Methodology of LooksMapping: Quantifying Superficiality via Neural Networks

walzr.com

March 31, 2025

Abstract

This paper outlines the methodology behind LooksMapping, a web-based analysis that quantitatively evaluates restaurant clientele attractiveness, age, and gender ratios using neural network techniques. Using publicly available profile photos from Google Maps reviews, we applied image processing and machine learning techniques to rate the superficial qualities of restaurant patrons in major US cities.

1 Data Collection

Data was collected from approximately 2.8 million Google Maps reviews across 9,834 restaurants located in New York, Los Angeles, and San Francisco. The reviews were authored by roughly 1.5 million unique accounts. From these accounts:

- Approximately 633,000 had no profile photo.
- Approximately 347,000 did not have a detectable face in their profile photo.
- Approximately 587,000 had at least one detectable face.

Square profile photos were downloaded at 1,000-pixel resolution and processed only if users had set a custom photo (excluding default Google avatars).

2 Image Processing and Face Detection

Face detection was performed using the `face_recognition` Python library, with each photo cropped to the first detected face with a 30-pixel padding. In cases where no face was detected, a center crop was performed. All processed images were standardized to RGB mode to maintain uniformity.

3 Attractiveness, Age, and Gender Assessment

The CLIP model employed was the LAION-trained "CLIP-ViT-B-32-laion2B-s34B-b79K" from 2022 [1]. Each cropped facial image was scored against a predefined set of descriptive phrases to determine attractiveness, age, and gender-related attributes. The phrases included:

- "She is attractive and beautiful"
- "He is attractive and handsome"
- "She is unattractive and ugly"
- "He is unattractive and ugly"
- "A young person"
- "An old person"

CLIP output a probability (0 to 1) indicating how closely each phrase matched the image content.

4 Individual Score Calculation

Relative attractiveness scores were calculated by subtracting each person's "unattractive" CLIP score from their corresponding "attractive" CLIP score. Similarly, relative age was computed by subtracting the "old person" score from the "young person" score. These relative scores were then standardized using Z-score normalization, repositioning the scores onto a Gaussian (bell curve) distribution. The normalization procedure involved calculating the mean and standard deviation for attractiveness and age scores separately, converting scores into Z-scores, and scaling these Z-scores onto a 1-to-10 scale:

$$\text{Normalized Score} = 5 + (\text{Z-score} \times \text{Scaling Factor}) \tag{1}$$

Scores were clipped to ensure a minimum of 1 and maximum of 10 and rounded to the nearest tenth.

Gender identification relied on comparing the CLIP-generated scores for:

- "He is a man"
- "She is a woman"

5 Restaurant-Level Score Aggregation

Restaurant-level scores for attractiveness, age, and gender ratio were calculated by averaging individual normalized scores for each attribute, provided the restaurant had at least 50 detected faces. Specifically:

- **Attractiveness score:** Average of normalized attractiveness scores of reviewers.
- **Age score:** Average of normalized age scores of reviewers.
- **Gender ratio score:** Calculated on a nonlinear scale from 1 to 10, where a score of 1 indicated exclusively male clientele, 10 indicated exclusively female clientele, and 5 represented a perfectly balanced ratio (1:1).

Restaurants with fewer than 50 detected faces were excluded from the analysis.

6 Acknowledgments

We acknowledge MSCHF's HotChat3000 [2] for inspiring the application of the CLIP model for evaluating superficial qualities in this analysis.

References

[1] LAION, CLIP-ViT-B-32-laion2B-s34B-b79K. *Hugging Face*, 2022. `https://huggingface.co/laion/CLIP-ViT-B-32-laion2B-s34B-b79K`

[2] MSCHF, HC3K Model Description. *Hot Chat 3000*, 2023. https://hotchat3000.com/tech-statement.pdf