Optimal rank 1 matrix completion

Louis Theran (Aalto / AScI, ICS)
joint work with
Franz Király

(Low-rank) Matrix Completion

- From a *partial* matrix Ω(A), reconstruct the *true matrix* A

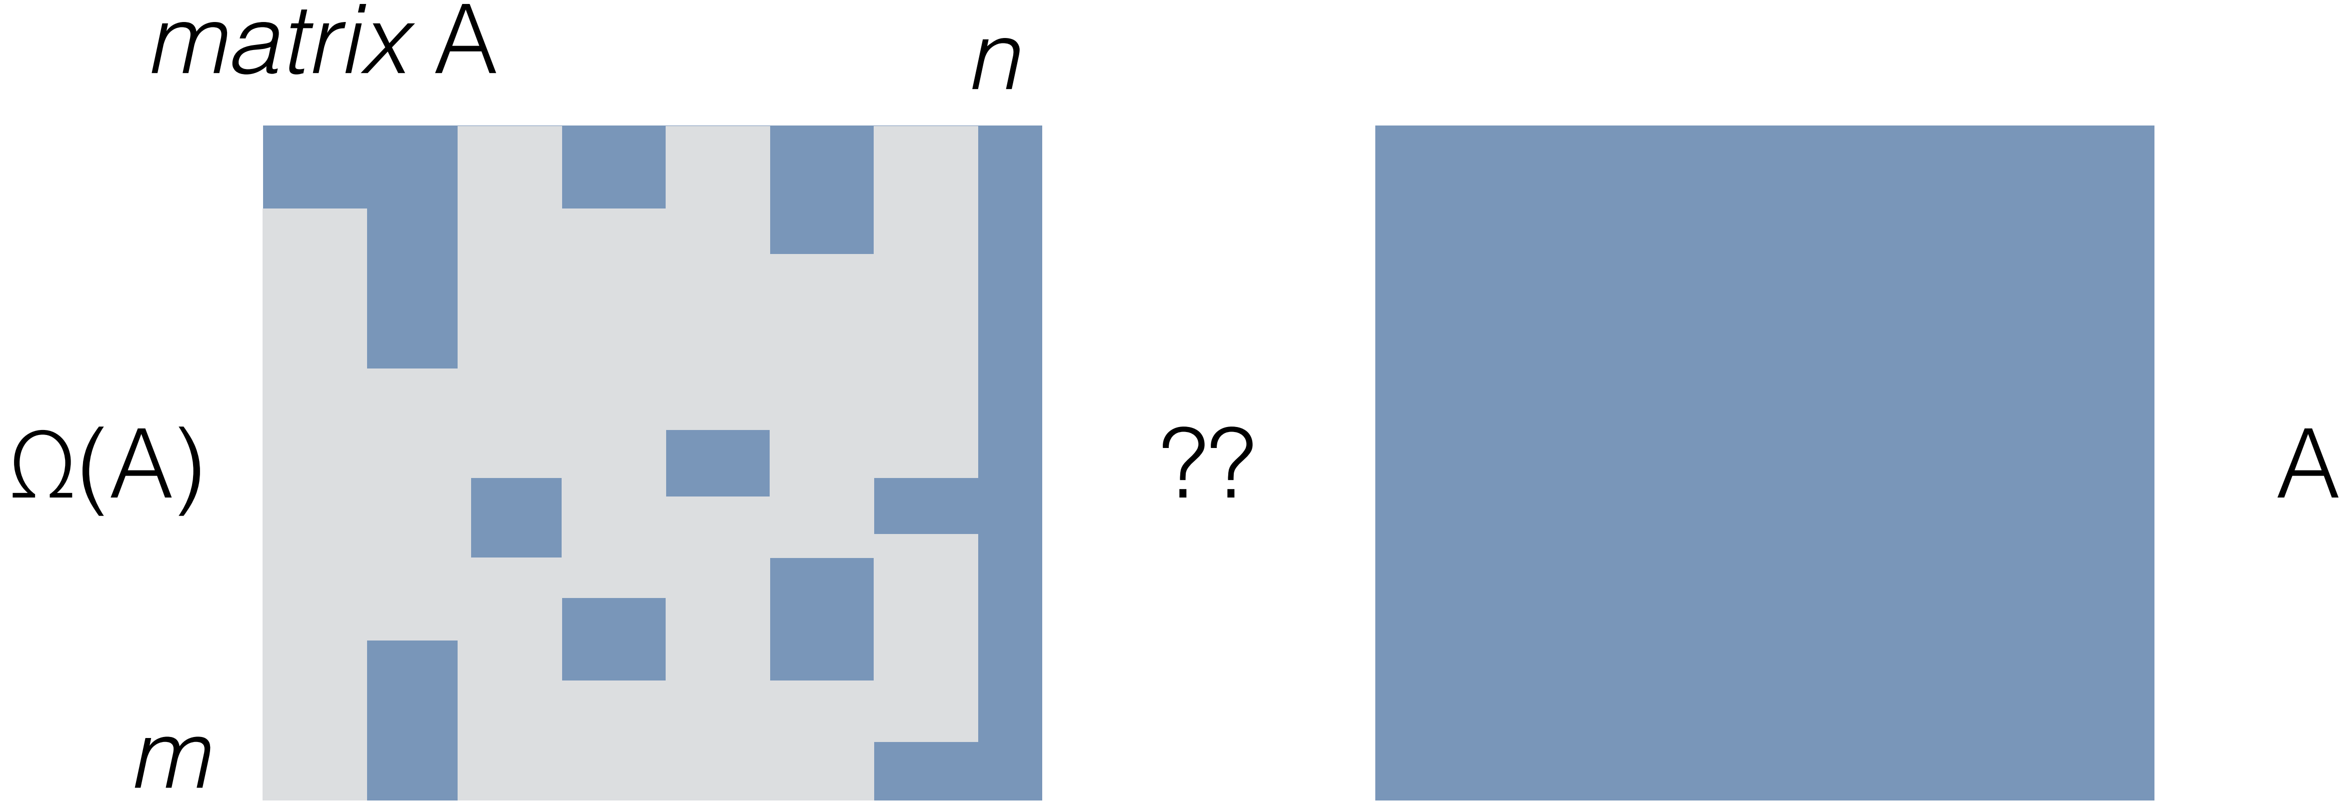

- Assuming that A has *known rank r*. ($r = o(m + n)$)

Connection to rigidity

- Intuitively, write $A = UV^T$. Task becomes finding U and V.
- Notions of rigidity matrices and stresses (Singer-Cucuringu '10)
- General (but worst-case impractical) scheme for completing entries using "circuit polynomials" (Király-Theran-Tomioka '15)
