# Social Networks 

Alexander Schönhuth

Bielefeld University
June 30, 2022

## Learning Goals Today / Overview

- Intro: Social Networks are Graphs
- How to Cluster Social Networks into Groups
- Non-overlapping communities: the Girvan-Newman Algorithm
- Overlapping communities: the Graph Affiliation Model
- Direct Discovery of Overlapping Communities


## Social Networks as Graphs

## Social Networks: Introduction

## BASIC EXAMPLES

- Facebook, Twitter, Google+


## Defining Properties

- Collection of entities participating in network
- Usually people, but other entities conceivable
- There is a relationship between the entities
- Being friends is frequent relationship
- Relationship can be of 0-1 type, or weighted
- Assumption of nonrandomness or locality
- Hard to formalize, intuition is that relationships tend to cluster
- If entity $A$ is related with both B and C, B and C are related with larger probability


## Social Network Graphs: Entities and Relationships



Adopted from mmds.org

- Entities: Nodes A to G
- Relationships: Represented by edges between nodes
- Example: A is "friends" with B and C


## Social Network Graphs: Locality



Adopted from mmds.org

- Locality: / $\binom{7}{2}$
- There are 9 out of 21 possible edges: $\frac{9}{21}=0.429$
- Given nodes $X, Y, Z$ such that there are edges $(X, Y),(Y, Z)$, random occurrence of $(X, Z)$ is $\frac{7}{19}=0.368$
- However, across all pairs of existing edges $(X, Y),(Y, Z)$, probability that $(X, Z)$ exists is $\frac{9}{16}=0.563$
Network exhibits locality


## Social Networks: Examples

- Telephone Networks:
- Nodes are phone numbers, edges exist if one number called another
- Edge weights: Number of calls (within certain period of time)
- Communities: Groups of friends, members of a club, people working at same company
- Email Networks:
- Nodes are email addresses, edges indicate exchange of emails
- Edge directionality may matter, so graph with directed edges
- Communities: Similar to telephone networks


## Social Networks: Examples

- Collaboration Networks:
- Nodes e.g. represent authors, edges indicate working on same document
- Alternatively: nodes represent documents, edges indicate that identical author contributed
- Communities: Groups interested in / working on same subjects; documents sharing related content
- Other:
- Information networks: Documents, web graphs, patents
- Infrastructure networks: Roads, planes, water pipes, power grids
- Biological networks: Genes, proteins, drugs
- Product co-purchasing networks: E.g. Groupon


## Several Types of Nodes



Adopted from mmds.org
Examples

- Figure: Users (U) put tags (T) on web pages (W): tri-partite network
- Put documents and authors into one bi-partite network


## Clustering Social Networks

## Clustering Social Networks: Introduction

- An important aspect of social networks are communities
- Communities reveal themselves as groups of nodes that share unusually many edges
- Clustering social networks relates to the discovery of such communities


## COMMUNITIES



Differently Colored Communities in Social Network

## Clustered Network



Adopted from mmds.org

## Distance Measures in Social Networks

- Standard clustering techniques work with distance measures
- Distance measures are not obvious to define in social networks
- Let $x, y \in V$ be two nodes in a social network $G=(V, E)$. The measure

$$
d(x, y)= \begin{cases}0 & (x, y) \in E \\ 1 & (x, y) \notin E\end{cases}
$$

violates the triangle inequality, hence is no distance measure

- Exchanging 0 with 1 , and 1 with $\infty$ does not help
- Other binary-valued measures (e.g. 1 and 1.5) agree with triangle inequality
- But: Additional issues apply


## Social Networks: Clustering Issues



Communities: A-B-C and D-E-F-G
Adopted from mmds.org

- Hierarchical Clustering: Randomly picks closest nodes/clusters
- Distance between clusters: distance between closest points
- As soon as clusters are joined on B and D, clusters not as desired
- Summary: Standard clustering techniques difficult/impossible to sensibly implement


## Betweenness

Idea: Identify edges that are least likely to be within community
Definition [BETWEENNESS]
The betweenness of an edge $(a, b)$ is


