Approximation vs Sensitivity: Stable Algorithms Under Input Perturbations

COCOON 2026

Yuichi Yoshida
National Institute of Informatics

1 / 4

Introduction

Motivation, definitions, and landscape

Toy Example: Exact Maximum Matching

Exact maximum matching: choose a maximum matching, breaking ties in favor of lower-indexed edges.

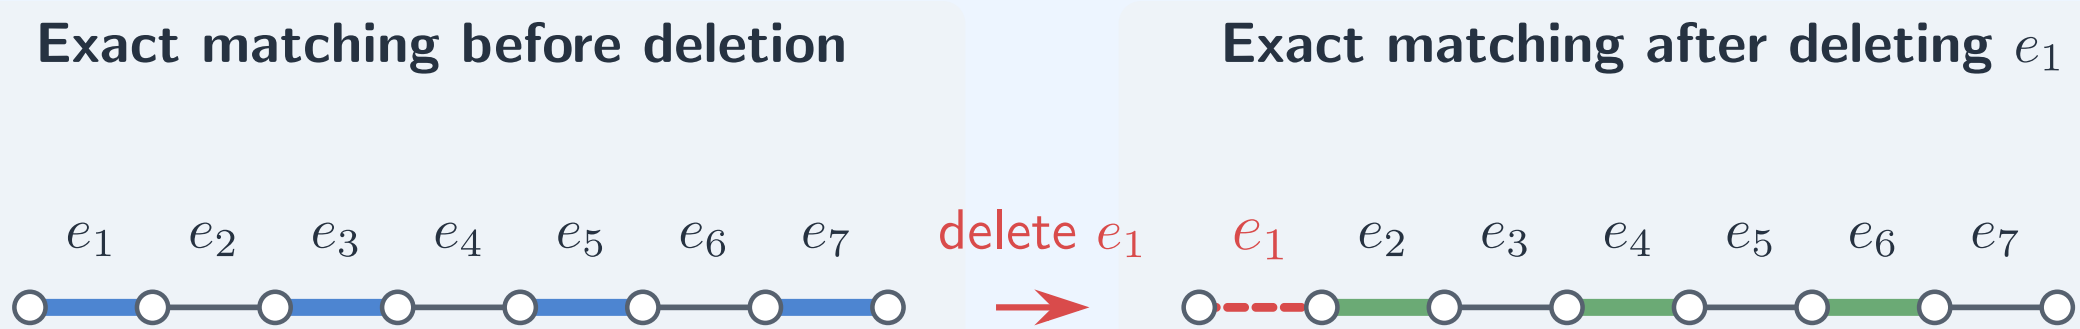

Exact maximum matching can be unstable: deleting e_1 changes every selected edge—four leave and three enter.

Why Stability Matters?

Algorithms should be stable, not only near-optimal and fast.

- Large-scale inputs are often incomplete, dynamic, and noisy.
- If the output changes drastically whenever the input changes slightly, it becomes hard to reuse solutions or make reliable decisions.
- Approximation quality does not tell us how much the chosen solution changes.

how good?
Approximation

how fast?
Running time

how much change?
Stability

Sensitivity of Deterministic Algorithms: Compare Outputs

For a deterministic algorithm A that outputs an edge set, write the change caused by deleting e as

$$\Delta_e(A, G) = |A(G) \triangle A(G - e)|.$$

Def. *Average/worst-case sensitivity* of **deterministic** A on G [VY23, YZ21]:

$$\mathrm{Sens}_{\mathrm{avg}}(A, G) = \mathbb{E}_{e \sim E}\big[\Delta_e(A, G)\big],$$
$$\mathrm{Sens}_{\mathrm{max}}(A, G) = \max_{e \in E} \Delta_e(A, G).$$

Sensitivity of Randomized Algorithms: Compare Distributions

Compare $P = A(G)$ and $Q = A(G - e)$ via the *Earth mover's distance*:

$$W_1(P, Q) = \inf_{\gamma \in \Gamma(P,Q)} \mathbb{E}_{(X,Y)\sim\gamma} |X \triangle Y|.$$

