General and Nested Wiberg Minimization

Dennis Strelow

General and nested Wiberg minimization

Sometimes we have to minimize a function of two sets of unknowns:

$$\min_{U,V} f(U, V)$$

Often these are chicken and egg problems.

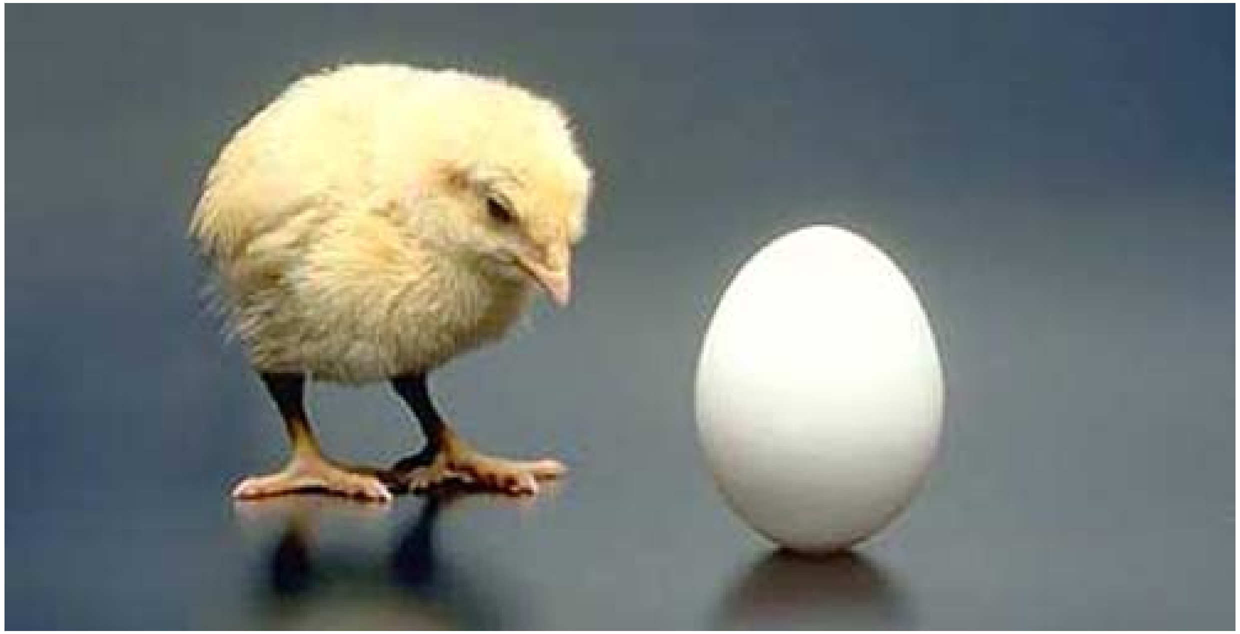

General and nested Wiberg minimization

In problems like these, we could alternately minimize wrt U, V, U, V, ...

- expectation-maximization
- alternating least squares
- alternating linear programming

General and nested Wiberg minimization

We could minimize wrt U, V simultaneously:

- Levenberg-Marquardt
- Newton-Raphson
- successive linear programming

General and nested Wiberg minimization

Or, we could eliminate *V* from the problem:

- given *U*, minimize *f* wrt *V*
- linearize the solution *V* wrt *U*
- minimize *f*(*U*, *V*(*U*)) wrt *U* only

(Wiberg 1976) proposed this idea for L_2 (least squares) matrix factorization when some matrix entries are not known

General and nested Wiberg minimization

Wiberg beats simultaneous and alternating methods?

- (Okatani and Deguchi 2007, Okatani *et al*. 2011) rediscover Wiberg factorization
  - convergence beats Levenberg-Marquardt
- (Eriksson and van den Hengel 2010) L_1-Wiberg factorization
  - more robust to outliers than L_2
  - convergence beats alternated quadratic programming
  - convergence beats successive linear programming

General and nested Wiberg minimization

But, so far Wiberg has only been matrix factorization

- linear in both *U* and *V*

This talk: general Wiberg minimization

- handles f nonlinear in both *U* and *V*
- focuses on the hardest case, L_1 minimization

General and nested Wiberg minimization

So, options for minimizing *f*(*U*, *V*):

	linear in U or V		
	minimize L_2	minimize L_1	MLE
simultaneous	✓	✓	✓
alternating	✓	✓	✓
Wiberg	Wiberg 1976	Eriksson 2010	this work

	nonlinear in both U and V		
	minimize L_2	minimize L_1	MLE
simultaneous	✓	✓	✓
alternating	✓	✓	✓
Wiberg	this work	this work	this work

Outline

→ Wiberg L_1 matrix factorization (Eriksson)

General Wiberg minimization (Strelow)

Nested Wiberg minimization (Strelow)

Wiberg L_1 matrix factorization

(Eriksson and van den Hengel 2010) extended Wiberg matrix factorization to L_1

1: Get the derivatives of a linear program solution wrt the linear program coefficients

2: Minimize the L_1 norm of a linear function using linear programming

- Get the solution's derivative using 1 + chain rule

3: Iteratively minimize $||Y - UV(U)||_1$ using successive linear programming, wrt U only

- To minimize wrt U only, use $V(U)$ and dV/dU from 2

Wiberg L_1 matrix factorization

Their method does factor the matrix, and can converge quadratically!

