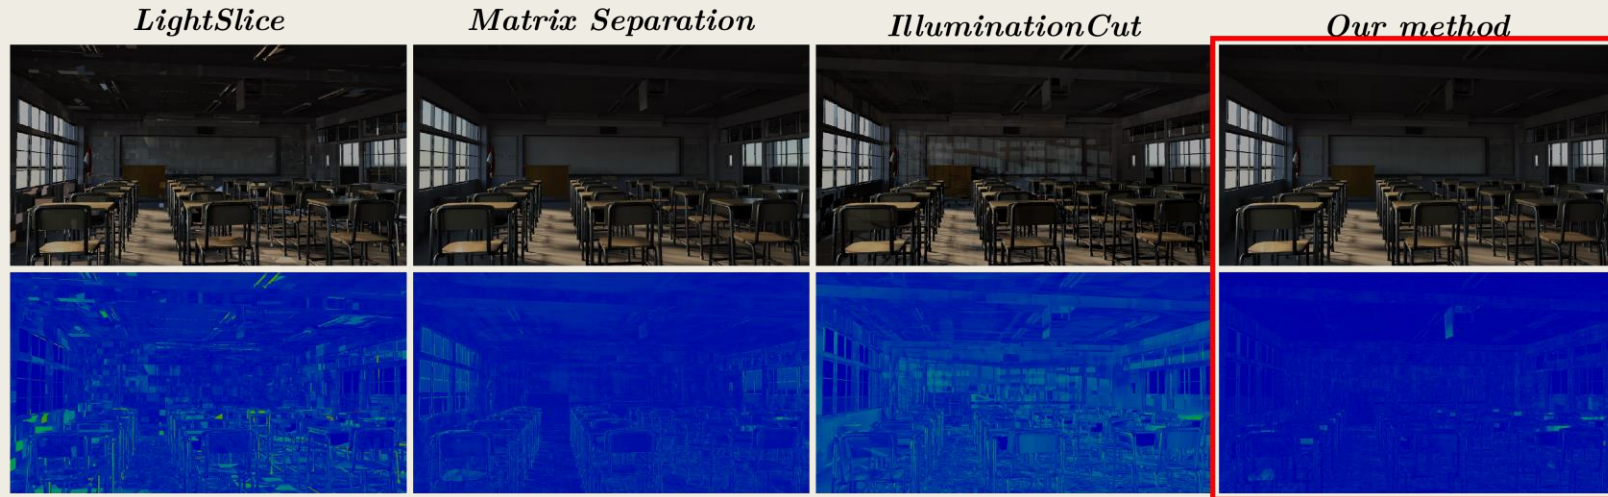


ADAPTIVE MATRIX COMPLETION FOR FAST VISIBILITY COMPUTATIONS WITH MANY LIGHTS RENDERING



Sunrise Wang and Nicolas Holzschuch

Many-Lights Problem

- Direct lighting problem
- Large numbers of point lights
 - Instant Radiosity (*Keller '97*)
 - VPLs approximate indirect lighting
- Can range from thousands to millions of lights
 - Expensive
 - Accurate visibility needs ray-casting




Contributions

- Fast Many-Lights framework
 - Clustering
 - Visibility Approximation
- Adaptive matrix completion to approximate visibility
- Improvements to Adaptive Matrix Completion for boolean visibility

Scalable Many-Lights Methods

- Hierarchical clustering (*Walter et al. '05, Walter et al. '06, Walter et al. '12, Bus et al. '15.*)
 - Hierarchically clusters the VPLs
 - Traverses tree to extract clusters
- Matrix-based methods (*Hasan et al. '07, Ou & Pellacini '11*)
 - Formulates problem as matrix
 - Clusters VPLs with information from sampling matrix

Many-Lights Matrix



- Lights as columns, receiving points as rows
- Final colour of receiver = sum of values in row

Many-Lights Matrix

- Approximately low-rank
 - Most information lies within few basis columns
- Some existing work
 - Generalized Nystrom method for inverse rendering (*Wang et al. '09*)
 - Low-rank and sparse matrix recovery on Many-Lights Matrix (*Huo et al. '15*)
 - Matrix completion on participating media Many-Lights Matrix (*Huo et al. '16*)