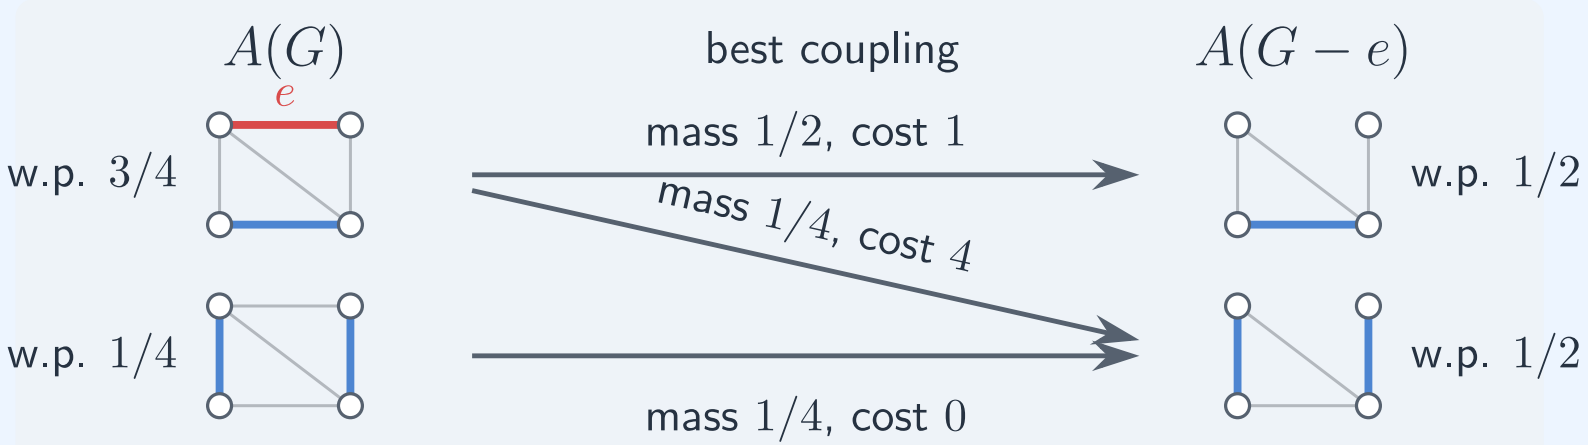

$$W_1 = (1/2) \cdot 1 + (1/4) \cdot 4 + (1/4) \cdot 0 = 3/2$$

Sensitivity of Randomized Algorithms: Aggregate Changes

The output change of a randomized algorithm A under deleting edge e is

$$\Delta_e(A, G) = W_1(A(G), A(G - e)).$$

(This extends the deterministic definition.)

Def. *average/worst-case sensitivity* of **randomized** A on G [VY23, YZ21]:

$$\mathrm{Sens}_{\mathrm{avg}}(A, G) = \mathbb{E}_{e \sim E}\big[\Delta_e(A, G)\big],$$
$$\mathrm{Sens}_{\max}(A, G) = \max_{e \in E} \Delta_e(A, G).$$

Why Not (Pure) Differential Privacy?

Pure DP asks that neighboring inputs induce almost indistinguishable output distributions.

> **Example: Matching**
>
> - On $G - e$, outputting e has probability 0.
> - Pure DP therefore forbids outputting e on G as well.
> - Since e was arbitrary, no edge can be output.

So pure DP is too strict for explicit combinatorial solutions.

Known Results

We now have low-sensitivity algorithms across several problem families.

Graphs
MST; global min cut; min s-t cut; max matching [VY23, YZ21, YZ26a, YZ26b]

Clustering
spectral clustering; Euclidean k-clustering; [PY20, YI22]

Dynamic programming
max weight chain; LIS; interval scheduling; LCS; RNA folding [KY22a,KY22b]

Learning
decision tree learning [HY23]

Today's focus.

- Techniques for low-sensitivity upper bounds via matching
- Sensitivity lower bounds for approximation algorithms.

Before We Dive In: Pick a Problem

Pick one problem and one algorithm for it.

What would sensitivity mean?

- Neighboring inputs?
- Output distance?

> Can one local input change affect the solution globally, or can the cascade be contained?

2 / 4

Upper Bounds

Stable matching algorithms under edge deletions

What Is Known for Maximum Matching?

