NMSL@NTHU

Networking and Multimedia Systems Lab

Mitigating Privacy Threats of HMD Users Without Degrading Visual Quality of VR Applications

YuSzu (weiyousz0328@gmail.com)

Advisor: Cheng-Hsin Hsu

Networking and Multimedia Systems Lab, ISA, National Tsing Hua University

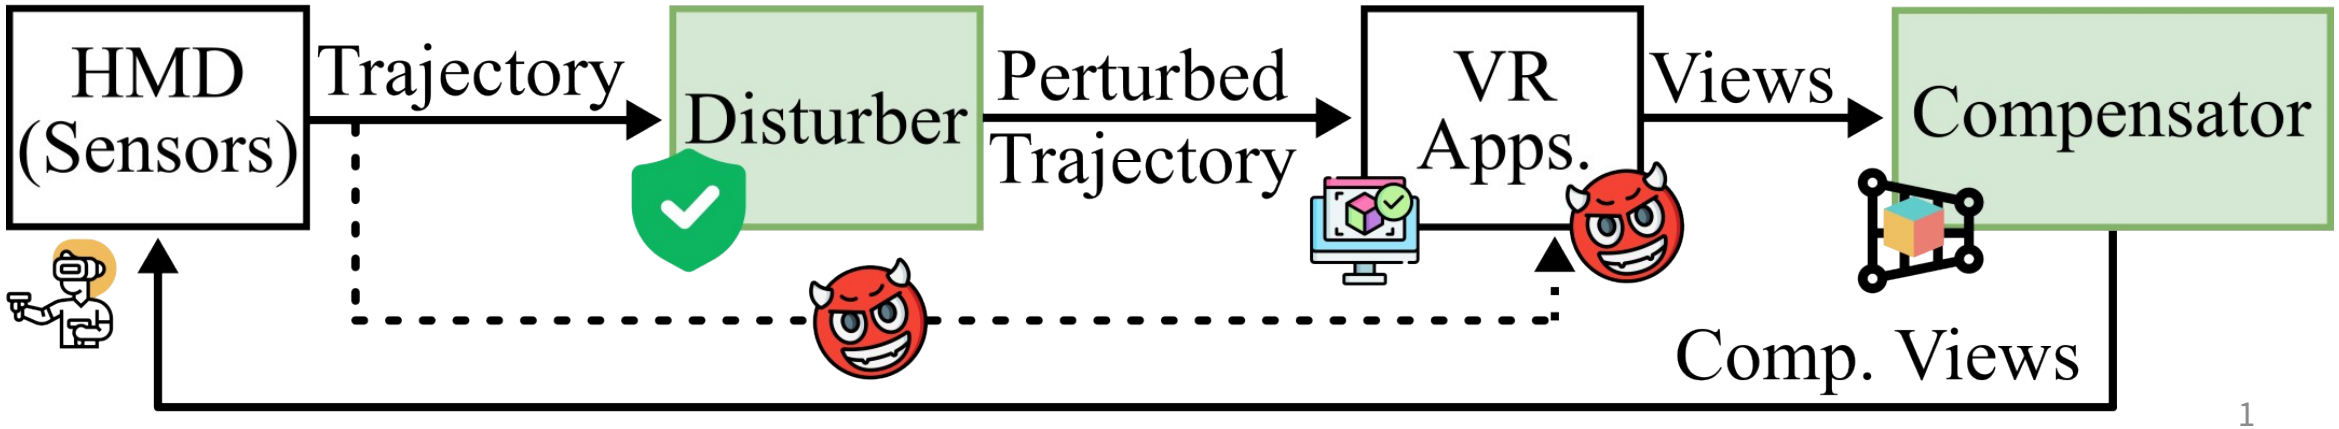

Importance of Privacy Protection in VR

- Detailed sensor data are collected and streamed in VR

Trajectory

- HMD locations and orientations
- Left/Right controller locations and orientations

Pose: The trajectory at a moment

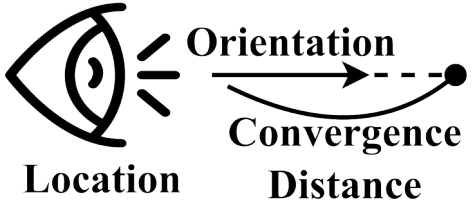

Eye Gaze

- Left/Right eye locations and orientations
- Convergence distance

Facial Expression

- Blendshapes [1]

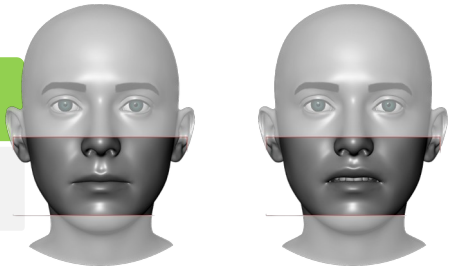

Physical Appearance

- RGBD video [2]

- The policies for protecting these data are unclear
- These data may reveal users' privacy

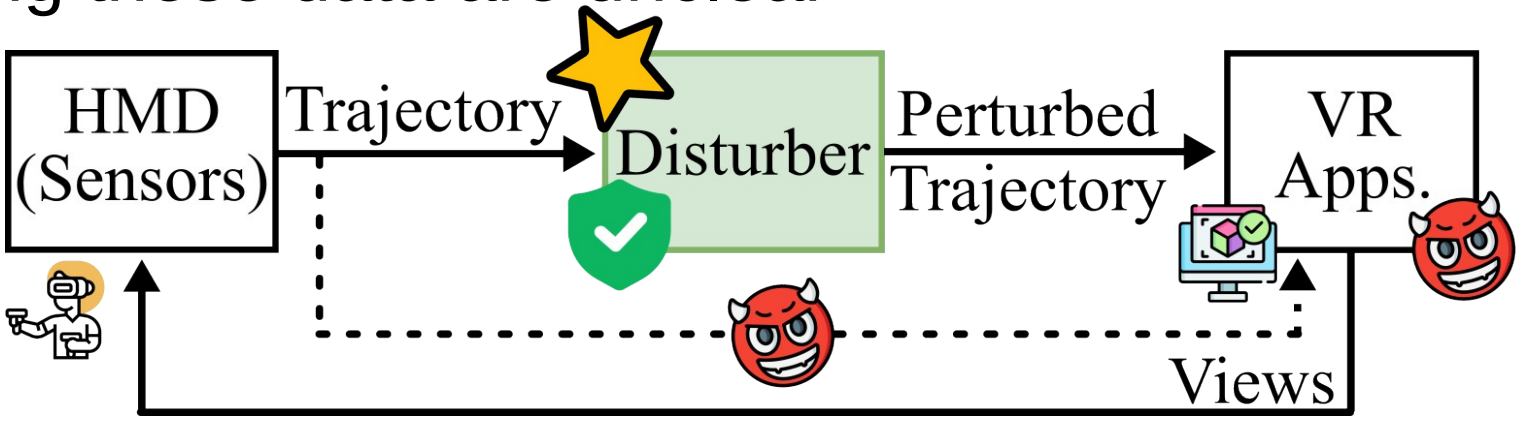

Disturber: An agent to add perturbation to the sensor data

penXR: Native | Oculus Developers, 2023. https://tinyurl.com/4rd49rp4
truction in immersive virtual reality applications. IEEE Transactions on Visualization and Computer Graphics, 2019

Compensation for the Perturbed Views

- The perturbed trajectory leads to shaky perturbed views

- The perturbed view with larger viewport may contain the original

Compensator: An agent to warp the perturbed view back to the original one

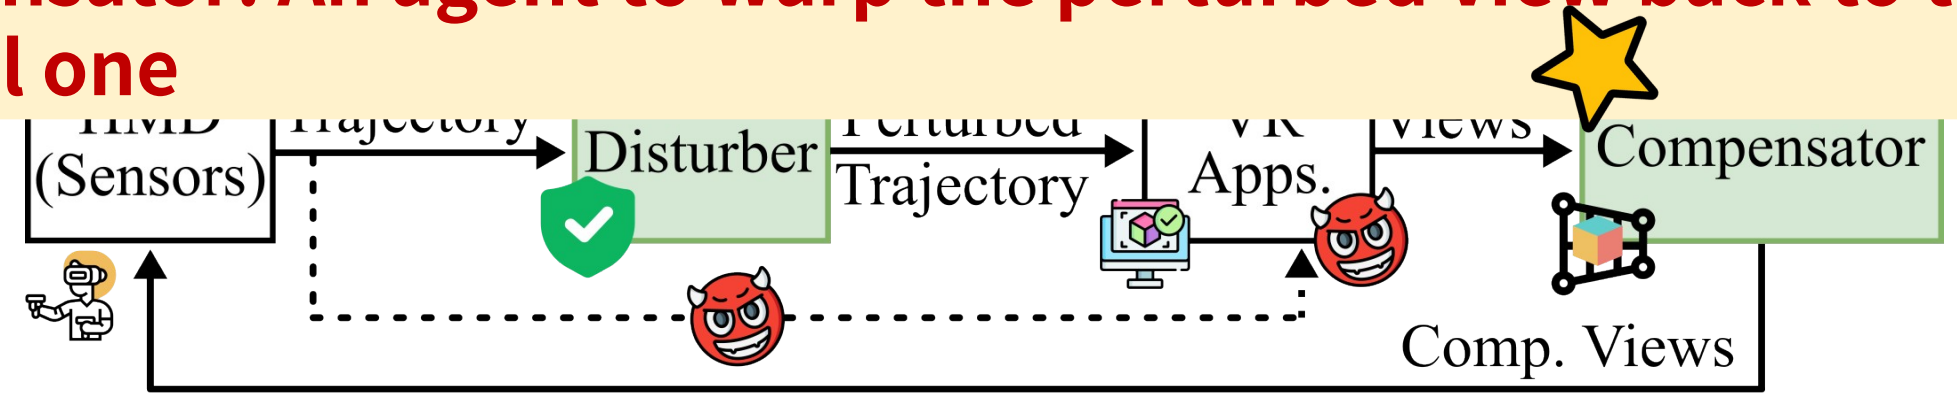

Challenges & Goal

Goal

Perturb the trajectory to mitigate the privacy threats while preserving high visual quality

Challenges

Perturbation for both temporal and spatial domain

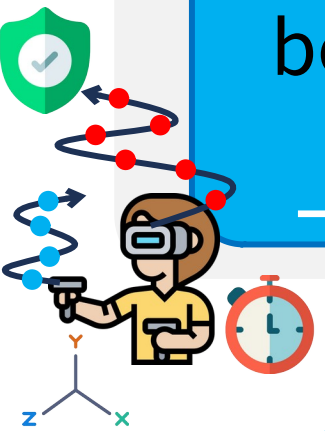

- Trajectory estimation
- Random Sampling

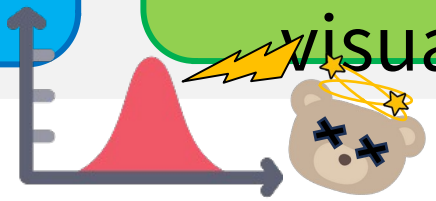

Random perturbation degrades the visual qua

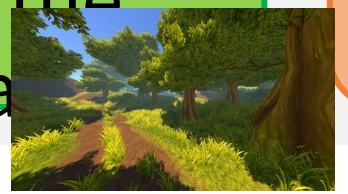

- Image warping
- Privacy-quality tradeoff

Diverse privacy threats

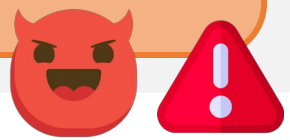

- General
- Temporal + Spatial
- Re-identification (Re-id) attack

Re-id: Identify the users with their trajectory

Outline

- Introduction
- Related Work
- 6DoF VR Dataset
- Privacy Threats Mitigation
- Evaluations
- Conclusion & Future Work

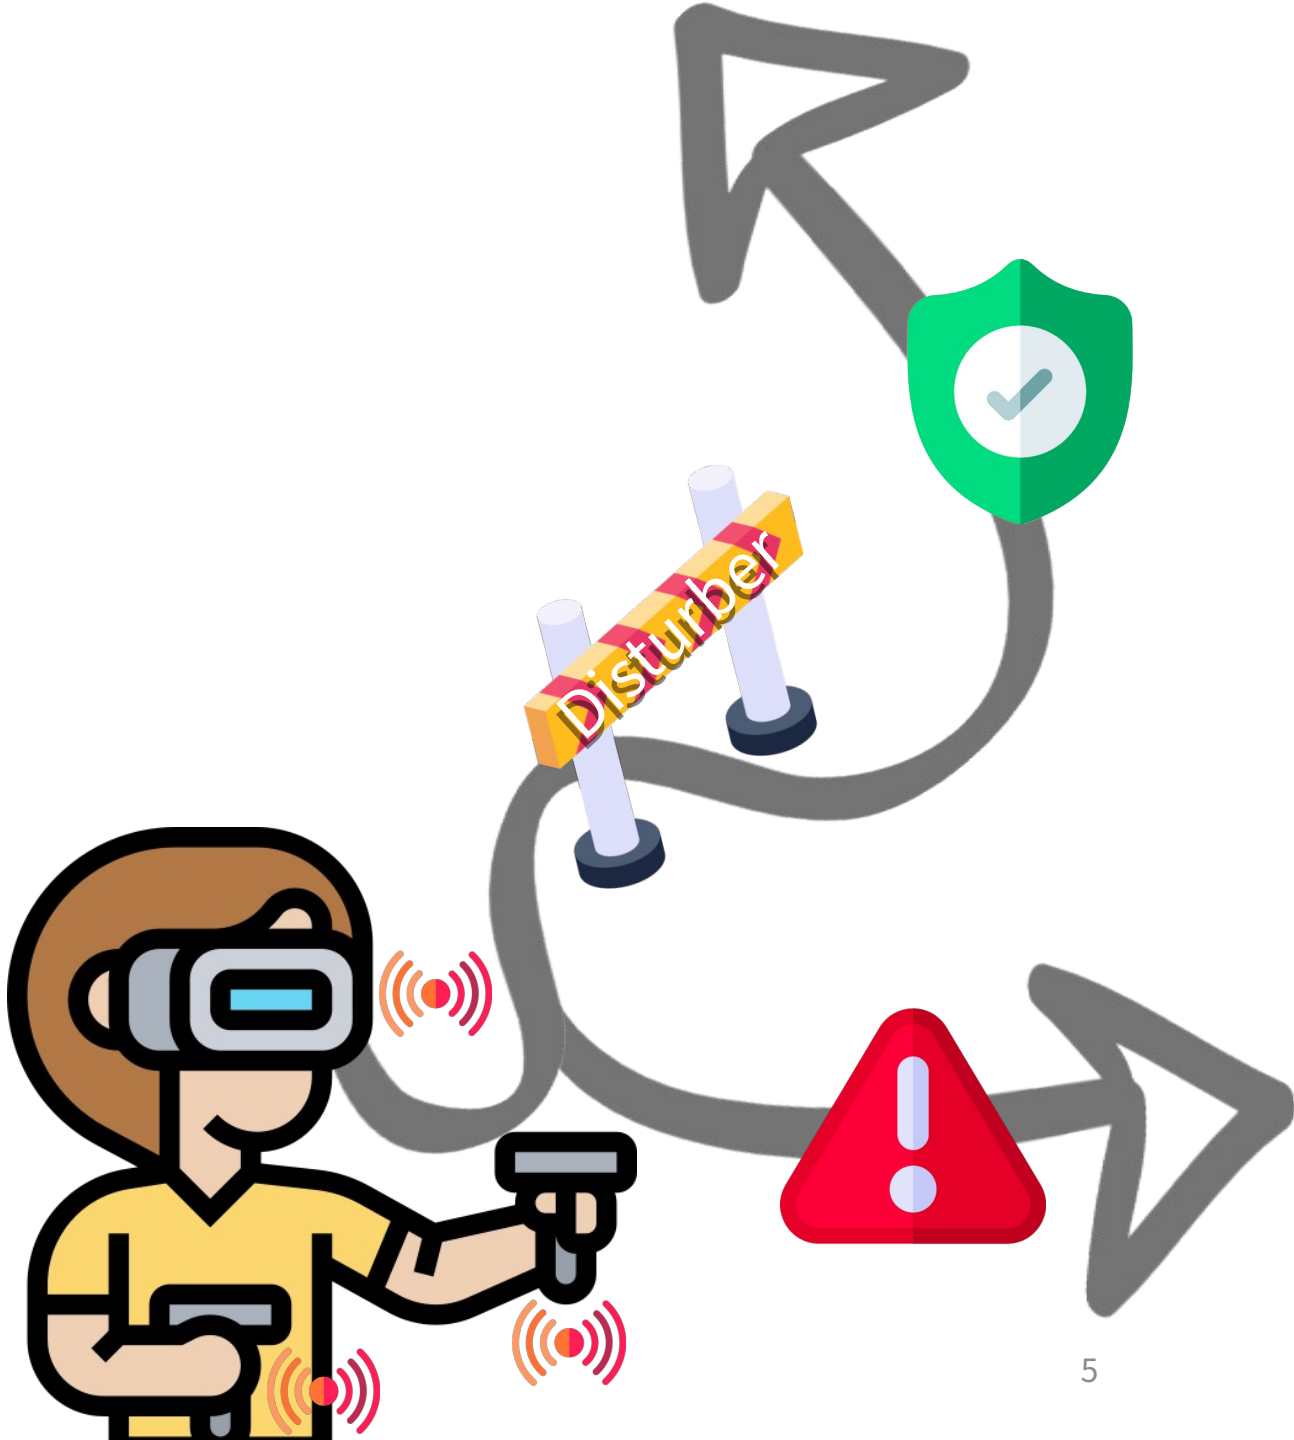

Perturbations for eye gaze

- Eye Features [RPI, SIC]

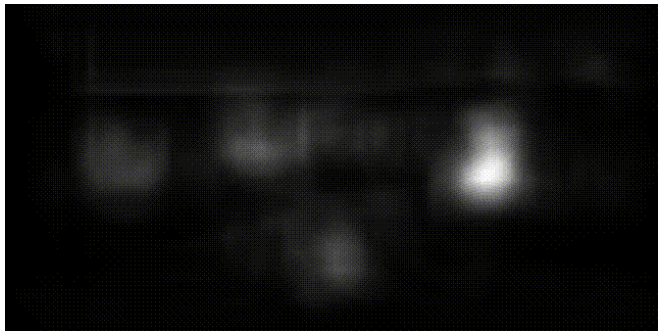

A saliency map

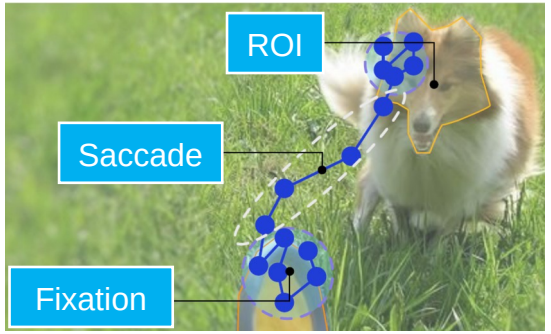

Common eye gaze features

2D Saliency maps, gaze maps [RPI]

- Eye movements [SIC]

Consider eye gazes only

- Eye Gaze Trace [IEEE, UW–Madison]

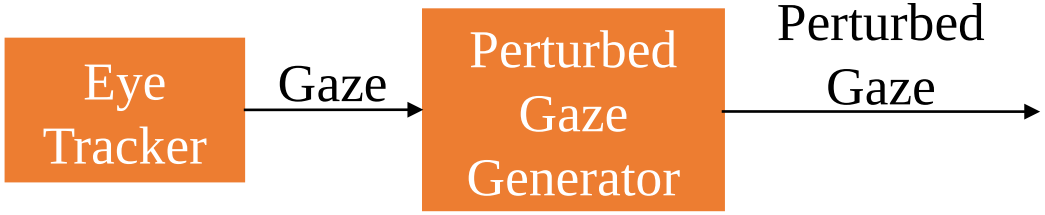

- Additive Gaussian noise, temporal downsampling, and spatial downsampling [UFL]

2D Geo-indistinguishability [UW–Madison]

Consider 2D content only

[RPI] A. Liu, L. rivacy for eye-tracking da ns (ETRA), Denver, USA, June 2019
[SIC] J. Steil, I. Hagestedt, X. Huang, and A. Bulling. Privacy-aware eye tracking using differential privacy. In Proc. of ACM Symposium on Eye Tracking Research & Applications (ETRA), Denver, USA, June 2019
[UFL] B. David, D. Hosfelt, K. Butler, and E. Jain. A privacy-preserving approach to streaming eye-tracking data. IEEE Transactions on Visualization and Computer Graphics, 2021
[UW–Madison] J. Li, A. Roy, K. Fawaz, and Y. Kim. Kalɛido: Real-time privacy control for eye-tracking systems. In Proc. of USENIX Security Symposium, Virtual, August 2021.