- Lots more (see talks for this workshop)

Things to think about

- Really want to complete at least one entry
    - Standard algorithm: convex relaxation of rank (not generic)
- There is going to be noise
    - Solutions to circuit polynomials don't "line up"

Rank 1 is easier

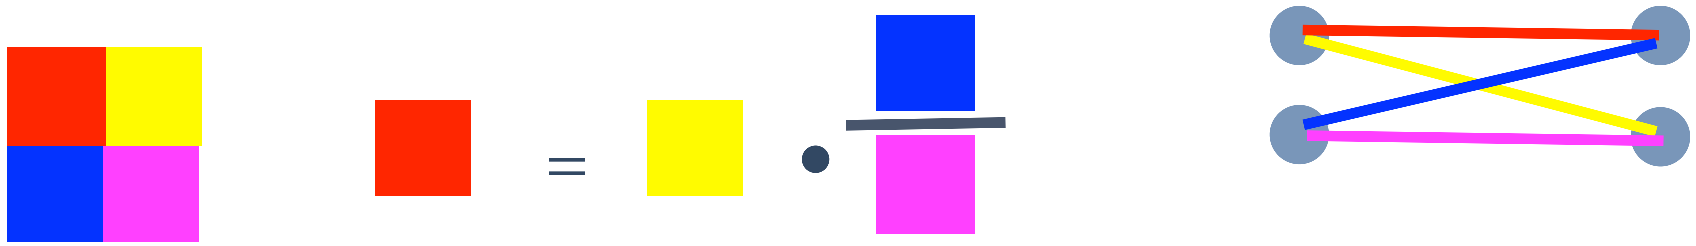

- Circuit polynomials are log-linear
- Can average them to produce *optimal single entry estimators* for log-centered noise
- Variance bounds depend on graph structure of the observation patterns (via graphic matroid)

