McMaster University

CVPR JUNE 3-7, 2026 DENVER COLORADO

LDP-Slicing: Local Differential Privacy for Images via Randomized Bit-Plane Slicing

Yuanming Cao, Chengqi Li, Wenbo He

McMaster University

Making local privacy practical for standard images

Centralized DP Assumes A Trusted Curator

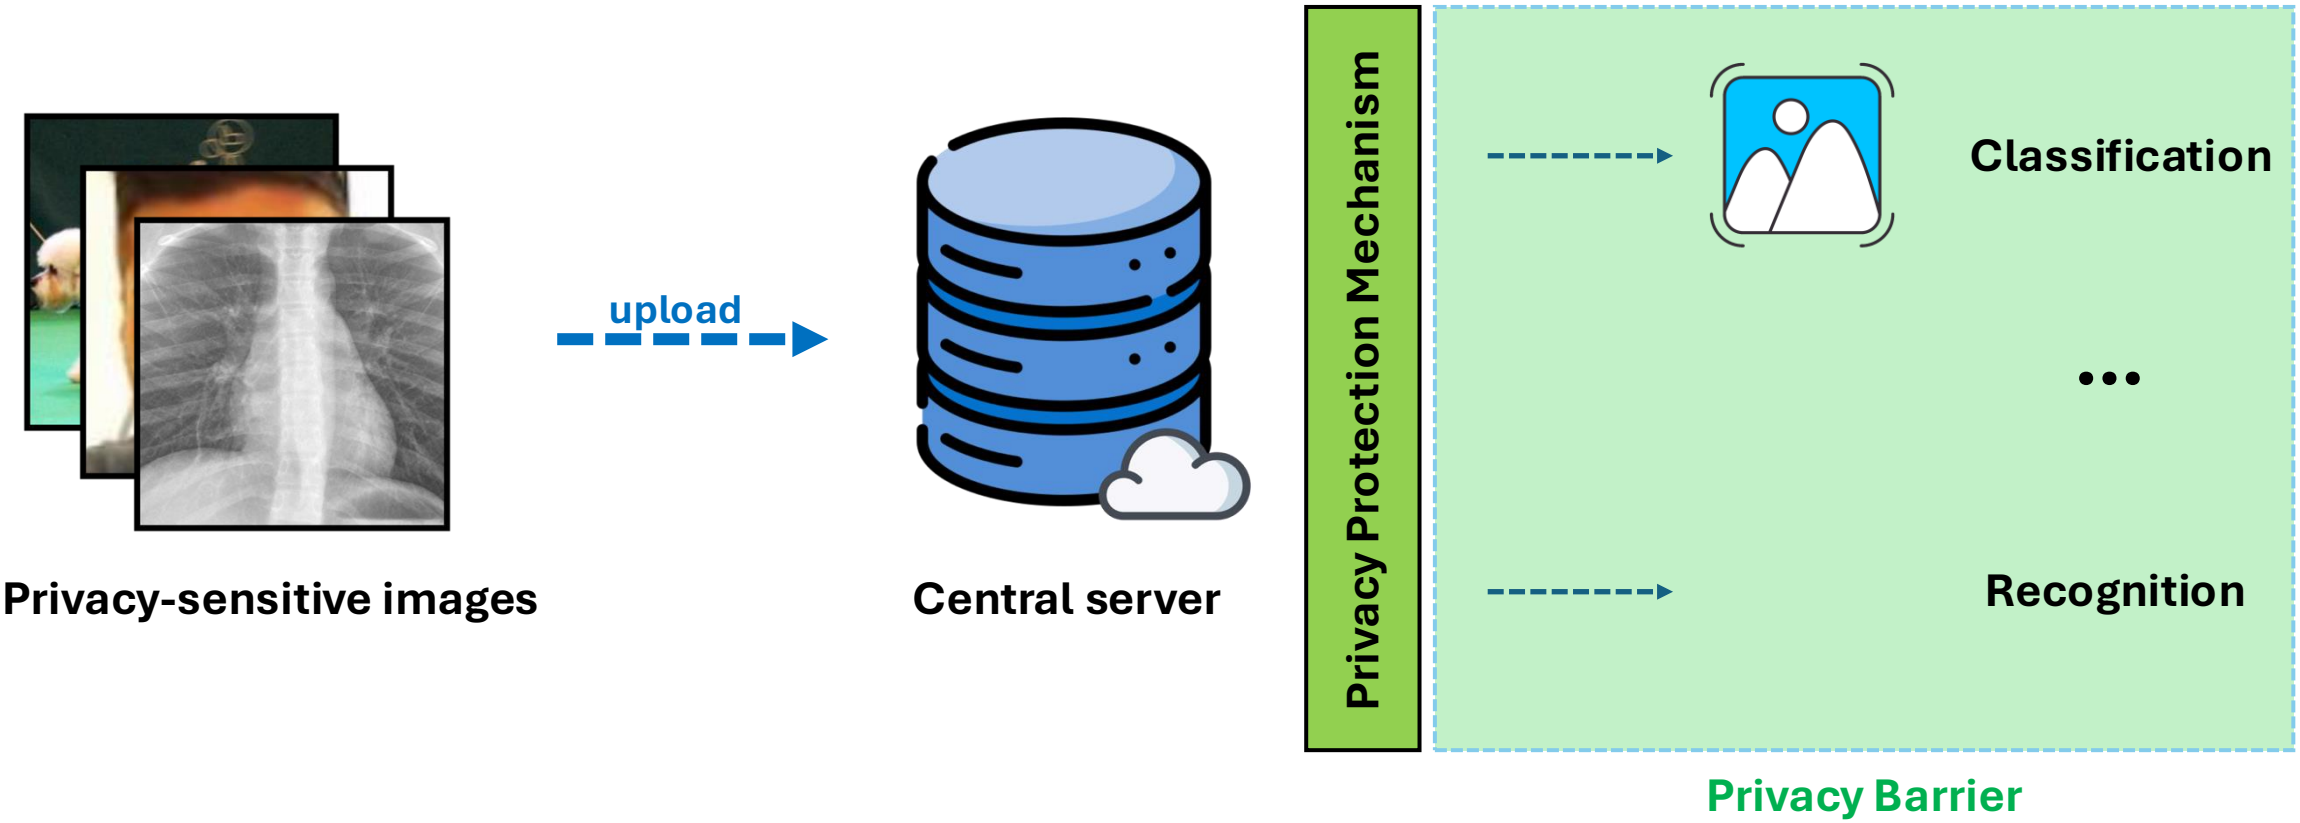

The curator can output a perturbed model /dataset —only after the raw image has been collected.

