Restyling Reality: How the Internet Shapes Facial Perception

Walter J. Scheirer
Department of Computer Science and Engineering
University of Notre Dame

Notre Dame Institute *for*
ADVANCED STUDY

Motivation and Background

Restyling Reality

Meitu XiuXiu App

Memeified Soldier in Ukraine

Are people perceiving their environments differently?

- Push to virtualize as many aspects of life as possible accelerates during the COVID-19 pandemic
- Two cultural undercurrents of the Internet
  - **The Imagination**
  - **Expressive Individualism**
- The average American now spends 4.2 hours per day on their phone[1], compared to 4.7 hours for the average Chinese[2]

https://www.justinpinkney.com/ukiyoe-yourself

1. Perez, Sarah. 2021. “Consumers Now Average 4.2 Hours Per Day in Apps, Up 30% from 2019.” *TechCrunch*. Apr. 8.
2. Liao, Rita. 2019. “Report: Chinese Spend Nearly 5 Hours on Entertainment Apps Daily.” *TechCrunch*. Jun. 12.

Implications of a shift in perception

- People are associating these filtered faces with their identity on the Internet
- Face recognition systems now need to adapt[1,2]
    - If the culture around filtered faces has changed, AI needs to adapt
    - Necessity of virtual spaces

https://blog.zoom.us/filters-reactions-lighting-features-zoom-meetings-2/

[1]Pontus Hedman, Vasilios Skepetzis, Kevin Hernandez-Diaz, Josef Bigun, and Fernando Alonso-Fernandez. 2022. On the effect of selfie beautification filters on face detection and recognition. *Pattern Recognition Letters* 163 (2022), 104–11.

[2]Nelida Mirabet Herranz, Chiara Galdi, and Jean-Luc Dugelay. 2022. Impact of Digital Face Beautification in Biometrics. In 2022 10th European Workshop on Visual Information Processing (EUVIP). IEEE, 1–6.

We have seen this before in human biometrics

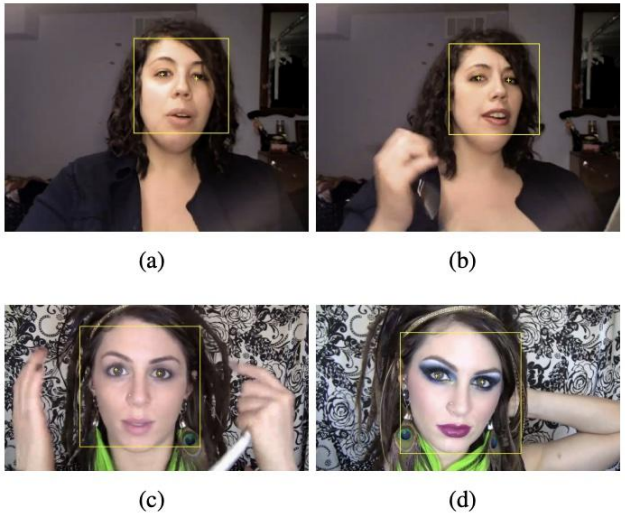

YMU Database[1]

- Similar to cosmetic use in the past[1]
  - Facial recognition systems had to adapt
  - A related scenario is facial recognition with changing facial hair[2]

[1]Antitza Dantcheva, Cunjian Chen, and Arun Ross. 2012. Can facial cosmetics affect the matching accuracy of face recognition systems? In 2012 IEEE Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS), 391–398.

[2]Haiyu Wu, Grace Bezold, Aman Bhatta, and Kevin W. Bowyer. 2023. Logical Consistency and Greater Descriptive Power for Facial Hair Attribute Learning. arXiv preprint arXiv:2302.11102.

By studying people first, we can gain a sense of any perceptual shift, while simultaneously understanding if we need to address this in automatic facial recognition

