A Shapley value Approach for Influence Attribution

Panagiotis Papapetrou

Aalto University, Finland

Joint work with:

Aris Gionis, Yahoo! Research, Spain

Heikki Mannila, Aalto University, Finland

Influential individuals

- People always intrigued by characterizing influential ideas, books, scientists, politicians, etc.
- Main question: who or what is influential?
- Examples
  - Who initiates the most influential "tweets"?
  - Who are the most influential scientists?
  - Which actors influence a movie rating the most?

Goal

- **We address a novel problem in the context of characterizing who is influential.**
- **Our setting:**
  - Individuals accomplish tasks in a collaborative manner.
  - **Influence attribution:** each individual is assigned a score based on his/her performance.

Outline

- Problem Formulation
- Proposed Solution
- Experimental Evaluation
- Conclusions

Example: author-publication

- Individual => author.
- Task => publication.
- Impact score =>
    - CC: Citation count of the publication.
    - PR: PageRank score of the publication.

Example: author-publication

- Two researchers A and B.
- Question: who is more influential?

Example: author-publication

- One common collaborator: Y.

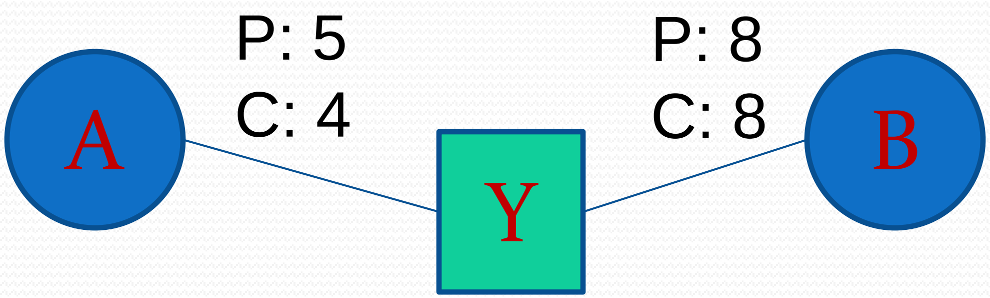

P: number of papers
C: number of citations per paper

Example: author-publication

- Three additional collaborators for A and B.

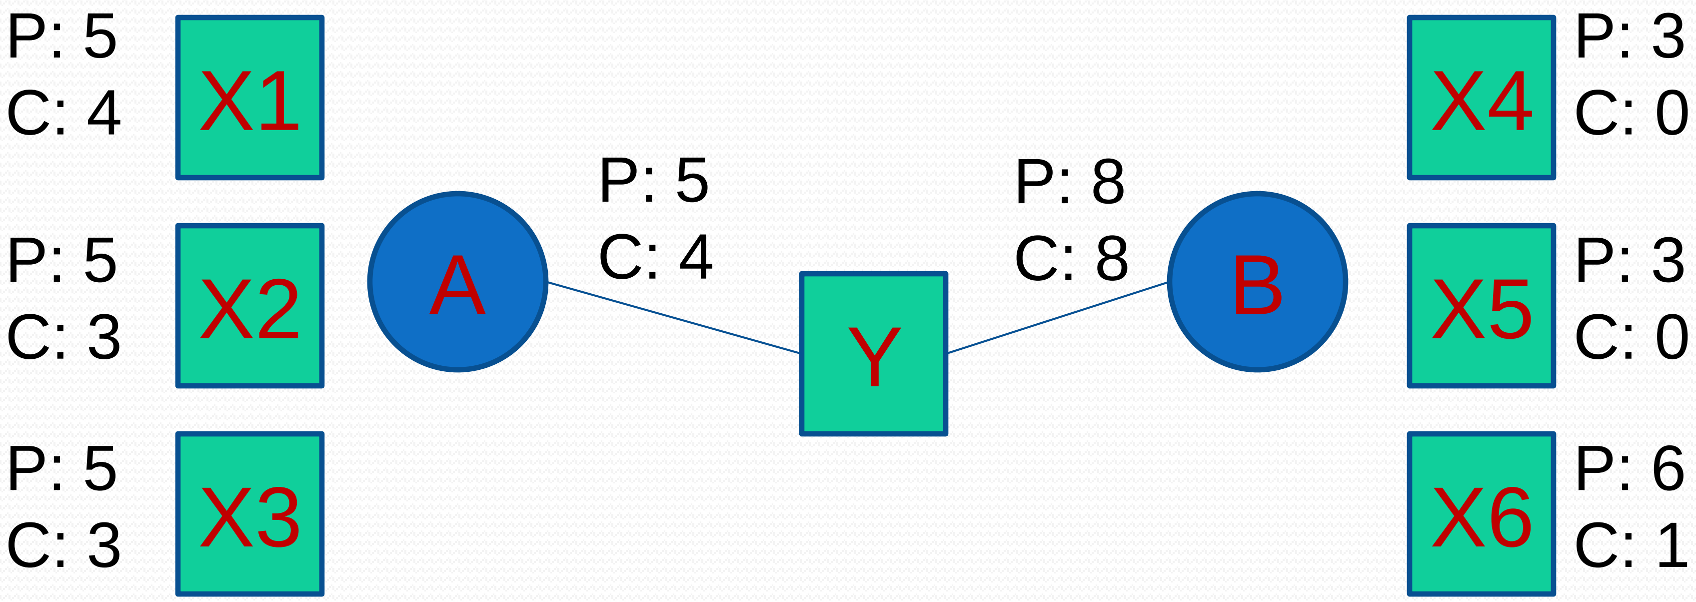

Example: author-publication

- Three additional collaborators for A and B.

