

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Categorizing Ambiguous Facial Expressions

Permalink

<https://escholarship.org/uc/item/08n1r6sq>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Nudnou, Ilya

Balas, Benjamin

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Categorizing Ambiguous Facial Expressions

Ilya Nudnou (ilya.nudnou@ndsu.edu)

Department of Psychology, 1210 Albrecht Blvd
Fargo, ND 58102 USA

Benjamin Balas (benjamin.balas@ndsu.edu)

Department of Psychology, 1210 Albrecht Blvd
Fargo, ND 58102 USA

Abstract

Categorical perception involves our perceptual system creating sharp boundaries along an objectively continuous stimulus property, such as the discrete colors of the rainbow being perceived despite continuous change in wavelength. The same mechanism is thought to take place in facial emotion perception. But how are emotions at these boundaries perceived? We presented participants with morphed emotional faces made by blending different emotional expressions in equal proportions. Next, we asked participants to respond freely to these ambiguous face morphs and examined these responses via natural language processing methods. The results showed that participants used many more labels than those related to the categories which went into the morphs. These results can inform theories on categorical facial perception as well as the mental representation of facial expressions.

Keywords: facial expression; categorical perception; free response; natural language processing; clustering

Introduction

Facial expressions are frequently described in terms of a small number of universal categories, such as anger, disgust, fear, happiness, sadness, and surprise (Ekman & Friesen, 1971; Ekman et al., 1969). Theories of emotion recognition based on these putative categories generally assume that facial images are evaluated for similarity to a prototype expression for each particular category, leading to the assignment of a category label based on maximum similarity (Fehr & Russell, 1984; Rosch 1999; Shaver et al., 1987). In the strictest version of such an account, any facial expression would be assigned one of the basic emotion labels, with images near the boundaries of category regions being especially sensitive to noise that could lead to variable categorization across observers or across multiple recognition opportunities.

Morphed images of facial expressions are commonly used as a means of exploring the nature of the boundary between candidate categories of facial emotion (Benton 2009; Fujimura et al., 2012; Gray et al., 2020; Korolkova, 2014; Lyons et al., 1998; Harris et al., 2012; Hsu & Young, 2004; Pallett & Meng, 2013; Young et al., 2016). In such studies, facial expressions from different categories (e.g. happy vs. angry) are blended together in different proportions (e.g. 60% happy, 40% angry) and observers categorize or rate faces across the full continuum to allow for a characterization of the transition from one parent category to the other. These

methods are especially useful for evaluating the effects of visual adaptation on facial emotion perception (Benton 2009; Hsu & Young, 2004; Pallett & Meng, 2013), or determining how manipulations like face inversion/negation (Benton 2009; Pallett & Meng, 2013), and image degradation (Lyons et al., 1998) affect the position and sharpness of the boundary between facial emotion categories. In standard forced choice paradigms, participants' responses do not change linearly with the proportion of each emotion present across a morph continuum, as may be expected by a fully low-level image-based account (Harris et al., 2012; Young et al., 2016). Rather, participant responses often resemble a step function, whereby "category A" responses abruptly change to "category B" responses after a specific threshold. This threshold also affects discrimination performance for image pairs across the continuum: Images that straddle this threshold are usually easier to discriminate than images that do not, even when physical similarity is closely matched (Harris et al., 2012). Some image manipulations, like contrast negation, disproportionately affect responses to the most ambiguous images (Pallett & Meng, 2013), suggesting that these boundary regions between categories have different properties than other parts of the space.

This examination of thresholds for facial emotion categorization is also useful in that they provide a way to examine perceptual biases exhibited by an individual. For example, anxious people who need less than 50% Fear in an ambiguous emotional image to categorize it as Fear likely have hypervigilance to fearful expressions (Bishop et al., 2015). Even when perceptual biases are not apparent, individuals may still react differently to emotional thresholds, such as perceiving a higher social cost for interacting with more disgusted faces (Schofield et al., 2007). Additionally, if a lot of negatively-valenced emotions were misclassified as fear, this likely meant that for that individuals, the "fear" label applied to a wider cluster of image prototypes within their mind. A healthy individual may separate the different category clusters according to the actual frequency with which these emotions were experienced by them in everyday life. Changes to a healthy individual's representation would likely occur based on situational or social context, not perceptual biases.

