UniCredit

UNIVERSITÀ DELLA CALABRIA
DIPARTIMENTO DI INGEGNERIA INFORMATICA, MODELLISTICA, ELETTRONICA E SISTEMISTICA
DIMES

UNIVERSITÀ DELLA CALABRIA

AI & Data Science Lab
@DIMES Dept.

MACHINES, LANGUAGES & NETWORKS TEAM

The views and opinions expressed in this paper are those of the author and do not necessarily reflect the official policy or position of the UniCredit group.

A Combinatorial Multi-Armed Bandit Approach to Correlation Clustering

Francesco Gullo*, Domenico Mandaglio, Andrea Tagarelli

UniCredit	DIMES – Univ. Calabria	DIMES – Univ. Calabria
Rome, Italy	Rende (CS), Italy	Rende (CS), Italy

ECML PKDD 2023

Main Contributions

- Novel formulation of Correlation Clustering (CC) within a reinforcement learning setting by designing a Combinatorial Multi-Armed Bandit (CMAB) framework for correlation clustering
  - The CMAB-CC problem
- Design and theoretical analysis of algorithms
  - We devise a principled regret definition for our problem
- Extensive experimental evaluation

Min-Disagreement Correlation Clustering (Min-CC)

Input:

- an *undirected graph* $G = (V, E)$, with vertex set V and edge set $E \subseteq V \times V$
- *weights* $w_{uv}^{+}, w_{uv}^{-} \in \mathbb{R}_0^{+}$ for all edges $(u, v) \in E$, where any w_{uv}^{+} (resp. w_{uv}^{-}) weight expresses the benefit of clustering u and v together (resp. separately)

Output:

- a *clustering* $\mathcal{C}^{*}: V \rightarrow \mathbb{N}^{+}$ that:

$$\mathcal{C}^{*} = \operatorname{argmin}_{\mathcal{C}}\, d(\mathcal{C}) = \operatorname{argmin}_{\mathcal{C}} \sum_{\substack{(u,v)\in E \\ \mathcal{C}(u)=\mathcal{C}(v)}} w_{uv}^{-} + \sum_{\substack{(u,v)\in E \\ \mathcal{C}(u)\neq\mathcal{C}(v)}} w_{uv}^{+}$$

Correlation Clustering with Unknown Edge Weights

- Traditionally, in correlation clustering it is assumed that the *edge weights are all given as input*, e.g. derived from past user-interaction history, experimental trials etc.
  - Disadvantage: clustering has to be performed after that all the weights are available
- We focus for the first time on a correlation-clustering setting where *the edge weights are not available and edge-weight assessment is carried out while performing (multiple rounds of) clustering*

Correlation Clustering with Unknown Edge Weights

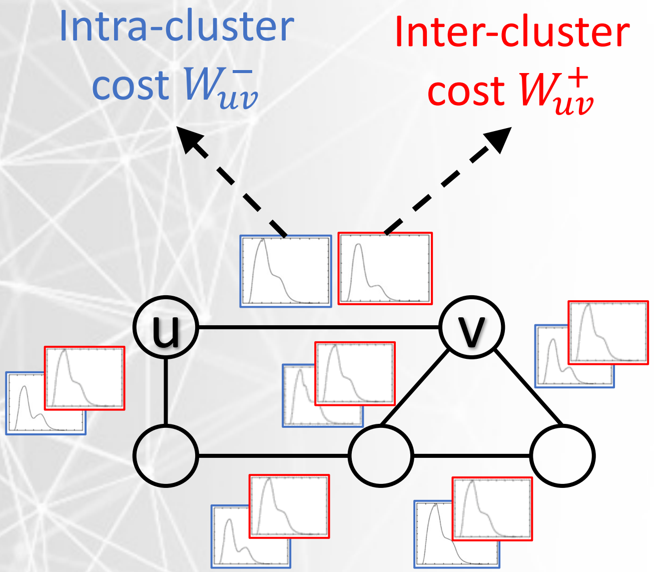
Intra-cluster cost W_{uv}^{-}

Inter-cluster cost W_{uv}^{+}

Edge weights w_{uv}^{-}, w_{uv}^{+} are modeled as **random variables** W_{uv}^{-}, W_{uv}^{+} with **means** $\mu_{uv}^{-} = \mathbb{E}[W_{uv}^{-}]$, $\mu_{uv}^{+} = \mathbb{E}[W_{uv}^{+}]$

Correlation Clustering with Unknown Edge Weights

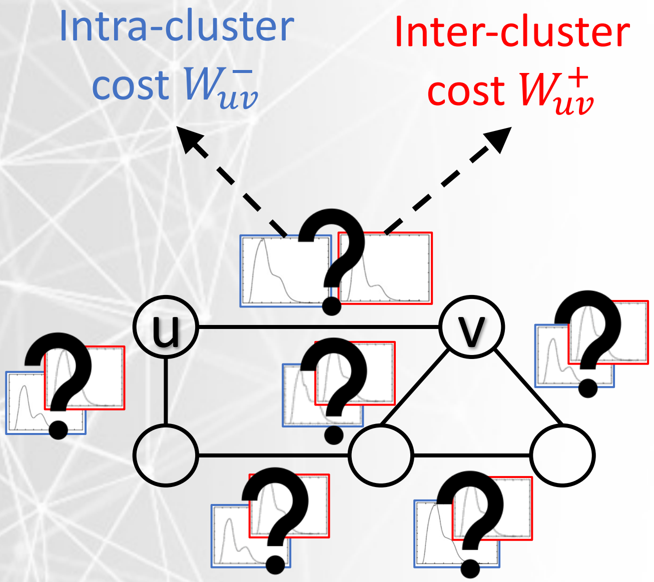

Intra-cluster cost W^-_{uv}

Inter-cluster cost W^+_{uv}

Random variables W^-_{uv}, W^+_{uv} and their means $\mu^-_{uv} = \mathbb{E}[W^-_{uv}]$, $\mu^+_{uv} = \mathbb{E}[W^+_{uv}]$ are **unknown**

Correlation Clustering with Unknown Edge Weights

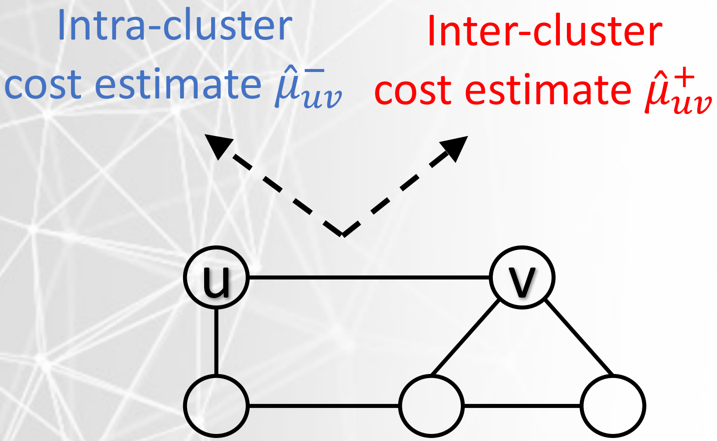

Intra-cluster cost estimate $\hat{\mu}_{uv}^{-}$

Inter-cluster cost estimate $\hat{\mu}_{uv}^{+}$

Estimates of the mean of the edge-weight distributions $\hat{\mu}_{uv}^{-}$, $\hat{\mu}_{uv}^{+}$ are maintained for each $(u, v) \in E$

Correlation Clustering with Unknown Edge Weights

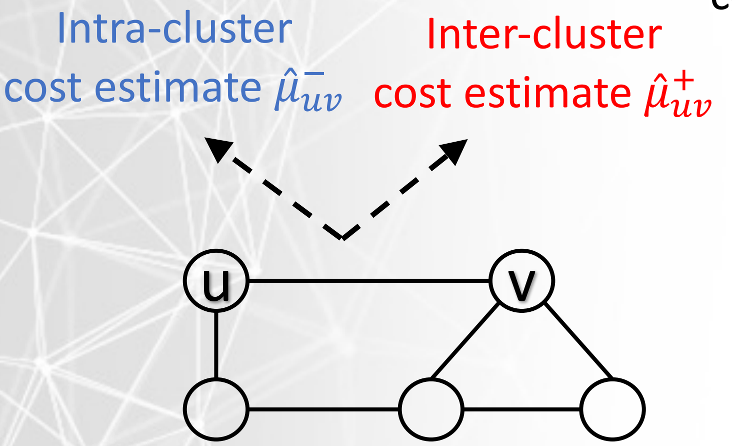

Intra-cluster cost estimate $\hat{\mu}_{uv}^{-}$

Inter-cluster cost estimate $\hat{\mu}_{uv}^{+}$

(i) Use an oracle $\boldsymbol{\mathcal{O}}$ (CC algorithm) with estimated weights to compute a clustering

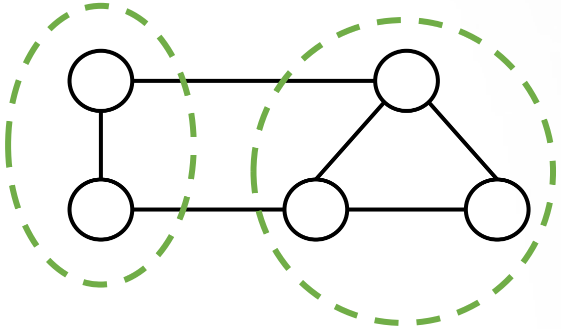

Estimates of the mean of the edge-weight distributions $\hat{\mu}_{uv}^{-}$, $\hat{\mu}_{uv}^{+}$ are maintained for each $(u, v) \in E$

Correlation Clustering with Unknown Edge Weights

Intra-cluster cost estimate $\hat{\mu}_{uv}^{-}$

Inter-cluster cost estimate $\hat{\mu}_{uv}^{+}$

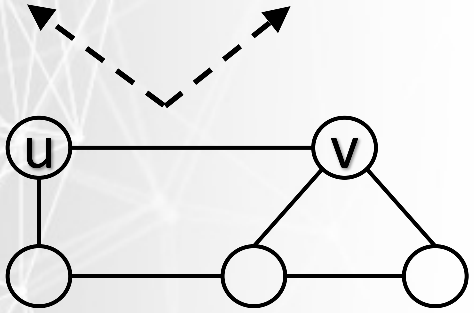

Estimates of the mean of the edge-weight distributions $\hat{\mu}_{uv}^{-}$, $\hat{\mu}_{uv}^{+}$ are maintained for each $(u, v) \in E$

(i) Use an oracle $\boldsymbol{\mathcal{O}}$ (CC algorithm) with estimated weights to compute a clustering

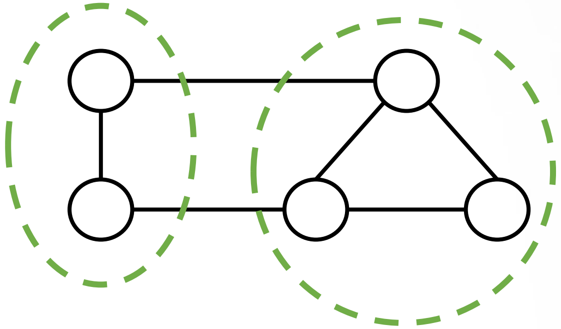

(ii) Placing the clustering gives ***feedback*** about the unknown edge weights W_{uv}^{-}, W_{uv}^{+} and the quality of the clustering