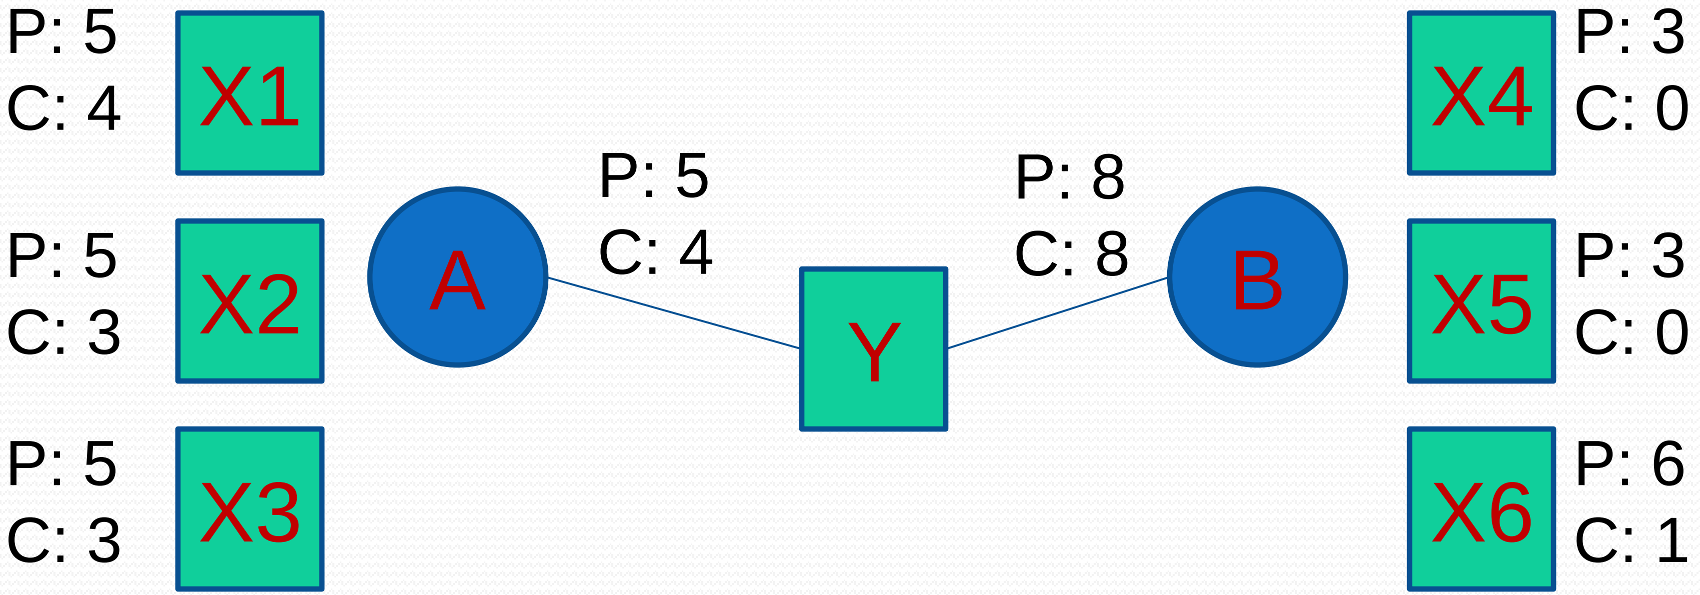

Researcher	Papers	Citations	H-index
A	20	70	4
B	20	70	8

Example: author-publication

- Three additional collaborators for A and B.

H-Index: a scientist's H-index is h, if h of his/her publications have at least h citations and the rest of his/her publications have at most h citations each.

Researcher	Papers	Citations	H-index
A	20	70	4
B	20	70	8

Example: author-publication

- Three additional collaborators for A and B.

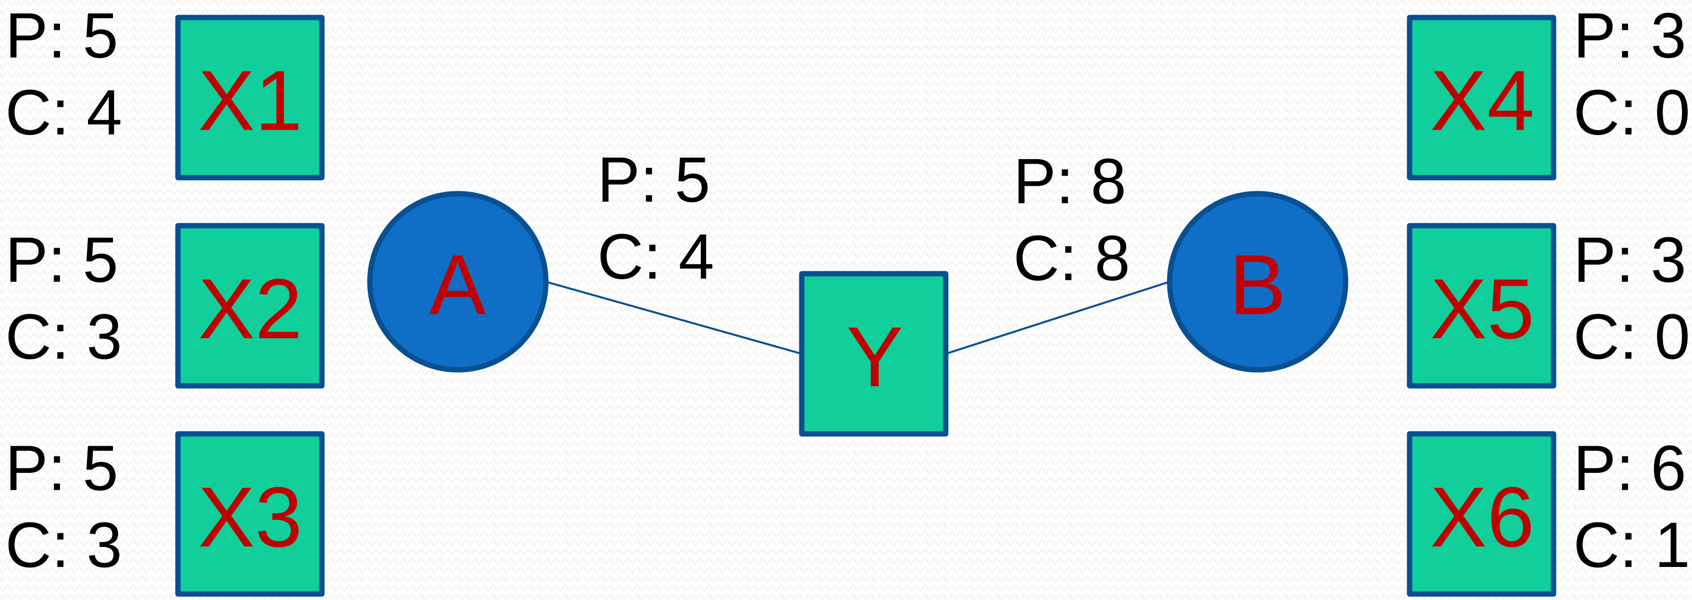

Researcher	Papers	Citations	H-index
A	20	70	4
B	20	70	8

Example: author-publication

- Three additional collaborators for A and B.

