# A Random Walk Around The Block 

Johan Ugander<br>Stanford University

Joint work with:
Isabel Kloumann (Facebook) \& Jon Kleinberg (Cornell)

Google Mountain View August 17, 2016

## Seed set expansion

- Given a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, goal is to accurately identify a target set $\mathbf{T} \subset \mathbf{V}$ from a smaller seed set $\mathbf{S} \subset \mathbf{T}$.



## Seed set expansion

- Given a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, goal is to accurately identify a target set $\mathbf{T} \subset \mathbf{V}$ from a smaller seed set $\mathbf{S} \subset \mathbf{T}$.



## Seed set expansion

- Given a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, goal is to accurately identify a target set $\mathbf{T} \subset \mathbf{V}$ from a smaller seed set $\mathbf{S} \subset \mathbf{T}$.



## Seed set expansion

- Given a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, goal is to accurately identify a target set $\mathbf{T} \subset \mathbf{V}$ from a smaller seed set $\mathbf{S} \subset \mathbf{T}$.



## Seed set expansion

- Given a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, goal is to accurately identify a target set $\mathbf{T} \subset \mathbf{V}$ from a smaller seed set $\mathbf{S} \subset \mathbf{T}$.
- Applications:
- Broadly: ranking on graphs, recommendation systems
- Spam filtering (Wu \& Chellapilla '07)
- Community detection (Weber et al. '13)
- Missing data inference (Mislove et al. '14)
- Common methods:
- Semi-supervised learning (Zhu et al. '03)
- Diffusion-based classification (Jeh \& Widom '03, Kloster \& Gleich '14)
- Outwardness, modularity and more (Bagrow '08, Kloumann \& Kleinberg '14)



## Seed set expansion

- Given a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, goal is to accurately identify a target set $\mathbf{T} \subset \mathbf{V}$ from a smaller seed set $\mathbf{S} \subset \mathbf{T}$.
- Applications:
- Broadly: ranking on graphs, recommendation systems
- Spam filtering (Wu \& Chellapilla '07)
- Community detection (Weber et al. '13)
- Missing data inference (Mislove et al. '14)
- Common methods:
- Semi-supervised learning (Zhu et al. '03)
- Diffusion-based classification (Jeh \& Widom '03, Kloster \& Gleich '14)
- Outwardness, modularity and more (Bagrow '08, Kloumann \& Kleinberg '14)



## Recall curves for seed set expansion



- Recall curve: true positive rate, as a function of the number of items returned based on small uniformly random seed set.
- Kloumann \& Kleinberg '14 tested many different methods on data, broadly found Personalized PageRank to be best.


## Recall curves for seed set expansion



- Recall curve: true positive rate, as a function of the number of items returned based on small uniformly random seed set.
- Kloumann \& Kleinberg '14 tested many different methods on data, broadly found Personalized PageRank to be best.
- Truncated PPR (first K steps) comparable to PPR from $\mathrm{K}=4$.
- Heat Kernel later found comparable to PPR.


## Diffusion-based node classification

- Classification based on random walk landing probabilities
- $r_{k}^{v}$, probability that a random walk starting in $\mathbf{S}$ is at $\mathbf{v}$ after $\mathbf{k}$ steps.
- $\left(r_{1}^{v}, r_{2}^{v}, \ldots, r_{K}^{v}\right)$, truncated vector of landing probabilities.
- Personalized PageRank and Heat Kernel ranking:

$$
\operatorname{PPR}(v) \propto \sum_{k=1}^{\infty}\left(\alpha^{k}\right) r_{k}^{v} \quad \operatorname{HK}(v) \propto \sum_{k=1}^{\infty}\left(\frac{t^{k}}{k!}\right) r_{k}^{v}
$$

- General diffusion score function:

$$
\operatorname{score}(v)=\sum_{k=1}^{\infty} w_{k} r_{k}^{v}
$$



## Diffusion-based node classification

- Personalized PageRank and Heat Kernel
= two parametric families of linear weights

$$
\operatorname{score}(v)=\sum_{k=1}^{K} w_{k} r_{k}^{v}
$$



PPR $w_{k}=\alpha^{k}$
$\mathrm{HK} \quad w_{k}=t^{k} / k$ !

