CITeR High-Quality Morph Datasets

Nasser Nasrabadi (JHU), Jeremy Dawson (WVU), Chen Liu (CU), Stephanie Schuckers (CU), David Doermann (UB), Srirangaraj Setlur (UB), Siwei Lyu (UB)

2

Problem

There is a lack of high-quality morph face image datasets that can be used to train detectors, evaluate performance, etc.

Research Goals:

- To generate high-quality sequestered face morphing datasets (including manually check/touchup, print & scan, compressed, perturbed) that can benefit and serve as the milestone for developing defense mechanisms against face morphing attacks (NIST FRVT – Morph Detection Report, 06/202023, **https://pages.nist.gov/frvt/html/frvt_morph.html**)
- Generate a high-quality synthetic face database for morph face generation.

Why is this important?

Single and differential morph detectors are used to identify morphed faces during the passport applications and at the border-crossing for entry and exit scenario.

Our Approach

Generate NIST-defined Tier 1, Tier 2 and Tier 3 Morph Images

- Generate low-quality to high-quality morph images.

  Assemble several databases with different image resolutions to represent the faces captured during border crossing, webcam, and high-quality mugshot passport photos. Generate morph images using WVU, CU, and BU advanced morphing techniques (e.g., combined land-mark and StyleGAN-based morphing, manipulating StyleGAN-based latent codes for morphing, transformer-based and diffusion-based morphing techniques.

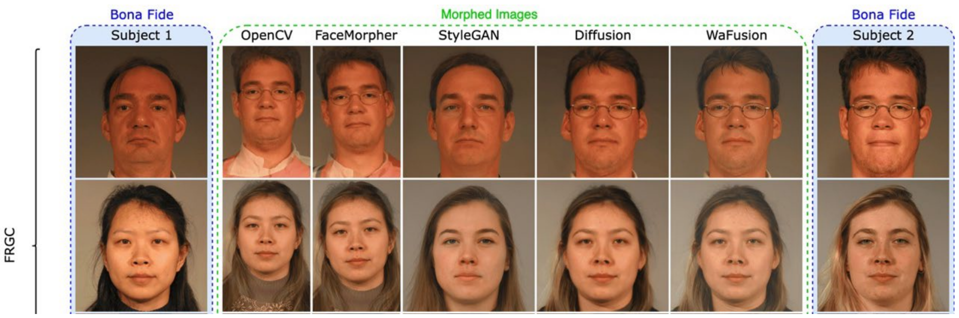

- Generate print & scan (P&S), compressed, perturbed morph images.

  Use several commercial printer & scanners to generate different quality P&S morph images. Generate compressed, adversarially perturbed, and manually touched up morph images. Evaluate the visual and quantitative quality of morph images.

Accomplishments - WVU

Morph Generation Using Wavelet Diffusion Method

- Combines wavelet decomposition's precision and diffusion autoencoders' generative power.
- Two bona fide images are decomposed into smaller sub-bands using wavelet transforms, representing information at different time-frequency, focusing on the low-low (LL) sub-bands for morphing
- LL sub-bands are then blended and reconstructed into the final morphed image using diffusion autoencoders,
- Finally, we use an inverse wavelet transform to generate a high-quality morphed image by combining the morphed image from the LL sub-bands with the high-frequency sub-bands averaged from the original bona fide images
- This approach ensures high perceptual quality while reducing computational complexity.