However, a critical assumption of the basic categorical account is the idea that emotion face space is partitioned into a small set of discrete regions (1 per category) and that limiting responses to these categories is an accurate reflection

of how facial expressions are evaluated. In the current study, we examine this assumption by asking participants to categorize ambiguous, morphed expressive face images in a task that allows a wider range of responses. Previous work primarily used forced choice categorization tasks to characterize emotion categorization. Participants would see an ambiguous face, and be presented with a limited number of category labels to pick from. If participants are instead allowed a much broader range of labels to use, what kinds of categorization judgments do we obtain across a morph continuum? More specifically, are ambiguous face morphs categorized primarily according to the parent categories that contributed to the morph, or do unrelated emotional expressions emerge from blending faces in this way? If we observe the latter result, this is a potentially important indicator that our conception of how emotional face space is partitioned is too limited and that constraining participant responses too closely may limit our ability to evaluate the cognitive architecture of emotional face space.

We presented participants with previously validated emotional faces, including ambiguous images made by morphing together different expressions in equal proportions. We asked participants to freely label these faces, and to provide short examples of situations within which making that face would be appropriate. We analyzed participant's responses with natural language processing methods in order to uncover representational clusters formed by the labels. If previous accounts are sufficient to explain category processing, we hypothesized that we would find primarily two clusters for each blended emotion image, and one cluster for each unambiguous emotion image. These clusters may include different words with which a particular category is described, but the distances between these words would be quite small in the representational space. Alternatively, we hypothesized that if the underlying space of emotion categories is more complex, then we should see new clusters representing different emotions in the data obtained from ambiguous images.

Methods

Participants

We recruited a total of 145 participants for the study. Seventy-four individuals participated in the ambiguous category labeling experiment, and 71 individuals participated in the unambiguous (parent) category labeling experiment. We recruited participants from the psychology undergraduate student pool at North Dakota State University (NDSU). Our study was approved by the NDSU IRB, and all participants gave informed consent prior to their participation. Participants were compensated with course credit for participation.

Materials

Stimuli We used images from the Real-world Affective Faces Database (Li & Deng, 2019) for both tasks. This database contains faces expressing the following emotions:




Figure 1. Examples of facial expression stimuli used in the study. The diagonal depicts parent category images. Non-diagonal cells depict blends of category pairs.

anger, disgust, fear, happiness, sadness, surprise, and neutral affect. All images were validated during database creation for expressed emotion by a minimum of 40 voters. We selected two exemplars of each emotion for further testing based on image quality, front-facing viewpoint, and good lighting conditions. We pre-processed the images to match size, inter-ocular distance, and normalize color histograms in CIELAB color space (see supplemental materials on OSF). Examples from the complete stimulus set can be seen in Figure 1.

Category Morphing We used WebMorph (DeBruine, 2018) to morph between all images within and across each category. We selected the physical/numerical midpoint images from each morph continuum (50% of each image in the morph). Morphing often produces unavoidable artifacts, such as overlapping shadows and dark spots, which we manually removed via the GIMP graphics application. This procedure resulted in 60 ambiguous paired-category images for the ambiguous labeling task, and six additional within-category average images for the parent labeling task (for a total of 15 images). The complete stimulus set for this study can be found on OSF.

Procedure

Participants performed the experiment online via Qualtrics. After reading the consent form, participants were given the




Figure 2. Plots of average BIC values (mean $\pm$ standard error) as a function of cluster numerosity for our semantic clustering analyses of participants' free response data. The BIC values were produced by a Gaussian Mixture Model when the depicted number of components was specified in the parameters. All plots include a clear minimum, indicating an optimal number of clusters, and the position of this minimum varies across conditions. Both unambiguous parent images considered singly (Parent Categories) and paired with another emotion (Paired Parent Categories) yield smaller optimal cluster numerosity than ambiguous images

option to begin the study. Participants were instructed to label faces according to emotion, and that they were free to use any words they wanted to do so (but to limit the number of words to one or two). On each trial, a face was presented at the top of the screen, with the caption "Which emotion does this face show?". A text box was presented underneath each image to record responses and participants had unlimited time to respond.