- the number of pairs of nodes $(x, y)$ such that $(a, b)$ makes part of the shortest path leading from $x$ to $y$
- If for $(x, y)$ there are several shortest paths, $(a, b)$ is credited the fraction of shortest paths leading through $(a, b)$ when computing its betweenness


## Betweenness



Telephone network:
Links between communities have great betweenness

## Adopted from mmds.org

## Explanation

- High betweenness means that $(a, b)$ is a bottleneck for shortest paths
- If nodes $(a, b)$ lie within community, there are too many options for shortest paths to circumvent $(a, b)$ (so $(a, b)$ gets credited only small fractions)


## Betweenness: Example



Adopted from mmds.org

- $(B, D)$ has the greatest betweenness, 12
- It is on any shortest path between $A, B, C$ and $D, E, F, G$
- $(D, F)$ has betweenness 4
- It lies on all shortest paths between $A, B, C, D$ and $F$


## The Girvan-Newman Algorithm

Calculating Betweenness

## Algorithmic Principle

- Visit each node $X$ once
- Compute shortest paths from $X$ to any other node $Y$
- To visit nodes $Y$ from $X$, perform breadth-first search (BFS)


Social Network; consider BFS from $E$ Adopted from mmds.org

## The Girvan-Newman Algorithm

Calculating Betweenness

## Algorithmic Principle

- Visit each node $X$ once
- Compute shortest paths from $X$ to any other node $Y$
- To visit nodes $Y$ from $X$, perform breadth-first search (BFS)


BFS starting from $E$ on social network from slide before Adopted from mmds.org

## The Girvan-Newman Algorithm

Calculating Betweenness


Intuition / Notation

- Length of shortest path from $X$ to $Y$ : level of BFS starting at $X$
- Edges within BFS level cannot be part of shortest paths from $X$
- Edges between different levels are referred to as DAG (directed acyclic graph) edges
- DAG edges are on at least one shortest path leaving from $X$


## The Girvan-Newman Algorithm

Calculating Betweenness

Level 1


Adopted from mmds.org

## Example Notation

- Solid edges $=$ DAG edges: e.g. $(D, B),(E, F)$
- Dashed edges = within level: e.g. $(D, F),(A, C)$
- For DAG edge $(Y, Z)$ where $Y$ is closer to root $X$ than $Z$ :
- $Y$ is said to be the parent
- Z is said to be the child


## The Girvan-Newman Algorithm <br> Calculating Betweenness

Two Stages

- Labeling: For each node, assign number of shortest paths from root to that node
- Proceed from root to leaves in BFS order
- Crediting: For each edge, compute contribution of shortest paths from root to betweenness of that edge
- Need to compute credits for nodes as well
- Proceed from leaves to root, bottom-up


## The Girvan-Newman Algorithm

Calculating Betweenness

Level 1

Level 3


Adopted from mmds.org

## Labeling NODES

- Label each node by the number of shortest path to the root
- Start by labeling the root with 1
- Top-down, label each node by the sum of labels of each parents


## The Girvan-Newman Algorithm

Calculating Betweenness

EXAMPLE LABELING

- Label the root $E$ with 1
- Level 1: Each D and F have only E as parent; label both with 1
- Level 2:
- B has only $D$ as parent, label with 1
- G has parents $D$ and $F$, label with 2
- Level 3: Both $A, C$ have only $B$ as parent, so both are labeled with 1


## The Girvan-Newman Algorithm

Calculating Betweenness

## Crediting Nodes

- Compute fraction of shortest paths from root passing through node
- Credit each leaf with 1
- If several shortest paths run to leaf, fractions add up to 1
- Each non-leaf node $v$ gets credit

$$
\begin{equation*}
1+\sum_{e \in \mathcal{D}(v)} c(e) \tag{1}
\end{equation*}
$$

where $\mathcal{D}(v)$ are the DAG edges leaving from $v$, and $c(e)$ is the credit of an edge $e$

How to credit edges?