Approach	Approx.	Sensitivity	Sens. type
Exact	1	$\Omega(n)$	avg.
Randomized greedy [VY23]	$1/2$	$O(1)$	worst
Local computation algorithms [YYI09, VY23]	$1-\varepsilon$	$\Delta^{O(1/\varepsilon^2)}$	avg.
Gibbs dist. [YZ26a]	$1-\varepsilon$	$\Delta^{O(1/\varepsilon)}$	worst
Gibbs + sparsifier [YZ26a]	$1-\varepsilon$	$\sqrt{n}(\varepsilon^{-1}\log n)^{O(1/\varepsilon)}$	worst
Auction + alternating forest [YZ26b]	$1-\varepsilon$	$\log^{2^{O(1/\varepsilon^2)}} n$	worst

Today: randomized greedy, local computation algorithms, and Gibbs distributions.

The common task: contain the **shock wave** from one deleted edge.

Approach 1: Randomized Greedy

A fixed priority can make one deletion propagate far.

Randomized greedy replaces it by a random priority.

```
choose a random priority order on E
for e in that order:
  if both endpoints are free:
    add e to M
```

- The output remains a maximal matching, hence a $1/2$-approximation.
- Long influence chains become unlikely: $\mathrm{Sens}_{\max} = O(1)$ [CHK16].

Why Random Priority Stops the Cascade

Focus on the local priority race around the deleted edge e.

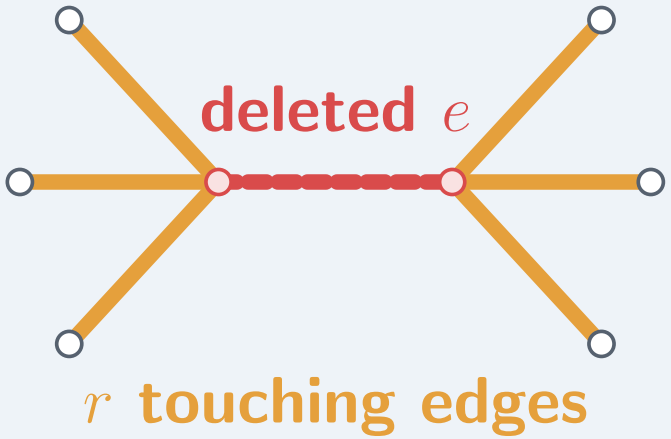

a touching edge comes first
e **is blocked** $\Rightarrow$ **no cascade**

e comes first
chance $1/(r+1)$

Randomness balances risk: chance $1/(r+1)$ × change at most $r+1$ gives expected change at most 1.

Approach 2: Local Computation Algorithms

A *local computation algorithm* (LCA) is an oracle view of a global output.

Rather than build all of M, it receives a query $\texttt{IsMatched}(f)$ and answers

$$\texttt{IsMatched}(f) \in \{\text{yes}, \text{no}\} \quad \text{meaning } f \in M?$$

The oracle's answers implicitly define one consistent matching M.

