Accelerated Evolving Set Processes for Local PageRank Computation

Binbin Huang [1] Luo Luo [1,2] Yanghua Xiao [3] Deqing Yang [1,3] Baojian Zhou [1,3]

[1] School of Data Science, Fudan University,
[2] Shanghai Key Laboratory for Contemporary Applied Mathematics,
[3] Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University

November 5, 2025

Background: PPR vector computation

Definition Given a source node vector e_s, damping factor $\alpha \in (0, 1)$ and an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with adjacency matricx $\boldsymbol{A}$ and degree matrix $\boldsymbol{D}$, the linear equation for computing the PPR vector π is defined as

$$\left(\boldsymbol{I} - (1 - \alpha)(\boldsymbol{I} + \boldsymbol{A}\boldsymbol{D}^{-1})/2\right) \boldsymbol{\pi} = \alpha \boldsymbol{e_s}. \quad (1)$$

Background: PPR vector computation

Definition Given a source node vector e_s, damping factor $\alpha \in (0,1)$ and an undirected graph $\mathcal{G}(\mathcal{V},\mathcal{E})$ with adjacency matricx $\boldsymbol{A}$ and degree matrix $\boldsymbol{D}$, the linear equation for computing the PPR vector π is defined as

$$\left(\boldsymbol{I} - (1-\alpha)(\boldsymbol{I} + \boldsymbol{A}\boldsymbol{D}^{-1})/2\right)\boldsymbol{\pi} = \alpha \boldsymbol{e}_s. \tag{1}$$

We solve the linear system by reformulating as:

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \left\{ f(\boldsymbol{x}) \triangleq \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} - \alpha \boldsymbol{x}^\top \boldsymbol{D}^{-1/2}\boldsymbol{b} \right\}, \tag{2}$$

where $\boldsymbol{Q} \triangleq \frac{1+\alpha}{2}\boldsymbol{I} - \frac{1-\alpha}{2}\boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2}$ and the condition number of f is $1/\alpha$.

Background: PPR vector computation

Definition Given a source node vector e_s, damping factor $\alpha \in (0,1)$ and an undirected graph $\mathcal{G}(\mathcal{V},\mathcal{E})$ with adjacency matricx $\boldsymbol{A}$ and degree matrix $\boldsymbol{D}$, the linear equation for computing the PPR vector π is defined as

$$\left(\boldsymbol{I} - (1-\alpha)(\boldsymbol{I} + \boldsymbol{A}\boldsymbol{D}^{-1})/2\right)\boldsymbol{\pi} = \alpha \boldsymbol{e}_s. \quad (1)$$

We solve the linear system by reformulating as:

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \left\{ f(\boldsymbol{x}) \triangleq \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} - \alpha \boldsymbol{x}^\top \boldsymbol{D}^{-1/2}\boldsymbol{b} \right\}, \quad (2)$$

where $\boldsymbol{Q} \triangleq \frac{1+\alpha}{2}\boldsymbol{I} - \frac{1-\alpha}{2}\boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2}$ and the condition number of f is $1/\alpha$.
The optimal solution of (2) is denoted by $\boldsymbol{x}_f^* := \alpha \boldsymbol{Q}^{-1}\boldsymbol{D}^{-1/2}\boldsymbol{b}$. When $\boldsymbol{b} = \boldsymbol{e}_s$, it implies $\boldsymbol{\pi} := \boldsymbol{D}^{1/2}\boldsymbol{x}_f^*$.

ϵ-approximation to (2) : $\quad \mathcal{P}(\epsilon,\alpha,\boldsymbol{b},\mathcal{G}) \triangleq \left\{ \boldsymbol{x} : \|\boldsymbol{D}^{-1/2}(\boldsymbol{x} - \boldsymbol{x}_f^*)\|_\infty \leq \epsilon \right\}. \quad (3)$

Key ingredients: Accelerated Evolving Set Processes

Accelerated Evolving Set Process (AESP) framework

- computes an ϵ-approximation for PPR using $\tilde{\mathcal{O}}(1/\sqrt{\alpha})$ short evolving set process;
- is built upon the inexact proximal point algorithm.

