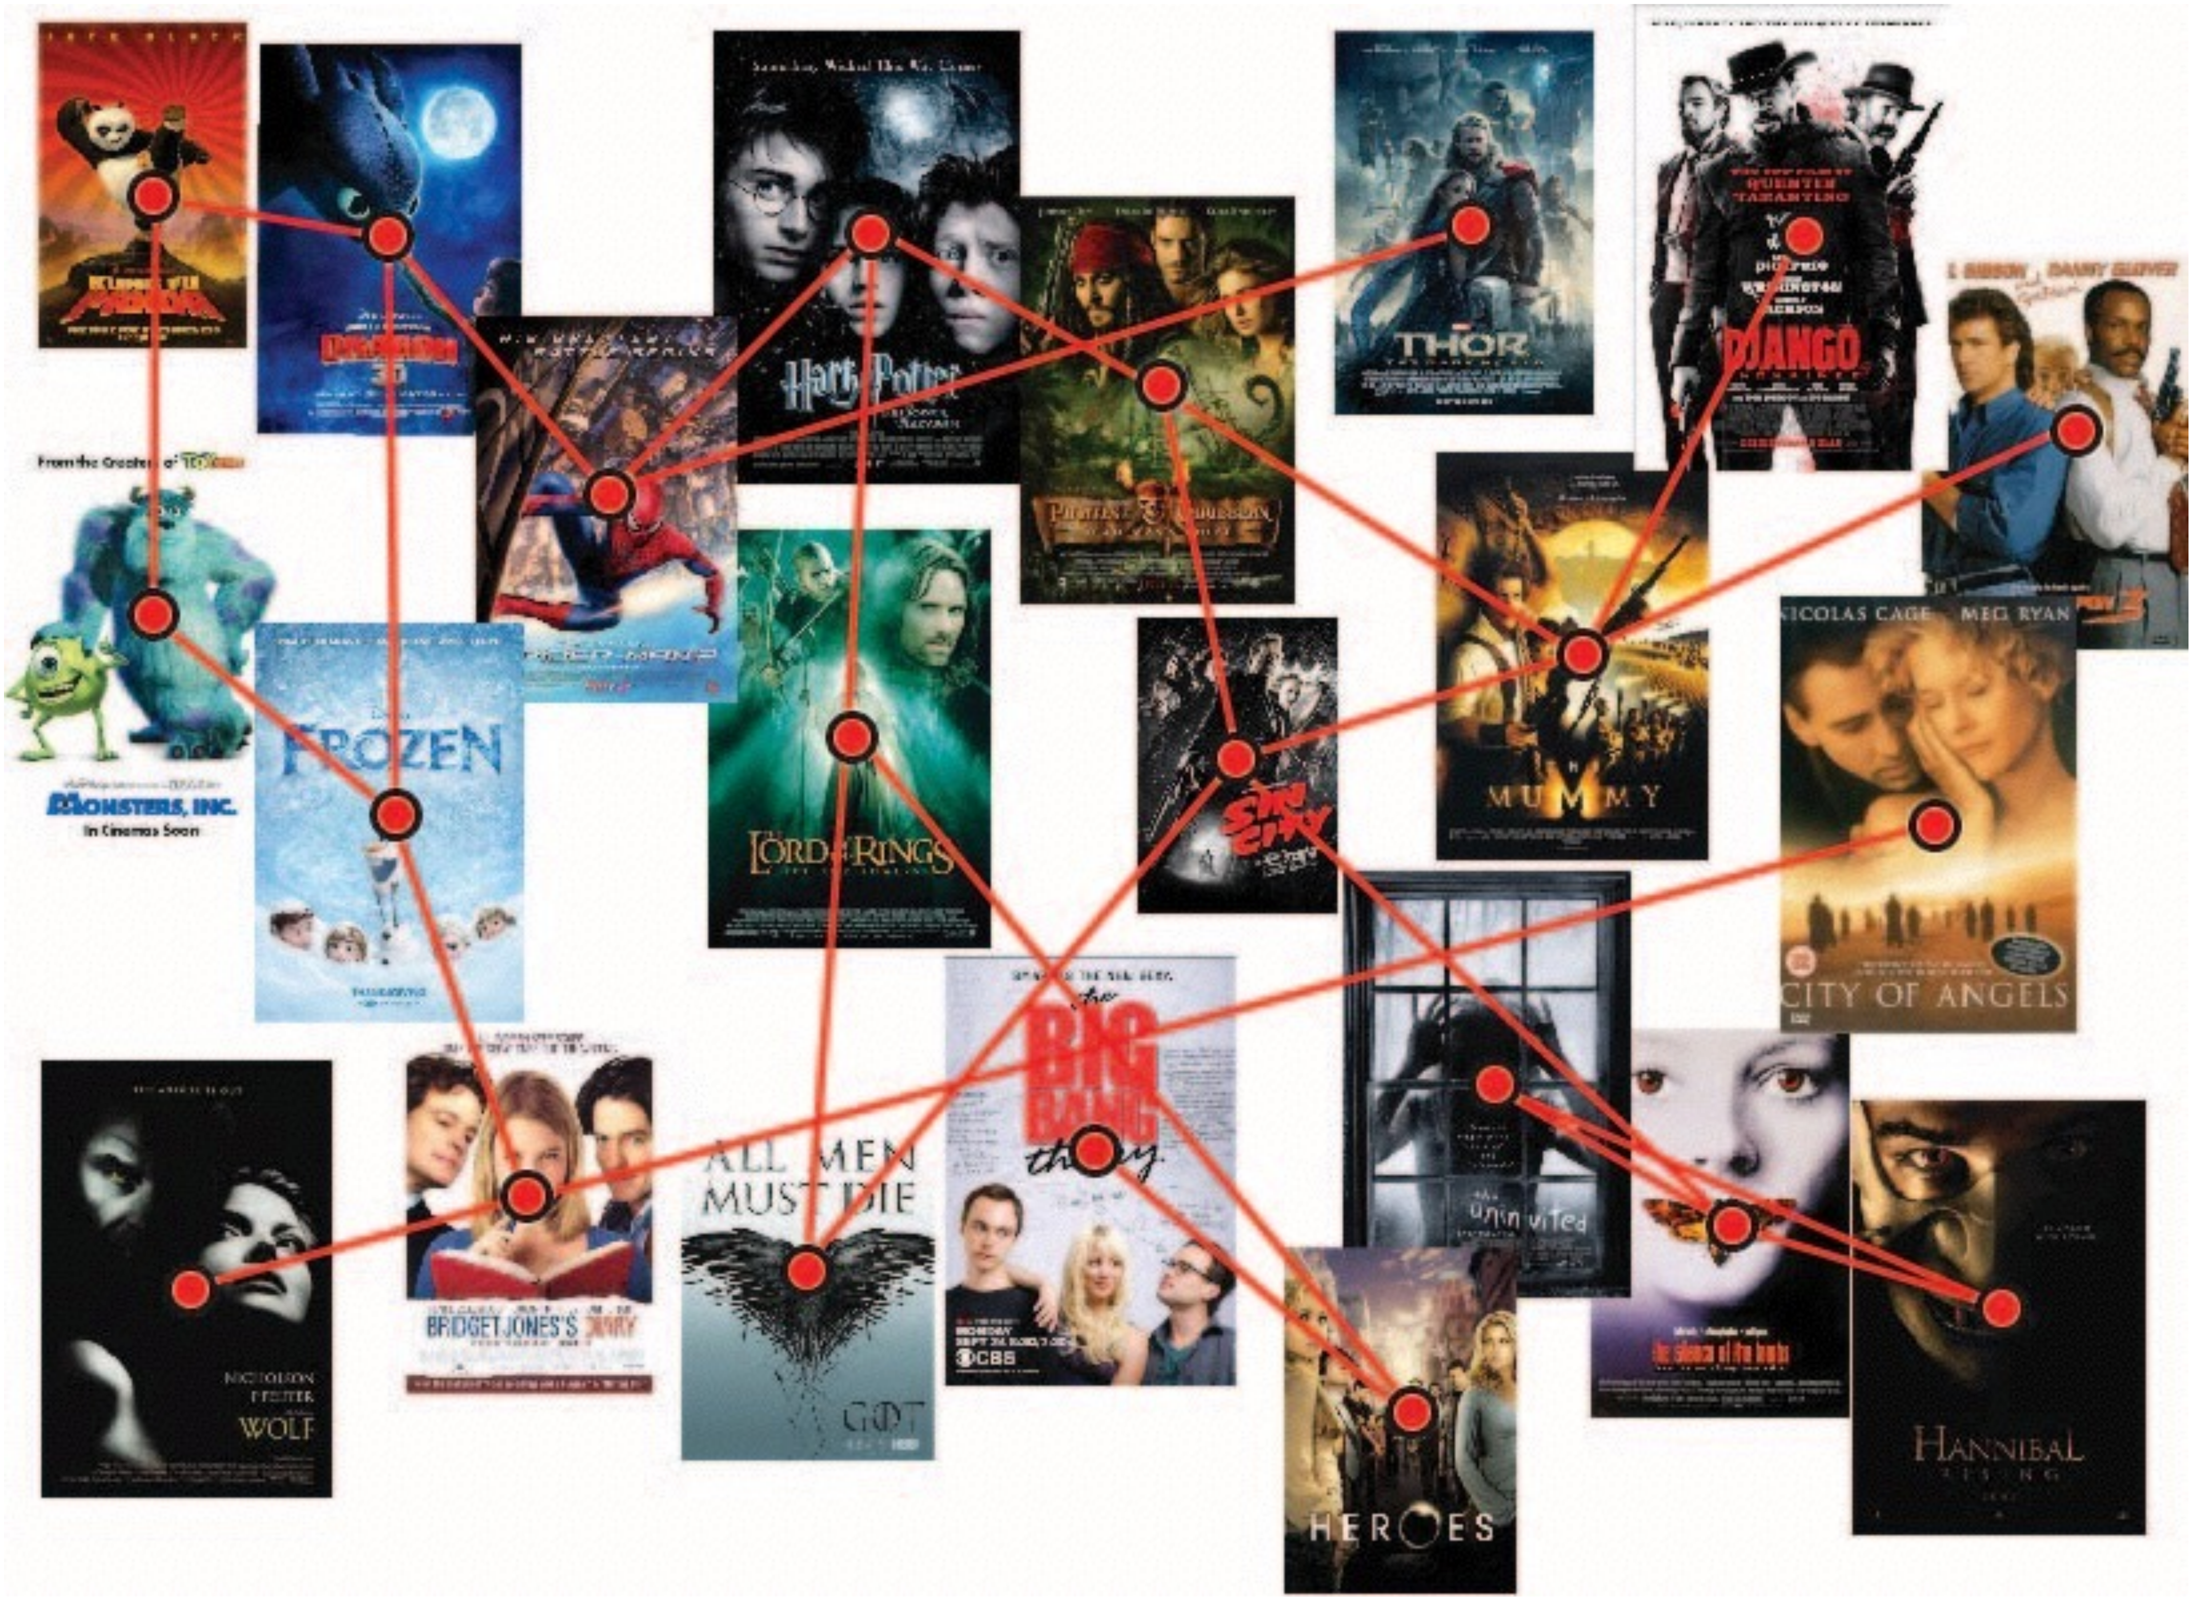

KUNG FU PANDA
MONSTERS, INC.
In Cinemas Soon
FROZEN
Harry Potter
THOR
DJANGO
QUENTIN TARANTINO
MUMMY
CITY OF ANGELS
MEG RYAN
THE LORD OF THE RINGS
ALL MEN MUST DIE
GOT
BRIDGET JONES'S DIARY
NICHOLSON
PFEIFFER
WOLF
the BIG BANG theory
the uninvited
HEROES
HANNIBAL