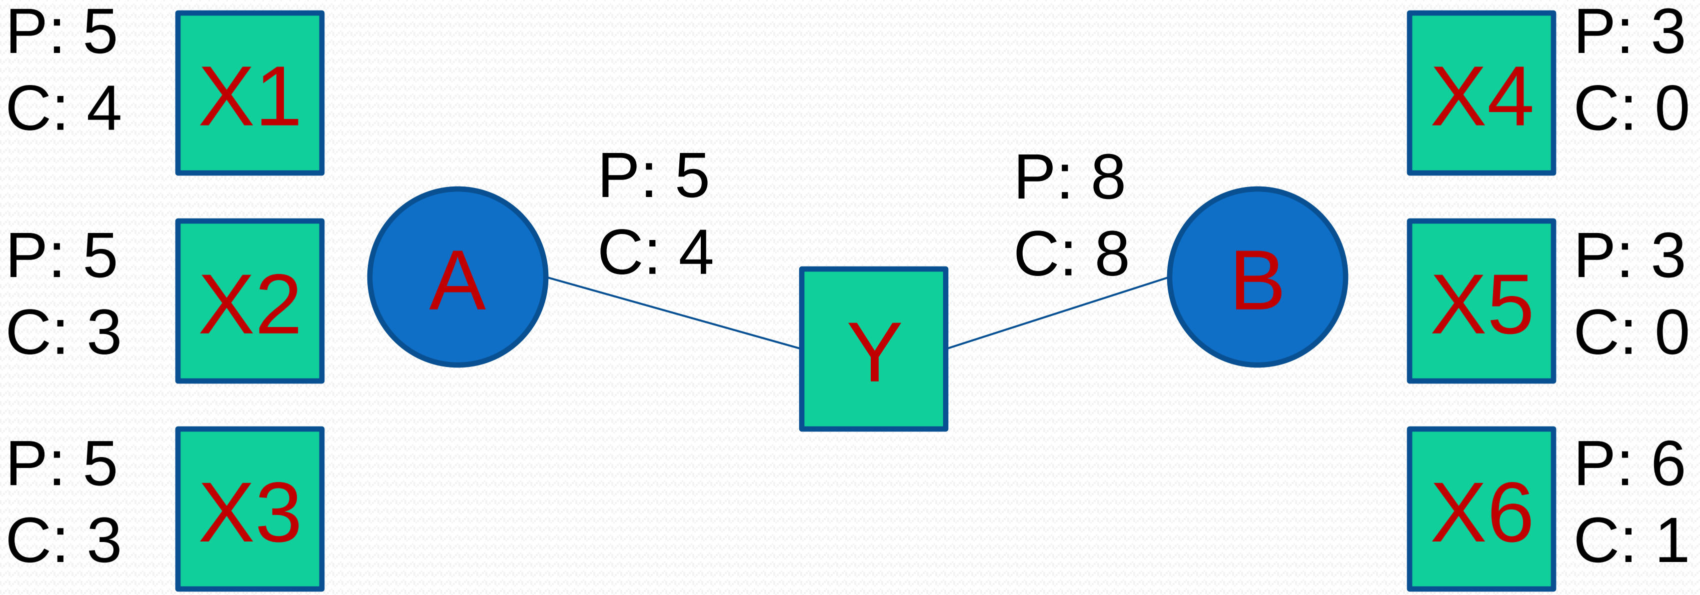

- But is B indeed that influential?
- Or is B just being favored due to the fame of Y?

Example: author-publication

- Drop Y out of the picture.

- The performance of A remains quite high.
- The performance of B is weakened a lot.

Example: author-publication

- Drop Y out of the picture.

Researcher	Papers	Citations	H-index
A	15	50	4
B	12	6	1

Background

- Existing measures in bibliometrics can be enriched.
- Social network analysis methods focus on finding *important* individuals based on in-degree or refinements.
- Information diffusion finds individuals who act as *good* initiators.
- Coalitional games: Shapley value.

Problem Definition

- Given
  - a set of individuals $V = \{V_1, \ldots, V_n\}$,
  - a set of tasks $T = \{T_1, \ldots, T_m\}$,
  - a set of impact scores $I = \{I_1, \ldots, I_m\}$.
- Goal:
  - Compute the set of influence scores $\phi = \{\phi_1, \ldots, \phi_n\}$.
- Φ_i is the influence score of individual V_i.

Outline

- Problem Formulation
- Proposed Solution
- Experimental Evaluation
- Conclusions

Shapley Value

- Consider an underlying set V.
- Assume for all possible subsets S of V we know $v(S)$.
- $v(S)$: gain function
  - expresses the gain achieved by the cooperation of the individuals in S.
- Shapley value: the share allocation to individual V_i.

$$\phi_i(v) = \sum_{\mathcal{S} \subseteq \mathcal{V}} \frac{|\mathcal{S}|!(|\mathcal{V}| - |\mathcal{S}| - 1)!}{|\mathcal{V}|!} (v(\mathcal{S} \cup \{V_i\}) - v(\mathcal{S})).$$

Shapley Value

- Can be shown theoretically that the resulting attribution satisfies natural fairness properties [Winter 2002].
- However, a direct application of the Shapley value definition in our setting is not possible:
  - it assumes an averaging over exponentially many sets,
  - it is not possible to probe arbitrary sets S and obtain $v(S)$,
  - we may not have available the impact score of papers for every possible subset of authors!

Our Approach

- We compute the marginal gains by averaging only over **coalitions for which we have available impact scores**.
- In order to average in a marginal contribution we need to have available both values $v(S \cup \{V_i\})$ and $v(S)$.
- In many cases we have available only one of the two.
- How shall we deal with such cases?
  - Ignore them? ➔ very sparse data.

Our Approach

- We choose to take into account all cases for which $S \cup \{V_i\}$ is available.