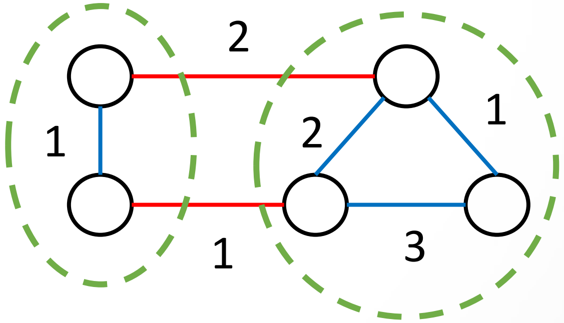

Correlation Clustering with Unknown Edge Weights

Intra-cluster cost estimate $\hat{\mu}^-_{uv}$

Inter-cluster cost estimate $\hat{\mu}^+_{uv}$

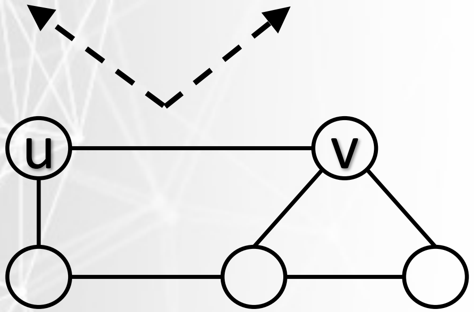

(i) Use an oracle $\mathcal{O}$ (CC algorithm) with estimated weights to compute a clustering

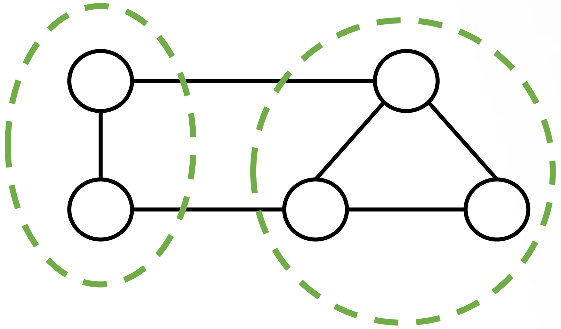

(ii) Placing the clustering gives ***feedback*** about the unknown edge weights W^-_{uv}, W^+_{uv} and the quality of the clustering

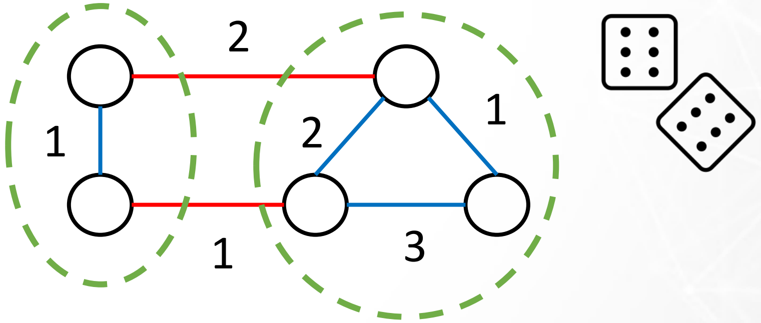

(iii) Update the mean estimates $\hat{\mu}^-_{uv}$, $\hat{\mu}^+_{uv}$ and repeat from (i) for T rounds

Estimates of the mean of the edge-weight distributions $\hat{\mu}^-_{uv}$, $\hat{\mu}^+_{uv}$ are maintained for each $(u, v) \in E$

Applications

- Team formation
- Recommendations in online social platforms
- Task allocation
- Commercial scheduling in slots
- Shelf space allocation

The CMAB-MIN-CC Problem

Given a graph $G = (V, E)$ and a number $T > 0$ of rounds, for every $t = 1, \dots, T$ find a clustering $\mathcal{C}_t \colon V \to \mathbb{N}^+$ so as to minimize

$$\mathbb{E}\left[\sum_{t=1}^{T} \bar{d}_{\boldsymbol{\mu}}(\mathcal{C}_t)\right]$$

$\bar{d}_{\boldsymbol{\mu}}(\mathcal{C}_t)$ is the expected disagreement (cost) of the clustering $\mathcal{C}_t$ according to the true (unknown) means $\boldsymbol{\mu} = \left\{\{\mu_{uv}^{+}\}_{(u,v)\in E}, \{\mu_{uv}^{-}\}_{(u,v)\in E}\right\}$

Exploration-Exploitation Trade-Off

A clustering at every round may be computed

1. by taking into account the current mean estimates based on an *approximate oracle* (***exploitation***)
2. without looking at the mean estimates, so as to get feedback on edge weights for which limited knowledge has been acquired so far (***exploration***)

Objective: getting the best *exploration-exploitation tradeoff* whose effectiveness is measured by the (expected) *cumulative quality of the clusterings produced in all the rounds*.

Combinatorial Multi-Armed Bandit[3] (CMAB)

- $\mathcal{A}$ is the set of m base slot-machines/arms to choose from
- Each arm i is assigned a set $\{X_{i,t} \mid 1 \leq i \leq m, 1 \leq t \leq T\}$ of random variables where each $X_{i,t}$ indicates the random "outcome" of playing base arm i in round t.
- At each step t the agent selects/plays a ***subset of base arms*** (super arm) $A_t \subseteq \mathcal{A}$ and the outcomes of the random variables $X_{i,t}$, for all the base arms $j \in A_t$, are observed.
- Playing a superarm A_t gives a reward $R_t(A_t)$, which is a random variable defined as a function of the outcomes of A_t's base arms.
- It is assumed the availability of an $(\boldsymbol{\alpha}, \boldsymbol{\beta})$***-approximation oracle*** that, for some $\alpha, \beta \leq 1$, it outputs a superarm A_t so that $$\Pr\left[\mathbb{E}\left[\hat{R}_t(A_t)\right] \geq \alpha \mathbb{E}\left[\hat{R}_t(A_t^*)\right]\right] \geq \beta$$
- The goal is to maximize the (expected) cumulative reward $\mathbb{E}[\sum_{t=1}^{T} R_t(A_t)]$ by a proper exploration/exploitation trade-off

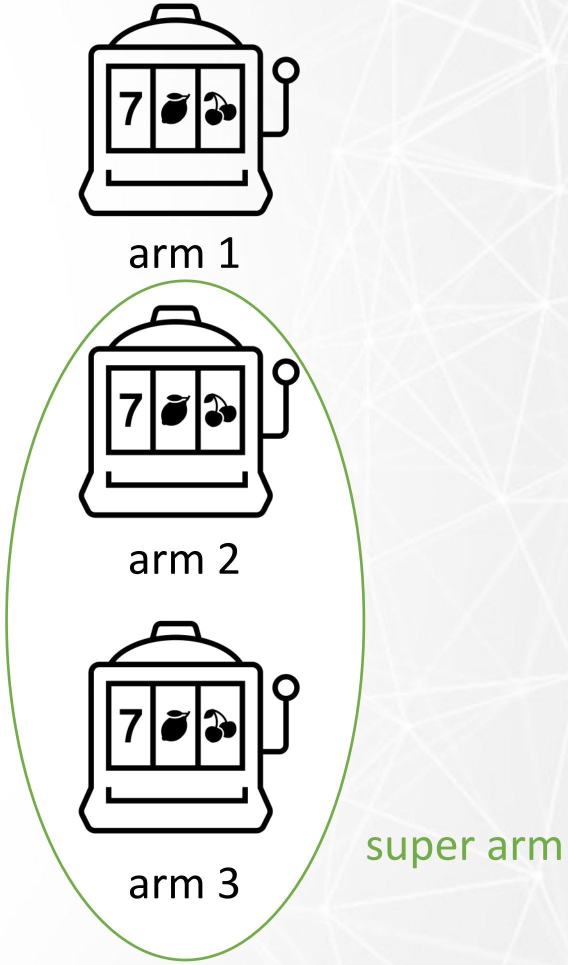

[3] Chen Wei et al. "Combinatorial multi-armed bandit and its extension to probabilistically triggered arms." The Journal of Machine Learning Research 17.1 (2016): 1746-1778.

CMAB-MIN-CC as a CMAB instance

- *base arms*: each edge $e = (u, v) \in E$ has a pair of replicas, e^{in} and e^{out} ($m = |\mathcal{A}| = 2|E|$)
- *superarms*: sets $A \subseteq \mathcal{A}$ that are consistent with the notion of clustering
  - for all $e \in E$, A only contains e^{in} or e^{out} ($|A| = |E|$)
  - for all $e_1 = (x, y), e_2 = (y, z), e_3 = (x, z) \in E$, if $e_1^{in}, e_2^{in} \in A$, then $e_3^{in} \in A$
- *base-arm outcome*:
  - *Intra-cluster base arm e^{in}:* a sample from W_{uv}^{-}
  - *Inter-cluster base arm e^{out}:* a sample from W_{uv}^{+}
- *loss*: $d(A) = \sum_{e \in A^{in}} W_{uv}^{-} + \sum_{e \in A^{out}} W_{uv}^{+}$

Algorithms for CMAB-MIN-CC

Adaptation of well-established CMAB algorithms to the correlation-clustering context:

1. focus on the context of *general oracles* for MIN-CC
2. case where the employed MIN-CC oracles achieve theoretical guarantees only if the input (unknown) weights meet certain properties
3. the special case of input edge-weight distributions satisfying specific constraints

The Correlation Clustering - Combinatorial Lower Confidence Bound (CC-CLCB) algorithm