Wiberg L_1 matrix factorization

They demonstrated affine SFM using the factorization:

- 36 images, 319 points, 18 minutes to solve
- artifical outliers added to 10% of the image observations

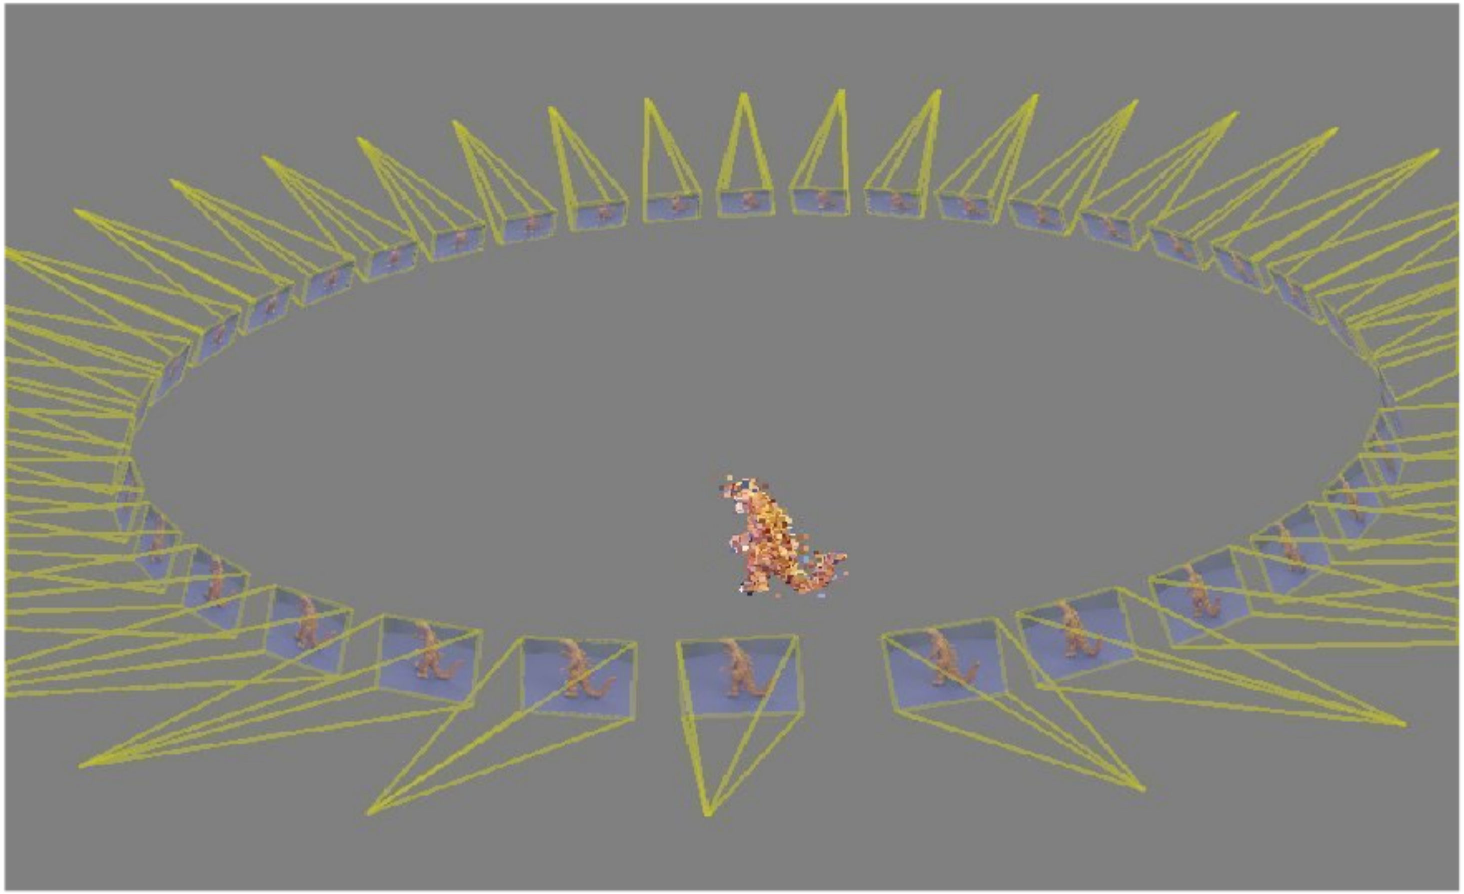

Wiberg L_1 matrix factorization

Their convergence beat alternated quadratic programming...