Low-rank Matrix Completion



- Complete matrix from partially observed coefficients
- Many or infinite solutions, recover lowest rank

Low-rank Matrix Completion

- Matrix Completion Algorithms
 - Nuclear-norm minimization (*Candes & Recht '09, Cai & Candes '10*)
 - Alternating Least Squares (*Haldar & Hernando '09, Wen et al. '12, Tanner & Wei '16*)
 - Low-rank approximation (*Goreinov et al. '97, Krishnamurthy & Singh '14*)
- Factors that influence required initial observations
 - Rank (number of linearly independent dimensions)
 - Coherence (localization of features)

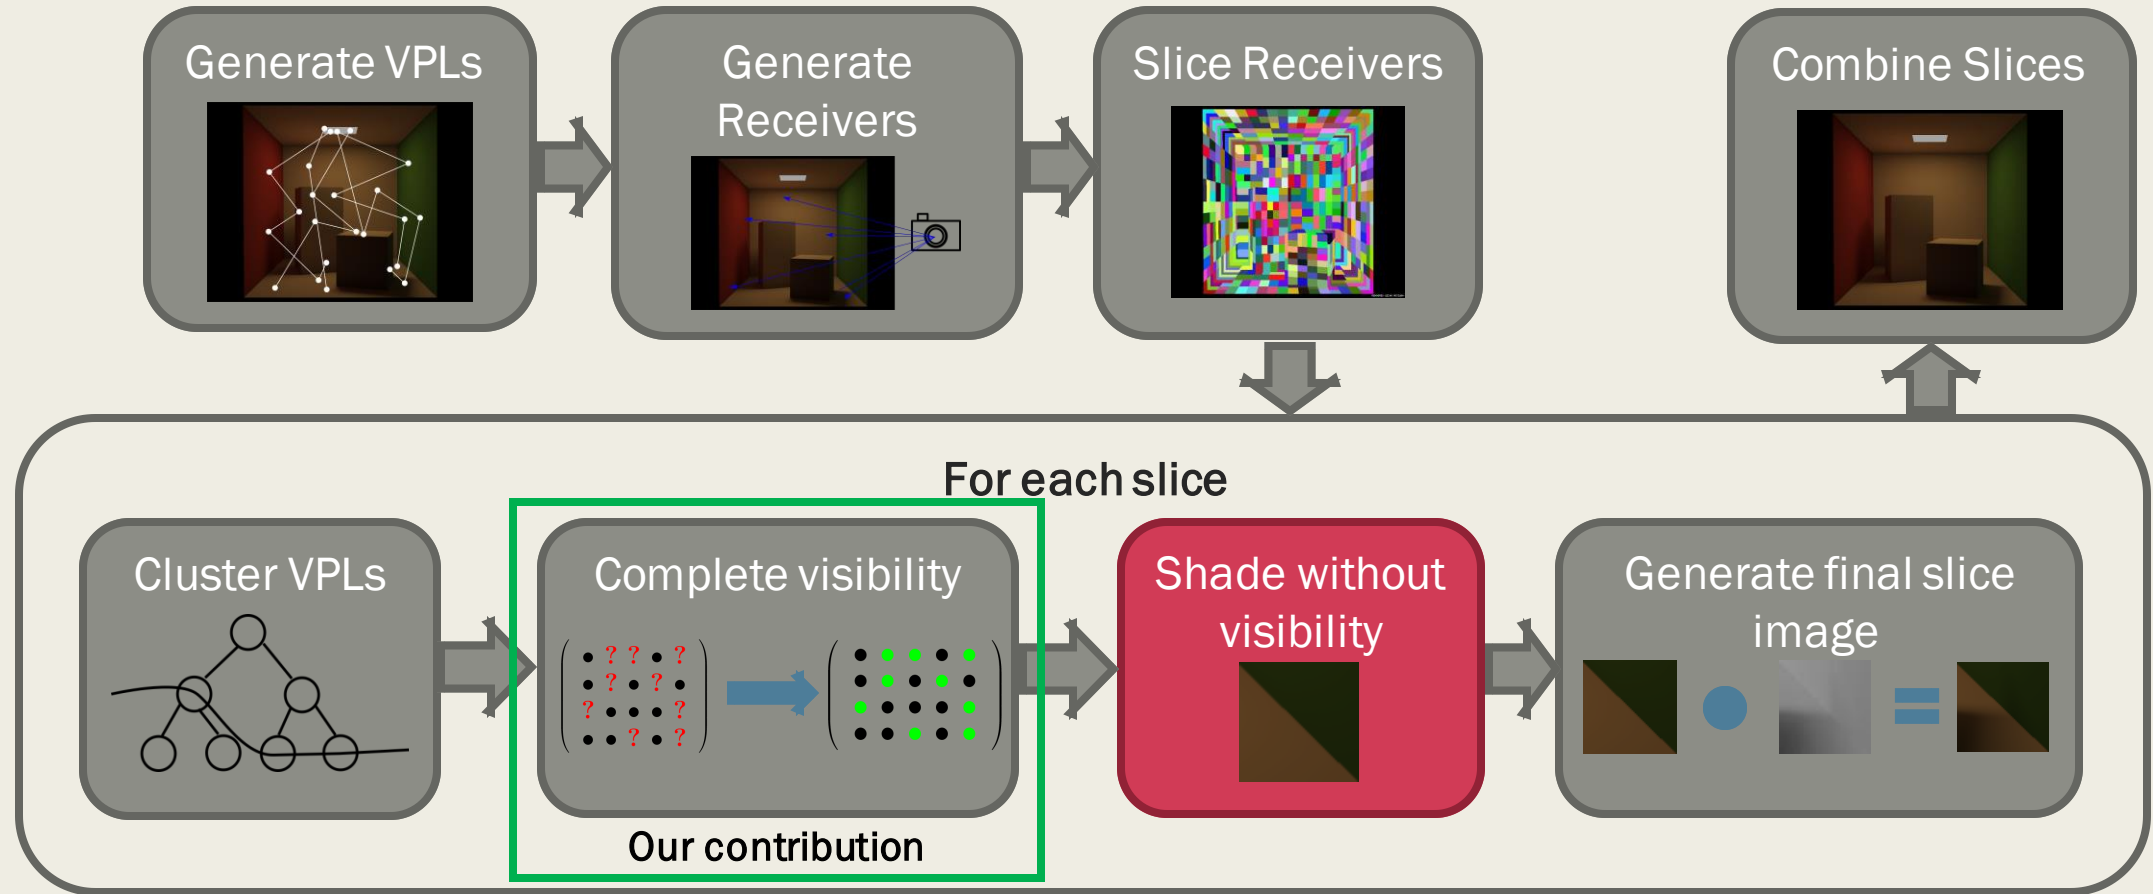
Naive Matrix Completion of the Many-Lights Matrix