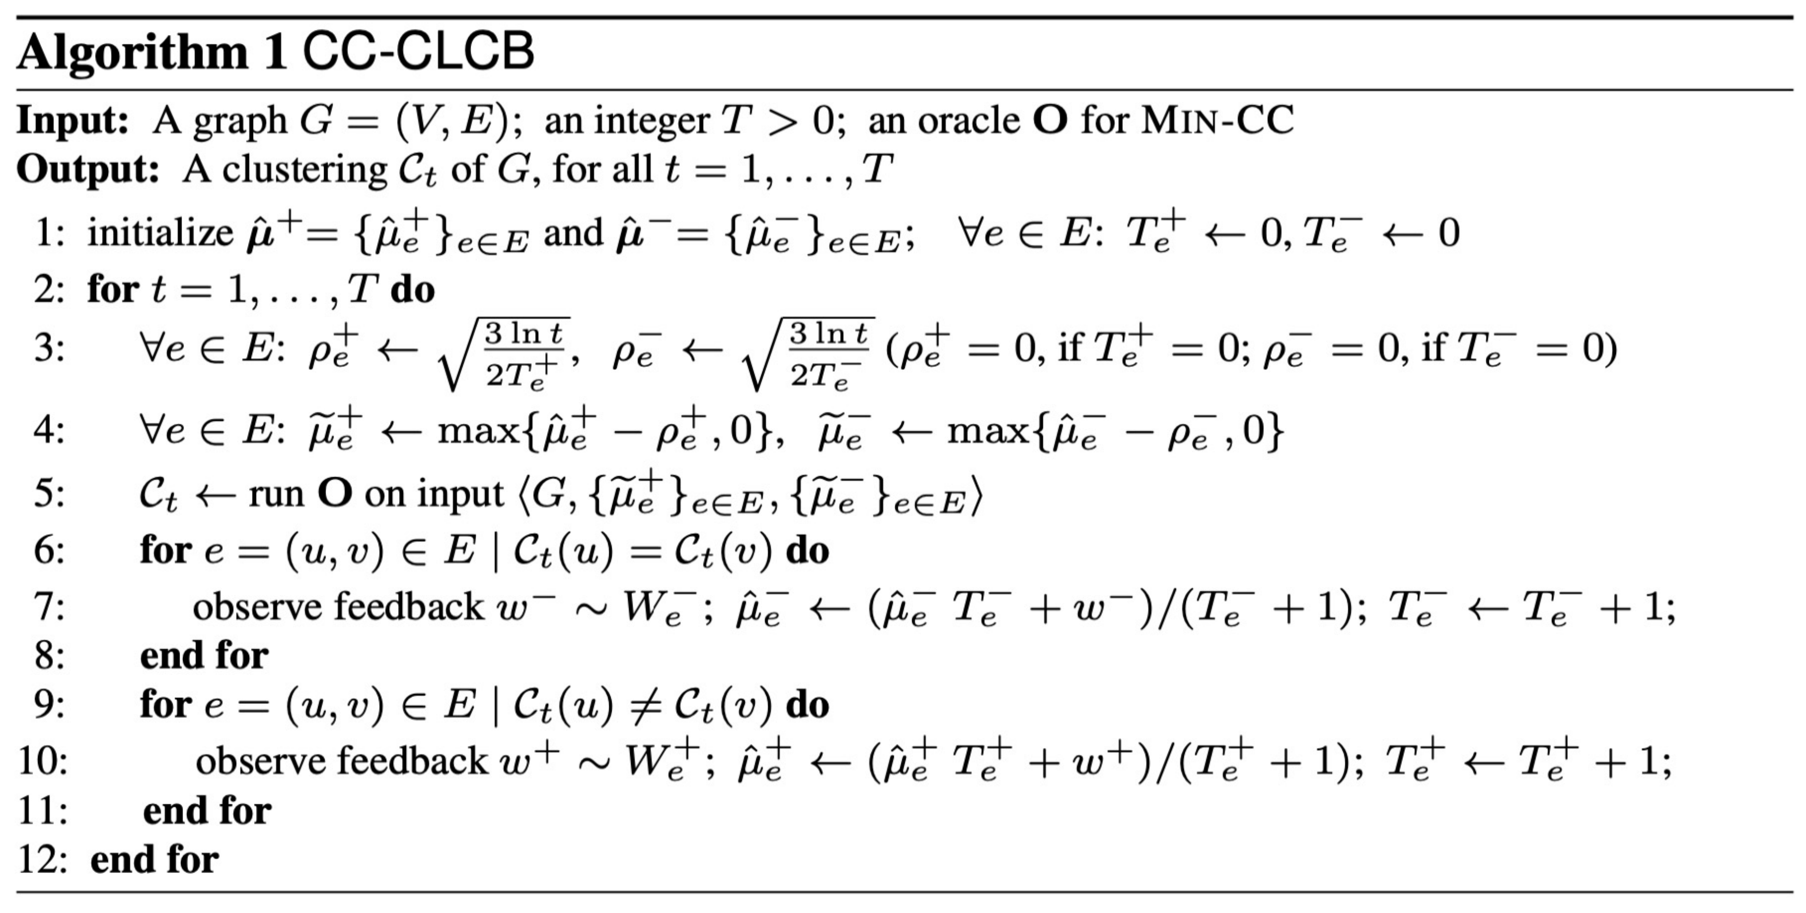

Algorithm 1 CC-CLCB

Input: A graph $G = (V, E)$; an integer $T > 0$; an oracle $\mathbf{O}$ for MIN-CC

Output: A clustering $\mathcal{C}_t$ of G, for all $t = 1, \dots, T$

```
initialize $\hat{\boldsymbol{\mu}}^+ = \{\hat{\mu}_e^+\}_{e \in E}$ and $\hat{\boldsymbol{\mu}}^- = \{\hat{\mu}_e^-\}_{e \in E}$;   $\forall e \in E$: $T_e^+ \leftarrow 0, T_e^- \leftarrow 0$
for $t = 1, \dots, T$ do
    $\forall e \in E$: $\rho_e^+ \leftarrow \sqrt{\frac{3 \ln t}{2T_e^+}}$, $\rho_e^- \leftarrow \sqrt{\frac{3 \ln t}{2T_e^-}}$ $(\rho_e^+ = 0$, if $T_e^+ = 0$; $\rho_e^- = 0$, if $T_e^- = 0)$
    $\forall e \in E$: $\tilde{\mu}_e^+ \leftarrow \max\{\hat{\mu}_e^+ - \rho_e^+, 0\}$, $\tilde{\mu}_e^- \leftarrow \max\{\hat{\mu}_e^- - \rho_e^-, 0\}$
    $\mathcal{C}_t \leftarrow$ run $\mathbf{O}$ on input $\langle G, \{\tilde{\mu}_e^+\}_{e \in E}, \{\tilde{\mu}_e^-\}_{e \in E} \rangle$
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) = \mathcal{C}_t(v)$ do
        observe feedback $w^- \sim W_e^-$; $\hat{\mu}_e^- \leftarrow (\hat{\mu}_e^- T_e^- + w^-)/(T_e^- + 1)$; $T_e^- \leftarrow T_e^- + 1$;
    end for
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) \neq \mathcal{C}_t(v)$ do
        observe feedback $w^+ \sim W_e^+$; $\hat{\mu}_e^+ \leftarrow (\hat{\mu}_e^+ T_e^+ + w^+)/(T_e^+ + 1)$; $T_e^+ \leftarrow T_e^+ + 1$;
    end for
end for
```

The Correlation Clustering - Combinatorial Lower Confidence Bound (CC-CLCB) algorithm

