Adaptive Annealing for Robust Geometric Estimation

Chitturi Sidhartha, Lalit Manam and Venu Madhav Govindu

Computer Vision Lab, Department of Electrical Engineering, Indian Institute of Science, Bengaluru, INDIA

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Graduated Non-Convexity (GNC)

- Minimization of robust cost:
  $$\underset{x}{\text{argmin}} \left\{ f(x) = \sum_{i=1}^{m} \rho_\sigma (\|r_i(x)\|) \right\}$$
- *Algorithm:* Iteratively Reweighted least squares (IRLS).
- GNC used to mitigate convergence to poor local minima caused by IRLS.
- Minimize a sequence of successively harder costs, $\rho_\sigma(\cdot)$, $\sigma = \sigma_0 > \sigma_1 > \sigma_2 > \cdots > \sigma_{min}$.
- $\gamma_k = \frac{\sigma_{k+1}}{\sigma_k}$,
- *Fixed annealing:* γ_k =constant $\forall\ k$.

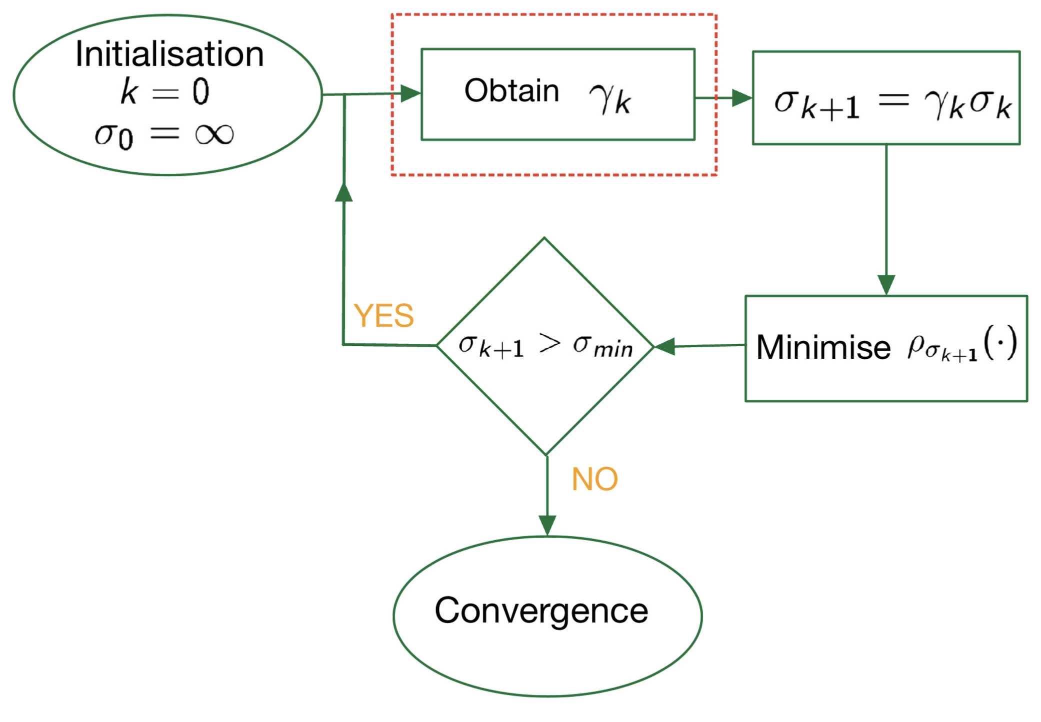

Figure: GNC Flowchart.

Why Adaptive Annealing?