Historical Context

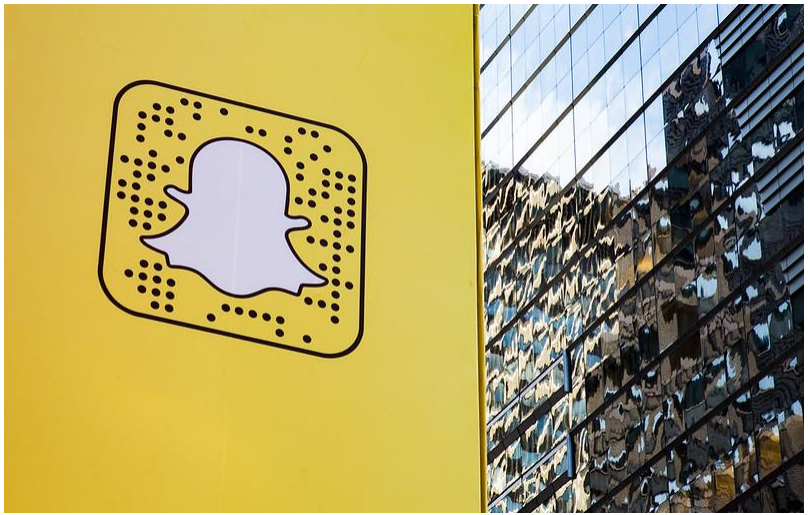

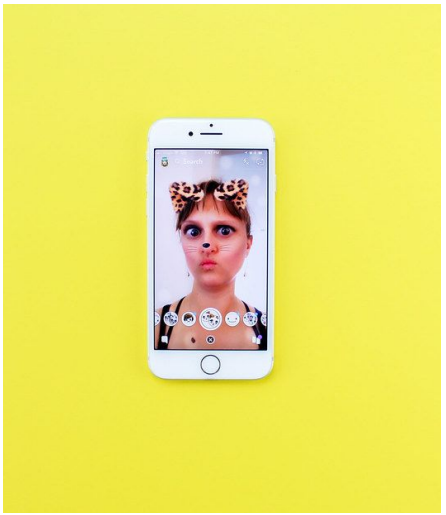

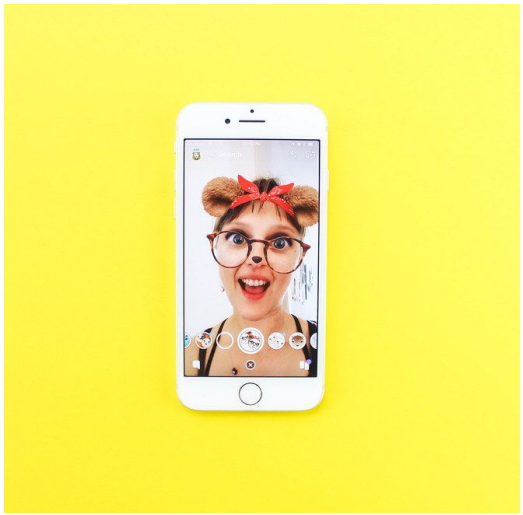

https://www.flickr.com/photos/quintanomedia/33770056655

Both images from http://www.quotecatalog.com

Why is facial filtering so compelling to users?

Facial perception is an important visual skill

- In addition to identifying individuals, the brain's facial recognition abilities allow us to deduce “a wealth of information that facilitates social communication”[1]

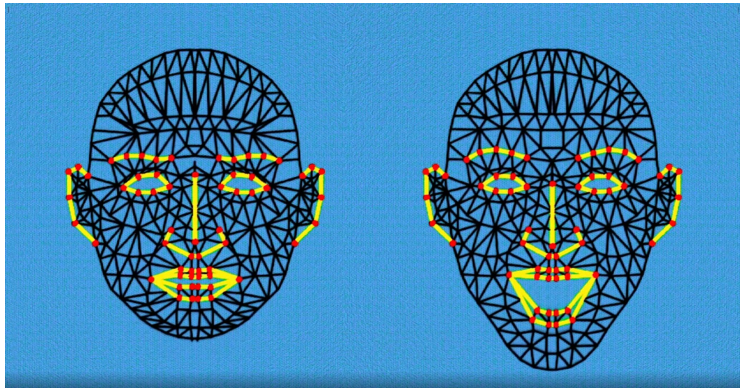

https://www.theatlantic.com/technology/archive/2021/04/artificial-intelligence-misreading-human-emotion/618696/

[1]James V Haxby, Elizabeth A Hoffman, and M Ida Gobbini. 2000. The distributed human neural system for face perception. *Trends in Cognitive Sciences* 4, 6 (2000), 223–233.

Facial perception is an important visual skill

- One’s particular facial traits affect how they are perceived with respect to personality[1]

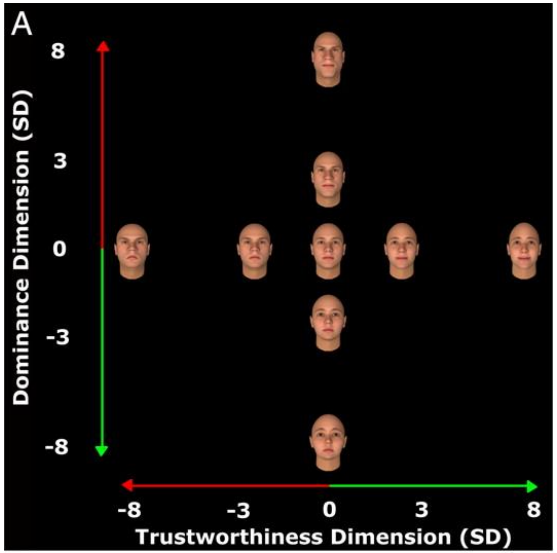

[1]Nikolaas N Oosterhof and Alexander Todorov. 2008. The functional basis of face evaluation. *Proceedings of the National Academy of Sciences* 105, 32 (2008), 11087–11092.

Facial perception is an important visual skill

- The judgments made from facial appearance, even if not accurate, have powerful social impacts ranging from electoral success[1,2,3] to sentencing decisions[4,5]

[1]Alexander Todorov, Anesu N Mandisodza, Amir Goren, and Crystal C Hall. 2005. Inferences of competence from faces predict election outcomes. *Science* 308, 5728 (2005), 1623–1626.

[2]Charles C Ballew and Alexander Todorov. 2007. Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences* 104, 46 (2007), 17948–17953.

[3]Anthony C Little, Robert P Burriss, Benedict C Jones, and S Craig Roberts. 2007. Facial appearance affects voting decisions. *Evolution and Human Behavior* 28, 1 (2007), 18–27.

[4]Irene V Blair, Charles M Judd, and Kristine M Chapleau. 2004. The influence of Afrocentric facial features in criminal sentencing. *Psychological Science* 15, 10 (2004), 674–679.

[5]Jennifer L Eberhardt, Paul G Davies, Valerie J Purdie-Vaughns, and Sheri Lynn Johnson. 2006. Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological Science* 17, 5 (2006), 383–386.

We can now modify how our faces look online

- 90% of young women (in London) have used a filters or edited their photos[1]

https://mymodernmet.com/teen-facetune-selfie-experiment-rankin/ https://medium.com/@markracette/snapchat-s-future-lies-in-augmented-reality-afbfe1834e7a

[1]R Gill. 2021. Changing the perfect picture: Smartphones, social media and appearance pressures. City, University of London (2021).

