

Learning and Generalizing Associations Between Social Cues and Outcomes

Rista Plate* (rplate@sas.upenn.edu), Emily R. Perkins, Rebecca Waller

Department of Psychology, University of Pennsylvania,
425 S. University Avenue, Philadelphia, PA 19104 USA

Martin Zettersten* (martincz@princeton.edu)

Department of Psychology, Princeton University
Princeton, NJ 08540 USA

*equal authorship contribution

Abstract

To succeed in social situations, we must learn how social cues predict subsequent events. How do we quickly form associations between a variety of social cues, such as individuals signaling their current emotion state, and social outcomes? To address this question, we developed a task in which participants viewed images of individuals conveying different emotions and searched among these images to gain rewards. Rewards were associated with either individuals' identities or emotion cues. Across four experiments ($N=720$), individuals learned about rewards more efficiently from individual identity cues versus a wide variety of emotion cues. Participants also generalized cue-outcome associations more easily for individuals versus emotions. Learning was worse if participants experienced a change in the association rule, especially when switching from learning individual-based associations to emotion-based associations. Overall, we show that social cue type influences how associations between cues and rewards are learned, with implications for understanding learning in social contexts.

Keywords: emotion, associative learning, active learning, multi-armed bandit task, facial cues

Introduction

The ability to make accurate predictions about other people is critical for social functioning. Humans attend to information conveyed by others and use social information to make and generalize inferences about social partners. However, social contexts include many potential sources of information. For example, if a social partner is smiling and offers to buy you coffee, you might form an association between the positive outcome (i.e., coffee) and the *individual*. Alternatively, you might form an association between the positive outcome and the *emotion* cue. Here, we ask how well and how quickly people form associations between social cues and reward outcomes. We focus on two types of cues: (1) emotion cues and (2) individual identity cues.

One source of information that humans track is the emotional signals conveyed by other people. Perceivers use emotion cues (e.g., smiling or looking sad) to make predictions about a social partner's internal state, past experiences, or likely future behaviors (Clément & Dukes, 2017; Walle et al., 2017). While emotion cues can, and do, convey information that informs predictions about others, emotion cues are highly complex. Emotion cues vary based on social partners, situational context, and cultural norms (Barrett et al., 2019). Additionally, myriad perceptual and

contextual features influence the discriminability and interpretation of emotion cues, and people often must make sense of this information without direct labels or feedback (Ruba & Repacholi, 2020).

There is increasing evidence that individuals attend to statistical regularities in others' emotion cues to build knowledge and make predictions (Plate, Woodard, & Pollak, 2022; Wu et al., 2017), suggesting that domain-general learning mechanisms underlie how people form associations between emotional signals and outcomes. However, given the complexity of emotion cues, more knowledge is needed to understand how regularities in the emotion domain are tracked (Ruba et al., 2022) and how we learn to associate emotional signals and outcomes compared to other cue-outcome associations in the social domain.

Often, people learn associations between a social cue and given outcome in the presence of many other social cues. However, little is known about whether a particular social cue (e.g., emotion or social partner identity) is prioritized or more easily associated with outcomes. In the opening example, it might be that you more readily form an association between the positive outcome (coffee) and the *individual* who bought it for you, as opposed to the emotional signal that the individual conveyed. In support of this possibility, people learn and generalize based on an individual's characteristics (Blanco & Sloutsky, 2020), suggesting that identity is a strong learning cue. However, learners may overgeneralize an individual-based association if the association changes over time (Sumner et al., 2019). Additionally, there may be differences in learning in the context of emotion. For example, emotion cues make it easier to track statistical regularities (Plate, Schapiro, & Waller, 2022) and children show more flexibility in learning and updating associations in the emotion domain than in other biologically relevant domains (Plate, Woodard, & Pollak, 2023), suggesting that learning in the context of emotion may be especially efficient. However, emotion cues are heterogeneous (Barrett et al., 2019), resulting in variability that may make it more difficult to form associations between emotion cues and rewards in unsupervised learning contexts (Ruba & Repacholi, 2020).