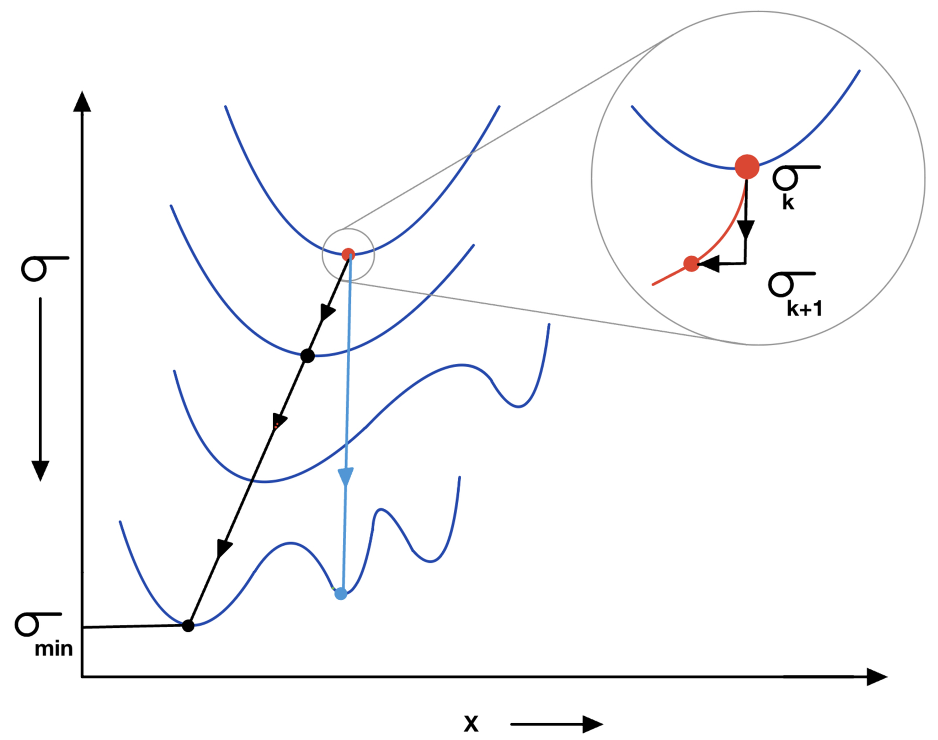

Figure: Cost landscape for varying σ.

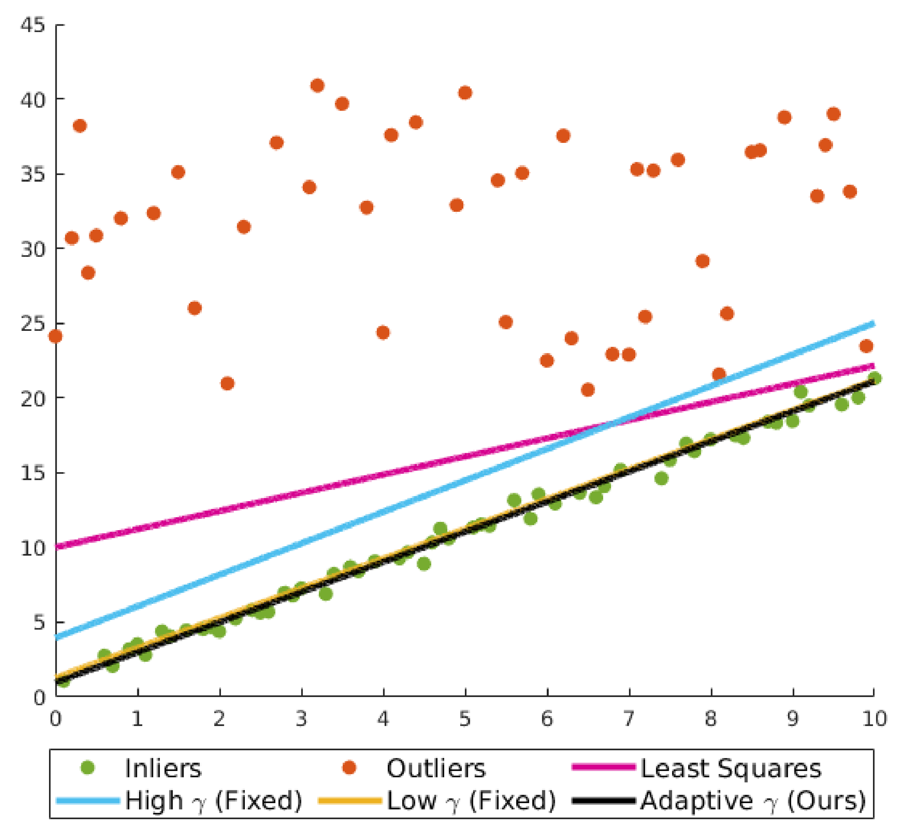

Figure: Line Fitting example.