Social Relevance

- People feel strong connections to external entities[1], including virtual reality avatars[2]

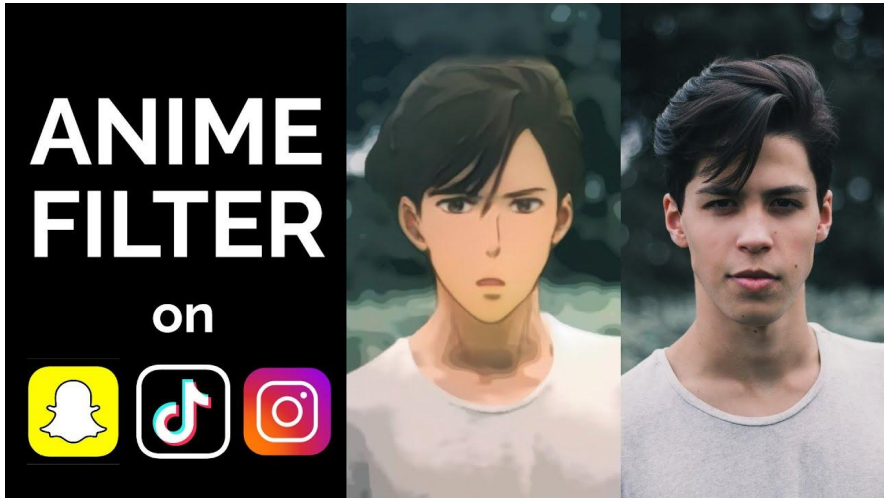

https://www.youtube.com/watch?v=3cBax0JbO4A

[1]Russell W Belk. 1988. Possessions and the extended self. *Journal of Consumer Research* 15, 2 (1988), 139–168.
[2]Russell W Belk. 2013. Extended self in a digital world. *Journal of Consumer Research* 40, 3 (2013), 477–500.

Social Relevance

- Augmented reality usage can change one’s self-perception
  - Users of AR filters may have trouble distinguishing real photos of themselves from photos where their facial features are modified to a preferred size[1]
  - Facial feature modification changed users’ perceptions of the personality traits their modified faces conveyed[2]

https://www.realself.com/news/how-filters-perform-digital-surgery-on-your-face

[1]Fatima M Felisberti and Kristina Musholt. 2014. Self-face perception: Individual differences and discrepancies associated with mental self-face representation, attractiveness and self-esteem. *Psychology & Neuroscience* 7, 2 (2014), 65–72.
[2]Rebecca Fribourg, Etienne Peillard, and Rachel McDonnell. 2021. Mirror, Mirror on My Phone: Investigating Dimensions of Self-Face Perception Induced by Augmented Reality Filters. In 2021 IEEE Intl. Symposium on Mixed and Augmented Reality (ISMAR). 470–478.

Social Relevance

- The less realistic (more “virtual”) a face looks the more likely it is to be perceived differently from how real faces look[1]

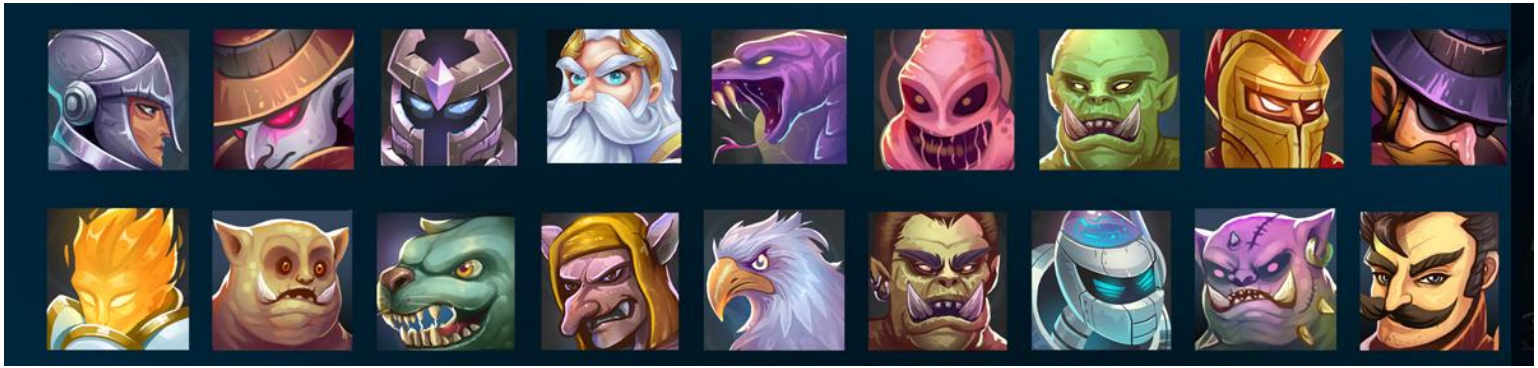

https://gameartpartners.com/downloads/rpg-game-avatar-icons/

[1]Ylva Ferstl and Rachel McDonnell. 2018. A perceptual study on the manipulation of facial features for trait portrayal in virtual agents. In Proceedings of the 18th International Conference on Intelligent Virtual Agents. ACM, 281–288.

Gap in the research we are addressing: have these augmented faces become our new normal?

Is the conditioning from seeing filtered faces online causing us to view filtered faces as normal, or is it changing our perception of human faces in some other way?

Louisa Conwill

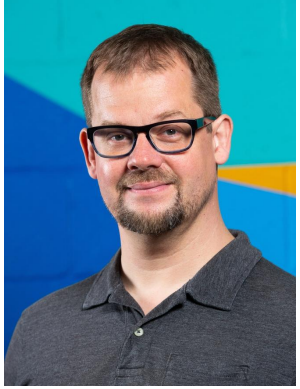

Sam Anthony

Human Behavioral Experiments

Research Questions

- RQ1: Do filtered faces look normal or strange to people? Does this vary by which filter is used?
- Responses to this question could be influenced by familiarity with that style of photo editing filter and ability to tell that a filter has been applied. Thus, we additionally consider the following two sub-questions:
    - RQ2: Do the different styles of facial filters look familiar to people?
    - RQ3: Can people tell that filtered images are digitally altered in the first place?

