Differentially Private 2D Human Pose Estimation

Kaushik Bhargav Sivangi, Paul Henderson, Fani Deligianni

School of Computing Science, University of Glasgow

Contact: {kaushik.sivangi, paul.henderson, fani.deligianni}@glasgow.ac.uk

Project Page

CVPR JUNE 3-7, 2026 DENVER COLORADO

Motivation

- **HPE relies on sensitive visual data.** Images can reveal identity, body traits, activity patterns, and private environments.
- **Anonymization is not enough.** Anonymization may hide visible identity cues, but they lack formal privacy guarantees and can remove pose-relevant details.
- **DP gives formal privacy but hurts accuracy.** DP-SGD adds noise to training, which strongly affects fine-grained keypoint localization.
- **Reducing DP noise through Feature-Projective DP.** Feature-Projective DP adds noise only to private image gradients and projects noisy updates into a pose-relevant subspace.

Contributions

- **First DP benchmark for 2D-HPE.** We evaluate private pose estimation across various privacy budgets, clipping thresholds, and multiple datasets.
- **Feature-Projective DP for pose estimation.** We combine feature differential privacy with subspace projection to reduce DP noise while preserving pose- relevant gradients.
- **Theory and empirical validation.** We theoretically show that the combined effect of projection and FDP is multiplicative in terms of signal-to- noise ratio and convergence speed.

Methodology

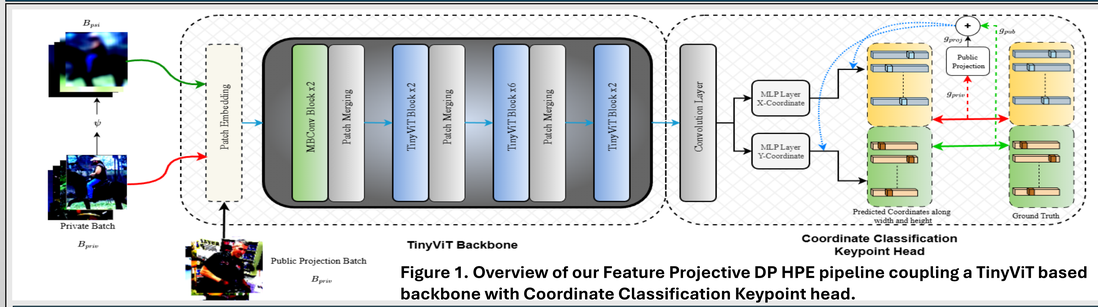

Figure 1. Overview of our Feature Projective DP HPE pipeline coupling a TinyViT based backbone with Coordinate Classification Keypoint head.

DP-SGD

DP-SGD ensures (ϵ, δ) - differential privacy by clipping per sample gradients and by adding gaussian noise

Per Sample Gradient Clipping:

$$\tilde{g}_i = clip(\nabla l(w, z_i), C)$$

Noisy Gradient Aggregation:

$$g = \frac{1}{B}\left(\sum_{i \in B} \tilde{g}_i + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})\right)$$

Projection based DP-SGD

The noisy gradients are restricted to a k – dimensional subspace spanned by top- k eigen vectors $V \in \mathbb{R}^{\mathbb{P} \times \mathbb{k}}$

Projection:

$$g_{proj} = (VV^T)g$$

The update direction is limited to informative subspace, reducing the noise variance by factor k/p

Feature Projective DP-SGD

The total loss is decomposed into public and private components and DP noise is added only to the gradient of private component:

Gradient on Public batch:

$$g^t_{pub} = \frac{1}{|B^t_{psi}|}\sum_{x \in B^t_{psi}} \nabla l_{pub}\left(w_{t-1}, \psi(x)\right)$$

Gradient on Private batch:

$$g^t_{priv} = \frac{1}{|B^t_{priv}|}\left(\sum_{x \in B^t_{priv}} \tilde{g} + \mathcal{N}(0, \sigma^2 C^2 I)\right)$$

Projection and Update:

$$g^t_{proj} = (V_t V_t^T) g^t_{priv};\ g_t = g^t_{pub} + g^t_{proj}$$

Qualitative Results

Evaluation on MPII (across private and non private baselines)

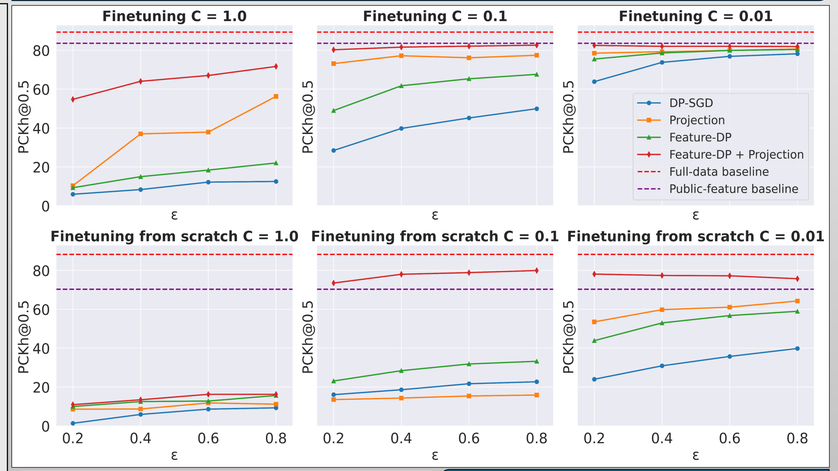

Evaluation on HumanART

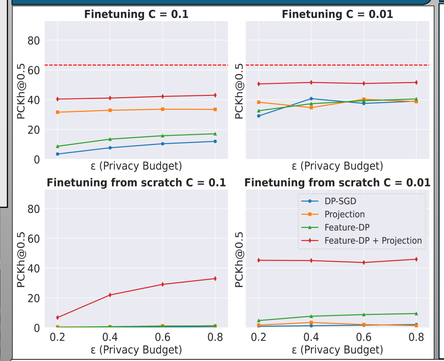

Takeaways

- Vanilla DP-SGD causes large accuracy drops in 2D-HPE because keypoint localization is highly sensitive to noisy gradients.
- Separating public visual cues from private raw-image information improves utility under the same privacy budget.
- Projecting noisy gradients into a low-dimensional pose-relevant subspace removes redundant noise.
- Combining feature DP and subspace projection gives stronger privacy-utility trade-offs than vanilla DP-SGD on MPII and HumanART.

- *Wu, K., Zhang, J., Peng, H., Liu, M., Xiao, B., Fu, J., & Yuan, L. (2022, October). Tinyvit: Fast pretraining distillation for small vision transformers. In European conference on computer vision (pp. 68-85). Cham: Springer Nature Switzerland.*
- *Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016, October). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security (pp. 308-318).*
- *Zhou, Y., Wu, Z. S., & Banerjee, A. (2020). Bypassing the ambient dimension: Private sgd with gradient subspace identification. arXiv preprint arXiv:2007.03813.*
- *Mahloujifar, S., Guo, C., Suh, G. E., & Chaudhuri, K. (2025, May). Machine learning with privacy for protected attributes. In 2025 IEEE Symposium on Security and Privacy (SP) (pp. 2640-2657). IEEE.*

Contact!