Natural Language Processing

To analyze participants' free responses, we used a similarity metric based on the WordNet database. WordNet is a large lexical database of English words that is freely available as part of the Python Natural Language Toolkit library (Bird et al., 2009). In WordNet, single words are represented as synsets, which are based on a hierarchical ontology derived from the word's usage within the database. Synsets can be considered as grouping of synonymous words that express the same concept. We used the first meaning of a word's synset, which avoided including unrelated word usage into our analysis. For example, "shock" is taken to mean "daze", and not "electric shock" or "shock absorber". We computed pairwise word similarity with the Lin similarity metric, which is based on information content (Lin, 1998). Briefly, the metric can be thought as the likelihood of the compared terms being related within WordNet, combined with the frequency of their most common ancestor within the same corpus (for example, the most common ancestor of "fear" and "anger"

could be "emotion"). Only words recognized by WordNet were used in the final analysis. In addition, Lin similarity requires compared words to be in the same form, which required us to convert verbs into nouns. The complete data analysis pipeline has been made available on OSF: https://osf.io/2myzs/?view_only=cbf55c1521af4aee8890713d008c369e.

Results

Category Labels

First, we computed the number of unique words participants used to label parent category and ambiguous category images. On average, 36 unique words ($SD = 11.41$) were used to label parent category images, and 67 unique words ($SD = 5.58$) were used to label ambiguous category images. This difference was statistically significant, $t(19) = 6.33$, $p < 0.01$, with an estimated effect size of $d = 3.47$. Participants do use more unique words to label ambiguous images – but do these unique words represent a substantial proportion of all the labels?

Participants did use the standard category labels to describe both ambiguous and parent images (for example, "anger" for angry faces and anger blends). We found that on average, 28% of the labels ($SD = 11\%$) for parent category images and 26% of the labels ($SD = 11\%$) for the ambiguous category images used the standard wording associated with that category. The following words were used for this analysis:

Table 1. Ambiguous emotion labels nearest to cluster centroids.

Category Pair	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
anger-disgust	disapproval	provoked	nervousness	confused	disbelief	boredom
anger-fear	upset	shocked	concentration	anguish	disaster	sorrow
anger-happiness	lust	funny	dominance	confused	comprehension	gross
anger-sadness	fear	tantrum	boredom	guilt	crying	disapproval
anger-surprise	wild	apprehension	tense	confusion	worried	surprised
disgust-fear	humiliation	shock	surprised	worry	grossed	confusion
disgust-happiness	confused	listening	nervousness	animated	disgusted	laughing
disgust-sadness	weary	surprise	disgusted	gross	upset	pain
disgust-surprise	confused	excitement	grossed	guilt	surprise	disapproval
fear-happiness	content	confidence	awe	confusion	hope	concealment
fear-sadness	surprised	ok	calm	guilt	judging	blank
fear-surprise	shock	worry	horror	embarrassed	stressed	scarred
happiness-sadness	exhaustion	betrayed	chill	pleasantness	confused	grossed
happiness-surprise	laughing	pride	chill	joy	traumatized	surprised
sadness-surprise	silliness	panic	alarm	confused	surprise	upset

Anger – “*anger*”, “*angry*”; Disgust – “*disgust*”, “*disgusted*”; Fear – “*fear*”, “*afraid*”, “*fright*”, “*frightened*”, “*scared*”; Happiness – “*happiness*”, “*happy*”, “*joy*”, “*joyful*”; Sadness – “*sadness*”, “*sad*”, “*sorrow*”, “*sorrowful*”, “*upset*”, “*unhappy*”; Surprise – “*surprise*”, “*surprised*”, “*shock*”, “*shocked*”, “*stunned*”. The difference between ambiguous and parent label percentages was not statistically significant. This means that participant used the same amount of standard category terms when describing both types of images, and it is the non-category words which account for the majority of responses. The larger number of unique labels for ambiguous category images suggests that these faces are represented in more variable manner than their parent categories. However, it is possible that the increased number of unique words represents just a larger use of close synonyms, not completely different words. We next examine this possibility within the representational space generated by clustering the scaled word similarities.

Clustering And Multidimensional Scaling

We used the similarity space generated by the Lin metric to scale pairwise similarity values between each word into three dimensions using multi-dimensional scaling. Within that low-dimensional space, we identified clusters via a gaussian mixture model. We determined the optimal number of clusters (mixture components) by minimizing the Bayesian information criterion (BIC) of the clustering solution. Briefly, the BIC value trades off variance explained by the model against model complexity (number of clusters) to yield an optimal number of clusters that captures variability in the data effectively, but is not overfit to the data (Figure 2).