- Question in this work:

What weights are "optimal" for diffusion-based classification?

## The stochastic block model

- C blocks
- Focus on C=2 blocks: 1="Target", 2="Other"
- $\mathbf{n}_{1}, \mathbf{n}_{2}$ nodes in blocks
- Independent edge probabilities:
- Edge probability within a block = $\mathbf{p i n}_{\text {in }}$
- Edge probability across blocks = pout
- (Results for C>2 as well, see paper)
- Model with many names:
- Stochastic Block Model (Holland et al. '83)
- Affiliation Model (Frank-Harary '82)
- Planted Partition Model (Dyer-Frieze '89)


## The SBM resolution limit

- Find true partition in poly(n) time w.h.p. as $\mathbf{n} \rightarrow \infty$ :
- Dyer-Frieze '89: If $p_{\text {in }}-p_{\text {out }}=O(1)$
- Condon-Karp '01: If $p_{\text {in }}-p_{\text {out }} \geq \Omega\left(n^{-1 / 2}\right)$
- McSherry '01: If $\mathrm{p}_{\text {in }}-\mathrm{p}_{\text {out }} \geq \Omega\left(\left(\mathrm{p}_{\text {out }}(\log \mathrm{n}) / \mathrm{n}\right)^{-1 / 2}\right)$



## The SBM resolution limit

- Find true partition in poly(n) time w.h.p. as $\mathbf{n} \rightarrow \infty$ :
- Dyer-Frieze '89: If $p_{\text {in }}-p_{\text {out }}=O(1)$
- Condon-Karp '01: If pin $^{-}$pout $^{2} \Omega\left(n^{-1 / 2}\right)$
- McSherry '01: If $p_{\text {in }}-p_{\text {out }} \geq \Omega\left(\left(p_{\text {out }}(\log n) / n\right)^{-1 / 2}\right)$

- Find partition positively correlated with true partition:
- Coja-Oghlan '06: If $\mathrm{p}_{\text {in }}-\mathrm{p}_{\text {out }} \geq \Omega\left(\left(\mathrm{p}_{\mathrm{out}} / \mathrm{n}\right)^{-1 / 2}\right)$,


## The SBM resolution limit

- Find true partition in poly(n) time w.h.p. as $\mathbf{n} \rightarrow \infty$ :
- Dyer-Frieze '89: If $\mathrm{p}_{\text {in }}-\mathrm{p}_{\text {out }}=\mathrm{O}(1)$
- Condon-Karp '01: If $p_{\text {in }}-p_{\text {out }} \geq \Omega\left(n^{-1 / 2}\right)$
- McSherry '01: If $\mathrm{p}_{\text {in }}-\mathrm{p}_{\text {out }} \geq \Omega\left(\left(\mathrm{p}_{\text {out }}(\log n) / n\right)^{-1 / 2}\right)$
- Find partition positively correlated with true partition:
- Coja-Oghlan '06: If $p_{\text {in }}-p_{\text {out }} \geq \Omega\left(\left(p_{\text {out }} / n\right)^{-1 / 2}\right)$,
- If and only if $(a-b)^{2}>2(a+b)\left(p_{\text {in }}=a / n, p_{\text {out }}=b / n\right)$ :
- Decelle et al '11: Conjecture and belief propagation numerics
- Mossel et al '12,'13, Massoulié '13, Abbe et al. '14: Proven
- Recent extensions:
- More than two blocks (e.g. Neeman-Netrapalli '14)
- Unequal block sizes (e.g. Zhang et al. '16)


## The SBM resolution limit

- Is block recovery/classification over? No!
- Unsupervised vs. semi-supervised
- Empirical graphs != SBMs
- Optimal algorithms not practical

- Beyond asymptotic limits, what are decay rates?
- Rather than being "problem down" (SBM classification), this talk will be "method up": how to tune diffusion weights to find seed sets?

$$
\operatorname{score}(v)=\sum_{k=1}^{K} w_{k} r_{k}^{v}
$$

- Possible variations: Diffusion weights for seed set expansion in core-periphery models? Latent space models (Hoff et al. 2002)? Etc.