Nested evolving set process

Given the configuration $\theta \triangleq (\alpha, \boldsymbol{b}, \mathcal{G})$, and a local method $\mathcal{M}$, the nested evolving set process at outer-loop iteration t generates a sequence of $\{\mathcal{S}_t^{(k+1)}, \boldsymbol{z}_t^{(k+1)}\}_{k\geq 0}$ according to $(\mathcal{S}_t^{(k+1)}, \boldsymbol{z}_t^{(k+1)}) = \boldsymbol{\Phi}_{\theta,\mathcal{M}}(\mathcal{S}_t^{(k)}, \boldsymbol{z}_t^{(k)})$, where $\mathcal{S}_t^{(k)} \subseteq \mathcal{V}$ is efficiently maintained using a queue data structure, avoiding accessing the entire graph. We say the process *converges* when $\mathcal{S}_t^{(K_t)} = \emptyset$ for some K_t. After T outer-loop iterations, the generated sequences of active sets and estimation pairs are

$$\begin{array}{ccccc} (\mathcal{S}_1^{(0)}, \boldsymbol{z}_1^{(0)}) & \rightarrow & \cdots & \rightarrow & (\mathcal{S}_1^{(K_1)} = \emptyset, \boldsymbol{z}_1^{(K_1)} = \boldsymbol{x}^{(1)}),\ t = 1; \\ \vdots & & & & \vdots \\ (\mathcal{S}_T^{(0)}, \boldsymbol{z}_T^{(0)}) & \rightarrow & \cdots & \rightarrow & (\mathcal{S}_T^{(K_T)} = \emptyset, \boldsymbol{z}_T^{(K_T)} = \boldsymbol{x}^{(T)}),\ t = T. \end{array}$$

Accelerated Evolving Set Processes

We propose an **Accelerated Evolving Set Process (AESP)** framework, which computes an ϵ-approximation for PPR using $\tilde{\mathcal{O}}(1/\sqrt{\alpha})$ short evolving set process and is built upon the inexact proximal point algorithm.

Accelerated Evolving Set Processes

We propose an **Accelerated Evolving Set Process (AESP)** framework, which computes an ϵ-approximation for PPR using $\tilde{\mathcal{O}}(1/\sqrt{\alpha})$ short evolving set process and is built upon the inexact proximal point algorithm.
localized Catalyst-style updates

$$\begin{aligned} \text{AESP} \quad \boldsymbol{x}^{(t)} &= \mathcal{M}(\varphi_t, \boldsymbol{y}^{(t-1)}, \eta, \alpha, \boldsymbol{b}, \mathcal{G}), \\ \boldsymbol{y}^{(t)} &= \boldsymbol{x}^{(t)} + \beta_t(\boldsymbol{x}^{(t)} - \boldsymbol{x}^{(t-1)}). \end{aligned}$$

At t-th iteration, proximal operator objective is

$$h_t(\boldsymbol{z}) \triangleq f(\boldsymbol{z}) + \frac{\eta}{2}\|\boldsymbol{z} - \boldsymbol{y}^{(t-1)}\|_2^2,$$

$$\boldsymbol{x}^{(t)} \in \mathcal{H}_t(\varphi_t) \triangleq \{\boldsymbol{z} \in \mathbb{R}^n : h_t(\boldsymbol{z}) - h_t^* \leq \varphi_t\},$$

where h_t^* is the minimal value of h_t.

Accelerated Evolving Set Processes

We propose an **Accelerated Evolving Set Process (AESP)** framework, which computes an ϵ-approximation for PPR using $\tilde{\mathcal{O}}(1/\sqrt{\alpha})$ short evolving set process and is built upon the inexact proximal point algorithm.

localized Catalyst-style updates

$$\text{AESP} \quad \boldsymbol{x}^{(t)} = \mathcal{M}(\varphi_t, \boldsymbol{y}^{(t-1)}, \eta, \alpha, \boldsymbol{b}, \mathcal{G}),$$
$$\boldsymbol{y}^{(t)} = \boldsymbol{x}^{(t)} + \beta_t(\boldsymbol{x}^{(t)} - \boldsymbol{x}^{(t-1)}).$$

At t-th iteration, proximal operator objective is

$$h_t(\boldsymbol{z}) \triangleq f(\boldsymbol{z}) + \frac{\eta}{2}\|\boldsymbol{z} - \boldsymbol{y}^{(t-1)}\|_2^2,$$
$$\boldsymbol{x}^{(t)} \in \mathcal{H}_t(\varphi_t) \triangleq \{\boldsymbol{z} \in \mathbb{R}^n : h_t(\boldsymbol{z}) - h_t^* \leq \varphi_t\},$$

where h_t^* is the minimal value of h_t.