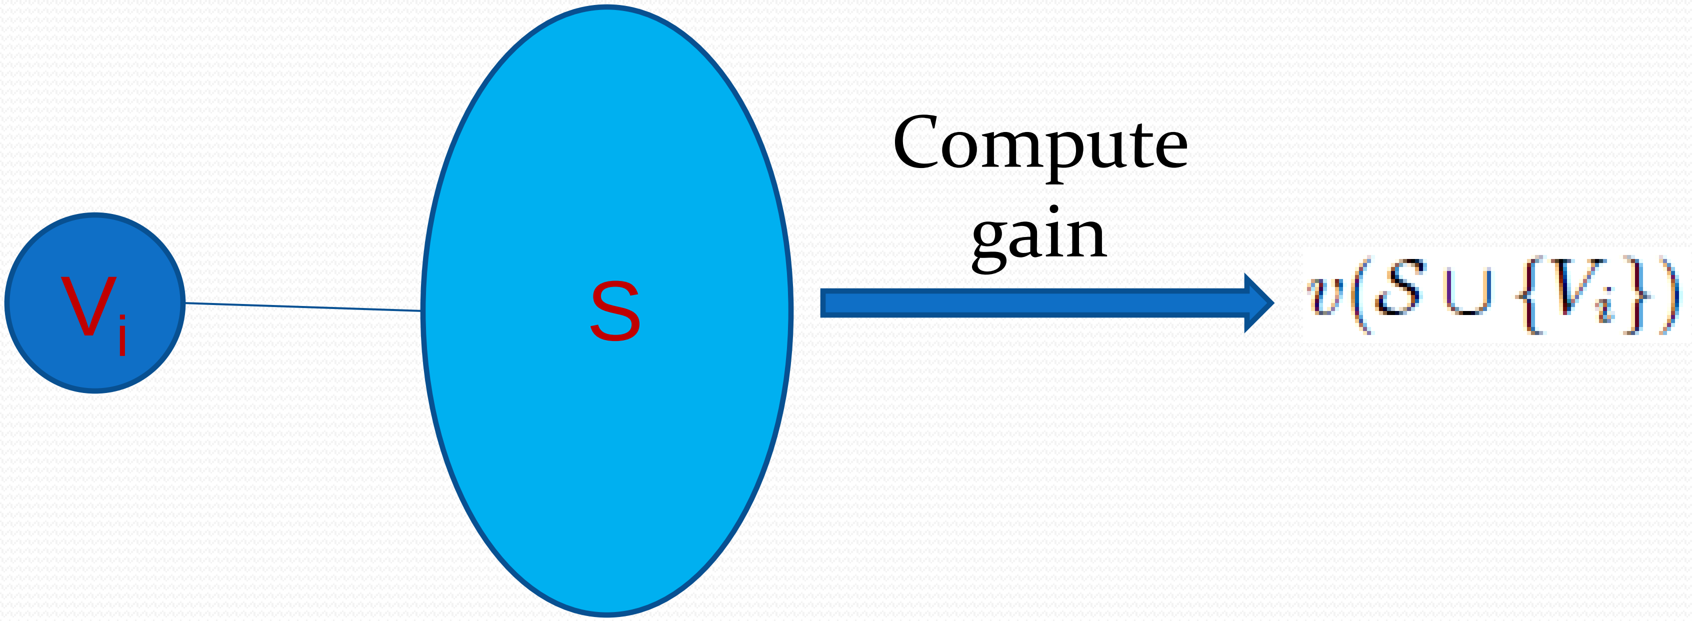

Our Approach

- We choose to take into account all cases for which $S \cup \{V_i\}$ is available.

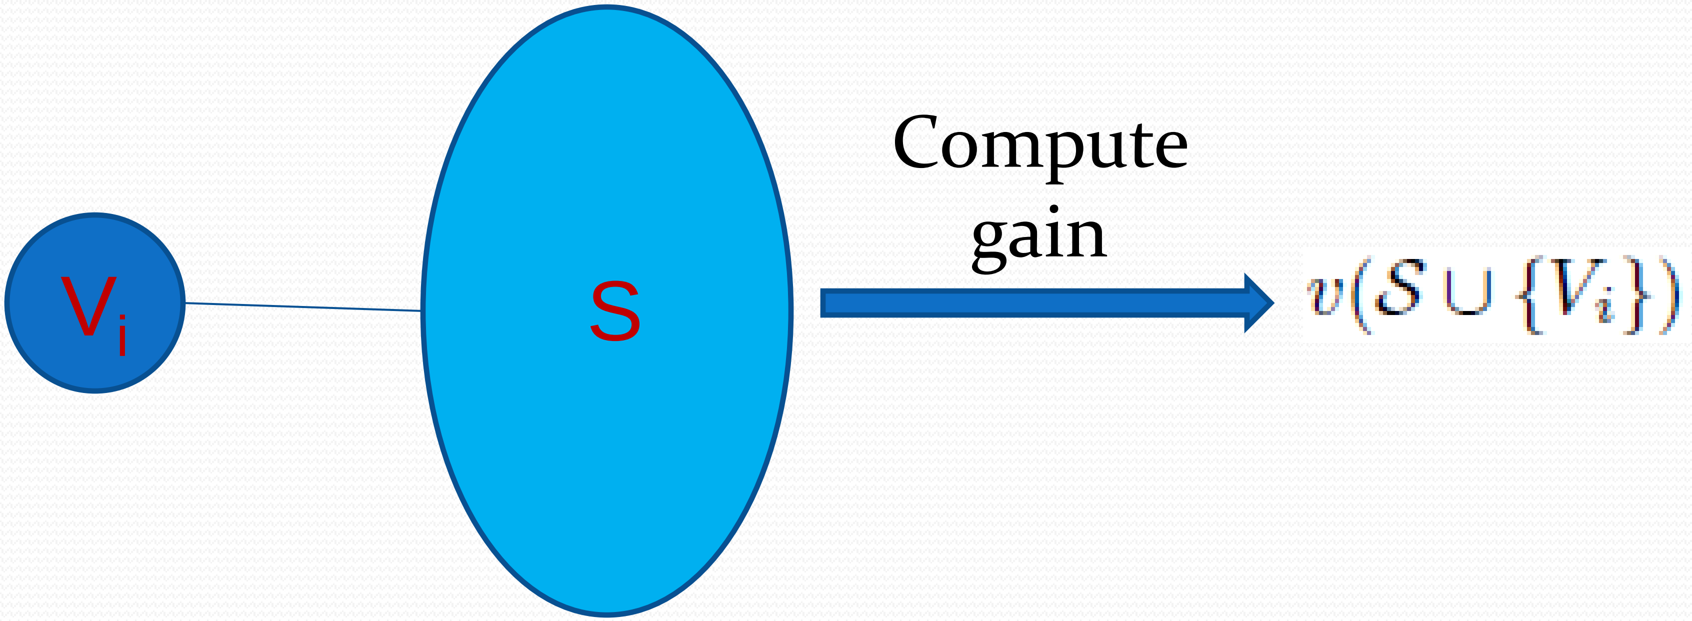

- **What about $v(S)$?**

Shared impact factor

- Let I_j be the impact factor of each common task T_j between a group of individuals S.
- Then the shared impact factor is the *average impact factor* among all their common tasks T_S.

$$v(S) = \frac{1}{|T_S|} \sum_{j=1}^{|T_S|} I_j.$$

Approximated Shared impact factor

- What if for some set S we have no complete information about the coalitions?