## Diffusion-based classification in SBMs

- SBMs present a natural binary classification problem.
- Recall notation:
- $r_{k}^{v}$, probability that a random walk starting in $\mathbf{S}$ is at $\mathbf{v}$ after $\mathbf{k}$ steps.
- $\left(r_{1}^{v}, r_{2}^{v}, \ldots, r_{K}^{v}\right)$, truncated vector of landing probabilities.
- Choices of $\left(w_{1}, \ldots, w_{K}\right)$ define sweep directions through space.
- Optimistically:



## The space of landing probabilities



- SBM: 2000 nodes, Target \& Other blocks, $\mathrm{p}_{\text {in }}=0.2, \mathrm{p}_{\text {out }}=0.05$
- One seed node (uniformly at random from Target set)


## The space of landing probabilities



- SBM: 2000 nodes, Target \& Other blocks, $\mathrm{p}_{\text {in }}=0.2, \mathrm{p}_{\text {out }}=0.05$
- One seed node (uniformly at random from Target set)


## The space of landing probabilities



- Geometric discriminant function: sweeps through the space of landing probabilities following vector from $\mathbf{b}$ to $\mathbf{a}$.


## The space of landing probabilities



- Fisher discriminant functions: Clearly exist better linear and quadratic functions. Forward pointer, will return.


## The space of landing probabilities





- Focus on deriving optimal Geometric discriminant function first.


## Geometric discriminant functions

- Let $\mathbf{r}=\left(r_{1}, \ldots, r_{K}\right)$ be the landing probabilities of a node
- Let $\mathbf{a}=\left(a_{1}, \ldots, a_{K}\right)$ be the Target class centroid
- Let $\mathbf{b}=\left(b_{1}, \ldots, b_{K}\right)$ be the $\mathbf{O t h e r}$ class centroid
- Then $f(\mathbf{r})=(\mathbf{a}-\mathbf{b})^{T} \mathbf{r}$ is the geometric discriminant function.
- Notice: $f(\mathbf{r})$ increases when $\mathbf{r}$ moves in direction of $\mathbf{a - b}$.
- Can classify nodes based on thresholds of $f(\mathbf{r})$.



## Personalized PageRank is "optimal"

- Main Theorem (informal version).

For 2-block SBM with equal sized blocks and edge densities $p_{i n}, p_{\text {out }}$ :

$$
a_{k}-b_{k}=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)^{k}
$$

and the optimal geometric classifier is therefore: which is $\operatorname{PPR}(!)$ with $\alpha_{*}=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)$.

$$
\sum_{k=1}^{K}\left(\alpha_{*}\right)^{k} r_{k}
$$



## Personalized PageRank is "optimal"

- Main Theorem (informal version).

For 2-block SBM with equal sized blocks and edge densities $p_{i n}, p_{o u t}$ :

$$
a_{k}-b_{k}=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)^{k}
$$

and the optimal geometric classifier is therefore: which is $\operatorname{PPR}(!)$ with $\alpha_{*}=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)$.

- Two main parts:

1. Centroids $\mathbf{a}, \mathbf{b}$ concentrate on quantities determined by the solution to a linear recurrence relation.
2. That linear recurrence relation can be solved and yields PPR.


## PPR is "optimal": Proof idea

## - Part 1: Concentration of landing probabilities

Lemma 1. For any $\epsilon, \delta>0$, there is an $n$ sufficiently large such that the random landing probabilities $\left(\hat{a}_{1}, \ldots, \hat{a}_{K}\right)$ and $\left(\hat{b}_{1}, \ldots, \hat{b}_{K}\right)$ for a uniform random walk on $G_{n}$ starting in the seed block satisfy the following conditions with probability at least $1-\delta$ for all $k>0$ :

$$
\begin{align*}
& N \hat{a}_{k} \in\left[(1-\epsilon) \frac{A_{k}}{A_{k}+B_{k}},(1+\epsilon) \frac{A_{k}}{A_{k}+B_{k}}\right] \text { and }  \tag{1}\\
& N \hat{b}_{k} \in\left[(1-\epsilon) \frac{B_{k}}{A_{k}+B_{k}},(1+\epsilon) \frac{B_{k}}{A_{k}+B_{k}}\right], \tag{2}
\end{align*}
$$