## The Girvan-Newman Algorithm

Calculating Betweenness
Crediting Edges

- Let $u_{j}, j=1, \ldots, k$ be the parents of $w$; so $\left(u_{j}, w\right)$ are the DAG edges entering $w$
- Let $N_{j}, j=1, \ldots, k$ be the number of shortest paths from root running through edges $\left(u_{j}, w\right)$
- Recall: $N_{j}$ agrees with the label of $u_{j}$, the number of shortest paths from root to $u_{j} \ldots$
- ... because every shortest path from root to $u_{j}$ is a shortest path from root to $w$
- Let $c(w)$ be the credit of $w$
- We compute the credit of $\left(u_{i}, w\right)$ as

$$
\begin{equation*}
c\left(u_{i}, w\right):=c(w) \times \frac{N_{i}}{\sum_{j=1}^{k} N_{j}} \tag{2}
\end{equation*}
$$

## The Girvan-Newman Algorithm

Calculating Betweenness


Example Crediting

- Level 3 Nodes: Credit each of nodes $A$ and $C$ with 1
- Level 2-3 Edges: Both A and C have only one parent, so full credit 1 is assigned to both $(B, A)$ and $(B, C)$

Crediting Nodes and Edges in Level 3 and 2
Adopted from mmds.org

## The Girvan-Newman Algorithm

Calculating Betweenness


## Example Crediting

Level 2 Nodes:

- G is a leaf, so gets credit 1
- $B$ is not a leaf, so gets credit $1+$ sum of credits 1 of DAG edges $(B, A),(B, C)$ leaving from it: credit 3 overall
- Intuitively, credit 3 for $B$ refers to all shortest paths from $E$ to $A, B, C$ going through $B$.
Crediting Nodes and Edges in Level 3 and 2
Adopted from mmds.org


## The Girvan-Newman Algorithm

Calculating Betweenness


Crediting Nodes and Edges
Adopted from mmds.org

## Example Crediting

Level 1-2 Edges:

- $B$ has only one parent, $D$, so the edge $(D, B)$ gets all of $B$ 's credit
- $(D, G),(F, G)$ : Both $D, F$ have label (not credit!) 1. So we credit both $(D, G),(F, G)$ with $1 /(1+1)=0.5$
- Example: If labels of $D$ and $F$ had been 3 and 5 , the credit of $(D, G)$ would be $3 /(3+5)=3 / 8$ and that of $(F, G)$ would be $5 / 8$.


## The Girvan-Newman Algorithm

CAlculating Betweenness

## Example Crediting



Crediting Nodes and Edges
Adopted from mmds.org

Level 1 Nodes / Edges:

- $D$ gets credit $1+$ credits of $(D, B),(D, G)=$ credit 4.5 overall
- $F$ gets credit $1+$ credit of $(F, G)=$ credit 1.5 overall
- Edges $(E, D),(E, F)$ receive credits of $D, F$ respectively, because $D, F$ each have only one parent

Summary: Credit on each edge is contribution to betweenness of that edge to shortest paths from $E$

## The Girvan-Newman Algorithm

SUMMARY

## Completing the Algorithm

- Repeat the calculation illustrated for $E$ for every other node
- Sum up the contributions for each edge across different roots
- Divide each edge weight by 2: each shortest path is counted twice, with each of its end points as root


Betweenness Scores
Adopted from mmds.org

## Finding Communities with Betweenness



Betweenness Scores
Adopted from mmds. org

Computing Communities: Principle

- Remove edges in decreasing order of betweenness
- Stop at reasonably chosen threshold
- Communities are the resulting connected components


## Finding Communities with Betweenness



Betweenness Scores
Adopted from mmds.org
Computing Communities: Example Threshold 4

- First, remove $(B, D)$ : communities $\{A, B, C\},\{D, E, F, G\}$
- Second, remove $(A, B),(B, C)$ : communities $\{A, C\},\{B\},\{D, E, F, G\}$
- Third, remove $(D, E),(D, G)$ : communities $\{A, C\},\{B\},\{D, E, F, G\}$
- Last, remove $(D, F)$ : communities $\{A, C\},\{B\},\{D\},\{E, F, G\}$