Privacy Needs To Be Enforced Before Upload

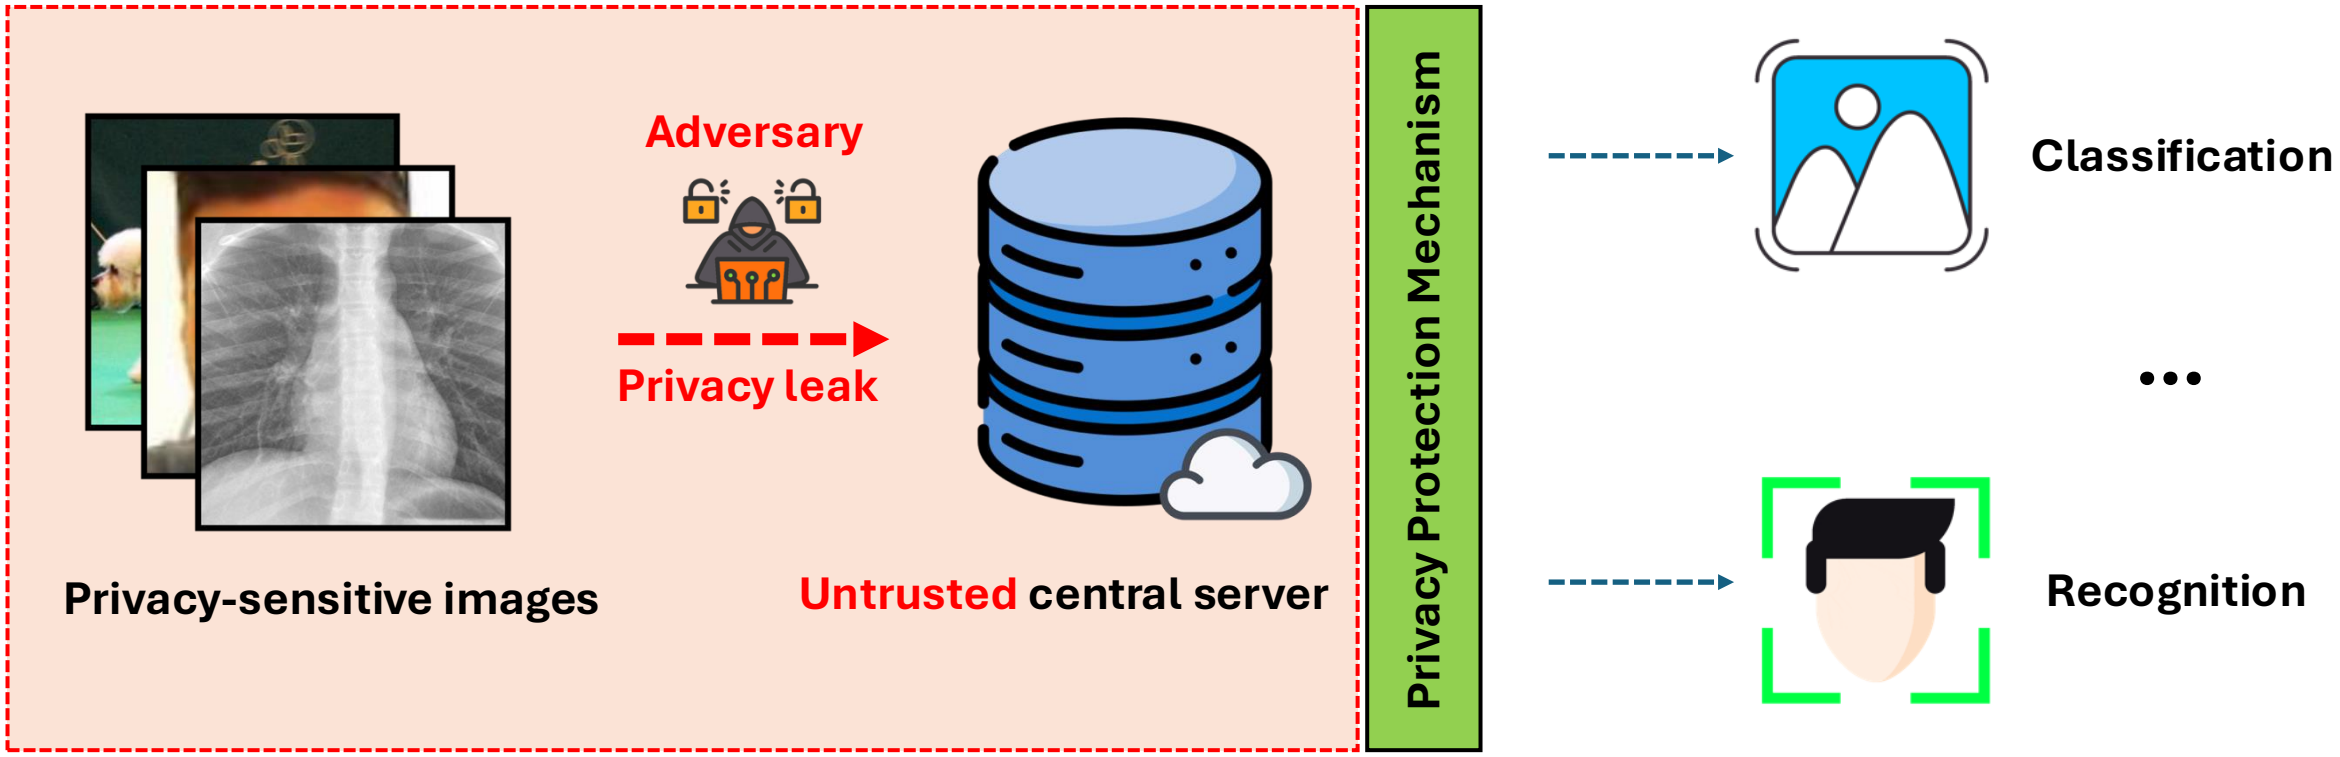

In a zero-trust setting, the raw image is already exposed.