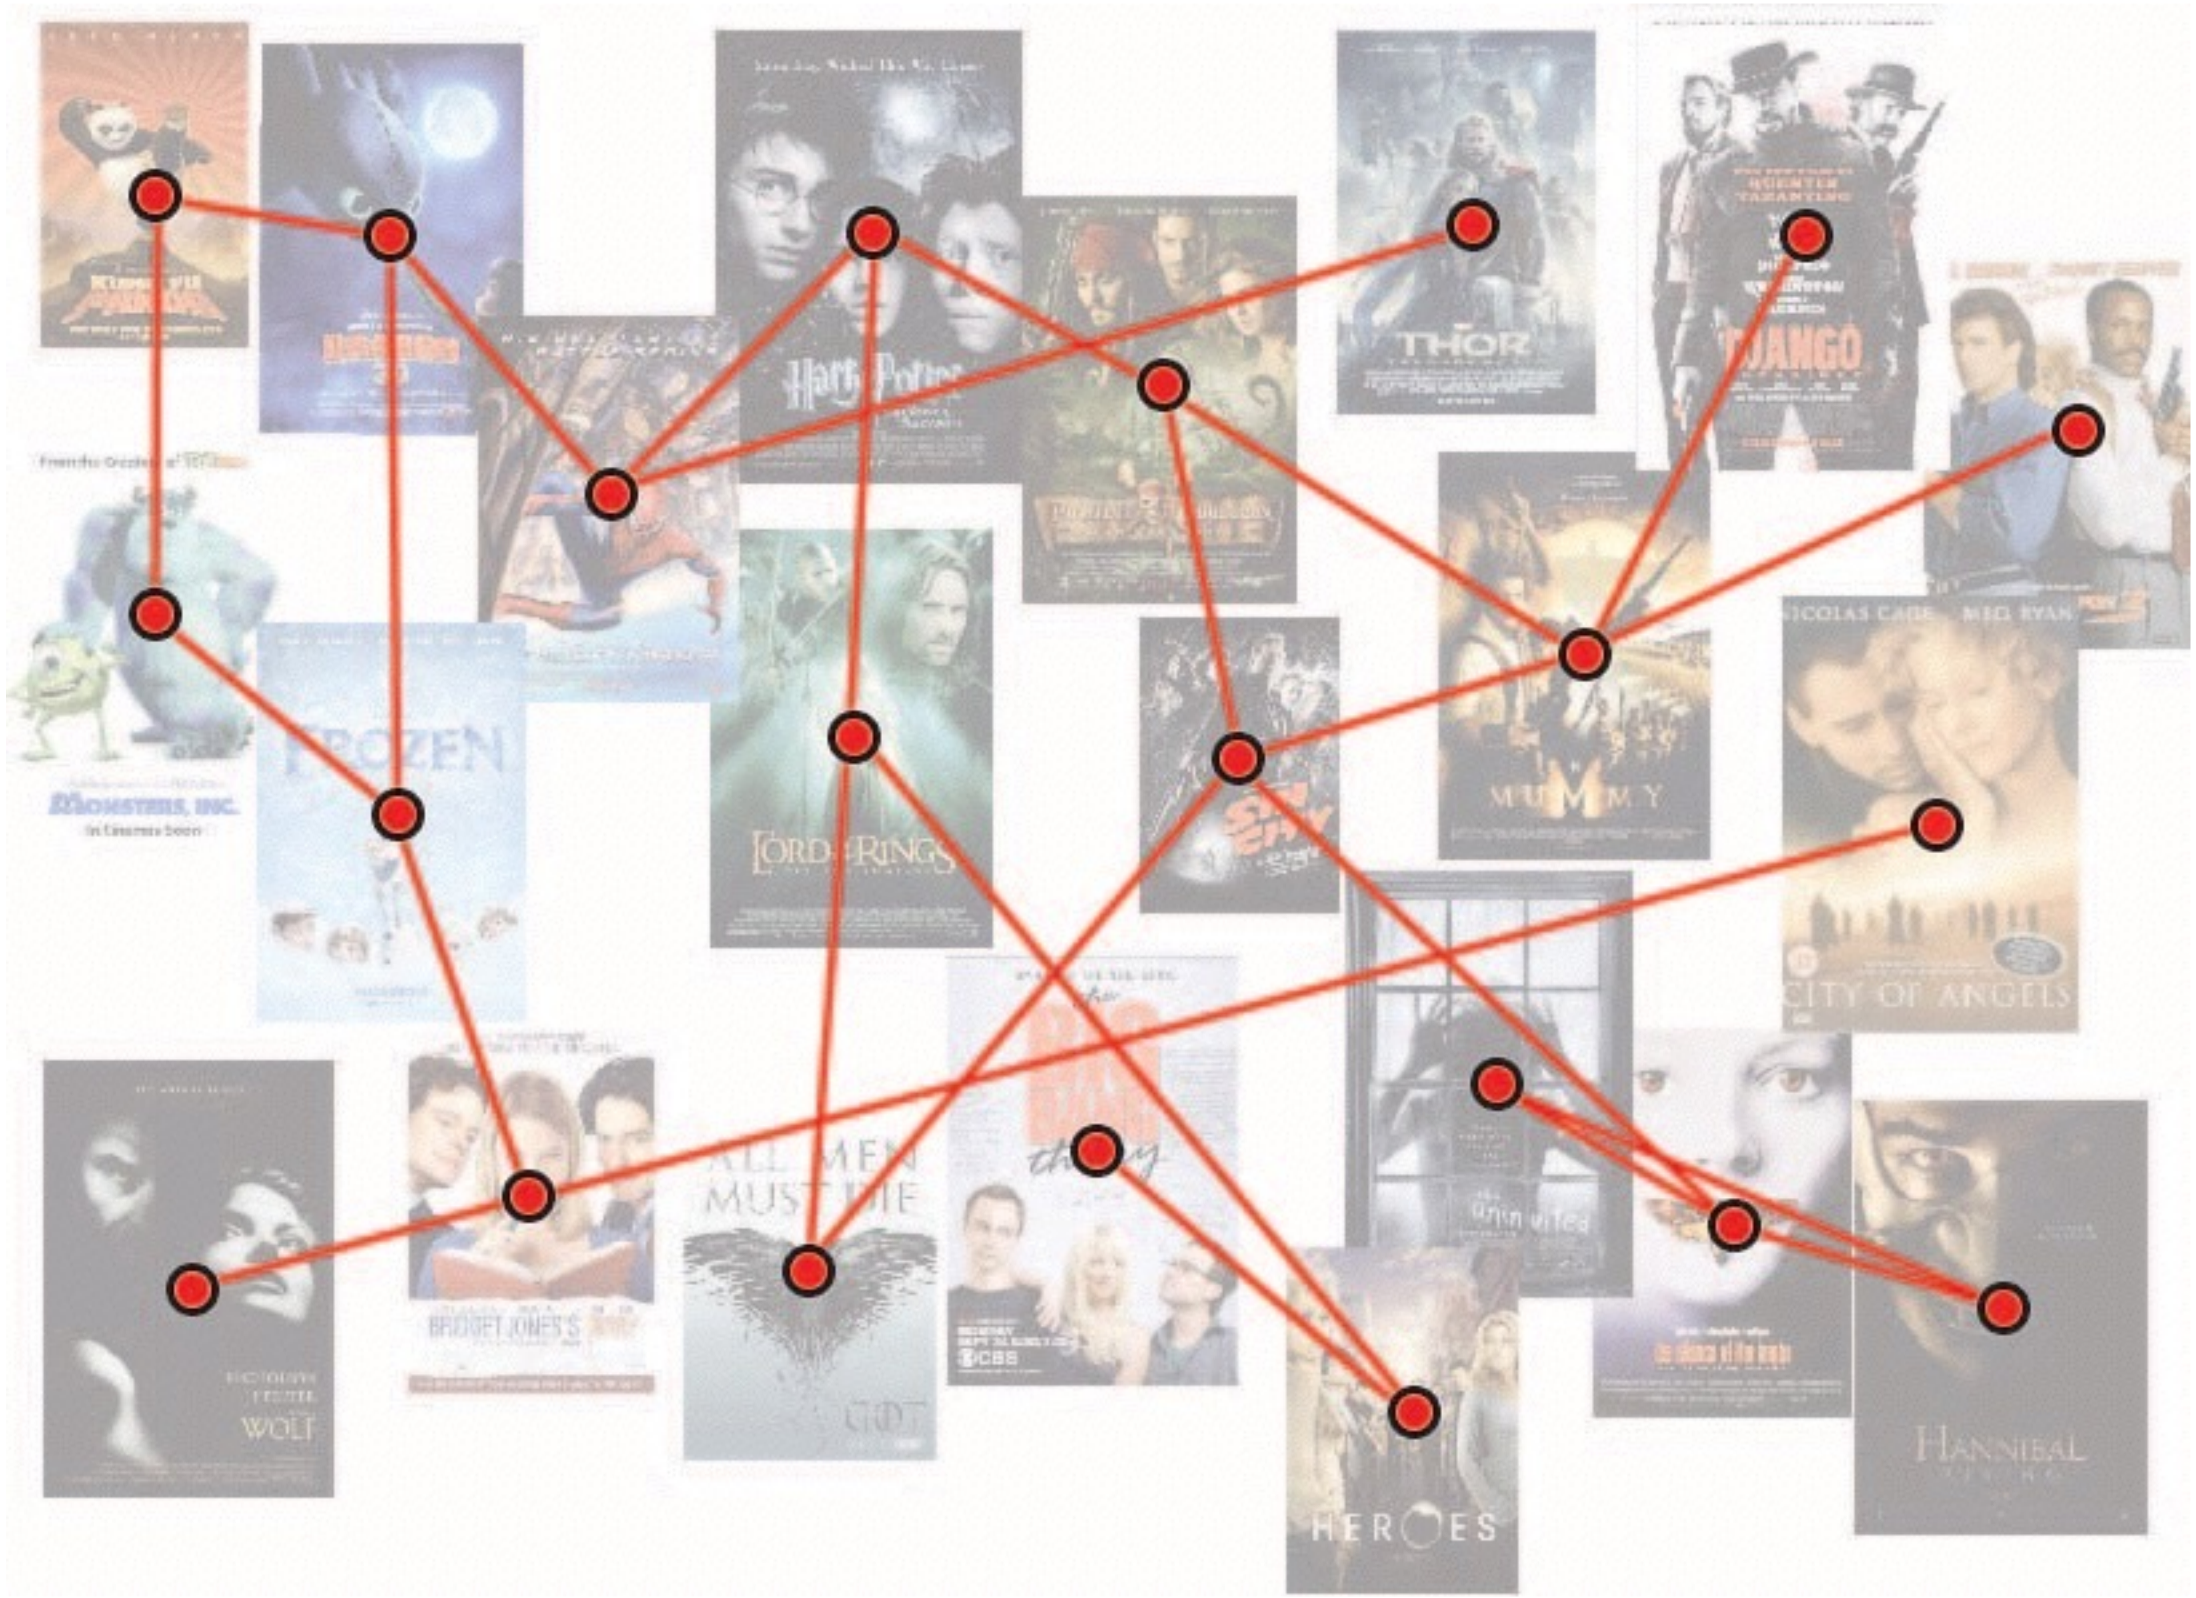
THOR
DJANGO
FROZEN
CITY OF ANGELS
HANNIBAL
HEROES
ALL MEN MUST DIE

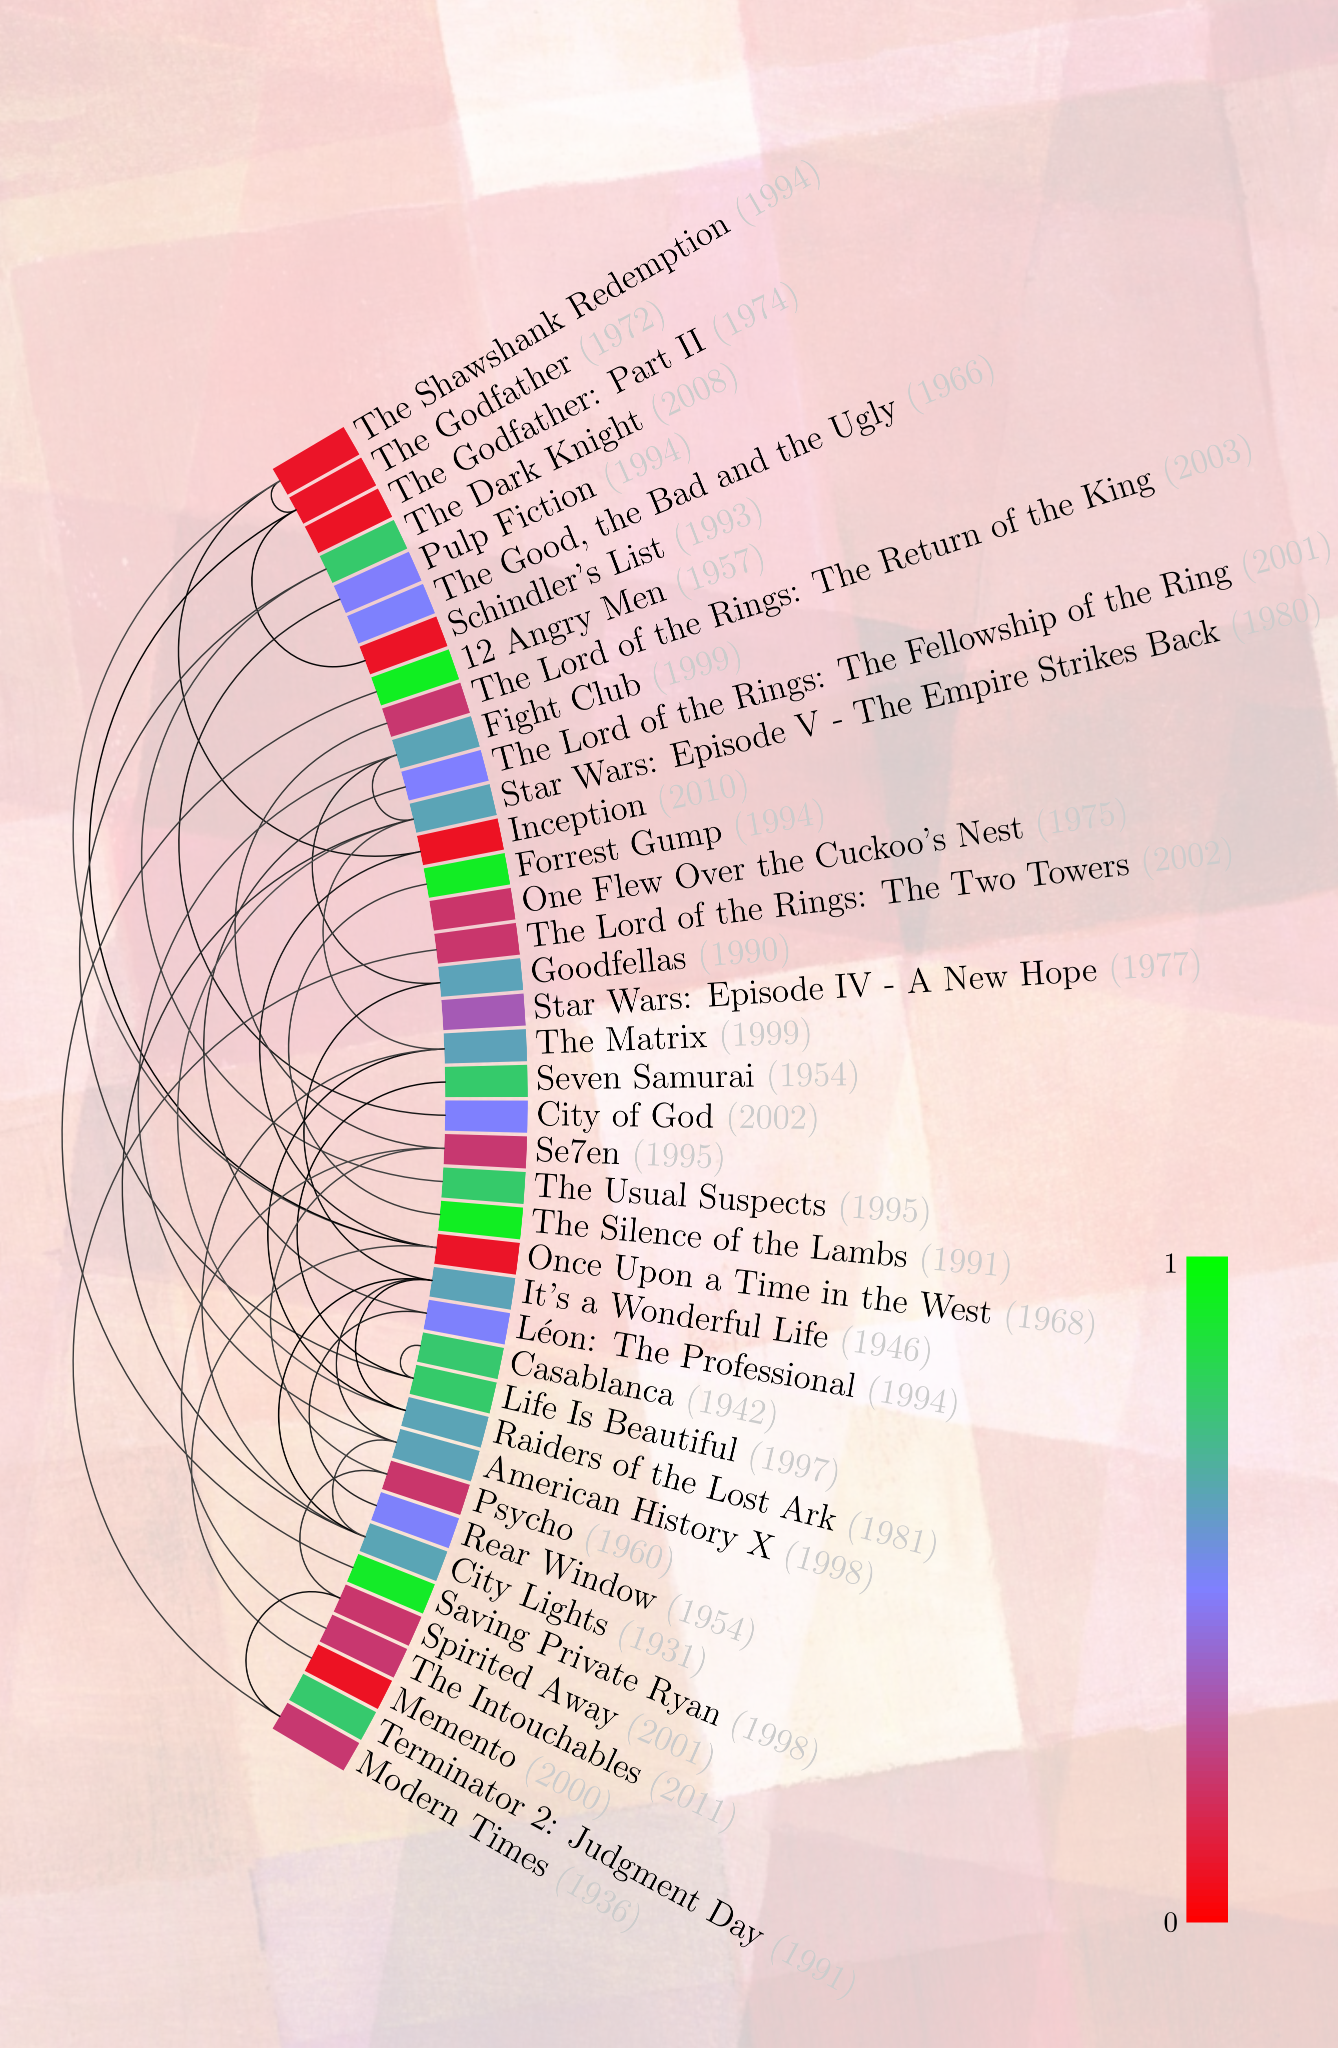

Example of a graph bandit problem

movie recommendation

- recommend movies to a **single** user
- **goal:** maximise the sum of the ratings (minimise regret)
- good prediction after just a few steps

$$T \ll N$$

- extra information
  - ratings are **smooth** on a graph
- main question: can we learn **faster**?

GETTING REAL

Let's be lazy and ignore the structure

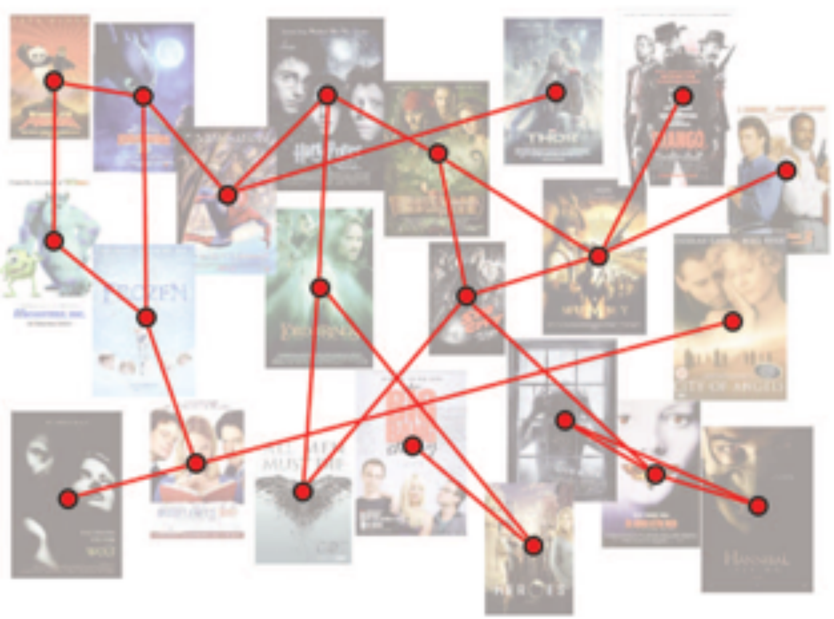

→

Multi-armed bandit problem!

Worst case regret (to the best fixed strategy)

$$R_T = \mathcal{O}\left(\sqrt{NT}\right)$$

#actions

#rounds

Matching lower bound (Auer, Cesa-Bianchi, Freund, Schapire 2002)

How big is N? Number of movies on http://www.imdb.com/stats: 4,513,842

Number of active users on FaceBook: https://newsroom.fb.com/company-info/ 2,017,822,735

Problem: Too many actions!

LEARNING FASTER

$$R_T = \mathcal{O}\left(\sqrt{NT}\right)$$

#actions

#rounds

#dimensions

- Arm independence is too strong and unnecessary
- Replace **N** with something much smaller
  - problem/instance/data dependent
  - example: linear bandits **N** to **D**
- **In this talk: Online Influence Maximization!**
  - sequential problems where **actions are nodes** on a graph
  - find strategies that replace **N** with a smaller **graph-dependent** quantity

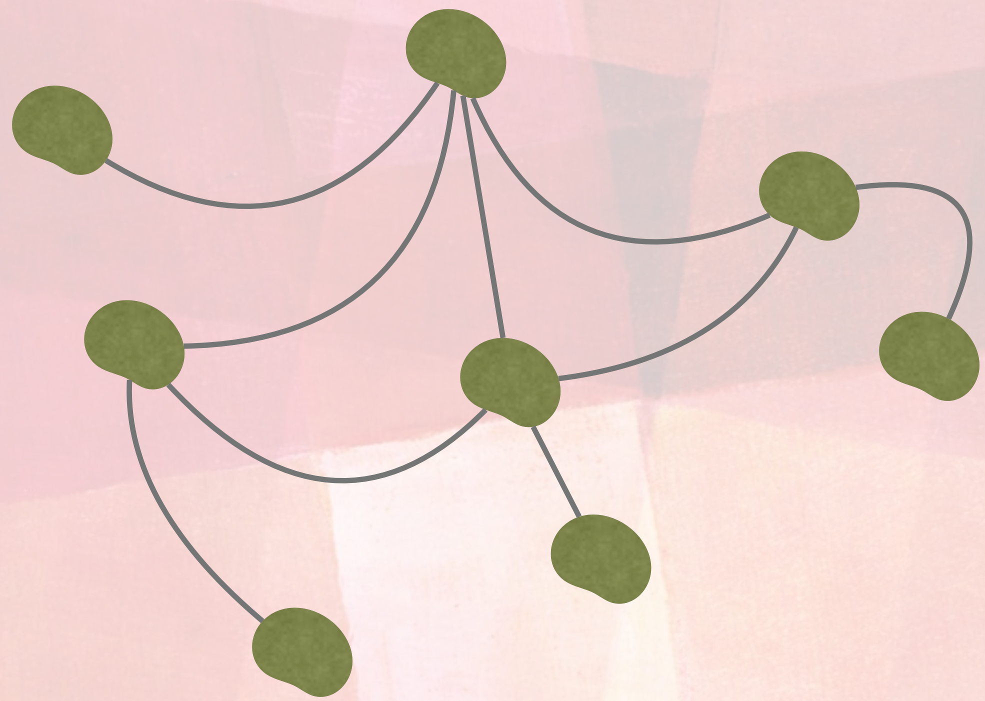

GRAPH BANDITS: GENERAL SETUP

Every round **t** the learner

- picks a node $I_t \in [N]$
- incurs a loss ℓ_{t,I_t}
- optional feedback

The performance is total expected regret

$$R_T = \max_{i \in [N]} \mathbb{E}\left[\sum_{t=1}^{T} (\ell_{t,I_t} - \ell_{t,i})\right]$$

Specific problems differ in

1. loss
2. feedback
3. guarantees

STRUCTURES IN BANDIT PROBLEMS

GRAPHS

KERNELS

POLYMATROIDS

BLACK-BOX FUNCTIONS

STRUCTURES WITHOUT TOPOLOGY

…

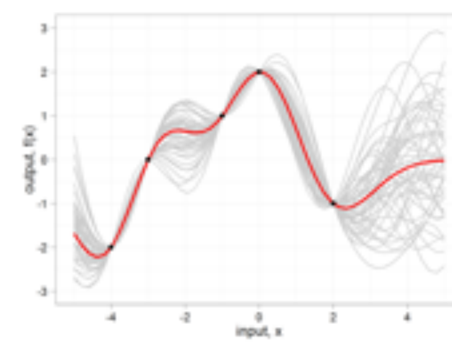

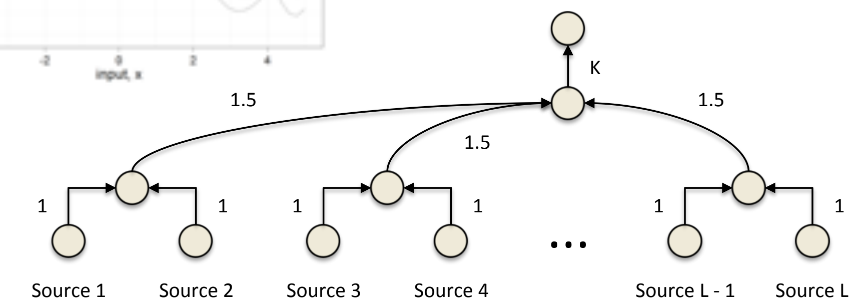

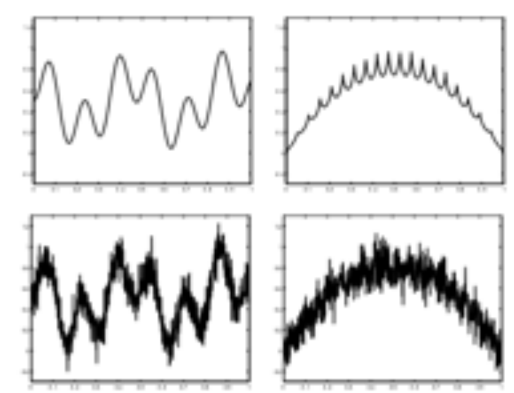

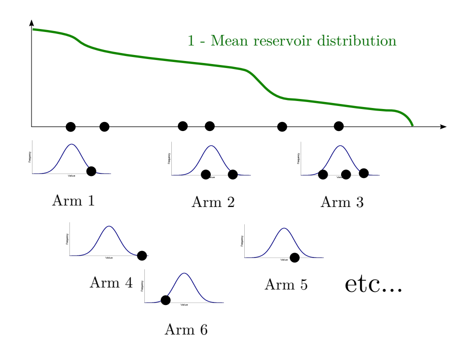

SPECIFIC GRAPH BANDIT SETTINGS

smoothness
spectral bandits
$$R_T = \widetilde{\mathcal{O}}\left(d\sqrt{T \ln T}\right)$$

#relevant eigenvectors

side observations
on graphs
$$R_T = \widetilde{\mathcal{O}}\left(\sqrt{\overline{\alpha} T \ln N}\right)$$

independence number

influence maximisation
revealing bandits
$$R_T = \widetilde{\mathcal{O}}\left(\sqrt{r_* T D_*}\right)$$

detectable dimension

Survey: http://researchers.lille.inria.fr/~valko/hp/publications/valko2016bandits.pdf

Decision Theory and Network Science: Methods and Applications,

STOR-i Workshop at Lancaster University, Sep 18th, 2017

ONLINE INFLUENCE MAXIMIZATION

Michal Valko, SequeL, Inria Lille - Nord Europe


2015-2016, AISTAST 2016




2016-2017, NIPS 2017

https://arxiv.org/abs/1605.06593

HOW TO RULE THE WORLD?

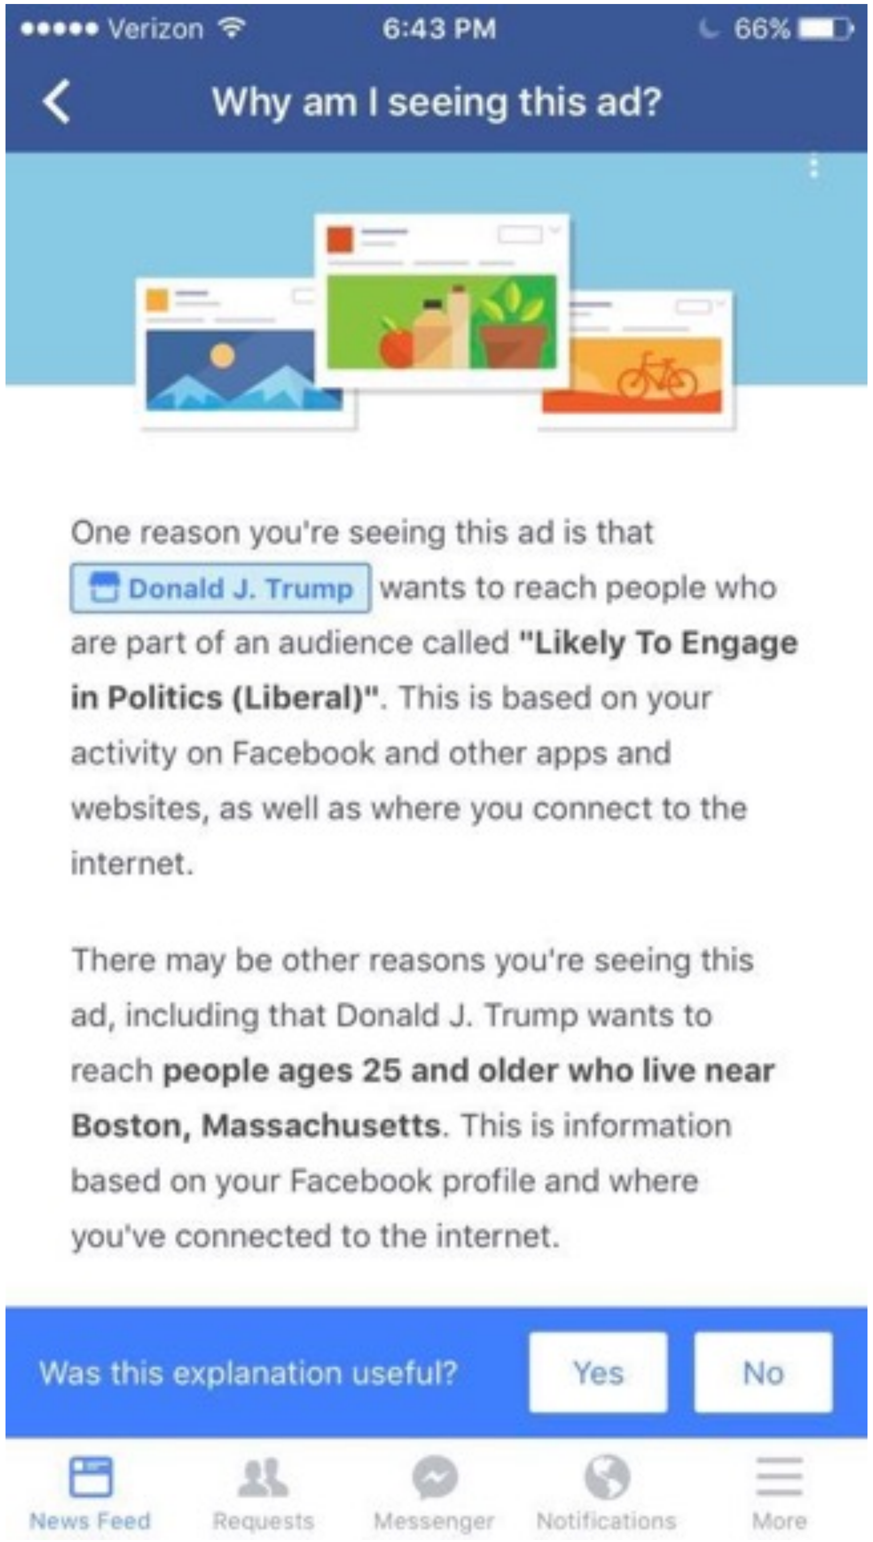

“IA” EST DÉJÀ LÀ

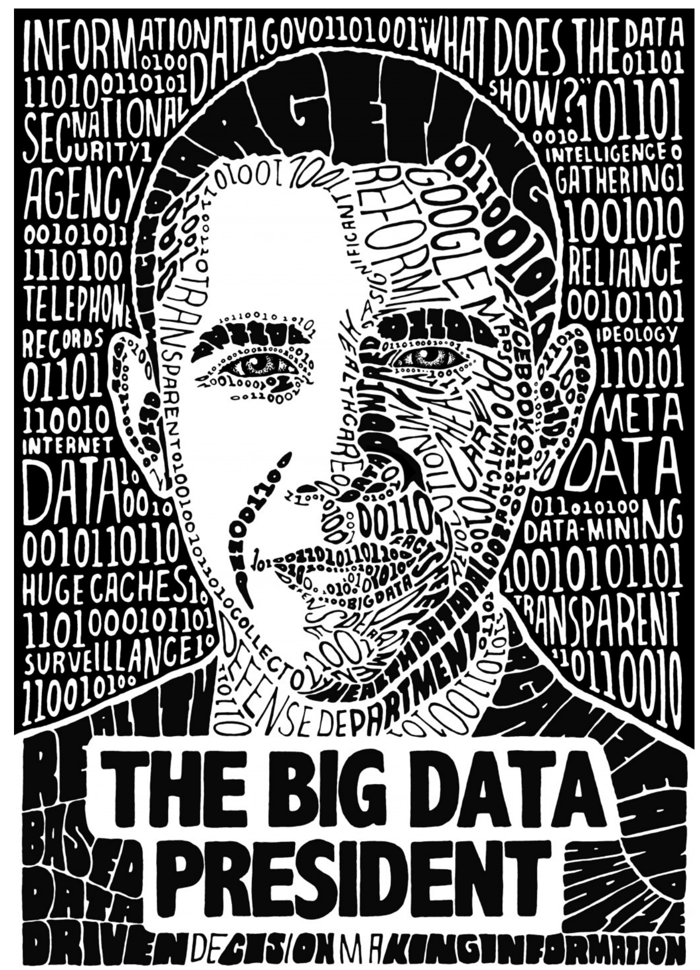

`https://www.washingtonpost.com/opinions/obama-the-big-data-president/2013/06/14/1d71fe2e-d391-11e2-b05f-3ea3f0e7bb5a_story.html`