## Finding Communities with Betweenness

Computing Communities: Example Threshold 4

- First, remove $(B, D)$ : communities $\{A, B, C\},\{D, E, F, G\}$
- Second, remove $(A, B),(B, C)$ : communities $\{A, C\},\{B\},\{D, E, F, G\}$
- Third, remove $(D, E),(D, G)$ : communities $\{A, C\},\{B\},\{D, E, F, G\}$
- Last, remove $(D, F)$ : communities $\{A, C\},\{B\},\{D\},\{E, F, G\}$


Final Communities
Adopted from mmds.org

## The Graph Affiliation Model

## Overlapping Communities



- Observation: Communities in social networks can overlap
- Graph partitioning does not help in these cases
- Would like to have a statistical interpretation of network data


## Nonoverlapping versus Overlapping COMMUNITIES



Left: Nonoverlapping communities Right: Overlapping communities

Adopted from mmds.org

- Communities may overlap or not
- Issue: How to determine communities correctly?


## Affiliation Graph Model: Introduction



Networks and their adjacency matrices
Adopted from mmds.org

- Left: No overlap, adjacency matrix sparse across communities
- Middle: Loose overlap, adjacency matrix less sparse in shared part
- Right: Tight overlap, adjacency matrix dense in shared part


## Community Discovery: Goal



Revealing (overlapping) communities
Adopted from mmds.org

- Goal: Discover communities correctly
- Regardless of whether they overlap or not Determine the statistically most likely community structure


## Affiliation Graph Model: Introduction

- Issue: Statistical control over community structure of a network
- Idea: Design generative probability distribution
- Given a number of nodes, this generative distribution generates edges
- The generative distribution represents a particular community structure
- The distribution knows about nodes belonging to communities
- It generates more edges within communities
- It generates less edges between communities


## Affiliation Graph Model: Introduction

- The generative distribution represents community structures
- The distribution knows about nodes belonging to communities
- It generates more edges within communities
- It generates less edges between communities


Distribution representing a community structure generating network
Adopted from mmds.org

## Affiliation Graph Model: Introduction



Distribution representing a community structure (left) generating network (right) Adopted from mmds.org

- We can generate networks when knowing community structure
- But: We would like to determine the community structure when knowing the network

Isn't that exactly the opposite?

## Generative Distributions



We can do this: generating network from distribution... Adopted from mmds.org

...but we want this: inferring distribution from network
Adopted from mmds.org

## Generative Distributions: Maximum Likelihood Inference



We want to infer distribution from network

> Adopted from mmds.org

## Maximum Likelihood Estimation

- Let $\Theta$ be a parameterized class of probability distributions that generate networks
- We identify the different distributions with the different parameterizations Formally not $100 \%$ correct, but doesn't matter here
- Let $\mathbf{P}(N \mid \theta)$ be the probability that distribution $\theta \in \Theta$ generates network $N$


## Generative Distributions: Maximum Likelihood Inference



We want to infer distribution from network
Adopted from mmds.org

Maximum Likelihood Estimation

- Let $\mathbf{P}(N \mid \theta)$ be the probability that distribution $\theta \in \Theta$ generates network $N$
- Maximum likelihood estimation: Determine distribution $\hat{\theta}$ that generated $N$ with greatest likelihood:

$$
\begin{equation*}
\hat{\theta}:=\underset{\theta \in \Theta}{\arg \max } \mathbf{P}(N \mid \theta) \tag{3}
\end{equation*}
$$

UNIVERSIIATTMis computes most reasonable distribution $\hat{\theta}$ for network $N$

## Affiliation Graph Model: Definition I

- An AGM $\theta$ generates a network $N=(V, E)$ by adding edges $E$ to a given set of nodes $V$
- For $u, v \in V$, edge $(u, v)$ is generated with probability $\mathbf{P}_{\theta}((u, v))$
- $\mathbf{P}_{\theta}((u, v))$ depends on the parameters $\theta$
- Recall that $\theta$ specifies community structure

So, what exactly is $\theta$ supposed to be?