where $A_{k}, B_{k}$ are the solutions to the matrix recurrence relation

$$
\left\{\begin{array}{l}
A_{k}=N\left(p_{\text {in }} A_{k-1}+p_{o u t} B_{k-1}\right) \\
B_{k}=N\left(p_{o u t} A_{k-1}+p_{\text {in }} B_{k-1}\right)
\end{array}\right.
$$

with $A_{0}=1, B_{0}=0$.

## PPR is "optimal": Proof idea

- Part 1: Concentration of landing probabilities

Lemma 1. For any $\epsilon, \delta>0$, there is an $n$ sufficiently large such that the random landing probabilities $\left(\hat{a}_{1}, \ldots, \hat{a}_{K}\right)$ and ( $\hat{b}_{1}, \ldots, \hat{b}_{K}$ ) for a uniform random walk on $G_{n}$ starting in the seed block satisfy the following conditions with probability at least $1-\delta$ for all $k>0$ :

$$
\begin{align*}
& N \hat{a}_{k} \in\left[(1-\epsilon) \frac{A_{k}}{A_{k}+B_{k}},(1+\epsilon) \frac{A_{k}}{A_{k}+B_{k}}\right] \text { and }  \tag{1}\\
& N \hat{b}_{k} \in\left[(1-\epsilon) \frac{B_{k}}{A_{k}+B_{k}},(1+\epsilon) \frac{B_{k}}{A_{k}+B_{k}}\right] \tag{2}
\end{align*}
$$

where $A_{k}, B_{k}$ are the solutions to the matrix recurrence relation

$$
\left\{\begin{array}{l}
A_{k}=N\left(p_{\text {in }} A_{k-1}+p_{\text {out }} B_{k-1}\right) \\
B_{k}=N\left(p_{\text {out }} A_{k-1}+p_{\text {in }} B_{k-1}\right),
\end{array}\right.
$$

$$
\text { with } A_{0}=1, B_{0}=0
$$

- $A_{k}, B_{k}$ interpretable as length-k walk count to nodes in block 1 vs. 2.
- For large n, block walk counts increase by factors of $\sim$ E[degree].


## More general SBMs

- For SBMs with $\mathbf{C}>\mathbf{2}$ blocks and/or with arbitrary P:
- Seed set expansion asks: identify nodes in a target block set.
- With conditions on equal expected degrees, PPR(!).
- Without conditions, still:
- Asymptotically optimal weights for geometric classification still obtainable from solutions to a matrix recurrence relation.



## Empirical vs. theoretical centroids

- 2048-node, 4-block SBM, empirical class centroids vs. theory:
- a, Target blocks
- b, Other blocks





## Empirical vs. theoretical centroids

- 2048-node, 4-block SBM, empirical class centroids vs. theory:
- a, Target blocks
- b, Other blocks


From matrix recurrence relation


## Theories of graph diffusion

- Other motivations for PPR:
- Random Surfer Model (Brin-Page '98)
- Cheeger inequalities for PPR, HK (Andersen et al '06, Chung '09)
- Local spectral algorithm with regularization (Mahoney et al. '12)
- Our work shows PPR can be derived as "optimal" geometric classifier.
- Also motivates how to choose $\operatorname{PPR} \alpha$, as $\alpha=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)$.



## Theories of graph diffusion

- Other motivations for PPR:
- Random Surfer Model (Brin-Page '98)
- Cheeger inequalities for PPR, HK (Andersen et al '06, Chung '09)
- Local spectral algorithm with regularization (Mahoney et al. '12)
- Our work shows PPR can be derived as "optimal" geometric classifier.
- Also motivates how to choose $\operatorname{PPR} \alpha$, as $\alpha=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)$.
- Most importantly: also opens door to methods beyond PPR.



## PPR is "optimal" in a narrow sense



- Discriminant functions that model higher moments of point clouds?