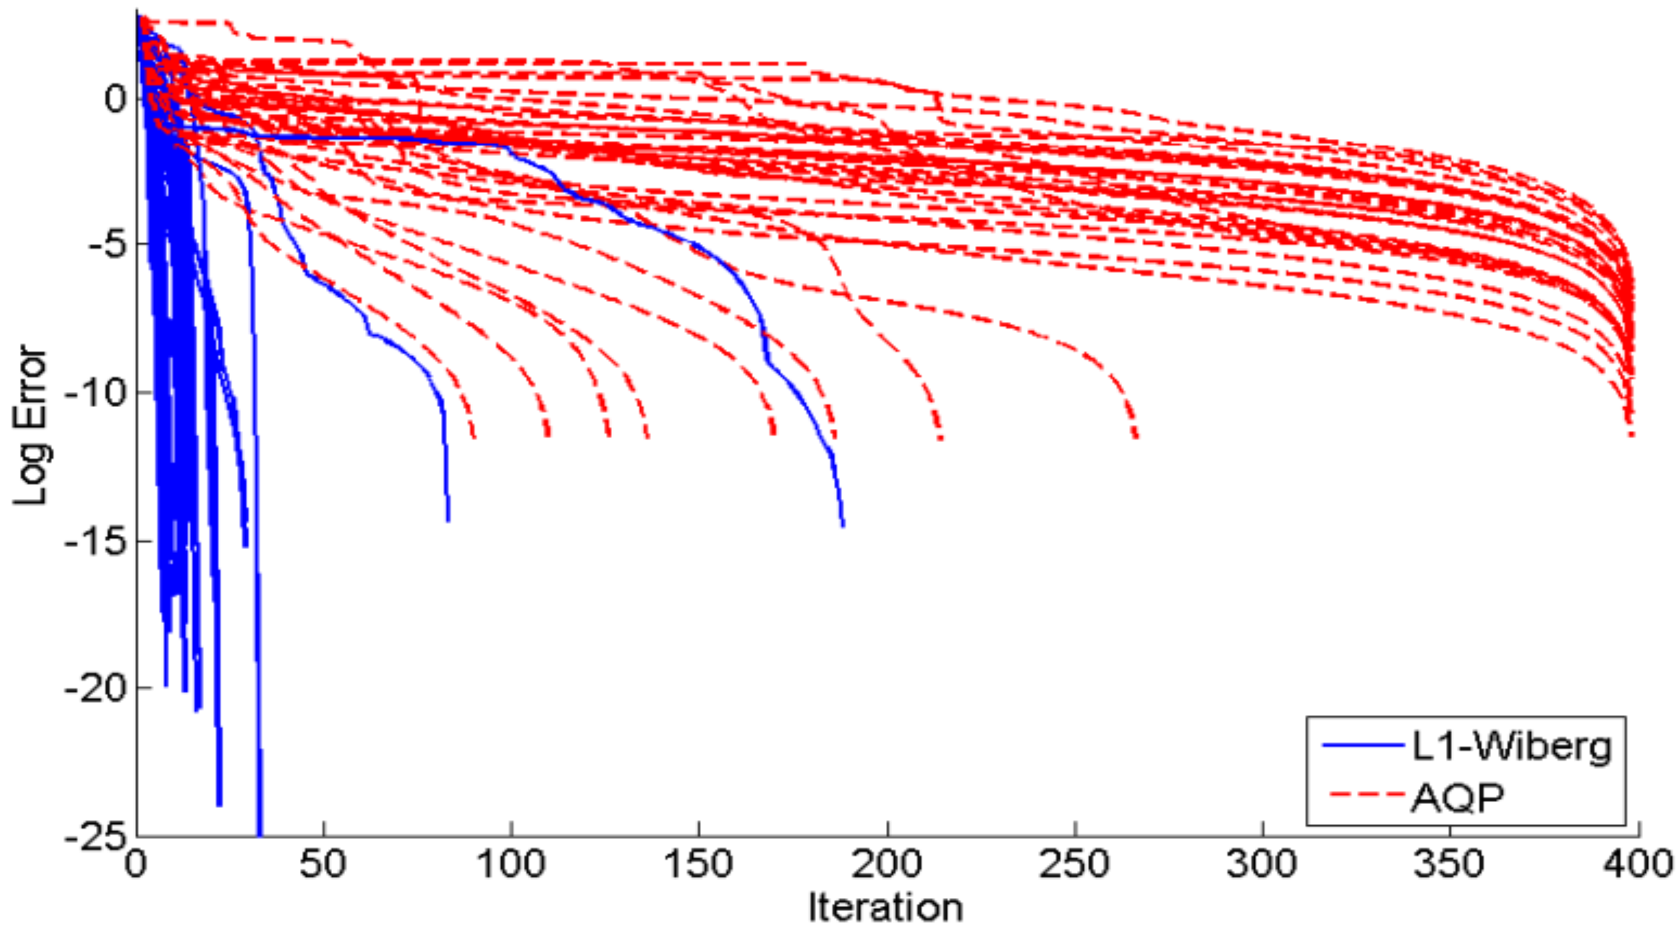

Wiberg L_1 matrix factorization

...and in our experiments, their method also beat successive linear programming:

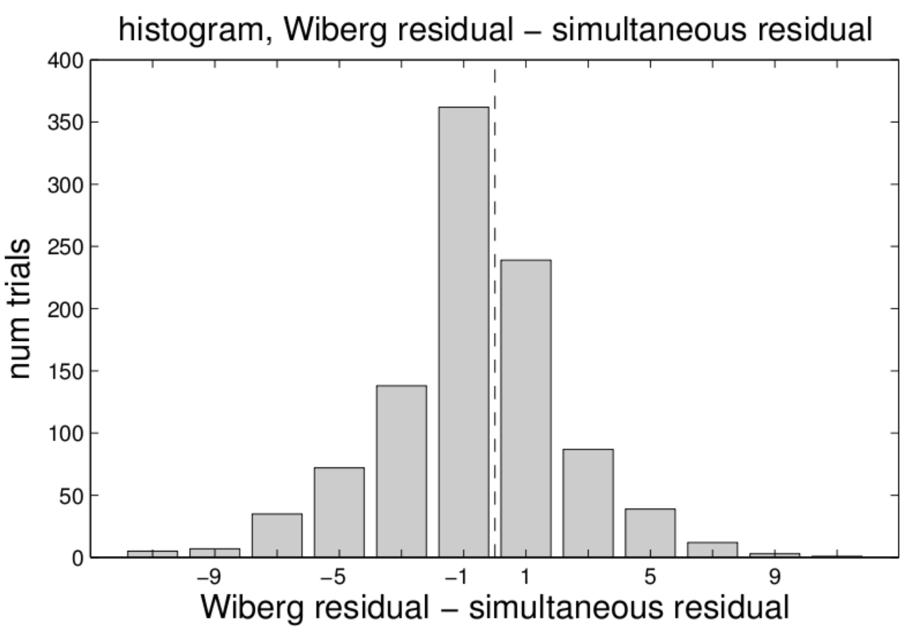

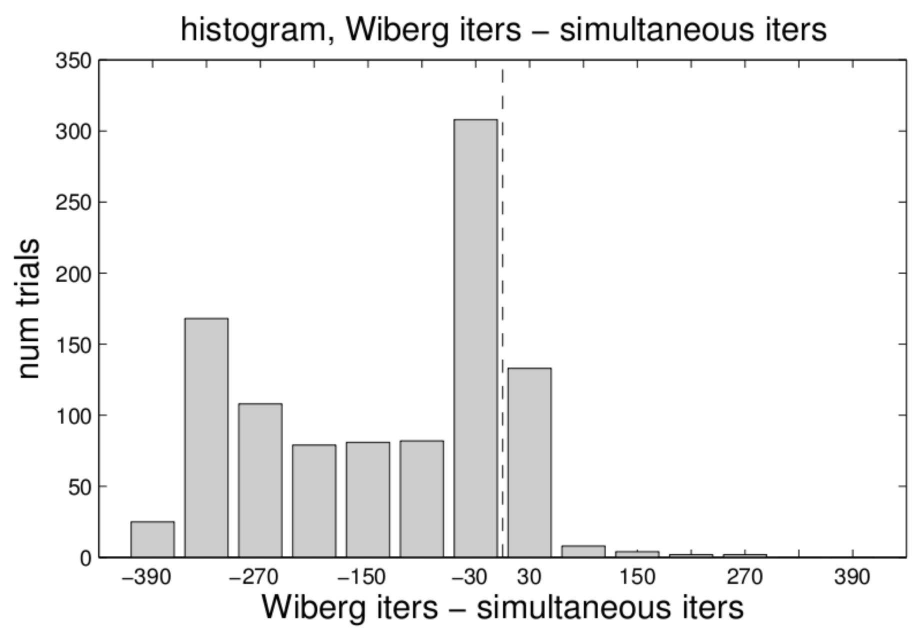

Wiberg L_1 matrix factorization

Is Wiberg faster than minimizing wrt *U*, *V* simultaneously?

- both methods add an unknown for every observation
  - In Wiberg, num remaining unknowns >> dim *V*
- Wiberg linear program is denser

So, successive linear programming iterations faster for > 20 rows in *U*

Outline

Wiberg L_1 matrix factorization (Eriksson)

→ General Wiberg minimization (Strelow)

Nested Wiberg minimization (Strelow)

General Wiberg minimization

Matrix factorization is linear in both *U* and *V*

Wiberg factorization:

- solves for *V* in closed form
- but, solves for *U* iteratively

So, we can easily tweak the method to minimize general nonlinear functions of *U*

General Wiberg minimization

But, *f* would still have to be linear in *V*.




General Wiberg minimization

To handle *f* nonlinear in both *U* and *V*:

- Find *V* iteratively
- But then what is *dV/dU*?
    - *V* is found iteratively, but each step *ΔV* is found in closed form
    - take *dV/dU = dΔV/dU* for the last *ΔV*

Like Wiberg matrix factorization, general Wiberg minimization works and can converge quadratically!

General Wiberg minimization

Example: Wiberg L_1 bundle adjustment

Real example, "rover"