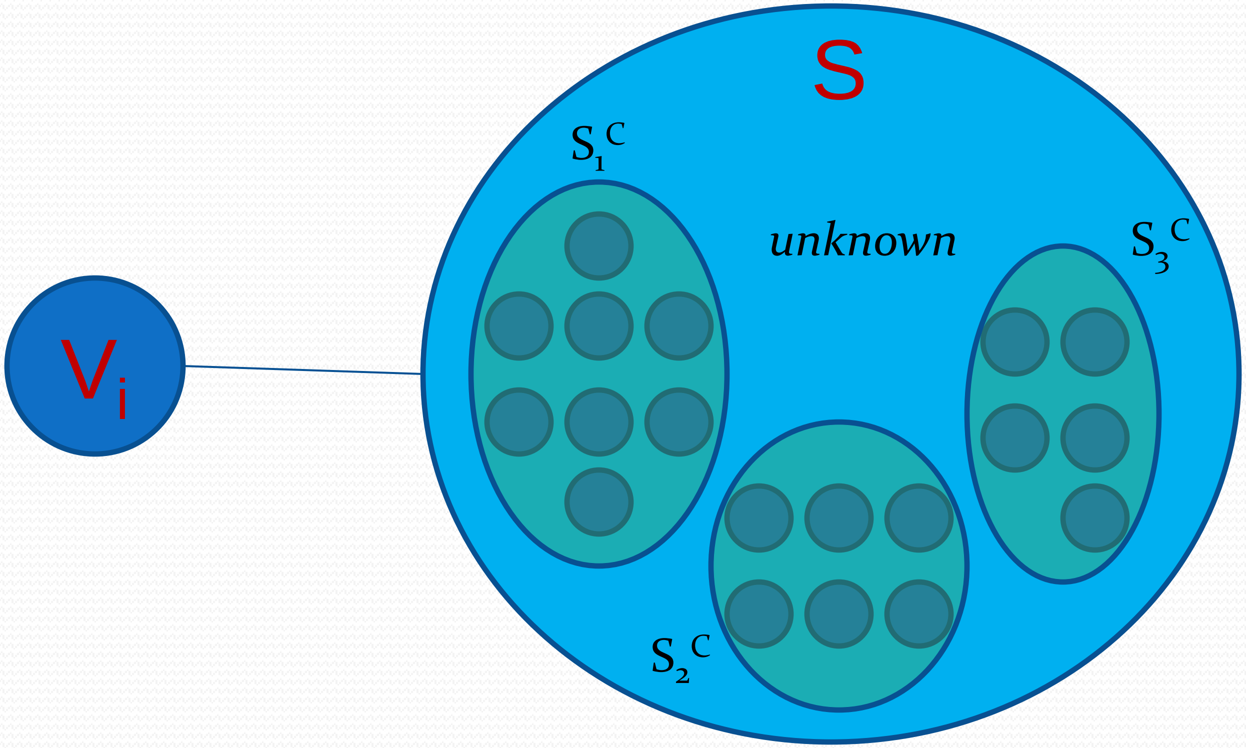

Approximated Shared impact factor

- What if for some set S we have no complete information about the coalitions?
- Take only the subsets S_i^C of S for which there is such information:

$$v'(S) = \frac{1}{|S^c| + 1} \left(\sum_{i=1}^{|S^c|} v(S_i^c) + \bar{v}(S \setminus S_i^c) \right)$$

Approximated Gain Function

- What about $\bar{v}(S)$?
- Assuming a monotonic behavior, i.e., teams are at least as good as the best individual in the team, we define:

$$\bar{v}(S) = \max_{V_i \in S} \phi_i(v).$$

The Iterative Algorithm

- Goal: compute the influence score ϕ_i of each individual.
- At each iteration t the Shapley value is computed using the original definition:

$$\phi_i(v) = \sum_{\mathcal{S} \subseteq \mathcal{V}} \frac{|\mathcal{S}|!(|\mathcal{V}| - |\mathcal{S}| - 1)!}{|\mathcal{V}|!} (v(\mathcal{S} \cup \{V_i\}) - v(\mathcal{S})).$$

The Iterative Algorithm

- Goal: compute the influence score ϕ_i of each individual.
- At each iteration t the Shapley value is computed using the original definition:
- Whenever we need to probe a coalition for which the impact factor is not available, use the approx. shared impact factor:

$$v'(S) = \frac{1}{|S^c| + 1} \left(\sum_{i=1}^{|S^c|} v(S_i^c) + \bar{v}(S \setminus S_i^c) \right)$$

The Iterative Algorithm

- Goal: compute the influence score ϕ_i of each individual.
- At each iteration t the Shapley value is computed using the original definition:
- Whenever we need to probe a coalition for which the impact factor is not available, use the approx. shared impact factor.
- Influence score is updated:

$$\phi_i^{t+1}(v') = \sum_{\mathcal{V}_{T_j} | V_i \in T_j} \frac{|\mathcal{V}_{T_j}|!(|\mathcal{V}| - |\mathcal{V}_{T_j}| - 1)!}{|\mathcal{V}|!} (v'(\mathcal{V}_{T_j}) - v'(\mathcal{V}_{T_j} \setminus V_i)).$$

Outline

- Problem Formulation
- Proposed Solution
- Experimental Evaluation
- Conclusions

Experimental Setup

- Datasets:
  - ISI Web of Science.
  - Internet Movie Database (IMDB).
- ISI Web of Science:
  - Part of the Thomson Reuters ISI Web of Science data.
  - ISI covers mainly journal publications.
  - We sampled data related to our institutions published within years 2003 and 2009.
  - Our dataset contains information about 1212 authors.

Experimental Setup

- Internet Movie DataBase:
  - We sampled a total of 2 000 male actors.
  - We restricted the movie genre type to comedy or action.
  - For each actor we considered only the movies where his credit position was among the top 3.

Experimental Evaluation

- We used two very common bibliometric indicators as the baseline:
  - H-Index, G-index.
- Impact score for a publication:
  - CC: Citation count of the publication.
  - PR: PageRank score of the publication.

Experimental Evaluation

- Each movie is assigned with an impact score defined as follows:

  average rating x *number of people*

- Performance measure:
  - **Rank of an individual:** *number of individuals who are at least as influential.*