Perturbation for VR trajectory data

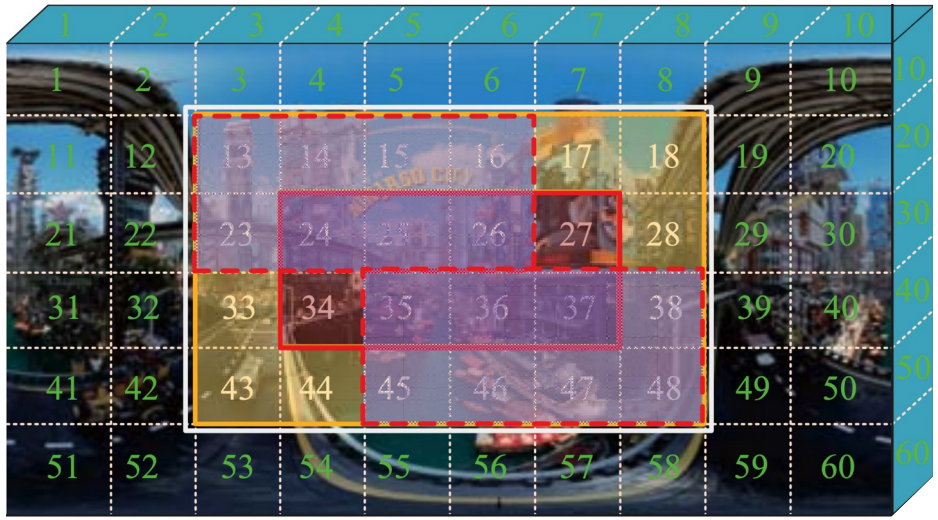

- Add noisy tiles around user consumed tiles in 360 video [BUAA]

Consider tiled 360 video only

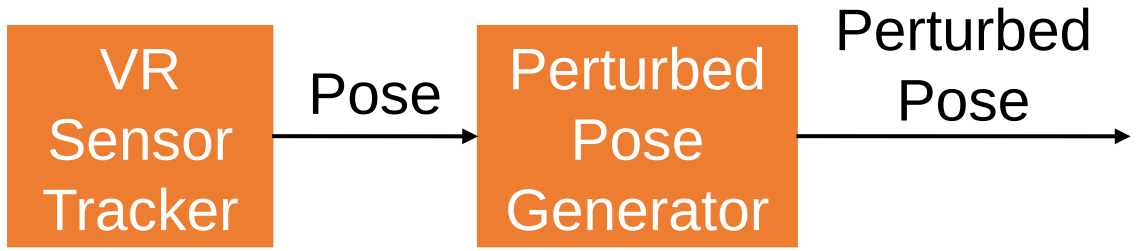

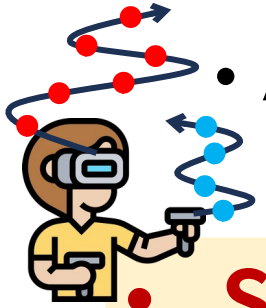

- Add perturbation to VR trajectories [UCB]

- **Sample one perturbation throughout the whole session**
- **Does not consider the temporal correlation of VR trajectory**

[BUAA] X. Wei and C. Yang. FoV privacy-aware VR streaming. In Proc. of IEEE Wireless Com
[UCB] V. Nair, M. Garrido, and D. Song. Going incognito in the metaverse. arXiv:2208.05604, 2

Outline

- Introduction
- Related Work
- 6DoF VR Dataset
- Privacy Threats Mitigation
- Evaluations
- Conclusion & Future Work

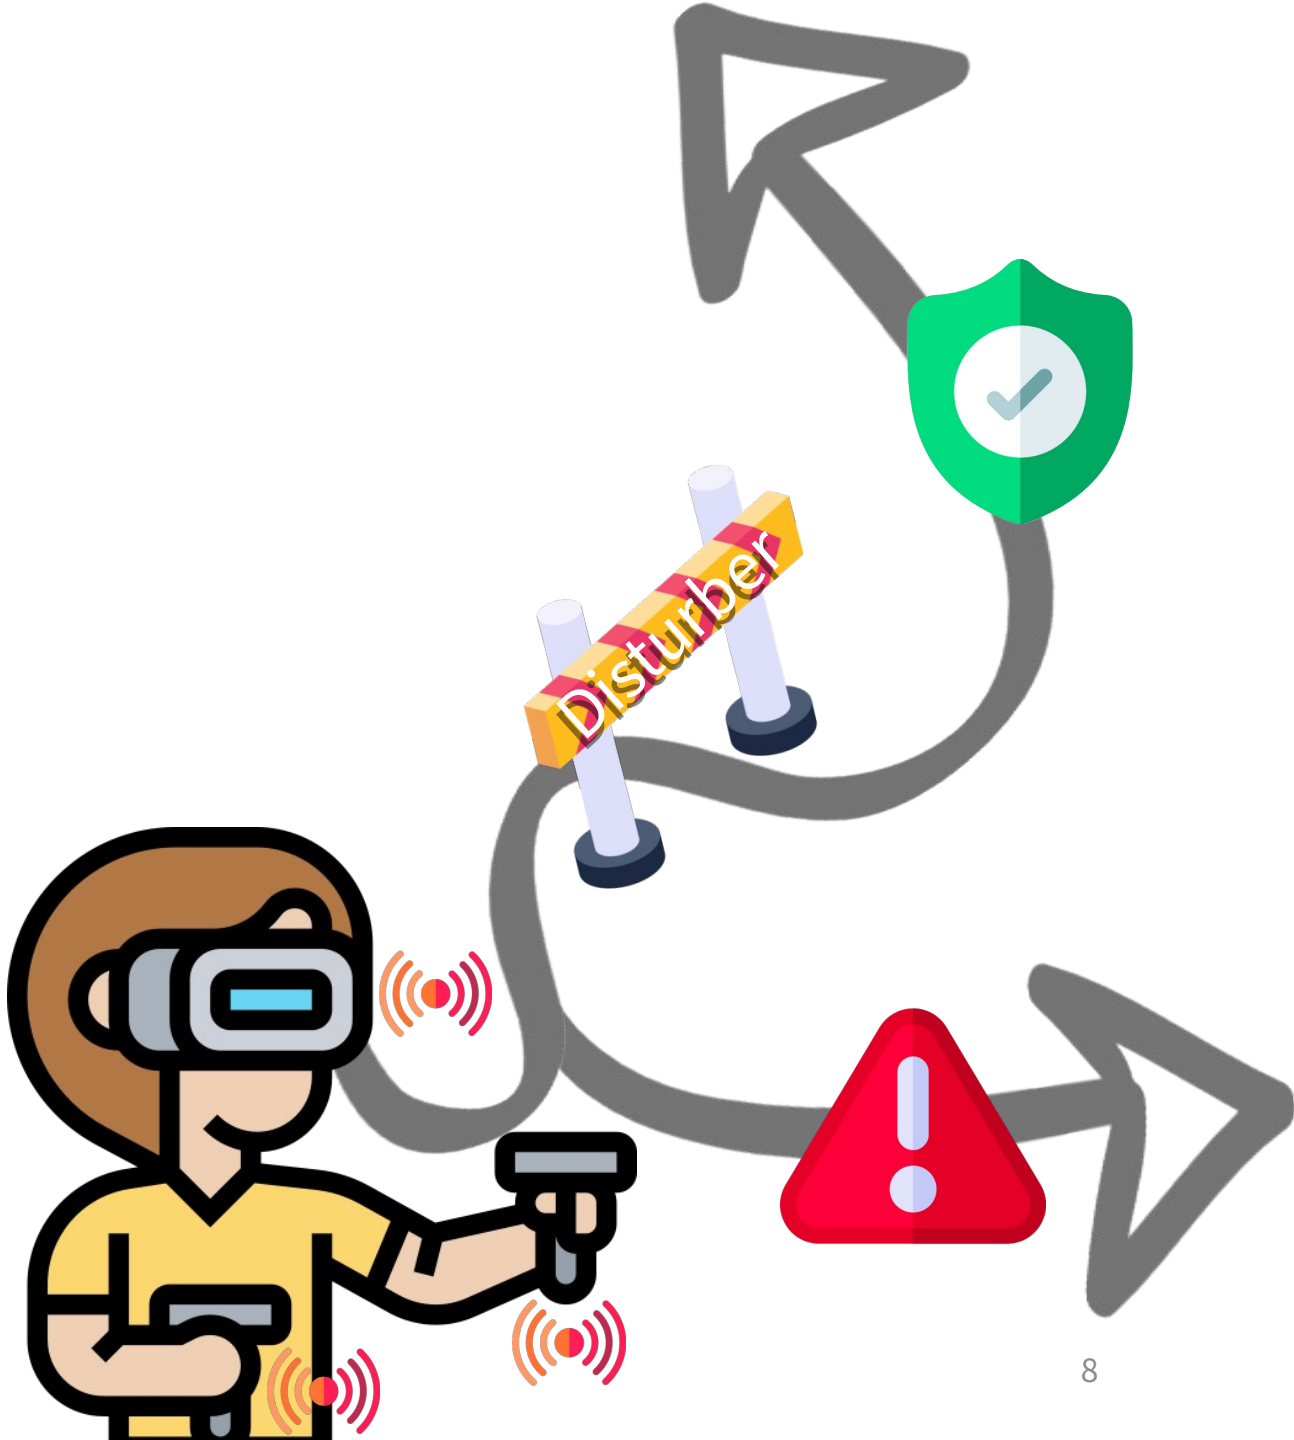

Dataset of 360 videos

HMD Movement

1: timestamp, raw x, raw y, raw z, raw yaw, raw pitch, raw roll
2: 1487571103.944,26.289,28.063,-15.581,-5.246,-4.298,-1.315
3: 1487571103.953,26.291,28.063,-15.567,-5.297,-4.287,-1.333
4: 1487571103.957,26.292,28.063,-15.559,-5.323,-4.284,-1.341
5: …

Motion Map

- Viewport/gaze prediction [NTHU]
- Cybersickness [TUI] SSQ
- Biometrics, head/eye movement, and user emotion correlation [BIT]
- Privacy [SU] Private

Saliency Map

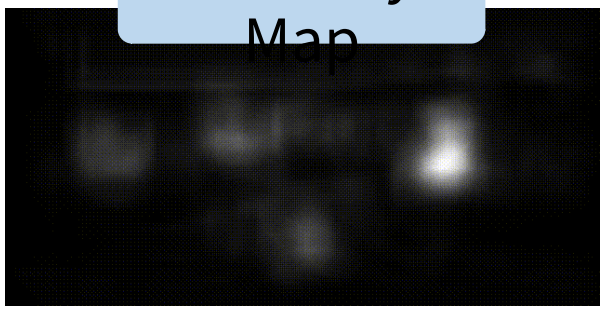

[NTHU] W.-C. Lo, C.-L. Fan, J. Lee, C.-Y. Huang, K.-T. Chen, and C.-H. Hsu. 360° video viewing dataset in head-mounted virtual reality. In Proc. of ACM on Multimedia Systems Conference (MMSys), Taipei, Taiwan, June 2017.

[TUI] S. Fremerey, A. Singla, K. Meseberg, and A. Raake. AVtrack360: An open dataset and software recording people's head rotations watching 360° videos on an HMD. In Proc. of ACM Multimedia Systems Conference (MMSys), Amsterdam, Netherlands, June 2018.

[BIT] T. Xue, A. El, T. Zhang, G. Ding, and P. Cesar. CEAP-360VR: A continuous physiological and behavioral emotion annotation dataset for 360 VR videos. IEEE Transactions on Multimedia, 2021

[SU] M. Miller, F. Herrera, H. Jun, J. Landay, and J. Bailenson. Personal identifiability of user tracking data during observation of 360-degree VR video. Scientific Reports, 2020

Dataset of 3D virtual world

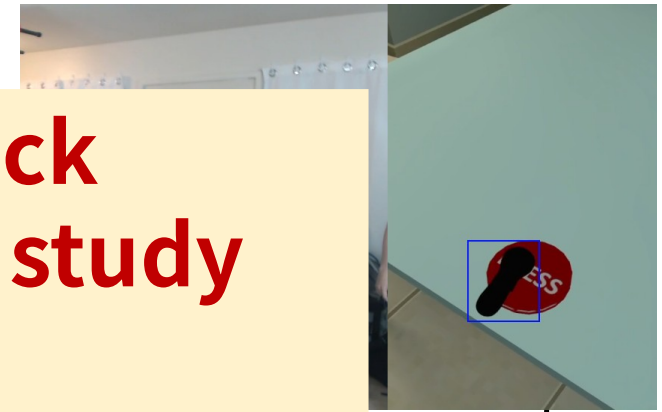

Existing public dataset lack privacy-sensitive data to study privacy issues

- Viewport/gaze
- Cybersickness [MuIT]
- Head/eye movement correlation [PKU]
- Privacy [UCB]

Heart Rate 、 SSQ

- Data related to the tasks
- Demographic

Private

- HMD Movement
- Controller Movement
- Eye Movement
- Scene

[Facebook] K. Emery, M. Zannoli, J. Warren, L. Xiao, and S. Talathi. OpenNEEDS: A dataset of gaze, head, hand, and scene signals during exploration in open-ended VR environments. In Proc. of ACM Symposium on Eye Tracking Research and Application (ETRA), Virtual, May 2021
[MuIT] J. Dong, K. Ota, and M. Dong. Why VR games sickness? an empirical study of capturing and analyzing VR games head movement dataset. IEEE MultiMedia, 2022.
[PKU] Z. Hu, C. Zhang, S. Li, G. Wang, and D. Manocha. SGaze: A data-driven eye-head coordination model for realtime gaze prediction. IEEE Transactions on Visualization and Computer Graphics, May 2019
[UCB] V. Nair, G. Garrido, and D. Song. Exploring the unprecedented privacy risks of the metaverse. arXiv:2207.13176, 2022

6DoF VR Dataset for Privacy Study [1]

[1] Y. Wei, X. Wei, S. Zheng, C. Hsu, and C. Yang. A 6DoF VR dataset of 3D virtualworld for privacy-preserving approach and utility-privacy tradeoff. In Proc. of ACM Multimedia Systems (MMSys), Vancouver, Canada, June 2023

- Collection Testbed

Software

- Rendering Engine: Unity
- Eye tracker: Tobii XR
- Sensor logger

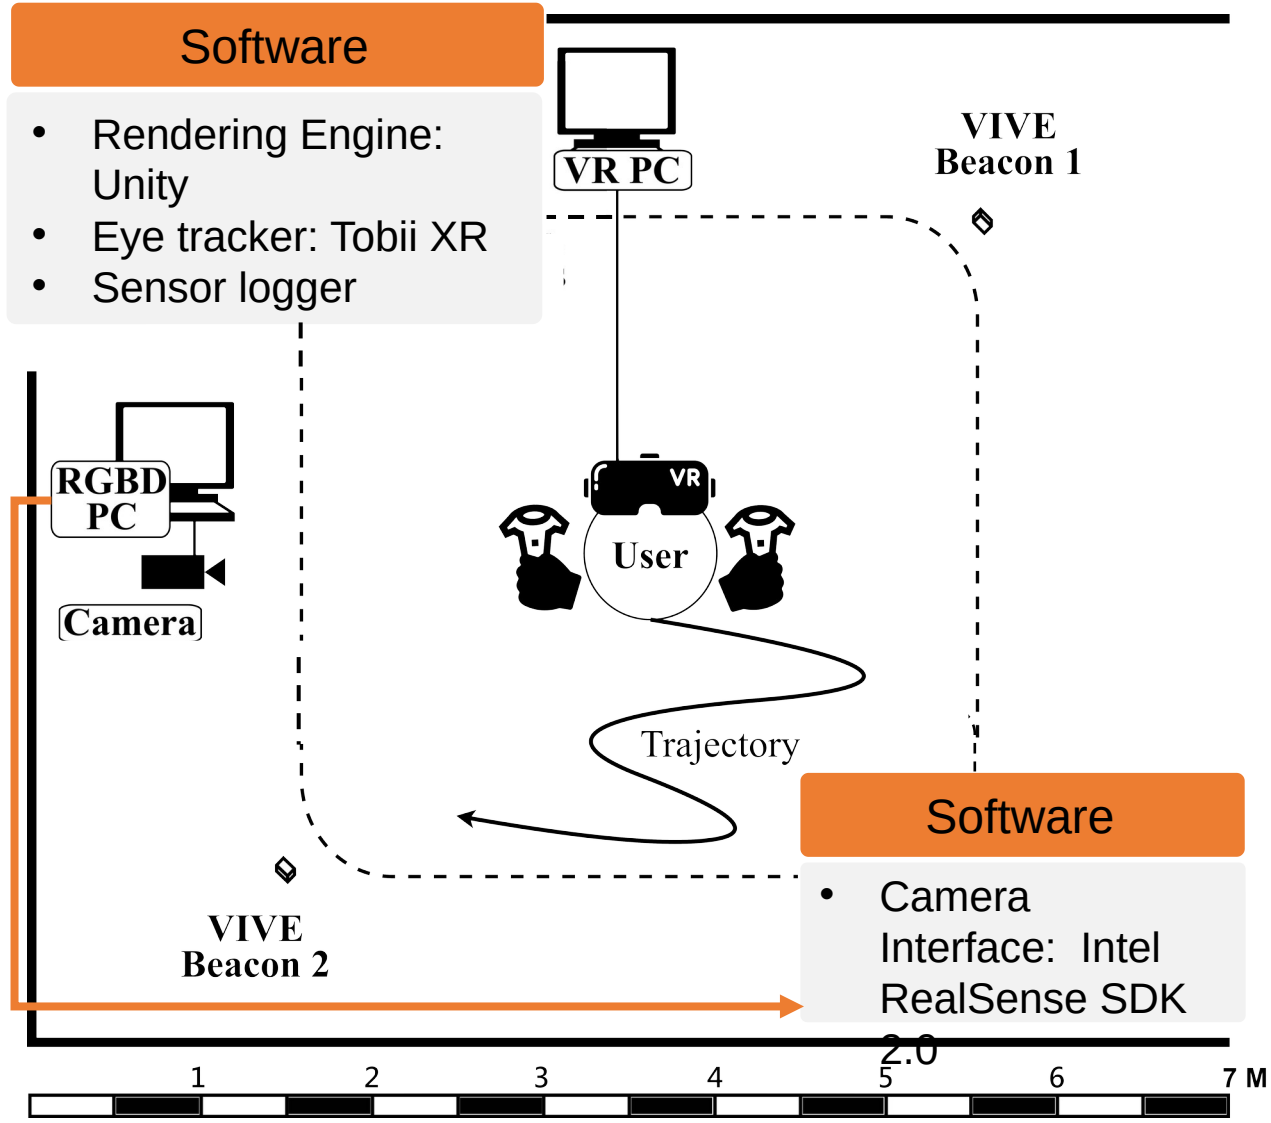

Software

- Camera Interface: Intel RealSense SDK 2.0

- Considered Scenes

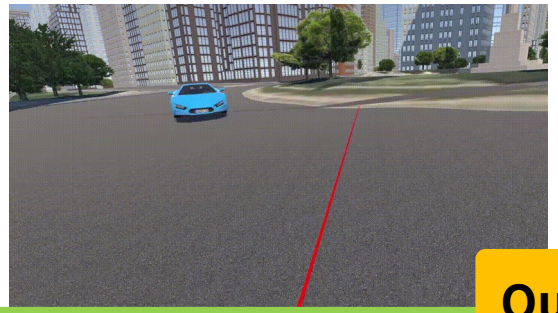

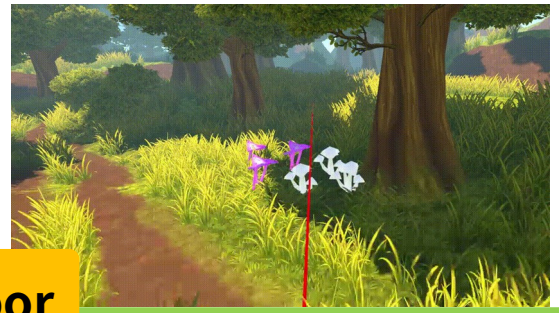

Outdoor

City

- 128*50*128 (m^3)
- Objects: cars
- City sightseeing

Nature

- 6*1.5*6 (m^3)
- Objects: plants
- Natural landscape

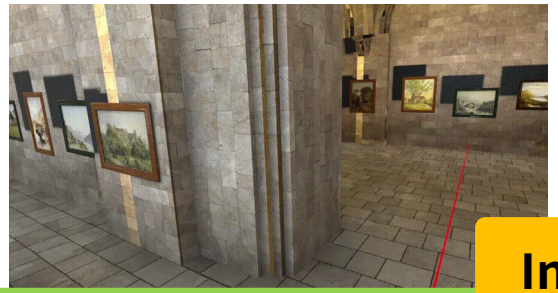

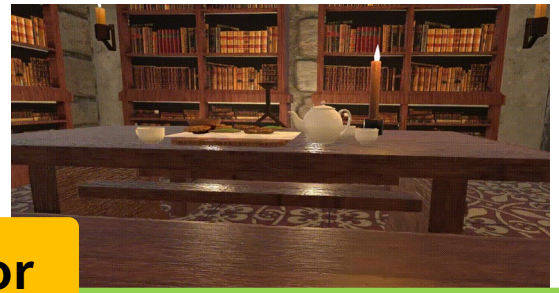

Indoor

Gallery

- 28*9*18 (m^3)
- Objects: statues
- Online museum

Office

- 16*7.5*12 (m^3)
- Objects: book, bench, and cookies
- VR gaming

Procedure

1. Obtain informed consent from all subjects
2. Have each subject answer the demographic and VR background questionnaire
3. For each subject
   a. Set up the VR headset for the subject
   b. Calibrate for the subject's eyes
   c. Launch the VR application
   d. For each VR scene
      I. Have the subject explore the VR scene once
      II. Help the subject out of the VR headset
      III. Have the subject answers the experience questionnaire

Demographic
Age
Gender
Height
Correlated eyesights
Handedness

VR Background
How many times have you used VR before?
How often did you experience motion sickness when using VR?

Experience
How is the overall quality?
How is the visual quality?
Are the objects moved as you expect?
How is the immersive level?
Will you continue exploring the scene under the current system quality and immersive level?

Collected Data

VR Devices Data

- HMD locations and orientations
- Controller locations and orientations
- Controller key strokes
- Eye gaze
- Object locations and orientations

1: Unity time, head pose, left/right ctrl. pose and key strokes, eye pose and conv., ntp time, user id
2: 0.08000000, -724.80690000, 26.66229000, -1040.00900000, …, 1673440396.17647076, 3
3: 0.09999999, -724.80690000, 26.66229000, -1040.00900000, …, 1673440396.23529434, 3
4: 0.12000000, -724.81010000, 26.66180000, -1040.01400000, …, 1673440396.29411793, 3
5: …

Questionnaires' Answers

- Demographic
- VR background
- Experience

1: user id, answers to the demographic questions and VR background questions
2: 0, 20-25, Male, 1.71-1.75, 0.1, 0.1, Right, 0, 1 (Never)
3: 1, 20-25, Male, 1.71-1.75, 1, 1, Right, 2-5, 1 (Never)
4: 2, 20-25, Male, 1.66-1.70, 1.2, 1.2, Right, 2-5, 1 (Never)
5: …

Videos

- Physical world

  RGBD videos

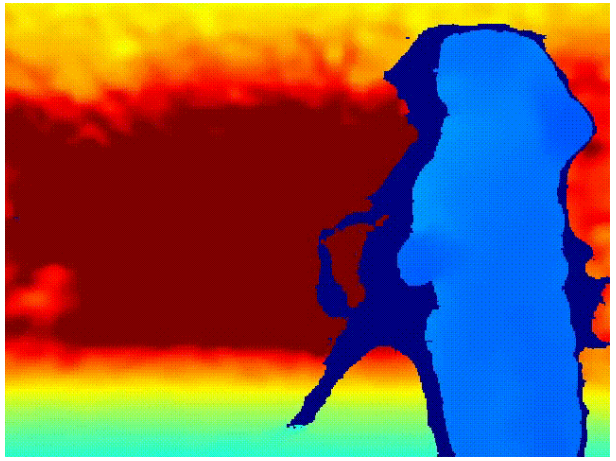

Outline

- Introduction
- Related Work
- 6DoF VR Dataset
- Privacy Threats Mitigation
- Evaluations
- Conclusion & Future Work