Properties: Adaptive Annealing vs Fixed Annealing

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

Table: Impact of different annealing strategies.

Adaptive Strategy to Find σ_{k+1} ($\equiv \gamma_k$)

- Given the current solution x_k^*, choose σ_{k+1} as:
  $$\sigma_{k+1} = \min_{\sigma \le \sigma_k} \left\{ \sigma \mid \lambda_{min}\left(H_{x_k^*}(\sigma)\right) > \lambda_T > 0 \right\} \quad (1)$$
- Given the least squares cost's gradient (g_{LSQ}) and Hessian (H_{LSQ}),
  $$H(\sigma) = \sum_{i=1}^{n} \left(-l_i \frac{g_{LSQ,i} g_{LSQ,i}^\top}{\|r_i\|^2} + m_i H_{LSQ,i} \right), \quad (2)$$
  where $l_i = \frac{\rho'(\|r_i\|)}{\|r_i\|} - \rho''(\|r_i\|)$, $m_i = \frac{\rho'(\|r_i\|)}{\|r_i\|}$ (3)
- $\lambda_{min}(H_{x_k^*}(\sigma))$ is monotonic in $\sigma \Rightarrow$ binary search is used to find σ_{k+1} in Eqn. 1.
- Approximate l_i's and m_i's used to make the Hessian (H) amenable for cheap evaluations of $\lambda_{min}(H)$.

Algorithm 1: Robust M-estimation with Adaptive GNC (GNCp)

Input: Observations $\{o_i\}$; residual function $r_i(x) = g(x, o_i)$; σ_{final}

Output: Optimal x

Initialization: $k = 0$, $\sigma = \sigma_0$, x_0 (random)

while $\sigma_k \ge \sigma_{final}$ **do**

- $r_i = g(x_k, o_i)$
- $\phi_i = \frac{1}{\left(1+\frac{\|r_i\|^2}{\sigma_k^2}\right)^2}$
- /* Solve for x_{k+1} using IRLS */
- $x_{k+1} = \underset{x}{\text{argmin}} \sum_{i=1}^{n} \phi_i \|r_i(x)\|^2$
- Obtain $H(\sigma)$ using Eqn. 2
- Perform *binary search* on $\lambda_{min}(H)$ to obtain σ_{k+1}
- $k = k + 1$

end

Illustration on Pairwise 3D Registration

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

Optimization problem: $$\min_{(R,t)} \sum_{i=1}^{N} \rho_\sigma(\|a_i - Rb_i - t\|) \quad (4)$$

Results on Synthetic Data

	Mean Rotation Errors (deg) ↓				Mean Translation Errors ($\times 10^{-3}$) ↓			
Dataset	FGR [5]	SE3Reg [1]	TEASER++ [3]	GNCp(Ours)	FGR [5]	SE3Reg [1]	TEASER++ [3]	GNCp(Ours)
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.

Results on Real Data

Dataset	Success % ↑				Time (in ms.) ↓			
3DMatch	FGR [5]	SE3Reg [1]	TEASER++ [3]	GNCp (Ours)	FGR [5]	SE3Reg [1]	TEASER++ [3]	GNCp (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**
KITTI	73.5	84.7	–	**85.6**	195	40	–	**13**

Table: Results on 3D Match dataset [4] and KITTI [2] datasets.

(a) SE3Reg (b) TEASER++ (c) GNCp (Ours) (d) Ground Truth

Figure: Two point clouds (red and green) with a low overlap in the MIT Lab sequence of 3D Match dataset.

Conclusion

- We propose a principled way to adaptively anneal the scale in Graduated Non-Convexity (GNC) by tracking the positive definiteness of the Hessian matrix of the cost function.
- We demonstrate our approach on pairwise registration of 3D point clouds. Our method yields superior results compared to the state-of-the-art both in accuracy and efficiency.

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

[1] 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)

[2] 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)

[3] Yang, H., Shi, J., Carlone, L.: Teaser: Fast and certifiable point cloud registration. IEEE Transactions on Robotics **37**(2), 314–333 (2020)

[4] 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)

[5] Zhou, Q.Y., Park, J., Koltun, V.: Fast global registration. In: European conference on computer vision. pp. 766–782. Springer (2016)

{chitturis,lalitmanam,venug}@iisc.ac.in