`https://www.technologyreview.com/s/509026/how-obamas-team-used-big-data-to-rally-voters/`

Talk of Rayid Ghaniy: `https://www.youtube.com/watch?v=gDM1GuszM_U`

INSOUMISE OU ENRACINÉE ?

Le "big data" ou la recette secrète du succès d'Emmanuel Macron?

https://www.rts.ch/info/sciences-tech/8580821-le-big-data-ou-la-recette-secrete-du-succes-d-emmanuel-macron-.html

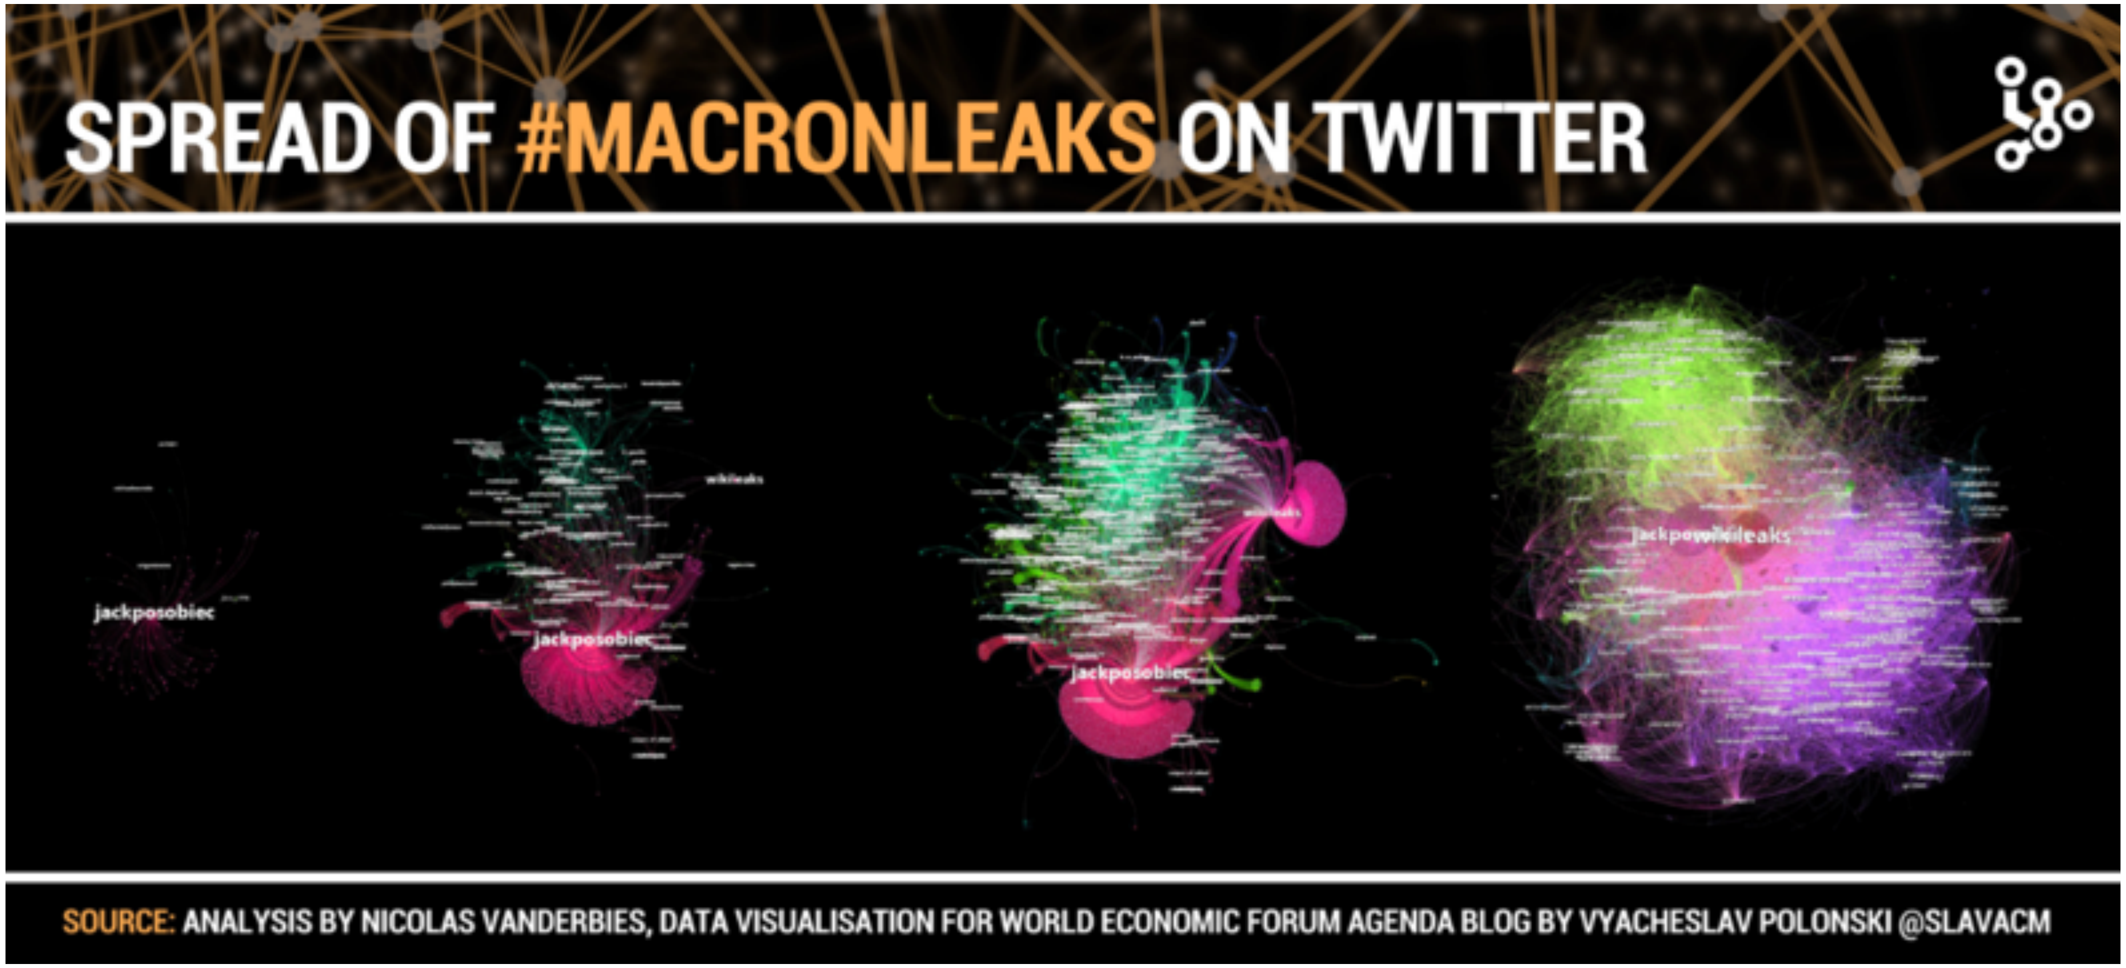

HOW TO RULE THE WORLD?

Influence the influential!

JULY 18, 2016

Religion

March 26, 2017

Politics

September 1, 2009

Culture

Influence the influential in England?

Religion

Politics

Culture

HOW TO RULE THE WORLD?

informatics mathematics

Influence the influential in England!

Religion ?

Politics ?

Culture

Influence the influential in England!

Religion ?

Politics ?