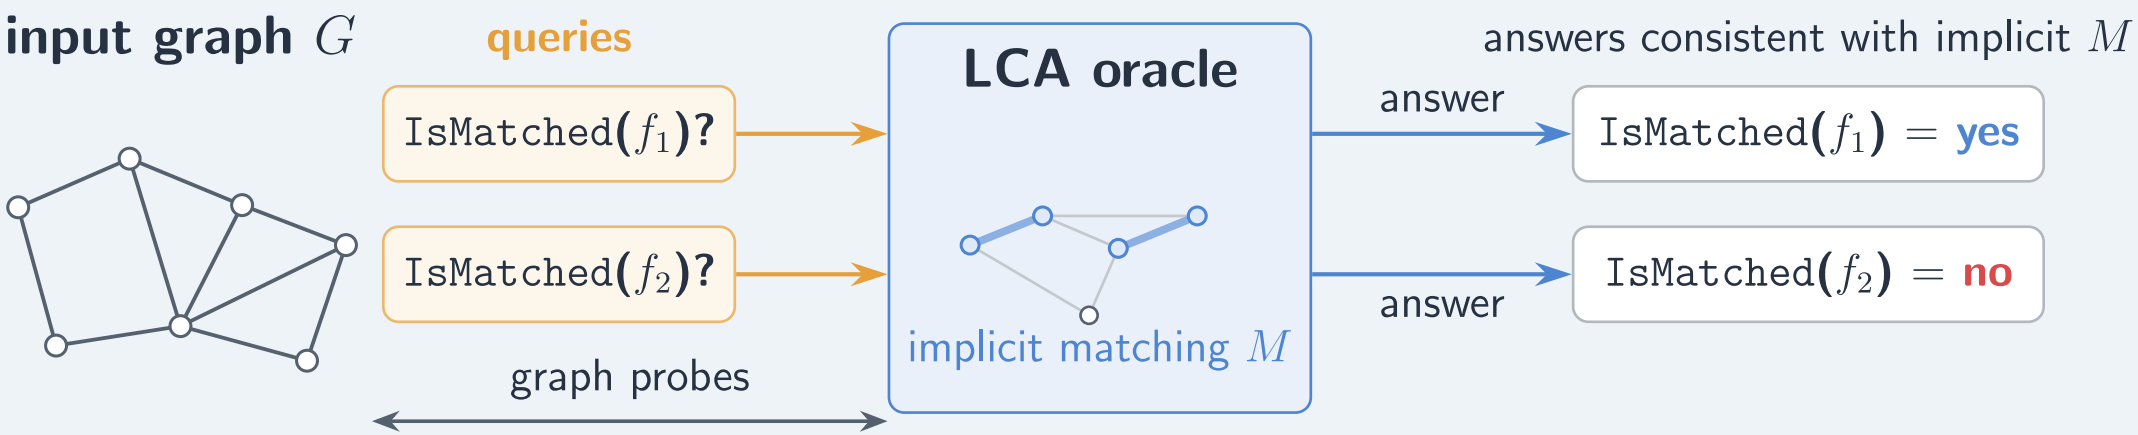

Approach 2: Local Views Control Changes

For a queried edge f, let $\mathrm{View}(f)$ be the edges inspected to decide whether $f \in M$.

Query complexity q means $|\mathrm{View}(f)| \leq q$.

Run the LCA with the same random coins on G and $G - e$:

$$e \notin \mathrm{View}(f) \implies \texttt{IsMatched}_G(f) = \texttt{IsMatched}_{G-e}(f).$$

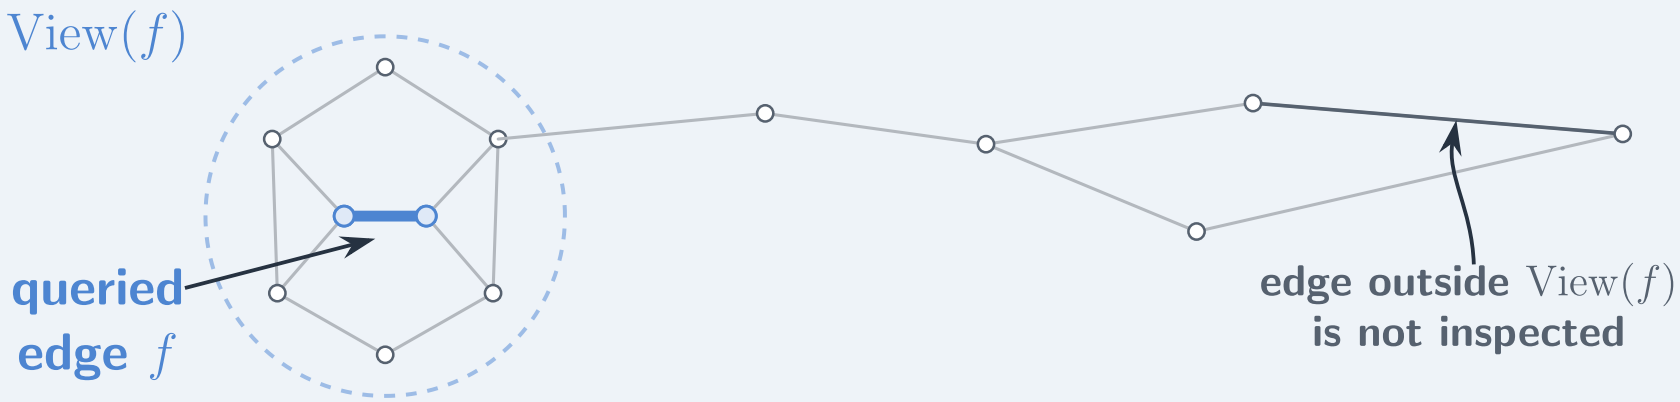

Approach 2: Query Complexity Gives Average Sensitivity

If $|\mathrm{View}(f)| \leq q$, then

$$\Pr_{e \sim E}[\text{answer to } f \text{ changes}] \leq \frac{|\mathrm{View}(f)|}{m} \leq \frac{q}{m}.$$