We developed a multi-armed associative learning task in which participants searched for rewards by clicking on images of different individuals conveying different emotions. Participants completed two task phases: a Learning Phase and a Generalization Phase. In the Learning Phase, rewards were structured to track with either an *emotion cue* (e.g., disgust)

or an *individual's identity*. In the Generalization Phase, participants learned a new set of cue-reward associations that were based on either the same type of cue (e.g., rewards tracked with emotion cues in both phases) or a different type of cue compared to the Learning Phase (e.g., rewards tracked with an emotion cue in the Learning Phase and an individual identity cue in the Generalization Phase).

Aim 1 focused on the Learning Phase, examining how participants form associations between reward outcomes and either emotion cues or individual identity cues. We predicted that participants would learn cue-outcome associations on the basis of each emotion and individual cues; however, we left open the possibility of differences in overall learning and rate of learning between conditions. In Aim 2, we investigated variability in learning across eight different emotion cues. For Aim 3, we examined generalization, specifically whether participants would form new cue-outcome associations faster if the predictive cue (i.e., emotion-based vs. individual-based) is the same as in the Learning Phase (Aim 3A) and whether this effect is moderated by whether the cue-outcome association is emotion-based or individual-based in the Generalization Phase (Aim 3B). Finally, we investigated characteristics of the emotion cue combinations across both the Learning and Generalization Phases, specifically how naming similarity between specific emotion cues affects reward learning (Aim 4). We predicted that emotion cues given similar names by participants would be easier to confuse with one another when paired in the task, making it more difficult to learn associations between emotion cues and rewards. Together, the goal of this work is to advance knowledge about how people form associations between social cues and outcomes to make predictions about others.

Method

The study was preregistered on AsPredicted for two of our four samples (sample 1: https://aspredicted.org/H6R_T5N; sample 2: https://aspredicted.org/blind.php?x=4T3_VB3). The experiment materials, data and analysis scripts are openly available on the Open Science Framework (<https://osf.io/vkd7t/>).

Participants

Participants were 720 adults ($M_{\text{age}} = 37.4$ years, $SD_{\text{age}} = 12.6$, range: 18 – 77 years; 316 female, 392 male, 3 non-binary, 9 did not report gender) recruited across four independent data collections (one sample was recruited from a university participant pool, three samples were recruited from Amazon Mechanical Turk using Cloud Research tools for improving data quality; Litman et al., 2017). We pool these four separate data collection efforts using the same task here to provide a compact overview of the results. The Institutional Review Board approved the research and participants provided informed consent and received course credit or \$2 for their participation. Participants were also given a small bonus of \$0.25 if they scored in the top 50 participants of their sample (split approximately equally across the four main experimental conditions to account for potential differences

in difficulty between conditions). We excluded an additional 14 participants for preregistered criteria, including missing at least one of two attention checks ($n = 10$) and choosing a single image location on more than 80% of trial ($n = 4$). The median completion time for the task was 9.0 minutes.

Design & Procedure

Participants completed the task online (unmoderated). Three aspects of the task varied: (1) task phase (Learning Phase vs. Generalization Phase; within-subjects), (2) reward structure condition (emotion-based vs. individual-based; between-subjects), (3) phase correspondence (match vs. mismatch of reward structure across phases; between-subjects).

In both the *Learning Phase* and *Generalization Phase*, participants saw four images (i.e., four unique individuals conveying four unique emotions) and were instructed to select the image that would provide the most reward (i.e., stars; Figure 1A). After selecting an image, participants were provided feedback on how many stars they received for that choice. On each trial in a given phase, participants saw the same four individuals and same four emotions, but the pairing between emotion and individual varied from trial to trial. Participants saw different individuals and different emotions in the Learning Phase and the Generalization Phase. For example, a participant may have seen individuals 1-4 conveying “surprise,” “calm,” “anger,” and “disgust” in the Learning Phase and individuals 5-8 conveying “sadness,” “happiness,” “fear,” and “exuberance” in the Generalization Phase. Participants completed 48 trials in each phase.

In the Learning Phase, participants were randomly assigned to one of two reward structure conditions: the *emotion-based* ($n = 364$) or *individual-based* ($n = 356$) condition. The conditions varied in how rewards were assigned to cues across trials based on emotion or target individual. Rewards included 2, 4, 6, or 8 stars (with a uniform distribution of noise +/- 2 stars). In the *emotion-based* condition, the reward values were associated with the emotional signals conveyed (Figure 1B, right). For example, for a given participant, “surprise” might be associated with 2 stars, “calm” might be associated with 4 stars, “angry” might be associated with 6 stars, and “disgust” might be associated with 8 stars, regardless of the individual's identity. In the *individual-based* condition, the reward values were associated with specific individuals (Figure 1B, left). For example, for a given participant, individual 1 might be associated with 2 stars, individual 2 with 4 stars, individual 3 with 6 stars, and individual 4 with 8 stars, regardless of the emotion cues.