Culture

338 ET 200

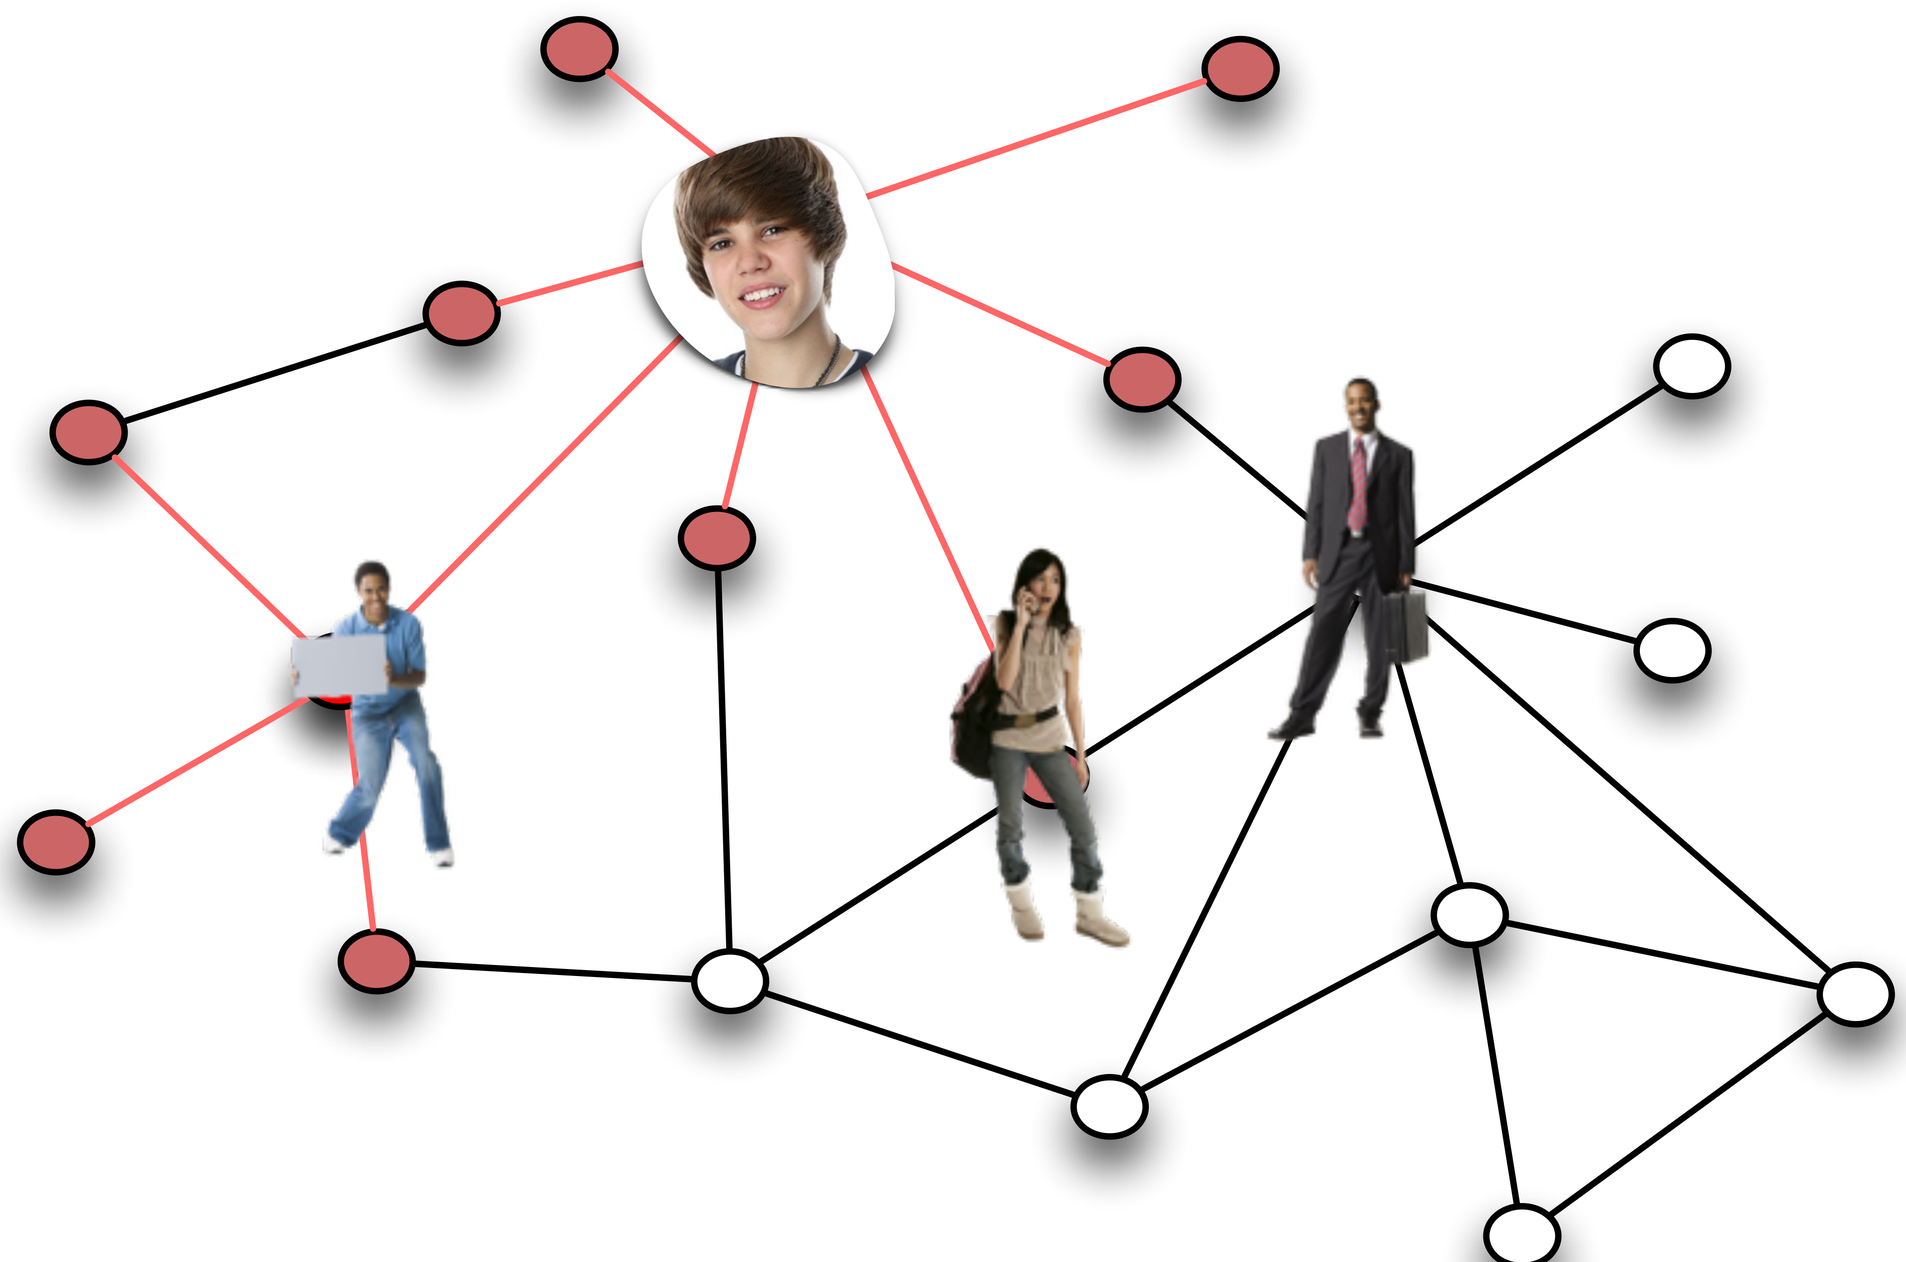

$F(S) = \text{spread}$

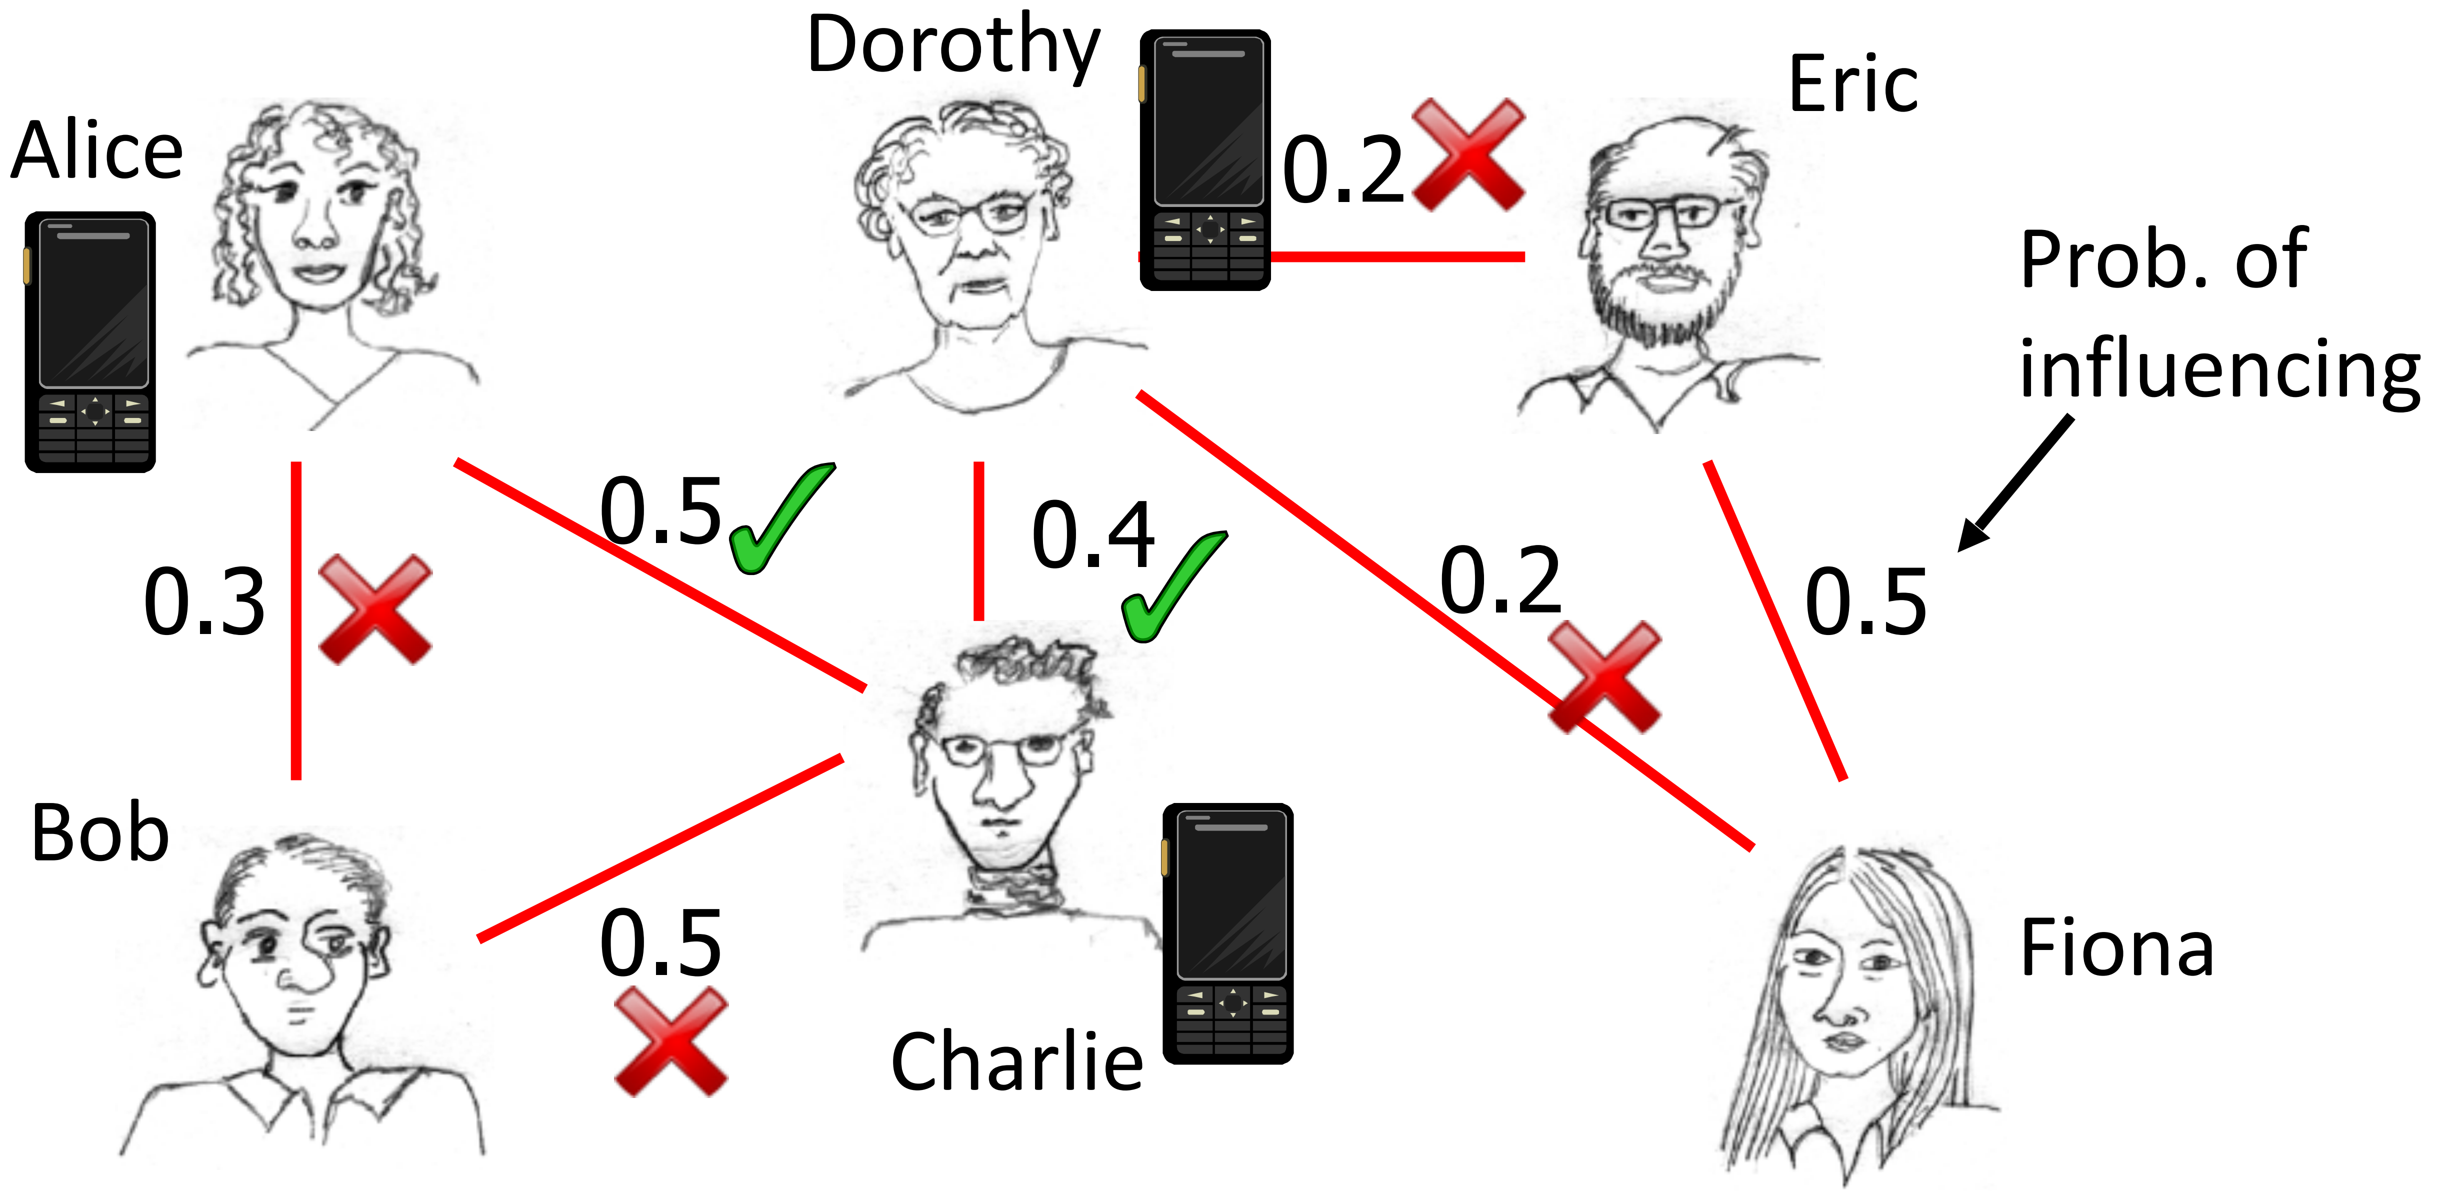

Who should get free cell phones?

V = {Alice,Bob,Charlie,Dorothy,Eric,Fiona}

F(A) = Expected number of people influenced when targeting A

MAXIMIZING INFLUENCE

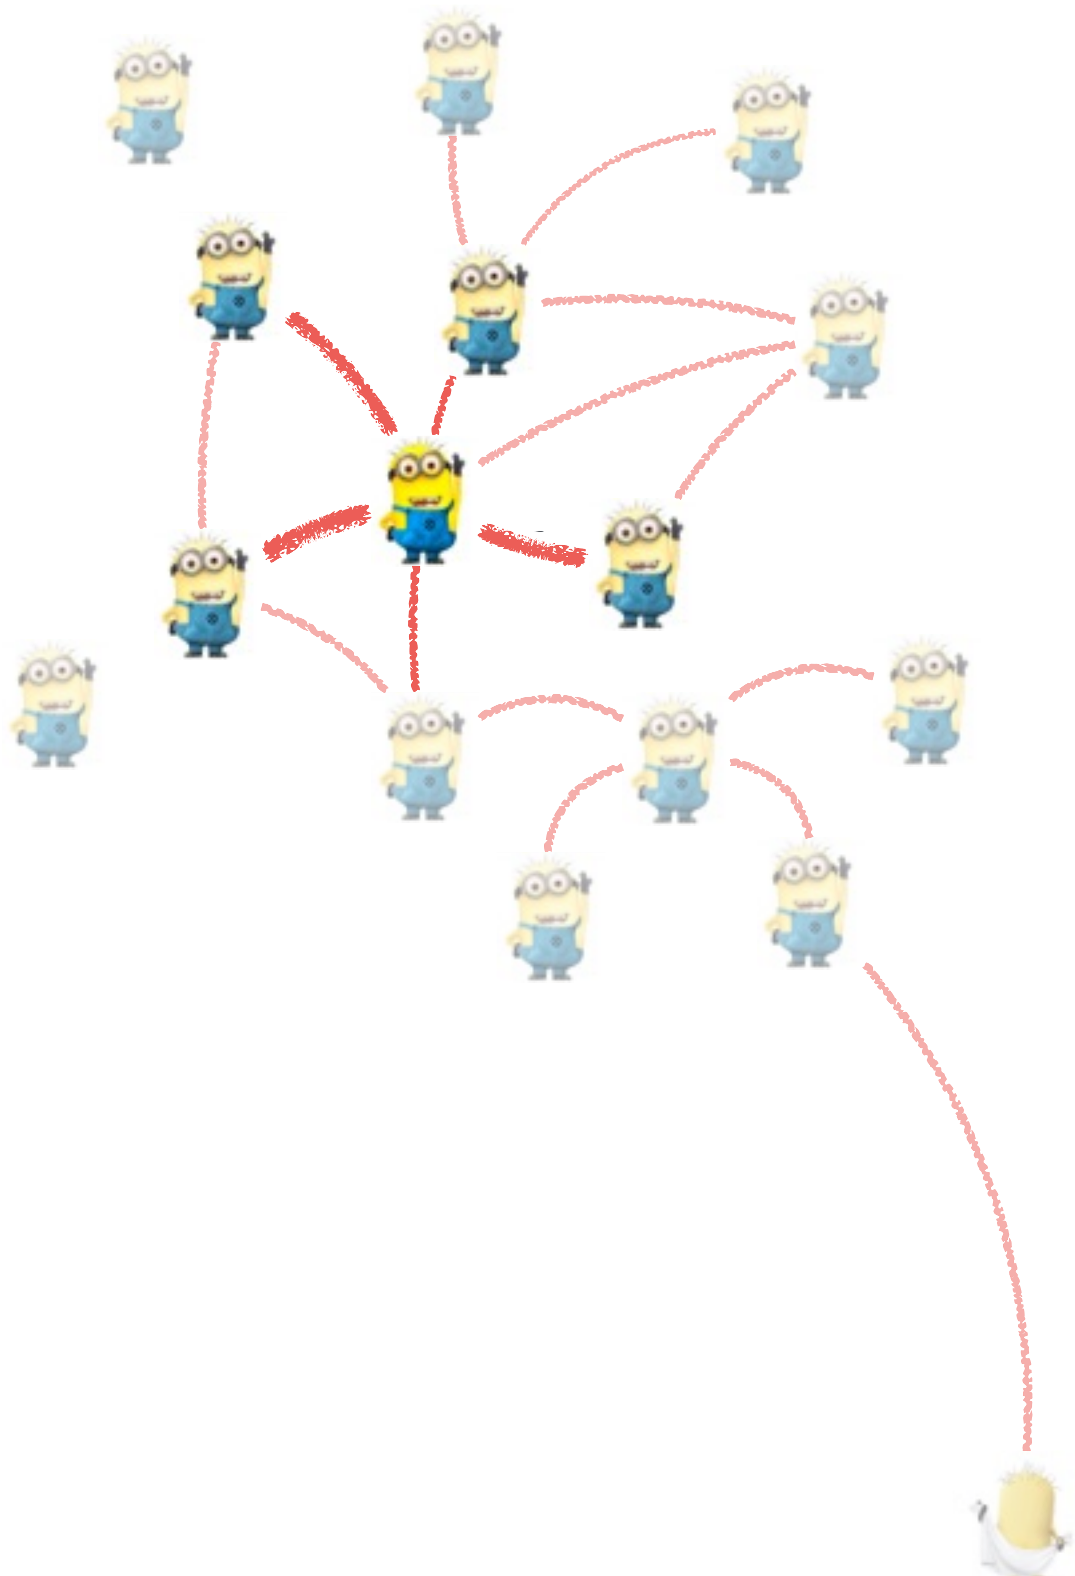

Product placement

- dispatch few to sell more
- target influential people

Gathering the information?

- likes on FB
- promotional codes

Unknown graphs

- all prior work needed to know the graph
- here: provably learning faster without it

REVEALING BANDITS FOR LOCAL INFLUENCE

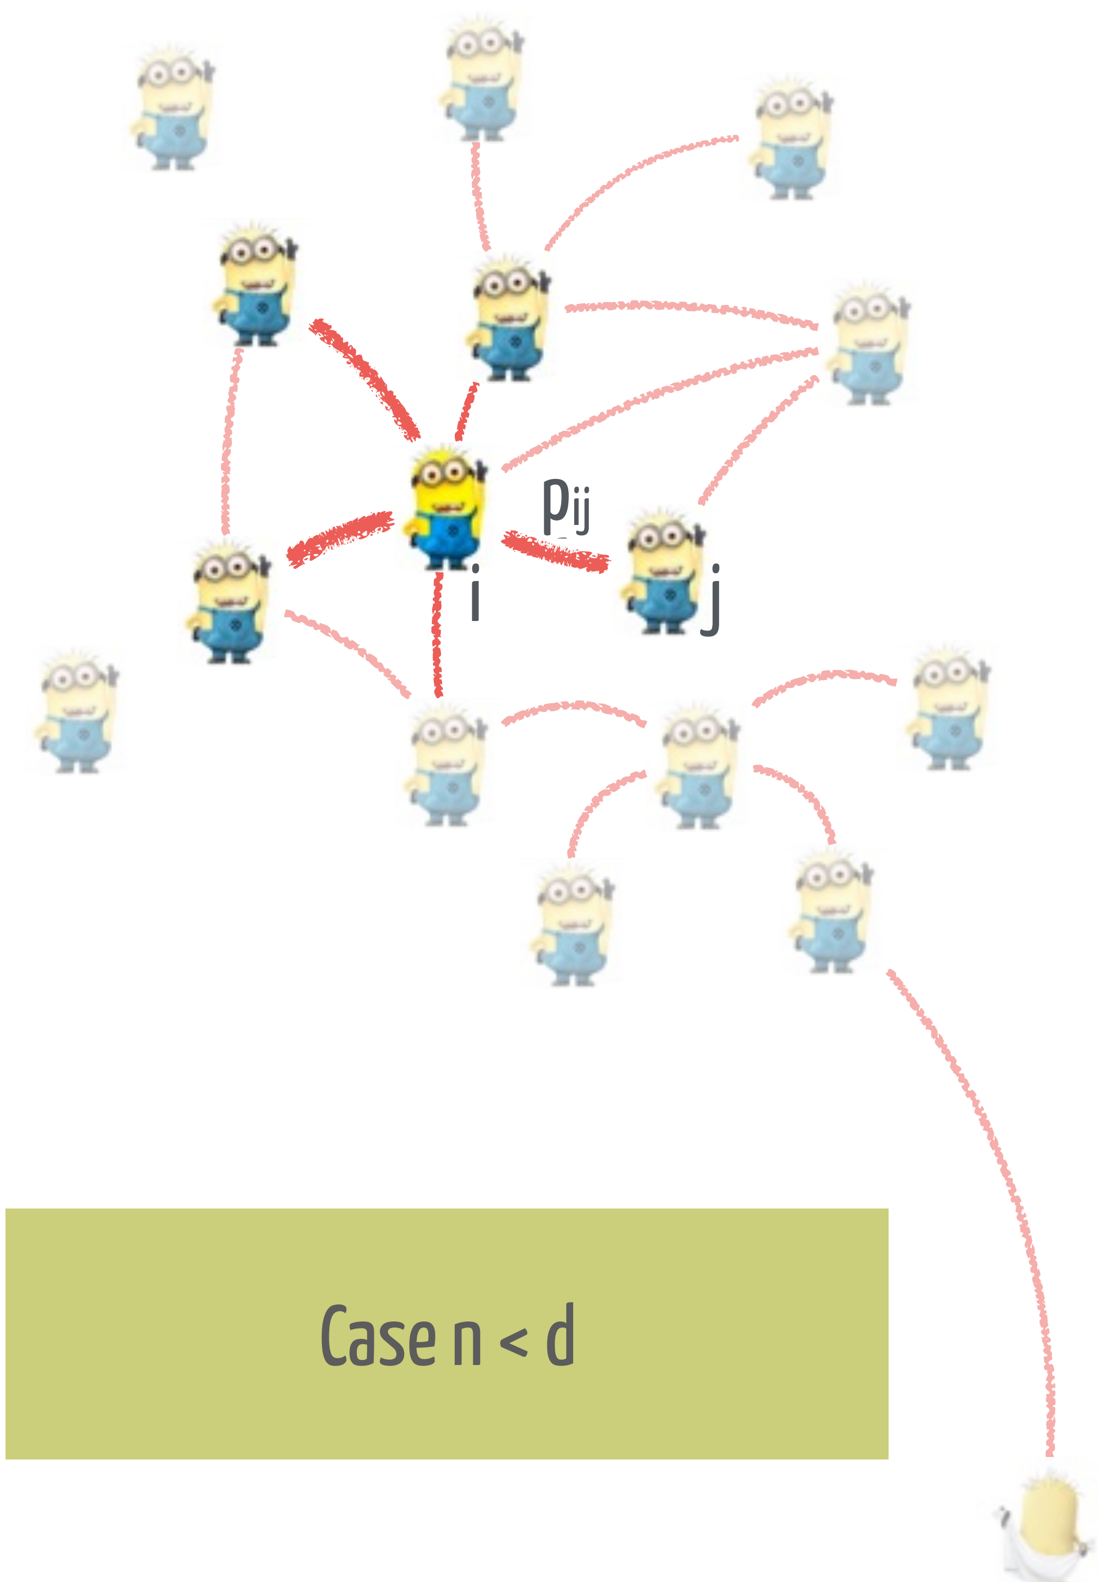

Unknown $(p_{ij})_{ij}$ — (symmetric) probability of influences

In each time step t = 1, …., n

learner picks a node k_t

environment **reveals** the set of influenced node S_{k_t}

Select influential people = Find the strategy maximising

$$L_n = \sum_{t=1}^{n} |S_{k_t,t}|$$

Why this is a bandit problem?

Case n < d

What are bandits anyway?

PERFORMANCE CRITERION

The number of expected influences of node **k** is by definition

$$r_k = \mathbb{E}\left[|S_{k,t}|\right] = \sum_{j \leq d} p_{k,j}$$

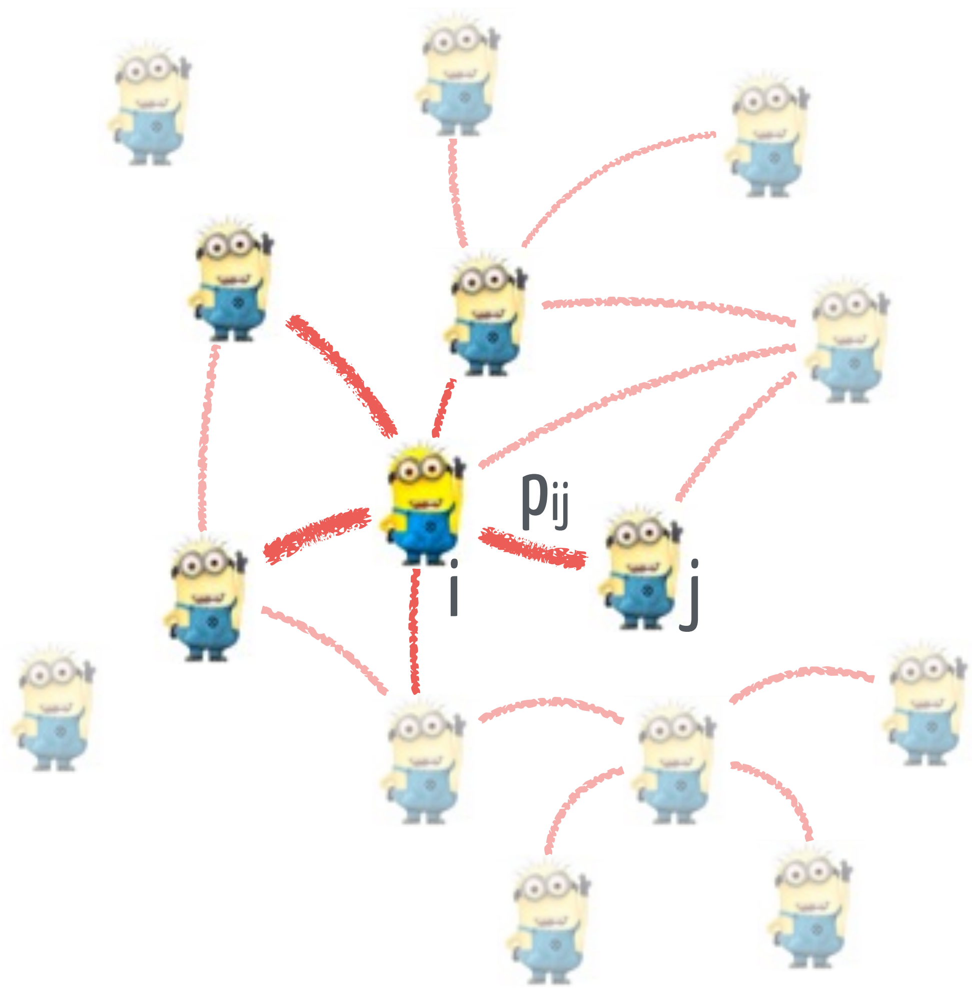

Oracle strategy always selects the best

$$k^\star = \arg\max_k \mathbb{E}\left[\sum_{t=1}^{n} |S_{k,t}|\right] = \arg\max_k n r_k$$

Expected reward of the oracle strategy

$$\mathbb{E}\left[L_n^\star\right] = n r_\star$$

Expected regret of any adaptive strategy unaware of $(p_{ij})_{ij}$

$$\mathbb{E}\left[R_n\right] = \mathbb{E}\left[L_n^\star\right] - \mathbb{E}\left[L_n\right]$$