- Limited to matrices with lower rank
 - Diffuse scenes with relatively simple occlusion
 - Participating media
- Difficult to determine rank and coherence for arbitrary scenes
- Materials and geometry influence rank and coherence
 - Glossier scenes & complex visibility = higher rank and coherence

Completing visibility

- Can separate the many-lights matrix as $M = S \bullet V$
 - S is shading matrix without visibility
 - V is a boolean visibility matrix
 - \bullet is element-wise multiplication
- Advantages
 - Shading easy to hardware accelerate
 - Fewer factors impacting rank and coherence of V
 - Allows for improvements to matrix completion algorithm

Our framework




Adaptive Matrix Completion

(Krishnamurthy A. & Singh A. '14)^[7]

- Iteratively completes sub-sampled columns
 - Using basis
 - If can't, full sample & expand basis
- Advantages
 - Relaxes row space coherence constraints
 - Operates on smaller basis -> faster

Adaptive Matrix Completion for Boolean visibility

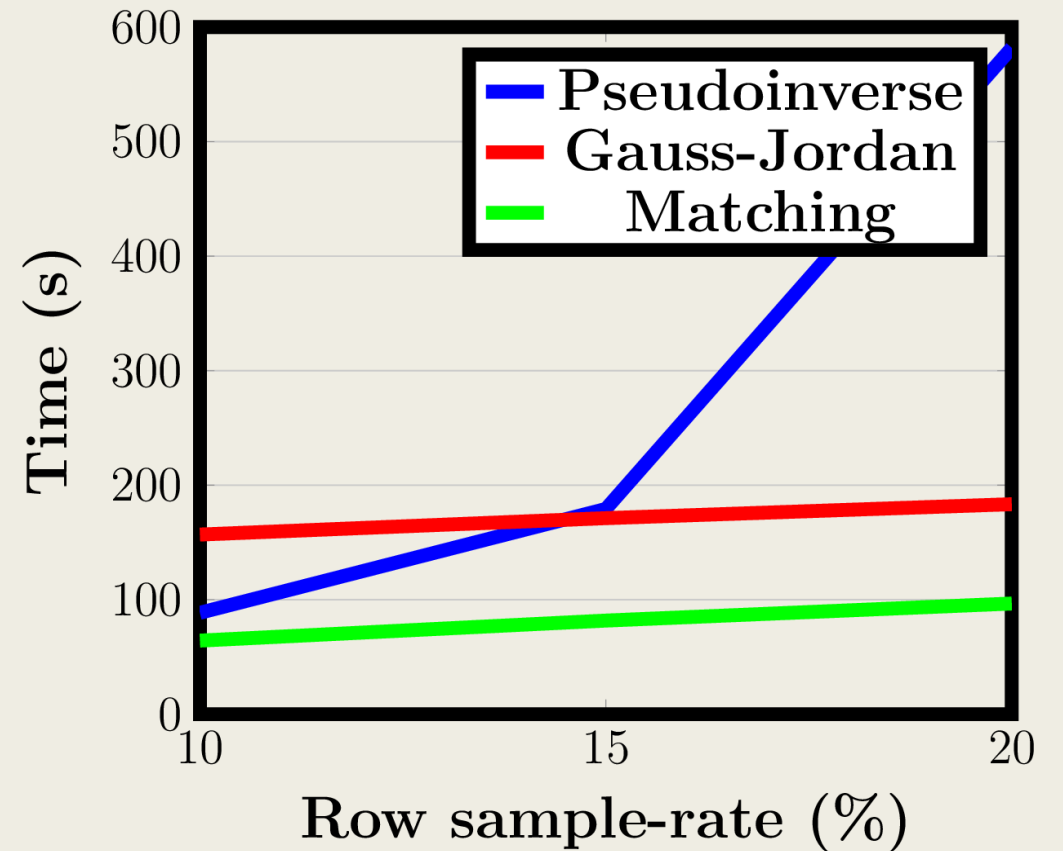
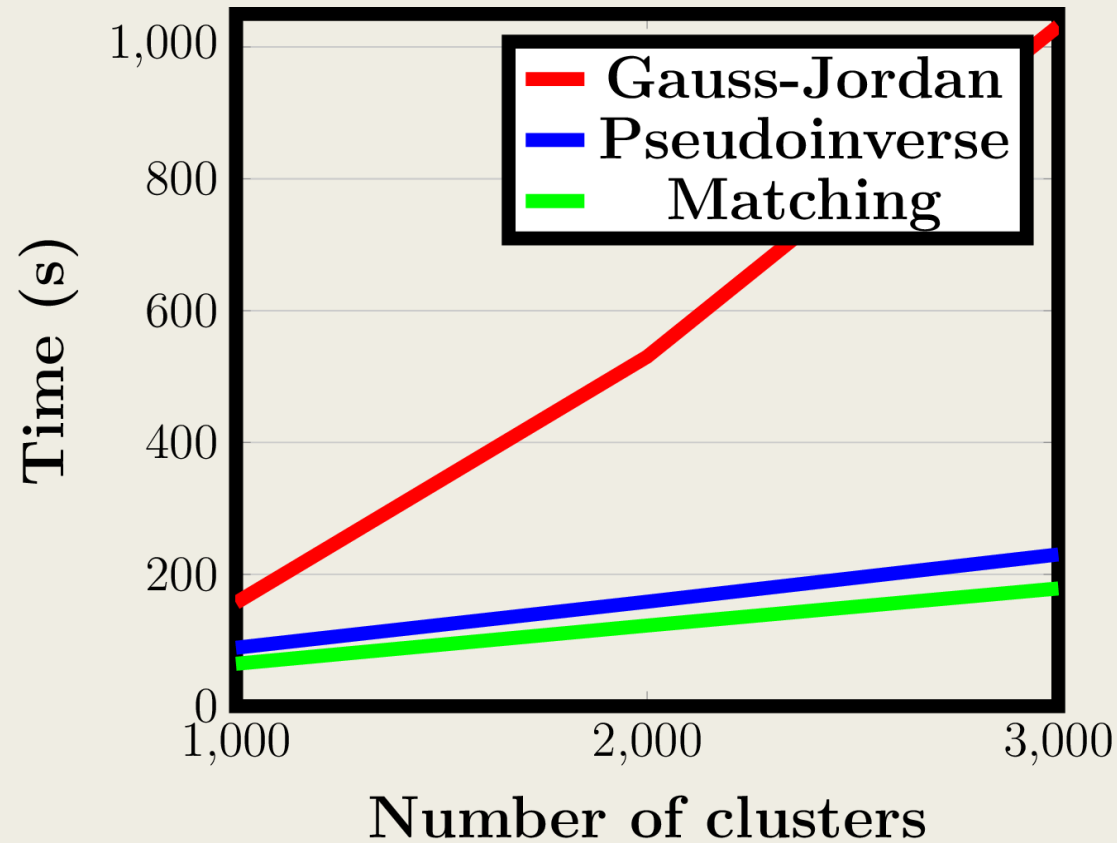