Participants were also randomly assigned to phase correspondence, in which the reward structure condition assignment across phases was either matched (i.e., a given participant was assigned to emotion-based rewards in both the Learning and Generalization Phases or individual-based rewards in both the Learning and Generalization Phases), or mismatched (e.g., a given participant was assigned to emotion-based rewards in the Learning Phase and individual-based rewards in the Generalization Phase or vice versa).

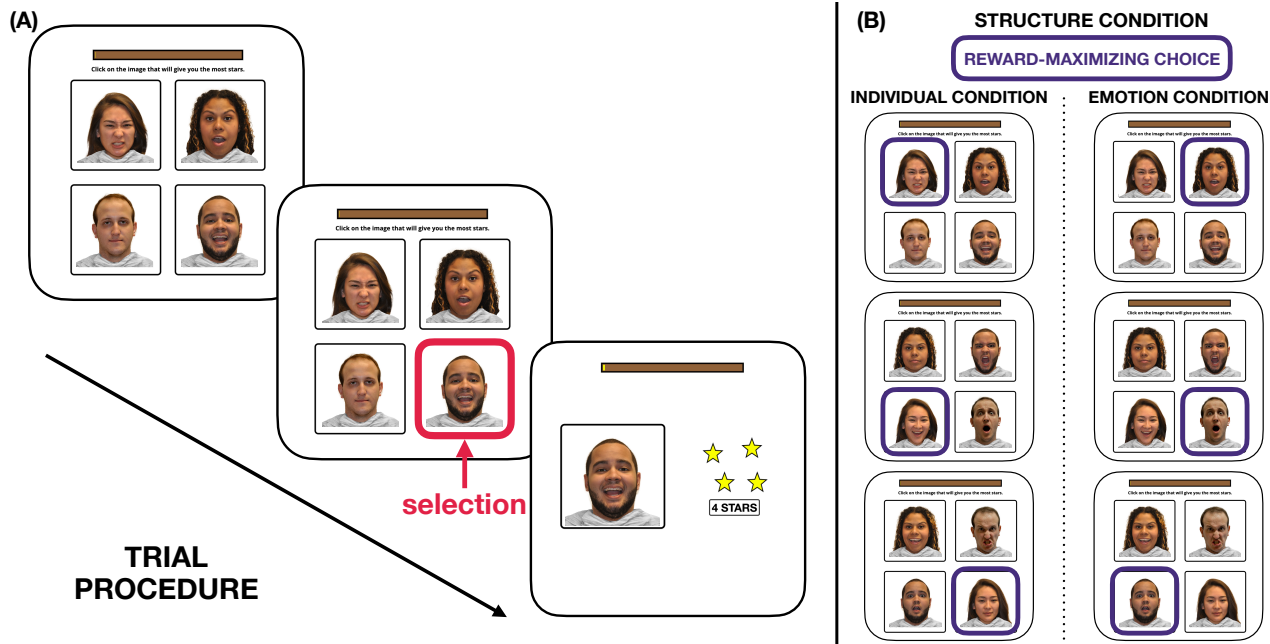


Figure 1: (A) Trial procedure. On each trial, participants could select one of four images and receive reward feedback. (B) Illustration of the reward structure condition manipulation. Each column shows an example of reward-maximizing choices across 3 trials of the individual-based condition and 3 trials of the emotion-based condition. Specific model identities and emotion configurations were randomized between participants.

Together, the reward structure condition and phase correspondence resulted in four between-subject conditions: (1) emotion-based rewards in the Learning Phase and emotion-based rewards in the Generalization Phase ($n = 183$), (2) individual-based rewards in the Learning Phase and individual-based rewards in the Generalization Phase ($n = 178$), (3) emotion-based rewards in the Learning Phase and individual-based rewards in the Generalization Phase ($n = 181$), and (4) individual-based rewards in the Learning Phase and emotion-based rewards in the Generalization Phase ($n = 178$).