## Affiliation Graph Model: Parameters

- $\mathcal{C}$, as a set of communities
- $M \in\{0,1\}^{\mathcal{C} \times V}$, specifying assignment of nodes $v \in V$ to communities $C \in \mathcal{C}$, where

$$
M_{C, v}= \begin{cases}1 & v \text { belongs to } C  \tag{4}\\ 0 & \text { otherwise }\end{cases}
$$

- $M$ specifies "affiliations" of nodes $v \in V$
- Note that one can vary $\mathcal{C}$, as a parameter, but not $V$
- $\left(p_{C}\right)_{C \in \mathcal{C}}$ as probabilities to generate edges $(u, v)$ because $u, v \in C$
- Summary: A particular AGM $\theta$ corresponds to

$$
\begin{equation*}
\theta=\left(\mathcal{C}, M,\left(p_{C}\right)_{C \in \mathcal{C}}\right) \tag{5}
\end{equation*}
$$

## Affiliation Graph Model: $\mathbf{P}_{\theta}((u, v))$

Several $C$ containing both $u, v$

- Let $M_{u}, M_{v} \subset \mathcal{C}$ be the subsets of communities that contain $u$ and $v$, respectively
- Existence of communities that contain both $u, v$ means

$$
M_{u} \cap M_{v} \neq \emptyset
$$

- Memberships in different communities have no influence on each other
- That is, we assume statistical independence


## Affiliation Graph Model: $\mathbf{P}_{\theta}((u, v))$

Several $C$ containing both $u, v$

- Statistical independence is expressed by

$$
\prod_{C \in M_{u} \cap M_{v}}\left(1-p_{C}\right)
$$

as probability of no edge ( $u, v$ ) in any community $C \in M_{u} \cap M_{v}$

- Hence, the probability to generate $(u, v)$ is

$$
\begin{equation*}
1-\prod_{C \in M_{u} \cap M_{v}}\left(1-p_{C}\right) \tag{6}
\end{equation*}
$$

Done? No: What about $M_{u} \cap M_{v}=\emptyset$ ?

## Affiliation Graph Model: $\mathbf{P}_{\theta}((u, v))$

No $C$ containing both $u, v$

- For $M_{u} \cap M_{v}=\emptyset$, computing (6) yields (empty product is 1 )

$$
1-\prod_{C \in \emptyset}\left(1-p_{C}\right)=1-1=0
$$

- No edges across communities makes no sense
- Let $\epsilon>0$ be small; we generate an edge $(u, v)$ with probability

$$
\mathbf{P}_{\theta}((u, v))=\epsilon \quad \text { if } \quad M_{u} \cap M_{v}=\emptyset
$$

## Affiliation Graph Model: $\mathbf{P}_{\theta}((u, v))$

## Affiliation Graph Model (AGM)

- An edge $(u, v)$ is generated with probability

$$
\mathbf{P}_{\theta}((u, v))= \begin{cases}1-\prod_{C \in M_{u} \cap M_{v}}\left(1-p_{C}\right) & M_{u} \cap M_{v} \neq \emptyset  \tag{7}\\ \epsilon & M_{u} \cap M_{v}=\emptyset\end{cases}
$$

- Edges $(u, v)$ are generated independently from one another
- Overall: The probability $\mathbf{P}_{\theta}(E)$ to generate edges $E$ given AGM $\theta$ computes as

$$
\begin{equation*}
\mathbf{P}_{\theta}(E)=\prod_{(u, v) \in E} \mathbf{P}_{\theta}((u, v)) \times \prod_{(u, v) \notin E} 1-\mathbf{P}_{\theta}((u, v)) \tag{8}
\end{equation*}
$$

where $\mathbf{P}_{\theta}((u, v))$ are computed following (7), with $\theta=\left(\mathcal{C}, M, p_{C}\right)$ determining $p_{C}$ and $M_{u}, M_{v}$ and so on.