“Light” Filtering: Portrait Pro Studio Max-Filtered Images

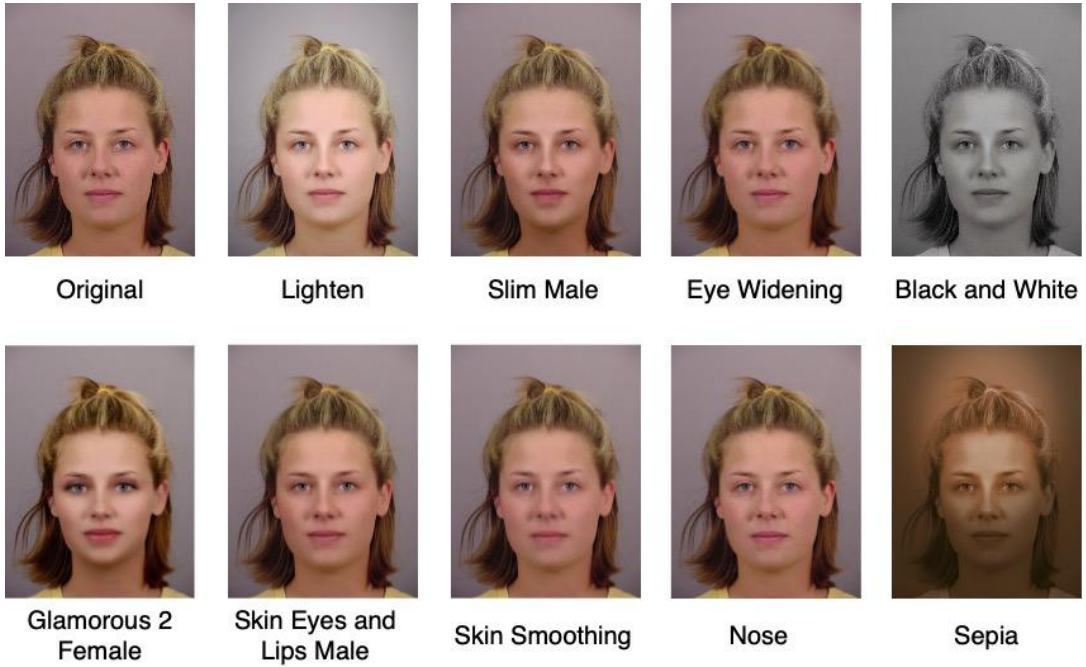

https://www.anthropics.com/portraitpro/editions/

“Heavy” Filtering: Deep AR-Filtered Images

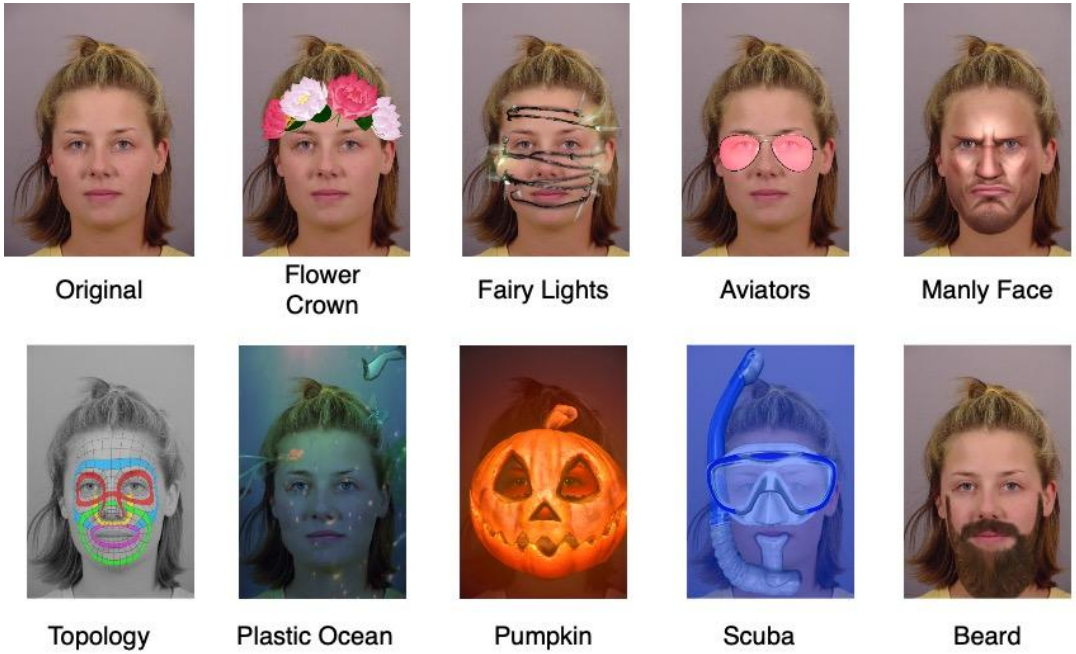

https://www.deepar.ai/

Survey Question Design

- Simply asking if an image looks normal is imprecise due to language ambiguities and confounding factors like familiarity with particular facial filters and ability to tell if an image is filtered
- To address these ambiguities, different surveys asked users if the style of image looked familiar, if the image looked strange, and if the image looked digitally altered.
    - Two-alternative forced choice and Likert scale versions of the surveys were conducted for each question type

Methods – Surveys

- 485 images, filtered with 9 PortraitPro and 9 DeepAR filters
- Questions:
  - Are you familiar with this style of image?
  - Does this image look strange?
  - Does this image look digitally altered?

2-Alternative Forced Choice

  - How similar is this style of image to what you've seen before?
  - How strange does this image look?
  - How digitally altered does the image of this person look?

Likert

Methods – Surveys

- Amazon Mechanical Turk
- Number of participants per question?
  - Familiar? *517 participants*
  - Strange? *477 participants*
  - Digitally altered? *479 participants*
  - How familiar? *503 participants*
  - How strange? *464 participants*
  - How digitally altered? *465 participants*

Directions

Overview

In this task, you will be asked if you have seen a style of image before. Here are two examples:

Are you familiar with this style of image?

Are you familiar with this style of image?

Your response times will be measured, but you should take as long as you need to provide the answer that accurately reflects what you think.

Structure

There are 100 questions in this HIT assignment, each should take only a few seconds to complete. *Additionally, there will be 5 control questions completely different from the task to ensure you are paying attention.* The control questions will ask you to press a particular letter on your keyboard. The total assignment length is ~10 minutes.

Responses

Please respond via keyboard - **[F] for yes, [J] for no.** These response options will appear above each pair of images as a reminder. Your response will automatically advance you to the next trial.