Following the experimental task, participants in three of the four samples ($n = 620$) were asked to name (in 1-2 words) the emotion expressed in each of the 64 images that they had seen. We used these naming responses to calculate a measure of *naming similarity* between image pairs. To derive a naming similarity measure, we computed the cosine similarity of the naming responses given to image pairs. For each image, we derived a vector representing the frequency count of each unique (lemmatized) word used to describe the emotion expressed in the image among all naming responses. Then, for each image pair, we computed the cosine similarity between each image's naming vector. This measure ranges from 0 (low naming similarity) to 1 (high naming similarity). For example, if image A was described as “happy” 4 times, “glad” two times and image B was described as “happy” 3 times, “glad” twice, and “content” once, then the resulting naming vectors would be $v_a = (4, 2, 0)$ for image A and $v_b = (3, 2, 1)$ for image B, resulting in a (high) cosine similarity of 0.96. Visual inspection confirmed that the naming similarity

measure captured meaningful variation in the relatedness of emotion cues: image pairs belonging to the same emotion category had high naming similarity ($M=0.89$), while image pairs from different emotion categories had low naming similarity ($M=0.10$). We used this measure to investigate whether higher naming similarity between emotion pairs affected learning in the emotion-based structure condition, by making it more difficult to distinguish emotion cues and pick out the consistent reward-maximizing emotion signal.

Stimuli

We selected 32 actors (four White, four Black, four Asian, and four Hispanic actors, with equal numbers of masculine and feminine gender presentation), each producing eight emotion configurations (anger, calmness, disgust, fear, sadness, surprise, happiness, and exuberance) from RADIATE, a validated stimulus set that includes a diverse representation of actors (Conley et al., 2018; Tottenham et al., 2009). Actors in the stimulus set were provided instructions for how to create each facial configuration and were given time to practice their expression using a mirror (see Conley et al., 2018 for details). The specific actors used were determined according to average reliability, kappa, and validity (i.e., proportion correct) statistics from Conley et al. ($>.70$ in each). If multiple models of the same race and gender met these criteria, we selected the model with the highest proportion correct. Over the course of the task, participants were exposed to eight actors and eight emotions.