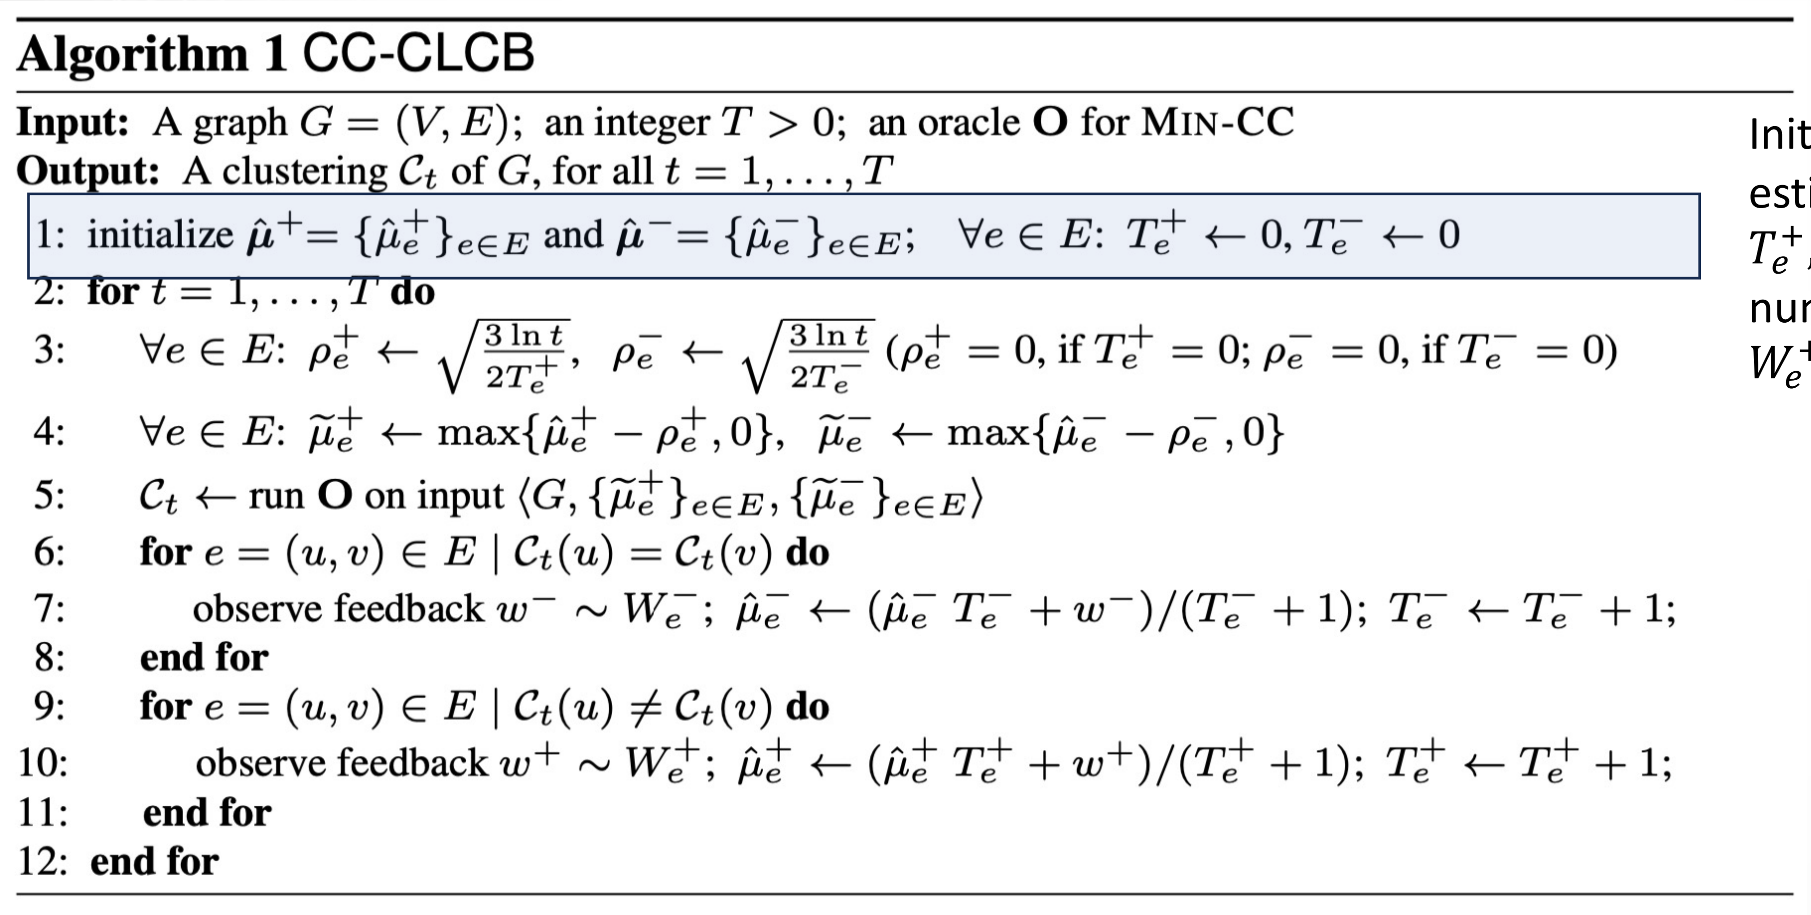

Algorithm 1 CC-CLCB

Input: A graph $G = (V, E)$; an integer $T > 0$; an oracle $\mathbf{O}$ for MIN-CC
Output: A clustering $\mathcal{C}_t$ of G, for all $t = 1, \dots, T$

```
initialize $\hat{\boldsymbol{\mu}}^+ = \{\hat{\mu}_e^+\}_{e \in E}$ and $\hat{\boldsymbol{\mu}}^- = \{\hat{\mu}_e^-\}_{e \in E}$;   $\forall e \in E$: $T_e^+ \leftarrow 0, T_e^- \leftarrow 0$
for $t = 1, \dots, T$ do
    $\forall e \in E$: $\rho_e^+ \leftarrow \sqrt{\frac{3 \ln t}{2T_e^+}}$,  $\rho_e^- \leftarrow \sqrt{\frac{3 \ln t}{2T_e^-}}$ ($\rho_e^+ = 0$, if $T_e^+ = 0$; $\rho_e^- = 0$, if $T_e^- = 0$)
    $\forall e \in E$: $\tilde{\mu}_e^+ \leftarrow \max\{\hat{\mu}_e^+ - \rho_e^+, 0\}$,  $\tilde{\mu}_e^- \leftarrow \max\{\hat{\mu}_e^- - \rho_e^-, 0\}$
    $\mathcal{C}_t \leftarrow$ run $\mathbf{O}$ on input $\langle G, \{\tilde{\mu}_e^+\}_{e \in E}, \{\tilde{\mu}_e^-\}_{e \in E} \rangle$
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) = \mathcal{C}_t(v)$ do
        observe feedback $w^- \sim W_e^-$; $\hat{\mu}_e^- \leftarrow (\hat{\mu}_e^- T_e^- + w^-)/(T_e^- + 1)$; $T_e^- \leftarrow T_e^- + 1$;
    end for
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) \neq \mathcal{C}_t(v)$ do
        observe feedback $w^+ \sim W_e^+$; $\hat{\mu}_e^+ \leftarrow (\hat{\mu}_e^+ T_e^+ + w^+)/(T_e^+ + 1)$; $T_e^+ \leftarrow T_e^+ + 1$;
    end for
end for
```

Initialization of the mean estimates $\hat{\boldsymbol{\mu}}^+, \hat{\boldsymbol{\mu}}^-$ and counters T_e^+, T_e^- which denote the number of times a sample from W_e^+, W_e^- has been observed

The Correlation Clustering- Combinatorial Lower Confidence Bound (CC-CLCB) algorithm

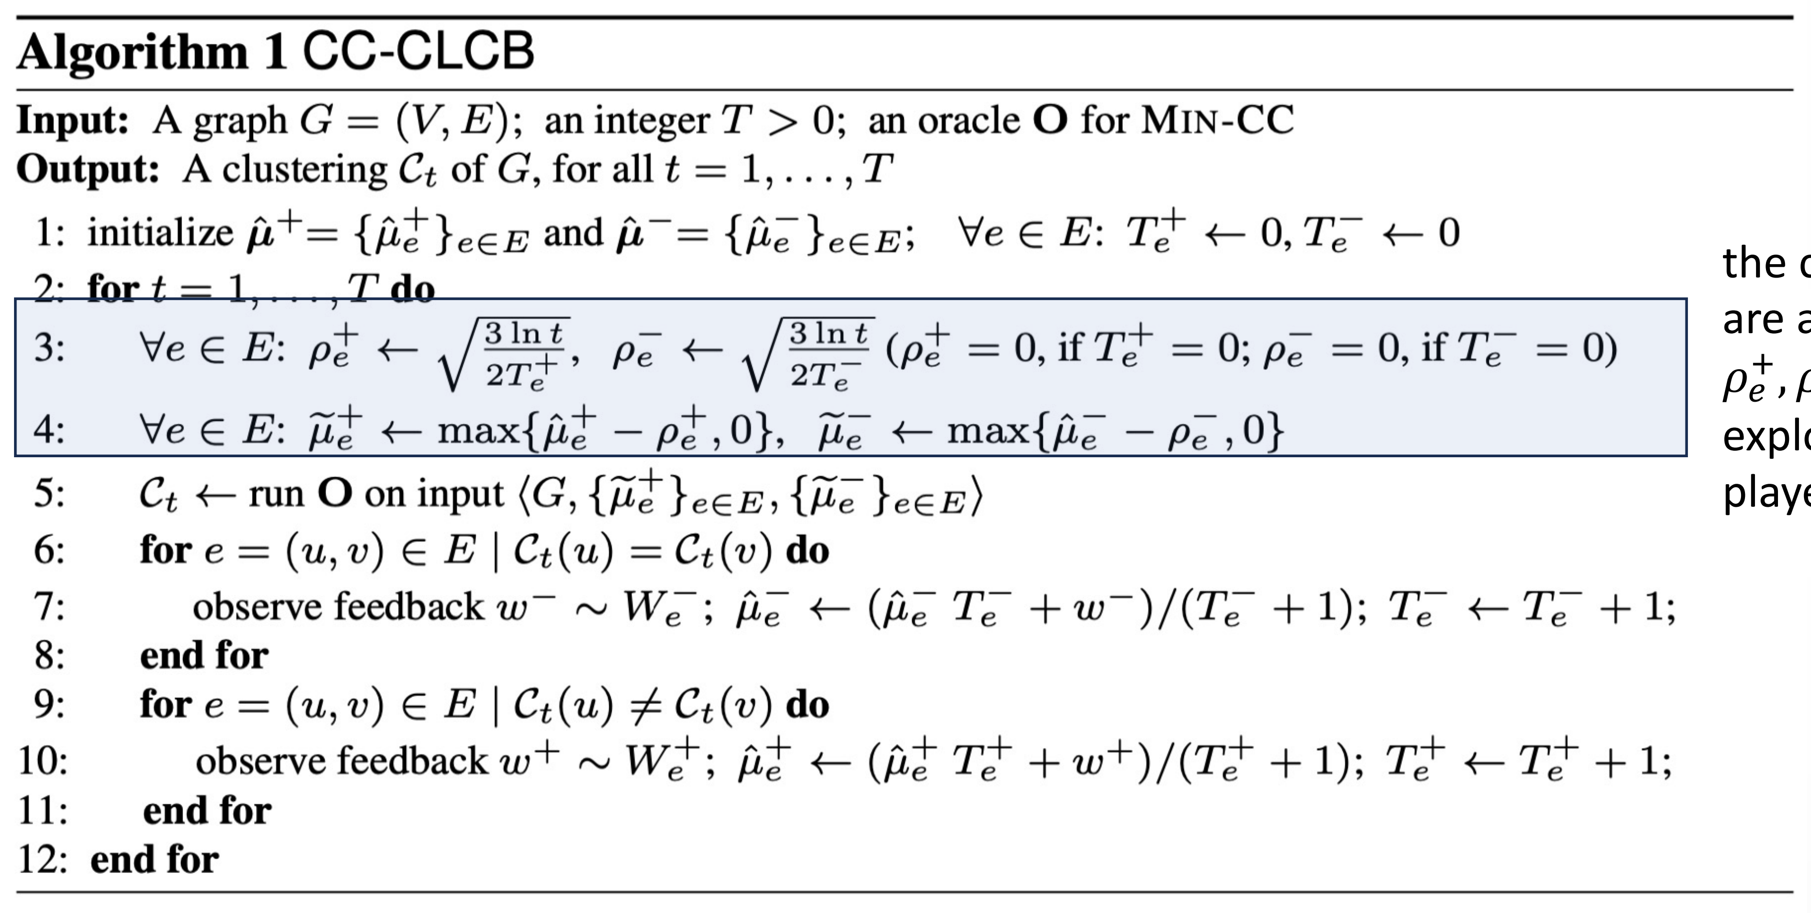

Algorithm 1 CC-CLCB

Input: A graph $G = (V, E)$; an integer $T > 0$; an oracle $\mathbf{O}$ for MIN-CC
Output: A clustering $\mathcal{C}_t$ of G, for all $t = 1, \ldots, T$

```
initialize $\hat{\boldsymbol{\mu}}^+ = \{\hat{\mu}_e^+\}_{e \in E}$ and $\hat{\boldsymbol{\mu}}^- = \{\hat{\mu}_e^-\}_{e \in E}$;   $\forall e \in E$: $T_e^+ \leftarrow 0, T_e^- \leftarrow 0$
for $t = 1, \ldots, T$ do
    $\forall e \in E$: $\rho_e^+ \leftarrow \sqrt{\frac{3 \ln t}{2T_e^+}}$,  $\rho_e^- \leftarrow \sqrt{\frac{3 \ln t}{2T_e^-}}$ $(\rho_e^+ = 0$, if $T_e^+ = 0$; $\rho_e^- = 0$, if $T_e^- = 0)$
    $\forall e \in E$: $\widetilde{\mu}_e^+ \leftarrow \max\{\hat{\mu}_e^+ - \rho_e^+, 0\}$,  $\widetilde{\mu}_e^- \leftarrow \max\{\hat{\mu}_e^- - \rho_e^-, 0\}$
    $\mathcal{C}_t \leftarrow$ run $\mathbf{O}$ on input $\langle G, \{\widetilde{\mu}_e^+\}_{e \in E}, \{\widetilde{\mu}_e^-\}_{e \in E} \rangle$
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) = \mathcal{C}_t(v)$ do
        observe feedback $w^- \sim W_e^-$; $\hat{\mu}_e^- \leftarrow (\hat{\mu}_e^- T_e^- + w^-)/(T_e^- + 1)$; $T_e^- \leftarrow T_e^- + 1$;
    end for
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) \neq \mathcal{C}_t(v)$ do
        observe feedback $w^+ \sim W_e^+$; $\hat{\mu}_e^+ \leftarrow (\hat{\mu}_e^+ T_e^+ + w^+)/(T_e^+ + 1)$; $T_e^+ \leftarrow T_e^+ + 1$;
    end for
end for
```

the current mean estimates are adjusted with the terms ρ_e^+, ρ_e^-, so as to foster the exploration of less often played base arms

The Correlation Clustering - Combinatorial Lower Confidence Bound (CC-CLCB) algorithm