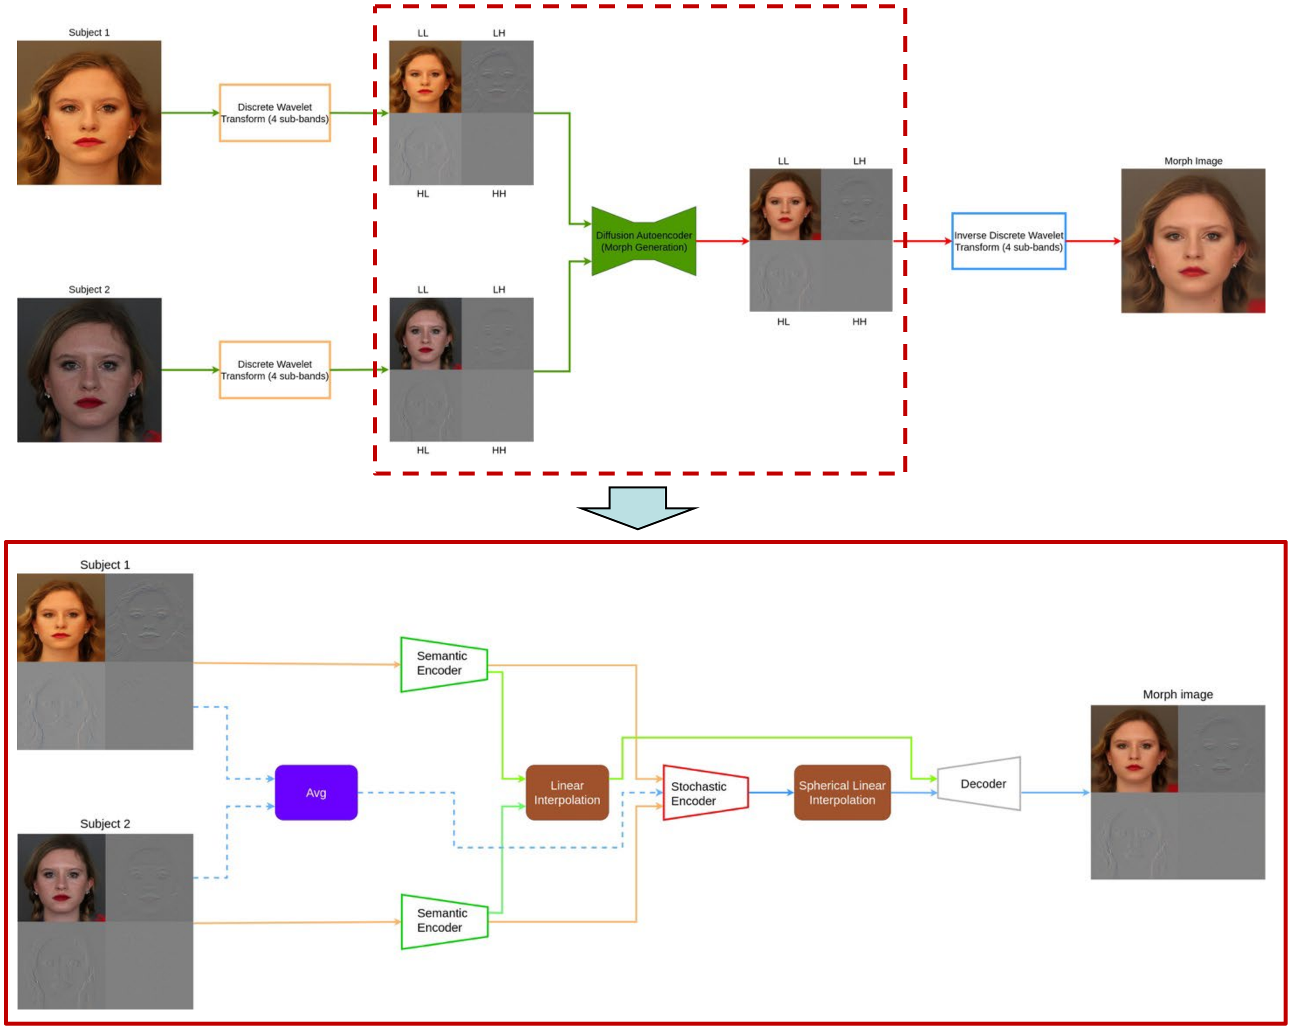

Accomplishments - WVU

Morph Generation Using Wavelet Diffusion Method

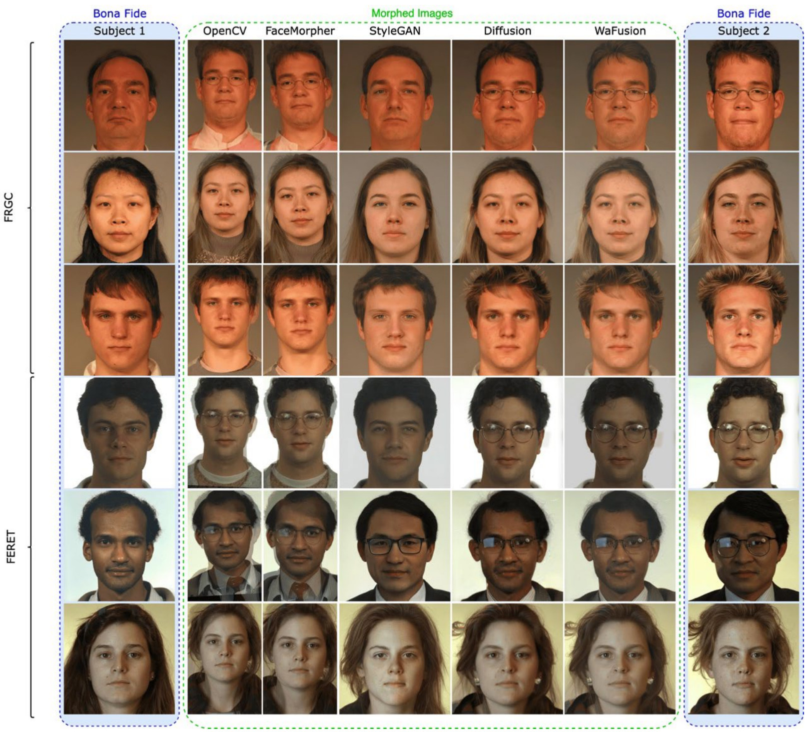

Datasets	Methods	AUC ↓	APCER@BPCER (%) ↑			BPCER@APCER(%) ↑			EER (%) ↑
			5%	10%	30%	5%	10%	30%	
FRGC	OpenCV [29]	0.7267	71.294	61.351	35.647	84.631	71.339	40.291	34.145
	FaceMorpher [39]	0.6735	77.861	69.231	46.341	91.069	80.166	47.04	37.523
	StyleGAN [23]	0.5073	94.746	88.555	68.292	94.496	88.681	71.131	50.281
	Diffusion [37]	0.4955	94.559	86.679	68.292	96.053	91.381	73.624	50.656
	WaFusion	**0.4843**	**95.717**	**89.453**	**71.422**	**96.148**	**92.487**	**75.861**	**52.593**
FERET	OpenCV [29]	0.7295	71.455	60.302	36.862	85.037	72.727	37.31	32.703
	FaceMorpher [39]	0.6798	77.315	68.241	45.557	88.825	76.515	43.56	37.618
	StyleGAN [23]	0.6295	83.742	76.37	52.741	91.477	82.765	50.568	41.398
	Diffusion [37]	0.501	93.383	90.359	67.296	97.348	90.909	69.696	50.85
	WaFusion	**0.4923**	**94.424**	**90.276**	**70.473**	**97.299**	**92.247**	**73.163**	**51.856**
FRLL	OpenCV [29]	0.5741	85.294	73.529	56.862	95.983	91.147	65.573	46.078
	FaceMorpher [39]	0.5848	88.235	77.45	56.862	95.331	88.615	61.179	44.117
	WebMorpher [9]	0.6	80.392	73.529	52.941	94.098	86.803	60.163	43.137
	StyleGAN [23]	**0.4793**	93.137	88.235	**71.568**	**97.788**	**93.529**	74.856	**51.96**
	Diffusion [37]	0.5149	93.137	88.235	61.764	95.61	89.64	70.676	49.019