Practice

Next you will be given a 5-question practice example. You will not be able to continue until you have successfully completed the practice example.

WorkerId: JCBCVUBG

Experiment: 2/100

Are you familiar with this style of image?

Press **[F]** on your keyboard for "yes". Press **[J]** on your keyboard for "no."

Experiment: 2/100

Does this image look strange?

Press **[F]** on your keyboard for "yes, strange". Press **[J]** on your keyboard for "no, not strange."

Experiment: 13/100

Does this image look digitally altered?

Press **[F]** on your keyboard for "yes, digitally altered". Press **[J]** on your keyboard for "no, not digitally altered."

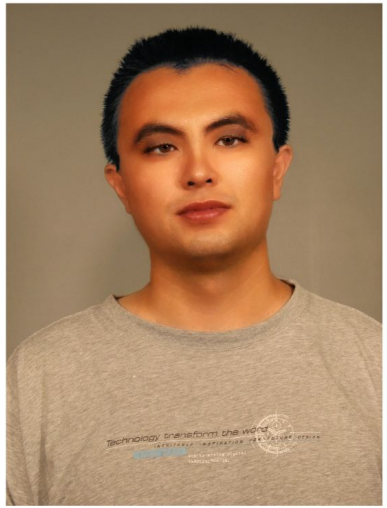

WorkerId: NJBVEQXJ

Experiment: 7/100

On a scale of 1 to 5, how similar is this style of image to what you've seen before?

Press a number from **[1]** through **[5]** on your keyboard where **[1]** means "this is completely unlike any style of image I've seen before" and **[5]** means "this is very similar to styles of images I've seen before."

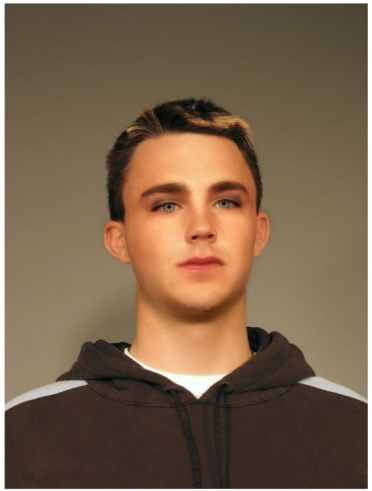

Experiment: 2/100

On a scale of 1-5, how strange does this image look?

Press a number from **[1]** through **[5]** on your keyboard where **[1]** means "not strange at all" and **[5]** means "very strange."

WorkerId: VQVKYPPR

Experiment: 1/100

On a scale of 1-5, how digitally altered does the image of this person look?

Press a number from **[1]** through **[5]** on your keyboard where **[1]** means "not digitally altered at all" and **[5]** means "heavily digitally altered."