Naïve PR vs. Shapley PR

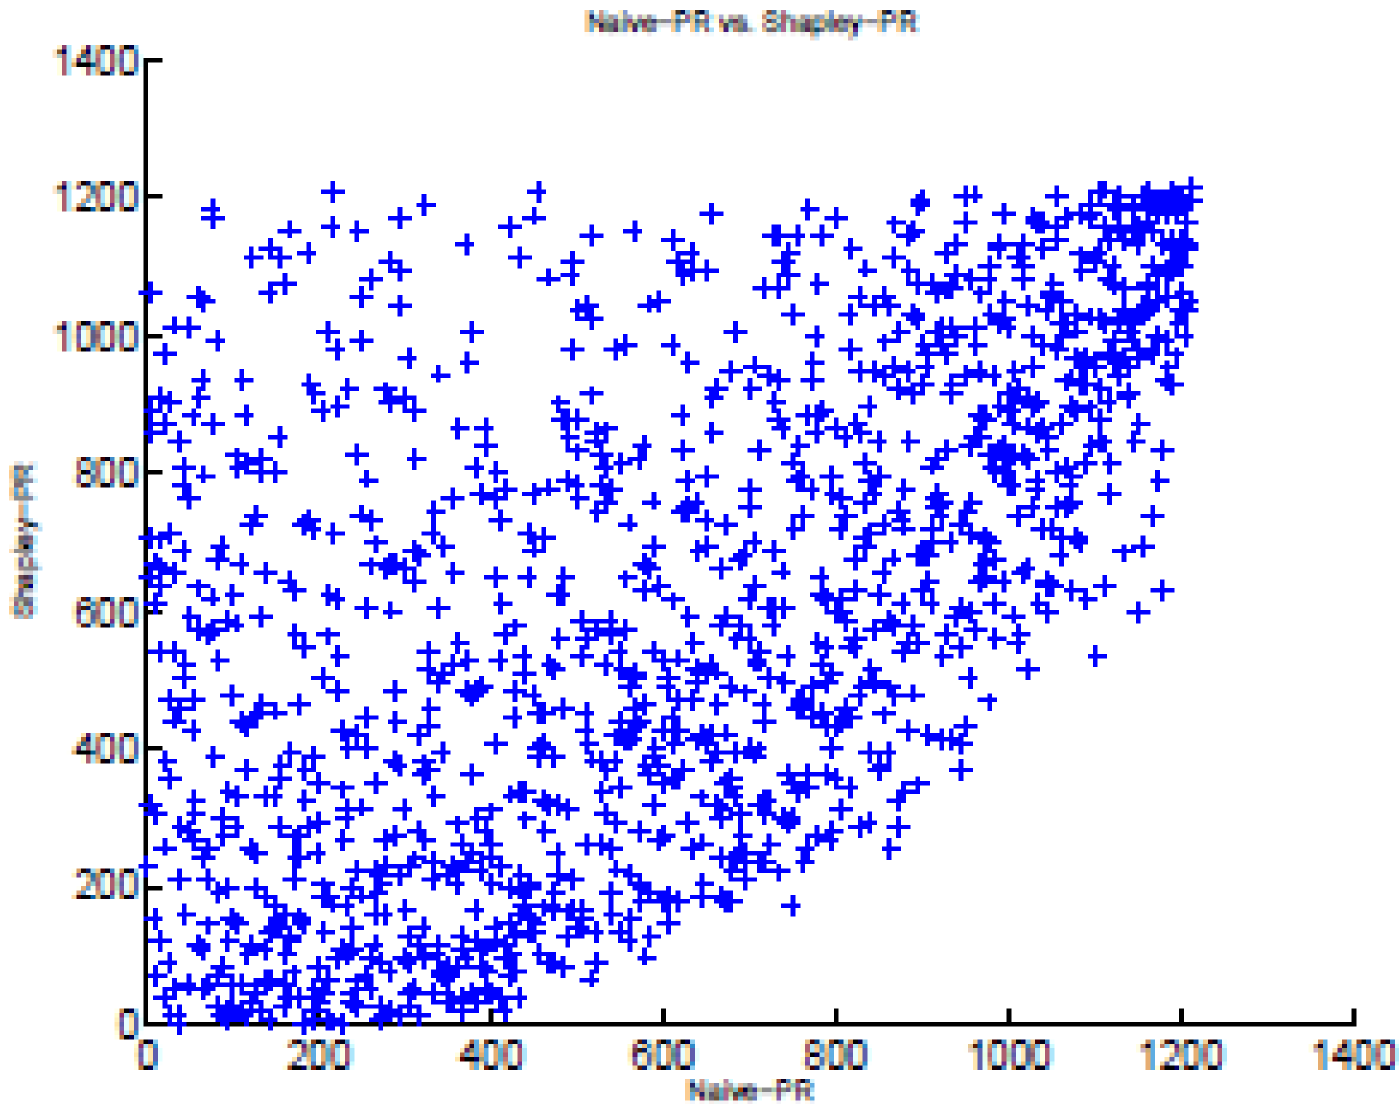

Naïve PR vs. Shapley PR

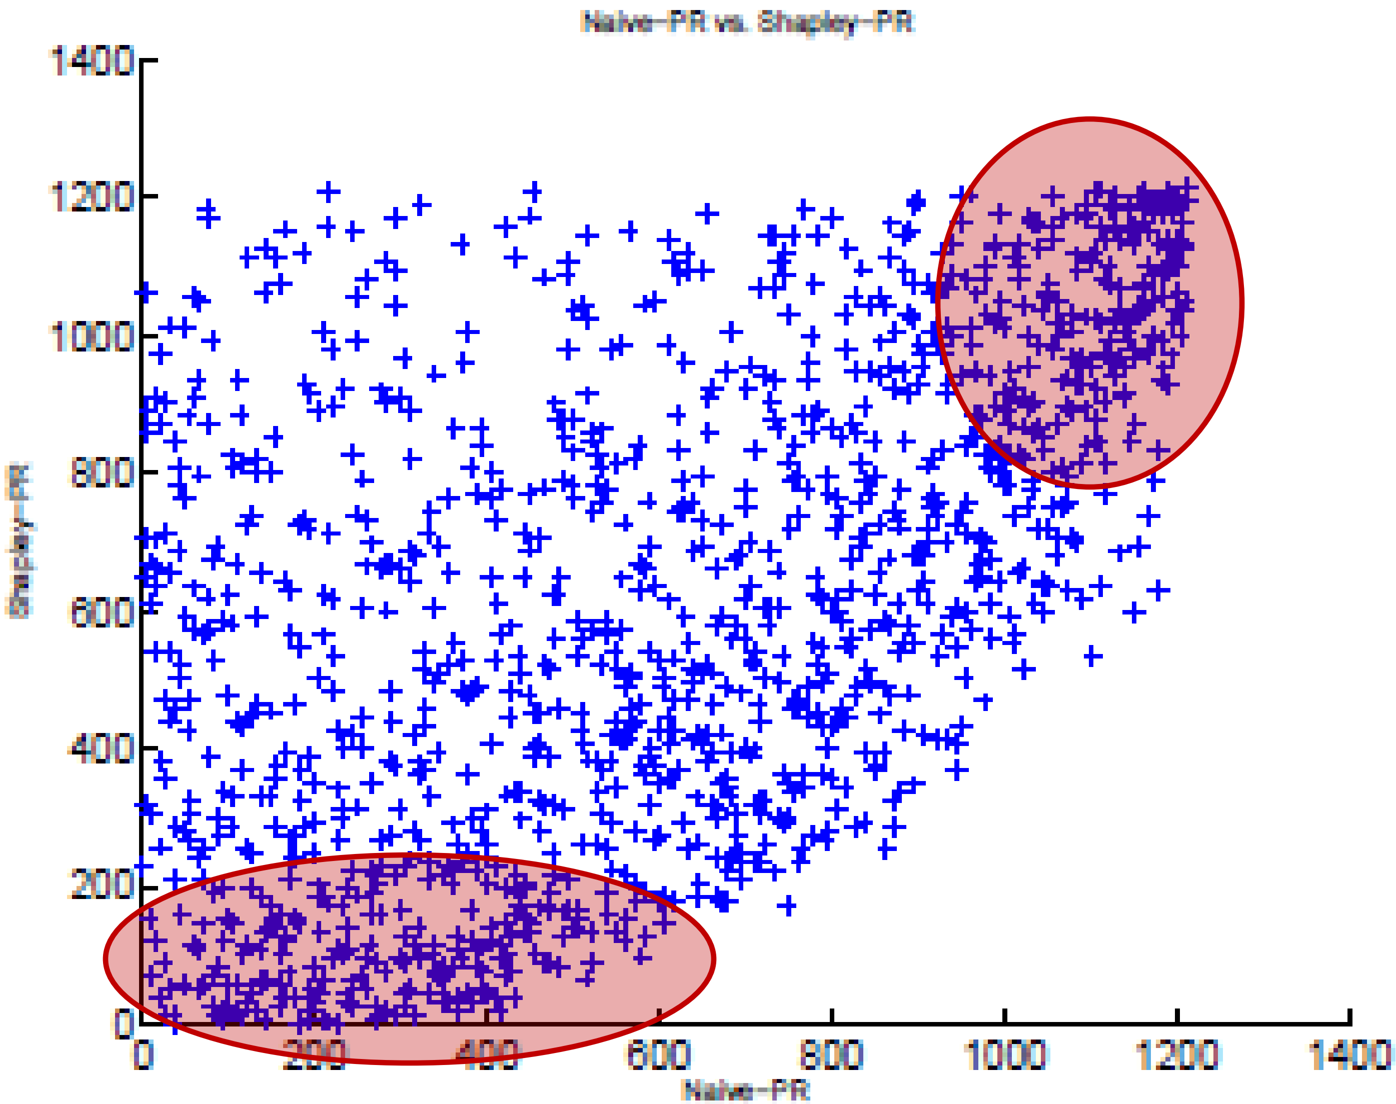

Experimental Evaluation

- Top-10 actors given by the Shapley method.

Actor Name	Shapley	Naïve	Actor Name	Naïve	Shapley
Robert De Niro	1	3	Peter Sellers	1	14
Al Pacino	2	8	Jack Nicholson	2	11
Brad Pitt	3	15	Robert De Niro	3	1
Bruce Willis	4	7	Adam Sandler	4	59
Arnold Schwarzenegger	5	24	Daniel Day-Lewis	5	36
Will Smith	6	13	Chris Farley	6	20
Eddie Murphy	7	10	Bruce Willis	7	4
Robin Williams	8	9	Al Pacino	8	2
Morgan Freeman	9	17	Robin Williams	9	8
Ben Stiller	10	29	Eddie Murphy	10	7

Experimental Evaluation

- Examples of actors with high ranking differences between Shapley and Naïve.

Actor Name	Shapley	Naïve	# of Movies in IMDB
Jim Carrey	11	79	34
Sylvester Stallone	12	41	46
Daniel Day-Lewis	36	5	27
Adam Sandler	59	4	39

Outline

- Problem Formulation
- Proposed Solution
- Experimental Evaluation
- Conclusions

Conclusions

- Addressed the problem of influence attribution
- Proposed a method that employs the game theoretic concept of Shapley value.
- Methodology can be applied to real scenarios:
  - Author-publication data.
  - Movie data.
- Experiments on two domains showed that the rankings produced by the proposed method and the naïve approach of equal division of influence differ highly.

Future Work

- Investigation of other domains such as:
  - user-blogs,
  - social media sites.
- How additional information about the individuals can affect/be taken into account.
- Further evaluate the quality of the obtained rankings by performing user studies.

Appendix

The Iterative Algorithm

Algorithm 1 The `Shapley` Algorithm

Input: a set of individuals $\mathcal{V}$, a set of tasks $\mathcal{T}$, and the corresponding set of impact scores $\mathcal{I}$.
Output: the influence score ϕ_i of each individual $V_i \in \mathcal{V}$
// Initialization: $\forall T_i, i = 1, \ldots, m$ assigned to individual V_i:
for $j = 1 : |\mathcal{V}|$ **do**
    $\phi_i^0 = \sum_{i=j}^{m} I_j$
end for
while convergence **do**
    Initialize $\phi_i^{t+1}(v') = 0$
    for $T_j \in \mathcal{T}$ **do**