UPPER CONFIDENCE BOUND BASED ALGOS

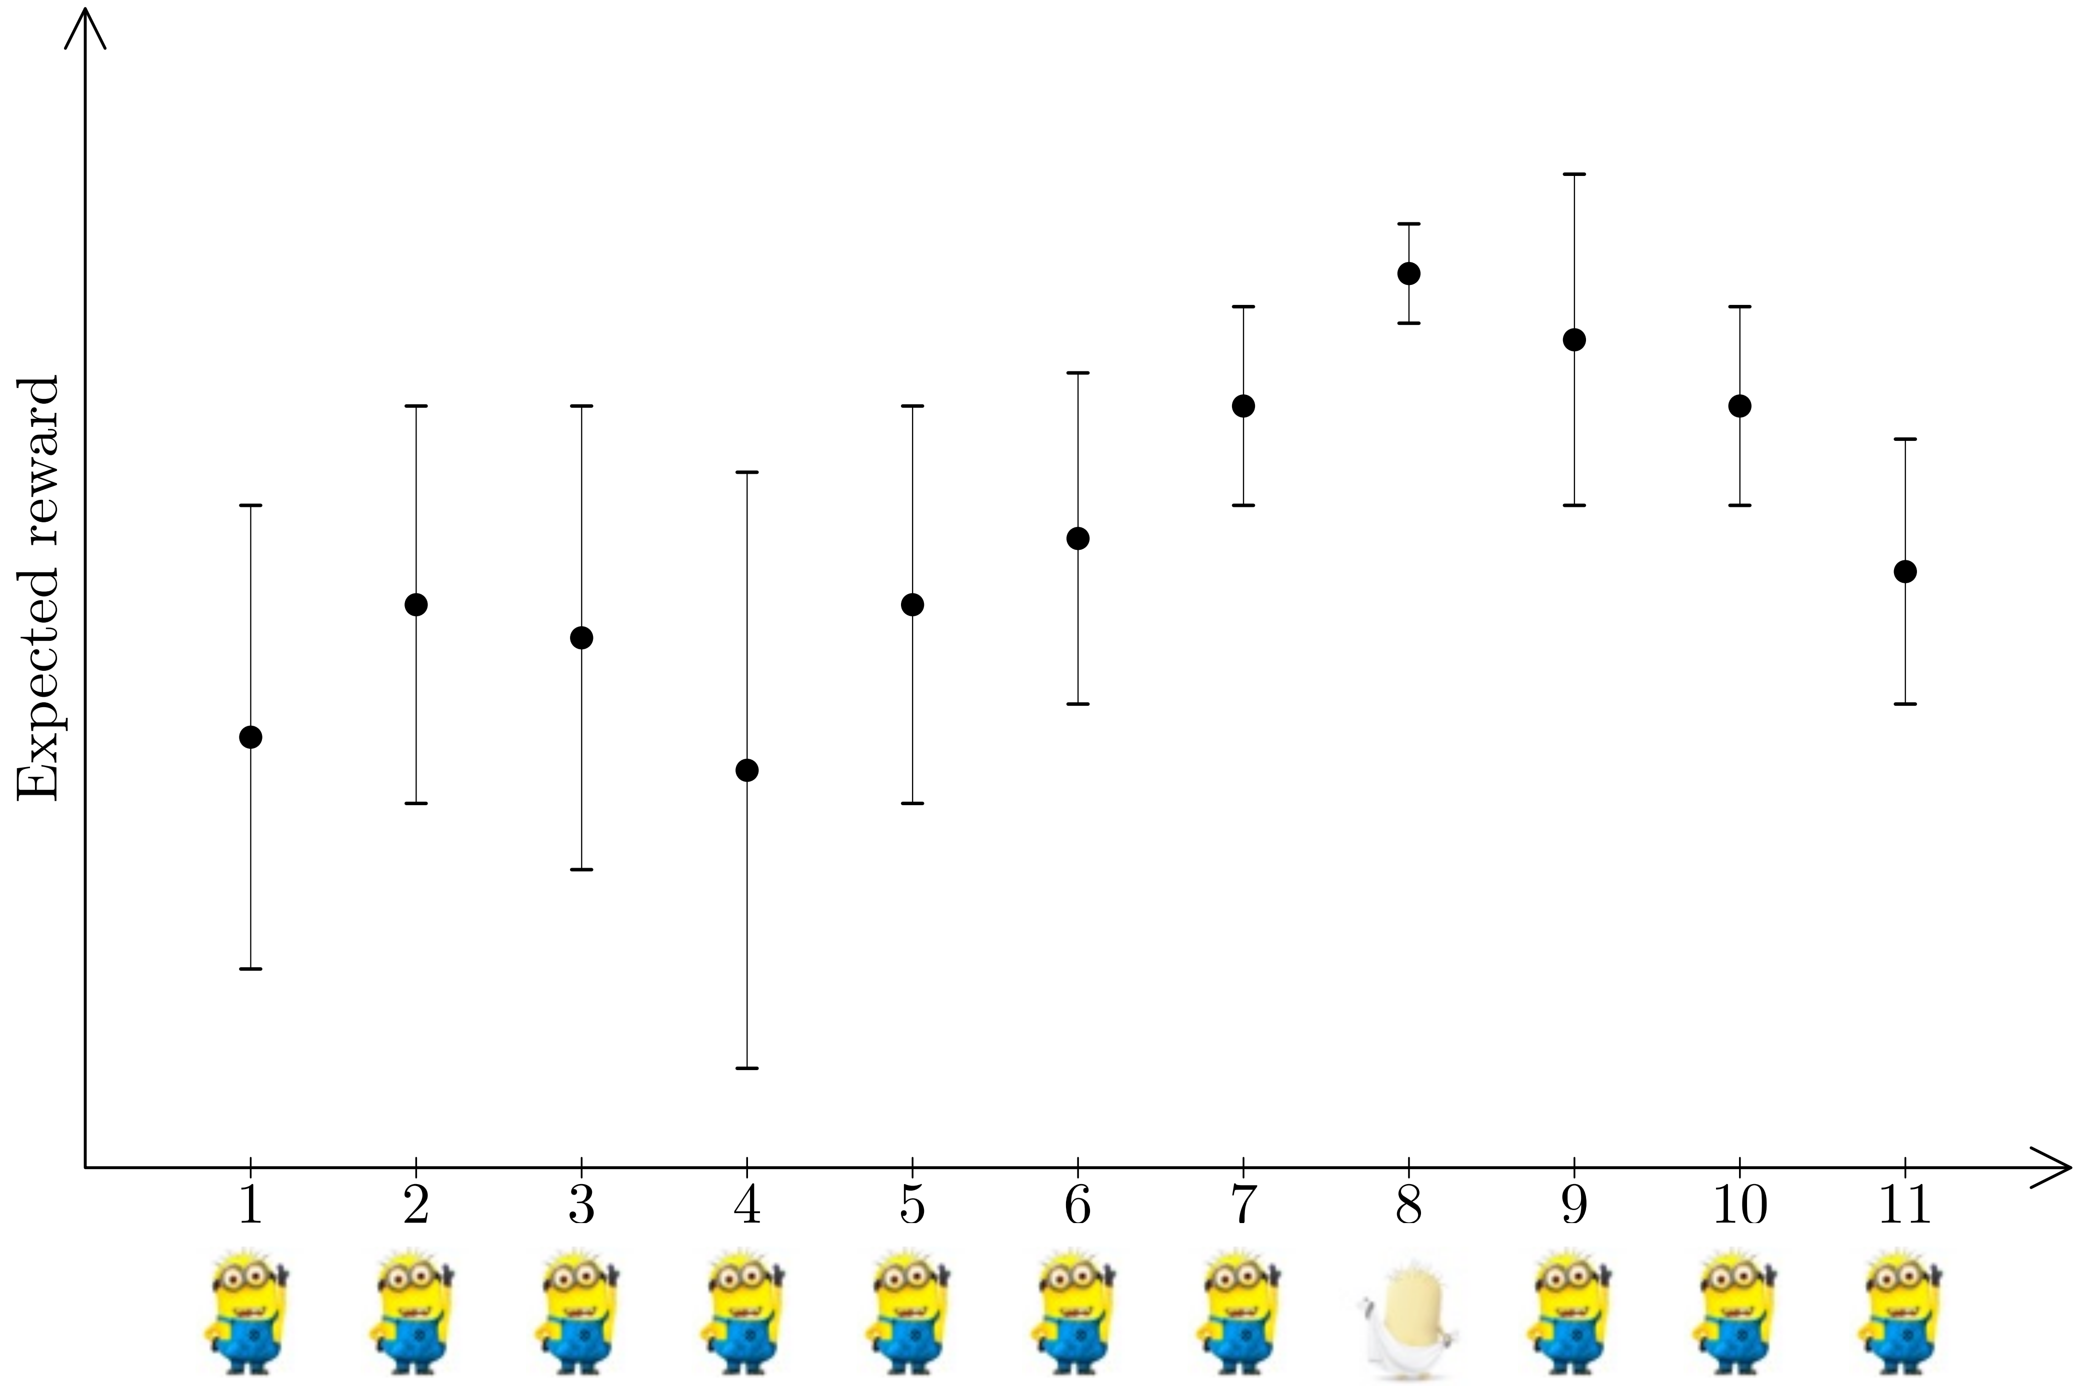

MULTI-ARM BANDITS IN CAFÉ CULTURE

Video recorded March 30th, 2017, 13h50,
Université de Lille, Susie & the Piggy Bones Band