On average, parent image and ambiguous image labels were separated into 28 ($SD = 4.71$) and 55 ($SD = 10.6$) clusters, respectively. This difference was statistically significant, $t(19) = 5.97$, $p < .01$, with an estimated effect size

of $d = 3.31$. This is potentially unsurprising if we assume that ambiguous images are close enough to the boundary between one emotion category and another to be interpreted differently by different participants. If this were the case, ambiguous images may elicit labels consistent with two different emotion categories, while any one unambiguous image would only yield labels consistent with one.

We examined this issue by combining labels generated for pairs of parent categories into one label set, which is analogous to the set of labels generated in response to the ambiguous category blend which contained the same parent categories. We determined the ideal cluster number for this new set of labels using the same iterative process with a gaussian mixture model. On average, paired parent category labels were separated into 50 clusters ($SD = 5.33$), as compared to 55 ($SD = 10.6$) clusters for ambiguous category labels. This difference did not reach statistical significance, $t(28) = 1.74$, $p = .096$. This suggests that at least in terms of the number of modes in the free-response data, ambiguous images and parent images do not differ when terms from both parent categories included in each morph sequence are considered as a unit. That is, the increase in unique labels for ambiguous categories could stem from the usage of terms related to two parent categories, as opposed to just one category in the parent labeling experiment. However, though the number of clusters estimated by this procedure does not differ across these two conditions, the BIC values associated with the best clustering solutions do. (see Figure 2). On average, paired parent category solutions resulted in a BIC value of -2230 ($SD = 472.61$), whereas ambiguous category solutions resulted in a BIC value of -1604 ($SD = 395.59$). This difference was statistically significant, $t(28) = 3.94$, $p < .01$, with an estimated effect size of $d = 1.44$. This suggests that the paired parent clustering solution achieved a better separation of the labels and partitioning of the variance within each cluster (recall that BIC values are minimized) while

Table 2. Paired parent emotion labels nearest to cluster centroids.

Category Pair	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
anger-disgust	disrespected	irritated	anger	tense	crying	disgust
anger-fear	pain	aggravated	hurting	doom	surprised	stunned
anger-happiness	satisfied	content	amusement	annoyance	funny	confused
anger-sadness	constipated	upset	fear	invested	concerned	letdown
anger-surprise	constipated	tense	agony	questioning	disapproval	amazement
disgust-fear	disgust	annoyed	embarrassed	surprised	confused	disagreement
disgust-happiness	disgusted	grief	drunk	cringe	confusion	content
disgust-sadness	disgusted	disbelief	disguised	questioning	disagreement	confusion
disgust-surprise	disgusted	smirk	embarrassed	awe	confusion	mortified
fear-happiness	funny	shocked	alarmed	drunk	affection	worry
fear-sadness	frightened	anger	alarmed	pain	surprised	doom
fear-surprise	frightened	amazement	disturbed	fear	stressed	distressed
happiness-sadness	gross	sadness	content	tired	blank	confusion
happiness-surprise	spooked	fear	bliss	awe	content	tired
sadness-surprise	dismay	help	confusion	awe	embarrassed	frightened

using a similar number of clusters. This is precisely what we would expect if the paired parent labels contained more words closely associated with the same concept, as these words would be partitioned into the same cluster. Therefore, we hypothesize that participants may have generated more unique words which weren't synonyms when labeling the ambiguous category images.

Visualizing Labels Within Clusters Although the BIC is a good criterion for determining the best clustering solution for a set of data, it does not lend itself well to understanding which labels are clustered together. In our next analysis, we selected a theoretically relevant number of clusters (six, the total number of putative emotion categories expressed by the parent images), and performed the gaussian mixture model prediction using this number. Within each cluster, we determined a cluster centroid coordinate which allowed us to characterize each cluster geometrically and in terms of a representative label. This set of labels is shown for ambiguous categories in Table 1, and for paired parent categories in Table 2. While the original parent category terms are sometimes represented in this list, most words reflect emotions which are only marginally similar in meaning to these terms. Therefore, category images are not labeled with just their parent category words – rather they are represented as a number of different, non-overlapping emotions. However, while the specific labels are clearly varied, is there a difference in the overall meaning of these terms?