	WaFusion	0.4909	**94.403**	**90.662**	65.177	97.496	91.891	**75.891**	50.798

[9] L. DeBruine. debruine/webmorph: Beta release 2. Zenodo https://doi. org/10, 5281, 2018.
[23] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4401–4410, 2019.
[29] S. Mallick. Face morph using opencv—c++/python. LearnOpenCV, 1(1), 2016.
[37] K. Preechakul, N. Chatthee, S. Wizadwongsa, and S. Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 10609–10619, 2021.
[39] A. Quek. Facemorpher, 2019.

APCER-BPCER Curves for FaceNet Verifier

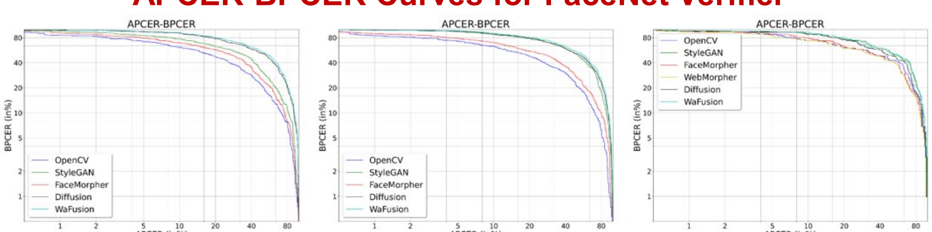

FERET FRGC FRLL

ROC Curves for FaceNet Verifier

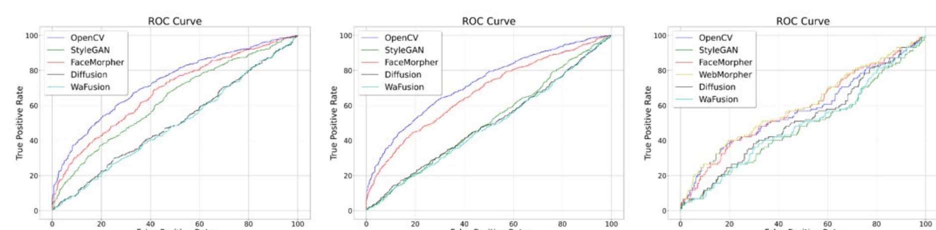

FERET FRGC FRLL

Accomplishments - WVU

Print & Scan

Datasets used include FRGC and FRGC-Compressed, with the following sub-directories:

1. FaceMorpher_StandardMorph
2. FaceMorpher_WaveletMorph
3. OpenCV_StandardMorph
4. OpenCV_WaveletMorph
5. Real Images
6. StyleGAN2_Warped

FRGC fully completed - 14,472 total images

- Printed and scanned on three printers (A-Canon PIXMA iX6820, B-Canon Pixma Pro 100, C-Epson SureColor P700) and scanned on one scanner (EPSON Perfection V700)

FRGC-Compressed dataset – the following fully printed & scanned on Epson SureColor P700 printer & EPSON Perfection V700 - 3,216 images

1. FaceMorpher_StandardMorph
2. FaceMorpher_WaveletMorph
3. OpenCV_StandardMorph
4. OpenCV_WaveletMorph

FRGC-Compressed dataset - fully printed using the following printers - 10,452 total images