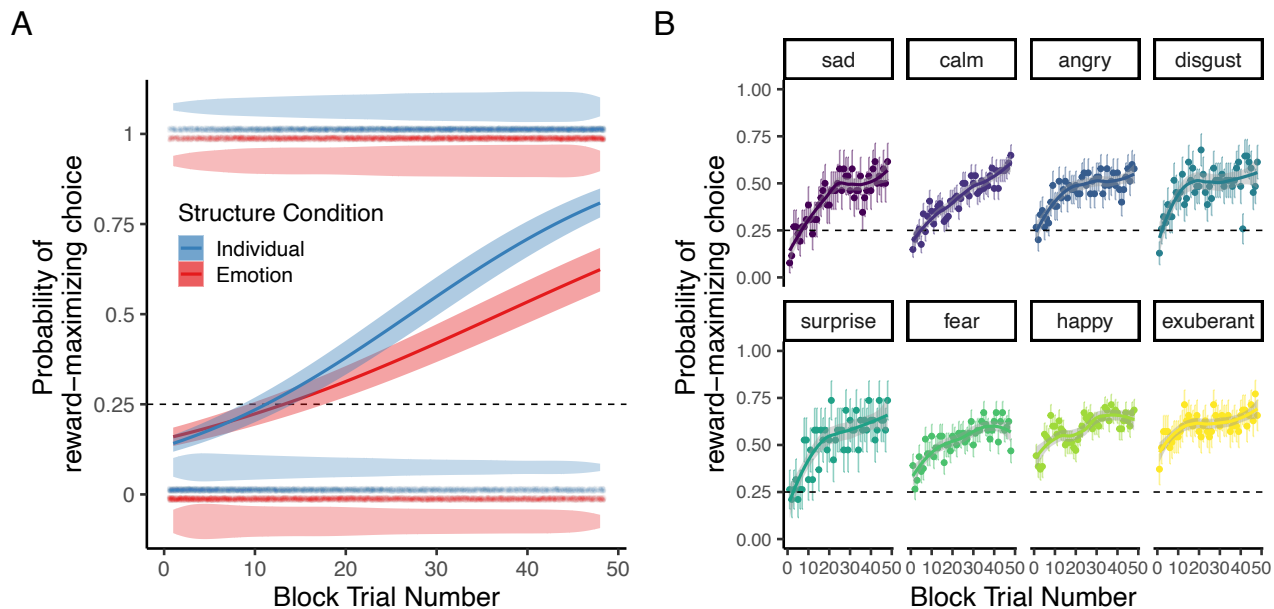


Figure 2: (A) Reward-maximizing choices (i.e., selections of the highest-rewarded image) by structure condition (blue: individual; red: emotion) in the Learning Phase. Smoothed lines represent model predictions with ± 1 SE confidence bands. Individual dots represent individual reward-maximizing or non-reward-maximizing choices on a given trial, with violin plots representing the density of each choice type across trials. The dashed line represents the probability of selecting the reward-maximizing image by chance. (B) Learning trajectory for each emotion cue. Points represent average probability of choosing the reward-maximizing option across participants, with ± 1 SEs. Lines represent smoothed loess fits with 95% confidence bands.

Results

Learning Phase

Overall performance. To compare how participants learned associations between cues and outcomes in each of the reward structure conditions, we fit a logistic mixed-effects model predicting whether or not participants selected the most highly rewarded image (based on the base reward assigned to each image) from structure condition (coded as individual-based = -0.5; emotion-based = 0.5), trial number (centered), and their interaction. We included by-participant and by-image random intercepts as well as by-participant random slopes for trial number. Participants' likelihood of making the reward-maximizing choice increased rapidly across trials, $b = 0.06$, Wald 95% CI = [0.05, 0.06], $z = 20.04$, $p < .001$, demonstrating that participants learned the association between specific cues and their assigned reward across the Learning Phase. Importantly, there was a significant structure condition by trial number interaction, $b = -0.02$, Wald 95% CI = [-0.03, -0.01], $z = -4.06$, $p < .001$. Participants were more likely to select reward-maximizing

choices in the *individual-based* condition ($M = 56.2\%$, 95% = [53.5%, 58.8%]) than in the *emotion-based* condition ($M = 50.5\%$, 95% = [47.3%, 53.7%]), $b = -0.40$, Wald 95% CI = [-0.76, -0.03], $z = -2.12$, $p = .03$, and this difference increased as the Learning Phase unfolded (Figure 2A).

Differences across emotion characteristics. In exploratory analyses, we also investigated whether there were differences in learners' tendency to make reward-maximizing choices depending on the specific emotion cue. Specifically, we fit a logistic mixed-effects model predicting whether participants made a reward-maximizing choice from trial number, emotion (dummy-coded, with "happy" as the reference level), and their interaction, including the same random effects structure as in the model above. There was an overall effect of emotion cue ($\chi^2(7) = 25.44$, $p < .001$), likely due to a bias to select some emotion cues over others (e.g., happy, exuberant, see Figure 2B). However, there was no overall interaction between emotion cue and trial number ($\chi^2(7) = 6.03$, $p = .54$), providing no evidence for differences in how rapidly participants increased their tendency to select reward-maximizing choices depending on emotion cue (i.e., there was no evidence for differential learning rates depending on the specific emotion cue).

