

Adaptive Annealing for Robust Geometric Estimation

THU-PM-126, June 22, 2023



C Sidhartha



L Manam



VM Govindu

COMPUTER VISION LAB
DEPARTMENT OF ELECTRICAL ENGINEERING
INDIAN INSTITUTE OF SCIENCE
BENGALURU

M-Estimation

- Minimization of robust cost:

$$\boldsymbol{\theta}^* := \operatorname{argmin}_{\boldsymbol{\theta}} \left\{ f(\boldsymbol{\theta}) = \sum_{i=1}^m \rho_{\sigma}(\|r_i(\boldsymbol{\theta})\|) \right\} \quad (1)$$

- Robust loss $\rho_{\sigma}(\cdot)$ is parametrized by σ .
- Solved using Iteratively Reweighted least squares (IRLS).

$$w_i := w(\|r_i(\boldsymbol{\theta}_k)\|) = \frac{\rho'(\|r_i(\boldsymbol{\theta}_k)\|)}{\|r_i(\boldsymbol{\theta}_k)\|}$$

$$g(\boldsymbol{\theta}) = \sum_{i=1}^m w_i \|r_i(\boldsymbol{\theta})\|^2$$

$$\boldsymbol{\theta}_{k+1} = \operatorname{argmin}_{\boldsymbol{\theta}} g(\boldsymbol{\theta})$$

Graduated Non-Convexity (GNC)

- IRLS is sensitive to initialization \Rightarrow convergence to poor local minimum.
- GNC alleviates this problem.
- Minimize a sequence of costs, $\rho_{\sigma}(\cdot)$, $\sigma = \sigma_0 > \sigma_1 > \sigma_2 > \dots > \sigma_{min}$.
- $\rho_{\sigma_k}(\cdot)$ is easier to solve than $\rho_{\sigma_{k+1}}(\cdot)$.
- $\gamma_k = \frac{\sigma_{k+1}}{\sigma_k}$, generally γ_k is chosen to be a constant $\forall k$.




Figure: Robust cost for varying σ

Flowchart for GNC

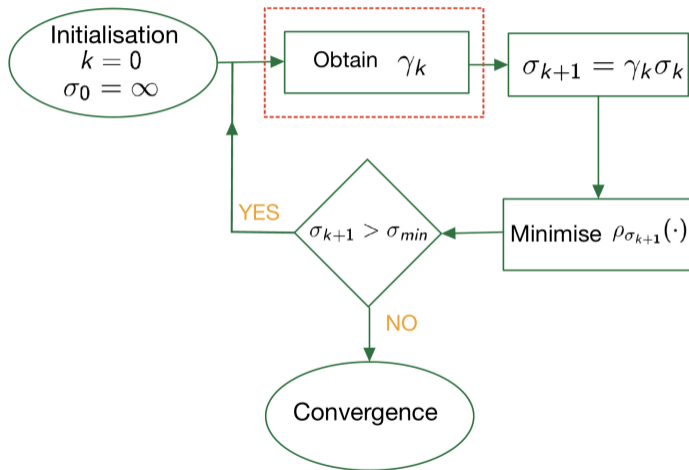


Figure: GNC Flowchart

Motivation: Effect of γ on GNC




Figure: Cost landscape for varying σ .




Figure: Line Fitting example

Motivation: Effect of γ on GNC

GNC strategy	# of stages in GNC	Accuracy
Small γ (fixed)	↑	↑
Large γ (fixed)	↓	↓
Adaptive γ [Ours]	↓	↑
Desirable	↓	↑

Table: Impact of different annealing strategies. Our adaptive approach achieves high accuracy with fewer annealing stages.

DESIRABLE: Best accuracy with least number of GNC stages.

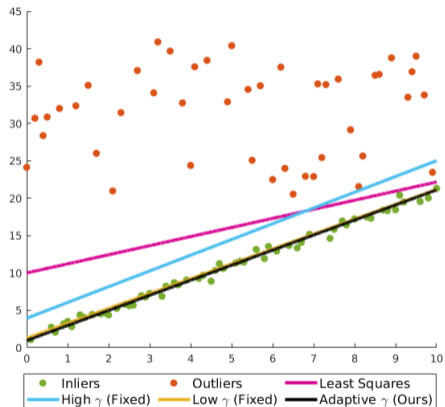



Figure: Line Fitting example

Adaptive strategy to find $\sigma_{k+1} (\equiv \gamma_k)$



- Ensure the same basin of the minimum throughout the GNC schedule starting from the global minimum at $\sigma = \infty$ till $\sigma = \sigma_{min}$.
- Achieved by tracking the positive (semi)definiteness of the Hessian matrix at the minimum throughout the annealing schedule.