We examined this difference by computing the largest Lin similarity metric between each centroid label and the classic terms used for the category of that cluster. For example, for the labels of the six “*anger-disgust*” centroids in both tables, we computed the similarity of those labels to “*anger*” and “*disgust*”, and kept the largest value for each centroid. On average, the similarity between the centroid labels and the

closest parent category terms was 40.7% ($SD = 11.7\%$) for ambiguous categories and 42.1% ($SD = 9.3\%$) for the paired parent categories. This difference did not reach statistical significance, $p = .709$. This implies that the labels for both ambiguous and parent categories are clustered around similar terms. When considering the BIC plots from Figure 2, we can conclude that both ambiguous and unambiguous expression representations are organized around similar prototypes, but that the immediate space around the prototypes is sparser for ambiguous faces, since a similar number of ideal clusters explains less variance in ambiguous category labels. Future work is needed to examine the content of specific clusters, and how it is affected by specific expressions in the blend.

Discussion

We were interested in examining the complexity of the facial emotion category space which was based on unconstrained response labels. We found that people used more unique labels for ambiguous categories than parent categories. The clustering solutions for these conditions showed that more clusters were needed to map the ambiguous category labels. This difference persisted after combining pairs of parent categories, which suggests that people used more labels of different emotions rather than close synonyms when describing ambiguous categories. While the number of clusters remained the same in this case, the best solution for the paired parent category labels explained significantly more variability than ambiguous labels as demonstrated by its lower BIC. When looking at the specific words within the clusters, we find emotions which are completely unrelated to either parent category within that blend. A lot of these labels occupy locations which are nearest to a cluster centroid in a six-cluster solution. Therefore, instead of having two similar modes of response per ambiguous category, or having six

modes of response overall, participants use a rich array of different emotion labels.

The idea that facial expressions are represented within a rich mental space is not new in the literature. Originally, emotion researchers found that the classic understanding of emotion categories was incomplete when they evaluated how people used these categories. For example, Russell and Fehr (1994) founds that different subcategories of anger (fury, jealousy) can be activated based on situational context, and people do not agree on these subcategories. That is, these emotion concepts are represented differently for different people based on their rating of their prototypicality. As well, different types of descriptive situations were generated for the same subcategories of anger (as well as anger itself). From our results, we can see that about a quarter of the labels used by our participants agreed on the category of the images, but the majority of labels were different, often substantially so.

More recently, Cowen and Keltner (2020) asked participants to freely label a large number of emotional stimuli, including faces, using dimensional and category terms. They found that faces are represented via as many as 28 distinct categories, with the overall space also organized according to 27 different dimensional metrics. Categories within this space function as clusters with thresholds, while dimensions account for ordering within clusters, as well as between partially overlapping categories. Our results suggest that even a smaller set of naturalistic emotions (those depicting only six categories) are also represented within a rich conceptual space. This space does not only exist in the mind – the brain is able to represent as many as 80 emotion categories and 25 dimensions via different activation patterns (Koide-Majima et al., 2020). In addition to the potential complexity of the emotional space, individuals are able to quickly modify their perception of emotions based on situational or social context. For example, Plate et al. (2019) and Woodard et al. (2021) were able to modify participants' response thresholds based on the proportion of "upset" emotions depicted by an actor, as well as information about the actor's traits. Clearly, future work in this area must consider more emotions and facial expressions than can be represented with six categories and their combinations. These representations must also be situated within a social context which provides information about the likelihood of experiencing a specific emotion from a particular individual.

The key contribution of our work is the observation that the specific approach of morphing between unambiguous facial expressions appears to frequently lead to the emergence of face images that are not reliably categorized as either of the parent image categories. To put it another way, rather than crossing a border between one emotional expression and the other, face morphing has the potential to lead us to another region entirely. This has theoretical implications related to our underlying model of the geography of emotional face space and also methodological implications for the use of morphed images in emotion recognition research. Instead of linearly connected categorical "islands", it is possible that our

emotional space is organized in a more complex non-linear manner. A large number of assorted potential emotions may be briefly cycled through by our perceptual system before we arrive at our categorical destination when perceiving an emotional face. In future work, our goal is to further quantify the richness of the conceptual space within which facial expressions of emotion are represented and categorized, ideally using tasks that allow participants to more fully express how they perceive emotions in face images.

