Network Dynamics I: Bayesian Learning, Information Cascades

Computational Social Systems I (VU) (706.616)

Elisabeth Lex

ISDS, TU Graz

May 7, 2020

Motivation

- Individuals easily influenced by decisions of others, especially in social and economic situations.
- Ex.: opinions, products to buy, political positions

Motivation

- Individuals easily influenced by decisions of others, especially in social and economic situations.
- Ex.: opinions, products to buy, political positions

Today: **Why** does such influence occur and how can we **model** this?

Example: Look Up

- In the 60s, some experimenters had groups of people 1-15 look up in the sky.
- They then watched to see if anyone else looked up as well.

Example: Look Up

- In the 60s, some experimenters had groups of people 1-15 look up in the sky.
- They then watched to see if anyone else looked up as well.
- With one looking up, most ignored him.

Example: Look Up

- In the 60s, some experimenters had groups of people 1-15 look up in the sky.
- They then watched to see if anyone else looked up as well.
- With one looking up, most ignored him.
- With 5 looking up, many stopped.

Example: Look Up

- In the 60s, some experimenters had groups of people 1-15 look up in the sky.
- They then watched to see if anyone else looked up as well.
- With one looking up, most ignored him.
- With 5 looking up, many stopped.
- With 15, 45% stopped and kept looking up.

Example: Restaurant

- You go to a unfamiliar city such as San Francisco.
- Where do you decide to eat?
- Restaurant A and B both look nice on the outside and have a similar menu.
- A has a small crowd.
- B is empty.
- Which do you go to?

Example: Restaurant

- You go to a unfamiliar city such as San Francisco.
- Where do you decide to eat?
- Restaurant A and B both look nice on the outside and have a similar menu.
- A has a small crowd.
- B is empty.
- Which do you go to?
- Probably A.

Example: Restaurant

- You go to a unfamiliar city such as San Francisco.
- Where do you decide to eat?
- Restaurant A and B both look nice on the outside and have a similar menu.
- A has a small crowd.
- B is empty.
- Which do you go to?
- Probably A.
- But: What if from the outside B looks slightly better?
- How about if a colleague said he heard B was good?

Observations

- In all these cases, decisions are made sequentially.
- People make decisions based on inferences from what earlier people have done
- Individuals may imitate behavior of others but not mindless

Observations

- In all these cases, decisions are made sequentially.
- People make decisions based on inferences from what earlier people have done
- Individuals may imitate behavior of others but not mindless
- Sometimes it is rational for an individual to follow the crowd even if the individual's own information suggests an alternative choice

Reasons for Imitation

- Direct-Benefit Effects: Actions of others affect you *directly*
- E.g. Becoming part of a social networking site - choosing an option that has a large user population
- Informational Effects: Actions of others affect you *indirectly* by changing your information

How to model this?

- Decision-based Models: adopt new behaviors if k others do it
- Probabilistic Models: adopt a behavior with some probability from neighbors in the network

Herding & Information Cascades

Herding - Basic Ingredients

- A **decision** needs to be made
- People make the decision **sequentially**
- Each person has some **private information** that helps with the decision
- This private information **cannot be directly observed** but one can see what people **do**
- This way, inferences about their private information can be made

A Simple Herding Example (1/2)

Large group of students - participants

- Consider an urn with 3 balls that is either:
  - Majority-blue: blue, blue, red or
  - Majority-red: red, red, blue
- Student sequentially guess whether urn is majority-blue or majority-red

A Simple Herding Example (2/2)

Experiment:

- One by one, each student:
  - Draws a ball
  - Privately looks at its color and puts it back
  - Publicly announces her guess
- If guess correct - receives monetary reward

What happens? (1/2)

Example:

- 1st student: draws blue
  - Guess: urn is majority-blue
- 2nd student: draws red
  - Guess: urn is majority-red
- 3rd person: draws blue
  - Guess: urn is majority-blue Why?
  - Student goes with her own color as students 1 & 2 made different guesses