Adaptive annealing ctd..

- Given the current solution x_k^* , choose σ_{k+1} as:

$$\sigma_{k+1} = \min_{\sigma \leq \sigma_k} \left\{ \sigma \mid \lambda_{\min} \left(H_{x_k^*}(\sigma) \right) > \lambda_T > 0 \right\} \quad (2)$$

- Given the least squares cost's gradient (g_{LSQ}) and Hessian (H_{LSQ}),

$$H(\sigma) = \sum_{i=1}^n \left(-l_i \frac{g_{LSQ,i} g_{LSQ,i}^T}{\|r_i\|^2} + m_i H_{LSQ,i} \right), \quad (3)$$

$$\text{where } l_i = \frac{\rho'(\|r_i\|)}{\|r_i\|} - \rho''(\|r_i\|), \quad m_i = \frac{\rho'(\|r_i\|)}{\|r_i\|} \quad (4)$$


- Computation time of $\lambda_{\min}(H)$ depends only on d , where $H \in \mathbb{R}^{d \times d}$,
irrespective of the problem type.

Efficient computation of σ_{k+1}


- Practically, $\lambda_{\min} \left(H_{x_k^*}(\sigma) \right)$ is an increasing function of σ .
- Perform binary search in σ space to obtain $\sigma_{k+1} \Rightarrow$ multiple computations of $H(\sigma) \Rightarrow$ expensive when the no. of samples n is high.
- σ enters the Hessian only through the terms l_i and m_i .
- Piecewise polynomial approximation to l_i and m_i to obtain the form of $H(\sigma) = \frac{C}{\sigma^2} + D$, where C and D are constant in piecewise constant intervals of l_i and m_i .
- Reduces the number of computations of C and D during the binary search making it **efficient**.

3D Registration as an Illustration

Given correspondences $\{a_i, b_i\}$ between two unaligned 3D scans, **find the transformation R, t that aligns the scans.**



Source: TEASER [6]
Figure: Before registration



Source: TEASER [6]
Figure: After registration

Optimization problem:

$$\min_{(R,t)} \sum_{i=1}^N \rho_{\sigma}(\|a_i - Rb_i - t\|) \quad (5)$$

Details

- The IRLS subproblem has the following form:

$$\min_{(R,t)} w_i \|a_i - Rb_i - t\|^2$$

- This admits a closed form solution derived as a variant of *Umeyama's* [4].

Results: Synthetic Data

Dataset	FGR	SE3Reg	TEASER++	GNCp	FGR	SE3Reg	TEASER++	GNCp
	[8]	[1]	[6]	(Ours)	[8]	[1]	[6]	(Ours)
	Mean Rotation Errors (deg) ↓				Mean Translation Errors ($\times 10^{-3}$) ↓			
armadillo	0.89	1.12	13.87	0.79	16.68	48.06	106.59	9.27
bunny	0.93	1.27	13.63	0.81	14.53	46.76	152.33	9.18
buddha	1.22	1.32	22.32	1.02	18.97	45.57	148.82	11.09
dragon	0.88	1.05	13.05	0.76	15.89	47.21	119.91	9.08

Table: Mean rotation and translation errors on synthetic datasets ($N = 10000$, high noise level) for 50% outliers.

Dataset	Success % ↑	MRE ↓	MTE ($\times 10^{-2}$) ↓	Time (ms)
FGR [8]	93.1	2.92	2.96	10
GORE [2]	93.3	2.87	2.86	2
SE3Reg [1]	93.0	2.60	2.59	3
TEASER++ [6]	94.2	2.40	2.28	7
GNCp (Ours)	94.1	2.26	2.26	8

Table: Evaluation on ModelNet dataset [5]. MRE: Mean Rotation Error (in degrees), MTE: Mean Translation Error (in metres). Our method has the least mean errors.