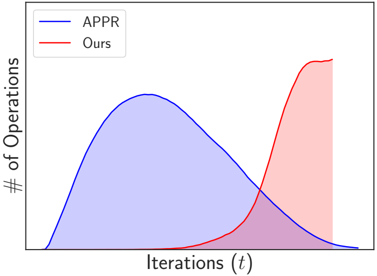

Figure: The comparison of volumes of ESP for APPR and Ours.

Accelerated Evolving Set Processes

We propose an **Accelerated Evolving Set Process (AESP)** framework, which computes an ϵ-approximation for PPR using $\tilde{\mathcal{O}}(1/\sqrt{\alpha})$ short evolving set process and is built upon the inexact proximal point algorithm.

localized Catalyst-style updates

$$\text{AESP} \quad \boldsymbol{x}^{(t)} = \mathcal{M}(\varphi_t, \boldsymbol{y}^{(t-1)}, \eta, \alpha, \boldsymbol{b}, \mathcal{G}),$$
$$\boldsymbol{y}^{(t)} = \boldsymbol{x}^{(t)} + \beta_t(\boldsymbol{x}^{(t)} - \boldsymbol{x}^{(t-1)}).$$

At t-th iteration, proximal operator objective is

$$h_t(\boldsymbol{z}) \triangleq f(\boldsymbol{z}) + \frac{\eta}{2}\|\boldsymbol{z} - \boldsymbol{y}^{(t-1)}\|_2^2,$$
$$\boldsymbol{x}^{(t)} \in \mathcal{H}_t(\varphi_t) \triangleq \{\boldsymbol{z} \in \mathbb{R}^n : h_t(\boldsymbol{z}) - h_t^* \leq \varphi_t\},$$

where h_t^* is the minimal value of h_t.

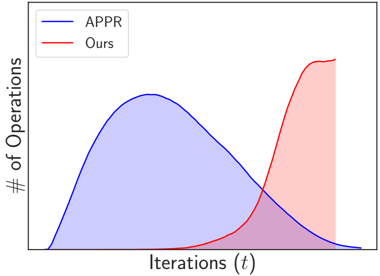

Figure: The comparison of volumes of ESP for APPR and Ours.

Condition number: $f(\boldsymbol{x}) : \frac{1}{\alpha} \longrightarrow h_t(\boldsymbol{z}) : \frac{\eta+1}{\eta+\alpha} \overset{\eta=1-2\alpha}{=} 2$. **A constant!**

Localized inexact proximal operators

- LocGD $\quad \boldsymbol{z}_t^{(k+1)} = \boldsymbol{z}_t^{(k)} - \frac{2\nabla h_t(\boldsymbol{z}_t^{(k)}) \circ \mathbf{1}_{\mathcal{S}_t^k}}{1+\alpha+2\eta}$, for $k \geq 0$,
- LocAPPR $\quad \boldsymbol{z}_t^{(k_{i+1})} = \boldsymbol{z}_t^{(k_i)} - \frac{2\nabla h_t(\boldsymbol{z}_t^{(k_i)}) \circ \mathbf{1}_{\{u_i\}}}{1+\alpha+2\eta}$, for $u_i \in \mathcal{S}_t^k = \{u_1, u_2, \ldots, u_{|\mathcal{S}_t^k|}\}$.

where $\mathcal{S}_t^k = \{u : |\nabla_u h_t^{-1/2}(\boldsymbol{z}_t^{(k)})| \geq \epsilon_t\}$.
stopping criterion: $\mathcal{S}_t^{K_t} = \emptyset$, which is $\|\nabla h_t^{-1/2}(\boldsymbol{z})\|_\infty < \epsilon_t$.