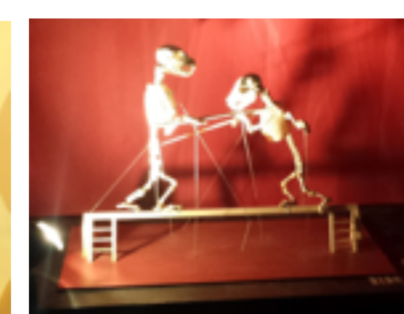

UPPER CONFIDENCE BOUND BASED ALGOS

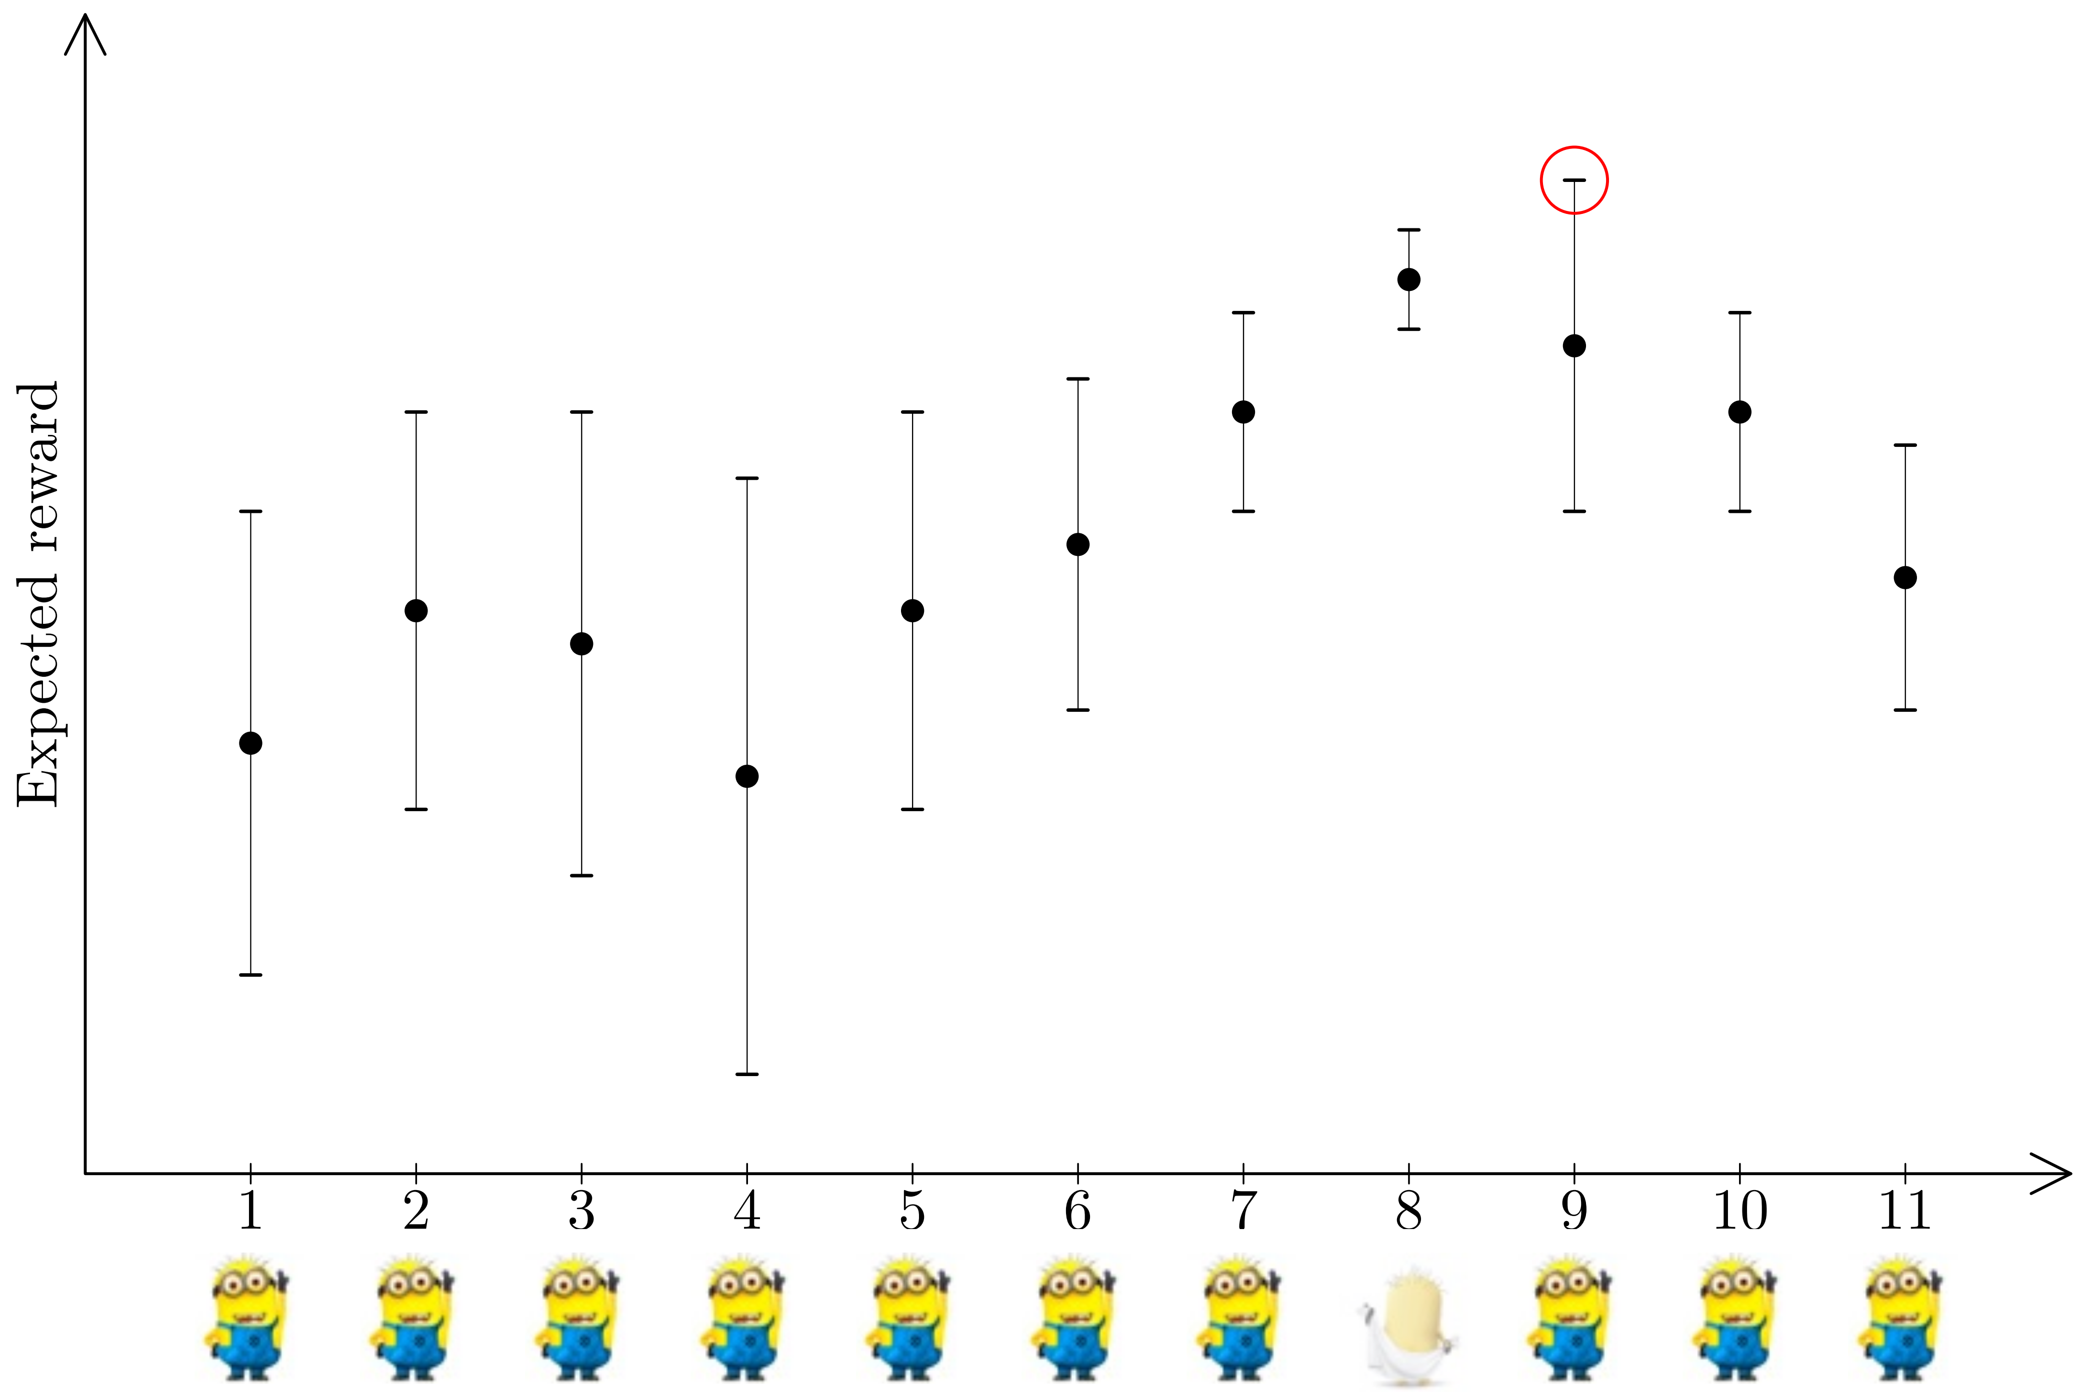

UPPER CONFIDENCE BOUND BASED ALGOS

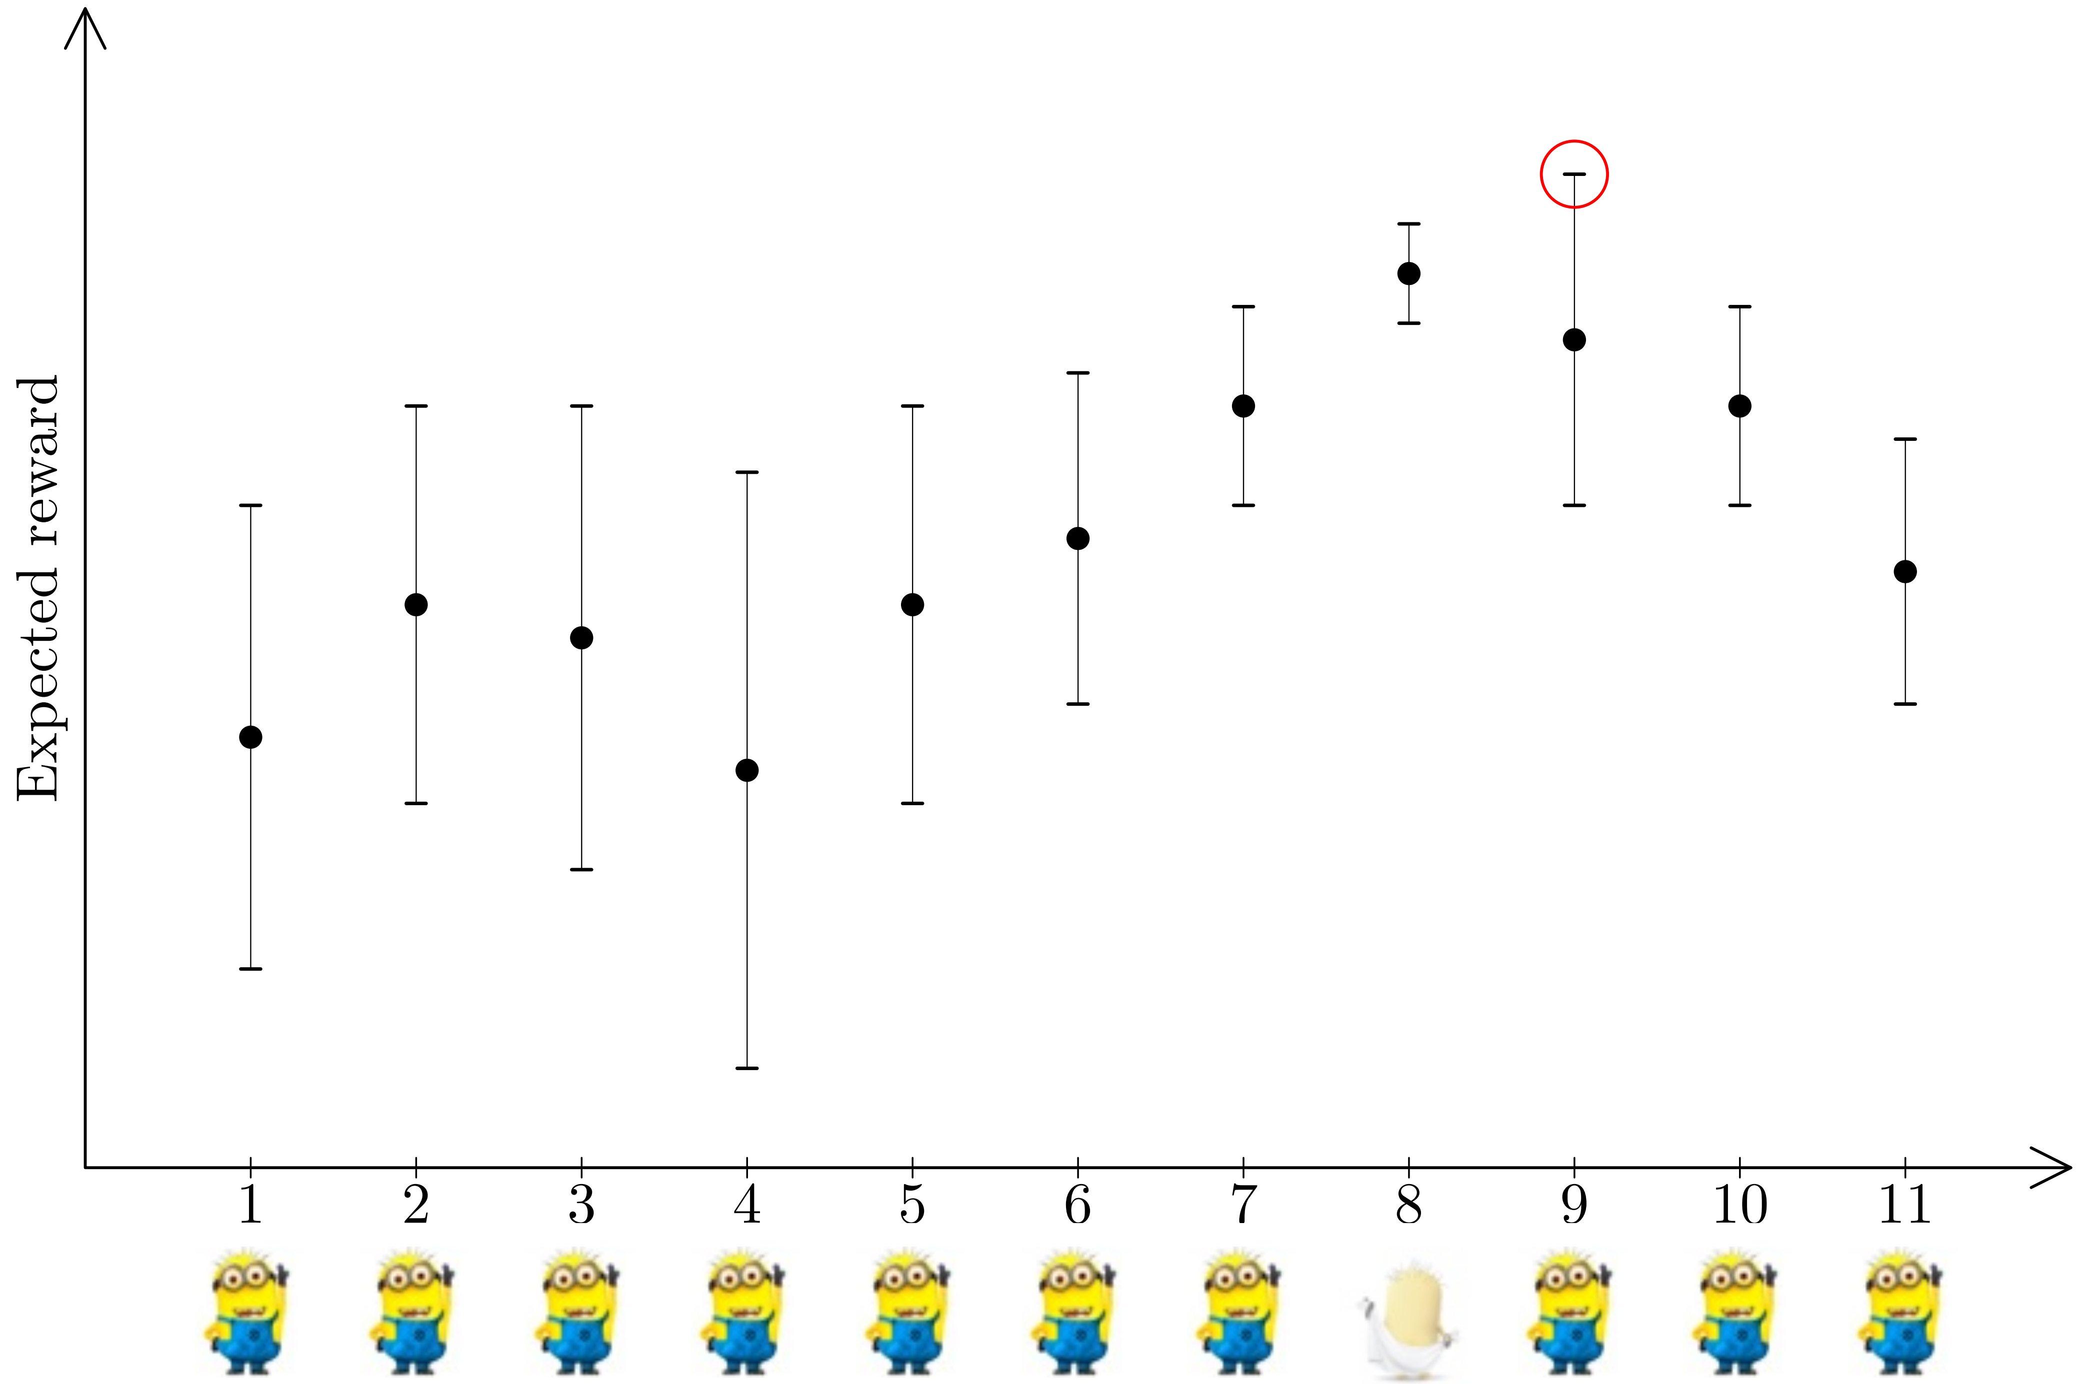

UPPER CONFIDENCE BOUND BASED ALGOS

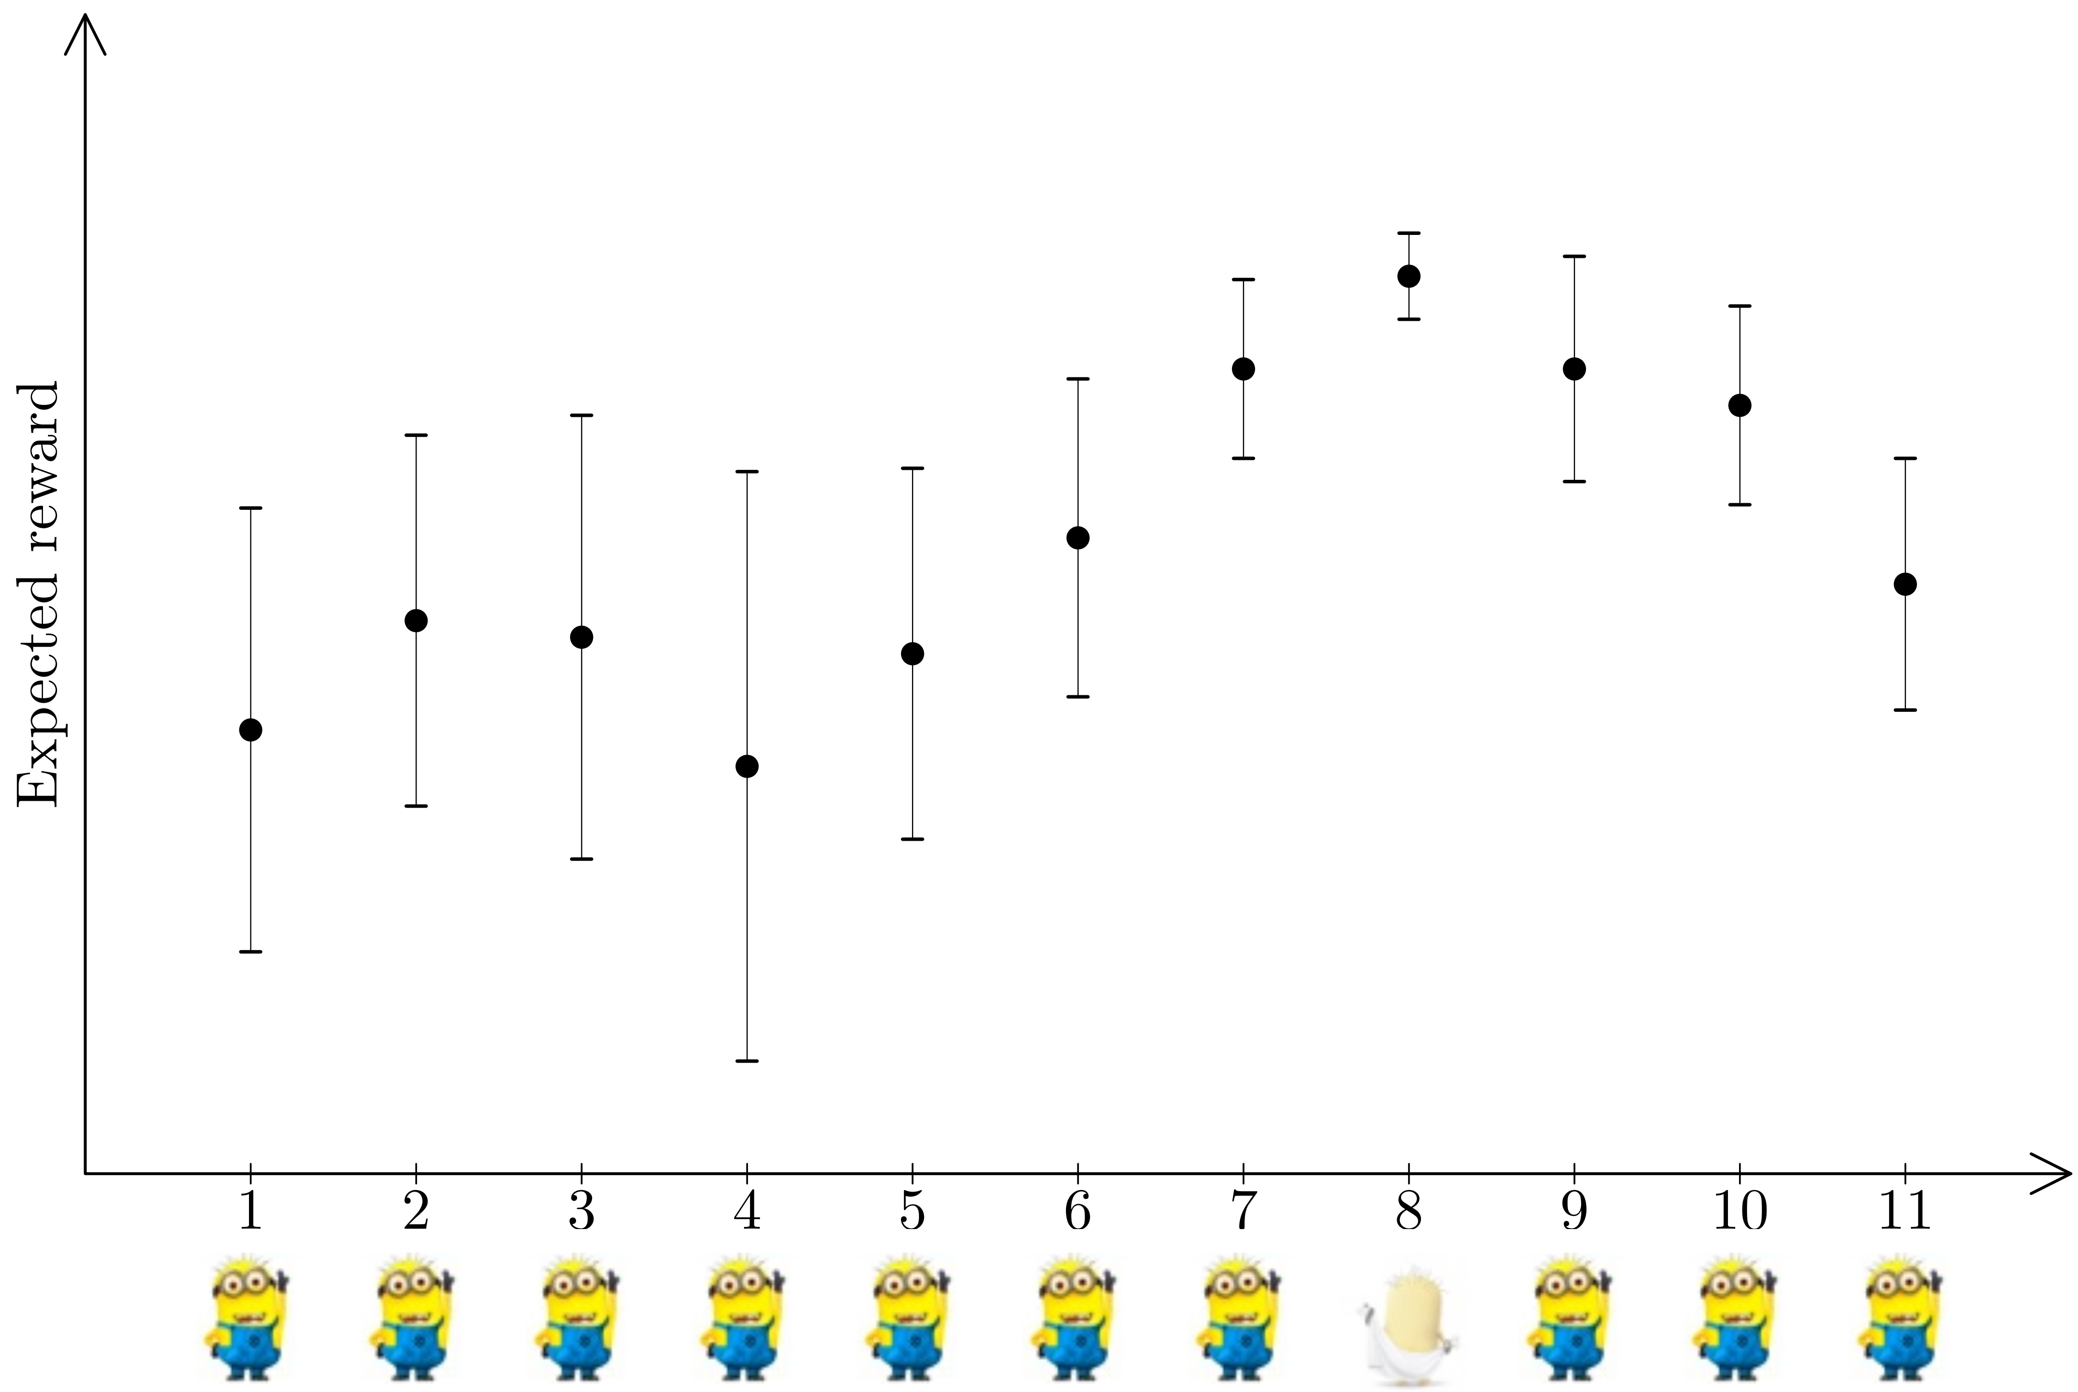

DETECTABLE DIMENSION

- number of nodes we can efficiently extract in less than n rounds
- function D controls number of nodes given a gap

$$D(\Delta) \stackrel{\text{def}}{=} |\{i \leq d : r_\star^\circ - r_i^\circ \leq \Delta\}|$$

- D(r) = d for r≥ r_* and D(0) = number of most influenced nodes
- **Detectable dimension D_*** = $D(\Delta_*)$
- Detectable gap Δ_* constants coming from the analysis and the Bernstein inequality

$$\Delta_\star \stackrel{\text{def}}{=} 16\sqrt{\frac{r_\star^\circ d \log(nd)}{T_\star}} + \frac{80 d \log(nd)}{T_\star}$$

- Detectable horizon T_*, smallest integer s.t. $T_\star r_\star^\circ \geq \sqrt{D_\star n r_\star^\circ}$
- Equivalently: D_* corresponding to smallest T_* such that

$$T_\star r_\star^\circ \geq \sqrt{D\left(16\sqrt{\frac{r_\star^\circ d \log(nd)}{T_\star}} + \frac{80 d \log(nd)}{T_\star}\right) n r_\star^\circ}$$

HOW DOES D* BEHAVE?

- For (easy, structured) star graphs $D_* = 1$ even for small n (big gain)
- For (difficult) empty graphs $D_* = d$ even for large n (no gain)
- In general: D_* roughly decreases with n and it is small when D decreases quickly
- For n large enough D_* is the number of the most influences nodes
- Example: D_* for Barabási–Albert model & Enron graph as a function of n

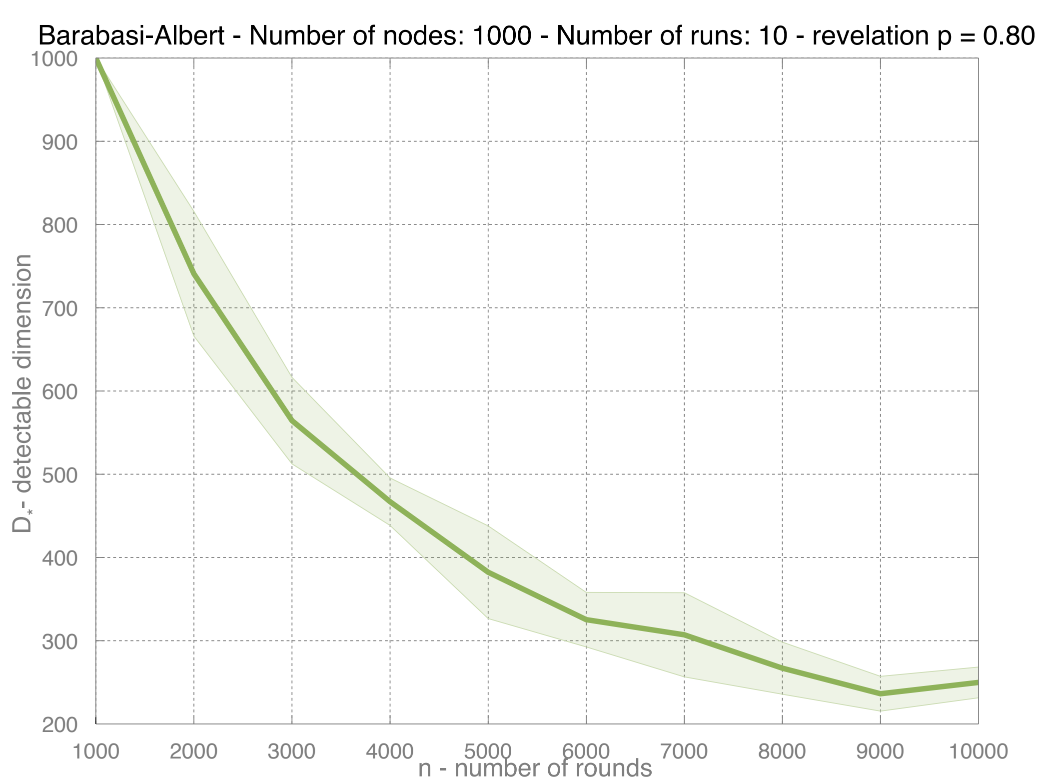

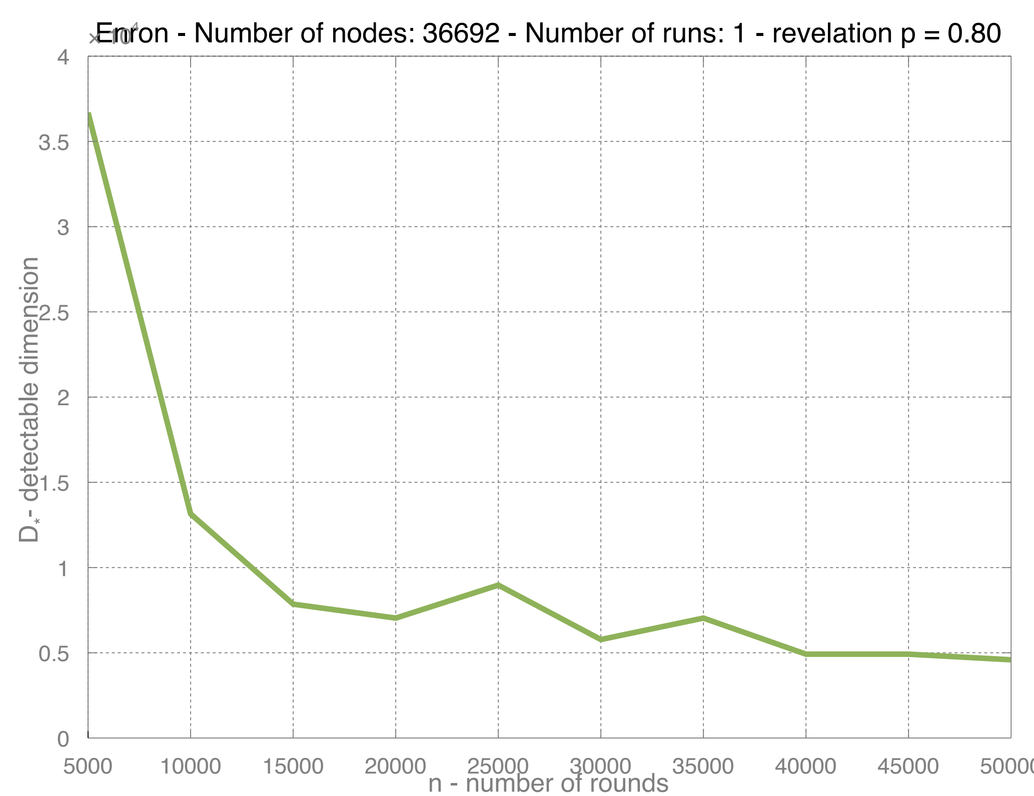

Algorithm

BARE - BAndit REvelator

Input

d: the number of nodes

n: time horizon

Initialization

$T_{k,t} \leftarrow 0$, for $\forall k \leq d$

$\widehat{r^{\circ}_{k,t}} \leftarrow 0$, for $\forall k \leq d$

$t \leftarrow 1$, $\widehat{T}_\star \leftarrow 0$, $\widehat{D}_{\star,t} \leftarrow d$, $\widehat{\sigma}_{\star,1} \leftarrow d$

Global exploration phase

while $t\left(\widehat{\sigma}_{\star,t} - 4\sqrt{d\log(dn)/t}\right) \leq \sqrt{\widehat{D}_{\star,t} n}$ **do**

Influence a node at random (choose k_t uniformly at random) and get $S_{k_t,t}$ from this node

$\widehat{r^{\circ}_{k,t+1}} \leftarrow \frac{t}{t+1}\widehat{r^{\circ}_{k,t}} + \frac{d}{t+1}S_{k_t,t}(k)$

$\widehat{\sigma}_{\star,t+1} \leftarrow \max_{k'} \sqrt{\widehat{r^{\circ}_{k',t+1}} + 8d\log(nd)/(t+1)}$

$w_{\star,t+1} \leftarrow 8\widehat{\sigma}_{\star,t+1}\sqrt{\frac{d\log(nd)}{t+1}} + \frac{24d\log(nd)}{t+1}$

$\widehat{D}_{\star,t+1} \leftarrow \left|\left\{k : \max_{k'} \widehat{r^{\circ}_{k',t+1}} - \widehat{r^{\circ}_{k,t+1}} \leq w_{\star,t+1}\right\}\right|$

$t \leftarrow t + 1$

end while

$\widehat{T}_\star \leftarrow t$.