So,

$$\mathrm{Sens}_{\mathrm{avg}} = \frac{1}{m} \sum_{e \in E} |A(G) \triangle A(G - e)| \leq \sum_{f \in E} \Pr_{e \sim E}[\text{answer to } f \text{ changes}] \leq q.$$

A $(1 - \varepsilon)$-approximate LCA with $q = \Delta^{O(1/\varepsilon^2)}$ is known [YYI09]. So,

$$\mathrm{Sens}_{\mathrm{avg}} \leq \Delta^{O(1/\varepsilon^2)}.$$

Approach 3: Gibbs Sampling over Matchings

Gibbs distribution over matchings:

$$\Pr[M] \propto \lambda^{|M|}.$$

- $\lambda = 1$ gives the uniform distribution over matchings.
- For $\lambda > 1$, changing λ creates a tradeoff:
  - **Larger λ:** more weight on larger matchings $\Rightarrow$ better approximation.
  - **Smaller λ:** closer to uniform $\Rightarrow$ better sensitivity.

Tune λ just large enough for a $(1 - \varepsilon)$-approximation.

The remaining question: how far can one deleted edge propagate through the sampled matching?

Approach 3: Couplings Control Cascades

Couple Gibbs distributions on G and $G - e$.

A disagreement starts at e, then either dies or moves to a neighboring edge.

One deleted edge can only trigger a local cascade

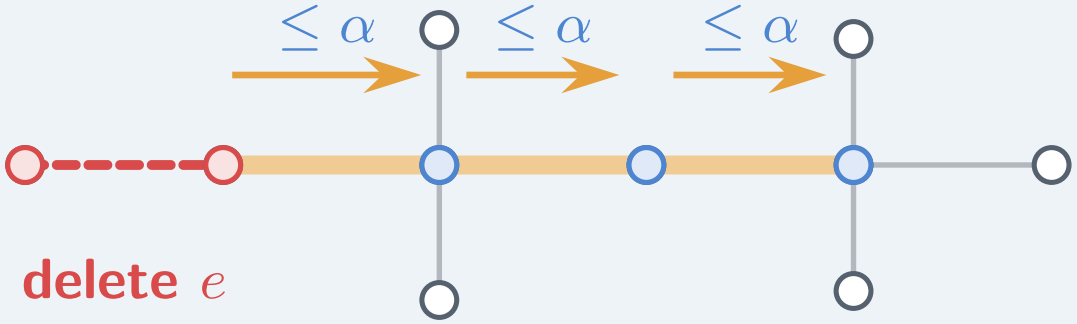

per-step continuation

$$\leq \alpha = \frac{\lambda\Delta}{1+\lambda\Delta}$$

The expected cascade length is $1 + \alpha + \alpha^2 + \cdots = O(\lambda\Delta)$.

With the tuned λ, we can obtain $\Delta^{O(1/\varepsilon)}$ worst-case sensitivity [YZ26a].

3 / 4

Lower Bounds

Sensitivity barriers for approximation algorithms

From Exact to Approximation Lower Bounds

Exact optimization: sensitivity lower bounds are easy to force.
Approximation: does the added slack permit stability?

The PCP theorem gives a reduction from 3SAT to Max 3SAT:

$$\text{3SAT} \xrightarrow{\text{PCP theorem}} \text{Max 3SAT}.$$

- **YES:** some assignment satisfies every clause.
- **NO:** every assignment violates at least an ε fraction of clauses.

Can this route also carry instability—and yield a **sensitivity lower bound for approximation algorithms**?

Polynomial Sensitivity Lower Bounds [FY26]

For CSPs, neighbors are I and $I - e$ (one constraint deleted); outputs are assignments with Hamming/Wasserstein distance.

Thm. Any $(1-\varepsilon)$-approximation algorithm for Max 3SAT needs worst-case sensitivity $\Omega(n^{\delta})$ for some $\varepsilon, \delta > 0$.