Localized inexact proximal operators

- LocGD $\quad z_t^{(k+1)} = z_t^{(k)} - \frac{2\nabla h_t(z_t^{(k)}) \circ \mathbf{1}_{\mathcal{S}_t^k}}{1+\alpha+2\eta}$, for $k \geq 0$,
- LocAPPR $\quad z_t^{(k_{i+1})} = z_t^{(k_i)} - \frac{2\nabla h_t(z_t^{(k_i)}) \circ \mathbf{1}_{\{u_i\}}}{1+\alpha+2\eta}$, for $u_i \in \mathcal{S}_t^k = \{u_1, u_2, \ldots, u_{|\mathcal{S}_t^k|}\}$.

where $\mathcal{S}_t^k = \{u : |\nabla_u h_t^{-1/2}(z_t^{(k)})| \geq \epsilon_t\}$.

stopping criterion: $\mathcal{S}_t^{K_t} = \emptyset$, which is $\|\nabla h_t^{-1/2}(z)\|_\infty < \epsilon_t$.

Algorithm 1 AESP$(\epsilon, \alpha, \boldsymbol{b}, \eta, \mathcal{G}, \mathcal{M})$

$\boldsymbol{y}^{(0)} = \boldsymbol{x}^{(0)} = \mathbf{0}, c = 1 - 0.9\sqrt{\mu/(\mu+\eta)}$
$T = \lceil \frac{10}{9}\sqrt{\frac{1-\alpha}{\alpha}} \log \frac{400(1-\alpha^2)}{\alpha^2 \epsilon^2} \rceil$
for $t = 1, 2, \ldots, T$ **do**
  $\varphi_t = (L+\mu)\|\boldsymbol{b}\|_1^2 c^t / 18$
  $\boldsymbol{x}^{(t)} = \mathcal{M}(\varphi_t, \boldsymbol{y}^{(t-1)}, \eta, \alpha, \boldsymbol{b}, \mathcal{G})$
  // $\mathcal{M}$ in LocAPPR or LocGD
  if $\{v : \epsilon\alpha\sqrt{d_v} \leq |\nabla_v f(\boldsymbol{x}^{(t)})|\} = \emptyset$ **then**
    break
  $\boldsymbol{y}^{(t)} = \boldsymbol{x}^{(t)} + \frac{\sqrt{\mu+\eta}-\sqrt{\mu}}{\sqrt{\mu+\eta}+\sqrt{\mu}}(\boldsymbol{x}^{(t)} - \boldsymbol{x}^{(t-1)})$
Return $\hat{\boldsymbol{x}} = \boldsymbol{x}^{(t)}$

Algorithm 2 AESP-PPR$(\epsilon, \alpha, s, \mathcal{G}, \mathcal{M})$

$\boldsymbol{y}^{(0)} = \boldsymbol{x}^{(0)} = \mathbf{0}$
$T = \lceil \frac{10}{9}\sqrt{\frac{1-\alpha}{\alpha}} \log \frac{400(1-\alpha^2)}{\alpha^2 \epsilon^2} \rceil$
for $t = 1, 2, \ldots, T$ **do**
  $\varphi_t = \frac{1+\alpha}{18}(1 - \frac{9}{10}\sqrt{\frac{\alpha}{1-\alpha}})^t$
  // $\mathcal{M}$ is LocAPPR or LocGD
  $\boldsymbol{x}^{(t)} = \mathcal{M}(\varphi_t, \boldsymbol{y}^{(t-1)}, 1-2\alpha, \alpha, \boldsymbol{b}, \mathcal{G})$
  if $\{v : \epsilon\alpha\sqrt{d_v} \leq |\nabla_v f(\boldsymbol{x}^{(t)})|\} = \emptyset$ **then**
    break
  $\boldsymbol{y}^{(t)} = \boldsymbol{x}^{(t)} + \frac{\sqrt{1-\alpha}-\sqrt{\alpha}}{\sqrt{1-\alpha}+\sqrt{\alpha}}(\boldsymbol{x}^{(t)} - \boldsymbol{x}^{(t-1)})$
Return $\hat{\boldsymbol{\pi}} = \boldsymbol{D}^{1/2}\boldsymbol{x}^{(t)}$