Local Differential Privacy — The Server Only Sees Randomized Data

Randomized Response [Warner'65]

For $v \in \{0, 1\}$, the (Binary) Randomized Response mechanism outputs $v' \in \{0, 1\}$ according to

$$\Pr[\mathcal{M}_{\mathrm{RR}}(v) = v'] = \begin{cases} p & \text{if } v = v', \\ 1 - p & \text{if } v \neq v', \end{cases} \qquad p = \frac{e^{\varepsilon}}{e^{\varepsilon} + 1}.$$

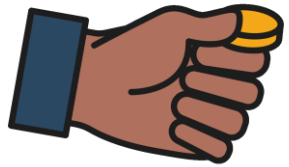

Why Raw-Pixel LDP Destroys Utility

Raw pixel RR treats intensity as a 256-way category

- A grayscale pixel has **256** possible intensity values.
- Under 256-class randomized response:
  - $P_{true} = \frac{e^{\varepsilon}}{e^{\varepsilon}+255}$
- For $\varepsilon = 1$:
  - $P_{true} \approx 0.011, P_{false} \approx 0.989$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Visualization of the 8-bit pixel value space

Prior LDP for images works in a lower-dimensional representation

- **Prior LDP releases derived features:** descriptors, eigenfaces, latent vectors.
- **Goal:**
  - Release a **standard image** with pixel-level ε-LDP.
  - Standard vision pipelines **stay intact.**

Prior Work

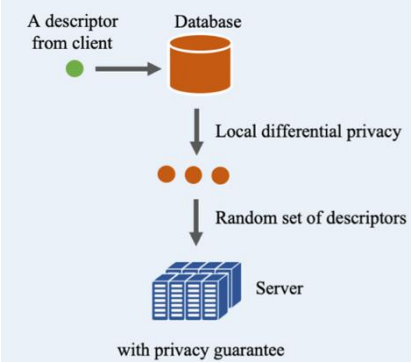

Local feature descriptor

Pittaluga & Zhuang, ICCV'23

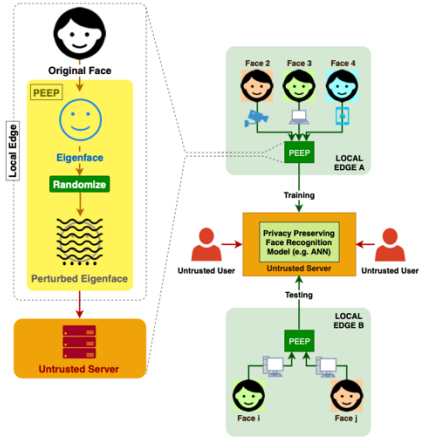

Eigenface projection

Chamikara et al., Computers & Security'20

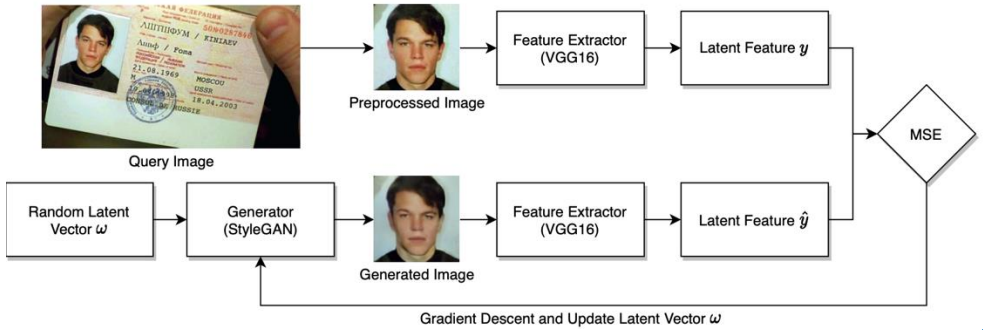

Generative latent representation

Li & Clifton, IEEE S&P'21

A 256-Way Pixel Value Is Actually An 8-Bit Code

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Visualization of the 8-bit pixel value space

A 256-Way Pixel Value Is Actually An 8-Bit Code

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Visualization of the 8-bit pixel value space

8-bit representation

b_7	b_6	b_5	b_4	b_3	b_2	b_1	b_0
1	0	1	0	1	1	0	1

$$x = \sum_{i=0}^{7} b_i 2^i$$

- Under binary RR for **flipping each bit:**
  - Reporting true pixel only if **all** bits are true.
  - $\boldsymbol{P_{true\ bit}} = \dfrac{e^{\varepsilon\ bit}}{e^{\varepsilon\ bit} + 1}$
  - Pixel: $\varepsilon_{total} = \sum_0^7 \varepsilon_{bit}$