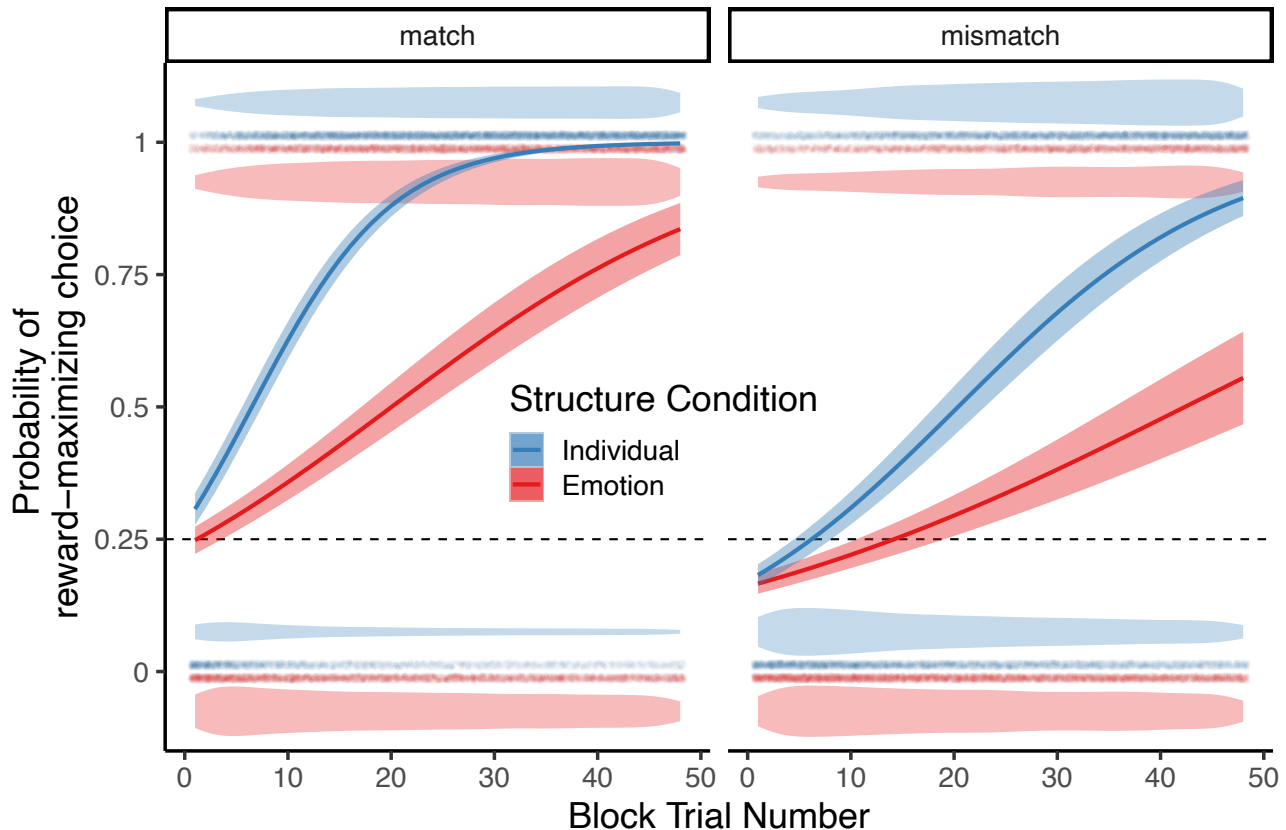


Figure 3: Reward-maximizing choices by structure condition (blue: individual; red: emotion) and phase correspondence condition (match vs. mismatch; faceted) in the Generalization Phase. Smoothed lines are model predictions with ± 1 SE confidence bands. Individual dots represent individual selections and violin plots represent the density of each choice type.

Generalization Phase

Overall performance. To investigate how participants performed in the Generalization Phase in comparison to the Learning Phase, we fit a logistic mixed-effects model predicting whether or not participants selected the most highly rewarded image in a trial from current structure condition (coded as individual-based = -0.5; emotion-based = 0.5), phase correspondence (coded as mismatch = -0.5; match = 0.5), task phase (Learning = -0.5; Generalization = 0.5), trial number (within block, centered), and all interactions between these predictors. We included by-participant random intercepts and by-participant random slopes for block, trial number, and their interaction, as well as by-image random intercepts.

In general, participants increased their reward-maximizing choices in the Generalization Phase (relative to the Learning Phase), $b = 0.51$, Wald 95% CI = [0.30, 0.73], $z = 4.62$, $p < .001$. Moreover, the difference between individual-based and emotion-based structure conditions increased substantially in the Generalization Phase, as indicated by an interaction between structure condition and task phase, $b = -1.40$, Wald 95% CI = [-1.90, -0.89], $z = -5.42$, $p < .001$. As expected,

there was also a general benefit to Generalization Phase performance when structure conditions were consistent across phases (e.g., emotion-based reward in both the Learning Phase and Generalization Phase) compared to when there was a mismatch (e.g., individual-based reward in Learning Phase and emotion-based reward in Generalization Phase), as indicated by a phase correspondence by block interaction ($b = 1.40$, Wald 95% CI = [0.98, 1.83], $z = 6.46$, $p < .001$). Finally, the degree to which matched phases promoted reward-maximizing choices in the Generalization Phase depended in turn on the structure condition, as indicated by a three-way interaction between structure condition, phase correspondence, and task phase ($b = -1.57$, Wald 95% CI = [-2.57, -0.57], $z = -3.07$, $p = .002$). Switching from the individual-based structure condition in the Learning Phase to the emotion-based structure condition in the Generalization Phase was more difficult than the reverse (switching learning from emotion-based cues to individual-based cues) (Figure 3).