Transfers to

clique / independent set
$n^{-\varepsilon}$-approx.

vertex cover
$(1+\varepsilon)$-approx.

max cut
$(1-\varepsilon)$-approx.

- Even with approximation, polynomial instability can be unavoidable.
- **Information-theoretic:** holds even for inefficient randomized algorithms.

What Must a Reduction Preserve?

Start from the ordinary reduction flow.

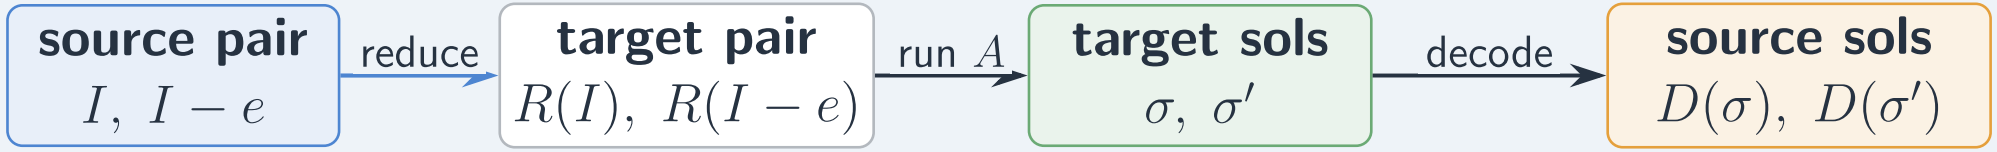

For sensitivity, the reduction and decoding maps must be stable.

Input stability

$$d(R(I), R(I-e)) \le C_R\, d(I, I-e)$$

Output stability

$$W_1\big(D(\sigma), D(\sigma')\big) \le C_D\, d(\sigma, \sigma')$$

$$\mathrm{Sens}_{\mathrm{target}} \ge \frac{\mathrm{Sens}_{\mathrm{source}}}{C_R C_D}.$$

PCP with Sensitivity Preserved

Start from E2LIN with $\Omega(n)$ sensitivity; obtain Label Cover with gap ε, then Max 3SAT.