- about 700 images, 10K 3-D points

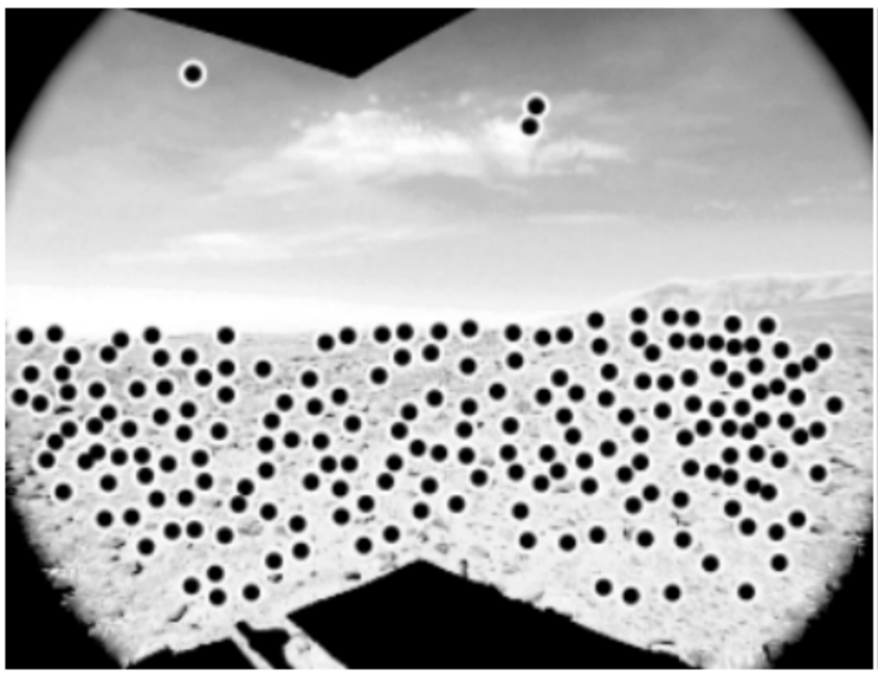

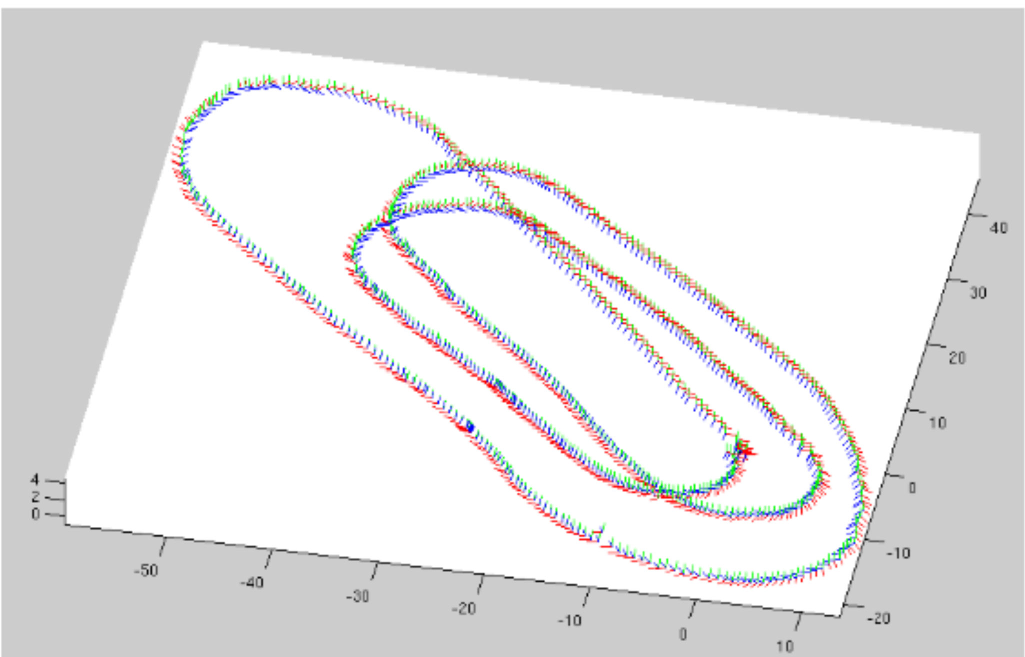

General Wiberg minimization

Two bundle adjustment experiments