[Király-T NIPS’13]

Pictures

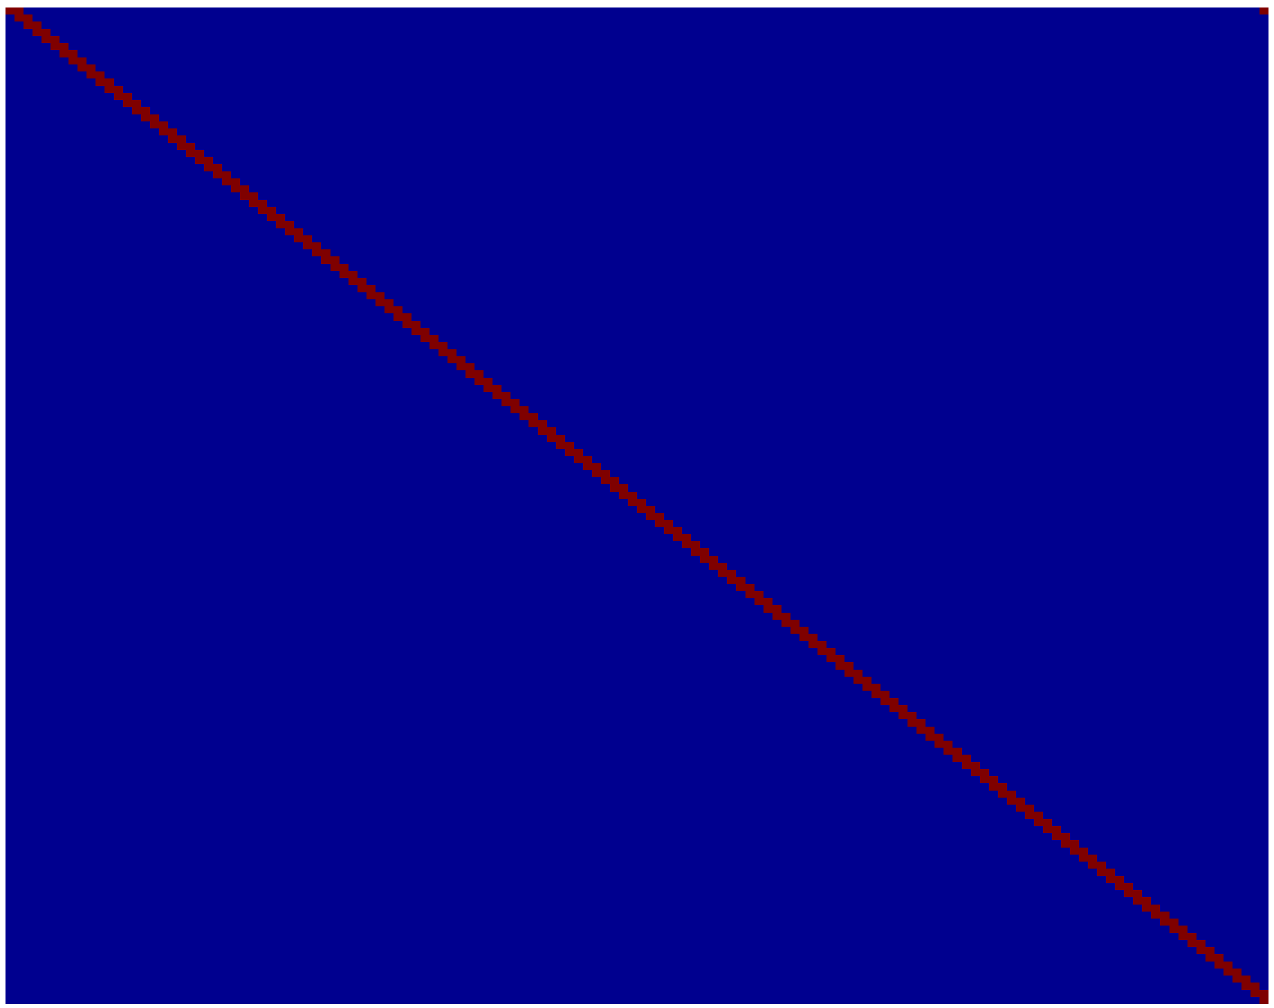