        for $V_i \in \mathcal{V}_{T_j}$ such that V_i is assigned with task T_j **do**
            $\phi_i^{t+1}(v') = \phi_i^{t+1}(v') + \frac{|\mathcal{V}_{T_j}|!(|\mathcal{V}|-|\mathcal{V}_{T_j}|-1)!}{|\mathcal{V}|!}(v'(\mathcal{V}_{T_j}) - v'(\mathcal{V}_{T_j} \setminus V_i))$
        end for
    end for
end while

The Iterative Algorithm

Algorithm 1 The `Shapley` Algorithm

Input: a set of individuals $\mathcal{V}$, a set of tasks $\mathcal{T}$, and the corresponding set of impact scores $\mathcal{I}$.

Output: the influence score ϕ_i of each individual $V_i \in \mathcal{V}$

// Initialization: $\forall T_i, i = 1, \ldots, m$ assigned to individual V_i:

for $j = 1 : |\mathcal{V}|$ **do**

  $\phi_i^0 = \sum_{i=j}^{m} I_j$

end for

while convergence **do**

  Initialize $\phi_i^{t+1}(v') = 0$

  for $T_j \in \mathcal{T}$ **do**

    for $V_i \in \mathcal{V}_{T_j}$ such that V_i is assigned with task T_j **do**

      $\phi_i^{t+1}(v') = \phi_i^{t+1}(v') + \frac{|\mathcal{V}_{T_j}|!(|\mathcal{V}|-|\mathcal{V}_{T_j}|-1)!}{|\mathcal{V}|!}(v'(\mathcal{V}_{T_j}) - v'(\mathcal{V}_{T_j} \setminus V_i))$

    end for

  end for

end while

The Iterative Algorithm

Algorithm 1 The `Shapley` Algorithm

Input: a set of individuals $\mathcal{V}$, a set of tasks $\mathcal{T}$, and the corresponding set of impact scores $\mathcal{I}$.
Output: the influence score ϕ_i of each individual $V_i \in \mathcal{V}$
// Initialization: $\forall T_i, i = 1, \ldots, m$ assigned to individual V_i:
for $j = 1 : |\mathcal{V}|$ **do**
    $\phi_i^0 = \sum_{i=j}^{m} I_j$
end for
while convergence **do**
    Initialize $\phi_i^{t+1}(v') = 0$
    for $T_j \in \mathcal{T}$ **do**
        for $V_i \in \mathcal{V}_{T_j}$ such that V_i is assigned with task T_j **do**
            $\phi_i^{t+1}(v') = \phi_i^{t+1}(v') + \frac{|\mathcal{V}_{T_j}|!(|\mathcal{V}| - |\mathcal{V}_{T_j}| - 1)!}{|\mathcal{V}|!}(v'(\mathcal{V}_{T_j}) - v'(\mathcal{V}_{T_j} \setminus V_i))$
        end for
    end for
end while

The Iterative Algorithm

Algorithm 1 The `Shapley` Algorithm

Input: a set of individuals $\mathcal{V}$, a set of tasks $\mathcal{T}$, and the corresponding set of impact scores $\mathcal{I}$.