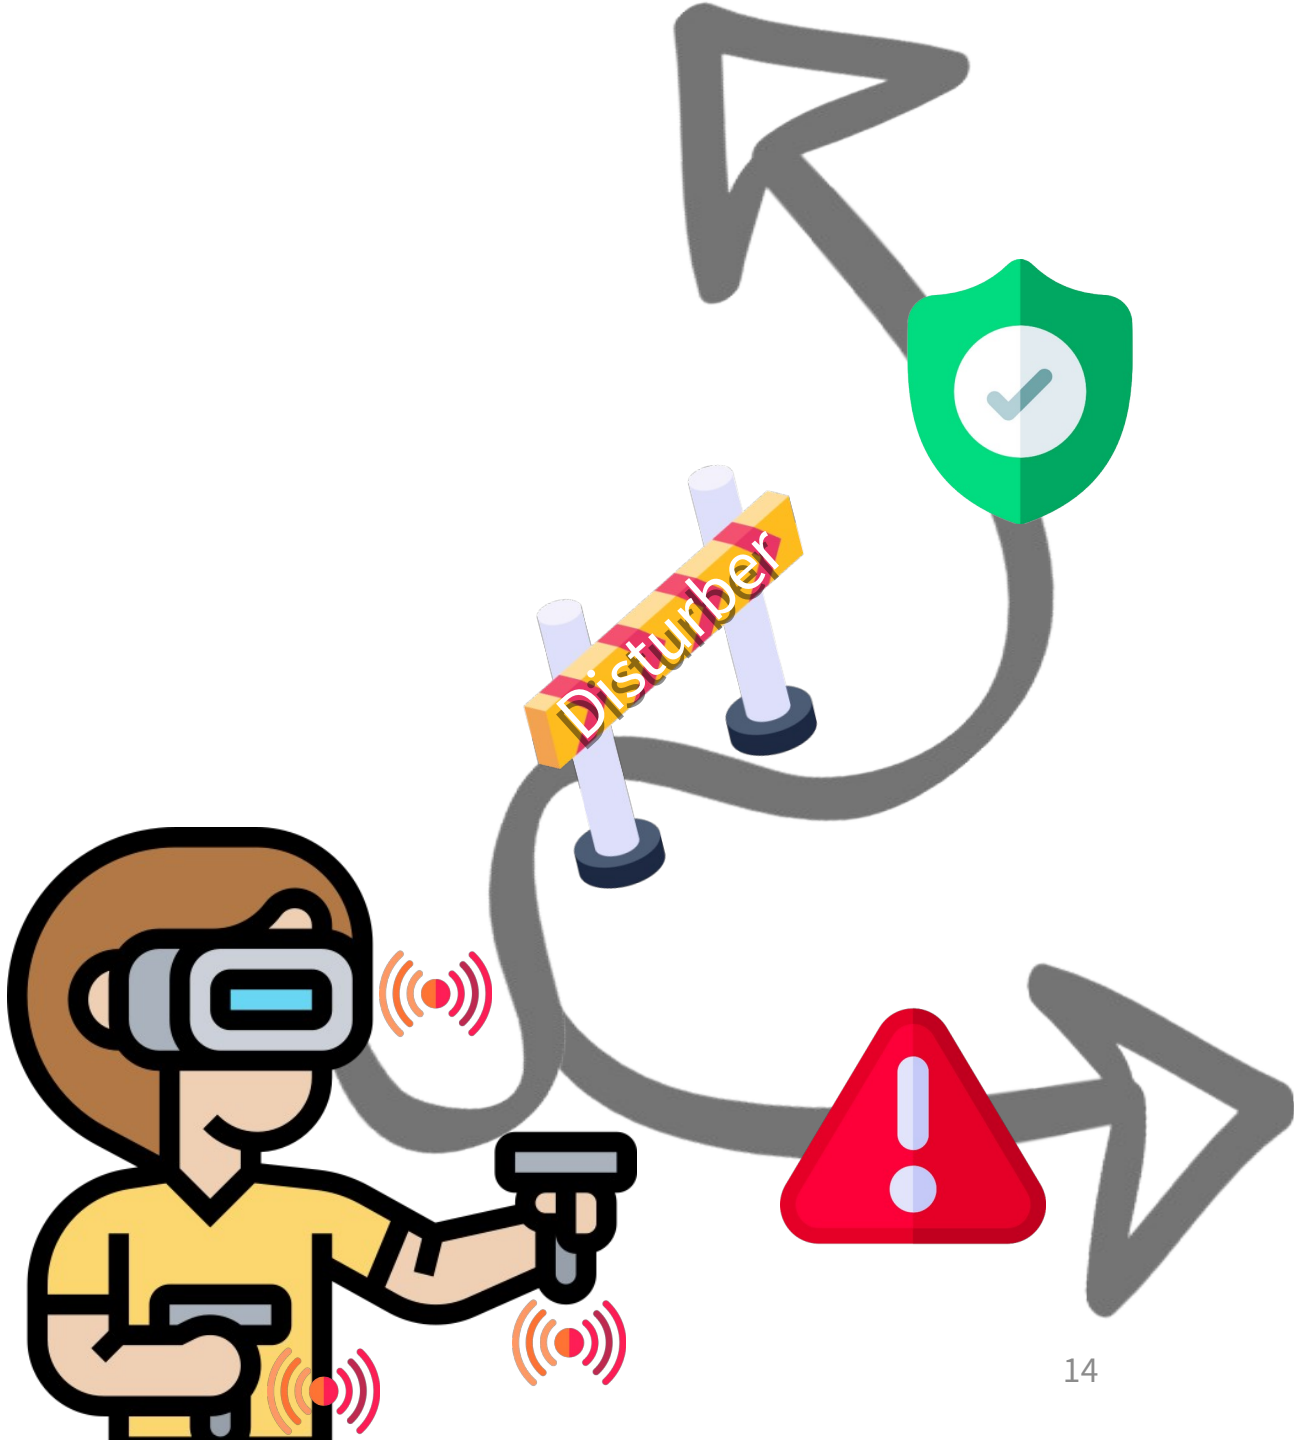

Privacy Threats

- Personal attributes inference [1]
- Re-identification attack [2]
- Typed text inference [3]

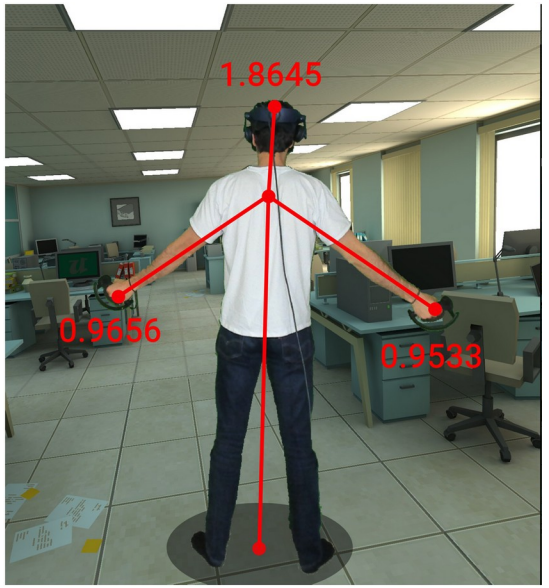

Height and Arm length [1]

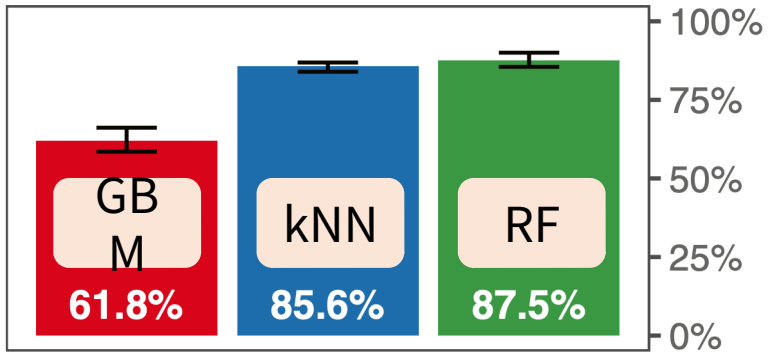

Squat Depth [1]

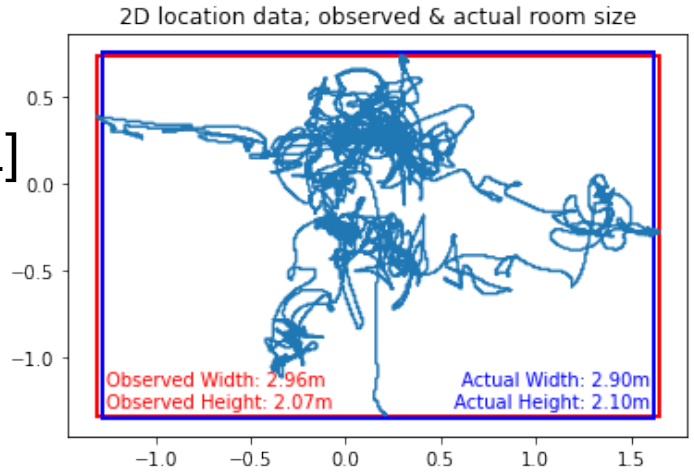

Room Size [1]

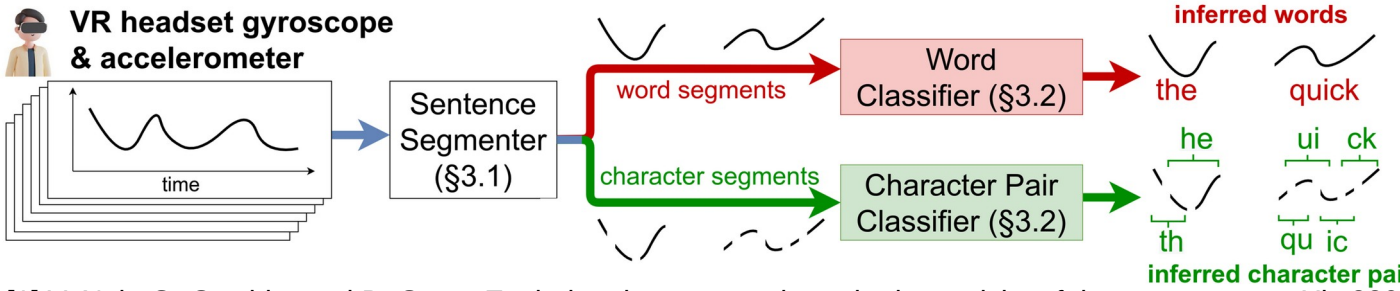

[1] V. Nair, G. Garrido, and D. Song. Exploring the unprecedented privacy risks of the metaverse. arXiv:2207.13176, 2022
[2] M. Miller, F. Herrera, H. Jun, J. Landay, and J. Bailenson. Personal identifiability of user tracking data during observation of 360-degree VR video. Scientific Reports, 2020
[3] C. Slocum, Y. Zhang, N. Abu, and J. Chen. Going through the motions: AR/VR keylogging from user head motions. In Proc. of USENIX Security Symposium, Anaheim, USA, August 2023

Networked VR System

- Each of the entities is a potential attacker

- Personal attributes inference
- Re-identification attack
- Typed text inference

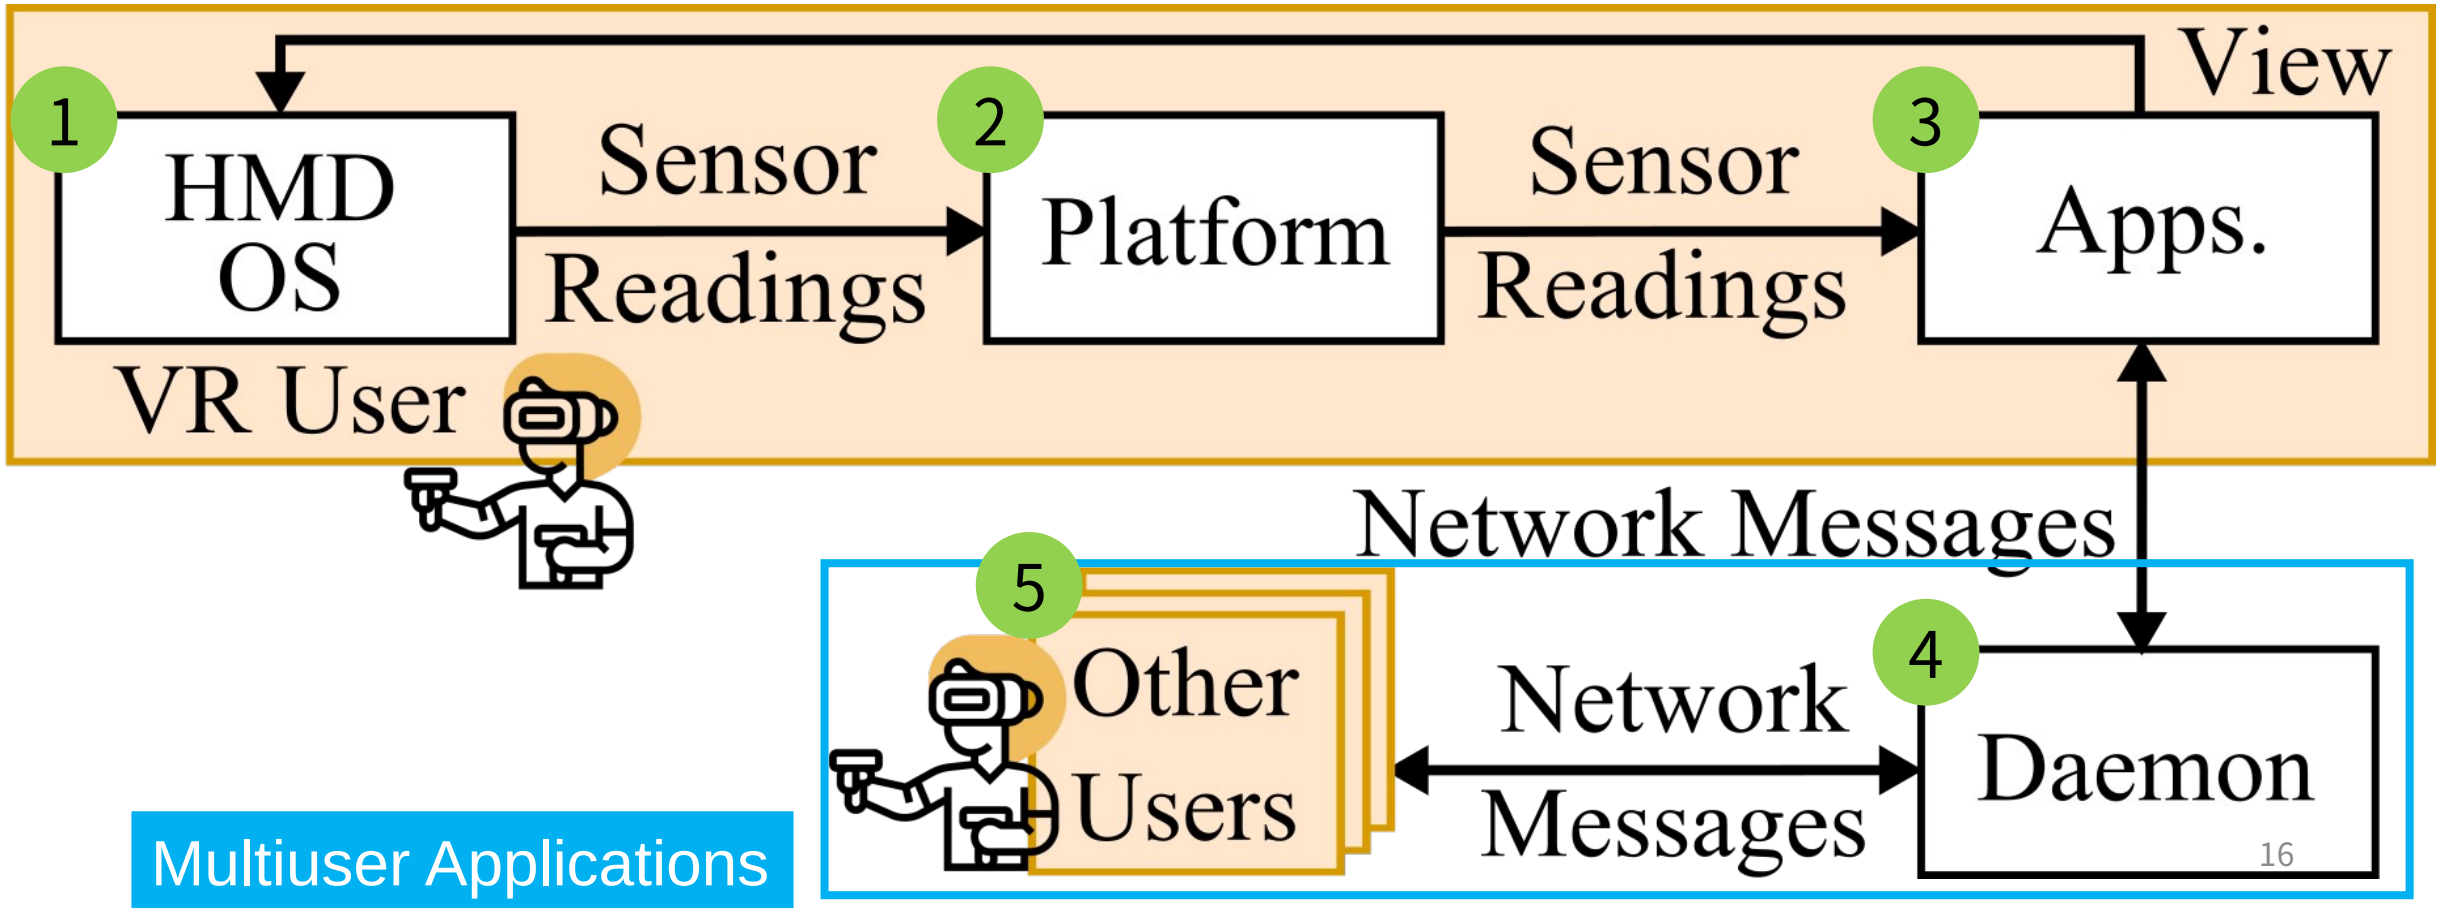

The Placement of the Disturber

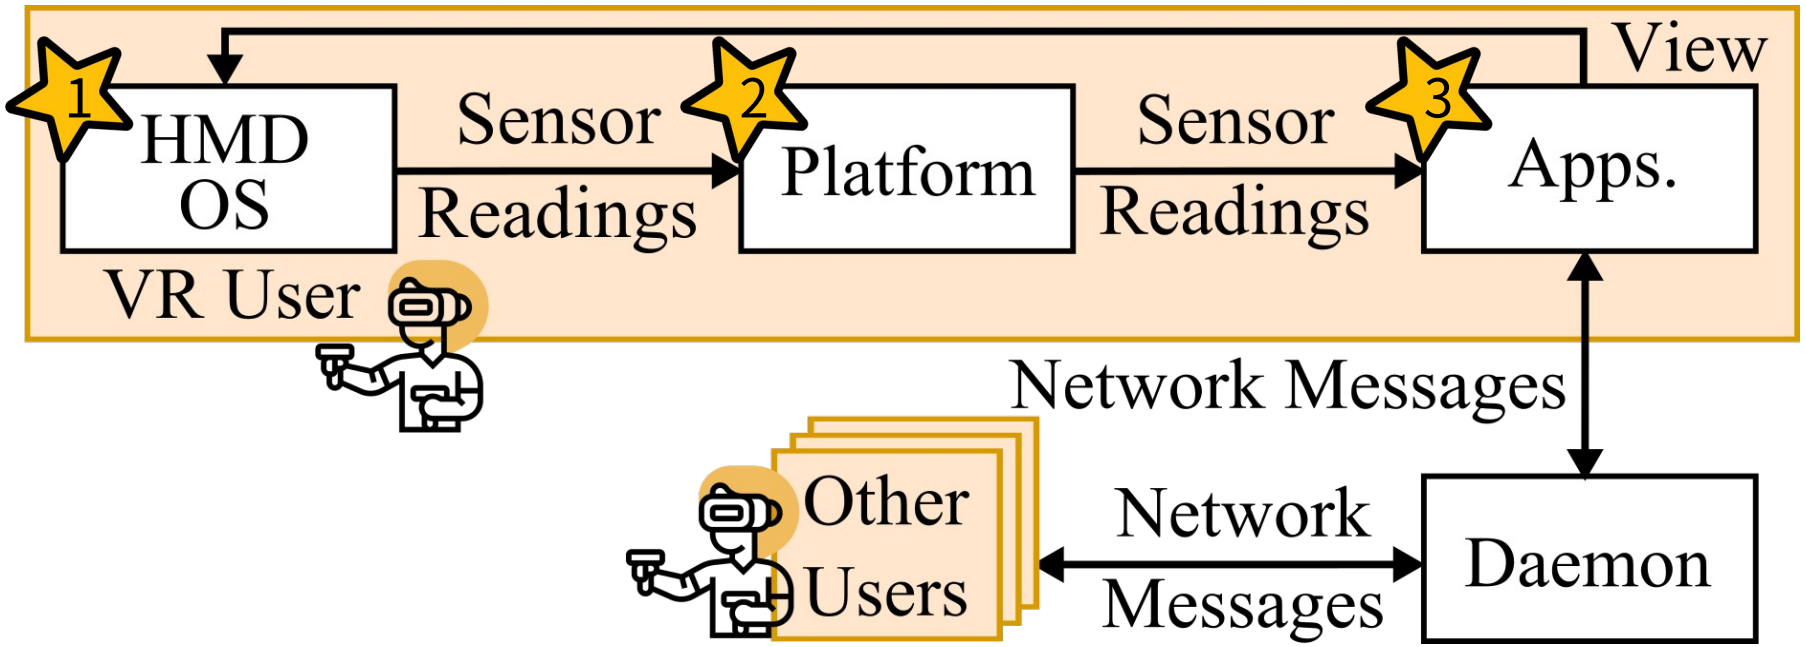

	Pros	Cons
HMD OS 1	• Provides the strongest protection • Supports all VR platforms and applications	• High engineering complexity
Platform 2	• Supports any VR applications on that platform • Provides medium engineering complexity	• Medium protection
Apps. 3	• Provides the lowest engineering complexity • Can incorporate application-specific optimization	• Less protection • Duplicated efforts for multiple VR applications