- FaceMorpher_StandardMorph, FaceMorpher_WaveletMorph, OpenCV_StandardMorph, OpenCV_WaveletMorph, Real Images printed on the EPSON
- FaceMorpher_StandardMorph, FaceMorpher_WaveletMorph, OpenCV_StandardMorph printed on the iX6820
- FaceMorpher_StandardMorph, FaceMorpher_WaveletMorph, OpenCV_StandardMorph, OpenCV_WaveletMorph, Real Images printed on the Pixma Pro 100

Data-set	Original	Printed & scanned
1A		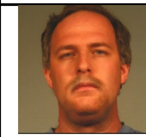
1B	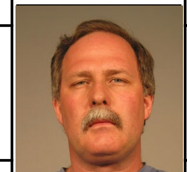	
1C		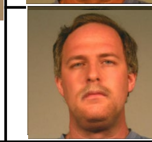
3A		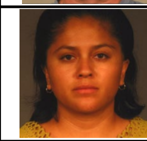
3B	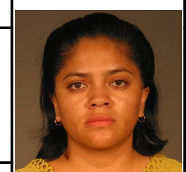	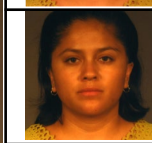
3C		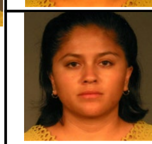
6A		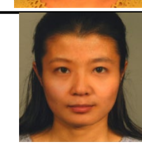
6B		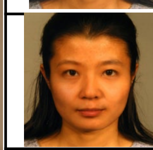
6C		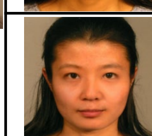

Data-set	Original	Printed & scanned
1		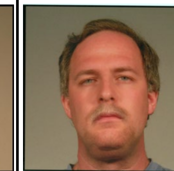
2	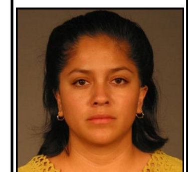	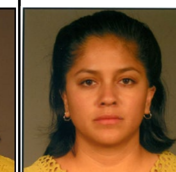
3	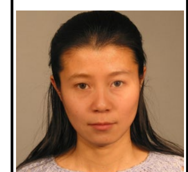	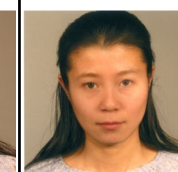
4	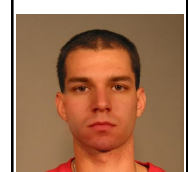	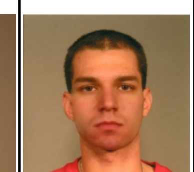

Sample Images

Accomplishments - Clarkson

Fast-DiM - Clarkson University

- Greedy guided generation for unreasonably effective face morphs
- High-order ODE solvers for faster morph creation
- Significant reduction in Network Function Evaluations (NFE)
- Published in IEEE Security & Privacy

Zander W. Blasingame and Chen Liu. “Fast-DiM: Towards Fast Diffusion Morphs”. In: IEEE Security & Privacy 22.4 (2024), pp. 103–114.

(a) identity a (b) DiM-A (c) Fast-DiM (d) Morph-PIPE (e) Greedy-DiM* (f) identity b

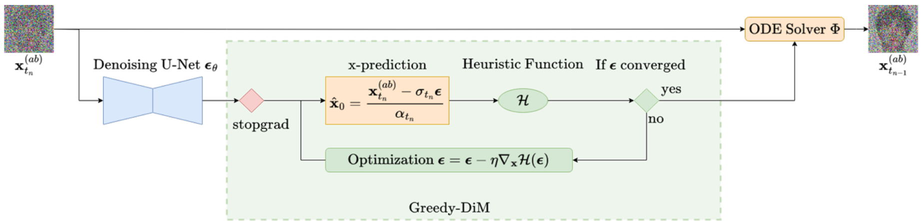

TABLE VI: Vulnerability of different FR systems across different morphing attacks on the SYN-MAD 2022 dataset.

		MMPMR(↑)		
Morphing Attack	**NFE(↓)**	**AdaFace**	**ArcFace**	**ElasticFace**
FaceMorpher [5]	-	89.78	87.73	89.57
Webmorph [5]	-	97.96	96.93	98.36
OpenCV [5]	-	94.48	92.43	94.27
MIPGAN-I [2]	150	72.19	77.51	66.46
MIPGAN-II [2]	150	70.55	72.19	65.24
DiM-A [3]	350	92.23	90.18	93.05
DiM-C [3]	350	89.57	83.23	86.3
Fast-DiM	300	92.02	90.18	93.05
Fast-DiM-ode	150	91.82	88.75	91.21

Accomplishments - Clarkson

Table 5. Vulnerability of different FR systems across different morphing attacks on the SYN-MAD 2022 dataset. FMR = 0.1%.