Bit-Plane Slicing converts pixels into binary LDP domains

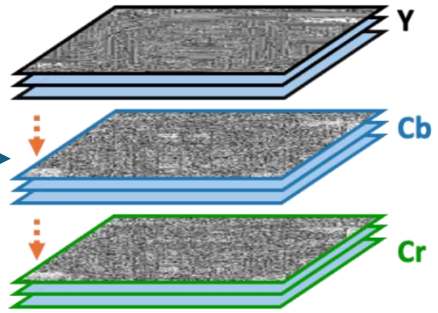

24 bit-planes

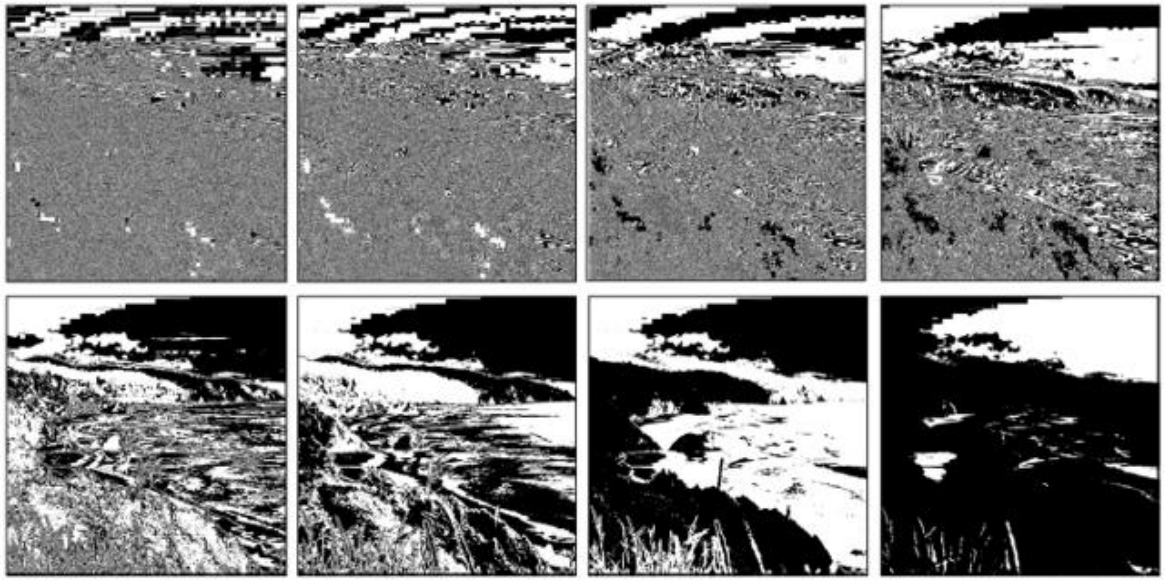

Y channel bit-plane

- **Naturally suited to binary RR.**
  - Each bit is a simple 0/1 value.
- **Expose useful image structure.**
  - High-order bits carry semantics.
  - low-order bits carry texture and noise.

Bit randomization erodes the spatial correlations behind semantics

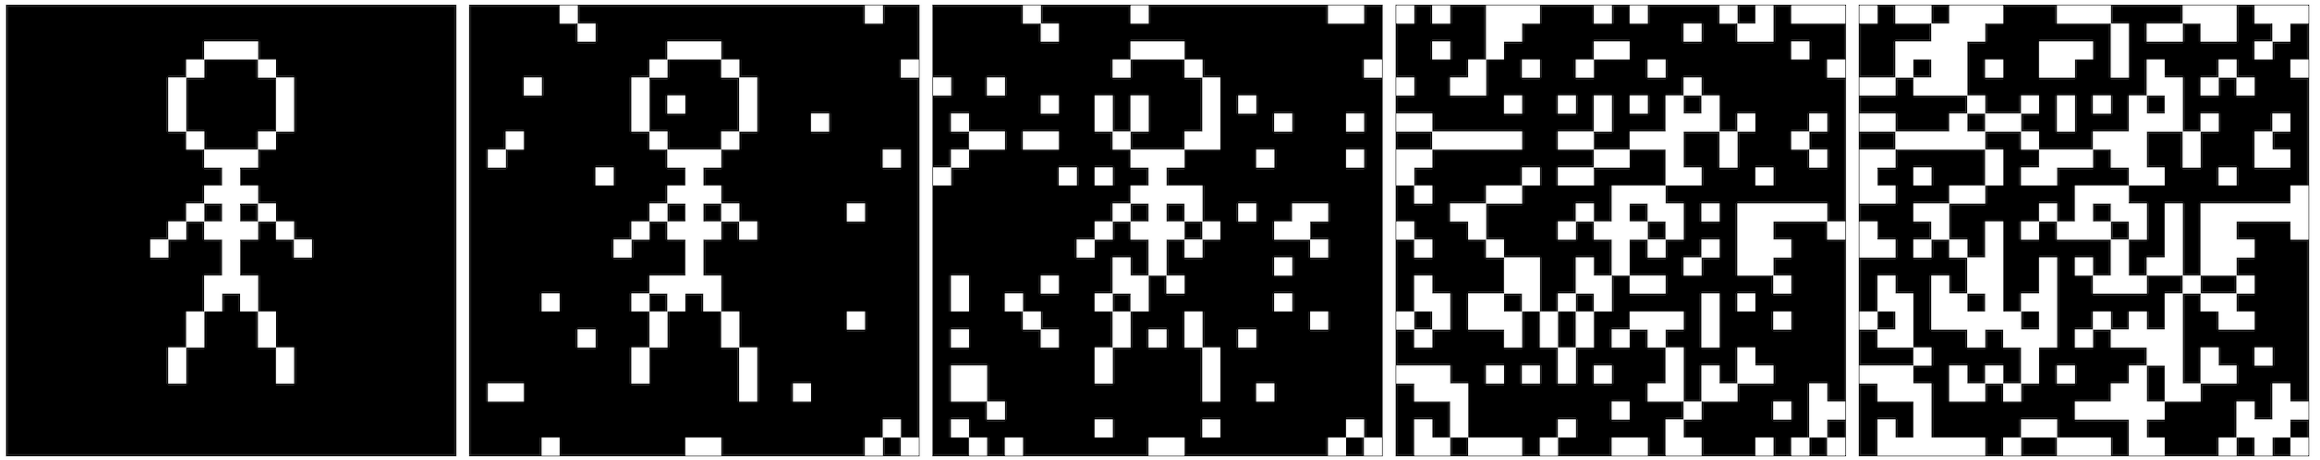

Binary image $\varepsilon = 3$ $\varepsilon = 2$ $\varepsilon = 1$ $\varepsilon = 0.5$

Image Utility Is Non-Uniform Across Channels and Bit-Planes

- **Uniform allocation is suboptimal.**

  $\varepsilon_\ell = \frac{\varepsilon_{total}}{24}$ treats all bit-planes equally.

- **Bit-planes carry unequal structure.**

  Higher-order bits dominate **coarse** image content.