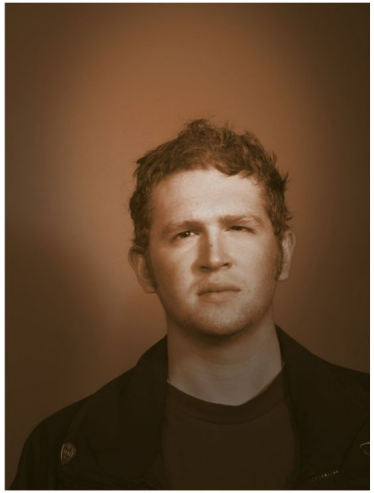

Results

Lightly altered images are familiar, heavily altered are not

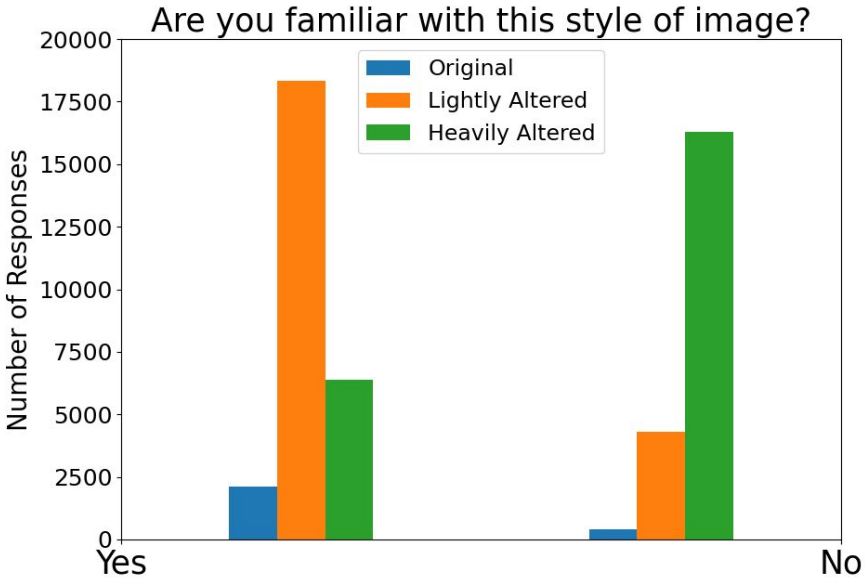

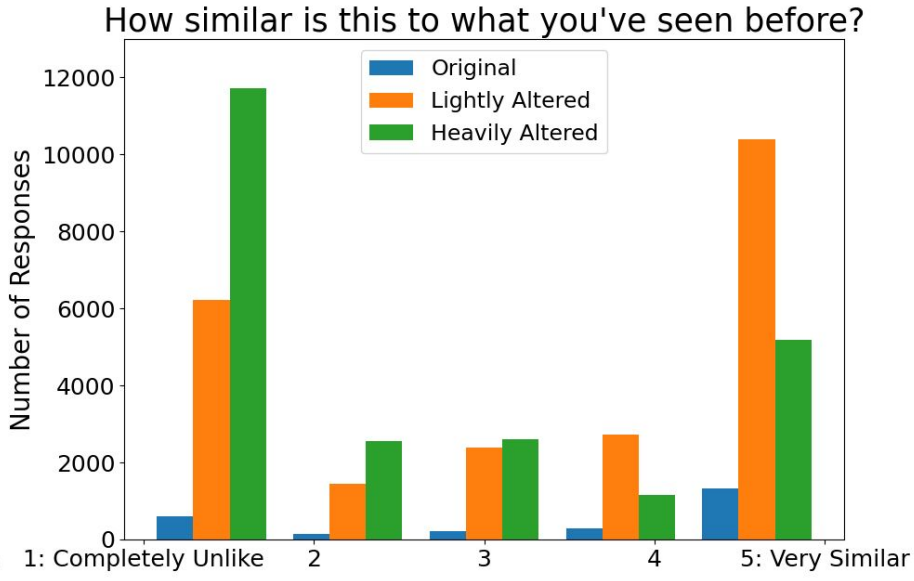

Heavily altered images look more strange

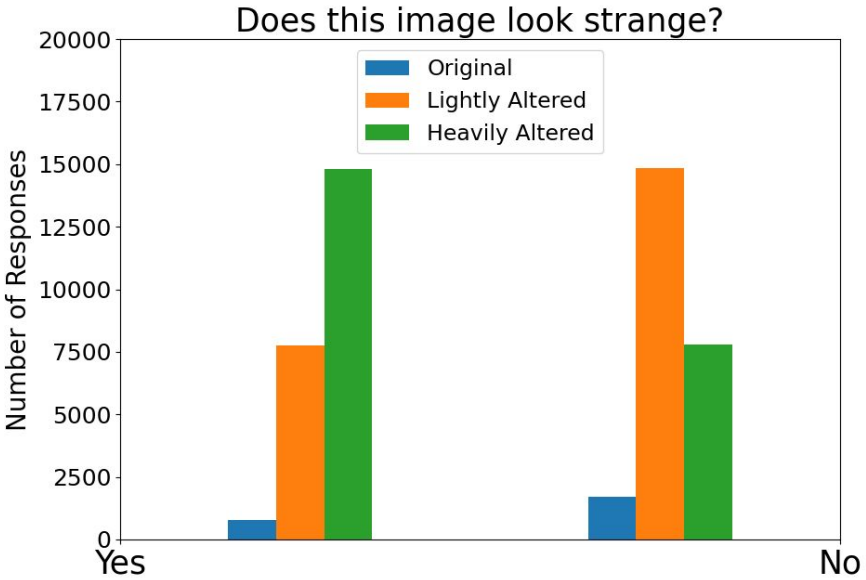

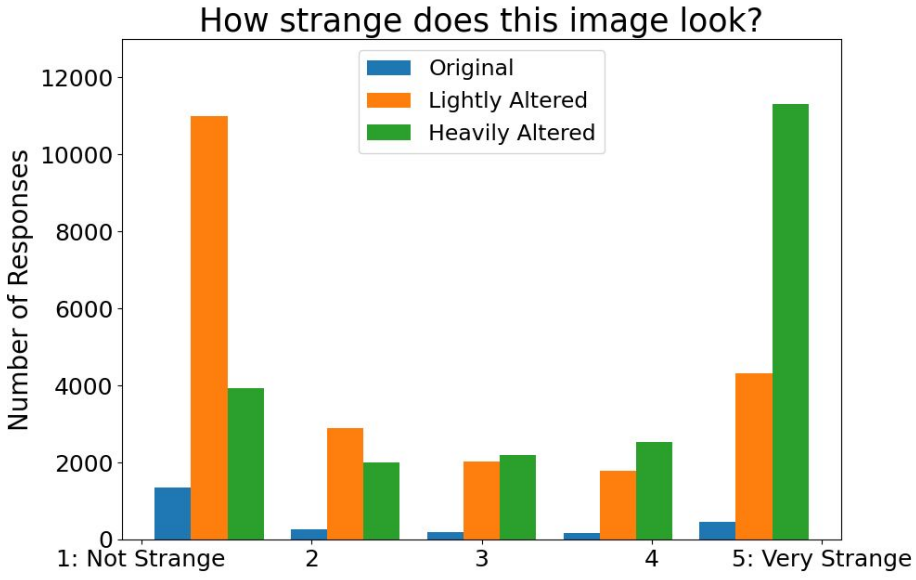

Lightly altered images don’t appear digitally altered, especially on a Likert scale

