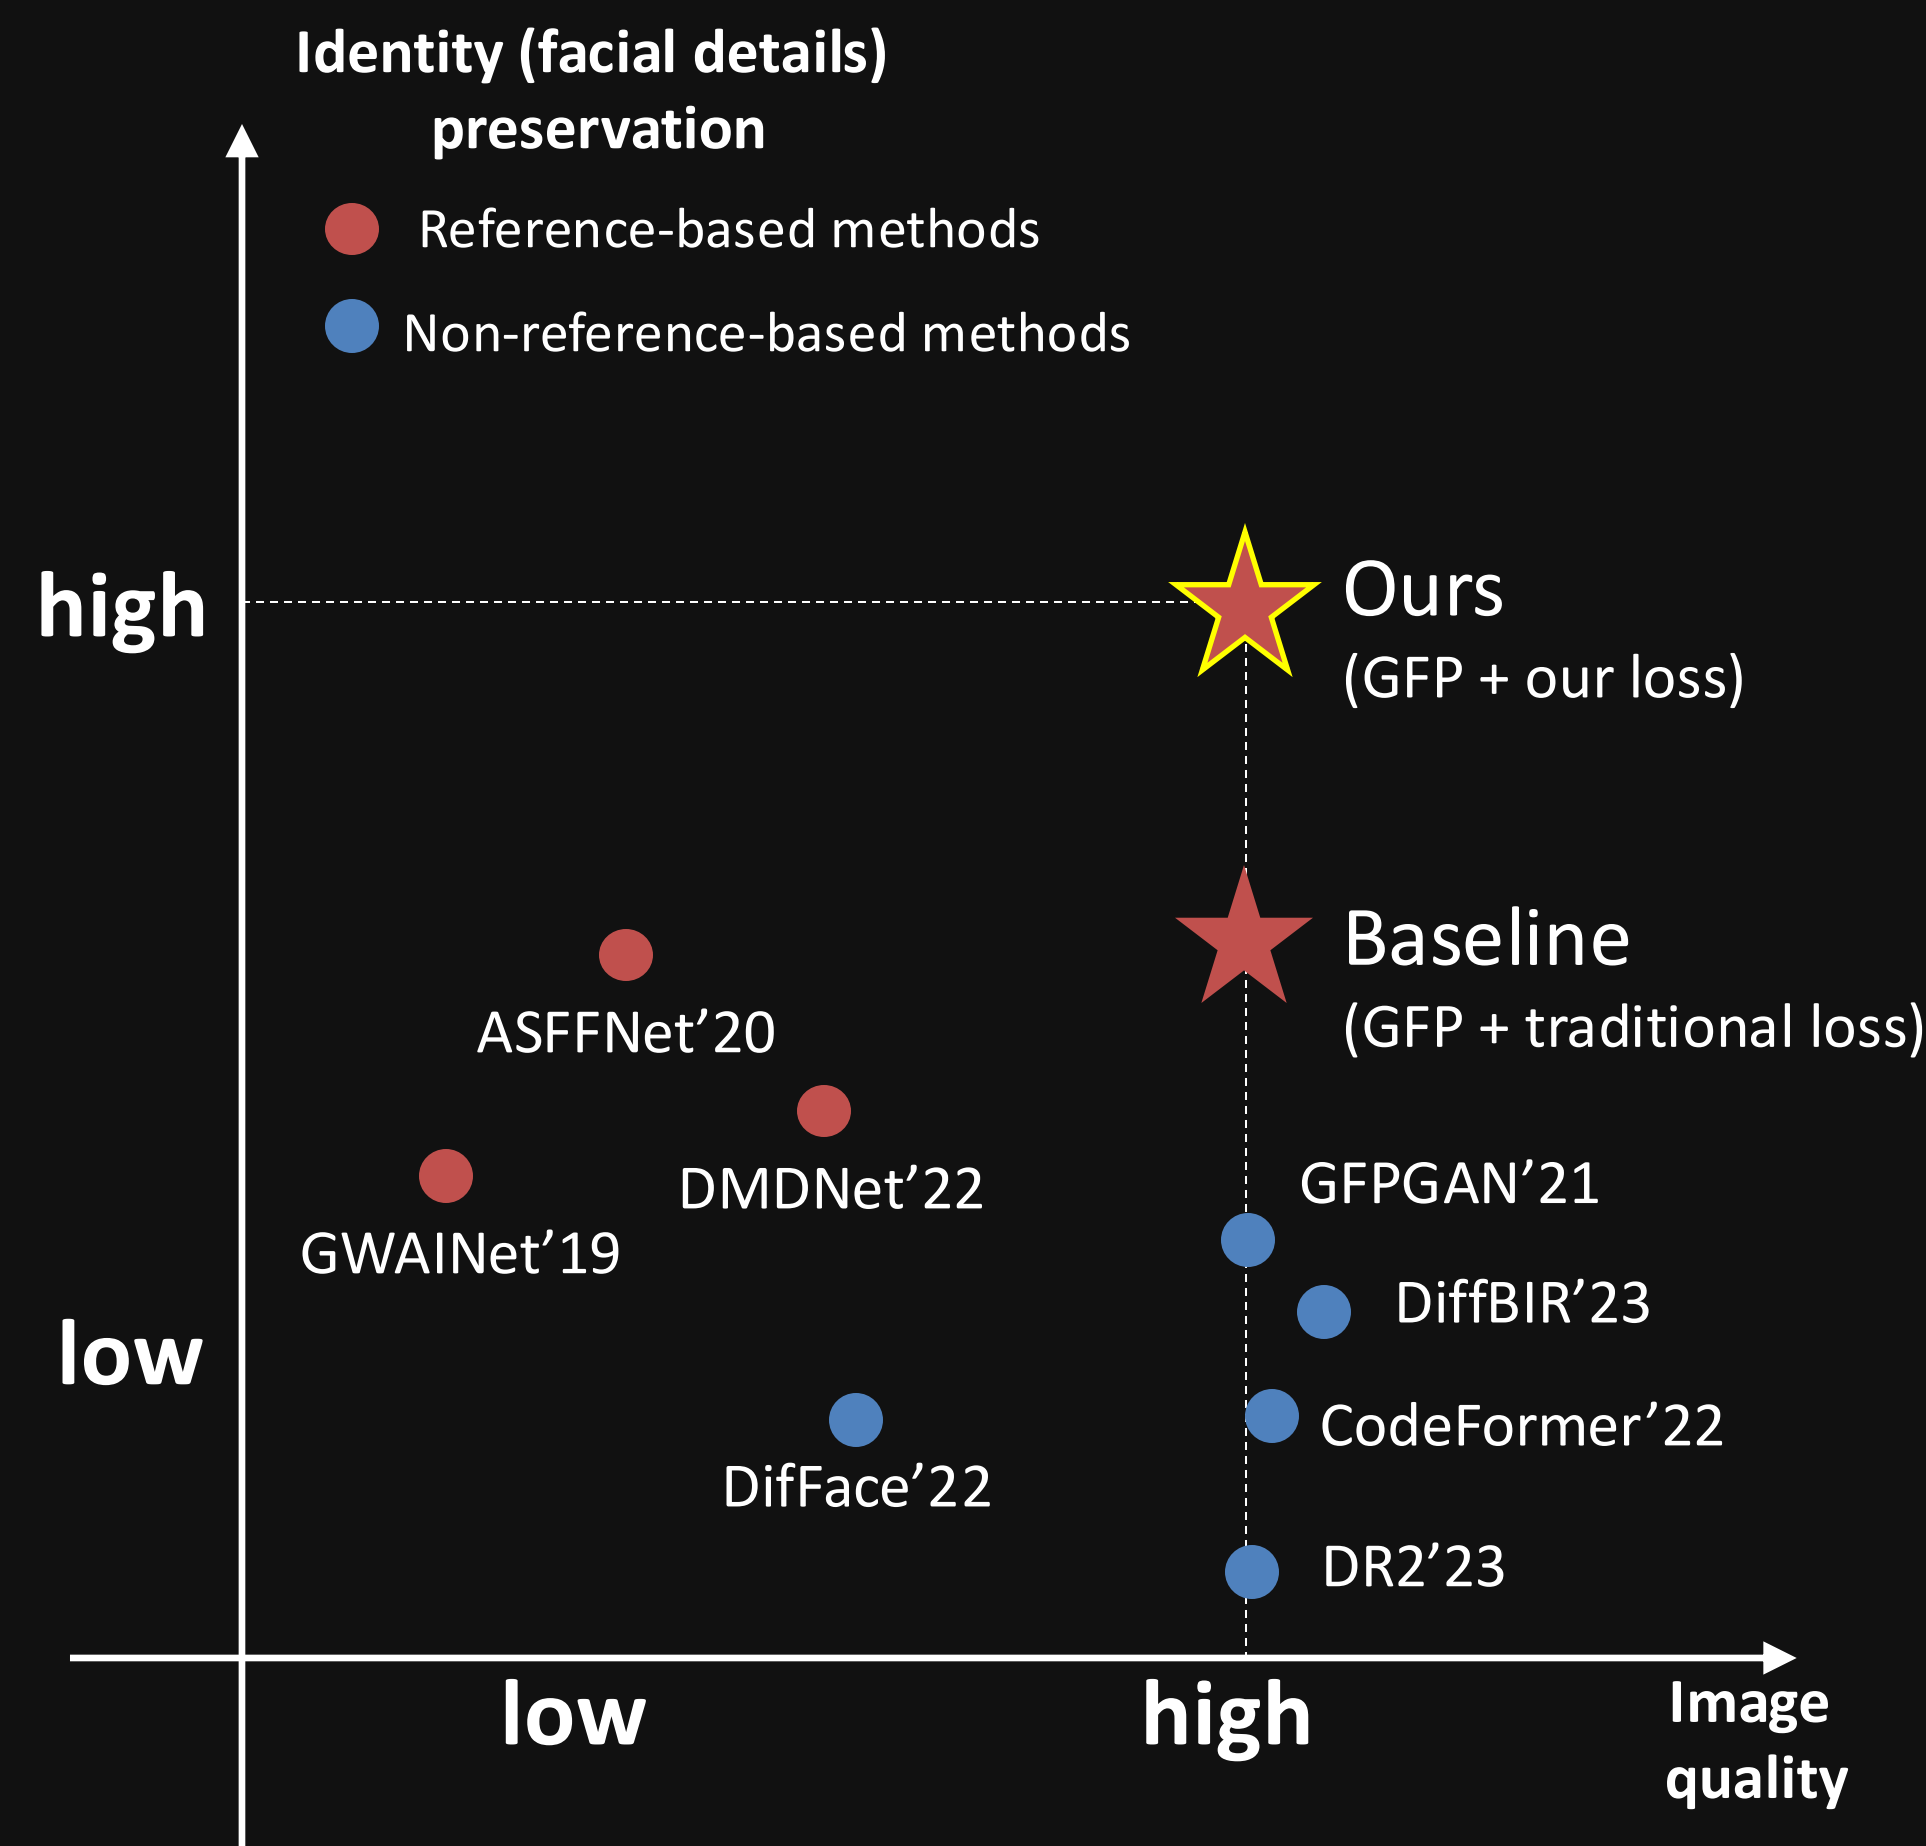



Copy or Not? Reference-Based Face Image Restoration with Fine Details

Min Jin Chong¹, Dejia Xu², Yi Zhang³, Zhangyang Wang², David Forsyth¹, Gurunandan Krishnan⁴, Yicheng Wu⁴, Jian Wang⁴

¹UIUC, ²UT Austin, ³CUHK, ⁴Snap Inc.

Background



Our Observation and Method



Age, makeup, etc. → Reference and ground truth can be different in facial details

If the reference image has a mole, but the ground truth doesn't, should the output have a mole?

Answer:

- If the input is too low quality to confirm a mole, the output should include one.
- If the input is low quality but clearly lacks a mole, the output should not include one.


Traditional goal:
Restore back the ground truth

→ Always no mole in the output

Our goal:

- (1) Copy features from reference as much as possible.
- (2) Generate details not available from reference.
- (3) Output semantically consistent with degraded input.
- (4) Face must be realistic.