## Affiliation Graph Model: Overall Probability

## Affiliation Graph Model (AGM)

- The probability $\mathbf{P}_{\theta}(E)$ to generate $E$ given $\theta$ is

$$
\begin{equation*}
\mathbf{P}_{\theta}(E)=\prod_{(u, v) \in E} \mathbf{P}_{\theta}((u, v)) \times \prod_{(u, v) \notin E} 1-\mathbf{P}_{\theta}((u, v)) \tag{9}
\end{equation*}
$$

- Reminder: For a given network $N=(V, E)$, the goal is to determine

$$
\hat{\theta}:=\underset{\theta \in \Theta}{\arg \max } \mathbf{P}_{\theta}(E)
$$

- That is, we need to vary $\theta=\left(\mathcal{C}, M, p_{C}\right)$ until $\mathbf{P}_{\theta}(E)$ is maximal

$$
\text { How to systematically vary } \theta=\left(\mathcal{C}, M, p_{C}\right) ?
$$

## Computing the MLE $\hat{\theta}$

## ISSUES

- Search space of combinations of
- Communities $\mathcal{C}$,
- Assignments of nodes to communities $M$, and
- Probabilities $p_{C}$ for communities
tends to be huge
- Concise formulas of (9) for $\mathbf{P}_{\theta}(E)$ as function of $\theta$ too difficult
- Analytical solution for determining $\hat{\theta}:=\arg \max _{\theta \in \Theta} \mathbf{P}_{\theta}(E)$ not available
- Moreover, parameters are both discrete $(\mathcal{C}, M)$ and continuous $\left(\left(p_{\mathcal{C}}\right)_{C \in \mathcal{C}}\right)$


## Computing the MLE $\hat{\theta}$

## Approach

1. Pick initial set of parameters $\theta_{0}$
2. Vary $\theta$ such that $\mathbf{P}_{\theta}(E)$ iteratively increases
3. Vary $\mathcal{C}$ or $M$ first

Partial derivates of $\mathbf{P}_{\theta}(E)$ wrt. $p_{C}$ computable on fixed $\mathcal{C}, M$
4. Determine optimal $\left(p_{C}\right)_{C \in \mathcal{C}}$, e.g. by gradient descent
5. Keep change if $\mathbf{P}_{\theta}(E)$ has increased, discard otherwise

## Computing the MLE $\hat{\theta}$

Iterative variations of $\mathcal{C}, M$

- Varying M:
- Delete node from community, i.e. for $M_{C, v}=1$, set $M_{C, v}=0$
- Add node to community, i.e. for $M_{C, v}=0$, set $M_{C, v}=1$
- Varying $\mathcal{C}$ :
- Merge two communities
- Split community
- Delete community
- Add new community, with initial random selection of members


## Computing the MLE $\hat{\theta}$

## Soft Community Membership

- Instead of $M_{C, v} \in\{0,1\}$, allow any real-numbered $M_{C, v} \geq 0$
- For $(u, v)$ to be generated because of $u, v \in C$, let

$$
\begin{equation*}
\mathbf{P}_{\theta}((u, v))=1-e^{-M_{C, u} M_{\mathcal{C}, v}} \tag{10}
\end{equation*}
$$

be the individual probability

- Proceeding exactly as before, we obtain

$$
\begin{equation*}
\mathbf{P}_{\theta}(E)=\prod_{(u, v) \in E}\left(1-e^{-\sum_{\mathrm{C}} M_{\mathcal{C}, u} M_{\mathcal{C}, v}}\right) \prod_{(u, v) \notin E} e^{-\sum_{\mathrm{C}} M_{\mathrm{C}, u} M_{\mathcal{C}, v}} \tag{11}
\end{equation*}
$$