- Blue = low rank, red = high rank, green = high coherence
- Low-rank areas have near identical columns

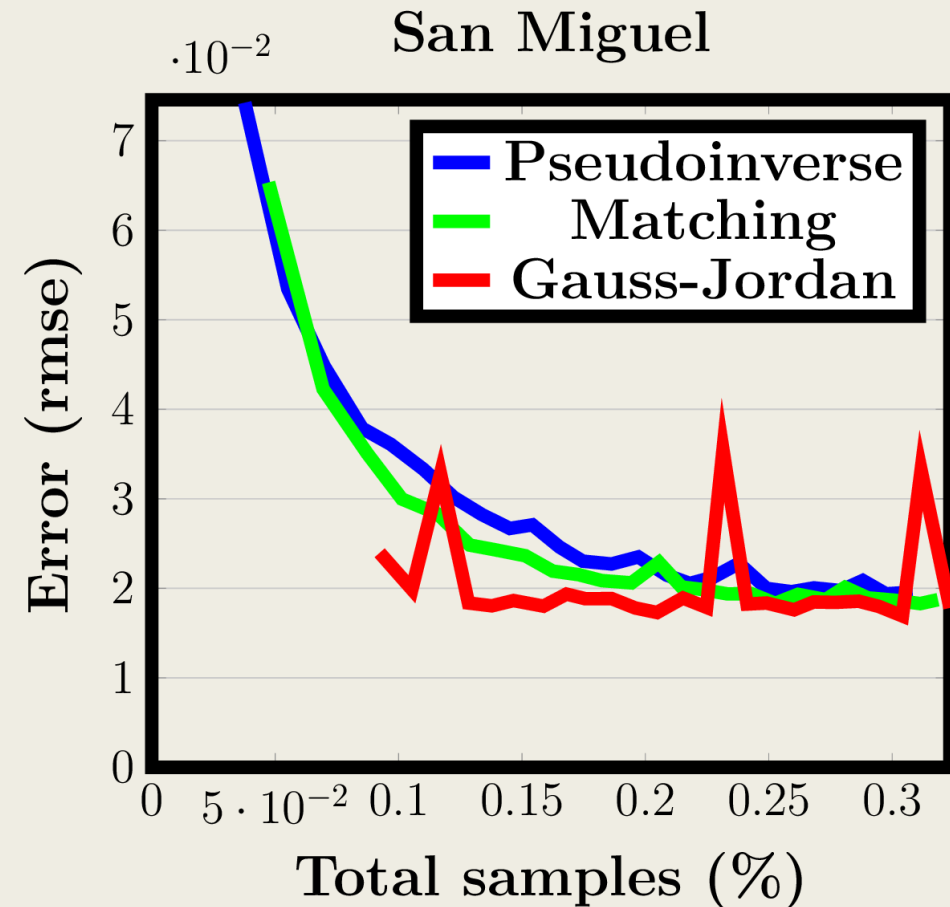
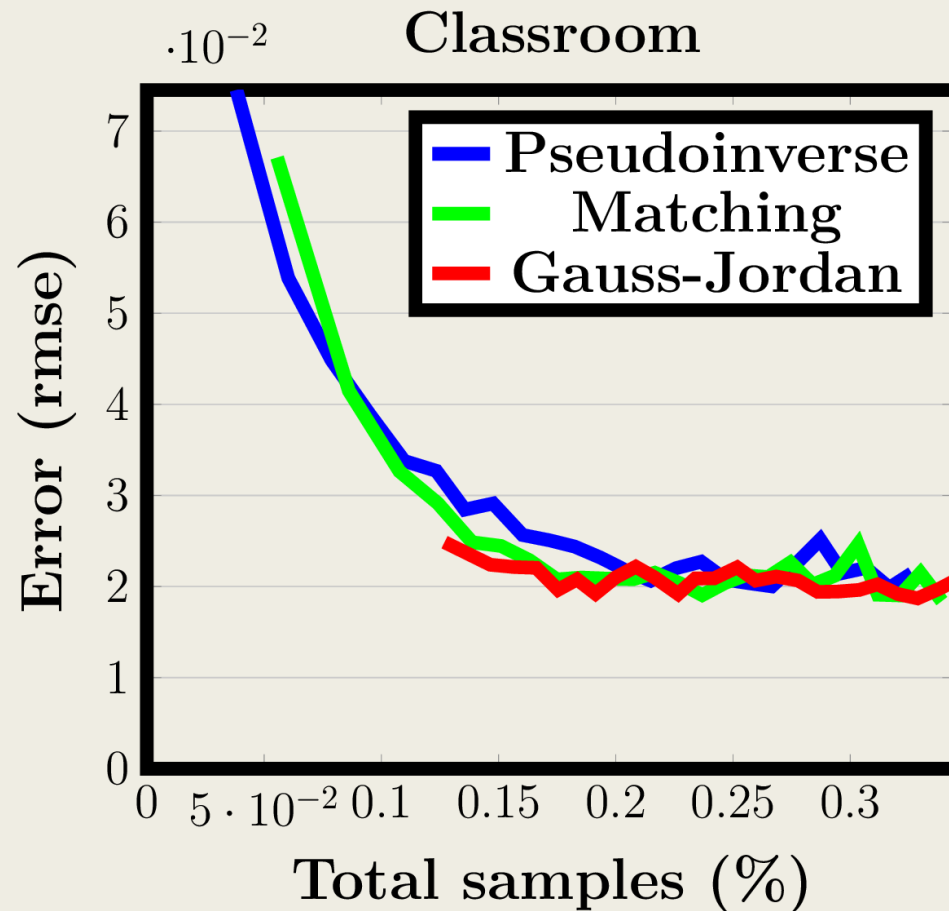
Adaptive Matrix Completion for Boolean visibility

- Can approximate column using a single basis vector
 - Replace pseudo-inverse with matching
 - $O(mn)$ rather than $O(mn^2)$
- Compared against pseudo-inverse of Q over \mathbb{R} and Gauss-Jordan elimination of matrix over $GF(2)$