## Fisher discriminant functions



- Discriminant functions that model higher moments of point clouds.


## Fisher discriminant functions

- Let $\mathbf{z}$ be the latent class of each node.
- Capture (mean, variance) of class point clouds:

$$
\begin{aligned}
& \operatorname{Pr}(\mathbf{r} \mid z=1) \propto\left|\Sigma_{a}\right|^{-\frac{1}{2}} \exp \left(-\frac{1}{2}(\mathbf{r}-\mathbf{a})^{T} \Sigma_{a}^{-1}(\mathbf{r}-\mathbf{a})\right) \\
& \operatorname{Pr}(\mathbf{r} \mid z=0) \propto\left|\Sigma_{b}\right|^{-\frac{1}{2}} \exp \left(-\frac{1}{2}(\mathbf{r}-\mathbf{b})^{T} \Sigma_{b}^{-1}(\mathbf{r}-\mathbf{b})\right)
\end{aligned}
$$

- Log-likelihood ratio as discriminant function:

$$
g(\mathbf{r})=\log \frac{\operatorname{Pr}(\mathbf{r} \mid z=1) \operatorname{Pr}(z=1)}{\operatorname{Pr}(\mathbf{r} \mid z=0) \operatorname{Pr}(z=0)}
$$



## Fisher discriminant functions

- Three approaches:

General: $\quad g_{2}(\mathbf{r}) \propto\left(\Sigma_{a}^{-1} \mathbf{a}-\Sigma_{b}^{-1} \mathbf{b}\right)^{T} \mathbf{r}+\frac{1}{2} \mathbf{r}^{T}\left(\Sigma_{b}^{-1}-\Sigma_{a}^{-1}\right) \mathbf{r}$
Assume $\Sigma_{a}=\Sigma_{b}=\Sigma: \quad g_{1}(\mathbf{r}) \propto \Sigma^{-1}(\mathbf{a}-\mathbf{b})^{T} \mathbf{r}$
Assume $\Sigma_{a}=\Sigma_{b}=I: \quad g_{0}(\mathbf{r}) \propto(\mathbf{a}-\mathbf{b})^{T} \mathbf{r}$

- We call the first two methods QuadSBMRank, LinSBMRank.
- Perhaps reasonable to assume equal covariances; effective.
- PPR follows from an assumption of uniform variance, no covariance.



## Fisher discriminant functions

- Three approaches:

General: $\quad g_{2}(\mathbf{r}) \propto\left(\Sigma_{a}^{-1} \mathbf{a}-\Sigma_{b}^{-1} \mathbf{b}\right)^{T} \mathbf{r}+\frac{1}{2} \mathbf{r}^{T}\left(\Sigma_{b}^{-1}-\Sigma_{a}^{-1}\right) \mathbf{r}$
Assume $\Sigma_{a}=\Sigma_{b}=\Sigma: \quad g_{1}(\mathbf{r}) \propto \Sigma^{-1}(\mathbf{a}-\mathbf{b})^{T} \mathbf{r}$
Assume $\Sigma_{a}=\Sigma_{b}=I: \quad g_{0}(\mathbf{r}) \propto(\mathbf{a}-\mathbf{b})^{T} \mathbf{r}$

- We call the first two methods QuadSBMRank, LinSBMRank.
- Perhaps reasonable to assume equal covariances; effective.
- PPR follows from an assumption of uniform variance, no covariance.
- Open challenge: Possible to show asymptotic normality and characterize covariance matrices?



## Evaluation: recall curves

- SBM with 2 blocks, 64 nodes/block, 1 seed node.
- Recall that Belief Propagation reaches resolution limit.


$$
\begin{array}{ll}
\hline-=- & \text { (1.0) Lin-SBMRank @ } \alpha_{e s t} \\
\cdots & \text { (1.0) Quad-SBMRank @ } \alpha_{*} \\
- & \text { (1.0) Belief Prop. } \\
=- & \text { (1.0) Quad-SBMRank @ } \alpha_{e s t} \\
\cdots & \text { (1.0) Lin-SBMRank @ } \alpha_{*} \\
- & \text { (0.93) Heat Kernel @ } 2 \\
\cdots & \text { (0.92) PageRank @ } \alpha_{*} \\
=- & \text { (0.92) PageRank @ } \alpha_{e s t} \\
\hline
\end{array}
$$

- Easy instance (pin >> pout):
- Everything does well.