Privacy-Threat Mitigation System Overview

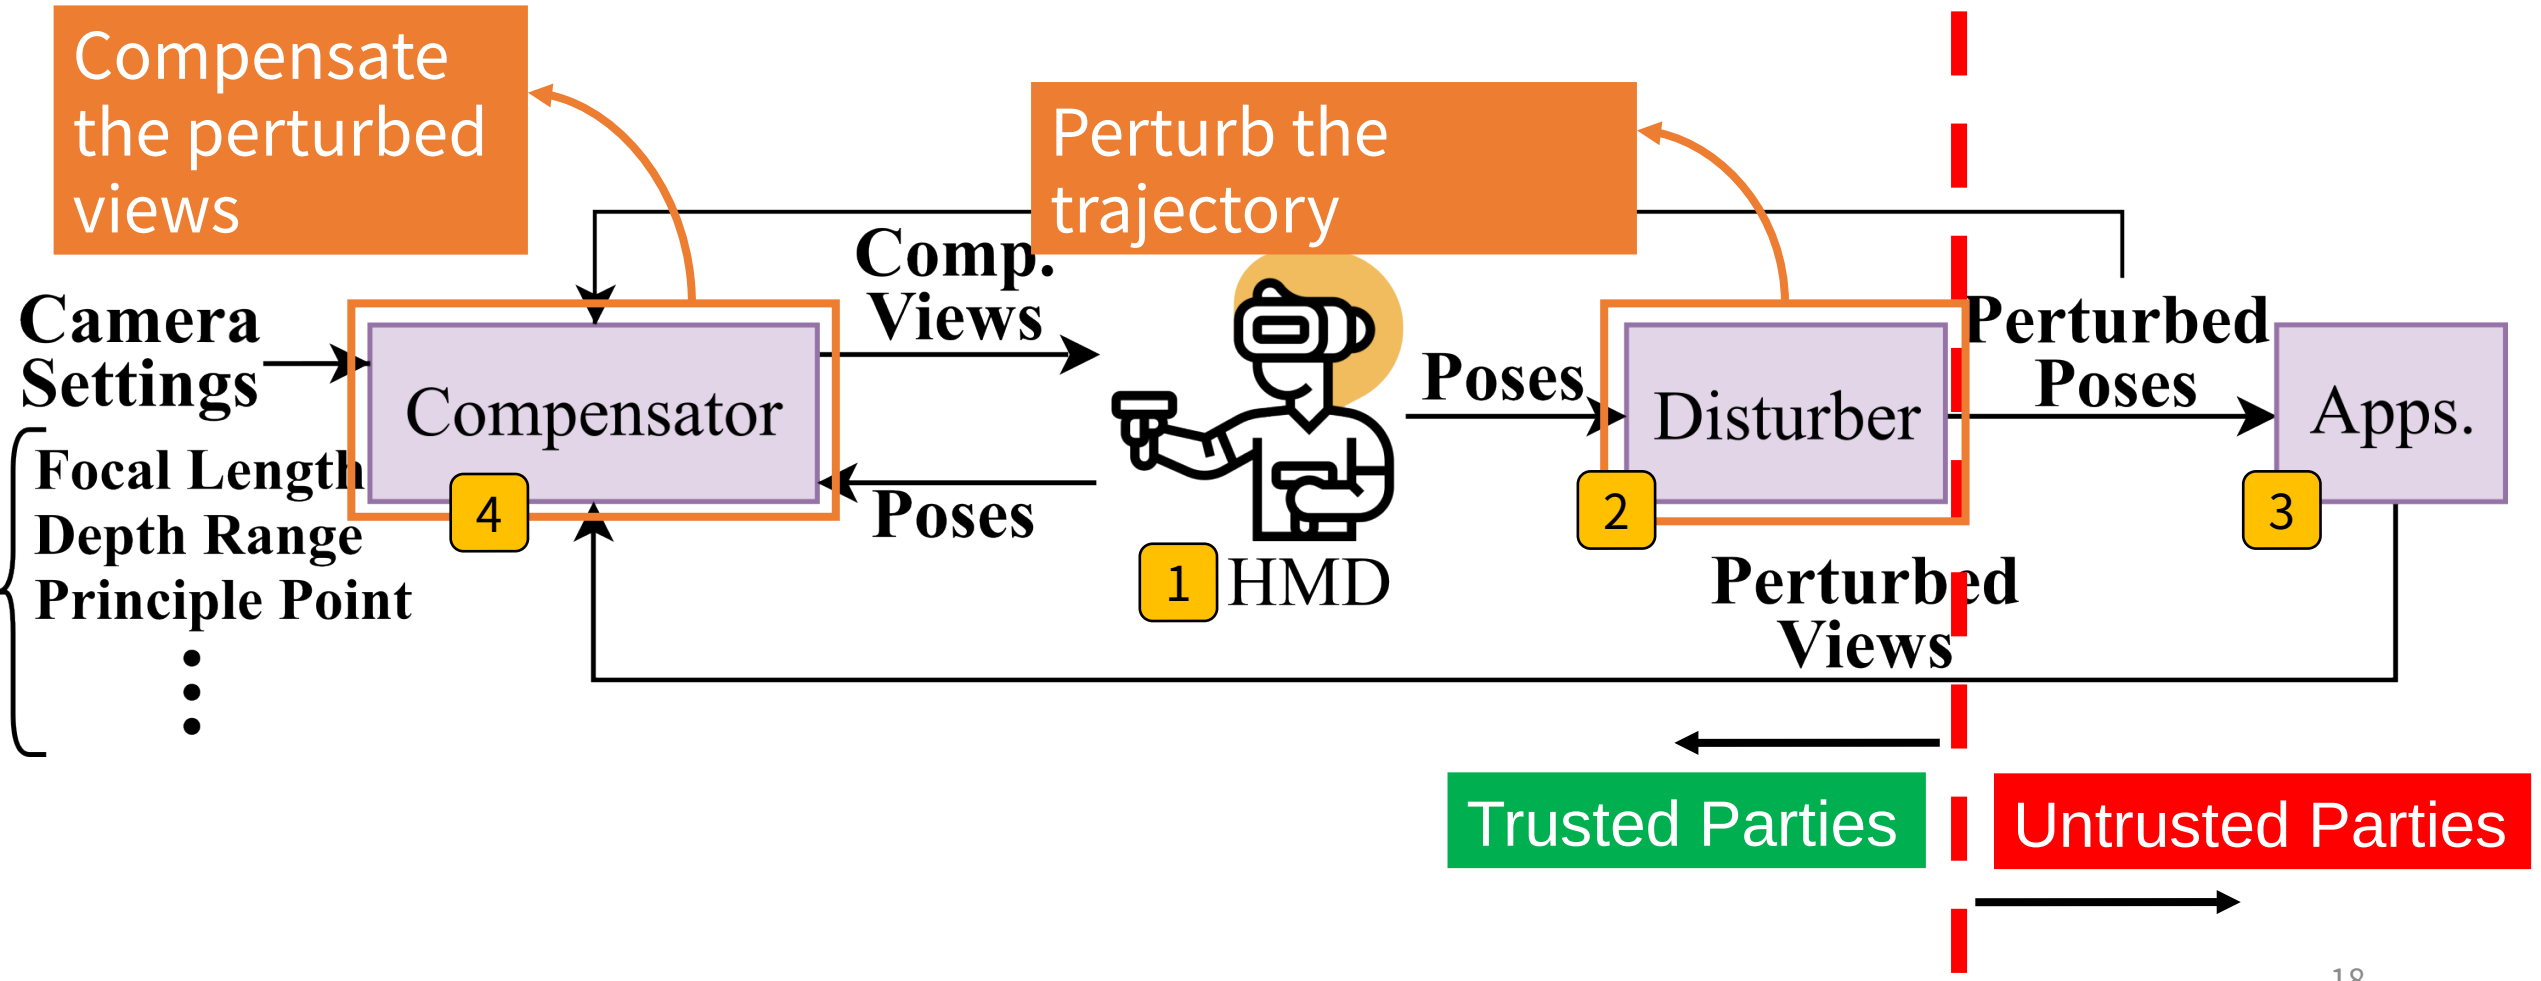

Design Objective

- Perturb a VR user's trajectory in both temporal (1, 2) and spatial (4) domains on-the-fly
- Find a good tradeoff (3) between the incurred perturbations and the degraded visual quality

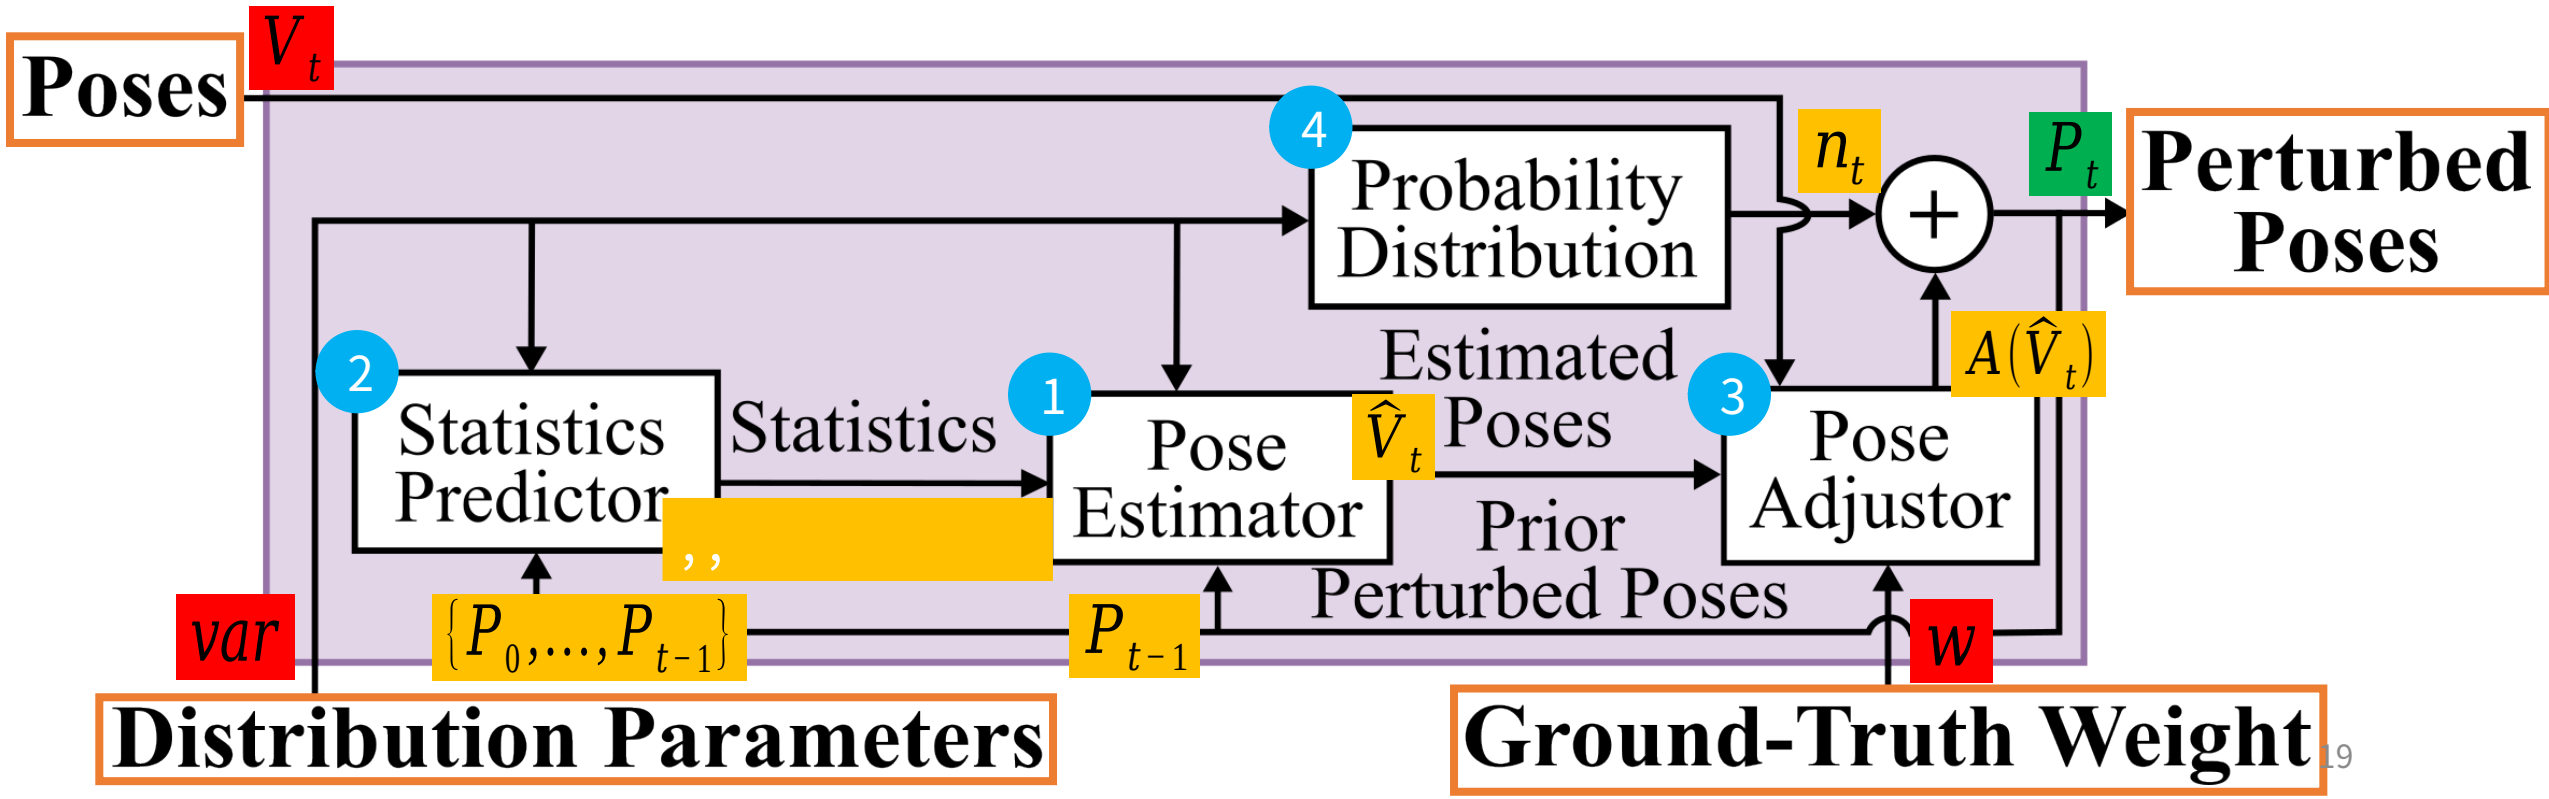

Disturber 2

Pose Estimation

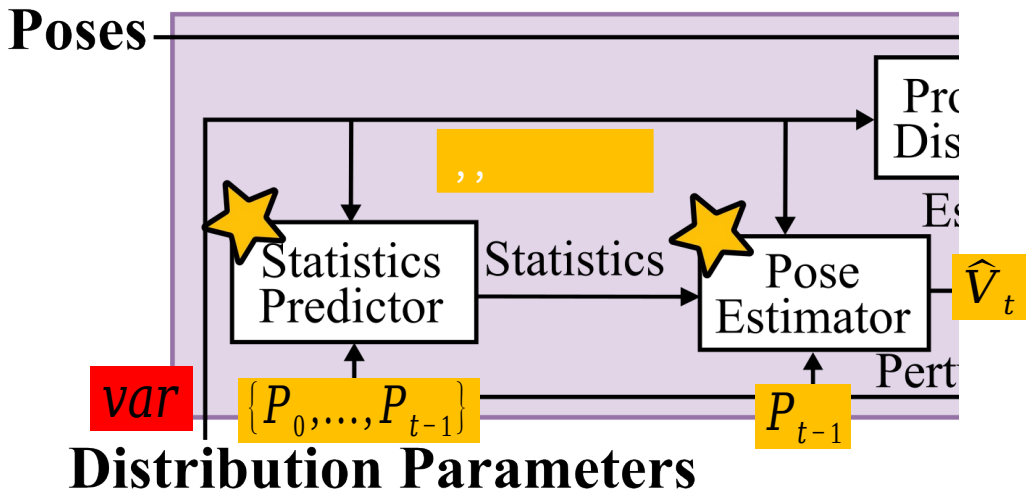

1 Modeling the trajectory with First-Order AutoRegressive Process [1]

- Time-series data modeling
- Linear model $V_{t-i}, i \in \{1,2,\cdots,p\}$
- previous data → current data (first-order:)

We use previous perturbed pose instead

2 Estimate poses with Linear Minimum Mean Squre Error (LMMSE)

In our case:

3 Predicted the statistics following Huitema and McKean [2]

$$\hat{V}_t = \hat{\mu}_{t-1}\left(1-\hat{\varrho}_{t-1}\frac{\hat{\sigma}^2_{t-1}}{\hat{\sigma}^2_{t-1}+var}\right)+\hat{\varrho}_{t-1}\frac{\hat{\sigma}^2_{t-1}}{\hat{\sigma}^2_{t-1}+var}P_{t-1}$$

Predicted statistics

Prior perturbed pose

Probability distribution variance

[1] X. Zhang, M. Khalili, and M. Liu. Differentially private real-time release of sequential data. ACM Transactions on Privacy and Security, 2022
[2] B. Huitema and J. McKean. Autocorrelation estimation and inference with small samples. Psychological Bulletin, 1991

Disturber 2

Estimated Pose Adjustment

- Weighted sum

$$A(\hat{V}_t) = (1 - w)\hat{V}_t + w V_t$$

Ground-truth weight

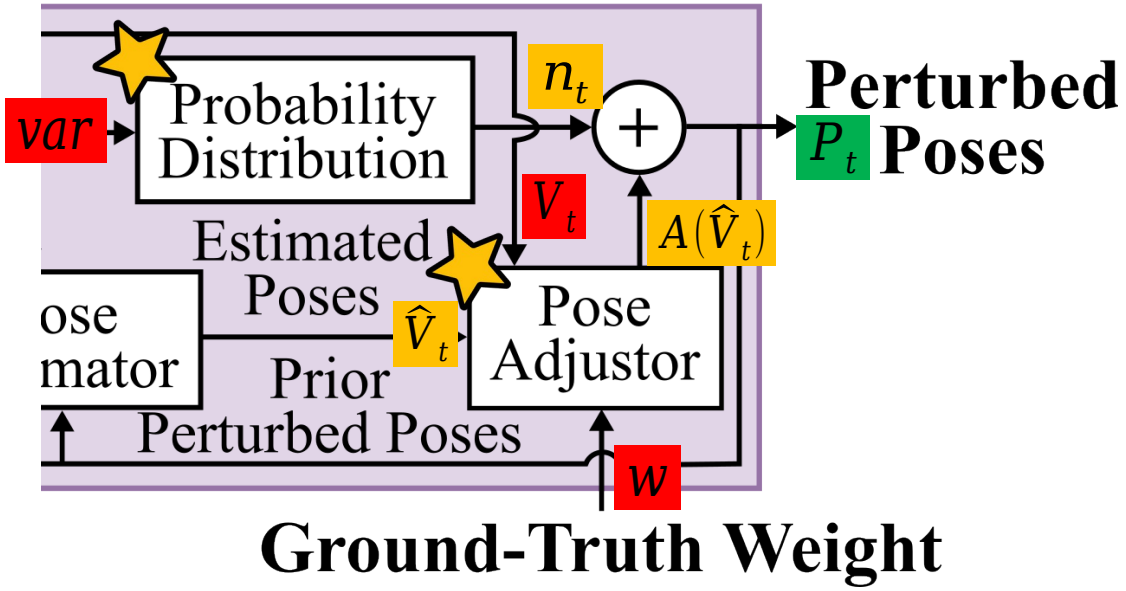

→ End of the temporal perturbation

Random Perturbations from Probability Distribution

A privacy framework to quantify the amount of privacy

- - Laplace mechanism
  - Laplace distribution

$$P_t = A(\hat{V}_t) + n_t \quad n_t \sim \mathcal{L}(0, var)$$

→ End of the spatial perturbation

Compensator 4 Design Objective

- Warp each rendered perturbed RGBD image to an RGB image viewed at the (original) pose with
  - short execution time
  - high visual quality

→ Find a good View Synthesizer

- Neural-network-based view synthesizers
  → large running time
- **Depth** Image-Based Rendering (DIBR)-based view synthesizers → **Reference View Synthesizer (RVS) [1]**

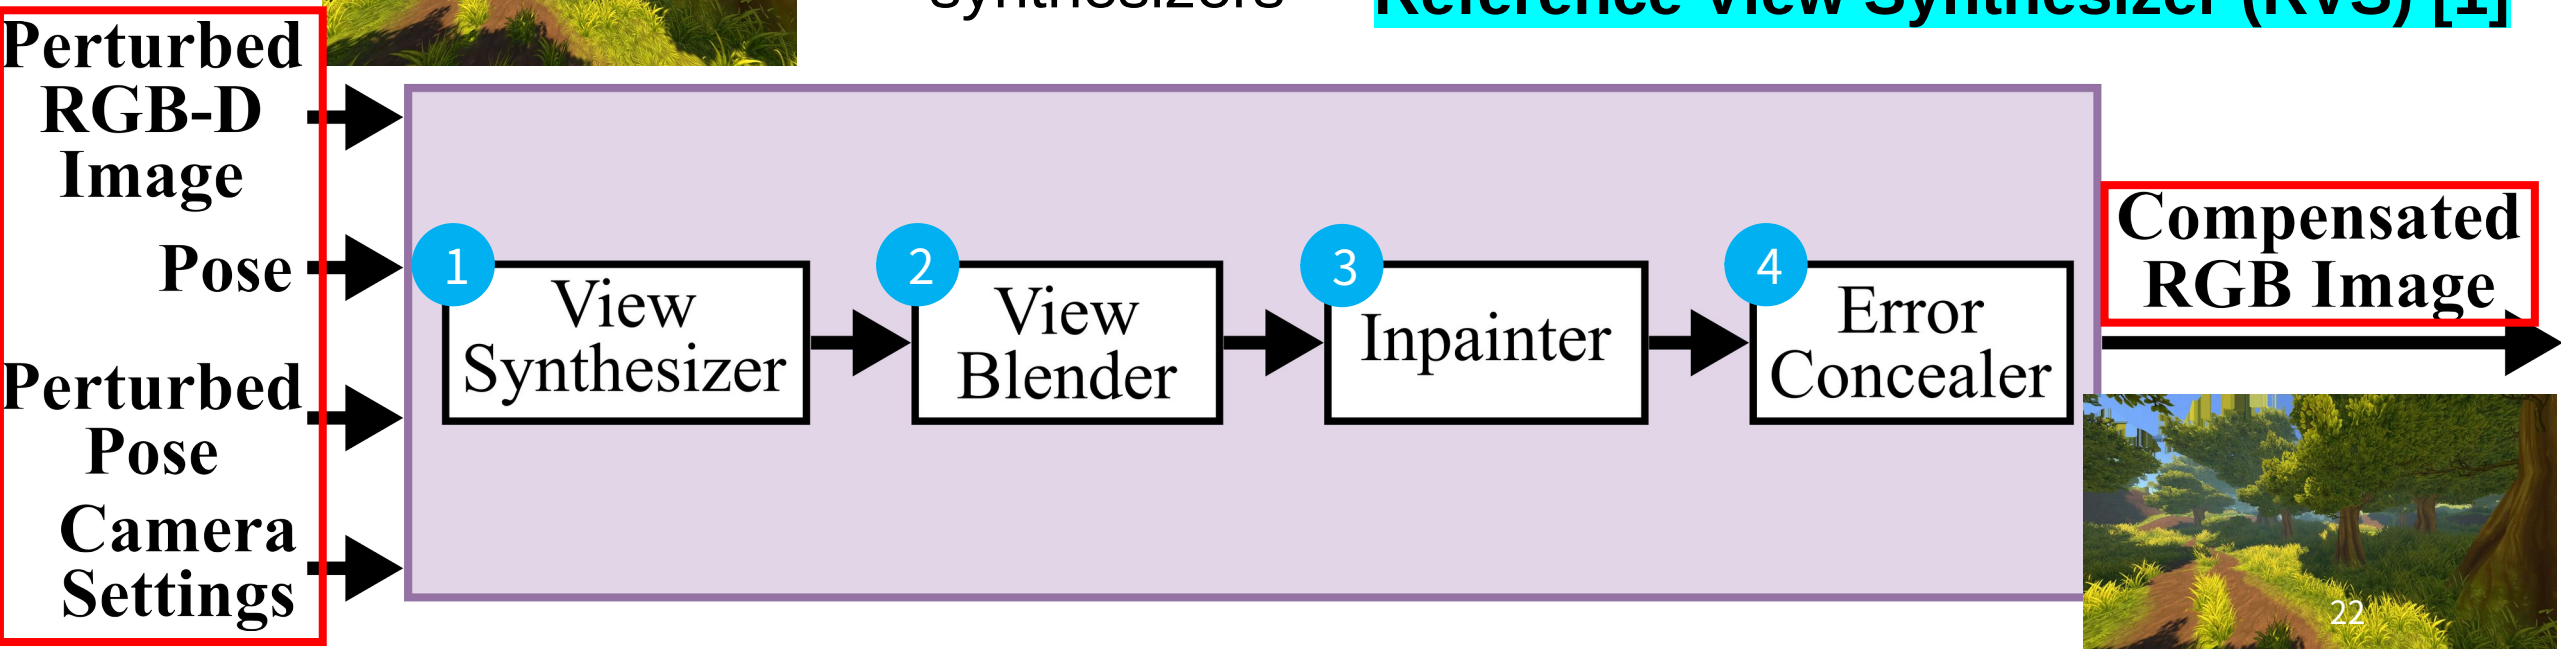

[1] B. Kroon and G. Lafruit. Reference view synthesizer (RVS) 2.0 manual, 2018.