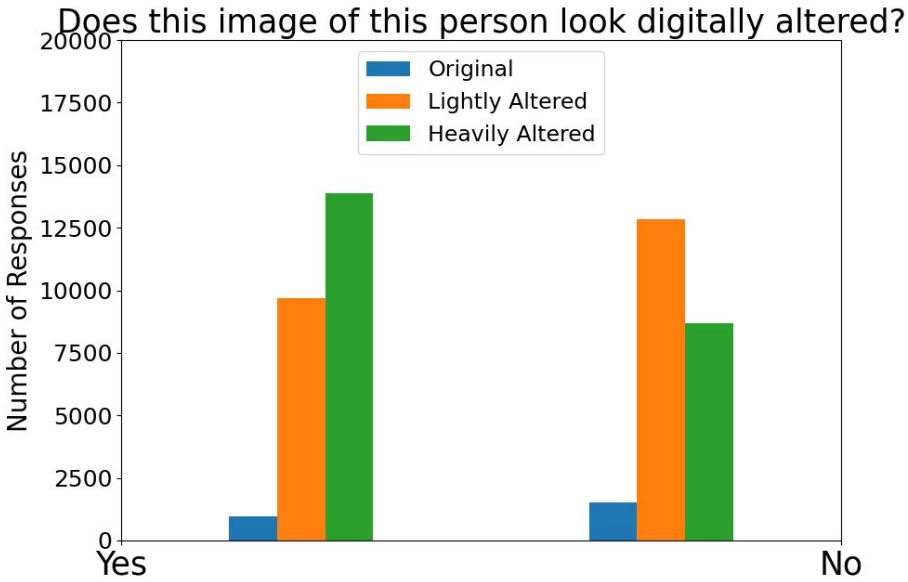

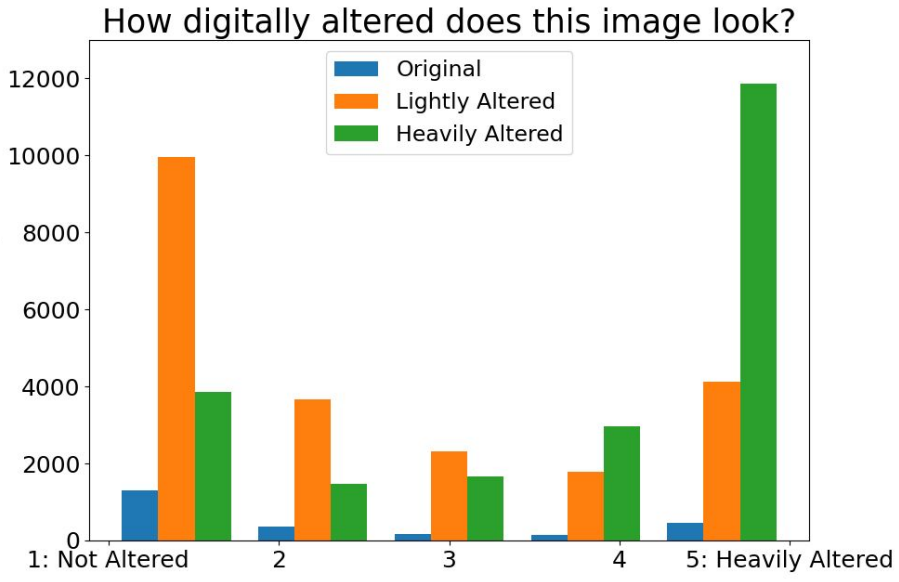

Results by Individual Filter Type

- Conducted individual significance tests comparing each individual filters' responses to the responses for the original images
    - Chi-squared test, $p < 0.05$ for significance
- On every survey, every heavy edit filter had significantly different responses from the originals
- **Light edit filters had mixed results**

Sepia and Glamorous Significantly Different from Original

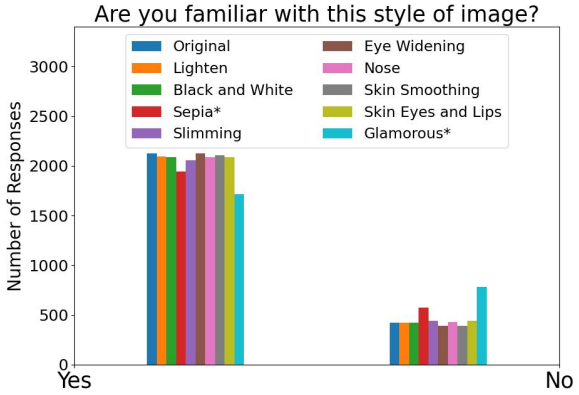

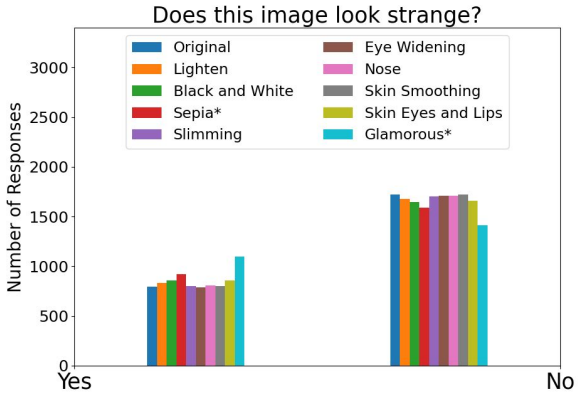

Original

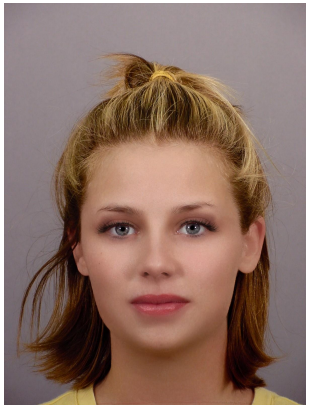

Glamorous

Sepia

Users may not be able to detect some digital alterations

- Lighten, eye widening, nose, and slimming did not have significantly different responses from the originals on the 2-AFC digital alteration survey
- Lighten, eye widening, nose, and skin smoothing did not have significantly different response from the originals on the Likert digital alteration survey
- Because these filters did not elicit significantly different responses from the responses to the original images, users may not be able to detect that they are digitally altered

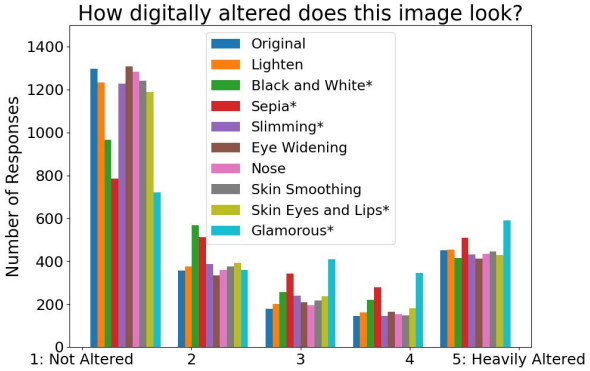

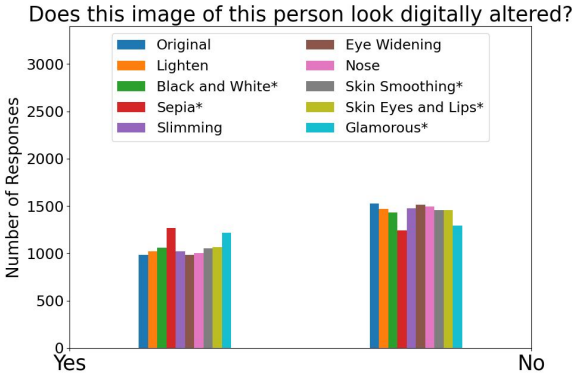

Filters that were not significantly different on digital alterations survey

Original
(for comparison)