Bandit phase

Run minimax-optimal bandit algorithm on the $\widehat{D}_{\star,\widehat{T}_\star}$ chosen nodes (e.g., Algorithm 1)

EMPIRICAL RESULTS

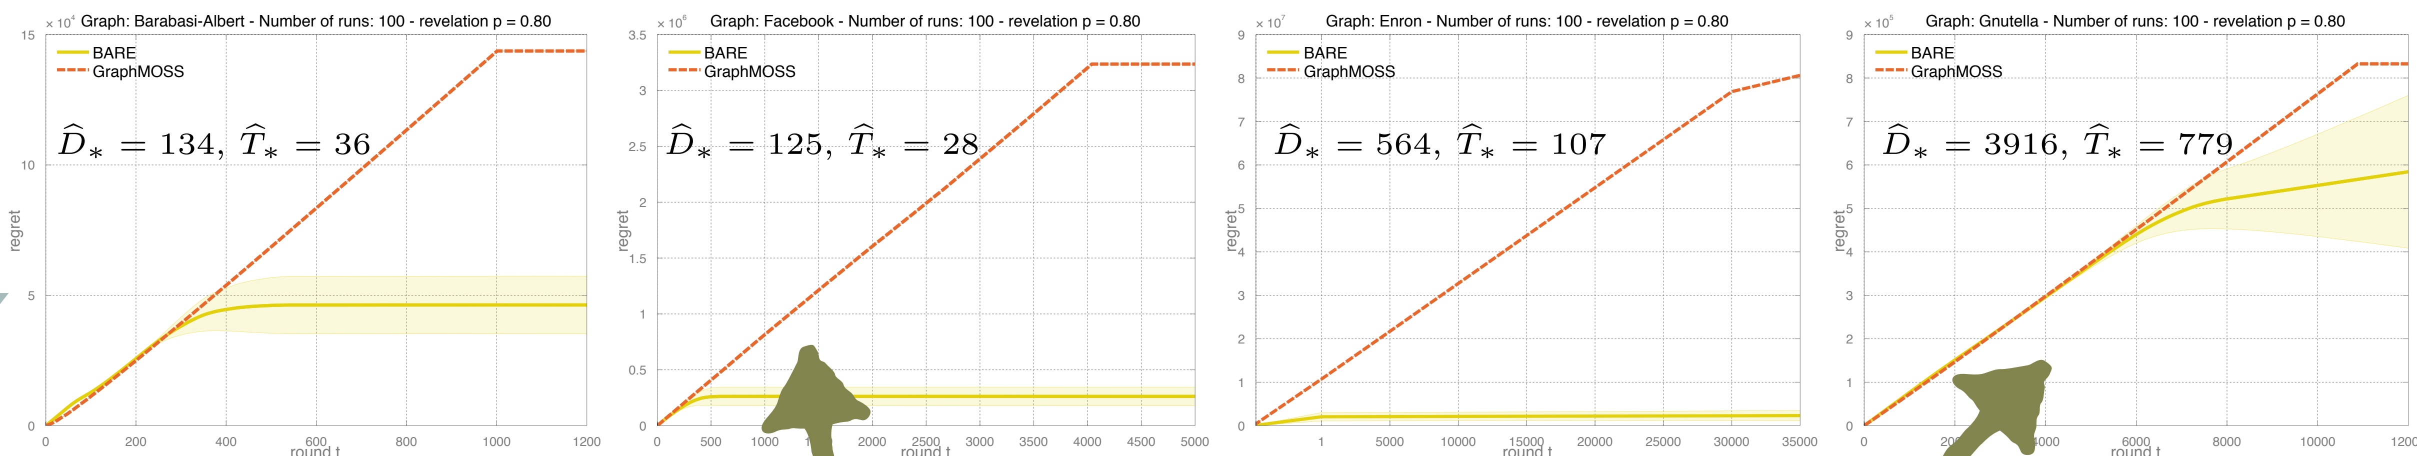

Figure 1: *Left*: Barabási-Albert. *Middle left*: Facebook. *Middle right*: Enron. *Right*: Gnutella.

Enron and Facebook **vs.** Gnutella (decentralised)

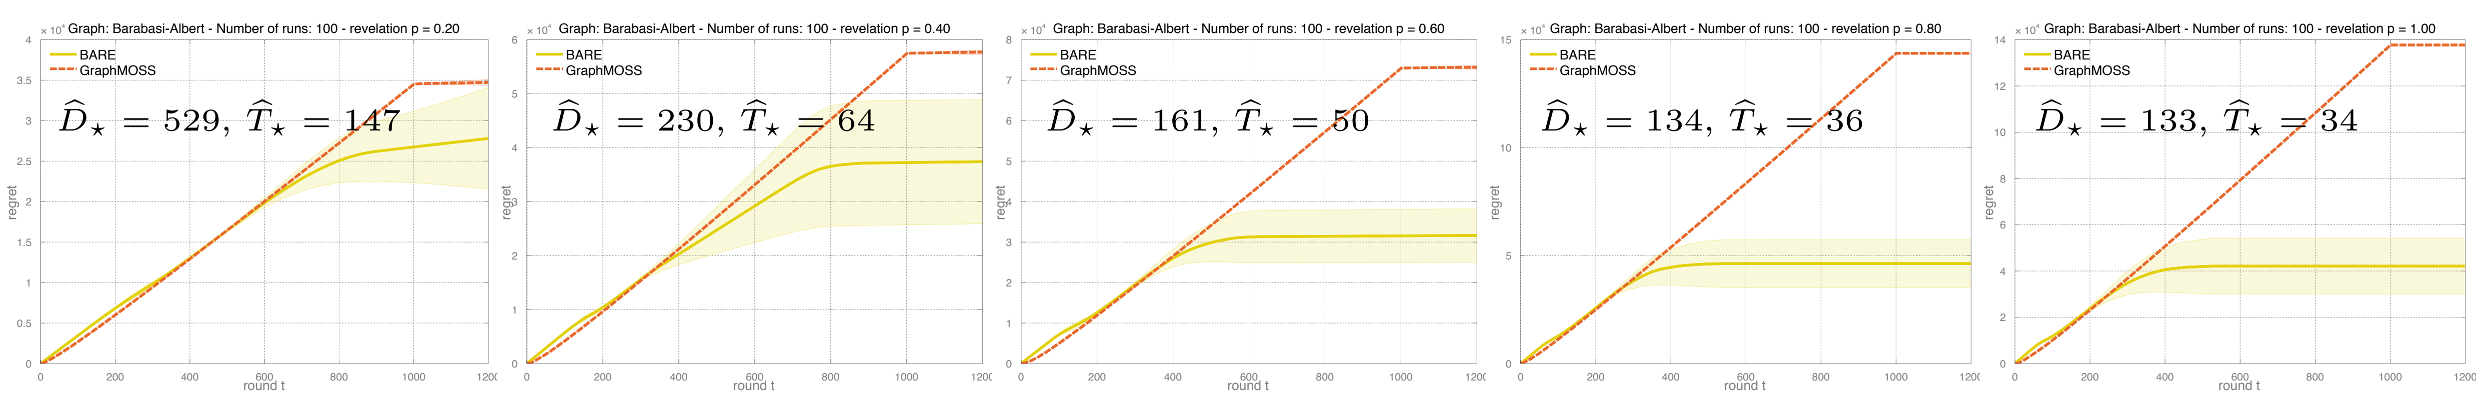

Varying a (constant) probability of influence

REVEALING BANDITS: WHAT DO YOU MEAN?

- Ignoring the structure? $\mathcal{O}(\sqrt{r_\star n d})$

  reward of the best node

- **BAndit REvelator:** 2-phase algorithm
- **global** exploration phase
  - super-efficient exploration
  - linear regret — needs to be short!
  - extracts D_* nodes
- **bandit** phase
  - uses a minimax-optimal bandit algorithm (GraphMOSS)
  - has a "square root" regret on D_* nodes
- D_* realizes the optimal trade-off !
  - different from exploration/exploitation tradeoff

Regret of BARE

$$\mathcal{O}(\sqrt{r_\star n D_\star})$$

- D_* - detectable dimension (depends on n and the structure)
  - **good case**: star-shaped graph can have $D^* = 1$
  - **bad case:** a graph with many small cliques.
  - **the worst case:** all nodes are disconnected except 2

NEXT: GLOBAL INFLUENCE MODELS

- Kempe, Kleinberg, Tárdos, 2003, 2015: **Independence Cascades**, Linear Threshold models
  - **global** and **multiple-source** models
- Different feed-back models
  - **Full bandit** (only the number of influenced nodes)
  - **Node-level semi-bandit** (identities of influenced nodes)
  - **Edge-level semi-bandit** (identities of influenced edges)
    - Wen, Kveton, Valko, Vaswani, **to appear at NIPS 2017**
    - preprint: https://arxiv.org/abs/1605.06593
    - IMLinUCB with linear parametrization of edge weights
    - Regret analysis for **general graphs, cascading model, and multiple-sources**

Online Influence Maximization under Independent Cascade Model with Semi-Bandit Feedback

Zheng Wen
Adobe Research
San Jose, CA 95110
zwen@adobe.com

Branislav Kveton
Adobe Research
San Jose, CA 95110
kveton@adobe.com

Michal Valko
Inria Lille-Nord Europe
59650 Villeneuve d'Ascq, France
michal.valko@inria.fr

Sharan Vaswani
University of British Columbia
Vancouver, B.C., Canada
sharanv@cs.ubc.ca

Abstract

We study the stochastic online problem of learning to influence in a social network with semi-bandit feedback, where we observe how users influence each other. The problem combines challenges of limited feedback, because the learning agent only observes the influenced portion of the network, and combinatorial number of actions, because the cardinality of the feasible set is exponential in the maximum number of influencers. We propose a computationally efficient UCB-like algorithm, IMLinUCB, and analyze it. Our regret bounds are polynomial in all quantities of interest; *reflect the structure of the network* and the *probabilities of influence*. Moreover, they do not depend on inherently large quantities, such as the cardinality of the action set. To the best of our knowledge, these are the first such results. IMLinUCB permits linear generalization and therefore is suitable for large-scale problems. Our experiments show that the regret of IMLinUCB scales as suggested by our upper bounds in several representative graph topologies; and based on linear generalization, IMLinUCB can significantly reduce regret of real-world influence maximization semi-bandits.

CHALLENGES AND SOLUTIONS

- **Already the offline problem is NP hard**
  - solution: **approximation/randomized algorithms**
- **Lots of edges**
  - lots of parameters to learn, if we want to scale, we need to reduce this complexity
  - solution: **linear approximation of probabilities**
- **Combinatorial size of possible seed-sets**
  - Combinatorial Bandits: IMLinUCB
- **Understanding what's going on?**
  - known analyses VERY loose (e.g., scaling with $1/p_{min}$, or only assymptotic)

seed set

seed size

$$\max_{\mathcal{S}:\,|\mathcal{S}|=K} f(\mathcal{S}, \overline{w})$$

APPROXIMATION ORACLE

- the optimal offline solution

seed size

$$\max_{\mathcal{S}:\, |\mathcal{S}|=K} f(\mathcal{S}, \overline{w})$$

- the oracle solution that is γ-optimal with probability at least α

$$\mathcal{S}^* = \texttt{ORACLE}(\mathcal{G}, K, \overline{w})$$

- γ-optimal

$$f(\mathcal{S}^*, \overline{w}) \geq \gamma f(\mathcal{S}^{\text{opt}}, \overline{w})$$

- γ-optimal with probability at least α

$$\mathbb{E}\left[f(\mathcal{S}^*, \overline{w})\right] \geq \alpha\gamma f(\mathcal{S}^{\text{opt}}, \overline{w})$$

- Our problem is a triple:

unknown to the agent

$$(\mathcal{G}, K, \overline{w})$$

topology

seed size

LINEAR GENERALIZATION

— learning the only network (weights) is VERY impractical

$$\rho \triangleq \max_{e \in \mathcal{E}} \left| \overline{w}(e) - x_e^\mathsf{T} \theta^* \right|$$

this is small

true weights

linear approximation

— by choosing the **d**imension (size of θ^*) we can reduce this complexity

— if we do not want to lose generality we set **d** to the number of edges

ALGORITHM AND PERFORMANCE MEASURE

Algorithm 1 `IMLinUCB`: Influence Maximization Linear UCB

Input: graph $\mathcal{G}$, source node set cardinality K, oracle `ORACLE`, feature vector x_e's, and algorithm parameters $\sigma, c > 0$,

Initialization: $B_0 \leftarrow 0 \in \Re^d$, $\mathbf{M}_0 \leftarrow I \in \Re^{d \times d}$

for $t = 1, 2, \ldots, n$ **do**

1. set $\overline{\theta}_{t-1} \leftarrow \sigma^{-2}\mathbf{M}_{t-1}^{-1}B_{t-1}$ and the UCBs as $U_t(e) \leftarrow \mathrm{Proj}_{[0,1]}\left(x_e^\top\overline{\theta}_{t-1} + c\sqrt{x_e^\top\mathbf{M}_{t-1}^{-1}x_e}\right)$ for all $e \in \mathcal{E}$
2. choose $\mathcal{S}_t \in \mathtt{ORACLE}(\mathcal{G}, K, U_t)$, and observe the edge-level semi-bandit feedback
3. update statistics:
   (a) initialize $\mathbf{M}_t \leftarrow \mathbf{M}_{t-1}$ and $B_t \leftarrow B_{t-1}$
   (b) for all observed edges $e \in \mathcal{E}$, update $\mathbf{M}_t \leftarrow \mathbf{M}_t + \sigma^{-2}x_e x_e^\top$ and $B_t \leftarrow B_t + x_e\mathbf{w}_t(e)$

$$R^\eta(n) = \sum_{t=1}^{n} \mathbb{E}\left[R_t^\eta\right]$$

$$R_t^\eta = f(\mathcal{S}^{\mathrm{opt}}, \mathbf{w}_t) - \frac{1}{\eta} f(\mathcal{S}_t, \mathbf{w}_t)$$

MAXIMUM OBSERVED RELEVANCE

$$N_{\mathcal{S},e} \triangleq \sum_{v \in \mathcal{V} \setminus \mathcal{S}} \mathbf{1}\left\{e \text{ is relevant to } v \text{ under } \mathcal{S}\right\} \quad \text{and} \quad P_{\mathcal{S},e} \triangleq \mathbb{P}\left(e \text{ is observed} \mid \mathcal{S}\right)$$

only depends on topology

depends on both

$$C_* \triangleq \max_{\mathcal{S}:\, |\mathcal{S}|=K} \sqrt{\sum_{e \in \mathcal{E}} N_{\mathcal{S},e}^2 P_{\mathcal{S},e}}$$

max (over) 2-norm of N weighted by P

- Worst-case upper bound:

#nodes

#edges

$$C_* \leq C_{\mathcal{G}} \triangleq \max_{\mathcal{S}:\, |\mathcal{S}|=K} \sqrt{\sum_{e \in \mathcal{E}} N_{\mathcal{S},e}^2} \leq (L-K)\sqrt{|\mathcal{E}|} = \mathcal{O}\left(L\sqrt{|\mathcal{E}|}\right) = \mathcal{O}\left(L^2\right)$$

seed size

WORST-CASE BOUNDS

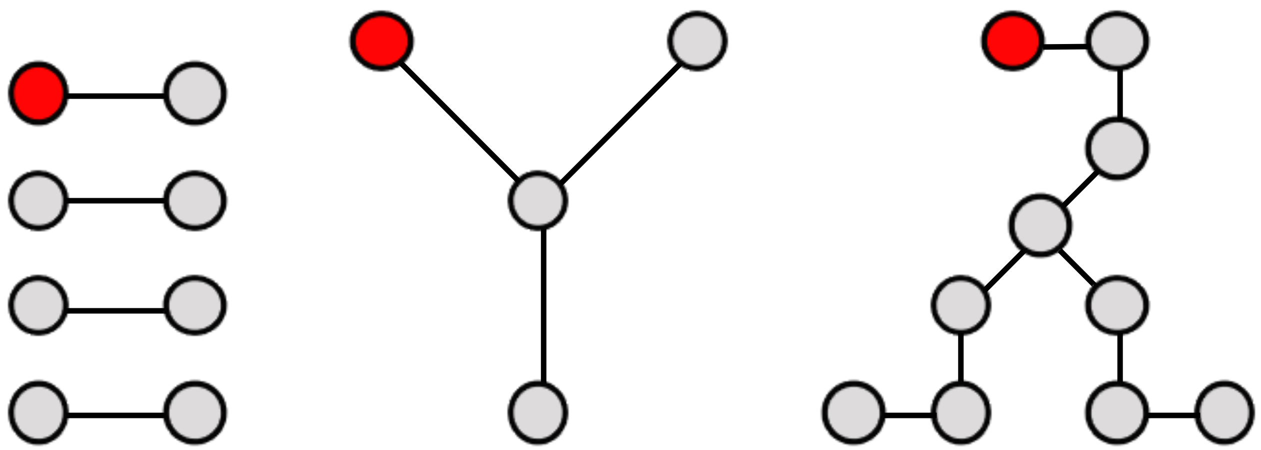

topology	$C_{\mathcal{G}}$ (worst-case C_*)	$R^{\alpha\gamma}(n)$ for general $\mathbf{X}$	$R^{\alpha\gamma}(n)$ for $\mathbf{X}=\mathbf{I}$
bar graph	$\mathcal{O}(\sqrt{K})$	$\widetilde{\mathcal{O}}\left(dK\sqrt{n}/(\alpha\gamma)\right)$	$\widetilde{\mathcal{O}}\left(L\sqrt{Kn}/(\alpha\gamma)\right)$
star graph	$\mathcal{O}(L\sqrt{K})$	$\widetilde{\mathcal{O}}\left(dL^{\frac{3}{2}}\sqrt{Kn}/(\alpha\gamma)\right)$	$\widetilde{\mathcal{O}}\left(L^2\sqrt{Kn}/(\alpha\gamma)\right)$
ray graph	$\mathcal{O}(L^{\frac{5}{4}}\sqrt{K})$	$\widetilde{\mathcal{O}}\left(dL^{\frac{7}{4}}\sqrt{Kn}/(\alpha\gamma)\right)$	$\widetilde{\mathcal{O}}\left(L^{\frac{9}{4}}\sqrt{Kn}/(\alpha\gamma)\right)$
tree graph	$\mathcal{O}(L^{\frac{3}{2}})$	$\widetilde{\mathcal{O}}\left(dL^2\sqrt{n}/(\alpha\gamma)\right)$	$\widetilde{\mathcal{O}}\left(L^{\frac{5}{2}}\sqrt{n}/(\alpha\gamma)\right)$
grid graph	$\mathcal{O}(L^{\frac{3}{2}})$	$\widetilde{\mathcal{O}}\left(dL^2\sqrt{n}/(\alpha\gamma)\right)$	$\widetilde{\mathcal{O}}\left(L^{\frac{5}{2}}\sqrt{n}/(\alpha\gamma)\right)$
complete graph	$\mathcal{O}(L^2)$	$\widetilde{\mathcal{O}}\left(dL^3\sqrt{n}/(\alpha\gamma)\right)$	$\widetilde{\mathcal{O}}\left(L^4\sqrt{n}/(\alpha\gamma)\right)$

Table 1: $C_{\mathcal{G}}$ and *worst-case* regret bounds for different graph topologies

RESULTS

$$R^{\alpha\gamma}(n) \leq \frac{2cC_*}{\alpha\gamma}\sqrt{dn|\mathcal{E}|\log_2\left(1+\frac{n|\mathcal{E}|}{d}\right)}+1 = \widetilde{\mathcal{O}}\left(dC_*\sqrt{|\mathcal{E}|n}/(\alpha\gamma)\right)$$
$$\leq \widetilde{\mathcal{O}}\left(d(L-K)|\mathcal{E}|\sqrt{n}/(\alpha\gamma)\right).$$

How good (tight) is this?

- comparison with linear bandits
- comparison with general combinatorial bandits
- (L-K) factor
- **How good is C*?**

PROOF SKETCH?