## Computing the MLE $\hat{\theta}$

## Soft Community Membership

- Probability for edges $E$ :

$$
\begin{equation*}
\mathbf{P}_{\theta}(E)=\prod_{(u, v) \in E}\left(1-e^{-\sum_{\mathrm{C}} M_{\mathcal{C}, u} M_{\mathcal{C}, v}}\right) \prod_{(u, v) \notin E} e^{-\sum_{\mathrm{C}} M_{\mathcal{C}, u} M_{\mathcal{C}, v}} \tag{12}
\end{equation*}
$$

- On fixed communities, include $M$ in gradient descent (or related) optimization step
- Advantages:
- Only one gradient descent run necessary
- Less prone to get stuck in unfavorable local optima
- If necessary, add or delete communities, and re-run


## Direct Discovery of Overlapping Communities

## INTRODUCTION

- Popular idea: Determine communities as (induced) subgraphs of a certain type
- Subgraphs should contain unusually large amount of edges
- Subgraphs are allowed to overlap
- Will treat two types briefly here:
- Cliques
- Complete bipartite subgraphs


## Finding Cliques

Definition [Induced Subgraph]
Let $G=(V, E)$ be a graph. A subgraph $C=\left(V^{\prime} \subset V, E^{\prime} \subset E\right)$ is induced iff

$$
\left(v^{\prime}, w^{\prime}\right) \in E \text { implies } \quad\left(v^{\prime}, w^{\prime}\right) \in E^{\prime}
$$

for any $v^{\prime}, w^{\prime} \in V^{\prime}$.
Definition [Clique]
Let $G=(V, E)$ be a graph.

- An induced subgraph $C=\left(V^{\prime}, E^{\prime}\right)$ is called a clique iff any pair of nodes in $C$ is connected by an edge.
- A clique $C=\left(V^{\prime}, E^{\prime}\right)$ is maximal iff extending the clique by any node and its edges implies that the clique property no longer holds.


## Communities as Cliques

- Possible idea: Determine communities as maximal cliques
- Caveat: The number of maximal cliques in a graph may be exponential in the number of nodes
- So, listing all maximal cliques is a computationally demanding problem
- Nevertheless, identifying communities as clique like arrangements is popular


## Complete Bipartite Graphs

## Definition [(Complete) Bipartite Graphs]

A graph $G=(V, E)$ with vertices $V$ and edges $E$ is referred to as bipartite iff

- there are $V_{1}, V_{2} \subset V$ such that

$$
V=V_{1} \dot{\cup} V_{2} \quad \text { and } \quad E \subset\left(V_{1} \times V_{2}\right)
$$

- A bipartite graph $G=(V, E)$ is complete iff

$$
V=V_{1} \cup V_{2} \quad \text { and } \quad E=\left(V_{1} \times V_{2}\right)
$$

that is iff each node from $V_{1}$ is connected with each node from $V_{2}$

- A complete bipartite graph where $\left|V_{1}\right|=s,\left|V_{2}\right|=t$ is referred to as $K_{s, t}$
- A complete bipartite graph is also referred to as biclique


## Complete Bipartite Graphs and Communities

- Strategy: Seek to discover all sufficiently large bicliques
- Treat them as "nuclei" (or seeds) of communities
- Theoretical Advantage over Cliques: While it is not possible to guarantee the existence of large cliques for graphs with many edges, one can guarantee the existence of large bicliques


## Finding Complete Bipartite Graphs

Frequent Itemset Mining Problem

- Let $G=(V, E)$ on $V=V_{1} \dot{\cup} V_{2}$ be a (large) bipartite graph
- Items are nodes from $V_{1}$
- Baskets are nodes from $V_{2}$
- Items in baskets are nodes from $V_{1}$ connected to basket node
- $K_{s, t}$ in $G$ is itemset of size $s$ that appears in $t$ baskets
- So mining for frequent itemsets at threshold $t$ dicovers all $K_{s, t}$


## General / Further Reading

Literature

- Mining Massive Datasets, Sections 10.1, 10.2, 10.3, 10.5 http://infolab.stanford.edu/~ullman/mmds/ ch10.pdf
- Next lecture: "Web Advertisements": sections 8.1-8.4 in Mining of Massive Datasets