Results: Real Data

Dataset	Success % \uparrow				Time (in ms.) \downarrow			
	FGR	SE3Reg	TEASER++	GNCp	FGR	SE3Reg	TEASER++	GNCp
<i>3D Match</i>	[8]	[1]	[6]	(Ours)	[8]	[1]	[6]	(Ours)
MIT lab	72.7	75.3	71.4	77.9	71.4	21.6	8216	7.4
home1	93.6	92.9	92.9	96.1	54.6	14.6	3964	5.6
home2	79.3	78.8	78.8	81.7	47.1	12.2	5555	6.0
hotel1	93.8	93.8	94.7	95.1	54.1	16	4416	6.4
hotel2	88.5	89.4	86.5	91.3	52.4	15.1	3969	7.6
hotel3	85.2	87.0	85.2	88.9	56.6	15.1	5849	6.3
kitchen	95.3	92.7	96.0	96.6	45.1	15.4	1978	6.0
study	79.8	82.5	86.0	84.6	55.2	16.2	3195	8.9
<i>KITTI</i>	73.5	84.7	-	85.6	195	40	-	13

Table: Results on 3D Match dataset [7], KITTI [3] datasets.

Dataset	Mean Rotation Errors (deg) \downarrow				Mean Translation Errors (m) \downarrow			
	FGR	SE3Reg	TEASER++	GNCp	FGR	SE3Reg	TEASER++	GNCp
<i>3D Match</i>	[8]	[1]	[6]	(Ours)	[8]	[1]	[6]	(Ours)
MIT lab	13.46	12.48	14.64	9.18	0.42	0.44	0.63	0.33
home1	5.91	6.21	8.74	4.21	0.19	0.19	0.29	0.15
home2	20.48	19.56	16.24	20.46	0.38	0.32	0.38	0.38
hotel1	6.95	7.13	7.14	6.63	0.18	0.18	0.18	0.18
hotel2	14.93	14.48	14.64	15.34	0.33	0.32	0.45	0.36
hotel3	23.38	20.98	13.43	20.81	0.46	0.40	0.38	0.40
kitchen	4.91	5.48	4.60	4.29	0.12	0.14	0.14	0.11
study	16.05	12.97	15.17	10.88	0.51	0.43	0.57	0.35
<i>KITTI</i>	0.94	0.88	-	0.74	0.38	0.32	-	0.29



(a) SE3Reg



(b) TEASER++



(c) GNCp (Ours)



(d) Ground Truth

Figure: Two point clouds (red and green) with low overlap.

Comparison with fixed γ annealing

Dataset	Small γ (Fixed)		Large γ (Fixed)		Adaptive γ (Ours)	
	S%	#Stages	S%	#Stages	S%	#Stages
MIT lab	89.3	14	80.4	7	89.3	6.5
home1	94.8	14	84.3	7	94.8	6.1
home2	99.3	14	95.8	7	99.3	5.4
hotel1	97.9	14	94.7	7	97.9	5.8
hotel2	100	14	98.8	7	100	5.9
hotel3	95.7	14	87	7	95.7	6.7
kitchen	97.7	14	94	7	97.7	5.4
study	90.9	14	75.5	7	90.9	7.9

Table: Comparison of different annealing schemes for instances with $< 50\%$ outliers. S% refers to the percentage of instances reaching global minimum.

Our method: Accuracy of 'small γ ' annealing + Number of stages of 'large γ ' annealing.





Conclusion

- Using parametric loss functions in a GNC framework mitigates convergence to poor local minima caused due to IRLS.
- Fixed factor (γ) annealing has a poor accuracy-speed tradeoff.
- Proposed adaptive annealing (GNC) approach by tracking the positive definiteness (*i.e.*, local convexity) of the Hessian.
- State-of-the-art results on 3D Registration are demonstrated using adaptive GNC.

Acknowledgments:

Chitturi Sidhartha and Lalit Manam are supported by the Prime Minister's Research Fellowship, Government of India. This research was supported in part by a Core Research Grant from the Department of Science and Technology, Government of India.

References I

-  Bhattacharya, U., Govindu, V.M.: Efficient and robust registration on the 3d special euclidean group. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5885–5894 (2019)
-  Bustos, A.P., Chin, T.J.: Guaranteed outlier removal for point cloud registration with correspondences. IEEE transactions on pattern analysis and machine intelligence **40**(12), 2868–2882 (2017)
-  Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3354–3361. IEEE (2012)
-  Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. IEEE Transactions on Pattern Analysis & Machine Intelligence **13**(04), 376–380 (1991)

References II

-  Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1912–1920 (2015)
-  Yang, H., Shi, J., Carlone, L.: Teaser: Fast and certifiable point cloud registration. IEEE Transactions on Robotics **37**(2), 314–333 (2020)
-  Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J., Funkhouser, T.: 3dmatch: Learning the matching of local 3d geometry in range scans. In: CVPR. vol. 1, p. 4 (2017)
-  Zhou, Q.Y., Park, J., Koltun, V.: Fast global registration. In: European conference on computer vision. pp. 766–782. Springer (2016)