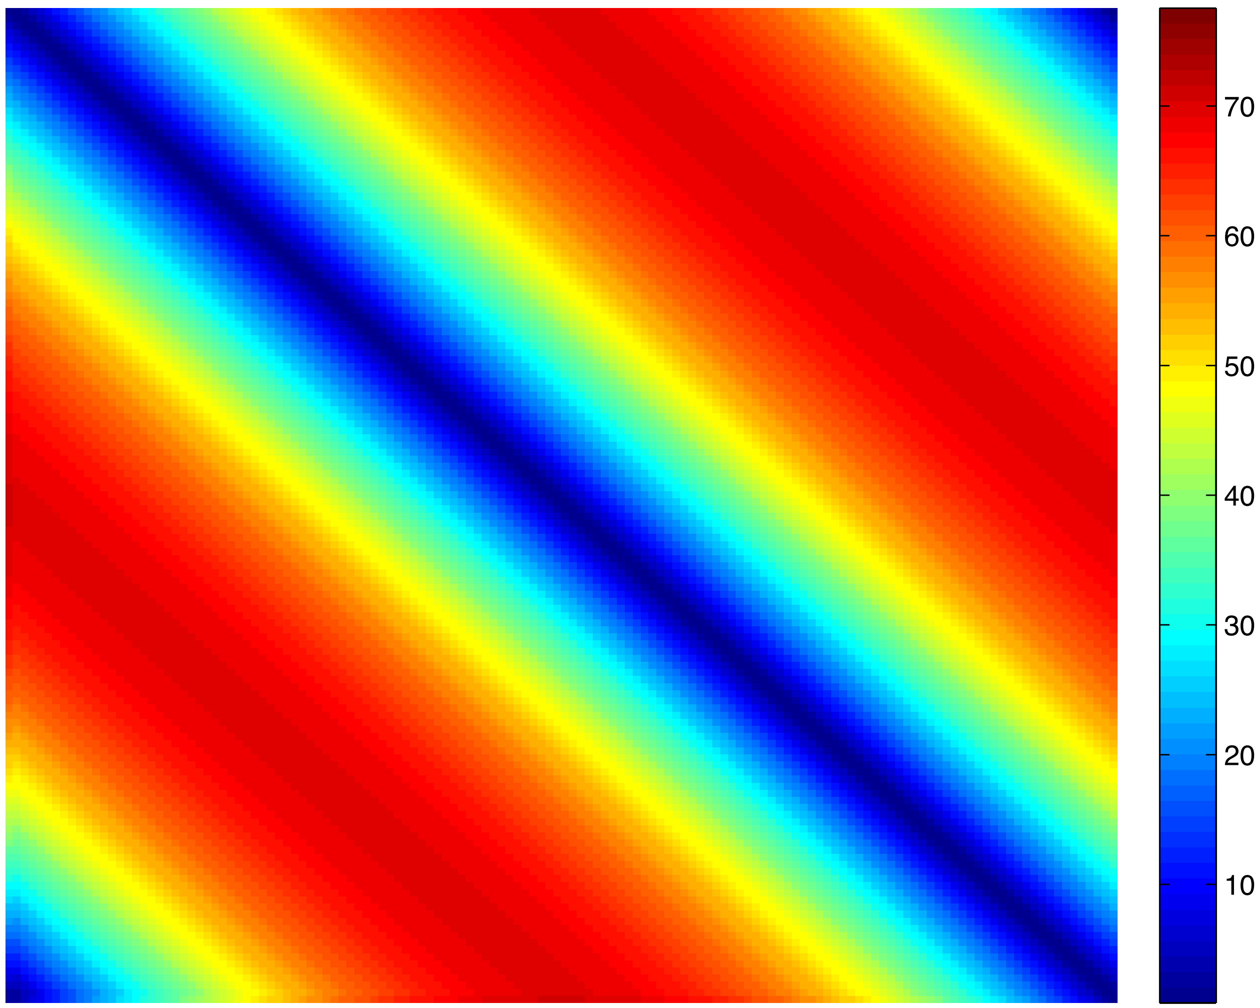

Observed entries