References

- Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- Bishop, S. J., Aguirre, G. K., Nunez-Elizalde, A. O., & Toker, D. (2015). Seeing the world through non rose-colored glasses: anxiety and the amygdala response to blended expressions. *Frontiers in human neuroscience*, 9, 152.
- conference on automatic face and gesture recognition (pp. 200-205). IEEE.
- Cowen, A. S., & Keltner, D. (2020). What the face displays: Mapping 28 emotions conveyed by naturalistic expression. *American Psychologist*, 75(3), 349.
- DeBruine, L. M. (2018). *debruine/webmorph: Beta release 2* (Version v0.0.0.9001). Zenodo.
- Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of personality and social psychology*, 17(2), 124.
- Ekman, P., Sorenson, E. R., & Friesen, W. V. (1969). Pan-cultural elements in facial displays of emotion. *Science*, 164(3875), 86-88.
- Fehr, B., & Russell, J. A. (1984). Concept of emotion viewed from a prototype perspective. *Journal of experimental psychology: General*, 113(3), 464.
- Fujimura, T., Matsuda, Y. T., Katahira, K., Okada, M., & Okanoya, K. (2012). Categorical and dimensional perceptions in decoding emotional facial expressions. *Cognition & emotion*, 26(4), 587-601.
- Gray, K. L., Flack, T. R., Yu, M., Lygo, F. A., & Baker, D. H. (2020). Nonlinear transduction of emotional facial expression. *Vision research*, 170, 1-11.
- Harris, R. J., Young, A. W., & Andrews, T. J. (2012). Morphing between expressions dissociates continuous from categorical representations of facial expression in the human brain. *Proceedings of the National Academy of Sciences*, 109(51), 21164-21169.
- Hsu, S. M., & Young, A. (2004). Adaptation effects in facial expression recognition. *Visual Cognition*, 11(7), 871-899.
- Koide-Majima, N., Nakai, T., & Nishimoto, S. (2020). Distinct dimensions of emotion in the human brain and their representation on the cortical surface. *Neuroimage*, 222, 117258.
- Korolkova, O. (2014). Categorical perception of facial expressions is not a homogeneous effect. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 36, No. 36).

- Li, S., & Deng, W. (2019). Blended emotion in-the-wild: Multi-label facial expression recognition using crowdsourced annotations and deep locality feature learning. *International Journal of Computer Vision*, 127(6), 884-906.
- Lin, D. (1998, July). An information-theoretic definition of similarity. In *Icml* (Vol. 98, No. 1998, pp. 296-304).
- Lyons, M., Akamatsu, S., Kamachi, M., & Gyoba, J. (1998, April). Coding facial expressions with gabor wavelets. In *Proceedings Third IEEE international*
- Pallett, P. M., & Meng, M. (2013). Contrast negation differentiates visual pathways underlying dynamic and invariant facial processing. *Journal of vision*, 13(14), 13-13.
- Plate, R. C., Wood, A., Woodard, K., & Pollak, S. D. (2019). Probabilistic learning of emotion categories. *Journal of Experimental Psychology: General*, 148(10), 1814.
- Rosch, E. (1999). Principles of categorization. *Concepts: core readings*, 189, 312-322.
- Russell, J. A., & Fehr, B. (1994). Fuzzy concepts in a fuzzy hierarchy: Varieties of anger. *Journal of personality and social psychology*, 67(2), 186.
- Schofield, C. A., Coles, M. E., & Gibb, B. E. (2007). Social anxiety and interpretation biases for facial displays of emotion: Emotion detection and ratings of social cost. *Behaviour Research and Therapy*, 45(12), 2950-2963.
- Shaver, P., Schwartz, J., Kirson, D., & O'connor, C. (1987). Emotion knowledge: further exploration of a prototype approach. *Journal of personality and social psychology*, 52(6), 1061.
- Woodard, K., Plate, R. C., & Pollak, S. D. (2021). Children track probabilistic distributions of facial cues across individuals. *Journal of Experimental Psychology: General*.
- Young, A. W., Rowland, D., & Calder, A. J. (2016). Facial expression megamix: tests of dimensional and category accounts of emotion recognition (1997). In *Facial Expression Recognition* (pp. 77-120). Psychology Press.