"Bundle Adjustment 1": low initial residuals

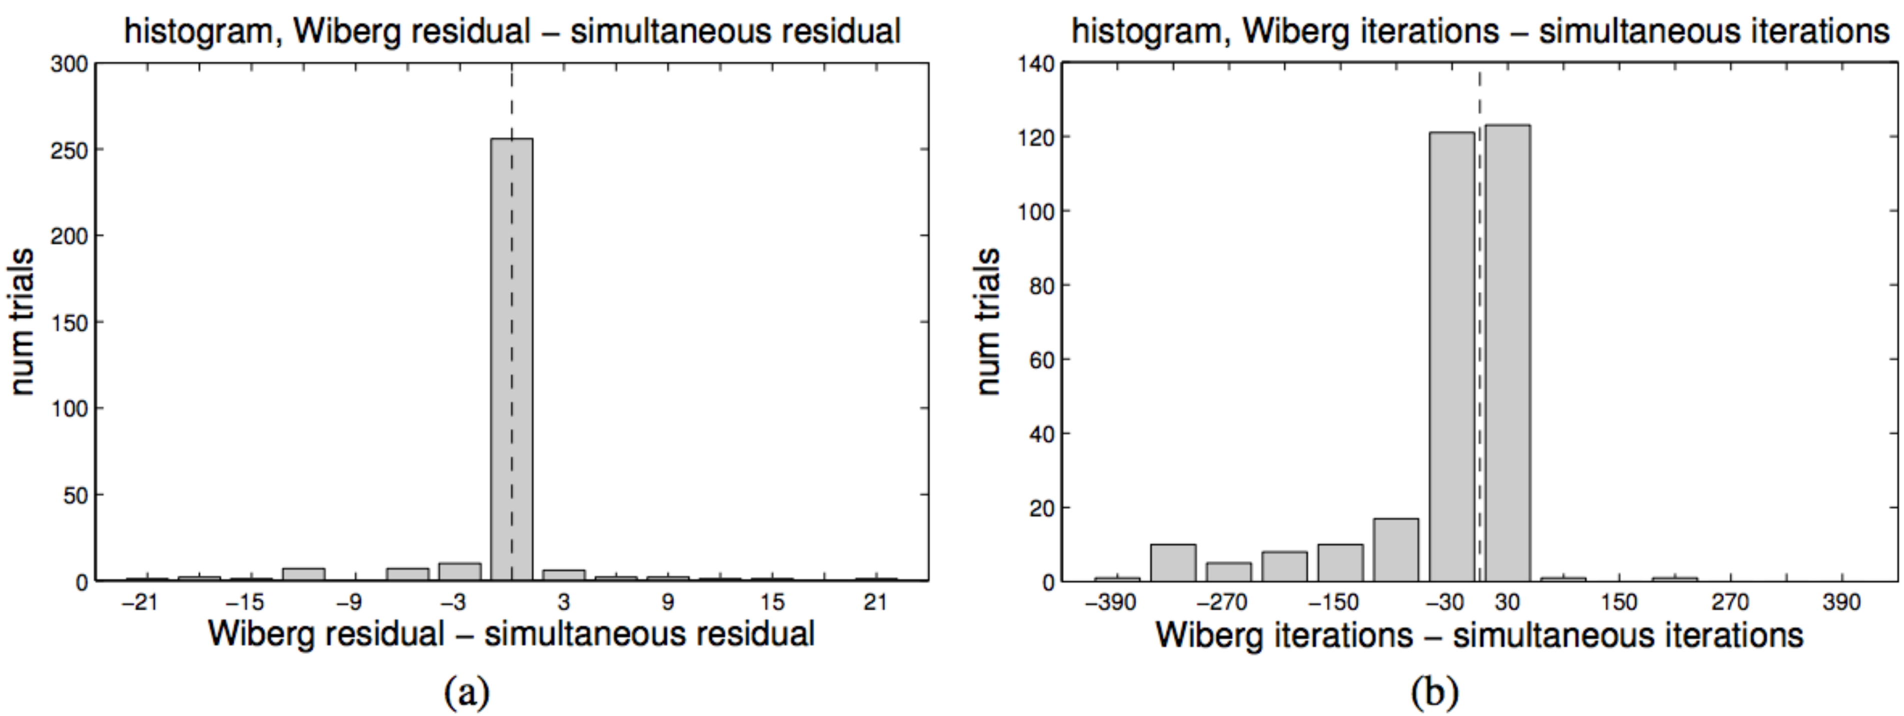

General Wiberg minimization

"Bundle adjustment 2": high initial residuals

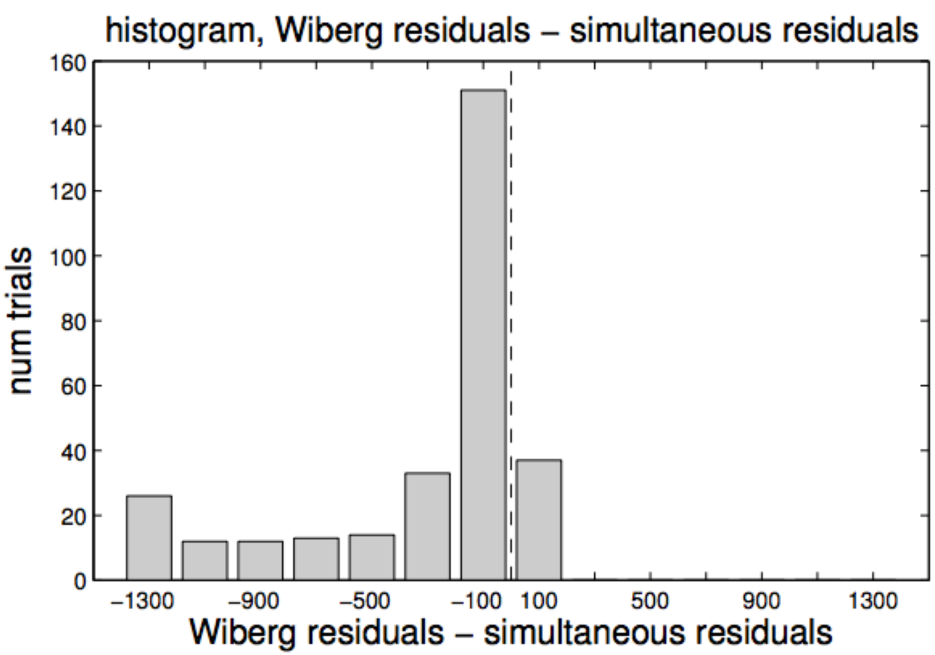

(c)