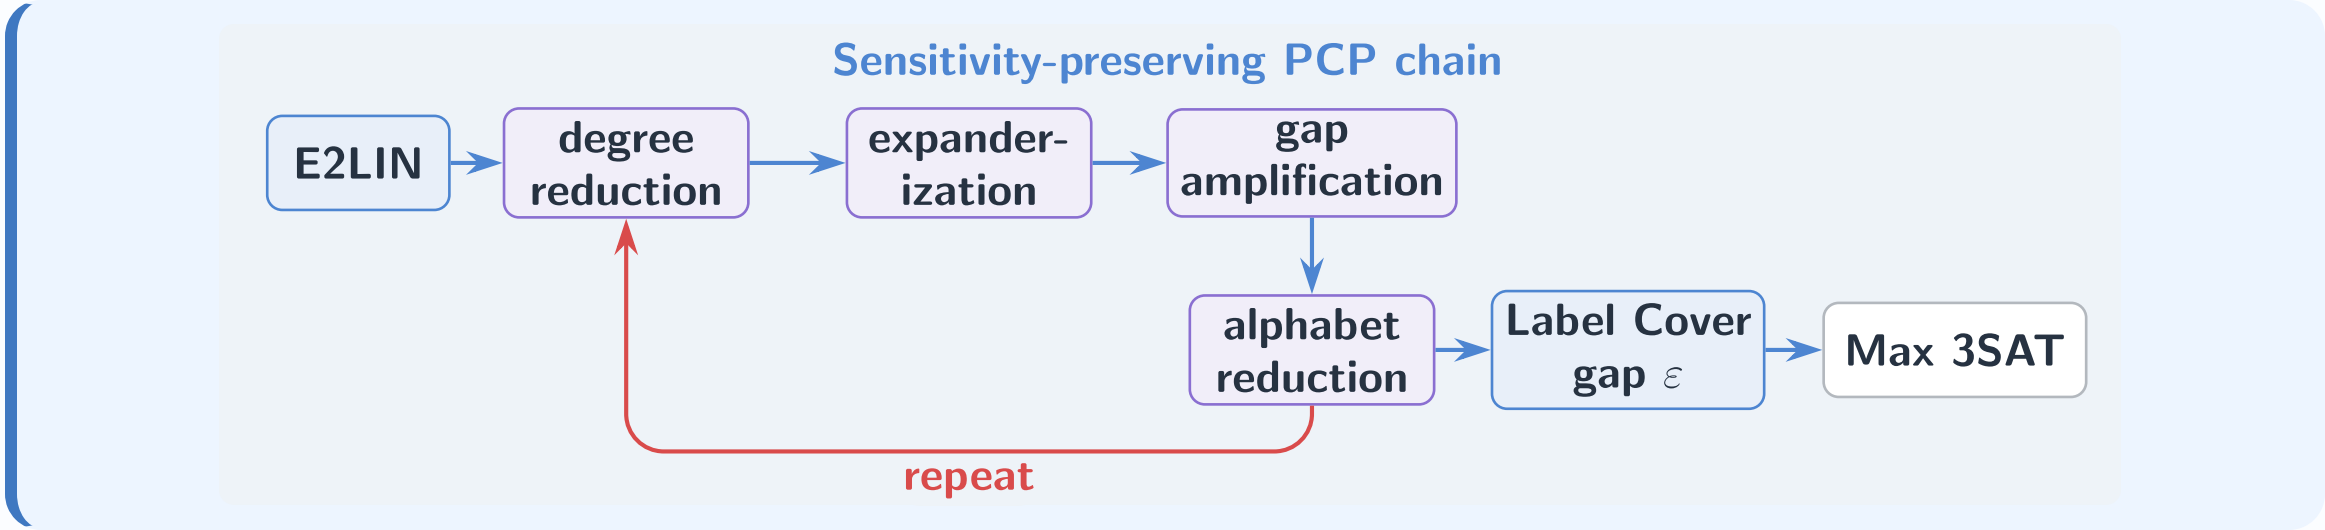

- **Input stability:** check that one change in the source instance causes only bounded changes in the target instance.
- **Output stability:** rebuild the decoder to be smooth: sample instead of hard majority; use randomized thresholds instead of brittle rounding.

Sensitivity-to-Locality Transfer [FY26]

LOCAL model. After t rounds, the output at v depends only on its radius-t neighborhood.

An input change at one edge can affect only outputs within distance t of its endpoints [FY26]:

t-local $\Longrightarrow \mathrm{Sens}_{\max} \leq \Delta^{O(t)}$.

$\mathrm{Sens}_{\max} \geq s \quad \Longrightarrow \quad t = \Omega(\log_\Delta s)$.

A sensitivity lower bound becomes a locality lower bound.

New Lower Bound for Dominating Set

For some fixed $\beta \in (0, 1)$, consider $O(\log^{\beta} \Delta)$-approximation.

Via Sensitivity Lower Bounds [FHY26]

$$\Omega\left(\frac{\log n}{\log \Delta}\right)$$

LOCAL rounds.

Known bound [KMW16]

$$\Omega\left(\min\left\{\frac{\log \Delta}{\log\log \Delta}, \sqrt{\frac{\log n}{\log\log n}}\right\}\right)$$

LOCAL rounds.

Combine the two bounds

$$\Omega\left(\sqrt{\frac{\log n}{\log\log n}}\right)$$

LOCAL rounds, **independent of Δ**.

4 / 4

Summary

Sensitivity as an algorithmic design objective

What Did We Learn?

Stability is a design objective, not a free consequence of approximation.

What we saw today

- exact optimum can be unstable; approximation creates slack;
- upper bounds: randomization, local decisions, and smoothing;
- lower bounds: sensitivity-preserving reductions.

Not covered today

- LP/SDP regularization: stabilize a relaxation, then round.
- Lipschitz algorithms [KY23]: weighted variant of sensitivity.
- $\Omega(n)$ sensitivity lower bounds for Max E3LIN2 / Max 3SAT via LTCs [YZ26c].

Where Is the Frontier?

Approximation–sensitivity tradeoff

For each problem, what is the optimal tradeoff between approximation and sensitivity?

When does structure help?

- **Problem structure:** which CSPs admit stable approximation?
- **Instance structure:** what changes under restrictions such as planarity?

Does computational power buy stability?

- **Randomness:** how much lower can sensitivity be with randomized algorithms?
- **Efficiency:** how much lower without a polynomial-time constraint?

Sensitivity turns stability into an algorithmic question.

Return to your problem: can one local input change be contained, or must it cascade?