Compensator 4

Relation Between Perturbed and Compensated Viewport

- Perturbed viewport Compensated viewport
- The rendered images' resolution () are the same
- Two key camera settings: **Perturbed FoV** and **Compensated FoV**

θ_p θ_c

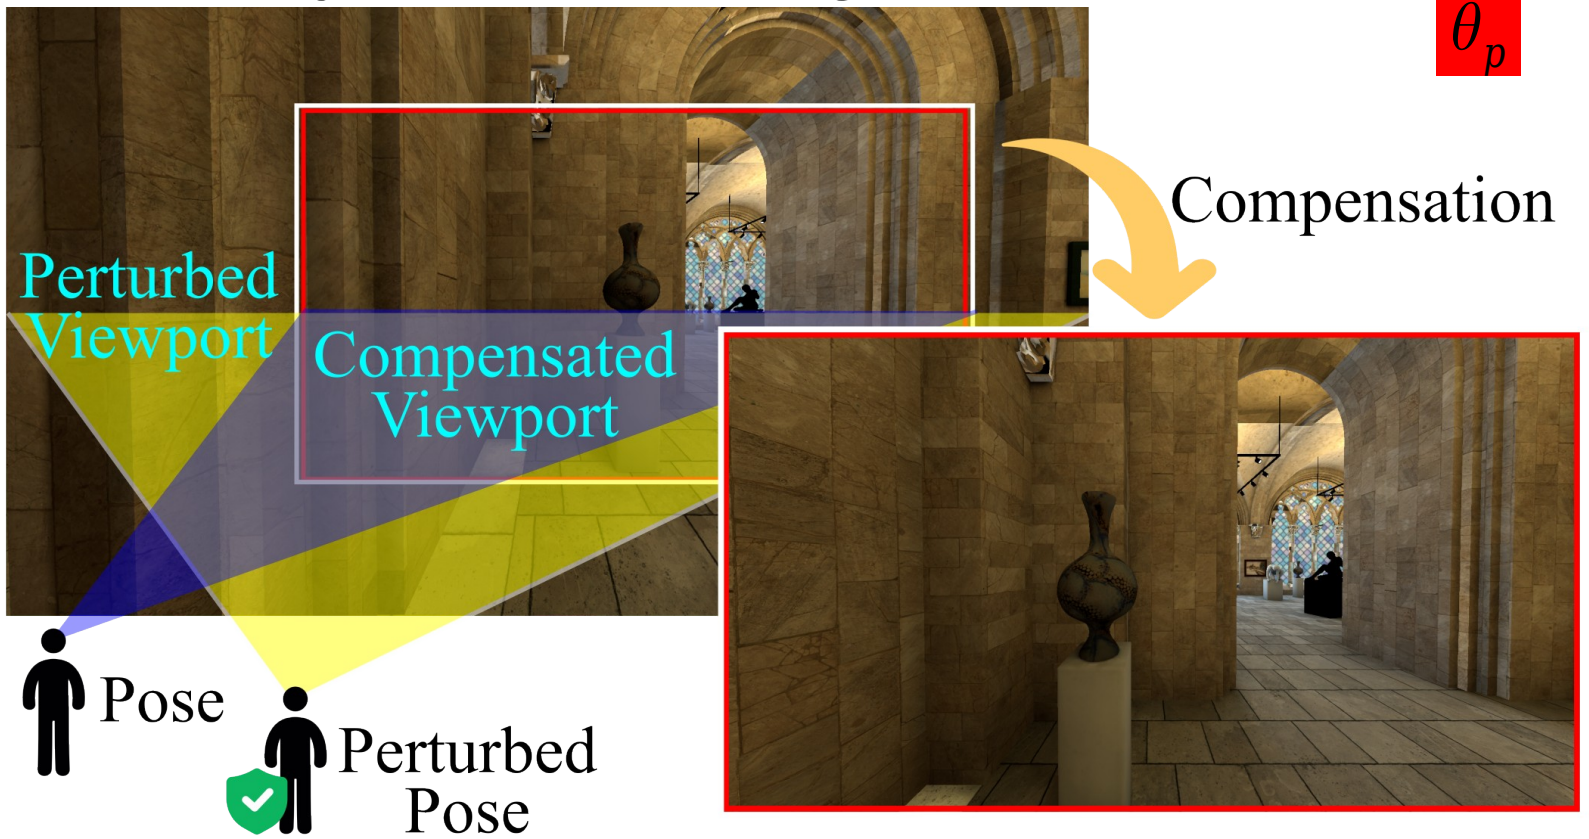

- Modify the **vertical FoV**
- Calculate the **horizontal FoV** with the vertical FoV and the aspect ratio
- Calculate the **Focal length** with the two FoV

Outline

- Introduction
- Related Work
- 6DoF VR Dataset
- Privacy Threats Mitigation
- Evaluations
- Conclusion & Future Work

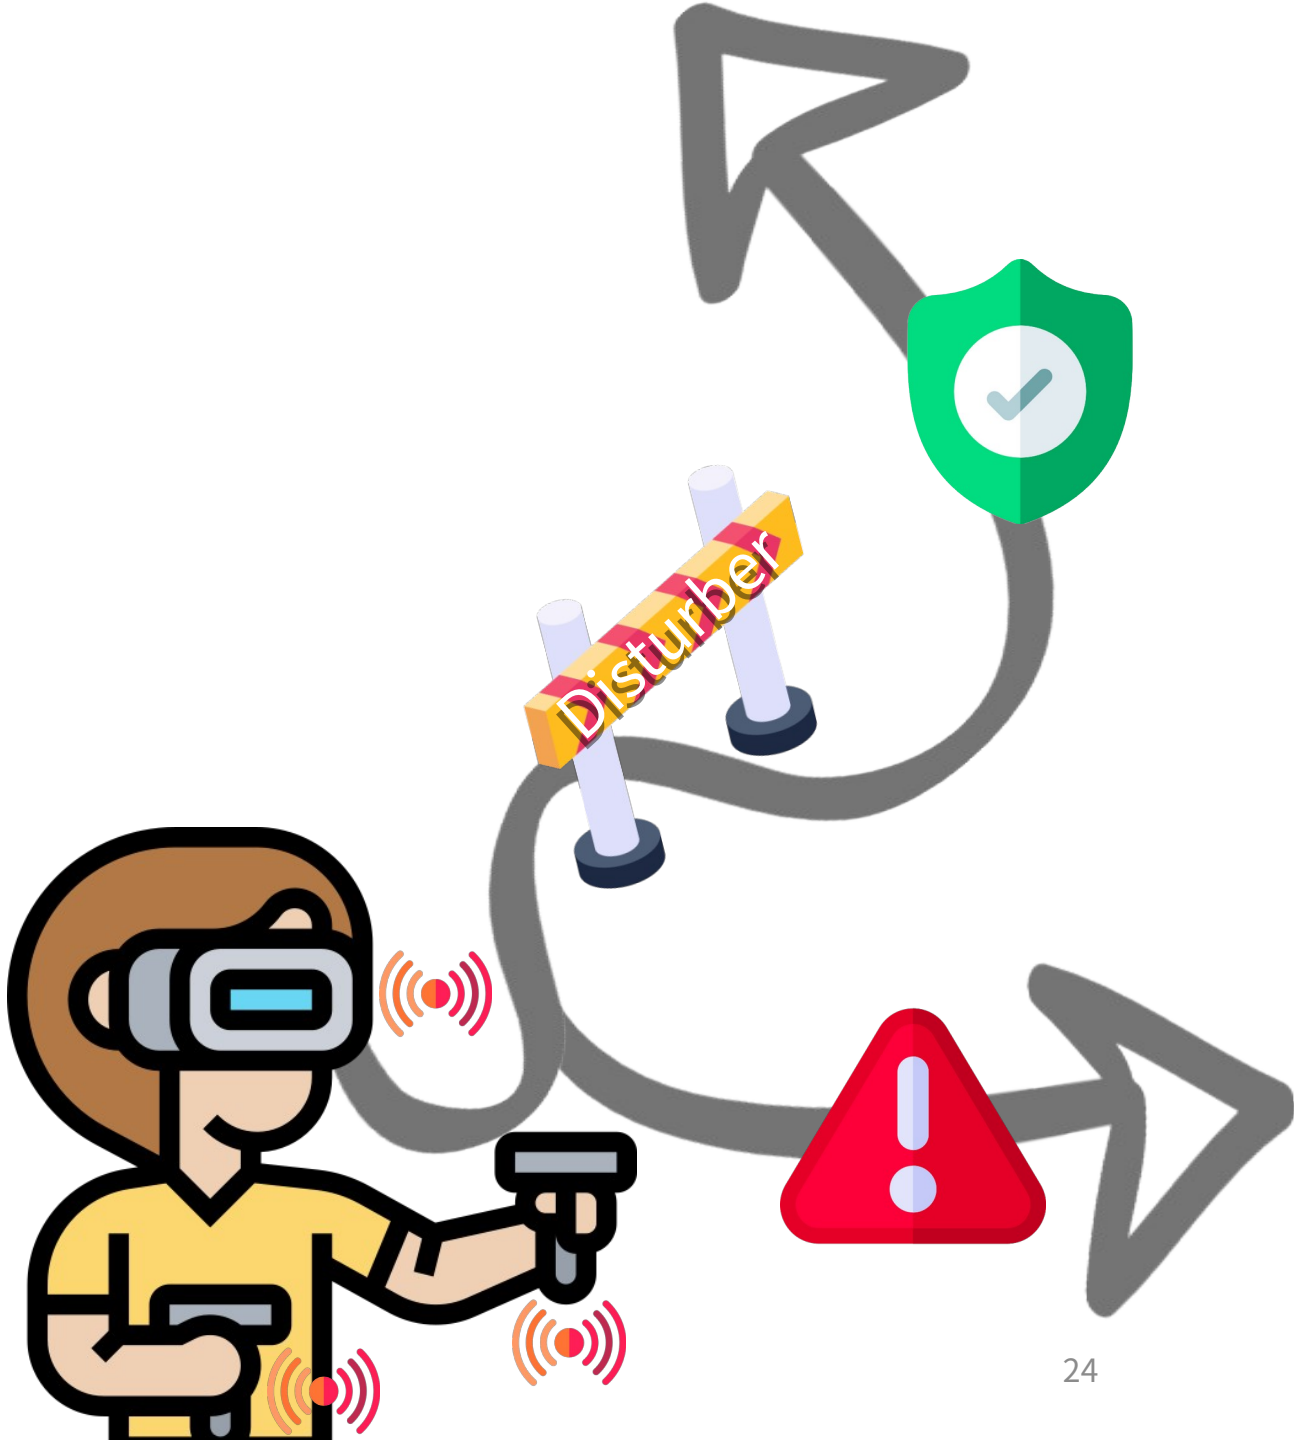

Compared Algorithms

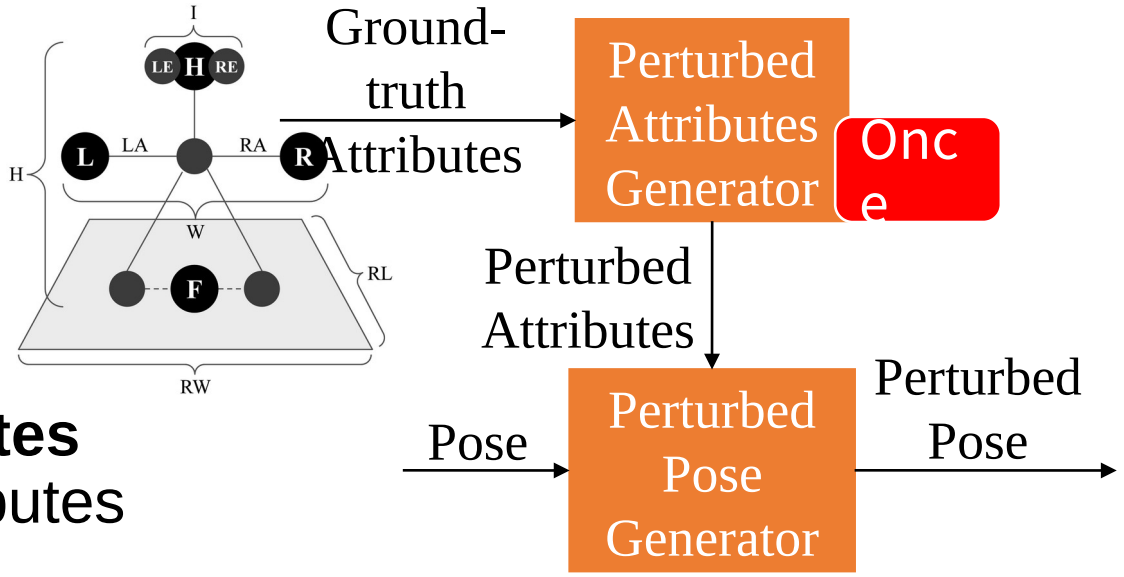

- MetaGuard (MG) [UCB]
  - Add perturbations to multiple **attributes** and then projects the perturbed attributes back to perturbed locations

-
  - Bounded Laplace Mechanism

Height, Arm length, Room size, etc.

$$\frac{Height'}{Height} = \frac{y'}{y}$$

Related to location

  - Selected attributes:

- Disturber Only (DO)
- Disturber with Compensator (DC)

[UCB] V. Nair, M. Garrido, and D. Song. Going incognito in the metaverse. arXiv:2208.05604, 2022

Evaluation Setup

- Testbed

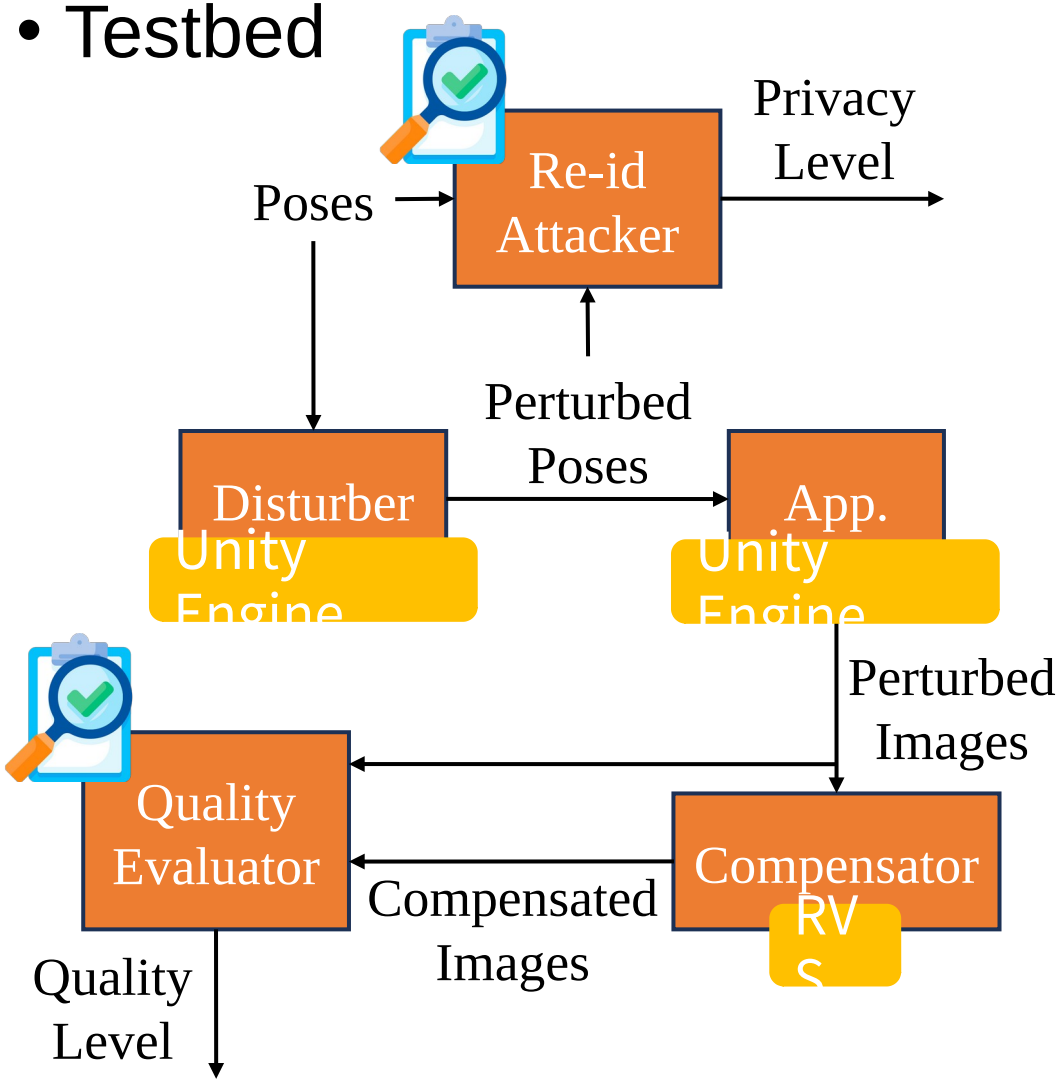

- Dataset
  - 6DoF VR dataset [1]
- Varied parameters
  - MetaGuard:
  - Disturber Only:
  - Disturber with Compensator:

[1] Y. Wei, X. Wei, S. Zheng, C. Hsu, and C. Yang. A 6DoF VR dataset of 3D virtualworld for privacy-preserving approach and utility-privacy tradeoff. In Proc. of ACM Multimedia Systems (MMSys), Vancouver, Canada, June 2023

Evaluation Metrics

Privacy

- Re-identification (Re-id) attack
  - Trajectory
    - 2000 poses
    - 50-poses sliding window (1s)
  - Random Forest
    - Maximal tree depth:
    - Number of trees:
  - Features
    - Velocity and angular velocity of each HMD and the controllers
    - Min, max, and average distances between each HMD and the controllers
    - Min, max, and average locations/orientation of each HMD and the controllers
  - Train-test split settings:
    - 5-fold cross-validation
    - Split each subject's trajectories into five folds
    - Each run picks one of the segments as the testing set

Visual Quality

- Trajectory
  - 2000 poses
  - Selecting sample users with diverse moving
- Metrics
  - Peak Signal-to-Noise Ratio (**PSNR**)
  - Structural Similarity (**SSIM**)
  - Video Multimethod Assessment Fusion (**VMAF**)

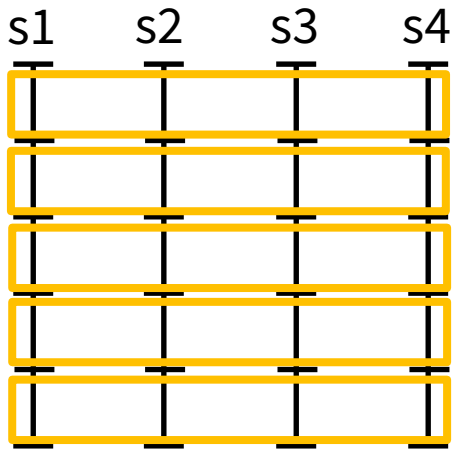

User Study

- 8 subjects
- Nature (outdoor) and Office (indoor) scene
- Compared algorithms: DO and DC
- Questions

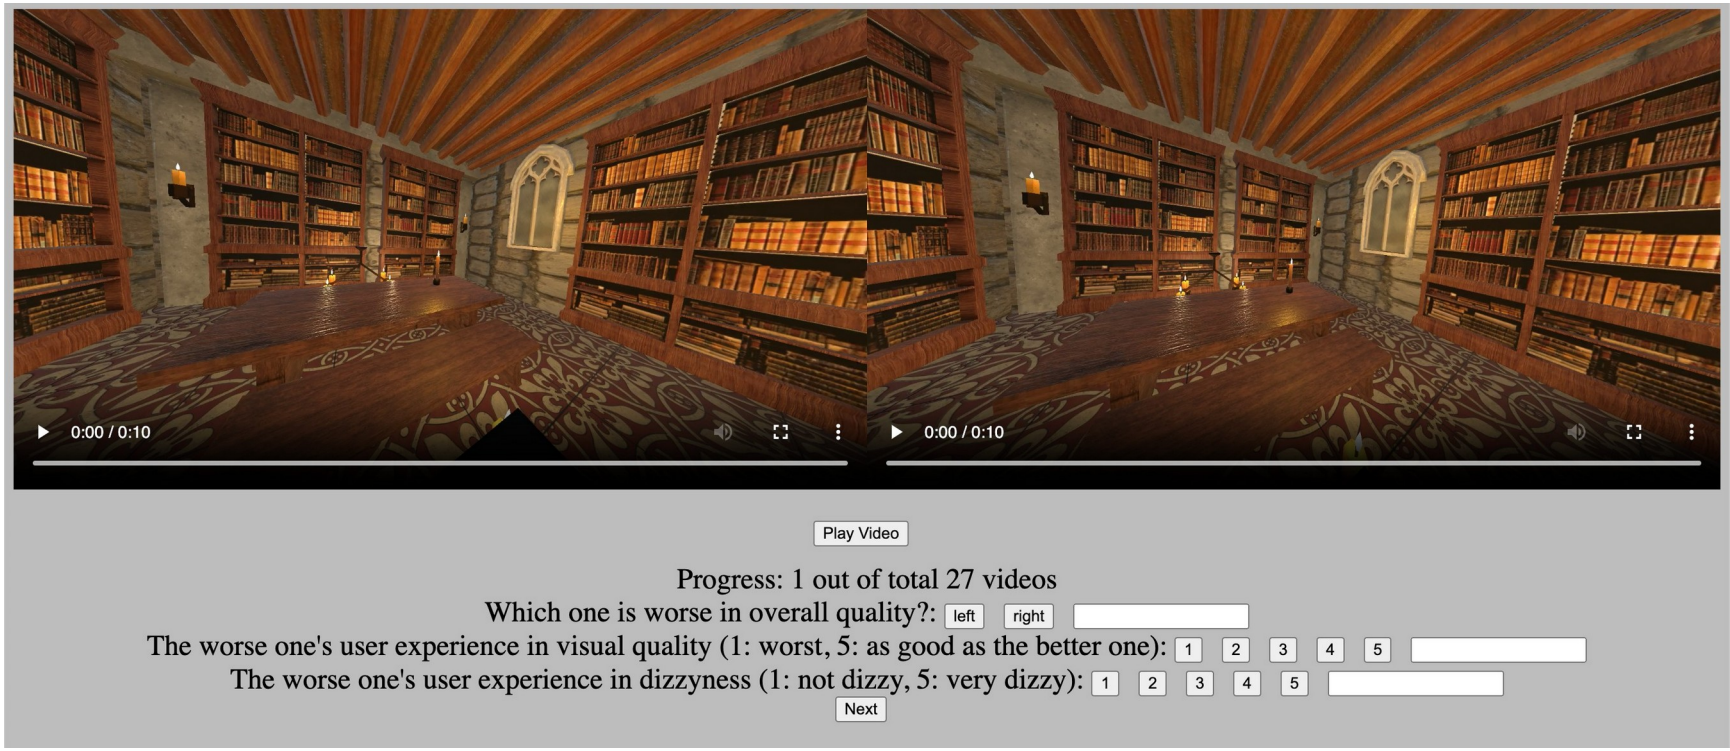

QoE Questions
Which viewport is worse in overall quality?
How would you rate the worse viewport's user experience in visual quality?
How is the dizziness when you watch the worse viewport?