Optimal variance

Pictures

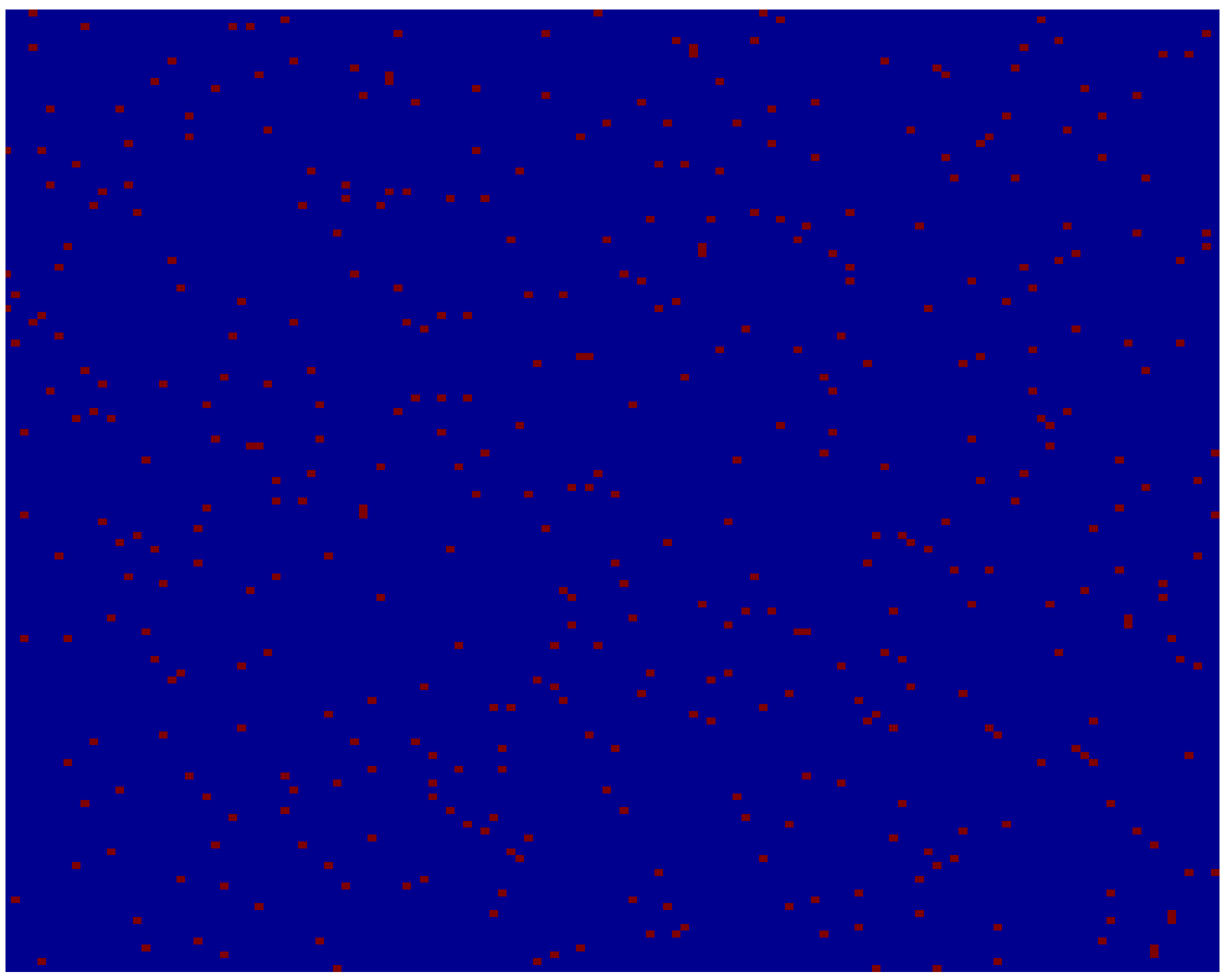