Localized inexact proximal operators

$$\epsilon_t \triangleq \max\left\{\sqrt{\frac{(\mu+\eta)\varphi_t}{m}}, \frac{2(\eta+\alpha)\varphi_t}{\|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(0)})\|_1}\right\} \Rightarrow \boldsymbol{z}_t^{(K_t)} \in \mathcal{H}_t(\varphi_t)$$

$$\overline{\mathrm{vol}}(\mathcal{S}_t) \triangleq \frac{1}{K_t}\sum_{k=0}^{K_t-1}\mathrm{vol}(\mathcal{S}_t^{(k)}), \overline{\gamma}_t \triangleq \frac{1}{K_t}\sum_{k=0}^{K_t-1}\left\{\gamma_t^{(k)} \triangleq \frac{\|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(k)}) \circ \mathbf{1}_{\mathcal{S}_t^{(k)}}\|_1}{\|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(k)})\|_1}\right\}.$$

Theorem (Convergence of LocGD)

LocGD *returns* $\boldsymbol{z}_t^{(K_t)} = \text{LocGD}(\varphi_t, \boldsymbol{y}^{(t-1)}, \eta, \alpha, \boldsymbol{b}, \mathcal{G}) \in \mathcal{H}_t(\varphi_t)$. *For* $k \geq 0$, *the scaled gradient satisfies*

$$\left\|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(k+1)})\right\|_1 \leq \left(1-\tau\gamma_t^{(k)}\right)\left\|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(k)})\right\|_1,$$

where $\tau := \frac{2(\alpha+\eta)}{1+\alpha+2\eta}$ *and* $\gamma_t^{(k)}$ *is the ratio, then the run time* $\mathcal{T}_t^{\text{LocGD}}$ *is bounded by*

$$\mathcal{T}_t^{\text{LocGD}} \leq \min\left\{\frac{\overline{\mathrm{vol}}(\mathcal{S}_t)}{\tau\overline{\gamma}_t}\log\frac{C_{h_t}^0}{C_{h_t}^{K_t}}, \frac{C_{h_t}^0 - C_{h_t}^{K_t}}{\tau\epsilon_t}\right\},$$

where $C_{h_t}^i = \|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(i)})\|_1$ *denote constants. Furthermore,* $\overline{\mathrm{vol}}(\mathcal{S}_t)/\overline{\gamma}_t \leq \min\left\{C_{h_t}^0/\epsilon_t, 2m\right\}$.

Time complexity analysis

Lemma (Outer-loop iteration complexity of AESP)

If each iteration of AESP, presented in Algorithm 1, finds $\boldsymbol{x}^{(t)} := \boldsymbol{z}_t^{(K_t)}$ *using* $\mathcal{M}$, *satisfying* $h_t(\boldsymbol{z}_t^{(K_t)}) - h_t^* \leq \varphi_t := (L+\mu)\|\boldsymbol{b}\|_1^2(1-\rho)^t/18$, *then the total number of iterations* T *required to ensure* $\hat{\boldsymbol{x}} = \mathrm{AESP}(\epsilon, \alpha, \boldsymbol{b}, \eta, \mathcal{G}, \mathcal{M}) \in \mathcal{P}(\epsilon, \alpha, \boldsymbol{b}, \mathcal{G})$ *as defined in Eq.* (3), *for solving* (2), *satisfies the bound*

$$T \leq \frac{1}{\rho} \log\left(\frac{4(L+\mu)\|\boldsymbol{b}\|_1^2}{\mu\epsilon^2(\sqrt{q}-\rho)^2}\right), \text{ where } \rho = 0.9\sqrt{q} \text{ and } q = \frac{\mu}{\mu+\eta}. \tag{4}$$

Time complexity analysis

Lemma (Outer-loop iteration complexity of AESP)