Privacy Protection Level of DO and MG Under Different

- DO ⬇ the re-id rate by up to 0.4

Perturbing the trajectories in the temporal domain preserves more user privacy

- When approaches 0, the re-id rate of DO approaches 0.1 while MG is still above 0.45

Bounded Laplace

DO protects user privacy more efficiently

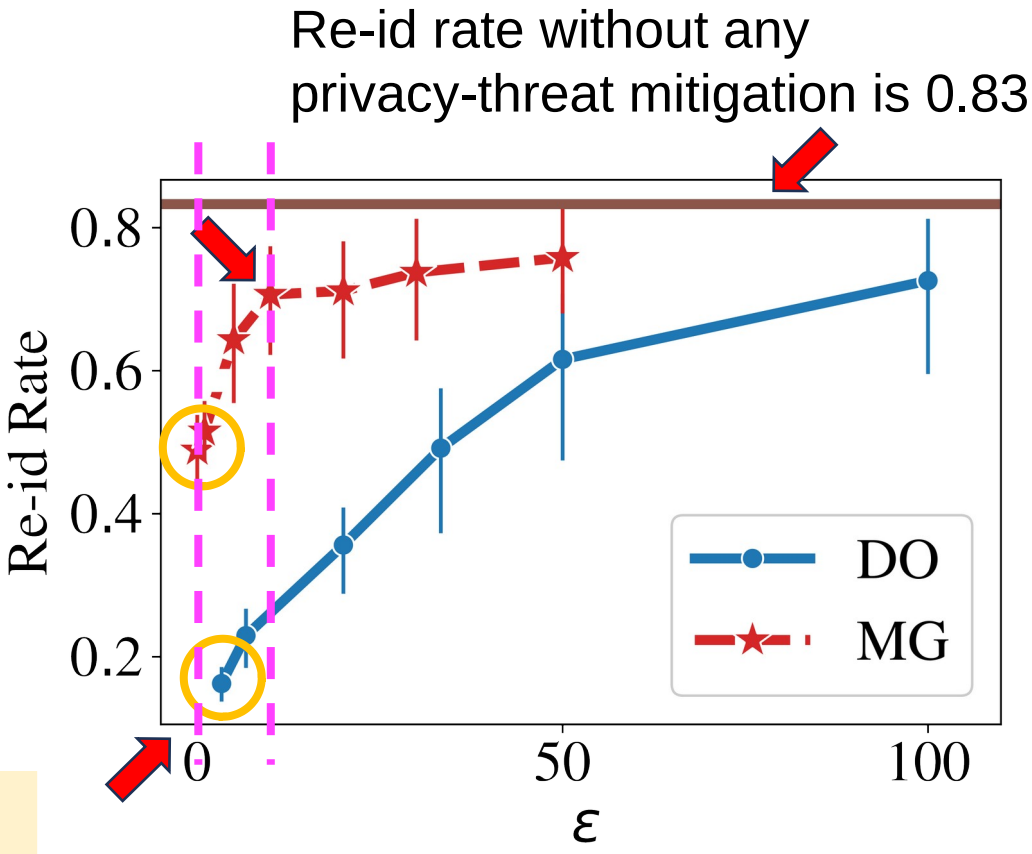

Visual Quality Improvement by DC Under Different

- DC ⇧ the visual quality by 6.8 dB at most and 5.9 dB in PSNR on average

DC alleviates the degradation of visual quality due to perturbation successfully

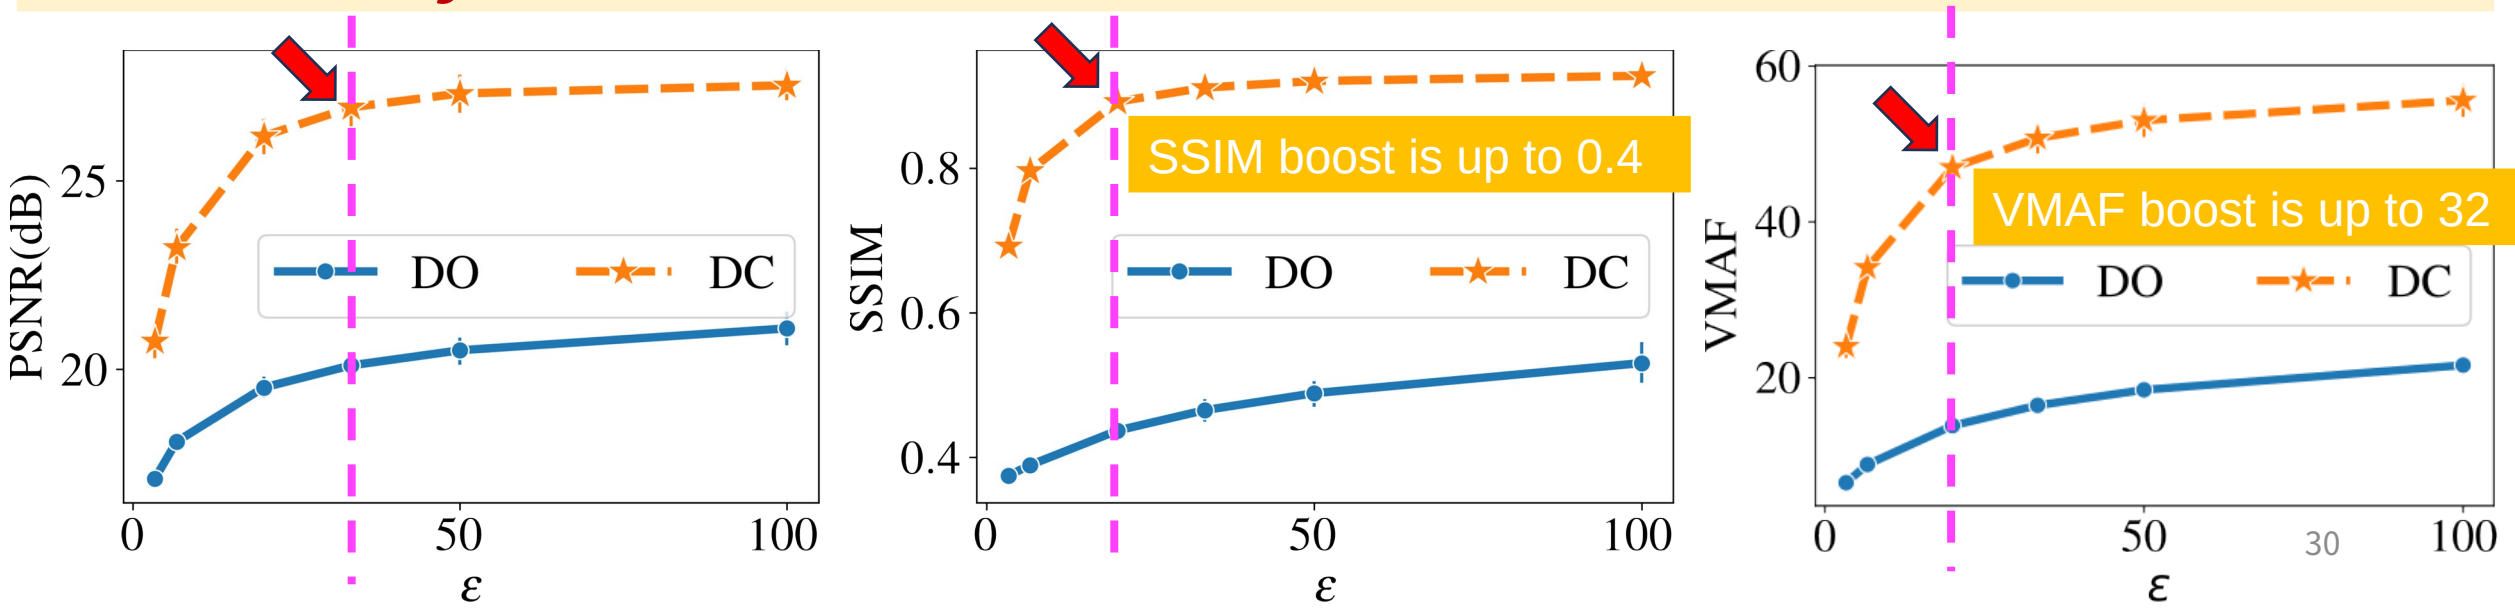

Privacy-Quality Tradeoff

- DO achieves better visual quality than MG when the re-id rate is between 0.10 and 0.73
- DO ⇓ the re-id rate by almost half compared to MG under the same visual quality
- DC ⇑ by up to 6.83 dB in PSNR at the same re-id rate compared to DO

Our solution provides strong protection while delivering good visual quality

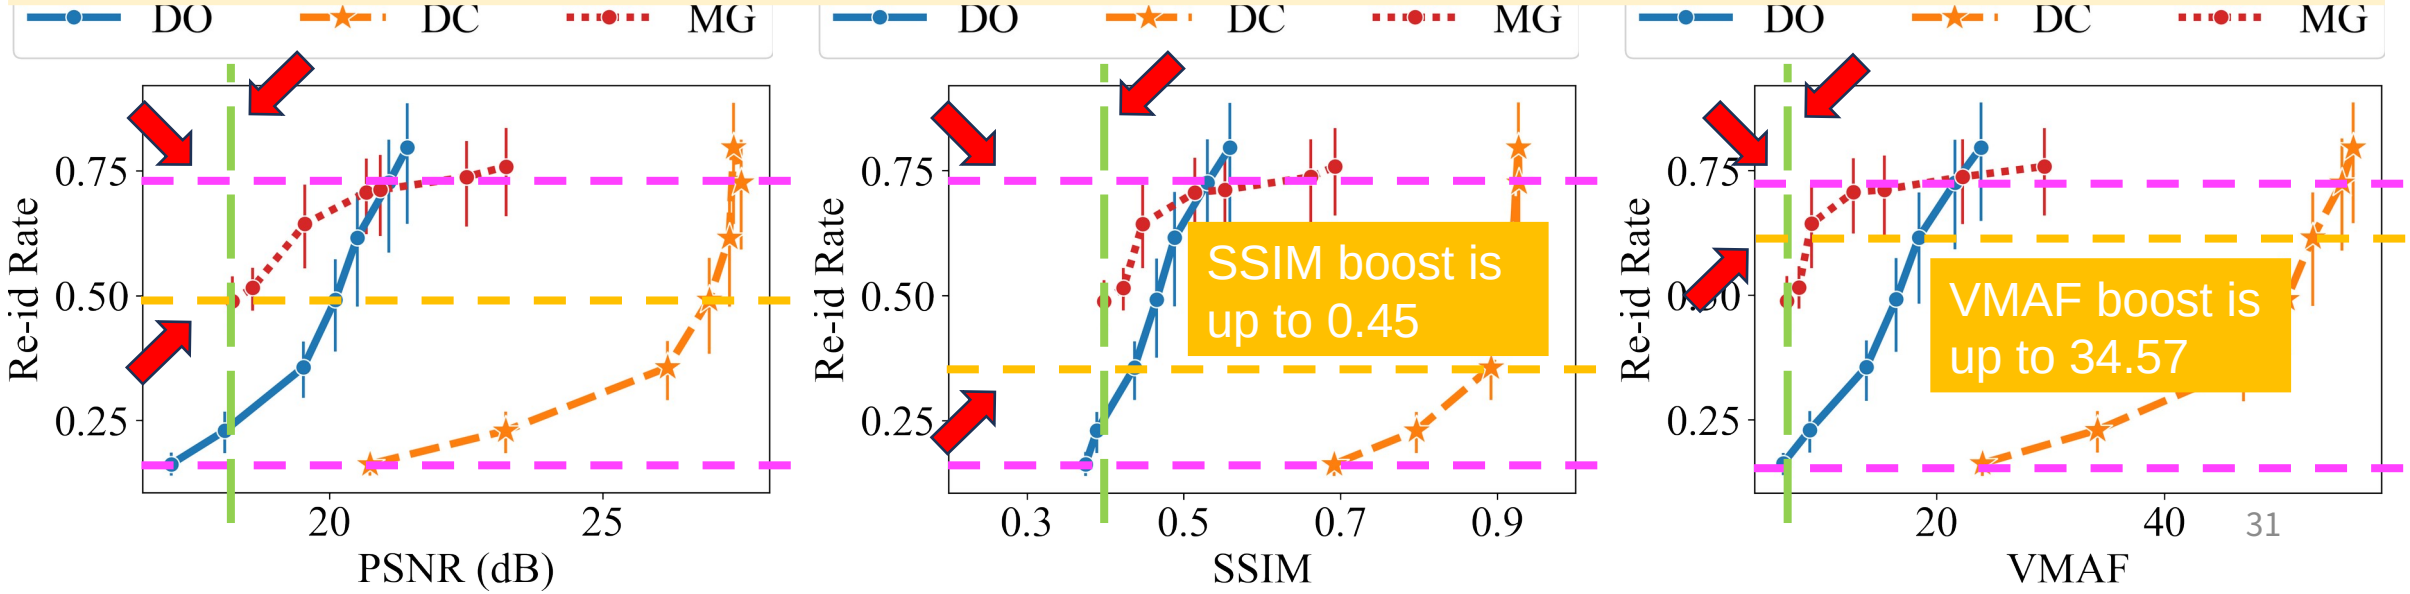

User Experience

- Visual quality and dizziness scores ⇧ with DC overall
- At most one subject reports a worse score with DC

Our compensator successfully improves user experience

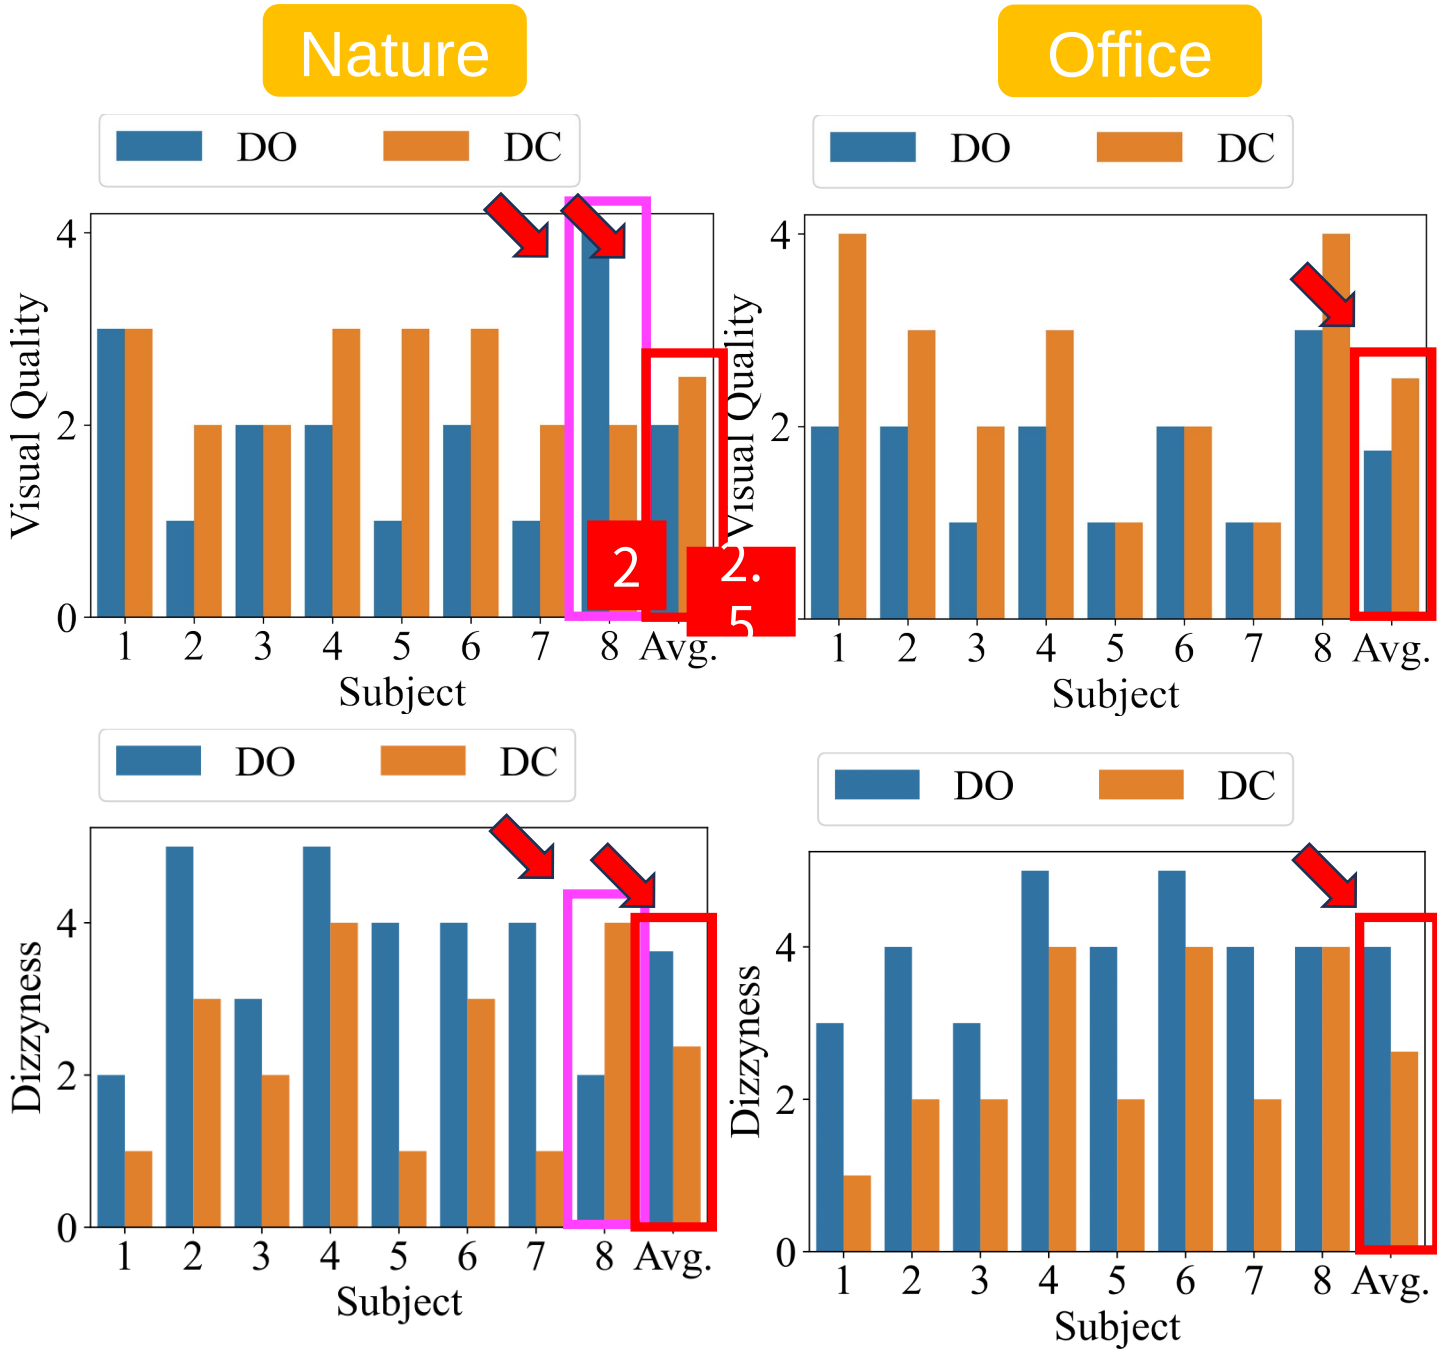

Outline

- Introduction
- Related Work
- 6DoF VR Dataset
- Privacy Threats Mitigation
- Evaluations
- Conclusion & Future Work

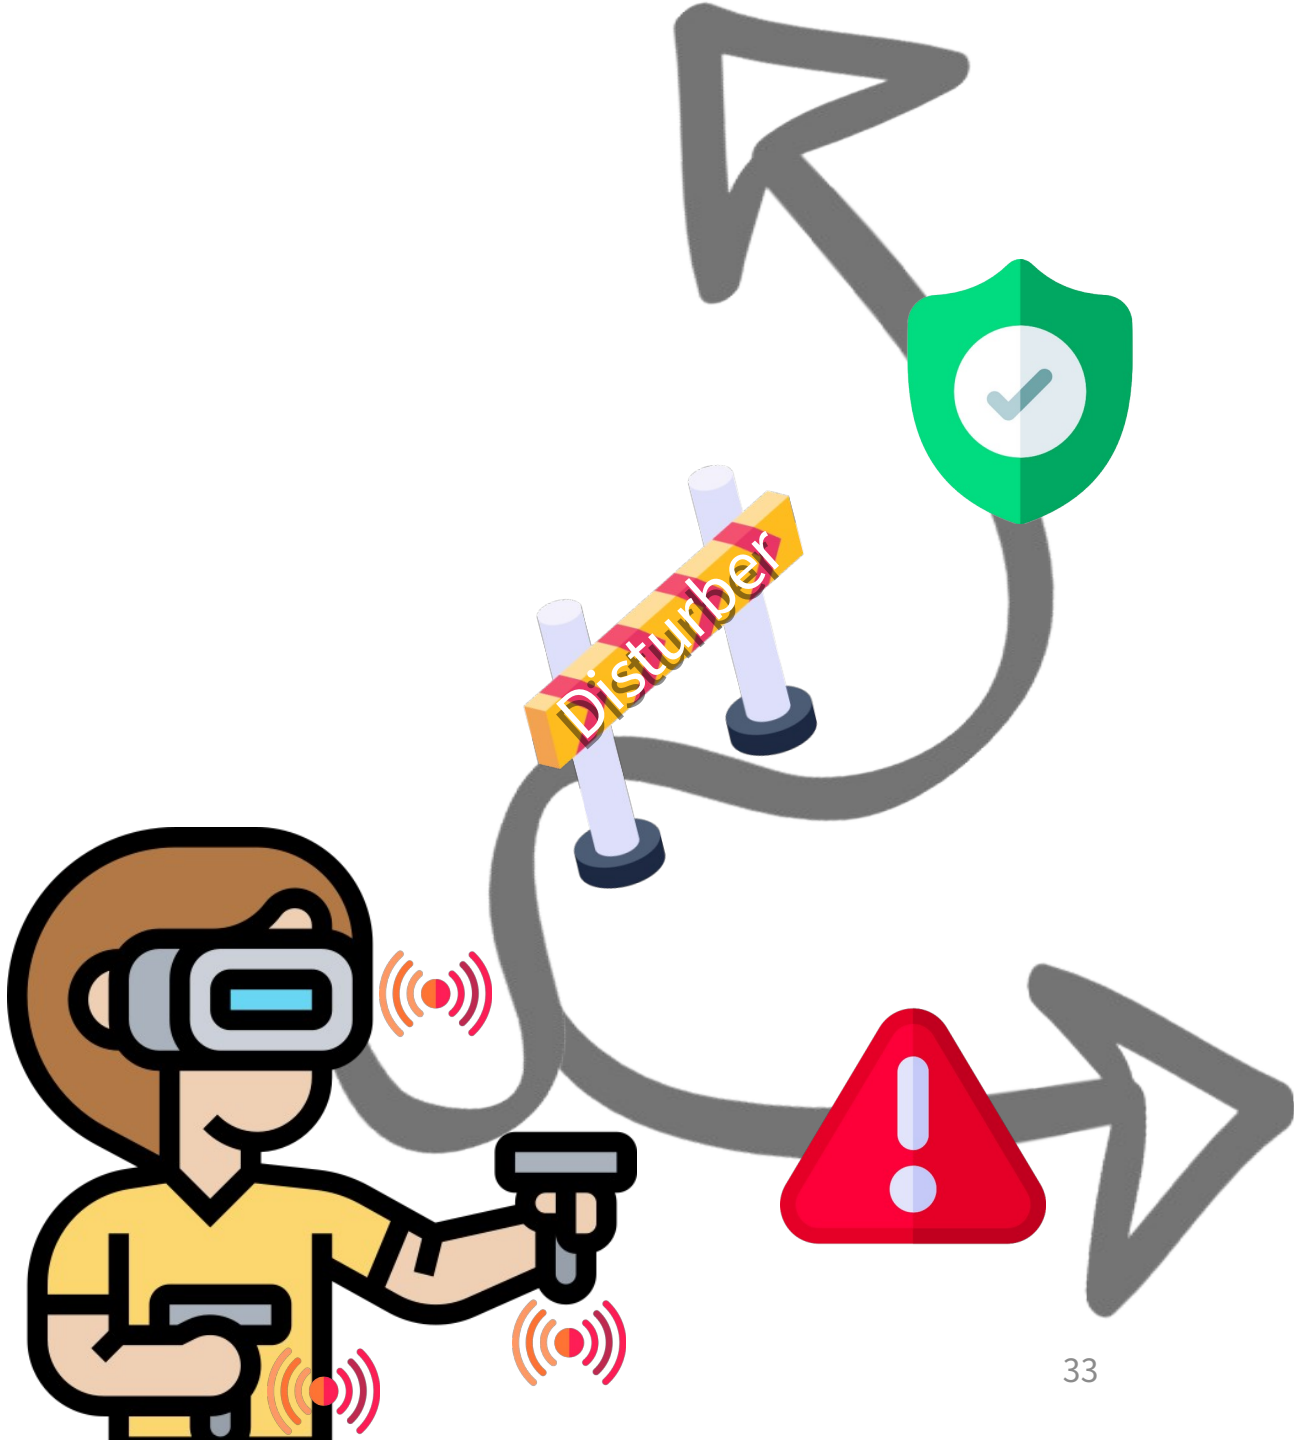

Conclusion

- Collect and release a 3D-virtual-world 6DoF VR dataset to study the privacy issues in VR
- Develop a privacy-preserving approach to mitigate privacy threats on-the-fly while retaining the visual quality
- DO reduces at most 0.4 re-id rate compared to MG under the same
- DC further improves the visual quality by at most 6.83 dB in PSNR, 0.45 in SSIM, and 34.57 in VMAF
- DC successfully improves user experience