		MMPMR(↑)		
Morphing Attack	**NFE(↓)**	**AdaFace**	**ArcFace**	**ElasticFace**
FaceMorpher [30]	-	89.78	87.73	89.57
Webmorph [30]	-	97.96	96.93	98.36
OpenCV [30]	-	94.48	92.43	94.27
MIPGAN-I [17]	-	72.19	77.51	66.46
MIPGAN-II [17]	-	70.55	72.19	65.24
DiM-A [20]	350	92.23	90.18	93.05
DiM-C [20]	350	89.57	83.23	86.3
Fast-DiM [21]	300	92.02	90.18	93.05
Fast-DiM-ode [21]	150	91.82	88.75	91.21
Morph-PIPE [22]	2350	95.91	92.84	95.5
Greedy-DiM-S	350	95.71	93.87	95.3
Greedy-DiM*	270	**100**	**100**	**100**

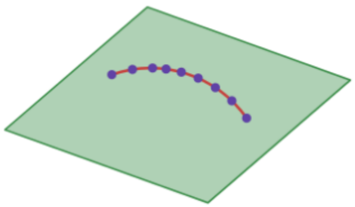

Figure 4. Illustration of the search space in $\mathbb{R}^2$ of different DiM algorithms at a single step. Purple denotes Morph-PIPE/Greedy-DiM-S, red denotes Greedy-DiM-S continuous, and green denotes Greedy-DiM*.

Theorem 1. *Given a sequence of monotonically descending time steps, $\{t_n\}_{n=1}^N$, from T to 0, the DDIM solver to the Probability Flow ODE, and a heuristic function $\mathcal{H}$, then the locally optimal solution admitted by Greedy-DiM* at time t_n is globally optimal.*

Theorem 2. *Let $\mathbb{P}$ be a probability distribution on a compact subset $\mathcal{X} \subseteq \mathbb{R}^n$ with full support on $\mathcal{X}$ which models the distribution of the optimal $\mathbf{x}_0^*$ and is absolutely continuous w.r.t. the n-dimensional Lebesgue measure λ^n on $\mathcal{X}$. Let $\mathcal{S}_P, \mathcal{S}_S, \mathcal{S}^*$ denote the search spaces of the Morph-PIPE, Greedy-DiM-S, and Greedy-DiM* algorithms. Then the following statements are true.*

1. $\mathbb{P}(\mathcal{S}_P) = \mathbb{P}(\mathcal{S}_S) = 0.$
2. $\mathbb{P}(\mathcal{S}^*) = 1.$

- Published in the proceedings of IJCB 2024
- 100% MMPMR on the studied data
- Strong theoretical results to justify the strong empirical performance

Code is available here!

Accomplishments - Clarkson

Print-Scan - Clarkson University

- Generated print-scanned morphs from FERET, FRGC and FRLL for the following morphing algorithms
  - MipGAN2, OpenCV, FaceMorpher, Diffusion, StyleGan2
- Generated print-scanned morphs for the CU Multimodal Dataset for StyleGan2, E4E, and Diffusion
- Total generated morphs: 31,786
- Total generated bona fides: 7,686

Table 3. MMPMR for all scenarios with FMR = 0.1%. A higher MMPMR value represents a stronger attack.

		FRLL			FRGC			FERET		
Morph	**Scenario**	**ArcFace**	**ElasticFace**	**AdaFace**	**ArcFace**	**ElasticFace**	**AdaFace**	**ArcFace**	**ElasticFace**	**AdaFace**
OpenCV	D-D	99.02	**98.69**	**99.26**	67.31	**50.99**	53.22	89.04	75.61	81.78
	D-PS	**99.18**	97.22	99.02	68.91	47.81	53.96	**89.97**	**81.66**	**83.51**
	PS-D	98.61	96.81	97.87	55.67	43.15	45.36	86.45	78.48	81.95
	PS-PS	98.85	94.19	99.02	**69.89**	41.61	**55.51**	88.82	78.58	77.13
StyleGAN2	D-D	5.89	3.27	6.55	**1.38**	1.21	1.25	0.82	0.32	0.72
	D-PS	3.44	**5.56**	4.66	0.67	**1.28**	**1.45**	**0.82**	**0.41**	**1.29**
	PS-D	5.32	1.31	**7.53**	1.00	1.00	0.56	0	0	0
	PS-PS	**6.63**	3.11	6.38	0.41	0.44	1.36	0	0	0
DiM-C	D-D	92.88	82.00	88.22	48.70	**43.24**	41.75	69.76	59.65	65.27
	D-PS	90.10	**88.95**	87.81	43.65	39.23	42.66	71.53	62.39	68.46
	PS-D	92.39	77.09	**91.33**	**49.11**	37.98	35.82	**74.03**	62.21	65.08
	PS-PS	**93.62**	83.22	90.83	37.47	28.30	**44.04**	66.91	**64.20**	**69.99**

Accomplishments - UB

Performance on UB FRGC StyleGAN2 Morphs

Differential Morph attack detection[1]:

- ArcFace for face representation extraction
- SVM-based decision

Good performance for in-domain dataset (left)

- BPCER at APCER = 0.01: 0.51%
- BPCER at APCER = 0.10: 0.00%

Bad performance for cross-domain datasets (right)

- BPCER at APCER = 0.01: 60.57%
- BPCER at APCER = 0.10: 40.50%

Training and testing on FRGC and our morphs

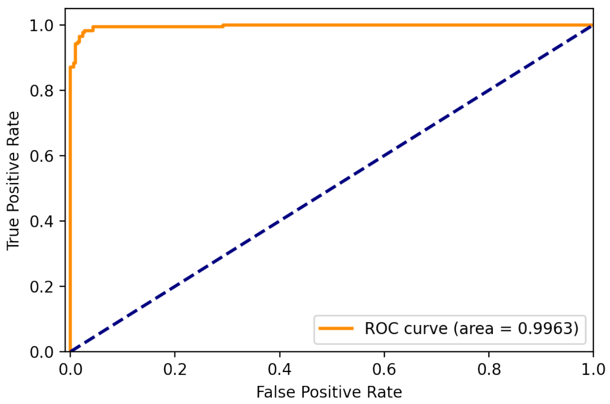

Training on Color FERET and FERET-Morphs, Testing on FRGC and our morphs

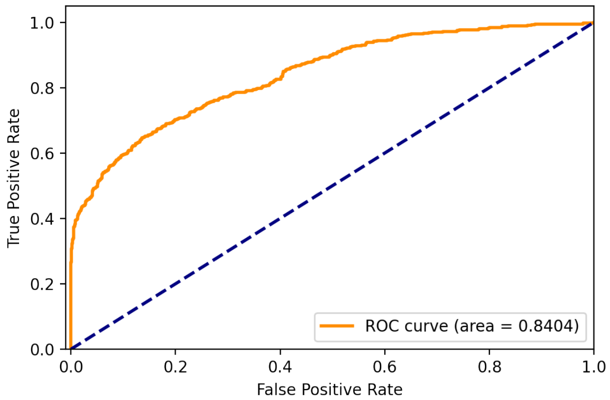

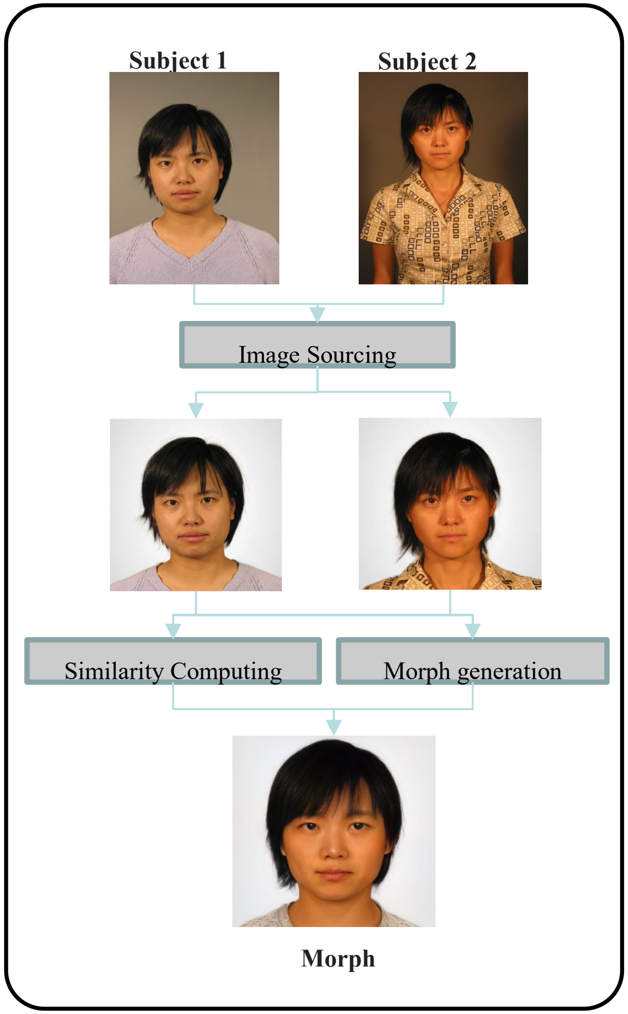

[1] Scherhag, U., Rathgeb, C., Merkle, J., & Busch, C. (2020). *IEEE Transactions on Information Forensics and Security*, *15*, 3625-3639.