What happens? (1/2)

Another example:

- 1st student: draws red
  - Guess: urn is majority-red
- 2nd student: draws red
  - Guess: urn is majority-red
- 3rd person: draws blue
  - Guess: urn is majority-red Why?
  - Go with guesses of 1st & 2nd student as both guessed urn is majority-red

What happens? (1/2)

Another example:

- 1st student: draws red
  - Guess: urn is majority-red
- 2nd student: draws red
  - Guess: urn is majority-red
- 3rd person: draws blue
  - Guess: urn is majority-red Why?
  - Go with guesses of 1st & 2nd student as both guessed urn is majority-red

Information Cascades

Definition

An information cascade develops when people abandon their own information in favor of inferences based on earlier people's actions

What happens? (2/2)

4th student and onward:

- Say first 2 guesses were both blue, 3rd person will also guess blue, regardless of what she saw
- 4th heard blue 3 times in a row, but knows that 3rd guess conveys no information
- 4th in the same situation as 3rd - should also guess blue
- Continues with all subsequent students since everyone's best strategy is to rely on limited genuine information available

Modeling this type of reasoning: Bayes's Rule (1/2)

- Mathematical model for when information cascades occur
- Based on computing probabilities of events (e.g. event is "the urn is majority-blue")
- Whether an event (not) occurs is the result of certain random outcomes (e.g. which urn was placed in the room)

Modeling this type of reasoning: Bayes's Rule (2/2)

We need to estimate the **conditional probability** of event A given that event B has occurred:

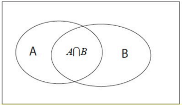

- The fraction of the area of region B occupied by the joint event $A \cap B$:

$$P[A|B] = \frac{P[A \cap B]}{P[B]} \tag{1}$$

- We can follow the Bayes' rule:

$$P[A|B] = \frac{P[A] * P[B|A]}{P[B]} \tag{2}$$

where $P[A]$ is the prior probability of A and $P[A|B]$ is the posterior probability of A given B

Our Example with Bayes's Rule (1/3)

Each person tries to estimate conditional probability that urn is majority-blue or majority-red given what she has seen and heard
She should guess

- $P_r[\text{majority-blue} | \text{what she has seen and heard}] > \frac{1}{2}]$
- majority-red otherwise

If both conditional probabilities 0.5 - doesn't matter what she guesses

Our Example with Bayes' Rule (2/3)

- Prior probabilities $P_r[\text{majority-blue}]$ and $P_r[\text{majority-red}] = \frac{1}{2}$
- Probabilities for the 2 urns $P_r[\text{blue}|\text{majority-blue}] = P_r[\text{red}|\text{majority-red}] = \frac{2}{3}$

Let's assume that 1st person draws blue ball

- $P_r[\text{majority-blue}|\text{blue}] = \frac{P_r[\text{majority-blue}] * P_r[\text{blue}|\text{majority-blue}]}{P_r[\text{blue}]} = \frac{2}{3}$

Ergo: Conditional probability greater than $\frac{1}{2}$, 1st person should guess urn is majority-blue

Our Example with Bayes' Rule (3/3)

Calculation for 3rd person:

- $P_r[\text{majority-blue}|\text{blue, blue, red}] = \frac{P_r[\text{majority-blue}] * P_r[\text{blue, blue, red}|\text{majority-blue}]}{P_r[\text{blue, blue, red}]} = \frac{2}{3}$

Ergo: Conditional probability greater than $\frac{1}{2}$, 3rd person should guess urn is majority-blue

Discussion

- Cascade can occur easily
- Can lead to non-optimal outcomes
- In our example, with prob $1/3x1/3 = 1/9$, the first two would see the wrong color, from then on, the whole population would guess wrong
- Cascades fragile despite their potential to produce long runs of conformity
  - Suppose, first 2 guesses are majority-blue
  - People 100 and 101 draw red and cheat, i.e. they show the drawn balls
  - Person 102 has then 4 pieces of honest information - she should guess based on her own color
  - Cascade is broken!