- **Color channels carry unequal structure.**

  The **luma Y channel** preserves **most spatial structure**;

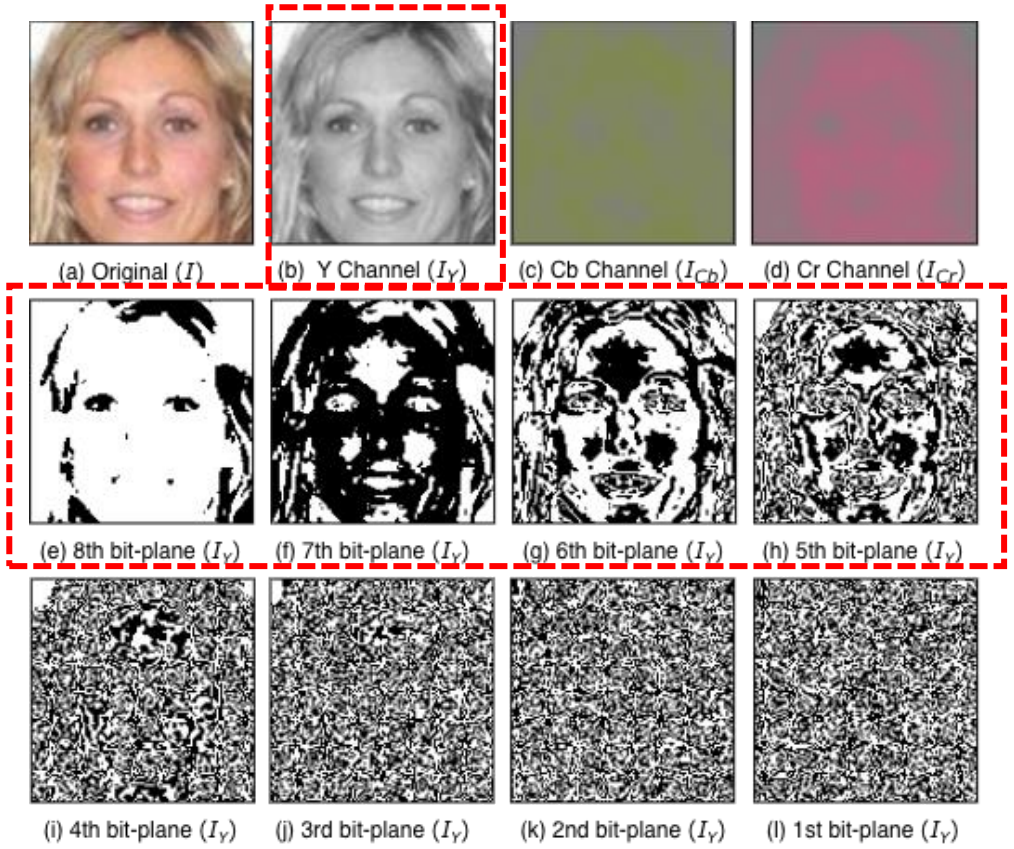

Privacy noise should be allocated according to structural importance, not uniformly.

Weighted Distortion Minimization Gives a Closed-Form Budget Allocation

Weighted distortion objective:

- **Channel sensitivity (W_c):**

  CNNs are robust to chroma loss → $W_y = 4$, $W_{cb} = W_{cr} = 1$.

- **Bit significance (W_b):**

  Proportional to the bit's pixel-value contribution → $W_b = 2^{b-1}$.

$$\min_{\{\varepsilon_{c,b}\}} \sum_{c,b} \frac{W_{c,b}}{\varepsilon_{c,b}}$$

$$\text{s.t.} \quad \sum_{c,b} \varepsilon_{c,b} = \varepsilon_{\text{total}}, \quad \varepsilon_{c,b} \geq 0.$$

Result:

$$\varepsilon_{c,b} = \varepsilon_{\text{total}} \cdot \frac{\sqrt{W_{c,b}}}{\sum_{i\in\{\text{Y,Cb,Cr}\}} \sum_{j=1}^{8} \sqrt{W_{i,j}}}$$

Higher-weight components receive larger ε, meaning less noise where utility is most sensitive.

Wavelet LL Pruning Removes Low-Frequency Identity Cues Before LDP

Wavelet LL pruning complements LDP at larger ε

- **Removes coarse cues:** shape, lighting, facial layout.
- **Reduces perceptual leakage:** less recognizable to humans.
- **Weakens reconstruction attack:** less low-frequency structure to recover