Observed entries

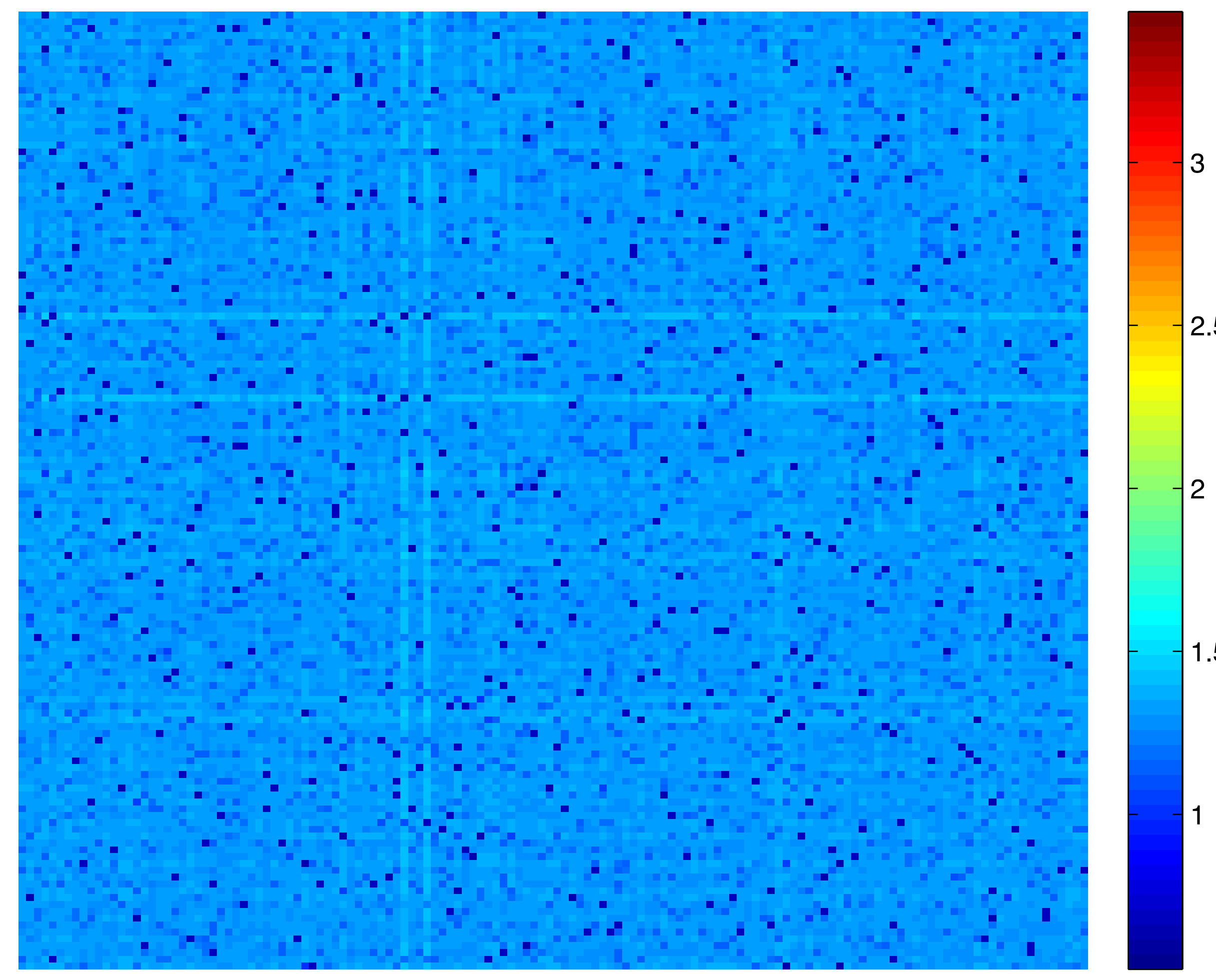

Optimal variance