Selected References

Average sensitivity and applications

[MY19] Shogo Murai and Yuichi Yoshida, **Sensitivity Analysis of Centralities on Unweighted Networks**, WWW 2019.

[PY20] Pan Peng and Yuichi Yoshida, **Average Sensitivity of Spectral Clustering**, KDD 2020.

[CHK16] Keren Censor-Hillel, Elad Haramaty, and Zohar Karnin, **Optimal Dynamic Distributed MIS**, PODC 2016.

[YYI09] Yuichi Yoshida, Masaki Yamamoto, and Hiro Ito, **Improved Constant-Time Approximation Algorithms for Maximum Matchings and Other Optimization Problems**, STOC 2009.

[VY23] Nithin Varma and Yuichi Yoshida, **Average Sensitivity of Graph Algorithms**, SODA 2021 / SICOMP 2023.

[KY22a] Soh Kumabe and Yuichi Yoshida, **Average Sensitivity of Dynamic Programming**, SODA 2022.

[KY22b] Soh Kumabe and Yuichi Yoshida, **Average Sensitivity of the Knapsack Problem**, ESA 2022.

[YI22] Yuichi Yoshida and Shinji Ito, **Average Sensitivity of Euclidean k-Clustering**, NeurIPS 2022.

[HY23] Satoshi Hara and Yuichi Yoshida, **Average Sensitivity of Decision Tree Learning**, ICLR 2023.

[DY23] Jing Dong and Yuichi Yoshida, **A Batch-to-Online Transformation under Random-Order Model**, NeurIPS 2023.

[EY26] Matthijs Ebbens and Yuichi Yoshida, **Average Sensitivity of Geometric Algorithms**, ITCS 2026.

[SY26] Shinsaku Sakaue and Yuichi Yoshida, **From Average Sensitivity to Small-Loss Regret Bounds under Random-Order Model**, arXiv 2026.

Worst-case sensitivity, lower bounds, and Lipschitz continuity

[AS98] Sanjeev Arora and Shmuel Safra, **Probabilistic Checking of Proofs: A New Characterization of NP**, JACM 1998.

[ALMSS98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy, **Proof Verification and the Hardness of Approximation Problems**, JACM 1998.

[CL23] Yi-Jun Chang and Zhiwei Li, **The Complexity of Distributed Approximation of Packing and Covering Integer Linear Programs**, PODC 2023.

[KMW16] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer, **Local Computation: Lower and Upper Bounds**, JACM 2016.

[YZ21] Yuichi Yoshida and Samson Zhou, **Sensitivity Analysis of the Maximum Matching Problem**, ITCS 2021.

[KY23] Soh Kumabe and Yuichi Yoshida, **Lipschitz Continuous Algorithms for Graph Problems**, FOCS 2023.

[KY24] Soh Kumabe and Yuichi Yoshida, **Lipschitz Continuous Allocations for Optimization Games**, ICALP 2024.

[LVYZ24] Wenzheng Liu, Grigoris Velegkas, Yuichi Yoshida, and Samson Zhou, **Pointwise Lipschitz Continuous Graph Algorithms via Proximal Gradient Analysis**, 2024.

[KY25] Soh Kumabe and Yuichi Yoshida, **Lipschitz Continuous Algorithms for Covering Problems**, SODA 2025.

[GKY25] Tatsuya Gima, Soh Kumabe, and Yuichi Yoshida, **Courcelle's Theorem for Lipschitz Continuity**, ESA 2025.

[YZ26a] Yuichi Yoshida and Zihan Zhang, **Low-Sensitivity Matching via Sampling from Gibbs Distributions**, SODA 2026.

[YZ26b] Yuichi Yoshida and Samson Zhou, **Approximate Maximum Matching with Polylogarithmic Sensitivity**, submitted.

[YZ26c] Yuichi Yoshida and Zihan Zhang, **Sensitivity Lower Bounds via Locally Testable Codes**, arXiv:2606.31657, 2026.

[FY26] Noah Fleming and Yuichi Yoshida, **Sensitivity Lower Bounds for Approximation Algorithms**, SODA 2026.

[FHY26] Noah Fleming, Max Hopkins, and Yuichi Yoshida, **Non-Signaling Locality Lower Bounds for Dominating Set**, arXiv 2026.