Output: the influence score ϕ_i of each individual $V_i \in \mathcal{V}$

// Initialization: $\forall T_i, i = 1, \ldots, m$ assigned to individual V_i:

for $j = 1 : |\mathcal{V}|$ **do**

  $\phi_i^0 = \sum_{i=j}^{m} I_j$

end for

while convergence **do**

  Initialize $\phi_i^{t+1}(v') = 0$

  for $T_j \in \mathcal{T}$ **do**

    for $V_i \in \mathcal{V}_{T_j}$ such that V_i is assigned with task T_j **do**

      $\phi_i^{t+1}(v') = \phi_i^{t+1}(v') + \frac{|\mathcal{V}_{T_j}|!(|\mathcal{V}|-|\mathcal{V}_{T_j}|-1)!}{|\mathcal{V}|!}(v'(\mathcal{V}_{T_j}) - v'(\mathcal{V}_{T_j} \setminus V_i))$

    end for

  end for

end while

Termination Criterion

- The iterative algorithm terminates when influence scores at two consecutive iterations converge:

$$\frac{\sum_{i=1}^{|\mathcal{V}|} |\phi_i^t - \phi_i^{t-1}|}{\sum_{i=1}^{|\mathcal{V}|} \phi_i^{t-1}} \leq \epsilon \in (0,1).$$

Enforcing Monotonicity

- Gain function should be
  - monotone, i.e., if $S_1 \subseteq S_2$ then $v(S_1) \leq v(S_2)$.
  - non-negative.
- Compute all pay-offs.
- Identify all pairs of pay-offs such that

  $S_1 \subseteq S_2$ and $v(S_1) > v(S_2)$.
- Set $v(S_1) = v(S_2)$.
- Repeat until all violations are eliminated.

Naïve CC vs. Shapley CC

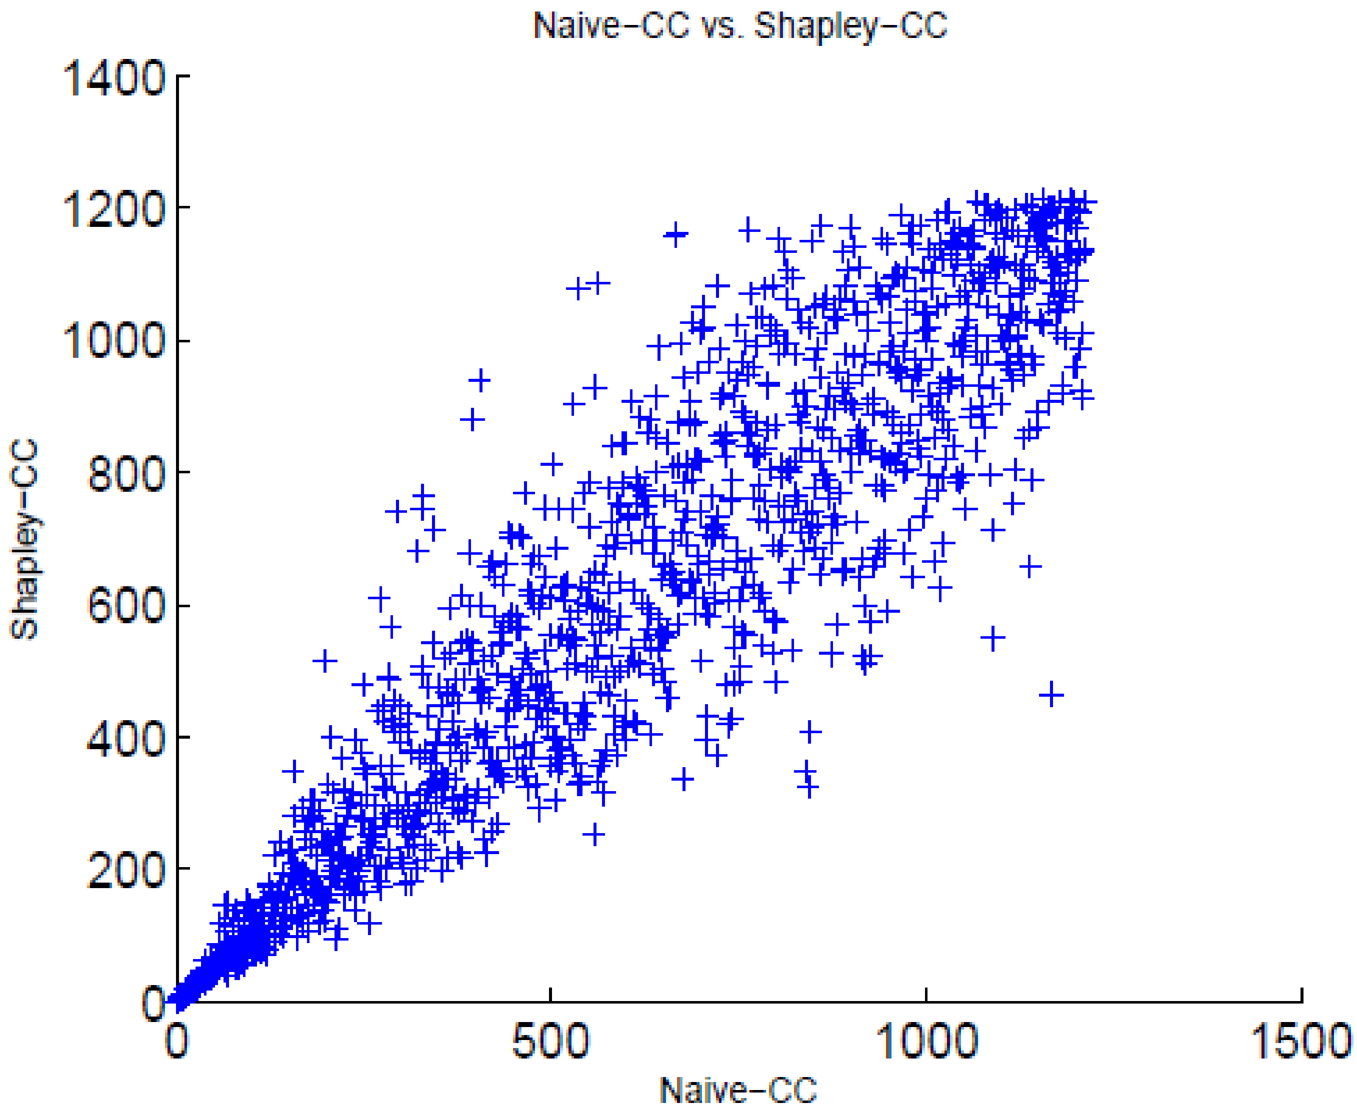

Naïve CC vs. Shapley CC