Summary

- People initially rely on own private information
- Observe what others decide
- If number of acceptances and rejections of decision is ≥ 2, people follow majority decision
- Over time, population ignores own information and follow the crowd - while still being fully rational

Problems of cascades

- Cascades can be wrong
- Cascades can be based on very little information
- Cascades are fragile

Problems of cascades

- Cascades can be wrong
- Cascades can be based on very little information
- Cascades are fragile

Careful if you want to follow the crowd :)

Applications and Practical Value

- Early adopters in Online Marketing: attempting to initiate a buying cascade

- Collaboration and consensus building: can be wise to make collaborators reach partial conclusions before entering phase of collaboration (e.g. hiring committees)

Models of influence

Models of influence

- Decision-based models: purchase decisions, adoption of innovation, joining riots
- Epidemic models: virus, infection, rumors, news

Diffusion of innovation

Diminishing returns vs threshold (critical mass)

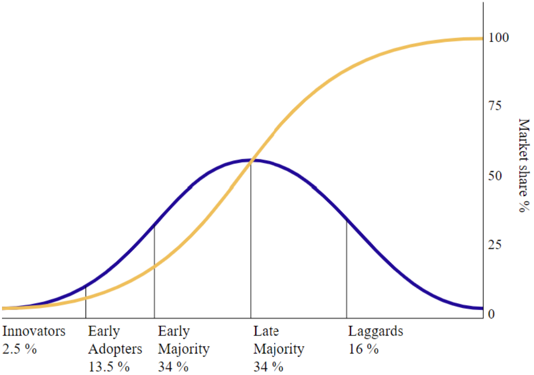

By Rogers Everett - Based on Rogers, E. (1962) Diffusion of innovations. Free Press, London, NY, USA., Public Domain, https://commons.wikimedia.org/w/index.php?curid=18525407

Underlying mechanisms of diffusion: Influence response

Diminishing returns vs threshold (critical mass)

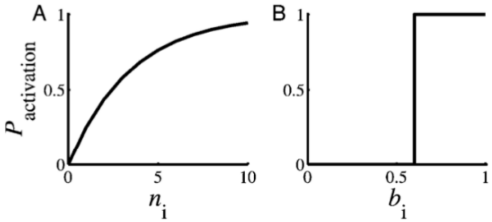

$P(n) = 1 - (1 - p)^n$ $\qquad$ $P(b) = \delta(b > b_0)$

Probabilities that person adopts behavior (becomes activated) based on behavior of neighbors

Influence response

Two Models: Probabilistic vs Threshold

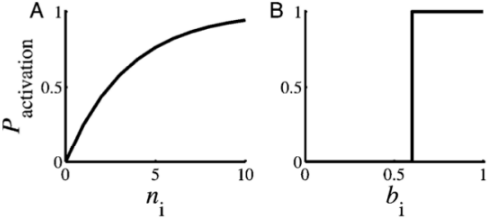

$$P(n) = 1 - (1 - p)^n \qquad P(b) = \delta(b > b_0)$$

Decision-based models

- Independent Cascade Model (diminishing returns)
- Linear Threshold Model (critical mass; threshold)

Goldenberg (2001): Cascades model
Granovetter (1978): Threshold models

Independent Cascade Model

Independent Cascade Model

- Initial set of active nodes S_0 (e.g. 2 activists)
- Discrete time steps t
- At each time step, active node v can activate connected neighbor w with probability $p_{v,w}$ - **single chance!** (e.g. p = 0.5)
- If activation succeeds, w will become active at $t + 1$
- Runs until no more activations possible
- Once activated nodes cannot be deactivated

Example

Joining groups (Backstrom, 2006)