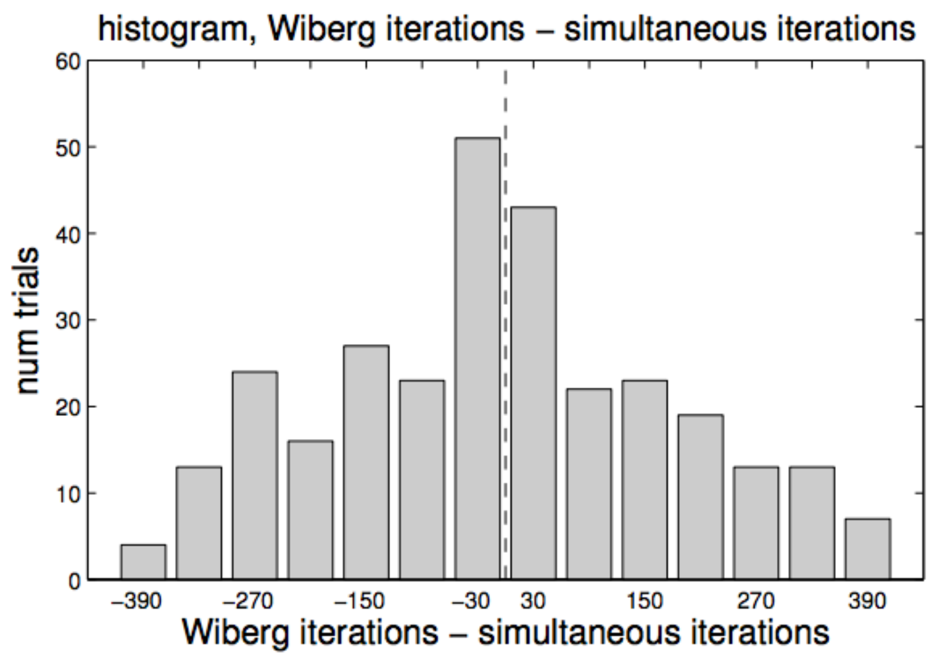

(d)

Speed: simultaneous iterations faster than Wiberg for > 5 images

Outline

Wiberg L_1 matrix factorization (Eriksson)

General Wiberg minimization (Strelow)

→ Nested Wiberg minimization (Strelow)

Nested Wiberg minimization

If we have three unknowns...

$$\min_{U,V,D} f(U, V, D)$$

...can general Wiberg be applied recursively to eliminate two?

$$\min_{U} f(U, V(U), D(V(U)))$$

Nested Wiberg minimization

Yes...nested Wiberg does work and can also converge quadratically!

Nested Wiberg minimization

The derivatives become exponentially more complicated.

Nested Wiberg minimization

Wiberg projective bundle adjustment

- uncalibrated camera
- additional set of unknowns: projective depths

Assignments:

- U: points
- V: cameras
- D: projective depths

Nested Wiberg minimization

Two projective bundle adjustment experiments

"Projective Bundle Adjustment 1": low initial residuals

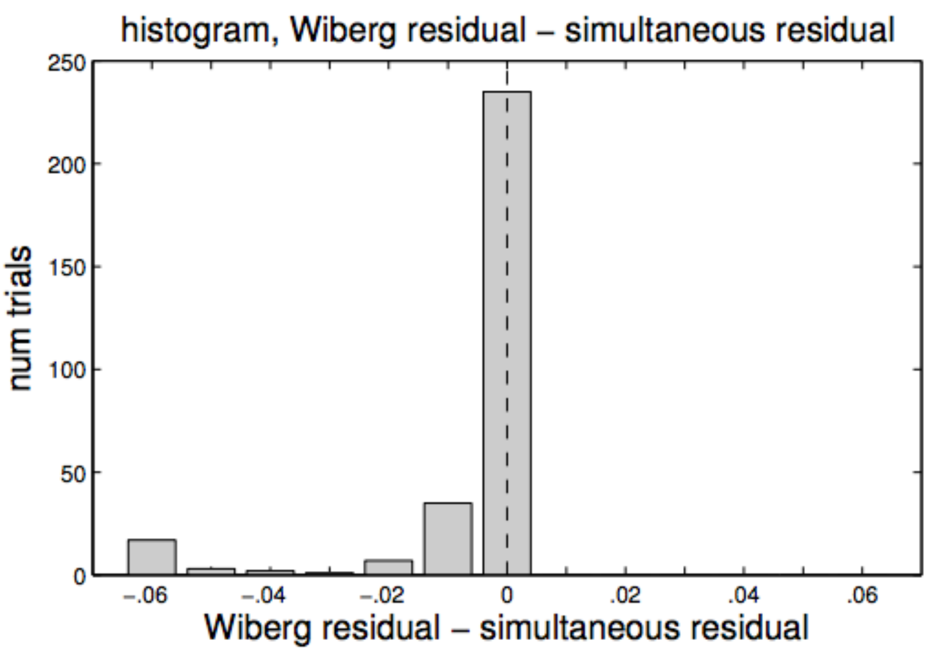

(a)

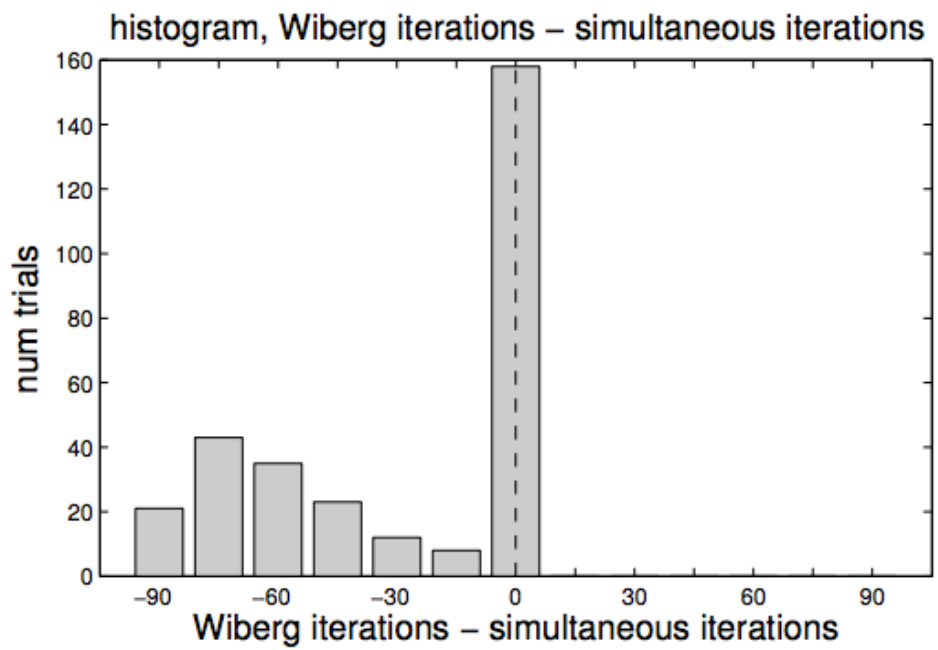

(b)