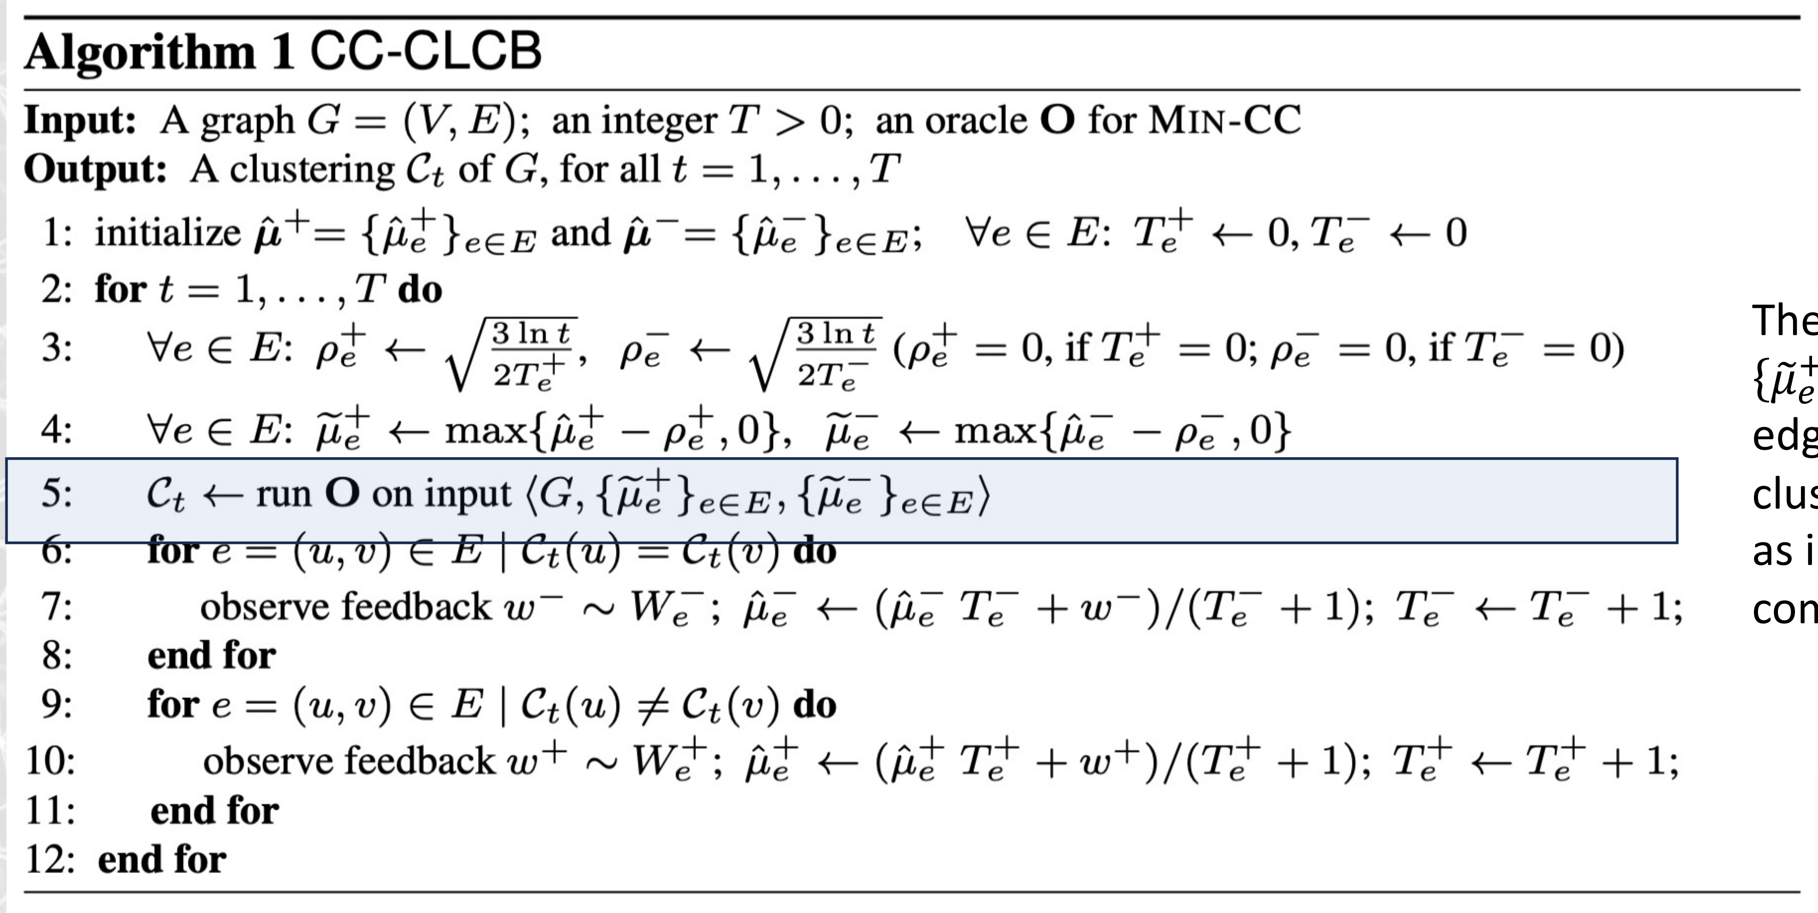

Algorithm 1 CC-CLCB

Input: A graph $G = (V, E)$; an integer $T > 0$; an oracle $\mathbf{O}$ for MIN-CC
Output: A clustering $\mathcal{C}_t$ of G, for all $t = 1, \dots, T$

```
initialize $\hat{\boldsymbol{\mu}}^{+} = \{\hat{\mu}_e^{+}\}_{e \in E}$ and $\hat{\boldsymbol{\mu}}^{-} = \{\hat{\mu}_e^{-}\}_{e \in E}$;   $\forall e \in E$: $T_e^{+} \leftarrow 0, T_e^{-} \leftarrow 0$
for $t = 1, \dots, T$ do
    $\forall e \in E$: $\rho_e^{+} \leftarrow \sqrt{\frac{3 \ln t}{2T_e^{+}}}$,  $\rho_e^{-} \leftarrow \sqrt{\frac{3 \ln t}{2T_e^{-}}}$ ($\rho_e^{+} = 0$, if $T_e^{+} = 0$; $\rho_e^{-} = 0$, if $T_e^{-} = 0$)
    $\forall e \in E$: $\tilde{\mu}_e^{+} \leftarrow \max\{\hat{\mu}_e^{+} - \rho_e^{+}, 0\}$,  $\tilde{\mu}_e^{-} \leftarrow \max\{\hat{\mu}_e^{-} - \rho_e^{-}, 0\}$
    $\mathcal{C}_t \leftarrow$ run $\mathbf{O}$ on input $\langle G, \{\tilde{\mu}_e^{+}\}_{e \in E}, \{\tilde{\mu}_e^{-}\}_{e \in E} \rangle$
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) = \mathcal{C}_t(v)$ do
        observe feedback $w^{-} \sim W_e^{-}$;  $\hat{\mu}_e^{-} \leftarrow (\hat{\mu}_e^{-} T_e^{-} + w^{-})/(T_e^{-} + 1)$;  $T_e^{-} \leftarrow T_e^{-} + 1$;
    end for
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) \neq \mathcal{C}_t(v)$ do
        observe feedback $w^{+} \sim W_e^{+}$;  $\hat{\mu}_e^{+} \leftarrow (\hat{\mu}_e^{+} T_e^{+} + w^{+})/(T_e^{+} + 1)$;  $T_e^{+} \leftarrow T_e^{+} + 1$;
    end for
end for
```

The adjusted means $\{\tilde{\mu}_e^{+}, \tilde{\mu}_e^{-}\}_{e \in E}$ are interpreted as edge weights of a correlation-clustering instance and are fed as input to an oracle O that computes a MIN-CC solution $\mathcal{C}_t$

The Correlation Clustering - Combinatorial Lower Confidence Bound (CC-CLCB) algorithm