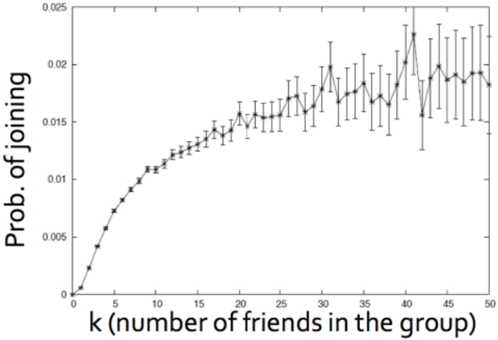

Linear Threshold Model

Core principles

- Group of people, each to make a decision
- Binary mutually exclusive decision (e.g. adopt/reject)
- Each person has personal preference; decision threshold

Example: Organize a revolt

- You live an in oppressive society
- You know of a rebel group that plans to fight against the empire tomorrow
- If a lot of people show up, the empire will fall
- If only a few people show up, the rebels will be arrested and it would have been better had everyone stayed at home

Organize a revolt: Model

- Personal threshold k: "I will show up if am sure at least k people in total (including myself) will show up"
- Each node only knows the thresholds and attitudes of all their direct friends.

Can we predict if a revolt can happen based on the network structure?

Ex.: Which network can have a revolt?

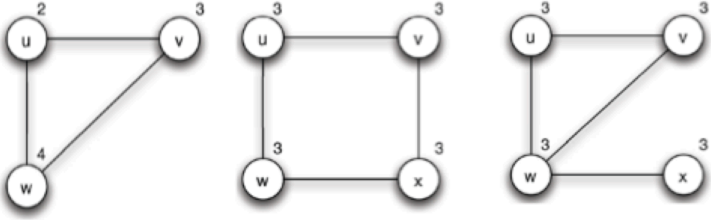

The last one. Why?

Because nodes u,v, and w share common knowledge about the thresholds of their friends

Linear Threshold Model

- initial set of active nodes as seeds
- discrete time steps t
- threshold θ for each node selected uniformly at random from interval between 0 and 1
- each edge has non-negative weight w
- at each time step, inactive node activated ifsum of the weights of the edges with active neighbors exceeds θ $\sum_{j \in N_i, \text{j is active}} w_{j,i} \geq \theta_i$

Impact of seed selection

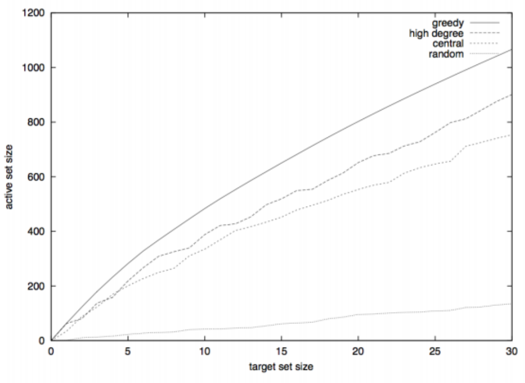

Kempe, Kleinberg, Tardos (2003)

References

- Granovetter (1978) Threshold Models of Collective Behavior. Journal of Sociology 83(6):1420-1443
- Goldenberg et al. (2001) Talk of the network: a complex systems look at the underlying process of word-of-mouth
- Kempe, Kleinberg, Tardos (2003) Maximizing the Spread of Influence through a Social Network
- Backstrom et al. (2006) Group Formation in Large Social Networks: Membership, Growth and Evolution

Epidemic models

Motivation

- Models of influence or disease spreading
- Randomly occur as a result of social contact - no decision making involved
- Ex.: Catching a disease with some probability from each active neighbor in the network
- Underlying networks affect spread
- Involves some randomness and given the initial conditions, different outcomes possible

Epidemics

- Model epidemic spread as random process on graph and study its properties
- Example questions:
  - Growth of infected populartion?
  - How much of the network will the epidemic take over?
- Underlying networks affect spread
- Involves some randomness and given the initial conditions, different outcomes possible