Accomplishments - UB

UB Morphs Manual Touch-up

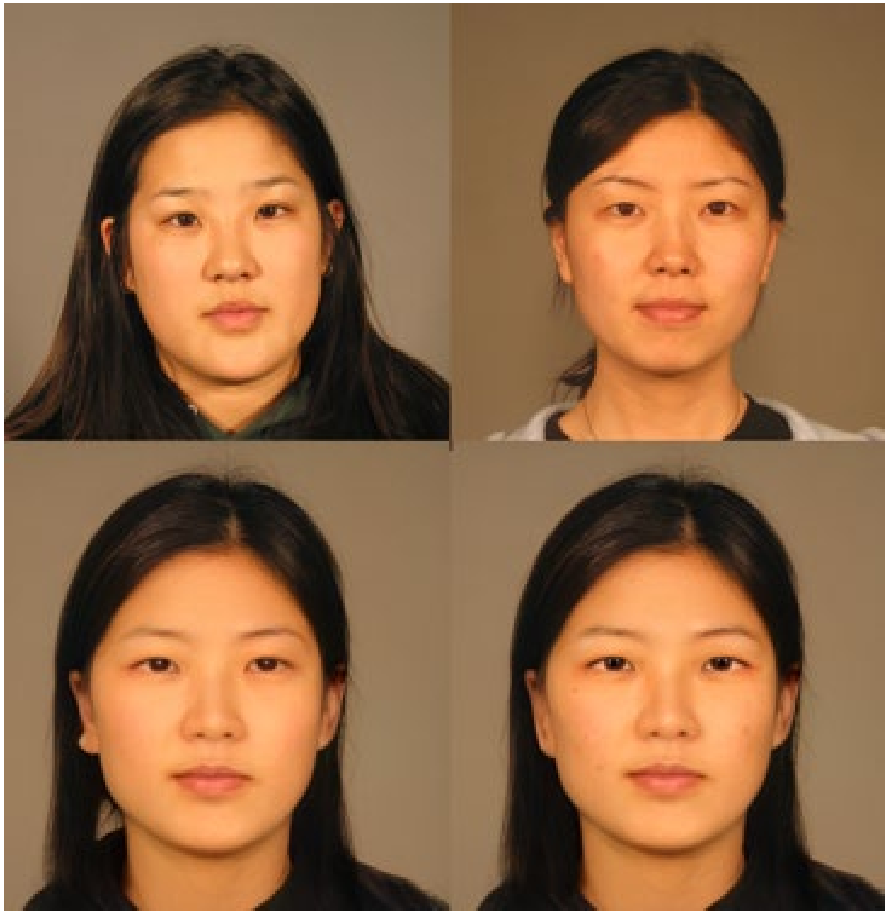

Objectives

- Fix morphing artifacts
- Introduce subtle human-observable features into morphs

Manual edit notes

- used liquify tool to change pupil shape in right eye
- used dodge tool to lighten whites of eyes and pupil reflection
- darkened edges of nose creases and bridge
- brought down nose a little
- used eyebrows from reference photos and blended them to help natural eyebrow shape and look
- used lasso tool to trim small section of the bottom of ear to match other
- copied the right ear and flipped it horizontally to place on left, then blend
- added acne/scarring/dark spots/freckles from reference images

- Need to analyze human subjective qualities of morphs as well
- Metrics are set up for Scenes or Faces, but not both
- General: Brisque,
- Face Specific: FaceQnet, FIQ, SER-FIQ

Kramer, R.S.S., Mireku, M.O., Flack, T.R. *et al.* Face morphing attacks: Investigating detection with humans and computers. *Cogn. Research* **4**, 28 (2019).

Accomplishments - UB

Data Generation -MorDIFF

- Pre-trained Diffusion AutoEncoder (DiffAE) model
  used to combine features of two images and generate a new image.
- Generates 256x256 images natively
- Adapting to generate 1024x1024 images
- Evaluating results on "verifiably close" pairs used to create previous GAN and Transformer images.

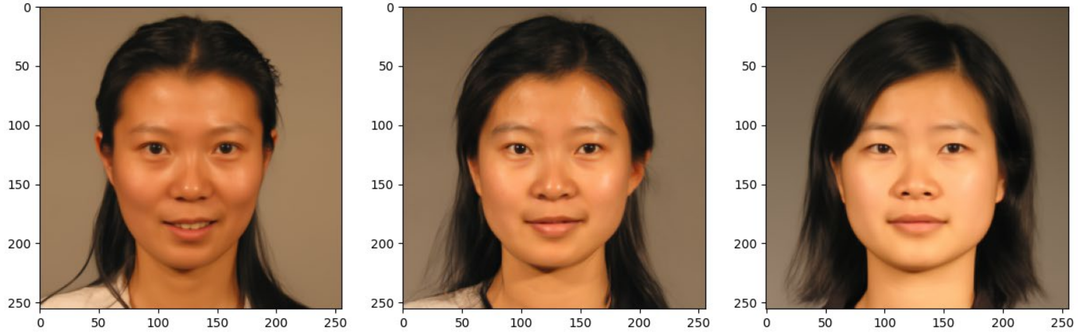

- K. Preechakul, N. Chatthee, S. Wizadwongsa, and S. Suwajanakorn, "Diffusion Autoencoders: Toward a Meaningful and Decodable Representation," presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 10619–10629
- N. Damer, M. Fang, P. Siebke, J. N. Kolf, M. Huber, and F. Boutros, "MorDIFF: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Diffusion Autoencoders." arXiv, Mar. 11, 2023