If each iteration of AESP, presented in Algorithm 1, finds $\boldsymbol{x}^{(t)} := \boldsymbol{z}_t^{(K_t)}$ *using* $\mathcal{M}$*, satisfying* $h_t(\boldsymbol{z}_t^{(K_t)}) - h_t^* \leq \varphi_t := (L+\mu)\|\boldsymbol{b}\|_1^2(1-\rho)^t/18$*, then the total number of iterations* T *required to ensure* $\hat{\boldsymbol{x}} = \mathrm{AESP}(\epsilon, \alpha, \boldsymbol{b}, \eta, \mathcal{G}, \mathcal{M}) \in \mathcal{P}(\epsilon, \alpha, \boldsymbol{b}, \mathcal{G})$ *as defined in Eq.* (3)*, for solving* (2)*, satisfies the bound*

$$T \leq \frac{1}{\rho} \log\left(\frac{4(L+\mu)\|\boldsymbol{b}\|_1^2}{\mu\epsilon^2(\sqrt{q}-\rho)^2}\right), \text{ where } \rho = 0.9\sqrt{q} \text{ and } q = \frac{\mu}{\mu+\eta}. \tag{4}$$

Theorem (Time complexity of AESP)

Assume damping factor $\alpha < 1/2$*. Applying* $\hat{\boldsymbol{x}} = \mathrm{AESP}(\epsilon, \alpha, \boldsymbol{b}, \eta, \mathcal{G}, \mathcal{M})$ *with* $\eta = L - 2\mu$ *and* $\mathcal{M}$ *be either* LocGD *or* LocAPPR*, then AESP presented in Algorithm 1, finds a solution* $\hat{\boldsymbol{x}}$ *such that* $\|\boldsymbol{D}^{-1/2}(\hat{\boldsymbol{x}} - \boldsymbol{x}_f^*)\|_\infty \leq \epsilon$ *with the dominated time complexity* $\mathcal{T}$ *bounded by*

$$\mathcal{T} \leq \sum_{t=1}^{T} \min\left\{\frac{\overline{\mathrm{vol}}(\mathcal{S}_t)}{\tau\overline{\gamma}_t} \log\frac{C_{h_t}^0}{C_{h_t}^{K_t}}, \frac{C_{h_t}^0 - C_{h_t}^{K_t}}{\tau\epsilon_t}\right\}, \text{ with } \frac{\overline{\mathrm{vol}}(\mathcal{S}_t)}{\overline{\gamma}_t} \leq \min\left\{\frac{C_{h_t}^0}{\epsilon_t}, 2m\right\},$$

where τ*,* ϵ_t*,* $C_{h_t}^0$ *and* $C_{h_t}^{K_t}$ *are defined in Theorem 1. Furthermore,* $q = \mu/(L-\mu)$ *and the number of outer iterations satisfies* $T \leq \frac{10}{9\sqrt{q}} \log\left(\frac{400(L+\mu)\|\boldsymbol{b}\|_1^2}{\mu\epsilon^2 q}\right) = \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{\alpha}}\right)$.

Time complexity analysis

$$\mathcal{T} = \tilde{\mathcal{O}}\left(\frac{\overline{\mathrm{vol}}(\mathcal{S}_t)}{\sqrt{\alpha}\overline{\gamma}_t}\right) = \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{\alpha}\epsilon_T}\right) = \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{\alpha}\epsilon^2}\right).$$

Theorem (Time complexity of AESP-PPR)

The PPR vector of $s \in \mathcal{V}$ is defined in Eq. (1), and the precision $\epsilon \in (0, 1/d_s)$. Suppose $\hat{\boldsymbol{\pi}} = \mathrm{AESP\text{-}PPR}(\epsilon, \alpha, s, \mathcal{G}, \mathcal{M})$ be returned by Algorithm 1. When $\mathcal{M}$ is either LocGD or $\mathrm{LocAPPR}$, then $\hat{\boldsymbol{\pi}}$ satisfies $\|\boldsymbol{D}^{-1}(\hat{\boldsymbol{\pi}} - \boldsymbol{\pi})\|_\infty \leq \epsilon$ and AESP-PPR has a dominated time complexity bounded by

$$\mathcal{T} \leq \min\left\{\tilde{\mathcal{O}}\left(\frac{\overline{\mathrm{vol}}(\mathcal{S}_{T_{\max}})}{\sqrt{\alpha}\overline{\gamma}_{T_{\max}}}\right), \tilde{\mathcal{O}}\left(\frac{\max_t C^0_{h_t}}{\sqrt{\alpha}\epsilon_T}\right)\right\} = \min\left\{\tilde{\mathcal{O}}\left(\frac{m}{\sqrt{\alpha}}\right), \tilde{\mathcal{O}}\left(\frac{R^2/\epsilon^2}{\sqrt{\alpha}}\right)\right\}, \tag{5}$$

where $T_{\max} := \arg\max_{t\in[T]} \overline{\mathrm{vol}}(\mathcal{S}_t)/\overline{\gamma}_t$ *and* $R := \max\left\{\|\nabla h_t^{1/2}(\boldsymbol{z}_t^{(0)})\|_1 / \|\nabla h_1^{1/2}(\boldsymbol{z}_1^{(0)})\|_1 : \forall t \in [T]\right\}$.