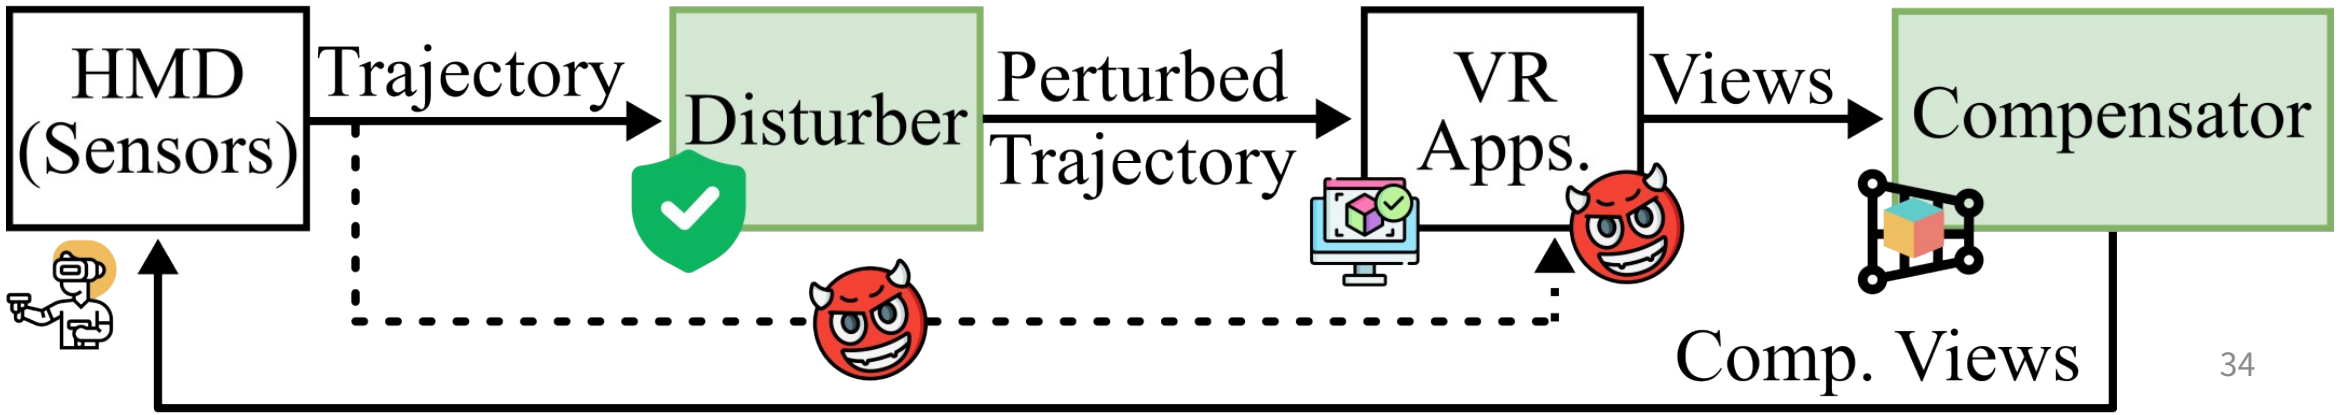

Future Work

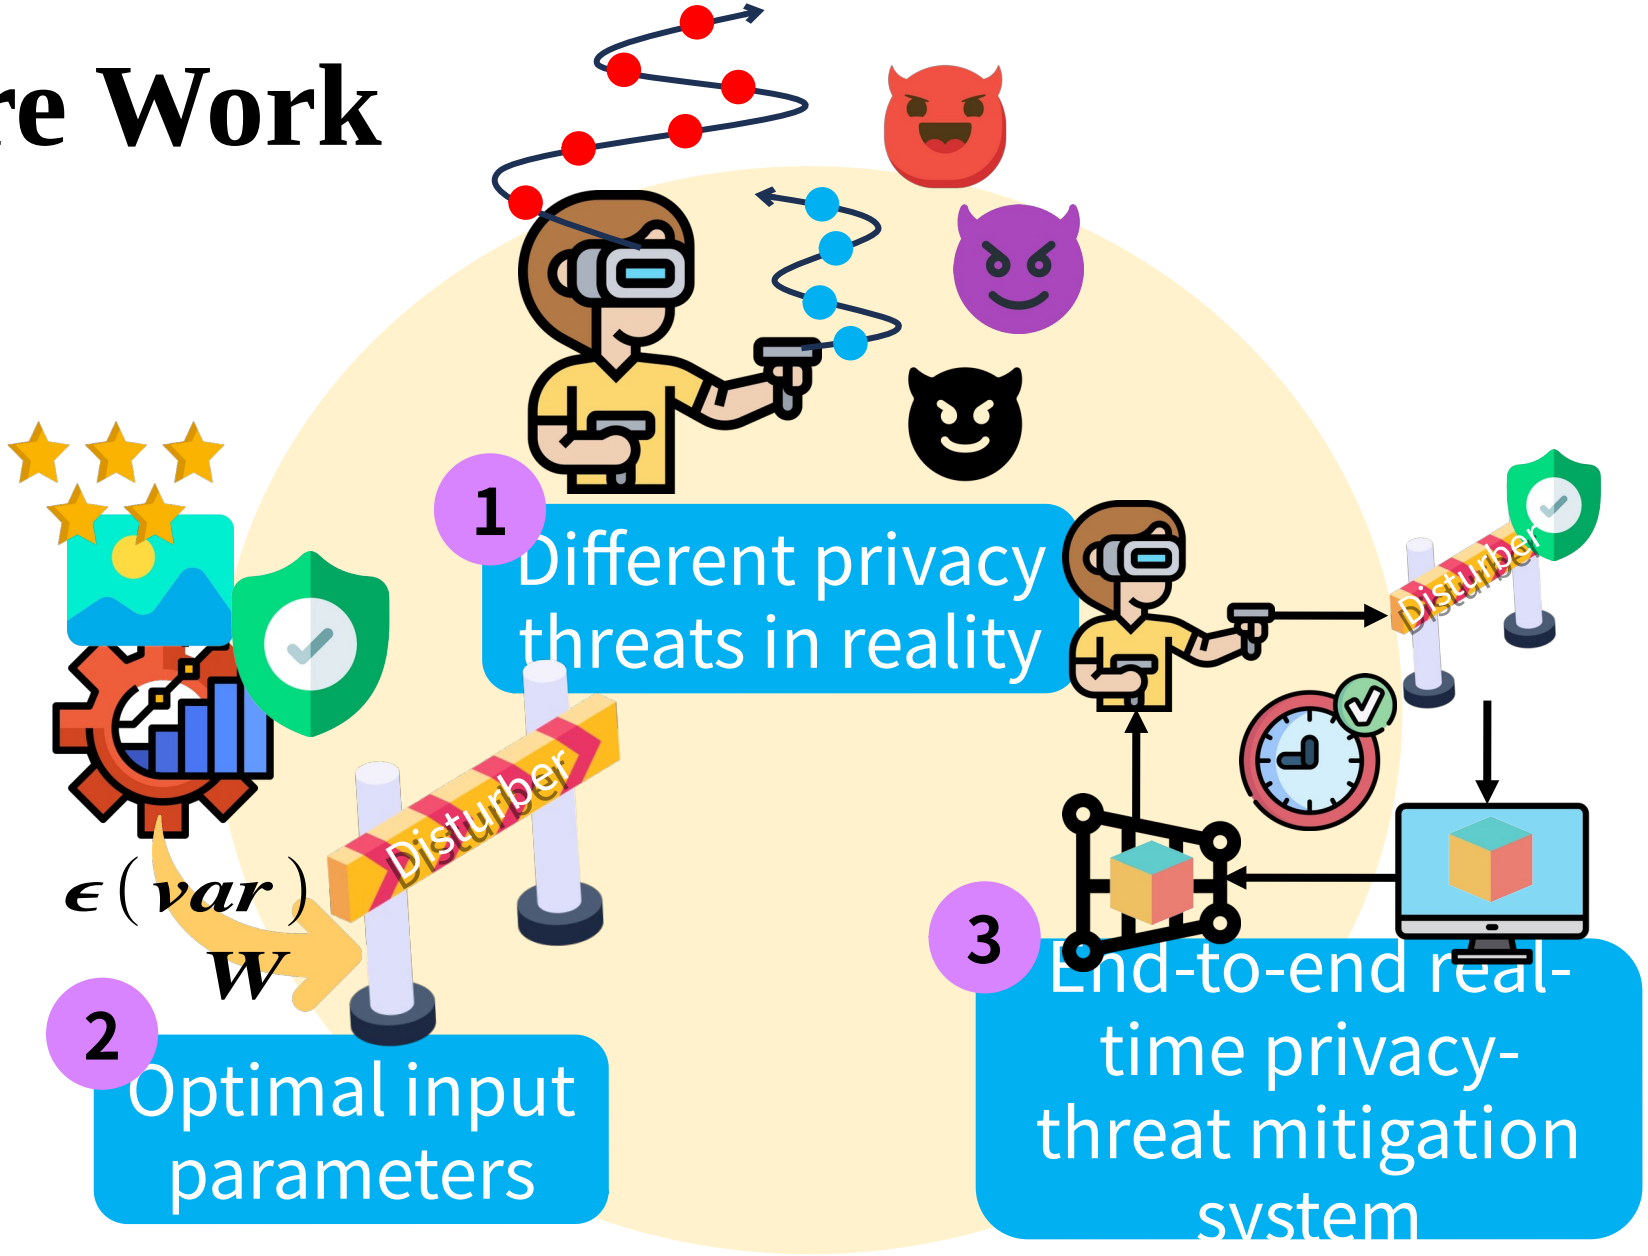
1
Different privacy threats in reality
Disturber
$\epsilon(var)$
W
2
Optimal input parameters
Disturber
3
End-to-end real-time privacy-threat mitigation system

Thank you for listening !

YU-SZU WEI (Email: weiyousz0328@gmail.com)
Thanks for the help of Prof. Hsu, Prof. Huang, Prof. Yang
Shin-Yi Zheng, Yuan-Chun Sun, Xing Wei, and all lab mates.

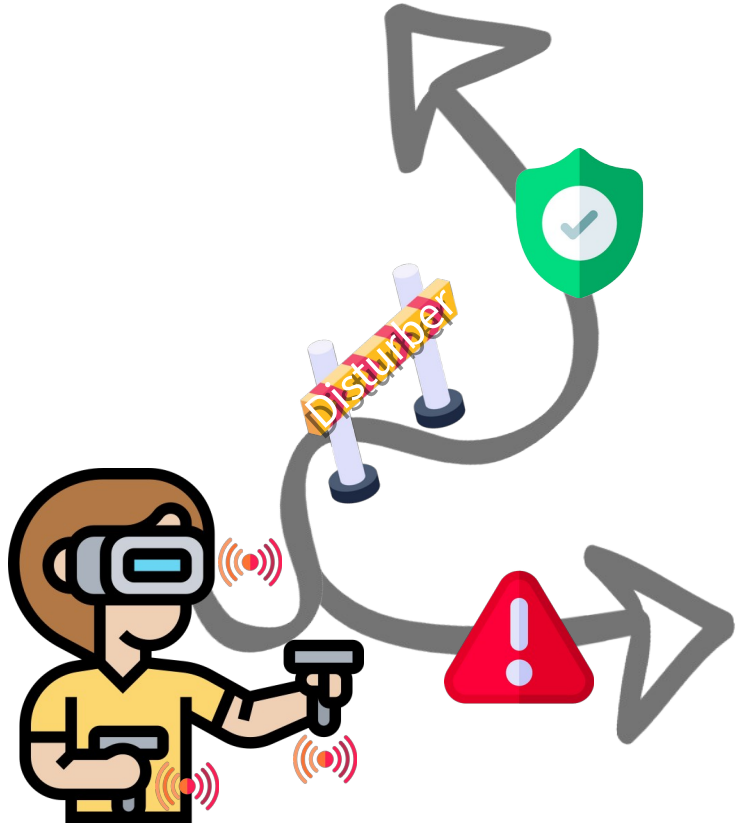

Publications:

- **Y. Wei**, X. Wei, S. Zheng, C. Hsu, and C. Yang. 2023. A 6DoF VR Dataset of 3D VirtualWorld for Privacy-Preserving Approach and Utility-Privacy Tradeoff. In Proc. of ACM Multimedia Systems (MMSys). Vancouver Canada, 444–450.
- **Y. Wei**, S. Zheng, Y. Sun, C. Huang, and C. Hsu. “Mitigating Privacy Threats Without Degrading Visual Quality of VR Applications”, in Proc. of ACM International Conference on Multimedia in Asia (MMAsia), December 6–8, 2023, Tainan, Taiwan. (**Under review**)

Differential Privacy

- A privacy framework that utilizes mathematics to quantify the amount of privacy that a privacy mechanism provides

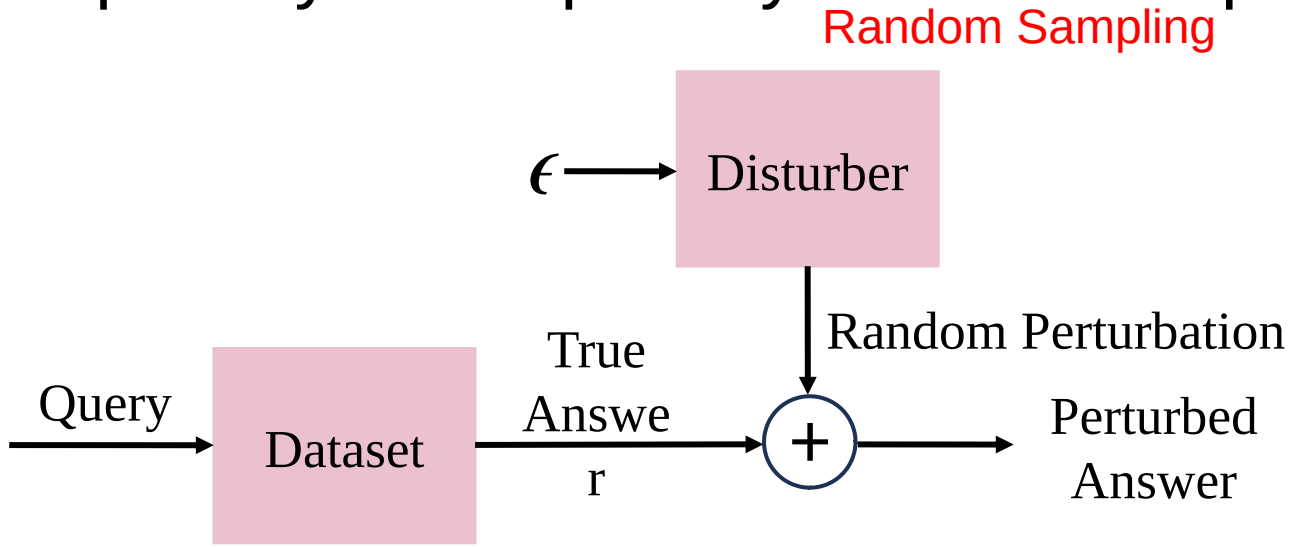

- -

$$P[M(D) \in S] \leq e^{\epsilon} \cdot P[M(D') \in S]$$

A dataset

The dataset that is one entry different from

A privacy mechanism

A subset of the output of the query

Directory Structure

```
Questionnaires
    demographic.pdf
    experience.pdf
    answers
        demographic.csv
        City_exp.csv
        Nature_exp.csv
        Office_exp.csv
        Gallery_exp.csv
Tools
    SensorLoggers
        Tracking.cs
        TrackingObj.cs
    realsense.py
    visualize.py
    README.md
```

```
City
    Trajectory
        city_$U.csv
        city_$U.json
    PhysicalRGBD
        city_$U.mp4
        city_depth_$U.mp4
        city_$U.txt
Nature
    ...
Office
    ...
Gallery
    ...
README.md
change.log
```

Dataset

Privacy Threats Investigation

Users' privacy can be threaten by revealing their raw trajectories

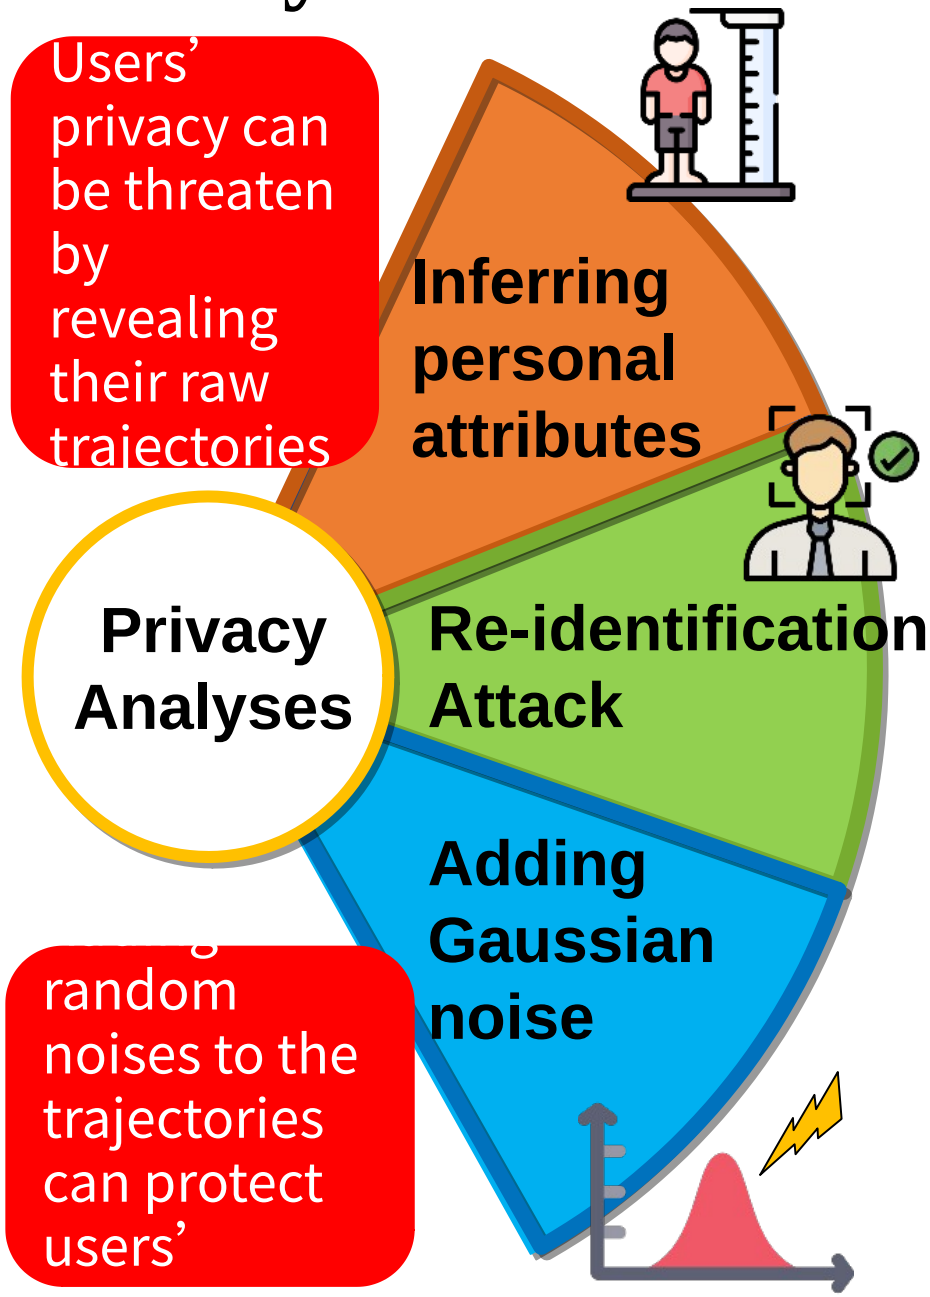

random noises to the trajectories can protect users'

- Height
- 75% accuracy

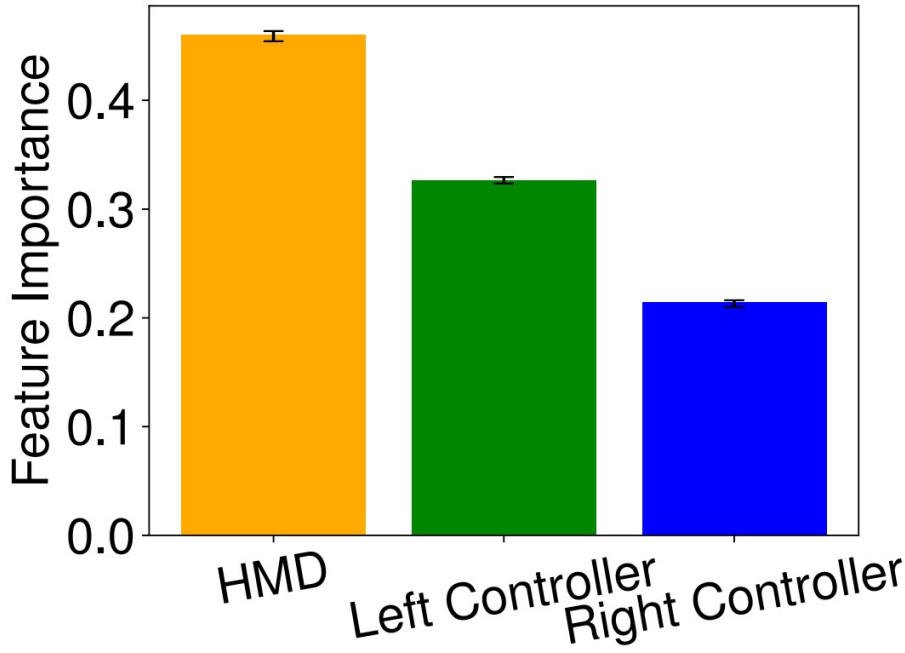

- Random Forest Classifier
- 80/20 train-test-split
- 96.41% accuracy
- Key features: HMD > Left ctrl. > Right ctrl.