w/ LL ,$\boldsymbol{\varepsilon_{total}} = \mathbf{58}$

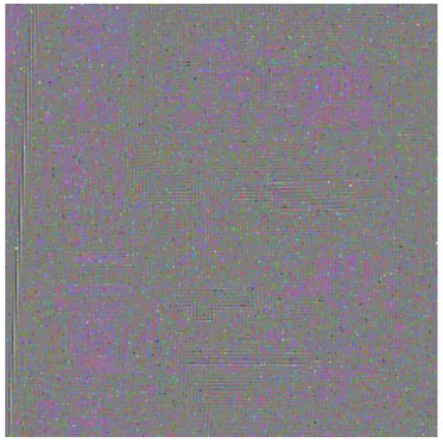

w/o LL, $\boldsymbol{\varepsilon_{total}} = \mathbf{58}$

w/ LL, $\boldsymbol{\varepsilon_{total}} = \mathbf{5.2}$

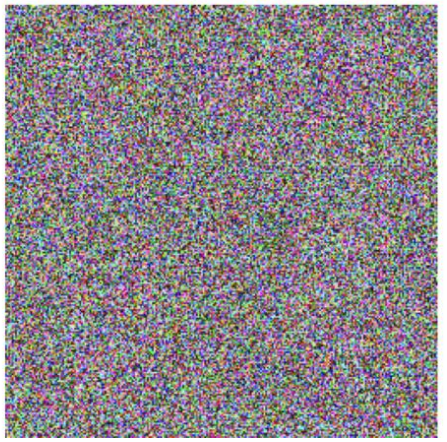

w/o LL, $\boldsymbol{\varepsilon_{total}} = \mathbf{5.2}$

The full pipeline outputs a standard image with pixel-level ε-LDP

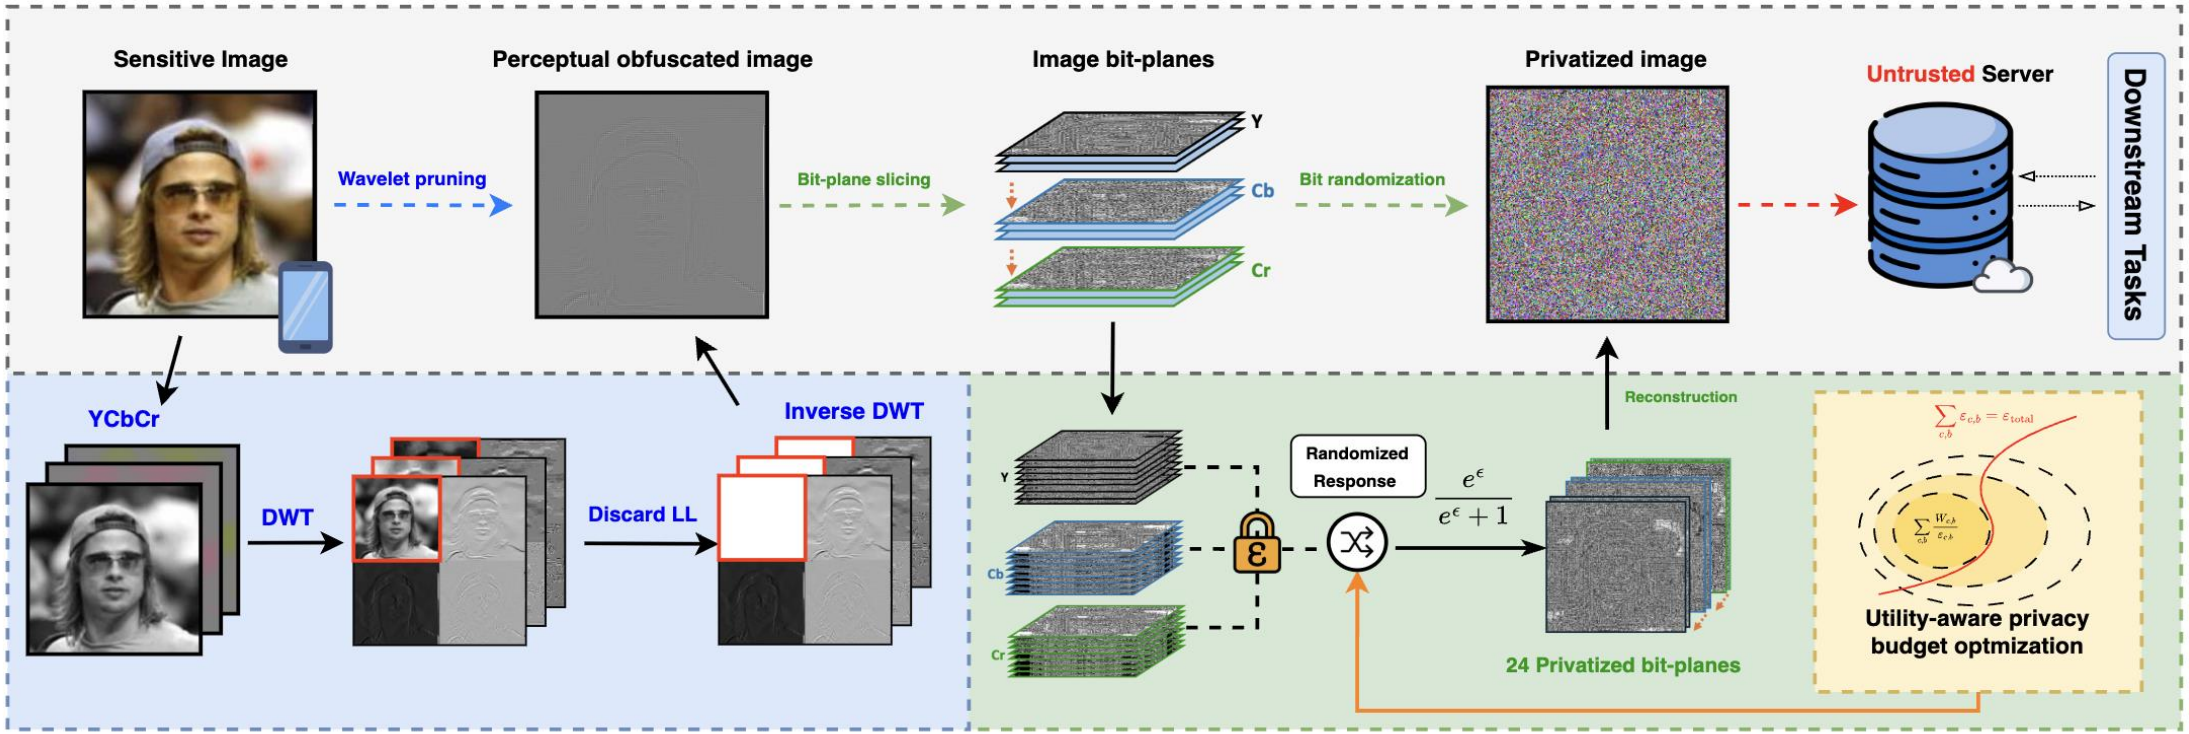

LDP-Slicing Preserves Face Recognition Utility Under Pixel-Level LDP

Main Results · Privacy-preserving face recognition

Method	DP/LDP Guarantee	AgeDB-30	LFW	CPLFW	CALFW
ArcFace (baseline)	None	97.88	99.77	92.77	96.05
InstaHide	None	79.58	96.53	81.03	86.24
Cloak	None	92.60	98.91	83.43	92.18
PPFR-FD	None	97.23	99.69	90.19	95.60
DuetFace	None	97.93	99.82	92.77	96.10
PartialFace	None	97.79	99.80	92.03	96.07
ProFace	None	92.81	98.27	88.17	93.20
AdvFace	None	92.57	98.45	83.73	93.62
MinusFace	None	97.57	99.78	91.90	95.90
PEEP	Feature-level	87.47	98.41	79.58	90.06
DCTDP	Block-level	94.37	99.48	90.60	93.47
LDP-Slicing (Ours)	**Pixel-level**	**96.68**	**99.75**	**91.08**	**96.02**