Experiments

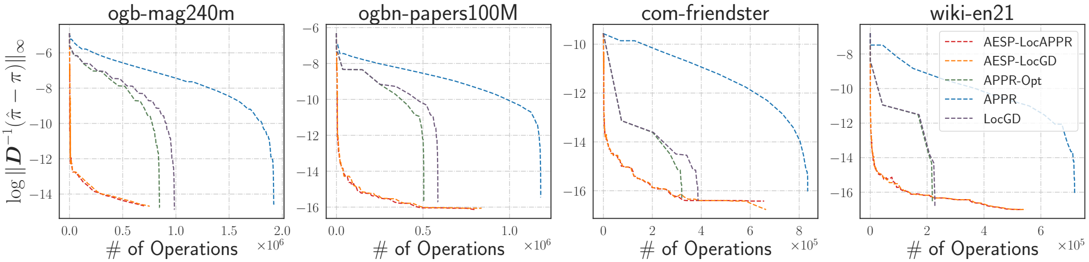

Figure: Performance of estimation error reduction, $\log \|\boldsymbol{D}^{-1}(\hat{\boldsymbol{\pi}} - \boldsymbol{\pi})\|_\infty$, as a function of operations $\mathcal{T}$, on the graph *ogb-mag240m*, *ogbn-papers100M*, *com-friendster* and *wiki-en21* with $\alpha = 0.01$ and $\epsilon = 10^{-6}$ where the graph can scale up to $n = 244M$ and $m = 1.728B$.

Experiment

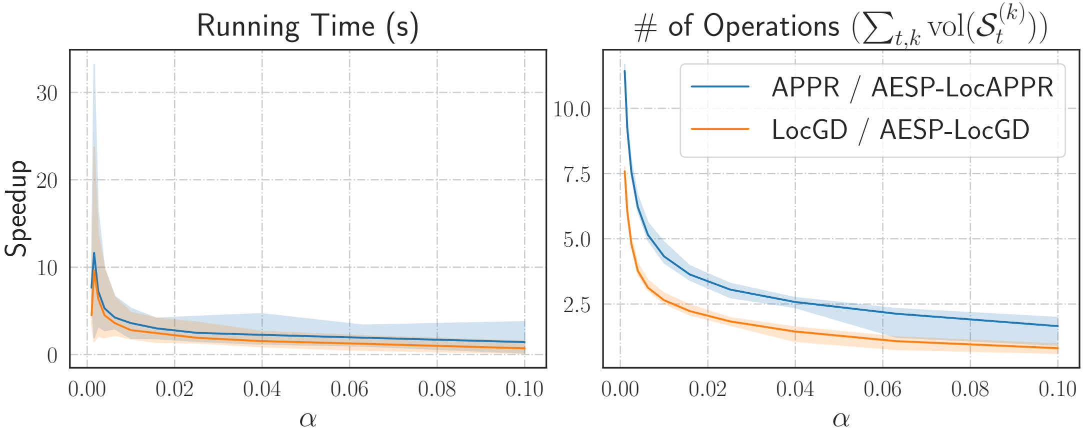

Figure: Speedup of AESP-based methods over standard local solvers (LocAPPR, LocGD) as a function of α, on the *com-dblp* graph with $\epsilon = 0.1/n$ and $\alpha \in (10^{-3}, 10^{-1})$.

Thanks!

- **Our code is publicly available at:**
  `https://github.com/Rick7117/aesp-local-pagerank`
- **If you have any questions, contact us:**
  Binbin Huang: bbhuang24@m.fudan.edu.cn
  Baojian Zhou: bjzhou@fudan.edu.cn