Results



One of the Product UIs

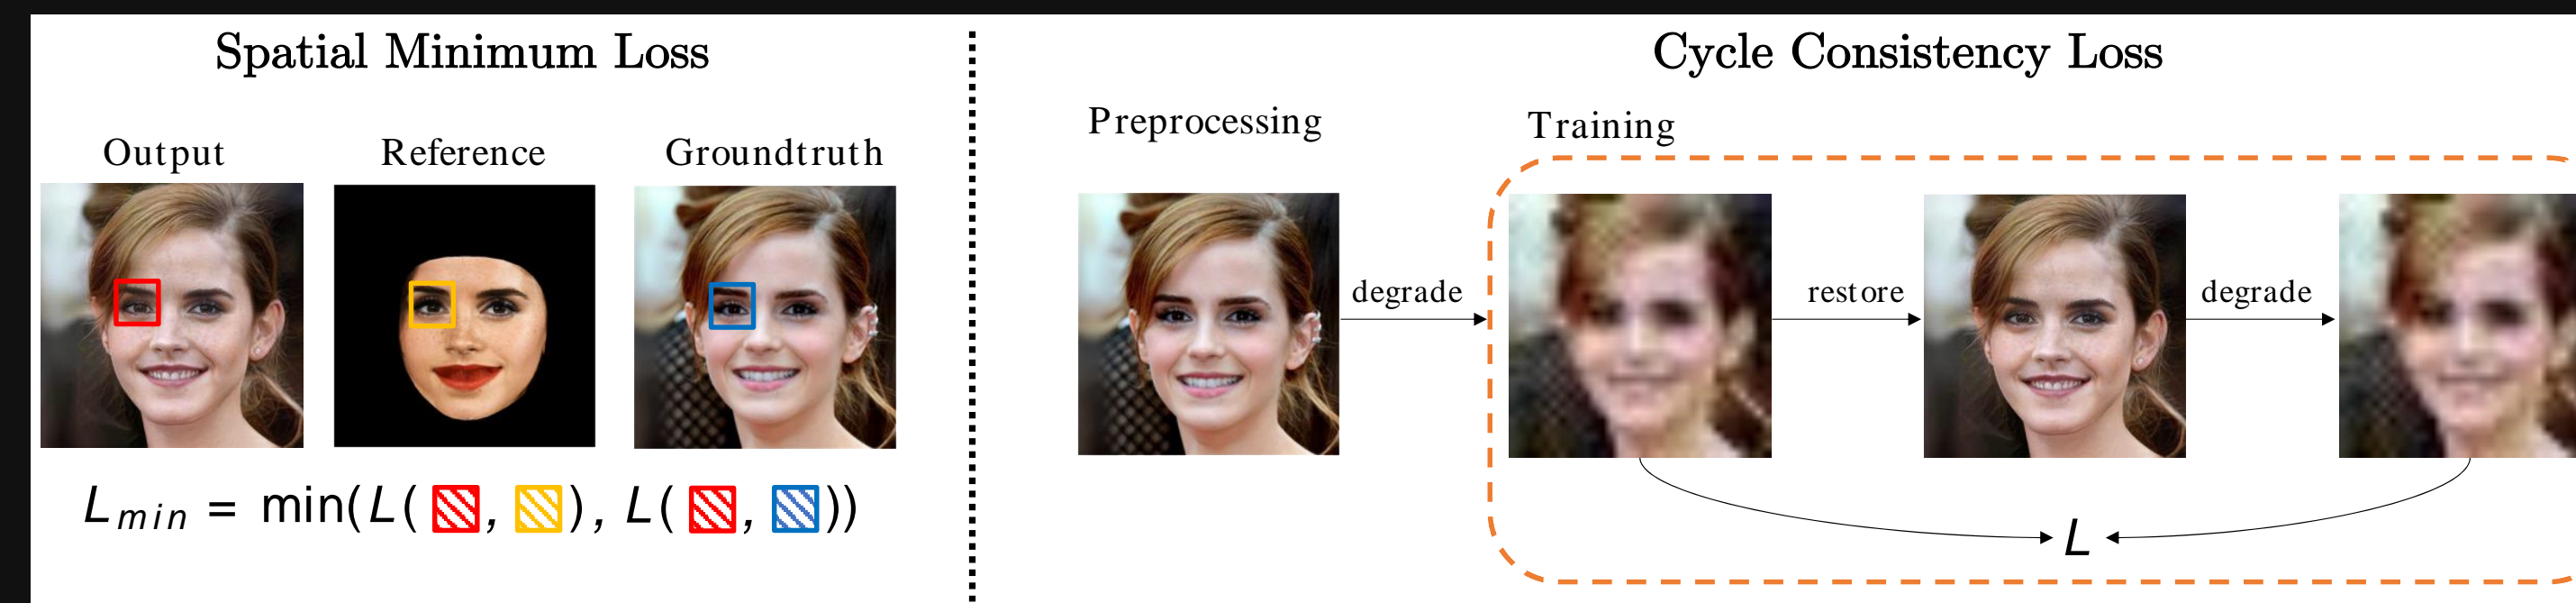


Distilled Student Model
Use Snapchat App to scan it



Credit:
Daniil Ostashev
@ Snap

Our loss



Our framework