Adaptive Matrix Completion for Boolean visibility



Adaptive Matrix Completion for Boolean visibility



Importance sampling

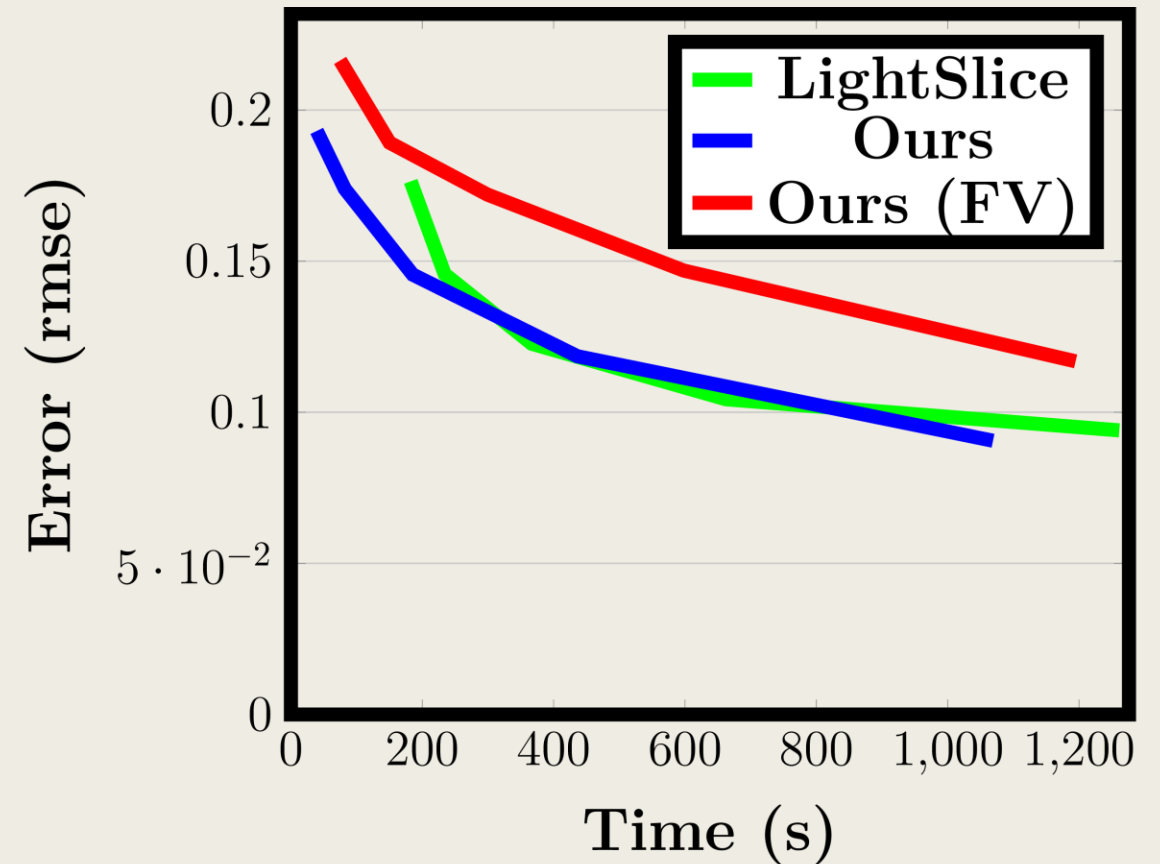
- More samples where visibility changes
- Use information from previous columns
- Maintain discrete distribution d
 - Update d when column fully sampled