Lighten

Eye Widening

Nose

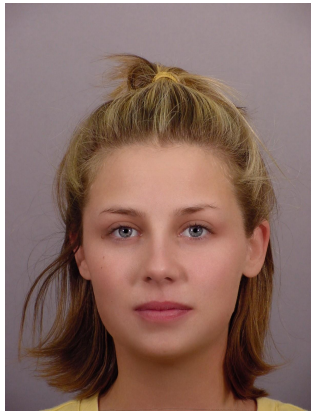

Slimming

Skin Smoothing

Likert version produces more significantly different results

- More significantly different responses than 2-AFC for all surveys, probably because there are more choice options in Likert
- How similar? Slimming, nose, skin eyes and lips **not** significantly different. Every other filter was significantly different
- This is in contrast to the 2-AFC version of the survey where only sepia and glamorous were significantly different

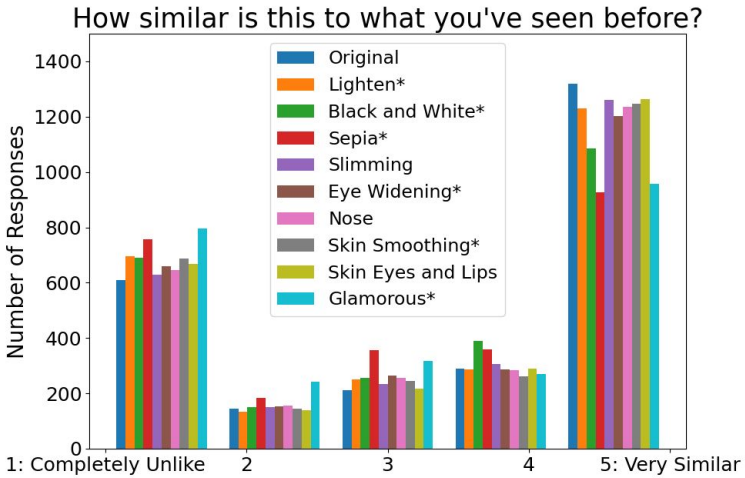

Strangeness Likert survey has one more significantly different filter than the corresponding 2AFC survey

- Black and white, sepia, and glamorous **are** significantly different
- Compared to the 2-AFC version, Black and White joins Sepia and Glamorous as a significantly different filter

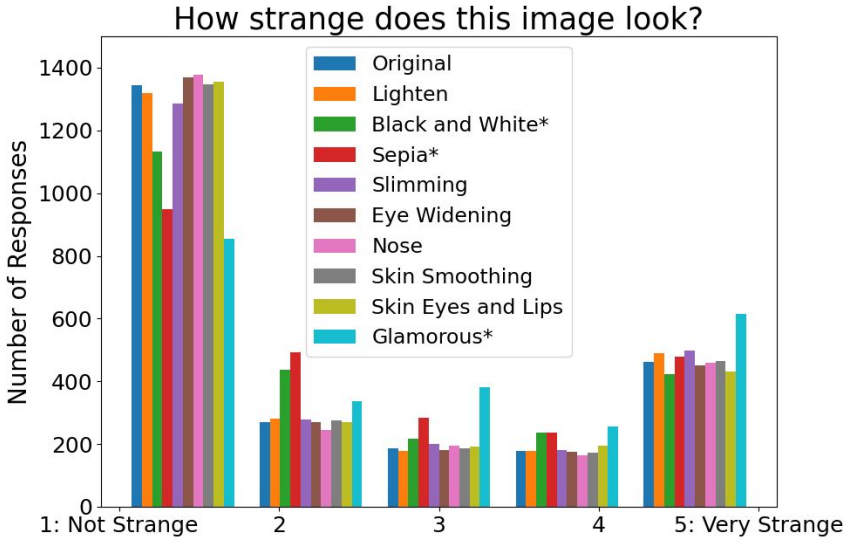

Discussion

Summary of Results

- All PortraitPro filters followed a similar distribution to the original images, DeepAR filters did not, and in some cases even had an inverse distribution
- Images filtered with color-changing, facial structure-changing, or other beautification filters are perceived similarly to unmodified images. Images filtered with AR filters are not
- Users may not realize that lighten, slimming, eye widening, and nose are digitally altered in the first place, because they didn't have significantly different responses from the originals on the surveys asking if they look digitally altered
- Inherently noisy measurements

Creep toward a virtualized self?

Recommendations for Automatic Facial Recognition

- Accommodate the lower-level filters, as these are the filters that people think are more normal
    - In some cases, they make look so normal to users that they cannot detect digital alterations
- Accommodating cartoon-ish filters is not necessary

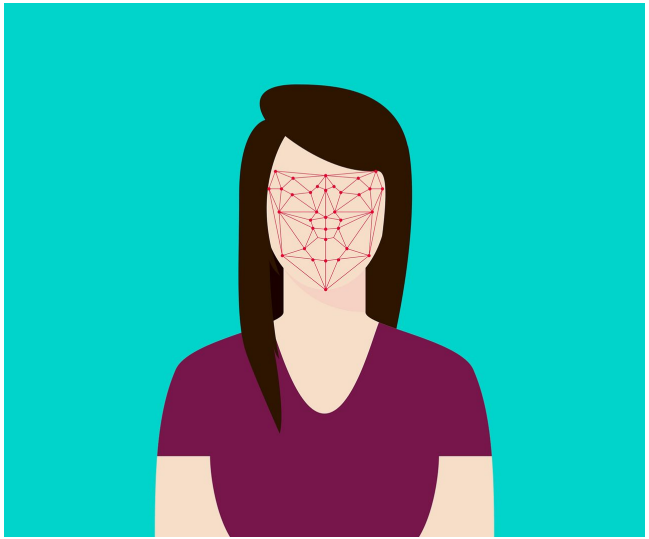

Questions?