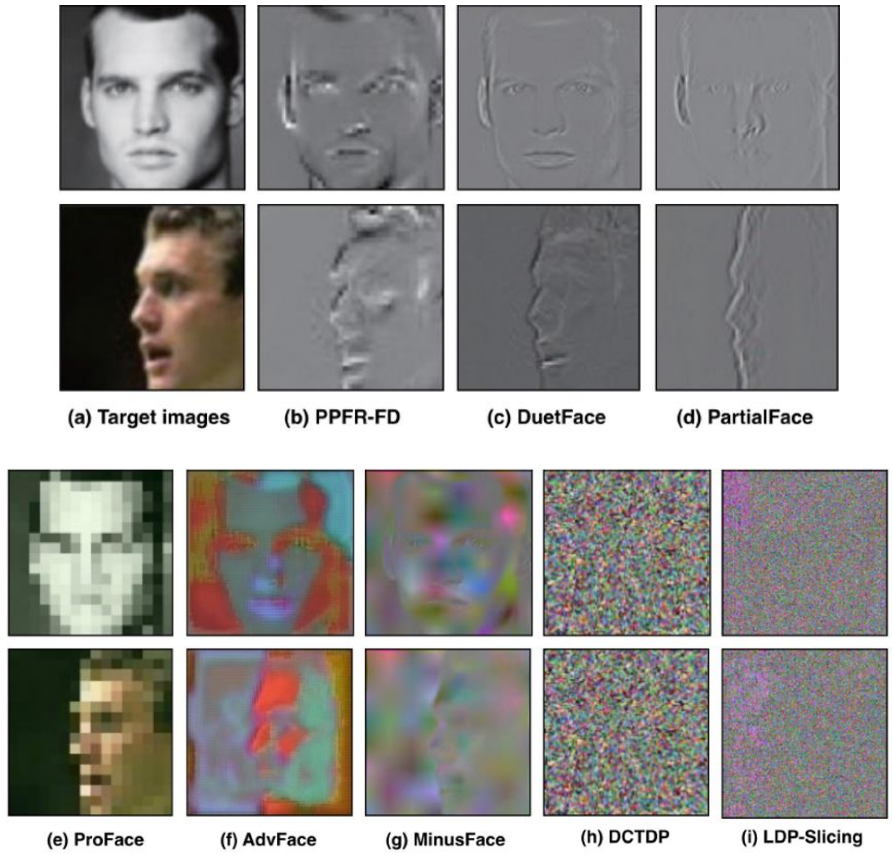

The Same Representation Also Preserves Classification Utility

Main Results · Privacy-preserving image classification

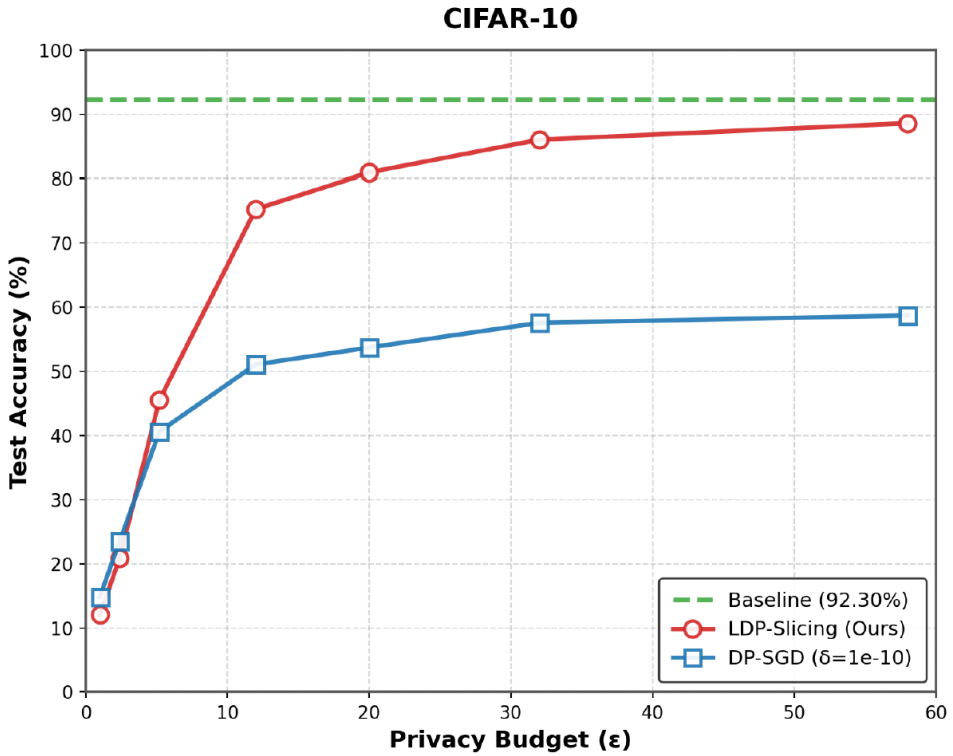

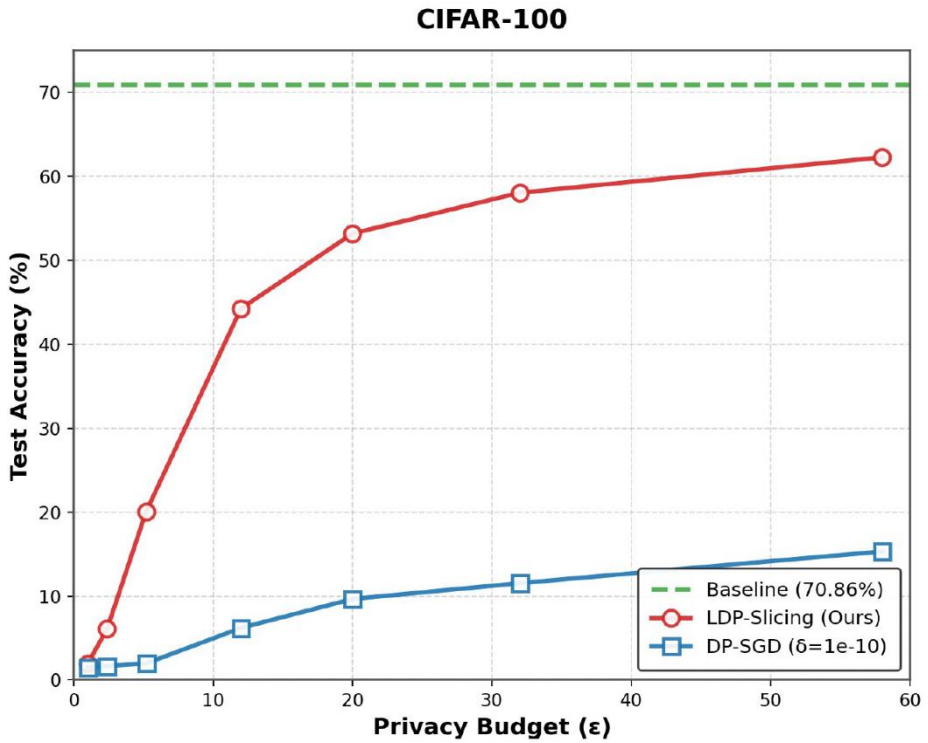

White-Box and Generative-Prior Attacks Fail to Recover Recognizable Identity

White-box reconstruction

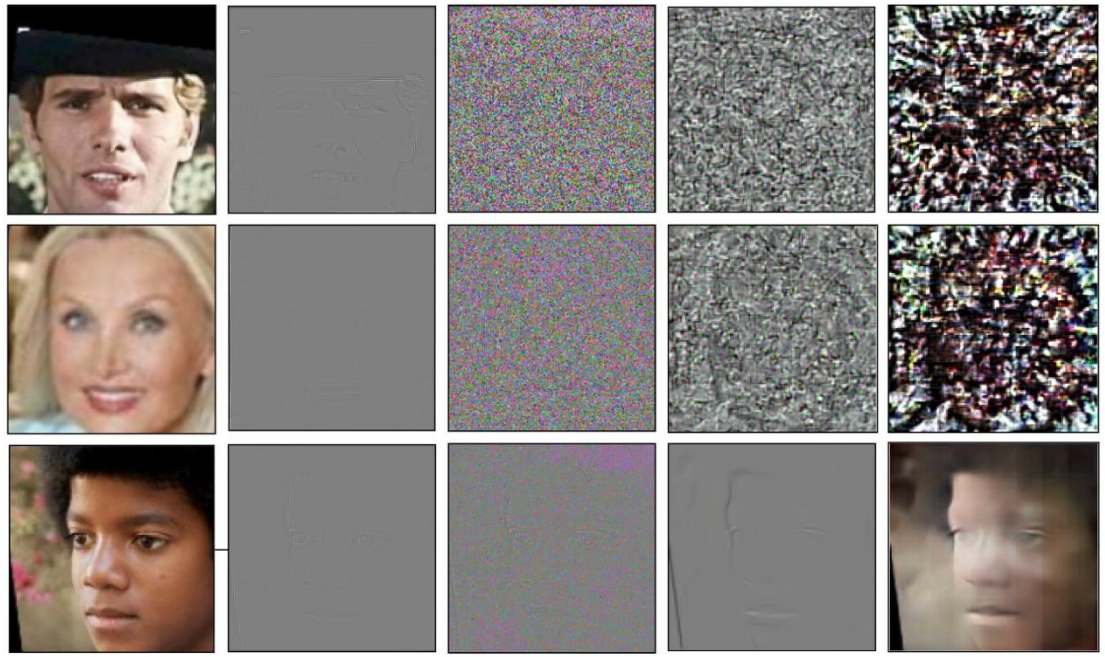

Map^{2}V reconstruction (CVPR'24)

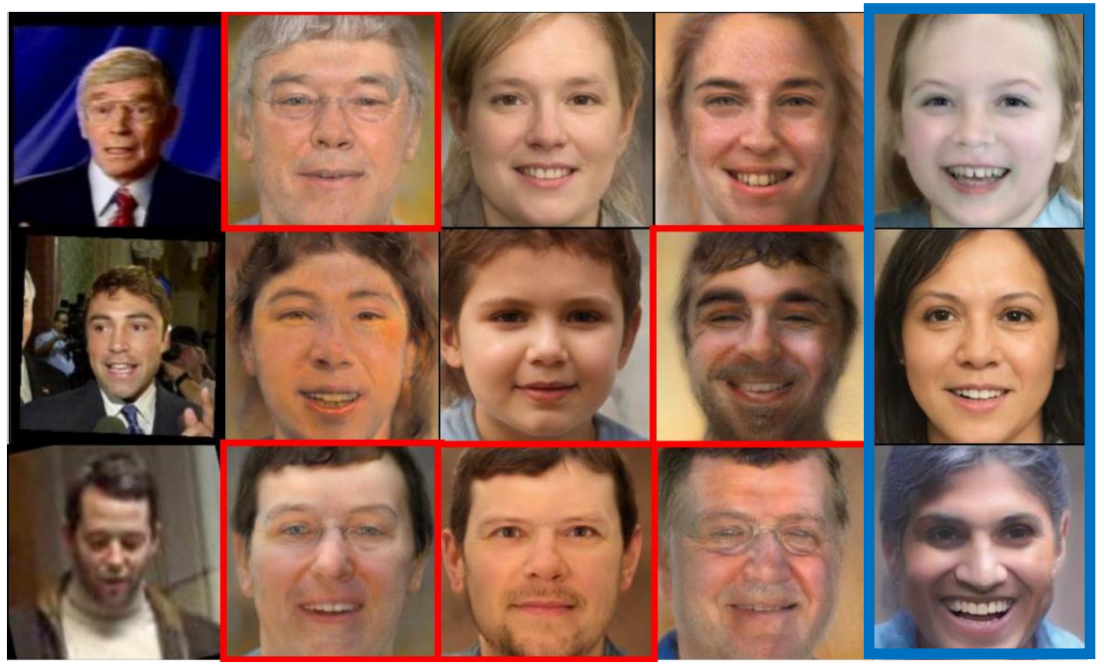

Representation makes image LDP practical

- **Problem:** raw image LDP suffers from a privacy–utility mismatch.
- **Method:** bit-plane slicing + binary randomized response gives pixel-level ε-LDP.
- **Utility:** budget allocation and LL pruning preserve task signal while reducing identity leakage.
- **Practicality:** outputs a standard image for existing vision pipelines.

Project page

GitHub

Questions?

McMaster University

CVPR JUNE 3-7, 2026 DENVER COLORADO