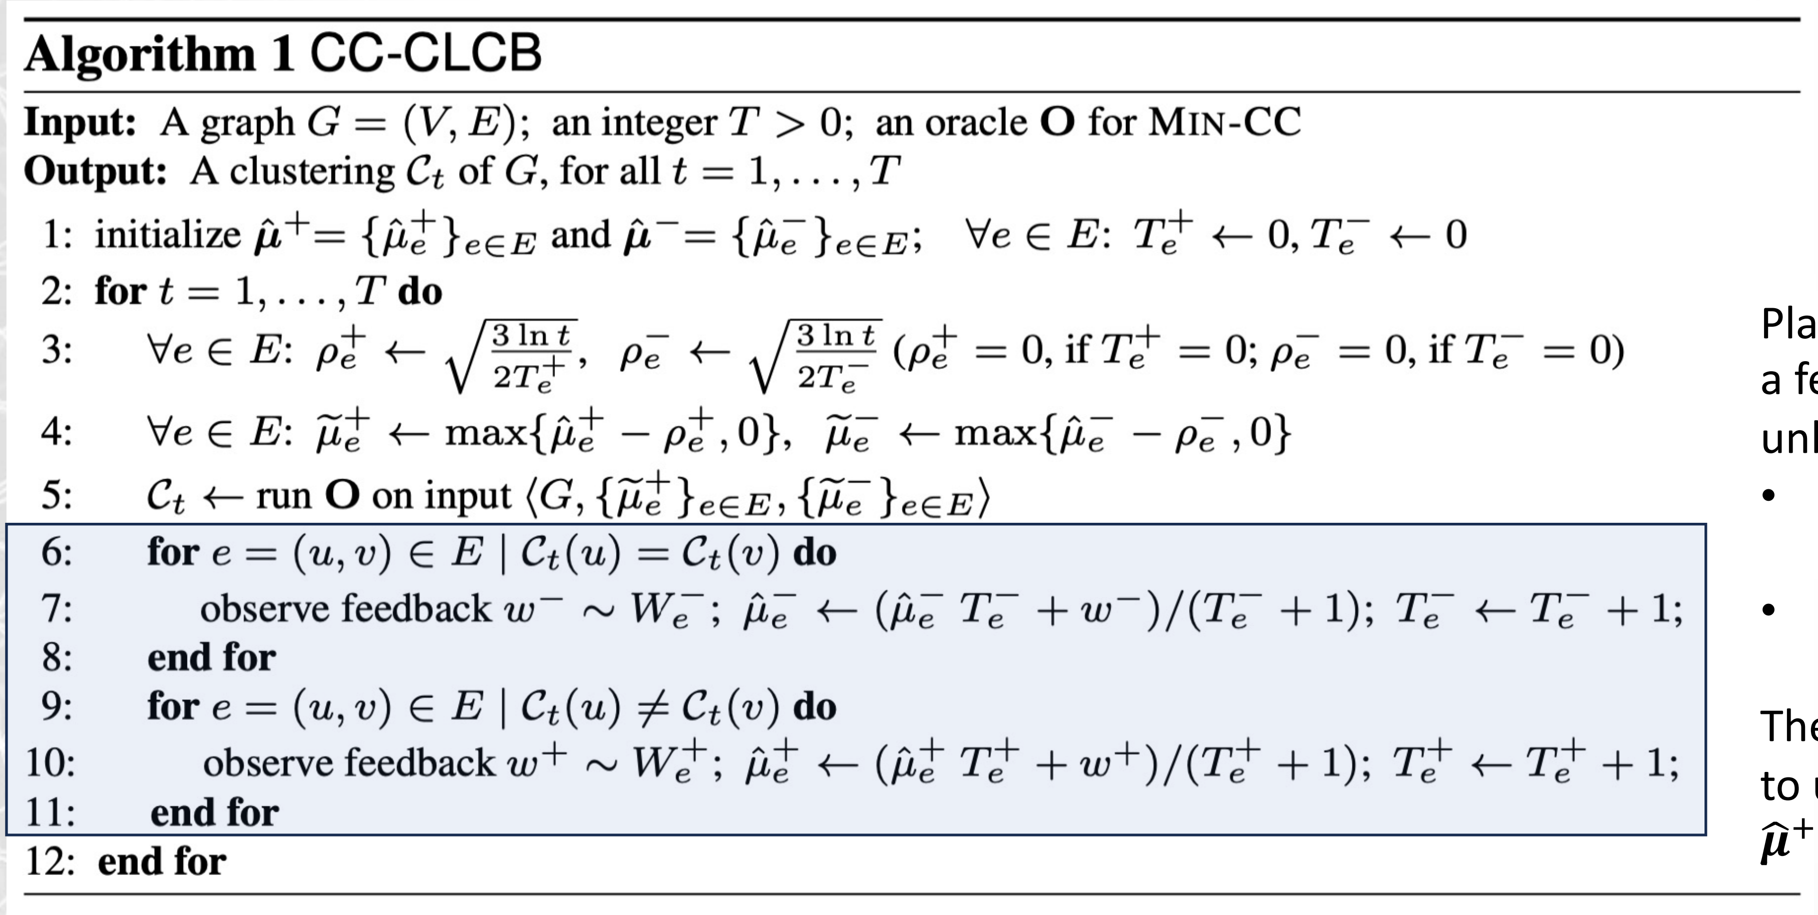

Algorithm 1 CC-CLCB

Input: A graph $G = (V, E)$; an integer $T > 0$; an oracle $\mathbf{O}$ for MIN-CC

Output: A clustering $\mathcal{C}_t$ of G, for all $t = 1, \ldots, T$

```
initialize $\hat{\boldsymbol{\mu}}^+ = \{\hat{\mu}_e^+\}_{e \in E}$ and $\hat{\boldsymbol{\mu}}^- = \{\hat{\mu}_e^-\}_{e \in E}$;   $\forall e \in E$: $T_e^+ \leftarrow 0, T_e^- \leftarrow 0$
for $t = 1, \ldots, T$ do
    $\forall e \in E$: $\rho_e^+ \leftarrow \sqrt{\frac{3 \ln t}{2T_e^+}}$,  $\rho_e^- \leftarrow \sqrt{\frac{3 \ln t}{2T_e^-}}$ ($\rho_e^+ = 0$, if $T_e^+ = 0$; $\rho_e^- = 0$, if $T_e^- = 0$)
    $\forall e \in E$: $\tilde{\mu}_e^+ \leftarrow \max\{\hat{\mu}_e^+ - \rho_e^+, 0\}$,  $\tilde{\mu}_e^- \leftarrow \max\{\hat{\mu}_e^- - \rho_e^-, 0\}$
    $\mathcal{C}_t \leftarrow$ run $\mathbf{O}$ on input $\langle G, \{\tilde{\mu}_e^+\}_{e \in E}, \{\tilde{\mu}_e^-\}_{e \in E} \rangle$
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) = \mathcal{C}_t(v)$ do
        observe feedback $w^- \sim W_e^-$; $\hat{\mu}_e^- \leftarrow (\hat{\mu}_e^- T_e^- + w^-)/(T_e^- + 1)$; $T_e^- \leftarrow T_e^- + 1$;
    end for
    for $e = (u, v) \in E \mid \mathcal{C}_t(u) \neq \mathcal{C}_t(v)$ do
        observe feedback $w^+ \sim W_e^+$; $\hat{\mu}_e^+ \leftarrow (\hat{\mu}_e^+ T_e^+ + w^+)/(T_e^+ + 1)$; $T_e^+ \leftarrow T_e^+ + 1$;
    end for
end for
```

Placing the clustering $\mathcal{C}_t$ gives a feedback about the unknown W_e^+, W_e^-:

- *Intra-cluster edge:* a sample from W_e^-
- *Inter-cluster edge:* a sample from W_e^+

The yielded samples are used to update the mean estimates $\hat{\boldsymbol{\mu}}^+, \hat{\boldsymbol{\mu}}^-$

Regret Analysis of CC-CLCB algorithm

MIN-CC-$(\boldsymbol{\alpha}, \boldsymbol{\beta})$-approximation regret. Let $\mathcal{C}_I^*$ be the clustering minimizing the expected loss $\bar{d}_{\boldsymbol{\mu}}(\cdot)$ on a CMAB-MIN-CC instance I, let $\mathcal{M}$ be the the expected loss of the worst possible clustering on I.

$$Reg_{\boldsymbol{\mu},\alpha,\beta}(T) = \mathbb{E}\left[\sum_{t=1}^{T} \bar{d}_{\boldsymbol{\mu}}(\mathcal{C}_t)\right] - T\left[\frac{1}{\alpha}\beta \bar{d}_{\boldsymbol{\mu}}(\mathcal{C}_I^*) + (1-\beta)\mathcal{M}\right]$$

Theorem. Given $\alpha, \beta \in (0,1]$, the MIN-CC-(α, β)-approximation regret of the CC-CLCB algorithm, when equipped with a MIN-CC-(α, β)-approximation oracle is upper-bounded by a function of T that is $\mathcal{O}(\log T)$.

Heuristic Variants of CC-CLCB

Rationale. Favor the fulfilment of some constraints on the MIN-CC instances to be processed by the oracle to make the latter perform better

- **PC+Exp-CLCB** computes the adjusted means $\tilde{\mu}^{+}_{uv}, \tilde{\mu}^{-}_{uv}$ such that the *local constraint* $\tilde{\mu}^{+}_{uv} + \tilde{\mu}^{-}_{uv} = 1$ holds for every $u, v \in V$
- **Global-CLCB** computes the adjusted means $\tilde{\mu}^{+}_{uv}, \tilde{\mu}^{-}_{uv}$ such that the *global constraint* $\binom{|V|}{2}^{-1} \sum_{u,v \in V} (\tilde{\mu}^{+}_{uv} + \tilde{\mu}^{-}_{uv}) \geq 1$ is satisfied

Special Edge-Weight Distributions

Symmetric distributions. [0, 1]-support random variables W_e^+, W_e^- have symmetric distributions if and only if $W_e^+(x) = W_e^-(1-x)$, for all $x \in [0,1]$.

Theorem. Given $\alpha, \beta \in (0,1]$, the MIN-CC-(α, β)-approximation regret of a *pure exploitation* (PE) strategy run on a CMAB-MIN-CC instance where all edge-weight distributions are symmetric, and equipped with a MIN-CC-(α, β)-approximation oracle is upper-bounded by a function of T that is $\mathcal{O}(1)$.

Evaluation

Data

- Real network data with artificially-generated edge weights

Assessment criteria

- *Average expected normalized cumulative* MIN-CC *loss* (up to round t) $f^{(t)}$ w.r.t. the true (unknown) edge weights $\boldsymbol{\mu}$:

$$f^{(t)} = \frac{1}{t}\sum\nolimits_{i=1}^{t} \mathbb{E}\left[\frac{\bar{d}_{\boldsymbol{\mu}}(\mathcal{C}_i)}{U}\right], \qquad U = \sum_{u,v\in V} \max\{\mu_{uv}^{+}, \mu_{uv}^{-}\}$$

- *Relative error norm* (at round t) $ren^{(t)}$:

$$ren^{(t)} = \sqrt{\frac{\sum_{e\in E}(\mu_e^{+} - \hat{\mu}_{e,t}^{+})^2 + \sum_{e\in E}(\mu_e^{-} - \hat{\mu}_{e,t}^{-})^2}{\sum_{e\in E}(\mu_e^{+})^2 + \sum_{e\in E}(\mu_e^{-})^2}}$$

Evaluation

Evaluation goals

- Assess the performance of the CMAB methods (CC-CLCB, EG, EG-fixed, PE, CTS[1]) in terms of $f^{(T)}$ and $ren^{(T)}$ and compare them to non-CMAB baselines (Adamic-Adar, Jaccard) and the reference Actual-weight method
- Evaluate the impact of varying the MIN-CC oracle (Pivot[2], LP+R[3]) on the performance of the various CMAB methods

1. Wang Siwei, and Wei Chen. "Thompson sampling for combinatorial semi-bandits." International Conference on Machine Learning. PMLR, 2018.
2. Ailon Nir, Moses Charikar, and Alantha Newman. "Aggregating inconsistent information: ranking and clustering." Journal of the ACM (JACM) 55.5 (2008): 1-27.
3. Charikar Moses, Venkatesan Guruswami, and Anthony Wirth. "Clustering with qualitative information." Journal of Computer and System Sciences 71.3 (2005): 360-383.

Data

Table: Main characteristics of the real-world datasets used in our evaluation.