Accomplishments - UB

Human Patch Quality Dataset

- Example Scoring – five-point scale
  - **Score 1 (Poor):** the image is highly pixelated/blurry, features are indistinct/difficult to recognize, severe artifacts, noise, etc. are present, Severe lighting issues (shadows, overexposure
- Images divided into nine patches that are analyzed without global context
- Currently five annotators
  - **Tasks:** Score 1-5, provide text annotation
- Currently annotating approximately 400 morph/source sets
- ~5% re-annotation for self and inter-annotator assessment

Observations:

- Consistency is reasonable
- Single patch consistency is typically very high
- Source/Pristine consistency tends to be higher
- Morph quality in general is lower than Source/Pristine images

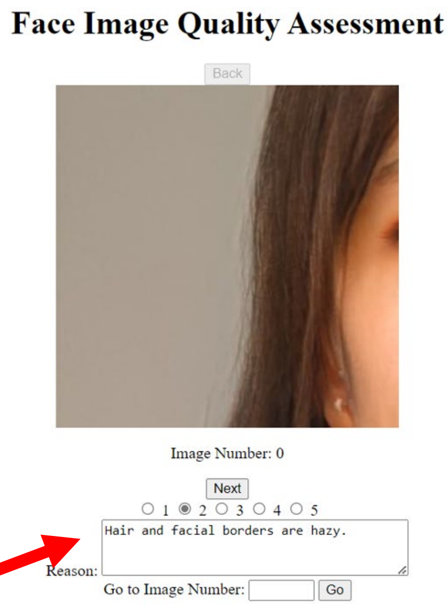

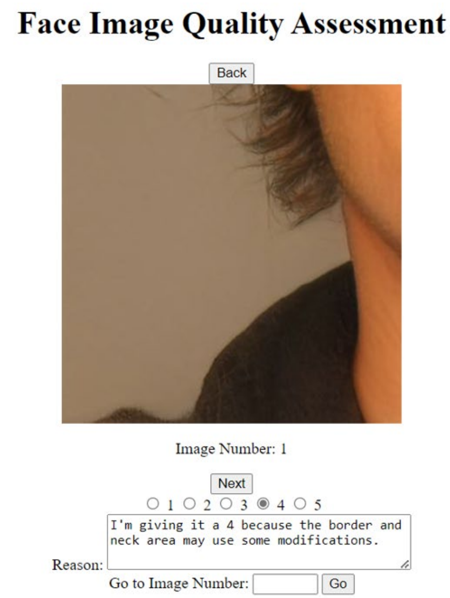

	Annotator 1	Annotator 2	Annotator 3
Annotator 1	1.00	0.58	0.53
Annotator 2	0.58	1.00	0.56
Annotator 3	0.53	0.56	1.00

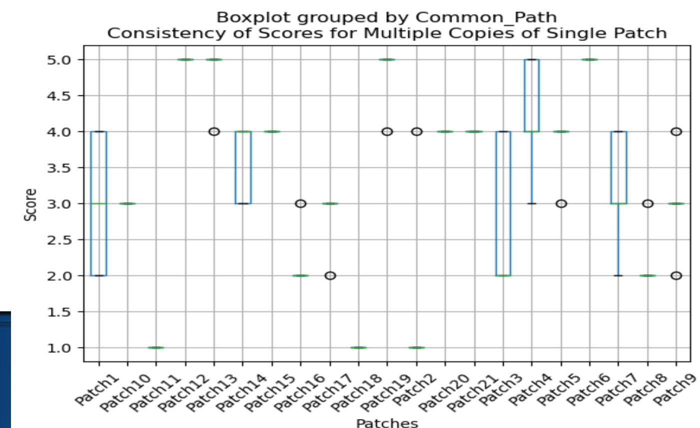

Conclusions

WVU Accomplishments:

- Generated morphs using a Wavelet Diffusion method and evaluated their quality as morph faces; combines wavelet decomposition's precision and diffusion autoencoders' generative
- Generated ~17K Print & Scan morph images using three different printers

CU Accomplishments:

- Developed Fast-DiM, which significantly reduces the computational demand in creating morphs (NFE reduction from 350 to 150) with only a minor loss in MMPMR
- Developed Greedy-DiM, a SOTA face morphing algorithm which beats representation-based AND landmark-based morphs with a 100% MMPMR on the studied FR systems on the SYN-MAD 2022 evaluation dataset
- Generated ~32K Print & Scan morph images; published paper in 2024 IEEE IJCB Proceedings

UB Accomplishments:

- Generated high quality StyleGAN2 morphs based on ArcFace similarity metrics; in domain performance: BPCER at APCER = 0.01: 0.51%, BPCER at APCER = 0.10: 0.00%, cross-domain BPCER at APCER = 0.01: 60.57%, BPCER at APCER = 0.10: 40.50%
- Performed manual touch up of morphs to remove artifacts
- Performed manual patch-based image quality assessment

Questions?

15