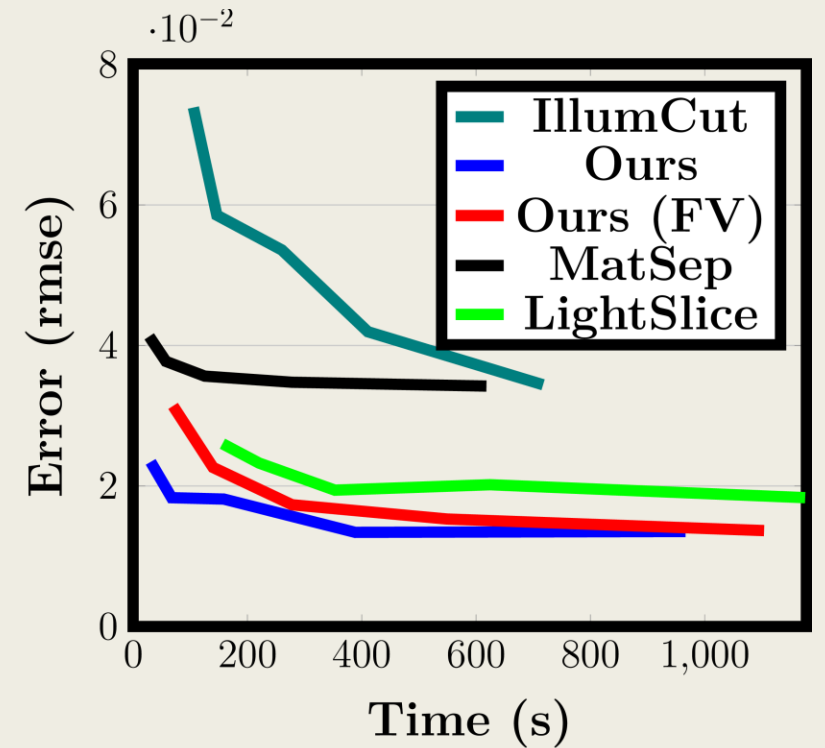
Results

- Over 12 scenes of varying complexity
 - VPLs in more diffuse scenes, VSLs in glossier scenes
- Ground truth obtained by brute force
- Compared to Low-rank and sparse separation (*Huo et al. '15*), IlluminationCut (*Bus et al. '15*), Lightslice (*Ou & Pellacini '11*), and our framework without matrix completion
 - only LightSlice and our method without matrix completion in VSL scenes

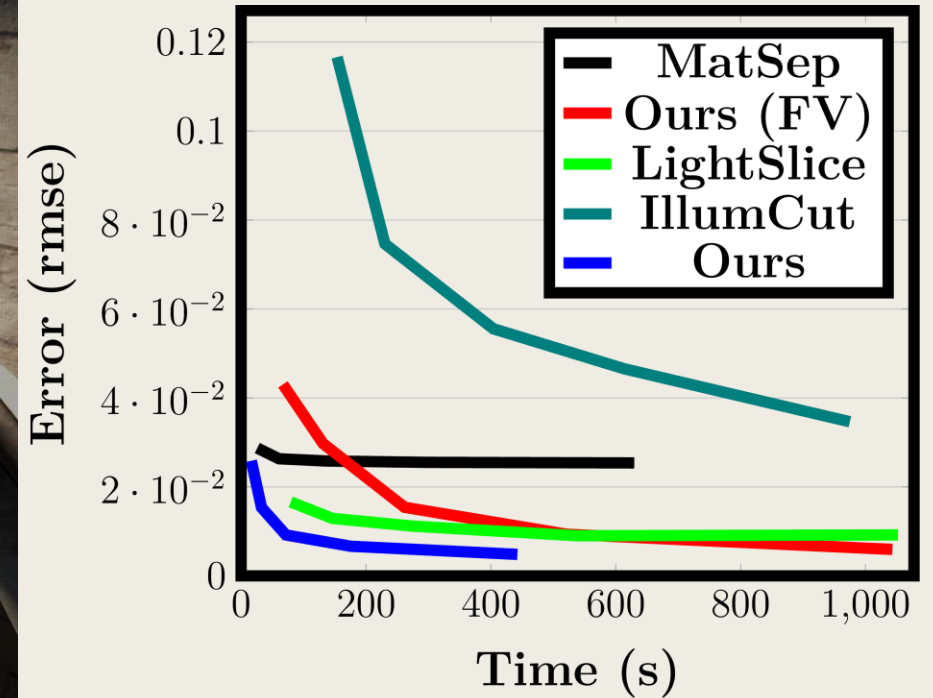
Results – Bathroom (185s, 23% samples)



Results– San Miguel (125s, 16.64% samples)




Results – Sponza (140s, 11.81%)




Results – San Miguel

-L1 equal time error images of Low rank and sparse recovery (*Huo et al. '15*) (first) versus our method (second)




Visibility matrix rank

- Rank is higher in penumbra
 - Receivers visible varies slightly per light
- Compared to Many-Lights matrix
 - Rank higher in diffuse scenes
 - Lower in glossier scenes



Drawbacks

- Inefficient column sampling



Drawbacks



- Clustering method does not consider visibility

Conclusions

- Method for speeding up Many-Lights renders
 - clustering & visibility approximation
- Approximates visibility with Adaptive Matrix Completion
- Improve Adaptive Matrix Completion for visibility
- Over 3 times faster for equal quality


Future Work

- Address to drawbacks
- Extend algorithm to area lights
- Integrating over other domains?
 - Time, Spectrum, Lens etc.
- Matrix completion in other contexts
 - Path-tracing, denoising etc.

Questions?



Adapting sample-rate



- Slices have different ranks
- Inefficient to use same sample-rate
- Use verification samples to adapt