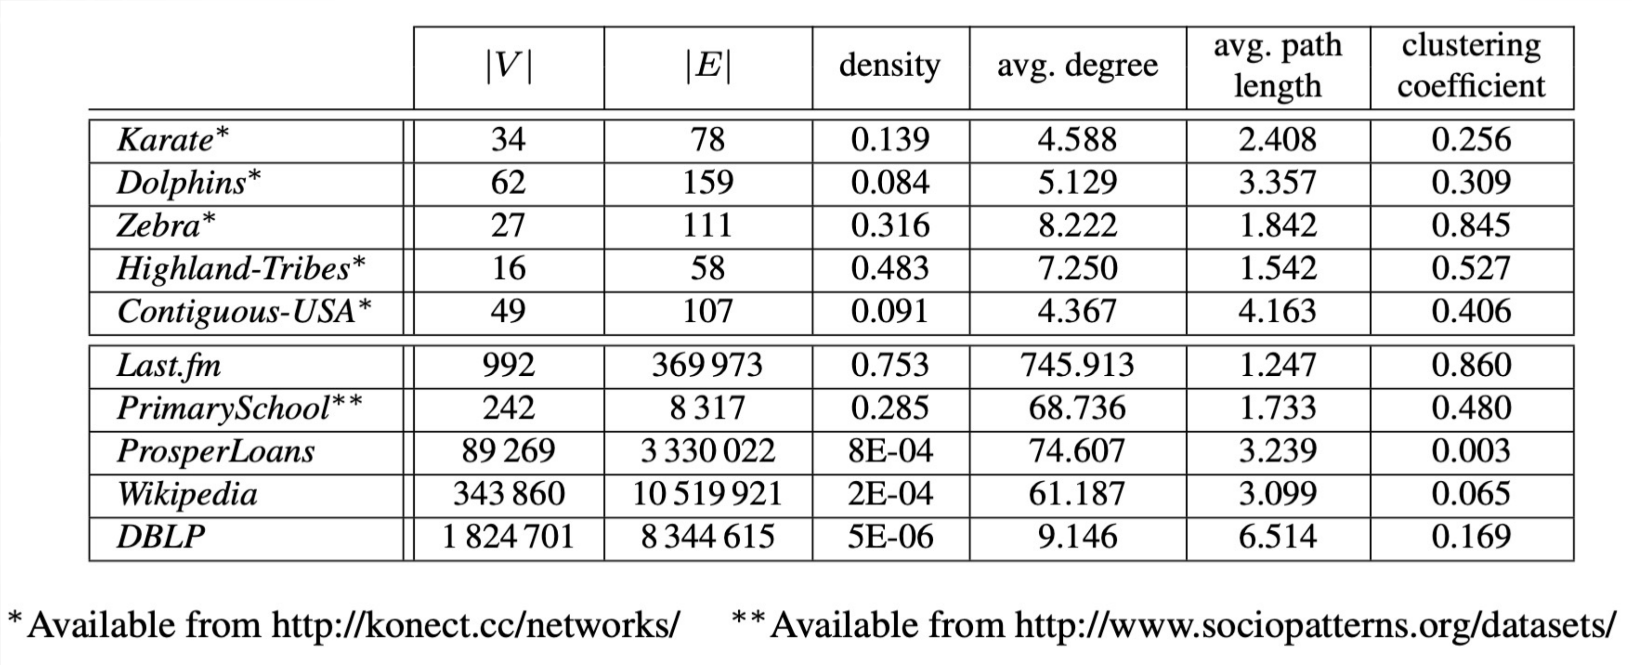

	$\lvert V \rvert$	$\lvert E \rvert$	density	avg. degree	avg. path length	clustering coefficient
*Karate**	34	78	0.139	4.588	2.408	0.256
*Dolphins**	62	159	0.084	5.129	3.357	0.309
*Zebra**	27	111	0.316	8.222	1.842	0.845
*Highland-Tribes**	16	58	0.483	7.250	1.542	0.527
*Contiguous-USA**	49	107	0.091	4.367	4.163	0.406
Last.fm	992	369 973	0.753	745.913	1.247	0.860
*PrimarySchool***	242	8 317	0.285	68.736	1.733	0.480
ProsperLoans	89 269	3 330 022	8E-04	74.607	3.239	0.003
Wikipedia	343 860	10 519 921	2E-04	61.187	3.099	0.065
DBLP	1 824 701	8 344 615	5E-06	9.146	6.514	0.169

*Available from http://konect.cc/networks/ **Available from http://www.sociopatterns.org/datasets/

Edge weight generation. The random variables W_e^+, W_e^- are assumed to follow a *Bernoulli* distribution, whose means are generated according to a scheme which ensures that the probability constraint holds on the generated means, i.e. first sample $\mu_e^+ \sim Uniform(0, 1)$, and then set $\mu_e^- = 1 - \mu_e^+$, for all $e \in E$.

Quality of the clusterings

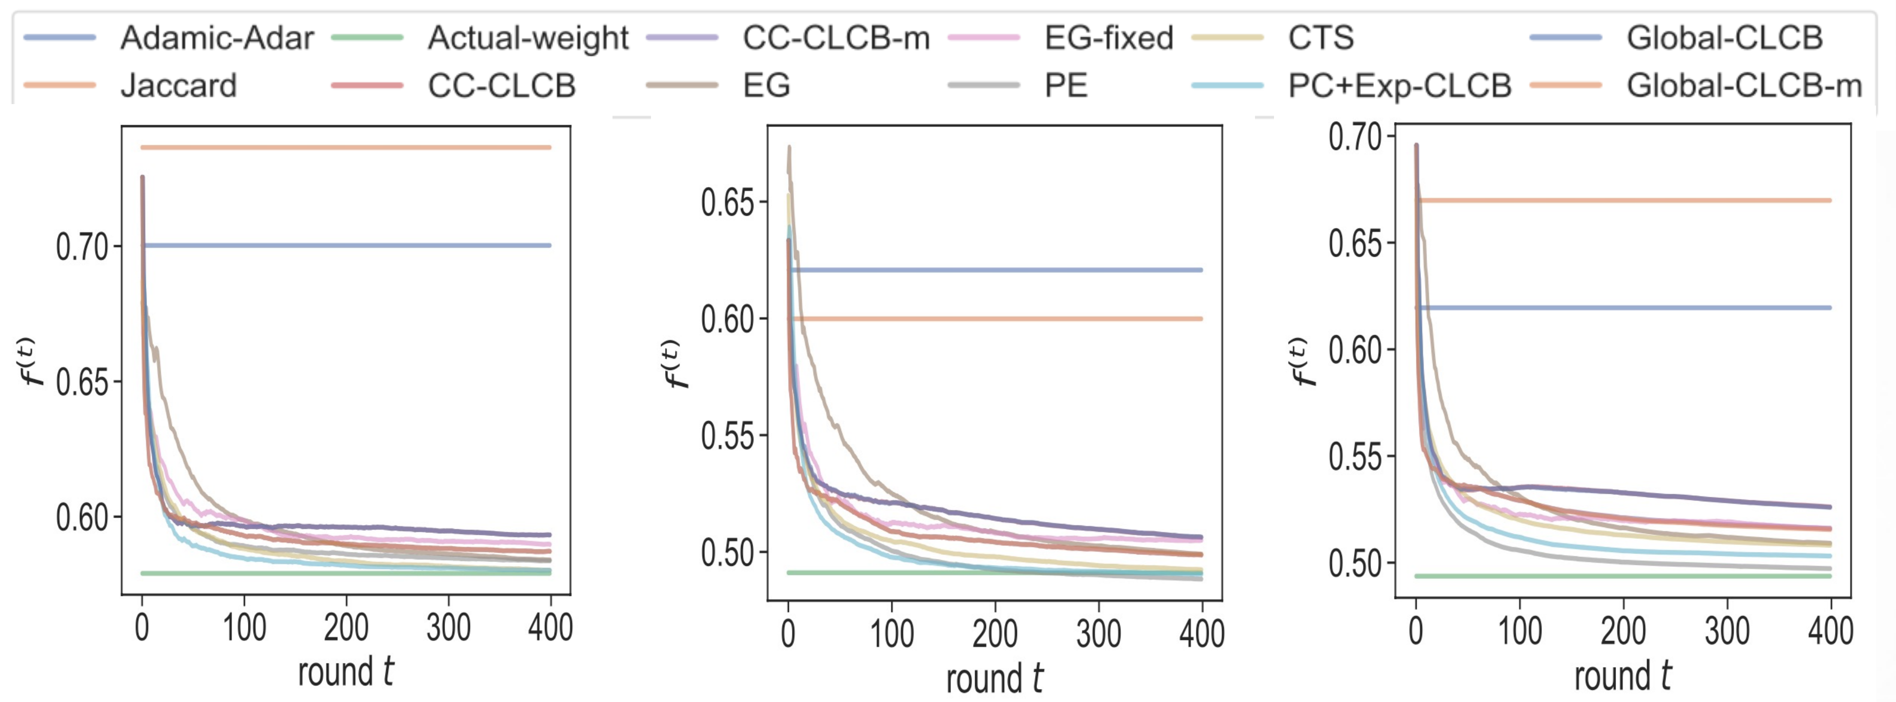

(a) Karate (b) Highland-Tribes (c) C-USA

- The loss values of all the CMAB methods follow a decreasing trend over the rounds since the CMAB algorithms learn how to cluster the vertices over time
- The non-CMAB baselines (Adamic-Adar, Jaccard) achieve the worst performance, Actual-weight is always the best method, the CMAB methods perform comparably or close to Actual-weight

Quality of the clusterings

Table. Performance in terms of $f(T)$ and growth rate (average amount of relative change between the initial and the final round over the span T, in percentage). All methods are equipped with Pivot as MIN-CC oracle.