## Evaluation: recall curves

- SBM with 2 blocks, 64 nodes/block, 1 seed node.
- Recall that Belief Propagation reaches resolution limit.

- Hard instance...
- PPR/HK lost all recall, LinSBMRank and QuadSBMRank near BP.


## Evaluation: recall curves

- SBM with 2 blocks, 64 nodes/block, 1 seed node.
- Recall that Belief Propagation reaches resolution limit.


> | $\cdots$ | $(0.72)$ Lin-SBMRank @ $\alpha_{*}$ |
| :--- | :--- |
| $-=-$ | $(0.71)$ Lin-SBMRank @ $\alpha_{e s t}$ |
| $\cdots$ | $(0.71)$ Quad-SBMRank @ $\alpha_{*}$ |
| $-=$ | $(0.69)$ Quad-SBMRank @ $\alpha_{e s t}$ |
| - | $(0.66)$ Belief Prop. |
| $\cdots$ | $(0.58)$ PageRank @ $\alpha_{*}$ |
| $=-$ | (0.58) PageRank @ $\alpha_{e s t}$ |
| - | (0.57) Heat Kernel @ 2 |

- Even harder instance...
- LinSBMRank and QuadSBMRank outperforming BP by a hair...?


## Evaluation: recall curves

- SBM with 2 blocks, 64 nodes/block, 1 seed node.
- Recall that Belief Propagation reaches resolution limit.


$$
\begin{array}{|ll|}
\hline \cdots & \text { (0.56) Quad-SBMRank @ } \alpha_{*} \\
\cdots & \text { (0.56) Lin-SBMRank @ } \alpha_{*} \\
\cdots & \text { (0.55) PageRank @ } \alpha_{*} \\
- & \text { (0.55) Heat Kernel @ } 2 \\
-= & \text { (0.53) Lin-SBMRank @ } \alpha_{e s t} \\
-- & \text { (0.53) PageRank @ } \alpha_{e s t} \\
- & \text { (0.53) Belief Prop. } \\
-= & \text { (0.53) Quad-SBMRank @ } \alpha_{e s t} \\
\hline
\end{array}
$$

- Impossible ( $\mathrm{p}_{\mathrm{in}}=\mathrm{p}_{\text {out }}$ ):
- Nothing works.


## Evaluation: resolution limit

- Pearson correlation $\mathbf{r}$ between true partition and inferred partition.
- Empirically, we see LinSBMRank and QuadSBMRank get very close to resolution limit (dotted line), with slower decay rate.


PPR, HK, LinSBMRank, QuadSBMRank, BP

## Conclusions

- Personalized PageRank with $\alpha=\left(\frac{p_{\text {in }}-p_{\text {out }}}{p_{\text {in }}+p_{\text {out }}}\right)$ is optimal geometric discriminant function for balanced 2-block SBM.
- Geometric discriminant functions for more general block models follow from recurrence relation.
- Landing probabilities are correlated; correcting for higher moments in the space of landing probabilities greatly improves classification.
- In practice: fit GMMs in space of landing probs.
- A new perspective on diffusion-based ranking that can hopefully open new doors.
- Pre-print:

Isabel Kloumann, Johan Ugander, Jon Kleinberg "Block Models and Personalized PageRank" arXiv:1607.03483


## Open directions

- Model covariance of landing probabilities?
- Currently requires at least ~logarithmic degrees (we think); possible to derive weights for bounded degree SBMs?
- Better classifiers in the space of landing probabilities for other random walks? (Non-backtracking, etc.)
- Not just SBM? Optimal weights for dcSBM, core-periphery, Hoff latent space model, etc, etc.
- Slow decay beyond resolution limit?
- Pre-print:

Isabel Kloumann, Johan Ugander, Jon Kleinberg "Block Models and Personalized PageRank" arXiv:1607.03483