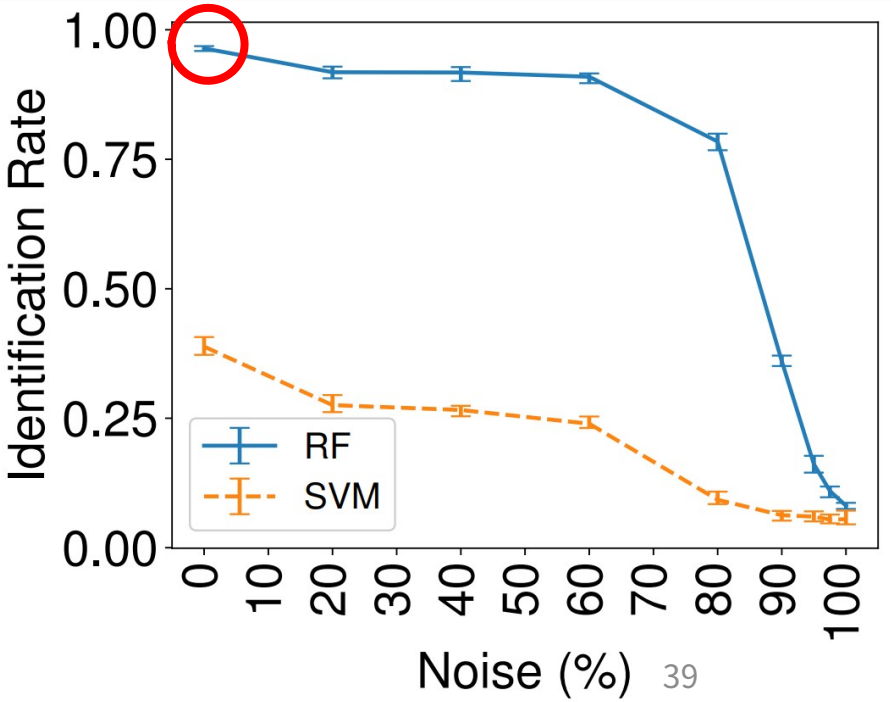

- Gaussian noise
- Noise percent
- The noise level ↑ the re-identification rate ↓

Disturber 2

Modeling the Trajectory with First-Order AutoRegressive Process [1]

- First-order AutoRegressive process
  - Time-series data modeling
  - Linear model
  - The current data is derived from the previous data

The order, first-order:

$$V_t = \sum_{i=1}^{p} \varphi_i V_{t-i} + U_t$$

White noise

Model perameters, autocorrelation when

- Model trajectory with Gaussian AR [2]
  - One of the most commonly used

,

when

$$V_t = \alpha + \varphi V_{t-1} + U_t$$

Autocorrelation,

Model parameters that needs to be estimated

[1] X. Zhang, M. Khalili, and M. Liu. Differentially private real-time release of sequential data. ACM Transactions on Privacy and Security, 2022.
[2] W. Wei. Time Siries Analysis, volume 2. Oxford University Press, 2013

Disturber 2

Trajectory Estimation

- Estimate the trajectory with Linear Minimum Mean Squre Error (LMMSE)

$$\hat{V}_t = \mu(1-\varrho) + \varrho V_{t-1}$$

- Estimate the trajectory with prior perturbed pose

$$\hat{V}_t = \mu\left(1-\varrho\frac{\sigma^2}{\sigma^2+var}\right) + \varrho\frac{\sigma^2}{\sigma^2+var}P_{t-1}$$

Prior perturbed pose,

Statistics of the whole trajectory, and needs to be estimated

Probability distribution variance

- Estimate the statistics following Huitema and McKean [1]

$$\hat{V}_t = \hat{\mu}_{t-1}\left(1-\hat{\varrho}_{t-1}\frac{\hat{\sigma}^2_{t-1}}{\hat{\sigma}^2_{t-1}+var}\right) + \hat{\varrho}_{t-1}\frac{\hat{\sigma}^2_{t-1}}{\hat{\sigma}^2_{t-1}+var}P_{t-1}$$

[1] B. Huitema and J. McKean. Autocorrelation estimation and inference with small samples. Psychological Bulletin, 1991

Implications of Diverse Characteristics of 3D Scenes (1/2)

- The visual quality of City outperforms the others at all time

- **The size of the City is considerably larger than that of the others**
- **City is a vast scene with wide roads and a large sky**
  - **The perturbed image may not be influenced too much when the poses are changed**

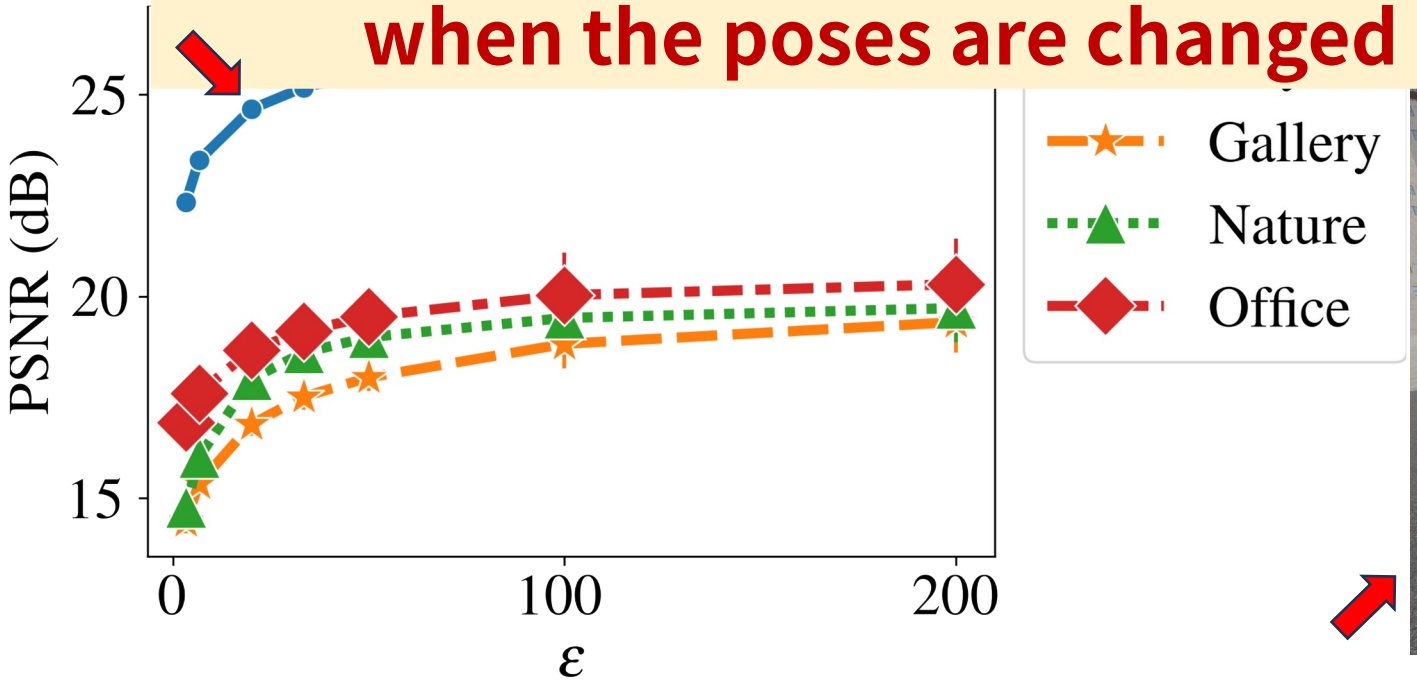

Implications of Diverse Characteristics of 3D Scenes (2/2)

- The quality of all the scenes ⬆ with DC except for City
- ⬆DC's visual quality of City ⬇

Most of the disoccluded sky areas in City are inpainted with buildings after DC is applied

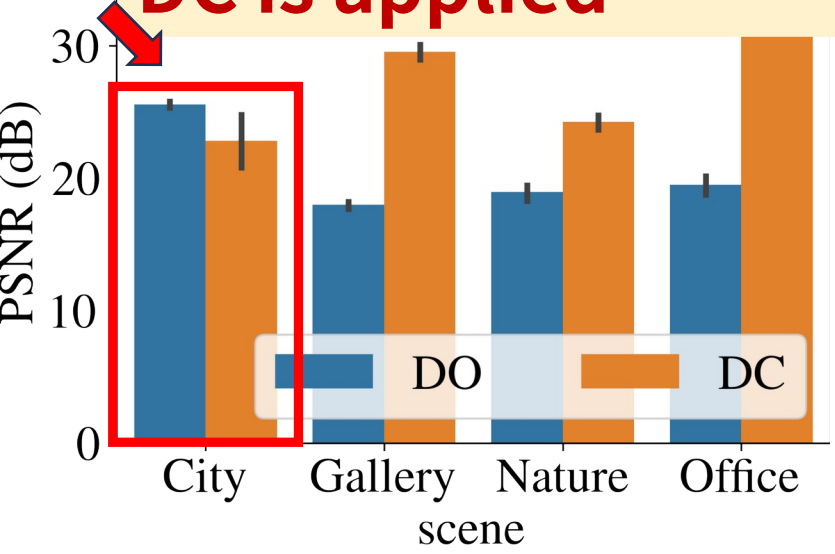

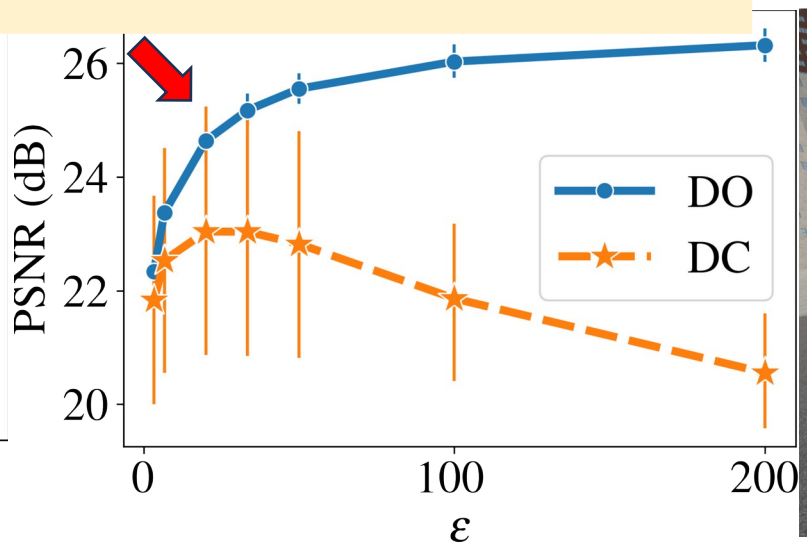

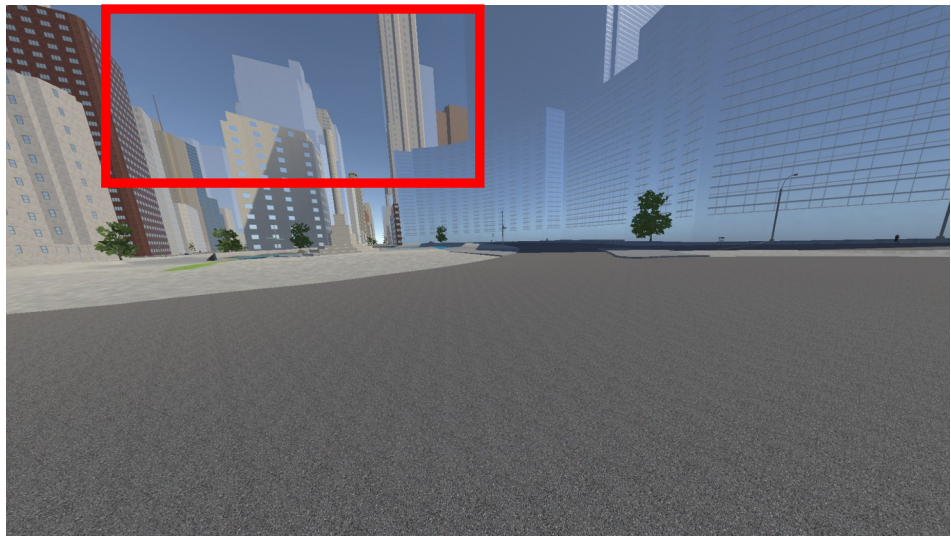

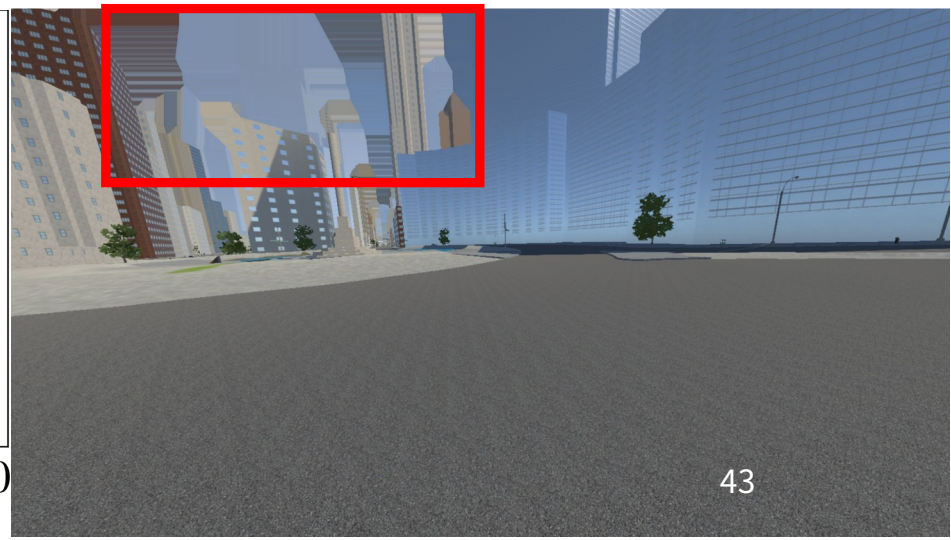

Implications of Different Parameters

- The re-id rates of is lower than that of

Weight	0.1	0.3	0.5	0.7
Re-id Rate	0.419 (±0.259)	0.483 (±0.253)	0.493 (±0.249)	0.495 (±0.244)

- The visual quality is slightly degraded (dB in PSNR)

Smaller leads to lower re-id rate

- The visual quality of DC with larger that with smaller

PSNR(dB)	25.62 (±4.57)	25.39 (±4.20)
SSIM	0.87 (±0.12)	0.86 (±0.12)
VMAF	45.86 (±21.43)	41.08 (±19.51)

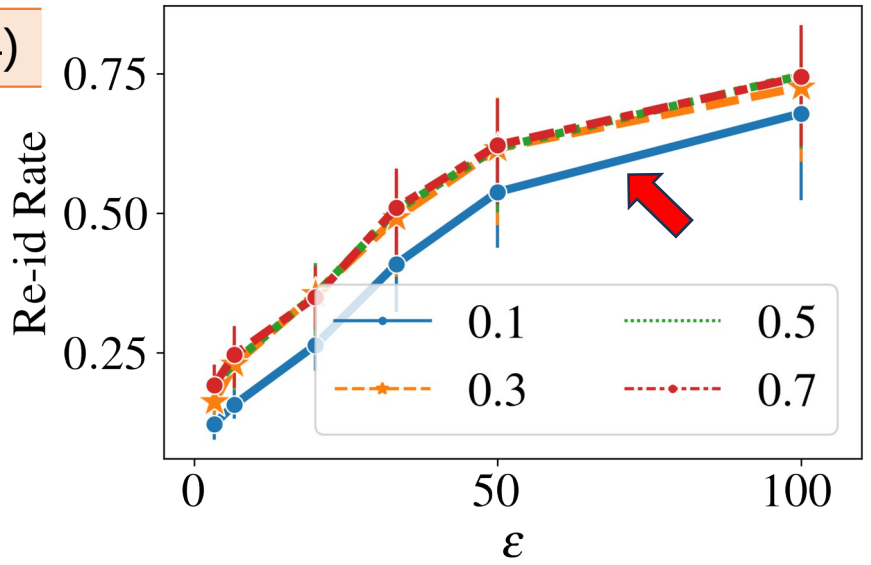

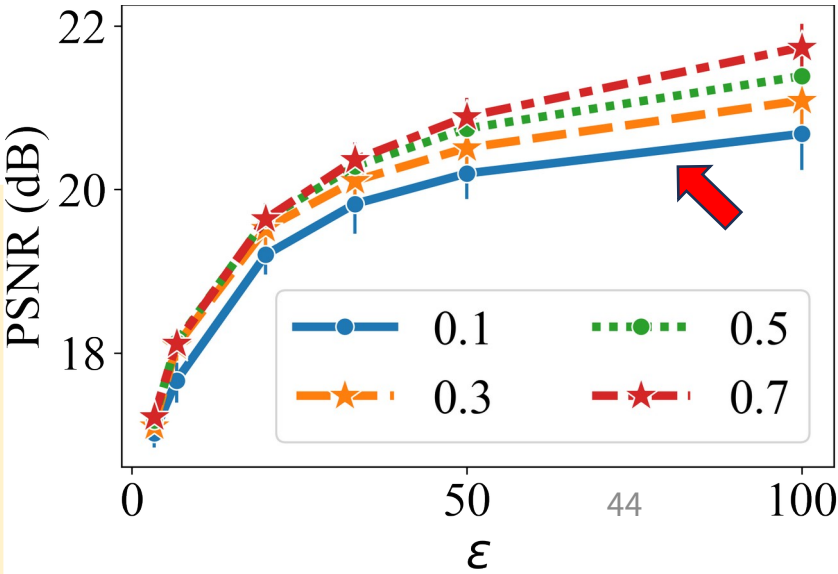

Larger doesn’ t lead to better performance of compensation

- **The inpainting method**
- **In a smaller scene, larger rendered FoV may only include more pixels that are ineffectual for**

Adding Perturbations to Both Locations and Orientations

- Adding perturbation to both locations and orientations degrades the visual quality drastically while only reducing the re-id rate a little compared to DO

Adding perturbations to location only is more efficient

Re-id rate without any privacy-threat mitigation is 0.83

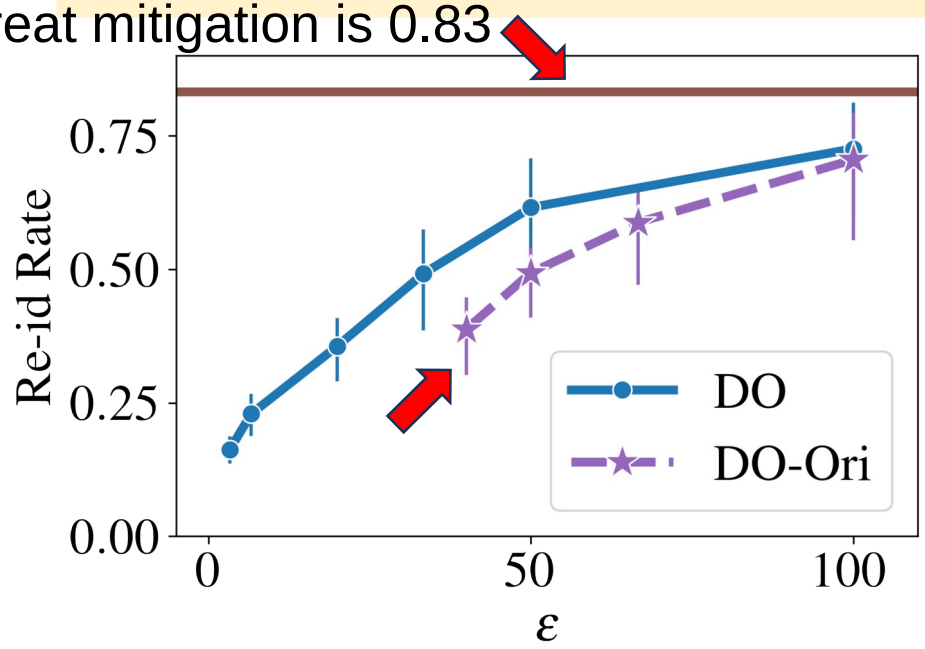

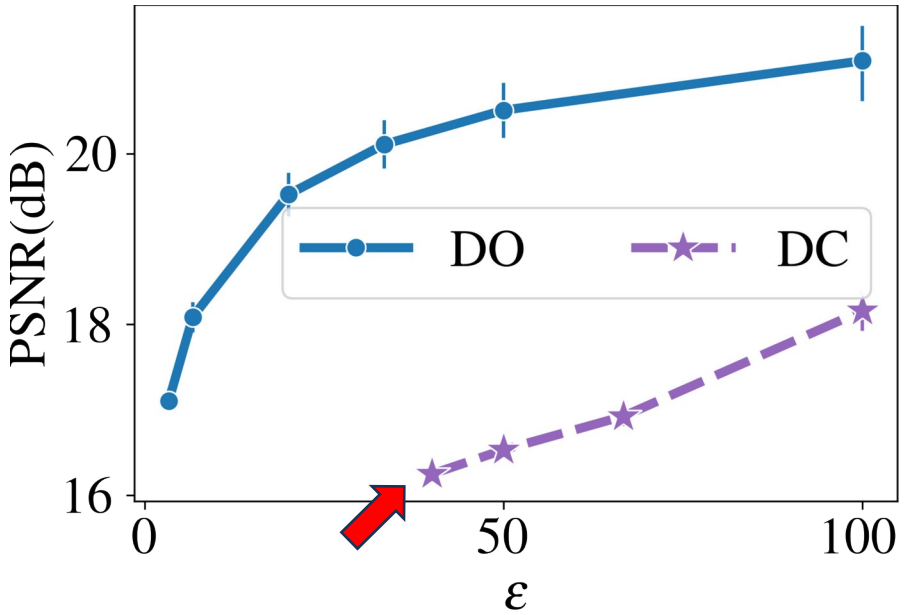

Future Work

- A large-scale user study with the system
- Other more secure placement