Nested Wiberg minimization

"Projective Bundle Adjustment 2": high initial residuals

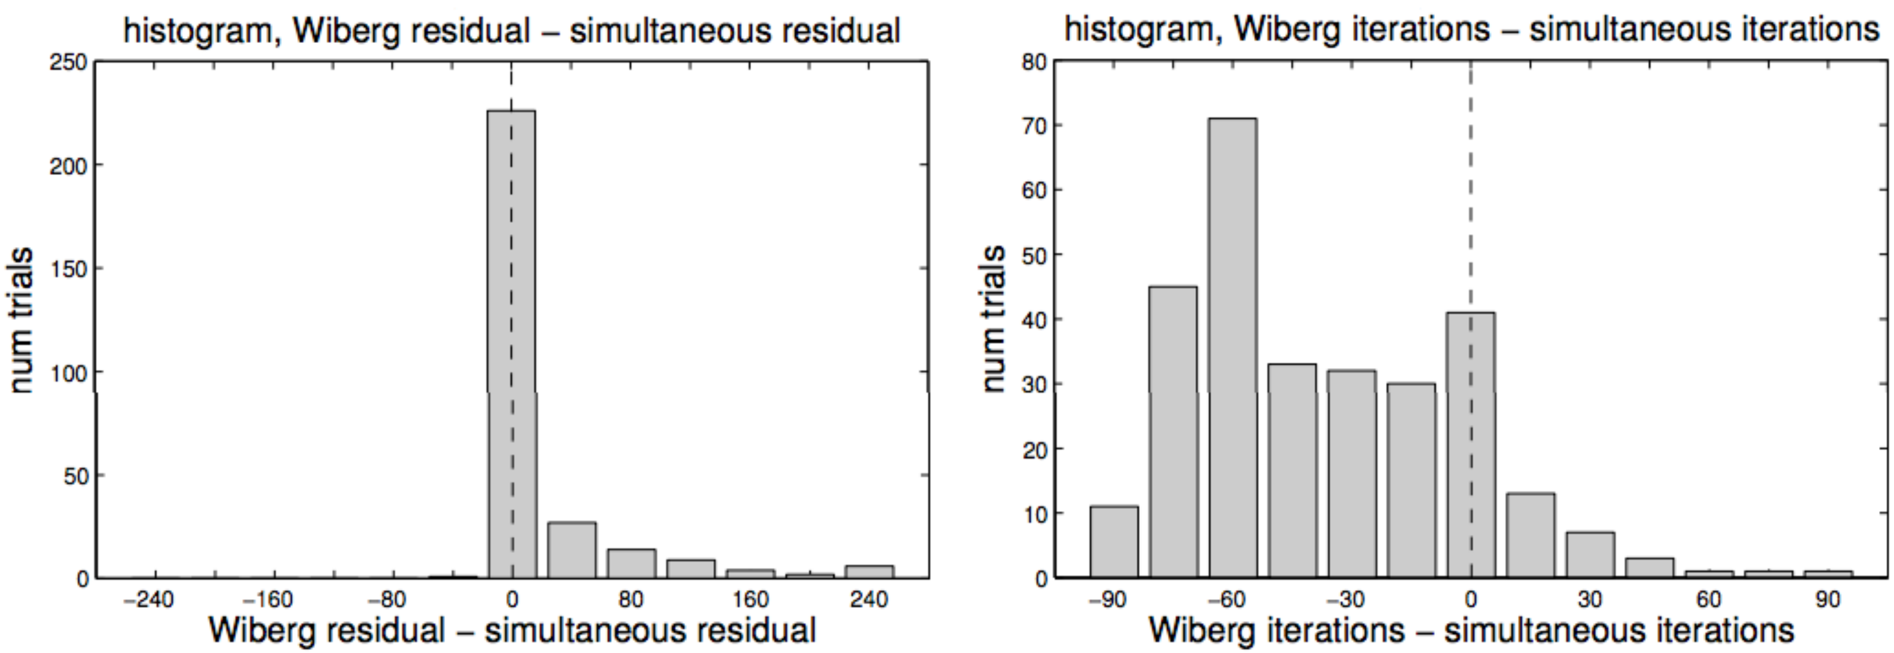

(c) (d)

Speed: Wiberg faster for < 51 points

Work since

Wiberg L_2 and maximum likelihood estimation

- Wiberg L_2 bundle adjustment
- Wiberg L_2 projective bundle adjustment
- Wiberg Poisson matrix factorization

Thanks!

- Emilie Danna
  - linear programming advice → success on real example
- Jay Yagnik, Mei Han, Luca Bertelli, Vivek Kwatra, Rich Gossweiler, Mohamed Eldawy
  - feedback on paper, talk
- Jim Teza, David Wettergreen, Chris Urmson, Mike Wagner
  - captured rover sequence
- anonymous CVPR reviewers
  - L_1 versus L_2 bundle adjustment