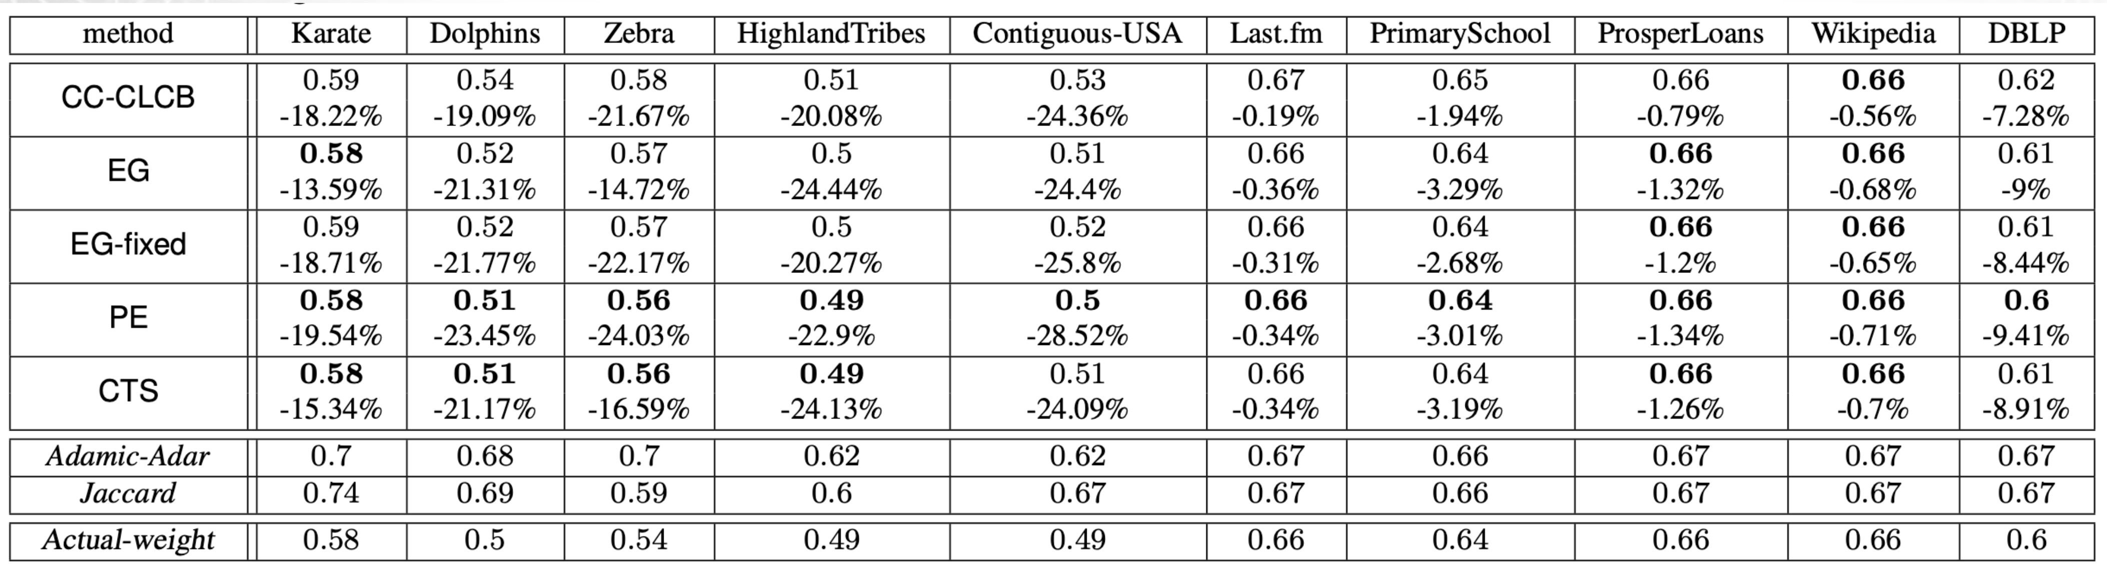

method	Karate	Dolphins	Zebra	HighlandTribes	Contiguous-USA	Last.fm	PrimarySchool	ProsperLoans	Wikipedia	DBLP
CC-CLCB	0.59 -18.22%	0.54 -19.09%	0.58 -21.67%	0.51 -20.08%	0.53 -24.36%	0.67 -0.19%	0.65 -1.94%	0.66 -0.79%	**0.66** -0.56%	0.62 -7.28%
EG	**0.58** -13.59%	0.52 -21.31%	0.57 -14.72%	0.5 -24.44%	0.51 -24.4%	0.66 -0.36%	0.64 -3.29%	**0.66** -1.32%	**0.66** -0.68%	0.61 -9%
EG-fixed	0.59 -18.71%	0.52 -21.77%	0.57 -22.17%	0.5 -20.27%	0.52 -25.8%	0.66 -0.31%	0.64 -2.68%	**0.66** -1.2%	**0.66** -0.65%	0.61 -8.44%
PE	**0.58** -19.54%	**0.51** -23.45%	**0.56** -24.03%	**0.49** -22.9%	**0.5** -28.52%	**0.66** -0.34%	**0.64** -3.01%	**0.66** -1.34%	**0.66** -0.71%	**0.6** -9.41%
CTS	**0.58** -15.34%	**0.51** -21.17%	**0.56** -16.59%	**0.49** -24.13%	0.51 -24.09%	0.66 -0.34%	0.64 -3.19%	**0.66** -1.26%	**0.66** -0.7%	0.61 -8.91%
Adamic-Adar	0.7	0.68	0.7	0.62	0.62	0.67	0.66	0.67	0.67	0.67
Jaccard	0.74	0.69	0.59	0.6	0.67	0.67	0.66	0.67	0.67	0.67
Actual-weight	0.58	0.5	0.54	0.49	0.49	0.66	0.64	0.66	0.66	0.6

- PE is the best-performing method since the adopted MIN-CC oracle (i.e., Pivot) is a randomized algorithm, thus, although with a pure-exploitation bandit strategy, it results in some implicit exploration
- CC-CLCB is comparable or close to the best methods in most cases

Quality of the learned edge weights

Table. Performance in terms of $ren^{(T)}$ and growth rate (average amount of relative change between the initial and the final round over the span T, in percentage). All methods are equipped with Pivot as MIN-CC oracle.

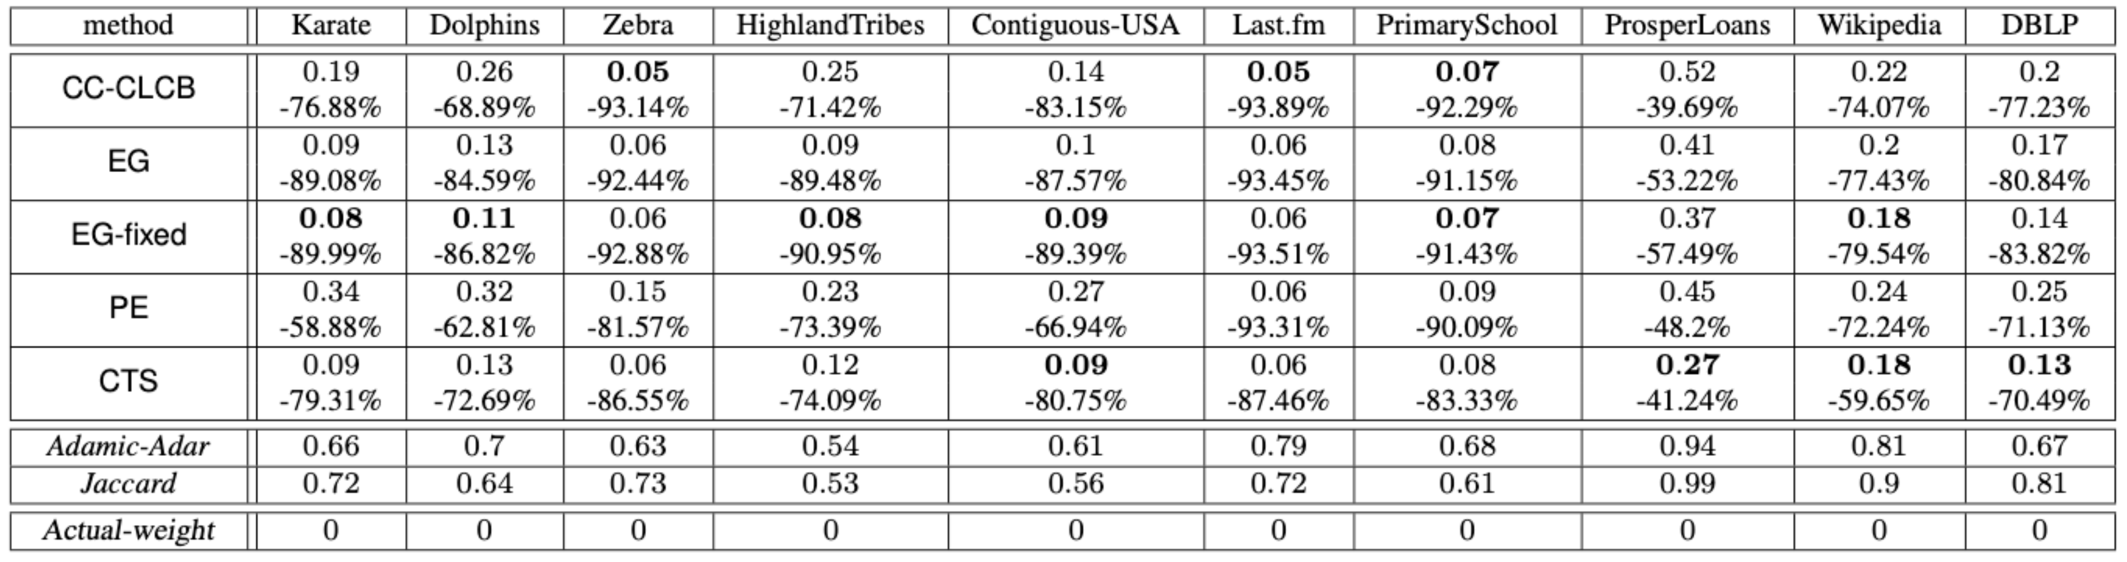

method	Karate	Dolphins	Zebra	HighlandTribes	Contiguous-USA	Last.fm	PrimarySchool	ProsperLoans	Wikipedia	DBLP
CC-CLCB	0.19 -76.88%	0.26 -68.89%	**0.05** -93.14%	0.25 -71.42%	0.14 -83.15%	**0.05** -93.89%	**0.07** -92.29%	0.52 -39.69%	0.22 -74.07%	0.2 -77.23%
EG	0.09 -89.08%	0.13 -84.59%	0.06 -92.44%	0.09 -89.48%	0.1 -87.57%	0.06 -93.45%	0.08 -91.15%	0.41 -53.22%	0.2 -77.43%	0.17 -80.84%
EG-fixed	**0.08** -89.99%	**0.11** -86.82%	0.06 -92.88%	**0.08** -90.95%	**0.09** -89.39%	0.06 -93.51%	**0.07** -91.43%	0.37 -57.49%	**0.18** -79.54%	0.14 -83.82%
PE	0.34 -58.88%	0.32 -62.81%	0.15 -81.57%	0.23 -73.39%	0.27 -66.94%	0.06 -93.31%	0.09 -90.09%	0.45 -48.2%	0.24 -72.24%	0.25 -71.13%
CTS	0.09 -79.31%	0.13 -72.69%	0.06 -86.55%	0.12 -74.09%	**0.09** -80.75%	0.06 -87.46%	0.08 -83.33%	**0.27** -41.24%	**0.18** -59.65%	**0.13** -70.49%
Adamic-Adar	0.66	0.7	0.63	0.54	0.61	0.79	0.68	0.94	0.81	0.67
Jaccard	0.72	0.64	0.73	0.53	0.56	0.72	0.61	0.99	0.9	0.81
Actual-weight	0	0	0	0	0	0	0	0	0	0

- the non-CMAB baselines yield the highest error values, while Actual-weight clearly achieves zero error everywhere
- EG-fixed is (comparable to) the best performer on the smaller datasets (Karate, Dolphins, Zebra, HighlandTribes, Contiguous-USA), while on the bigger datasets, CTS is (comparable to) the best method

Varying the MIN-CC oracle

- The general trend in terms of clustering quality is that LP+R leads to an increase in performance at the cost of higher running times
- The best performing method in terms of clustering quality is:
  - PE when equipped with Pivot
  - CTS when equipped with LP+R
- In terms of learned edge weights the advantage of using LP+R is less evident due to the high randomization of Pivot
- The best performing method in terms of edge weights estimation is:
  - EG-fixed when equipped with Pivot
  - CLCB when equipped with LP+R

Conclusion & Future Work

Summary:

- we have focused on the novel setting of correlation clustering where edge weights are unknown, and they need be discovered while performing multiple rounds of clustering.
- we have provided a Combinatorial Multi-Armed Bandit (CMAB) framework for correlation clustering, algorithms for it, analyses of the theoretical guarantees of these algorithms, more practical heuristics, and extensive experiments.

Future Work:

- we plan to investigate the theoretical properties of our heuristics, advanced CMAB settings, and clustering problems other than correlation clustering

Thank you!
Questions?

Related Work

Query-efficient correlation clustering[1,2]

- edge weights are discovered by querying an oracle
- the goal is to cluster the input graph by using a limited budget of Q queries ($Q \ll O(|V|^2)$)
- the oracle provides *true edge weights* for any query, at any time. Instead, in our setting, the feedback consists in a sample of the weight distributions
- existing approaches handle binary weights only (i.e., $w^+_{uv}, w^-_{uv} \in \{0, 1\}$)

1. Bressan M, Cesa-Bianchi N, Paudice A, Vitale F (2019) Correlation clustering with adaptive similarity queries. In: Proceedings of the NIPS conference, pp. 12531–12540.
2. García-Soriano D, Kutzkov K, Bonchi F, Tsourakakis C (2020) Query-efficient correlation clustering. In Proceedings of the WWW conference, pp 1468–1478.