- when are our upper bounds on the estimates right?

$$\xi_{t-1} = \{|x_e^\mathsf{T}(\overline{\theta}_{\tau-1} - \theta^*)| \leq c\sqrt{x_e^\mathsf{T}\mathbf{M}_{\tau-1}^{-1}x_e}, \forall e \in \mathcal{E}, \forall \tau \leq t\}$$

- .... decomposes the regret at round t

$$\mathbb{E}[R_t^{\alpha\gamma}] \leq \mathbb{P}\left(\xi_{t-1}\right)\mathbb{E}\left[R_t^{\alpha\gamma}|\xi_{t-1}\right] + \mathbb{P}\left(\overline{\xi}_{t-1}\right)[L-K]$$

- monotonicity of f

decomposed into nodes

$$\mathbb{E}\left[R_t^{\alpha\gamma}|\xi_{t-1}\right] \leq \mathbb{E}\left[f(\mathcal{S}_t, U_t) - f(\mathcal{S}_t, \overline{w})|\xi_{t-1}\right]/(\alpha\gamma)$$

- studying second-order derivatives of f
  - monotonicity and concavity of f wrt w
  - sub-modularity of f wrt newly added edge

$$f(\mathcal{S}_t, U_t, v) - f(\mathcal{S}_t, \overline{w}, v) \leq \sum_{e \in \mathcal{E}_{\mathcal{S}_t, v}} \mathbb{E}\left[\mathbf{1}\left\{O_t(e)\right\}\left[U_t(e) - \overline{w}(e)\right]|\mathcal{H}_{t-1}, \mathcal{S}_t\right]$$

probability that node v is influences

EXPERIMENTS

- **Objective:** "Check" how good is our C*
- Tabular case, K = 1, exact comparison possible, all weights are same = ω

Star $\widetilde{\mathcal{O}}(L^2)$ vs. $\mathcal{O}(L^{2.040})$ and $\mathcal{O}(L^{2.056})$

Ray $\widetilde{\mathcal{O}}(L^{\frac{9}{4}})$ vs. $\mathcal{O}(L^{2.488})$ and $\mathcal{O}(L^{2.467})$

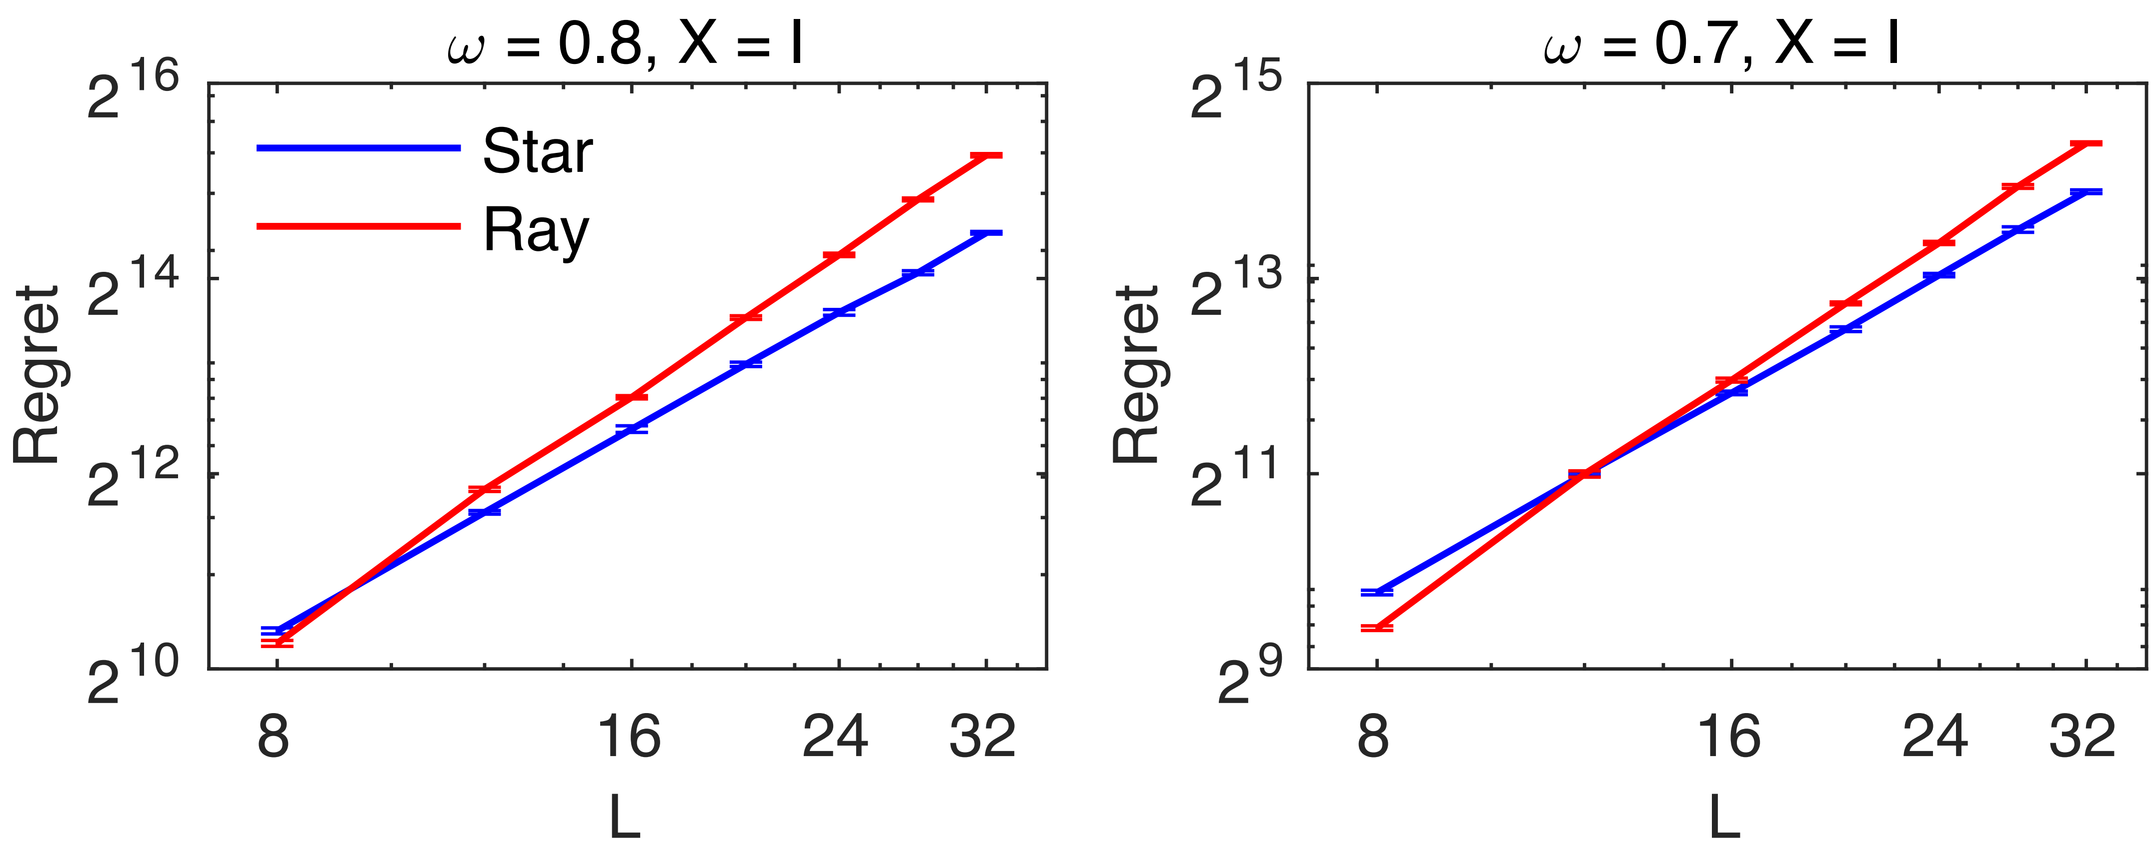

- **Conclusion:** evidence that our C* is a reasonable complexity measure

FACEBOOK EXPERIMENT

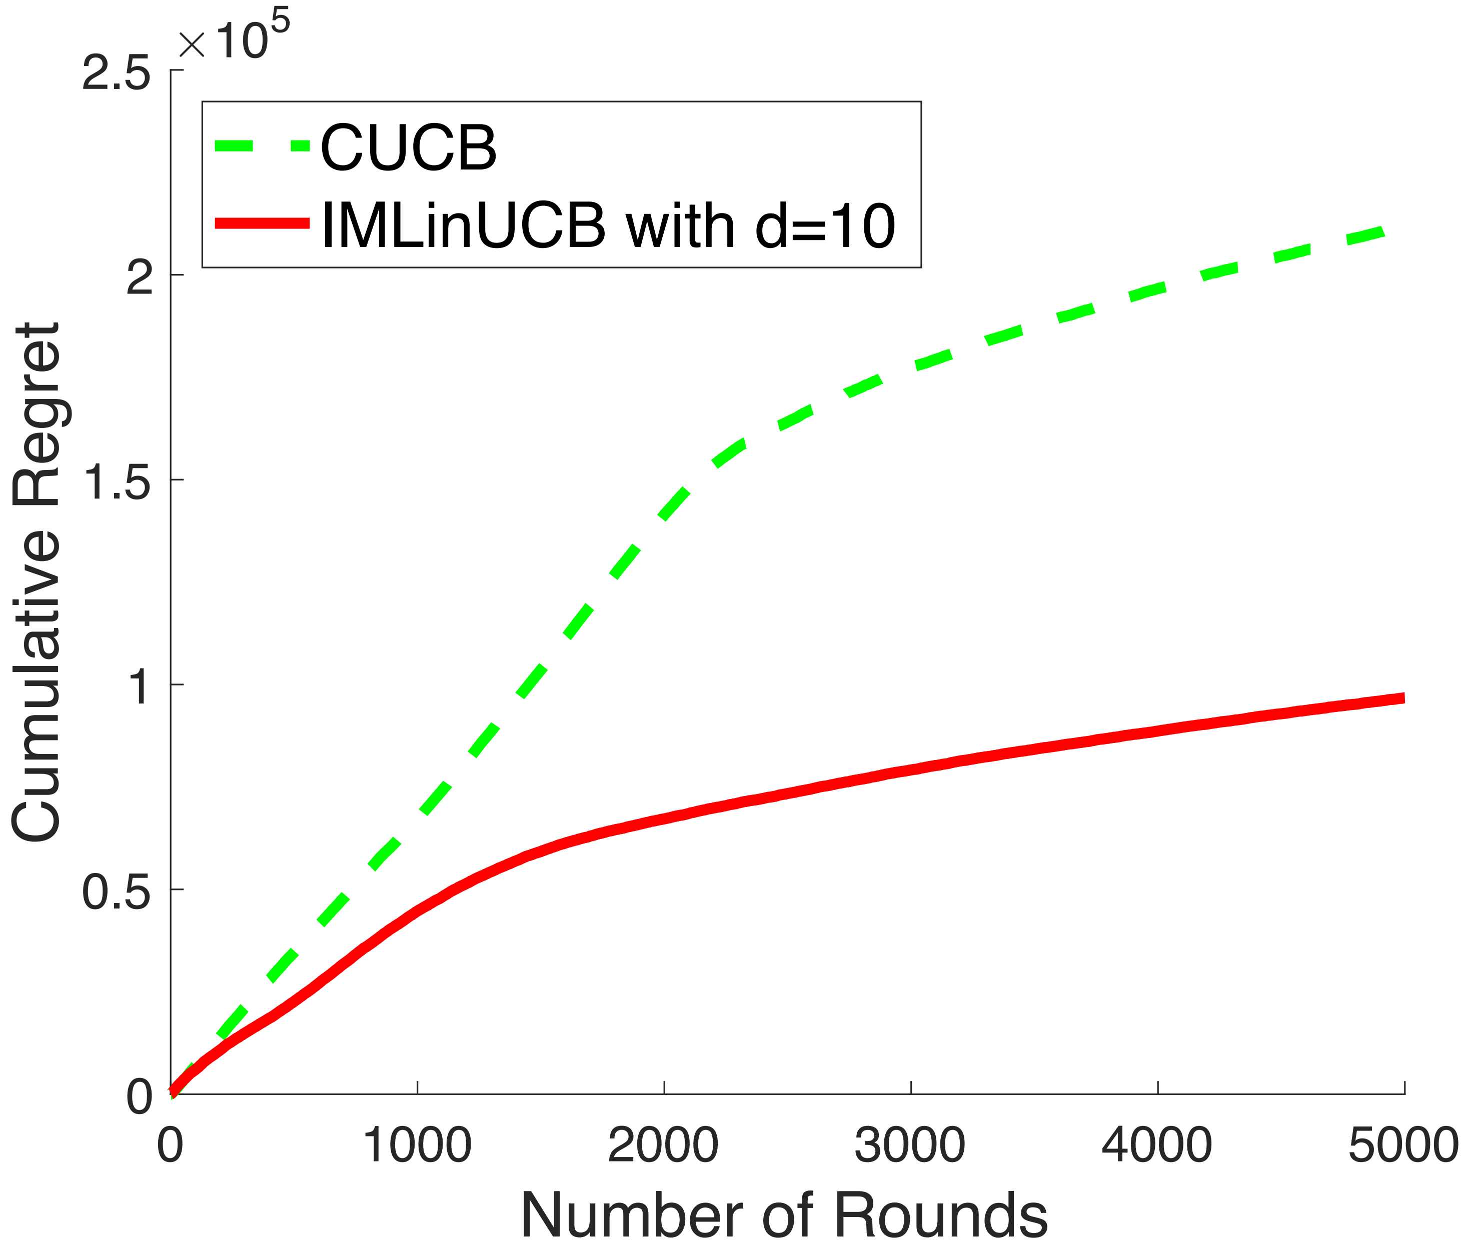

- real Facebook (a small subgraph)
- weights from U(0,0.1)
- **nodetovec** with d=10
  - imperfect
- K = 10
- CUCB with no linear generalisation

CONCLUSION AND NEXT STEPS

- **Active learning on graphs**
  - learning the graph **while** acting on it optimal
  - **difficulty of the problem** and scaling with it
  - online influence maximization
    - local model (minimax optimal algorithm)
    - global cascading model
- **What is next?**
  - dynamic/evolving graphs
  - realistic accessibility constraints

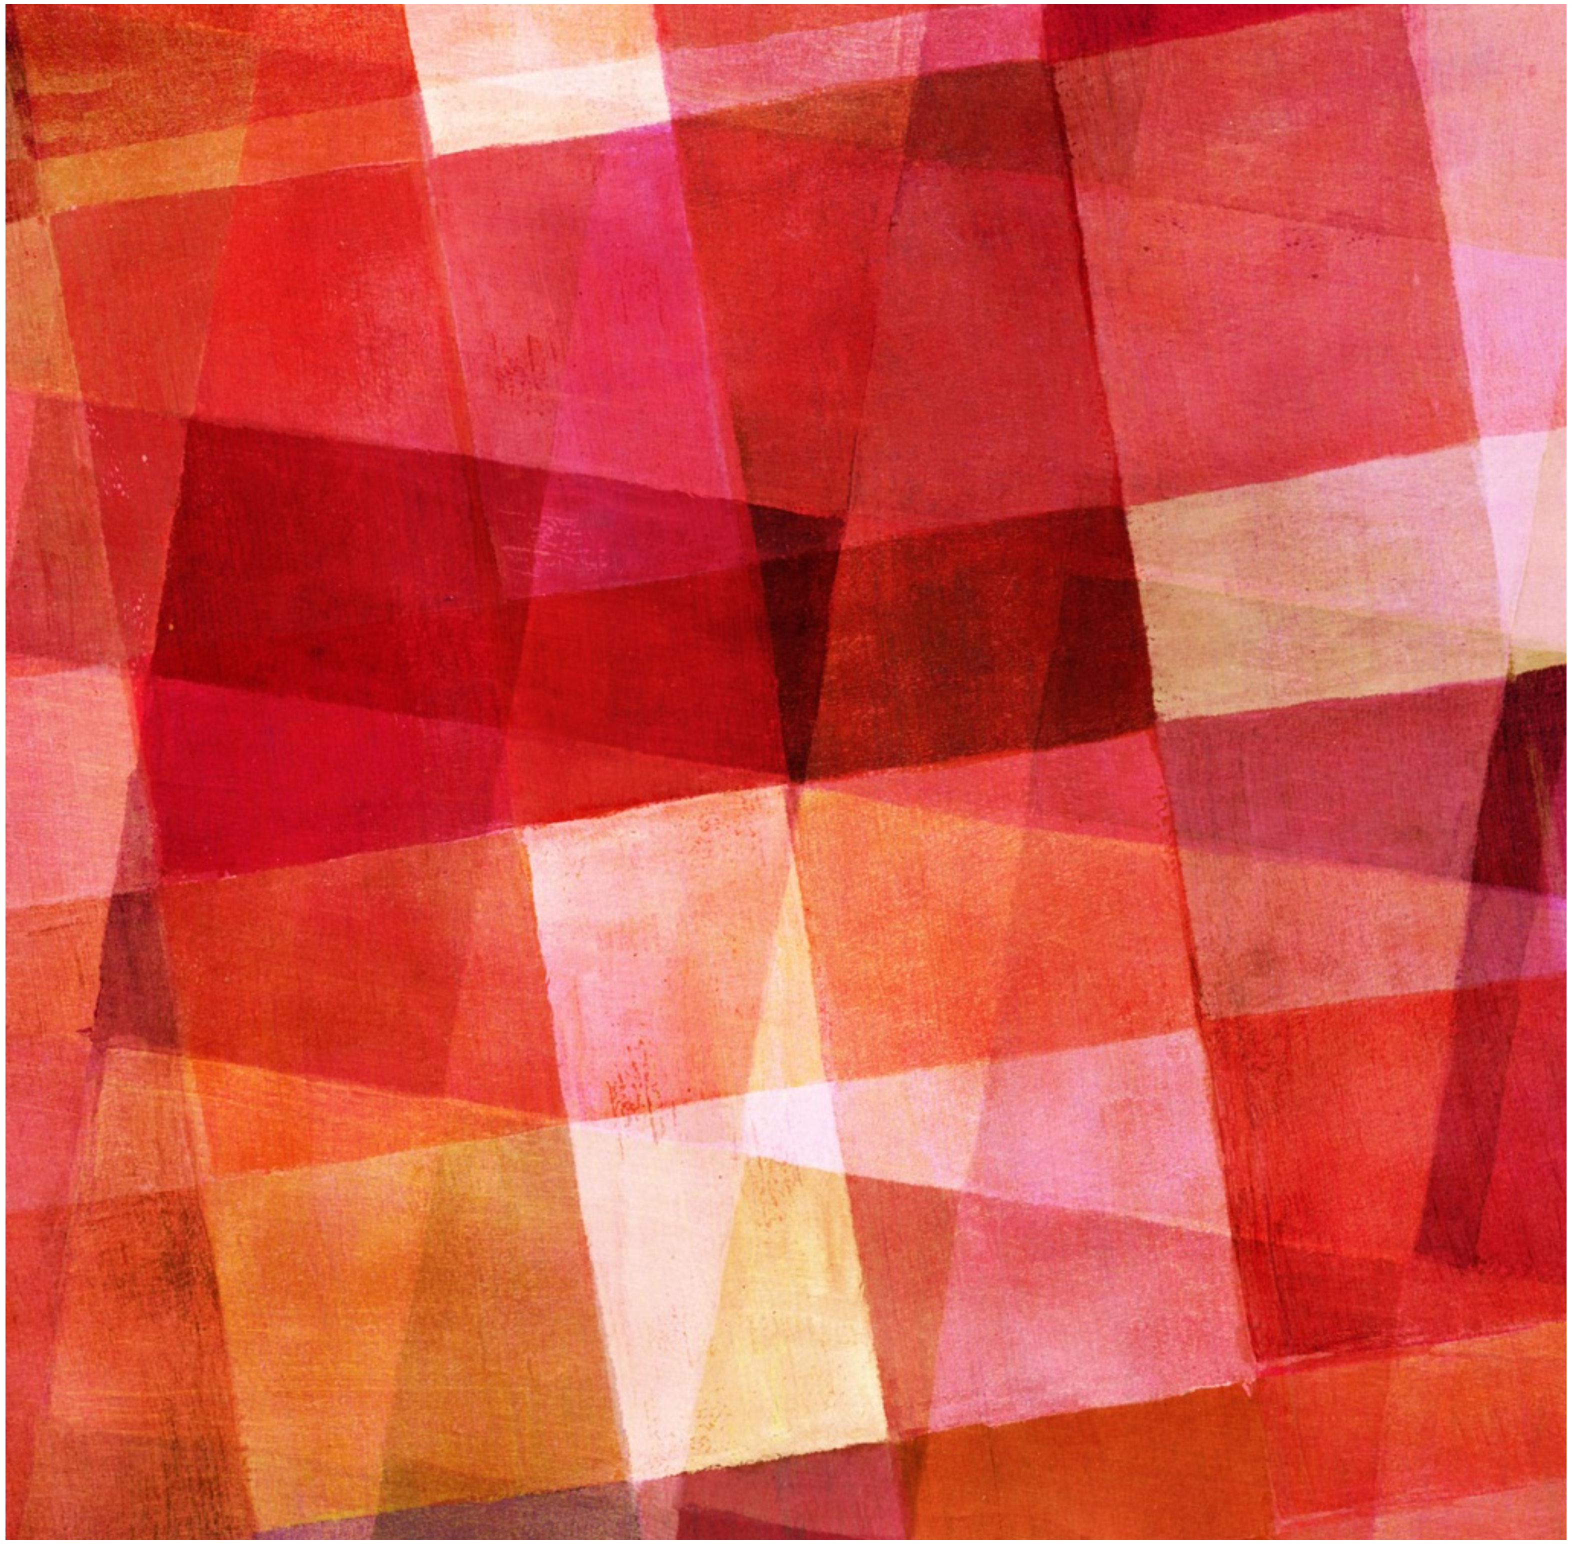

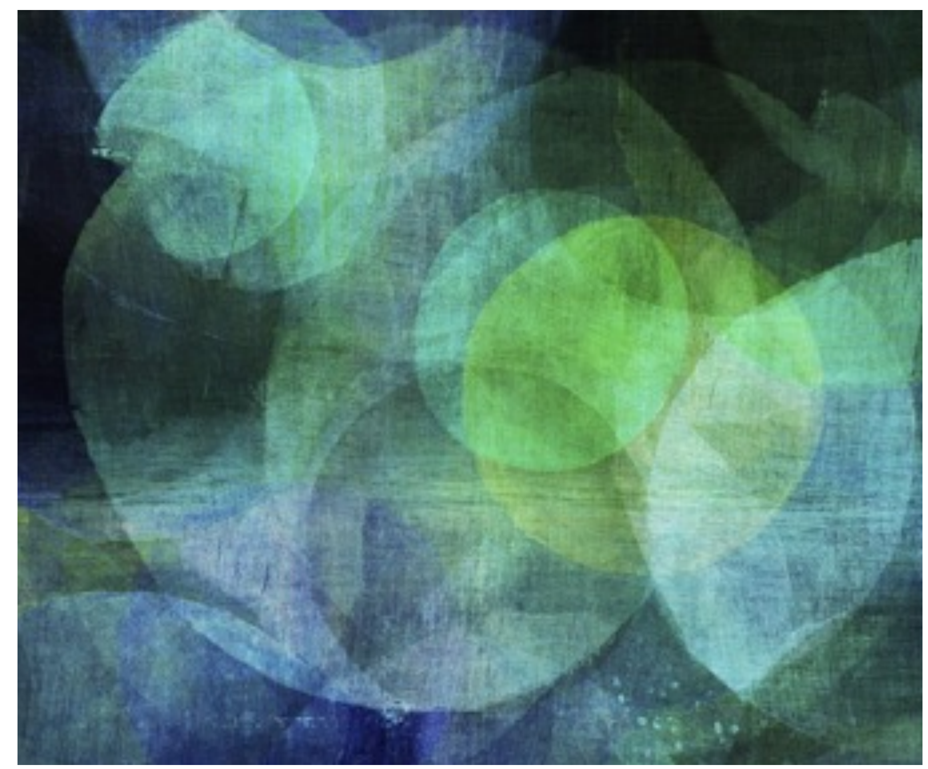

ExtraLearn

Michal Valko, SequeL, Inria Lille - Nord Europe, michal.valko@inria.fr
http://researchers.lille.inria.fr/~valko/hp/