Epidemic Models

- SIR: Susceptible-Infective-Recovered (e.g., chickenpox)
- SIS: Susceptible-Infective-Susceptible (e.g., flu)

The SIR model

SIR Model

- Model to study propagation of diseases, epidemics (Kermack and McKendrick, 1927)
- Can also be used to study propagation of information, rumors, falsehoods,..
- Suitable to model infections that result in immunity (or death) such as chickenpox

SIR: Susceptible-Infective-Recovered Model (1/2)

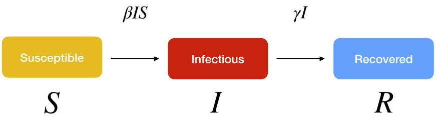

- Each node of a population N can be in one of the following states:
  - Susceptible: node healthy but not immune
  - Infective: has the virus and can actively propagate it
  - Recovered: had the virus and is no longer active (immune or dead)
- Infection rate β: how much virus can be transmitted through exposure
- Recovery rate γ: probability of node to recover at time step

SIR Model: Algorithm

- Initially: a fraction of nodes is in state I, all other nodes in the state S
- Each node in state I remains infectious for period of time t_I
- We iterate over a period of time
- We take a node from the network
- Say the node is e.g. infected, it can
  - recover and enter state R
  - infect neighbors in state S
  - stay infected for a period of time based on infection rate β - after that period, node enters state R
- Nodes with state R cannot be infected again and stay in R

SIR Model: Parameters

- Changes in the population by natural births and deaths not considered
- Reason: model assumes period of virus much shorter than the human lifetime
- Via β and γ, we can estimate two important values:
  - Average days to recover $D = \frac{1}{\gamma}$
  - Base reproduction $R_0 = \frac{\beta}{\gamma}$ denotes average number of people infected from one other person
    - $R_0 < 1$ infection dies out
    - $R_0 = 1$ infection becomes endemic
    - $R_0 > 1$ infection becomes pandemic (exponential growth)

Example

- Assume new disease
- Probability that one infected person infects healthy person 20%
- Say, one infected meets 5 people per day, who are infected with 20% prob. Thus: $\beta = 1$ $(20\% \cdot 5 = 1)$
- Say $D = 7$. Since $D = \frac{1}{\gamma}$ $R_0 = \beta \cdot D = 7$

Using the SIR model

- Aim: get number of people in S, I, R using β, γ, N
- Done by describing change per day of S, I, R
  - $S'(t) = -\beta \cdot I(t) \cdot \frac{S(t)}{N}$
  - $I'(t) = \beta \cdot I(t) \cdot \frac{S(t)}{N} - \gamma \cdot I(t)$
  - $R'(t) = \gamma \cdot I(t)$

Example: SIR

https://www.public.asu.edu/~hnesse/classes/sir.html

The SIS model

SIS: Susceptible-Infective-Susceptible Model

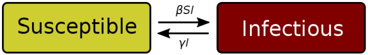

- Each node of a population N can be in one of the following states:
  - Susceptible: node healthy but not immune
  - Infective: has the virus and can actively propagate it
- Cured nodes become susceptible again

SIS model equations

- Infection rate β; recovery rate $\gamma = \frac{1}{D}$
- $S'(t) = -\beta S(t)I(t) + \gamma I(t)$
- $I'(t) = \beta S(t)I(t) - \gamma I(t)$

thus, one infected causes βSI infections per time unit

Summary

- Herding and Information Cascades, Bayes Rule as model
- Models of influence: Independent Cascade Model, Linear Threshold Model
- Epidemic models: SIR, SIS

Thanks for your attention

elisabeth.lex@tugraz.at

Slides use figures from Chapter 16 and Chapter 19 of Networks, Crowds and Markets by Easley and Kleinberg (2010)
http://www.cs.cornell.edu/home/kleinberg/networks-book/
http://web.stanford.edu/class/cs224w/handouts.html
http://snap.stanford.edu/na09/11-viral-annot.pdf