Differences based on emotion naming similarity. In exploratory analyses, we investigated whether the naming similarity of emotion pairs impacted learning, with the hypothesis that this effect would be observed specifically in the emotion structure condition. We computed the average naming similarity of the reward image with the other three

image options using the cosine similarity of the word count vectors for image pairs (see Methods). We fit a logistic mixed-effects model across both phases predicting reward-maximizing choice from the interaction of trial number, naming similarity, and structure condition, as well as all lower-order interactions. We included a by-participant random slope for average naming similarity and by-participant and by-image random intercepts. The presence of emotion cues that were described more similarly made participants less likely to select the reward-maximizing choices in the emotion structure condition ($b = -2.29$, Wald 95% CI = [-2.99, -1.63], $z = -6.64$, $p < .001$). There was an interaction between structure condition and naming similarity, suggesting that the effect was substantially larger in the emotion structure condition compared to the individual structure condition ($b = -3.15$, Wald 95% CI = [-3.89, -2.42], $z = -8.41$, $p < .001$). Overall, learners' probability of selecting the reward-maximizing choice in the emotion structure condition decreased when there was higher naming similarity between emotion cues, i.e., when it was likely more difficult to distinguish emotion cues and pick out the consistent reward-maximizing emotion signal.

Discussion

In a newly developed task, we demonstrated differences in how participants learn to associate rewards based on an individual's identity versus specific emotional signals in the presence of both cues. First, learning occurred more rapidly when rewarded outcomes were associated with an individual's identity compared to an emotional signal. This finding could reflect expectations that individual characteristics tend to be stable (McCrae & Costa, 1994; Srivastava, Guglielmo, & Beer, 2010) and/or that emotional information is variable both in terms of how an emotional signal is conveyed (making it a more difficult cue to track) and in terms of correspondence with outcomes more generally (Barrett et al., 2019). This explanation is consistent with our finding that, specifically when learning from emotion cues, image sets with higher naming similarity — and therefore likely containing emotion cues that were more confusable — decreased the likelihood that participants would successfully select the reward-maximizing option. While these differences may reflect higher difficulty in distinguishing emotion cues from individual identity cues, it is important to note that learning robustly occurred in *both* the individual and emotion conditions, including across all emotion signals. Observing similar patterns of learning across varied emotional configurations emphasizes the important role of associative learning in the social domain (Plate, Woodard, & Pollak 2022).

In Aim 3, we assessed generalization and found that learning according to a given social pattern — that rewards tracked with either individuals or emotions — facilitated more efficient learning when the pattern was maintained in a new learning scenario. Switching patterns came with a cost to learning, particularly when initial learning tracked with individual characteristics and subsequent learning tracked

with emotion cues. Testing learning across task phases allows us to advance knowledge about any priors that learners bring to social interactions, the strength of those priors, and how learners flexibly adjust to signals in a learning environment that violate initial expectations.

In developing a learning task to assess rapid acquisition of reward-relevant patterns in a social context, it is important to highlight limitations. First, in curating a controlled experimental setting, we sacrifice features of live social interactions. That is, we used a finite set of static facial images that conveyed posed emotional content. While this approach is consistent with much research in the emotion field of psychology, it limits ecological validity. Second, each participant saw a unique combination of individual faces and emotions. It may be that the contrasts between images held importance in conjunction with, or in addition to, other features of the task. A key future direction is therefore unpacking the influence of task-relevant contextual features. Implementing computational models of choice behavior (e.g., reinforcement learning models) may aid in further investigating task dimensions and cue features that drive differences in learning.

In sum, our findings show that learners robustly pick out and generalize rewarding associations across a wide array of social cues, even as instantiations of these social cues vary across individuals and emotion signals. These skills are critical for our ability to make accurate predictions about others in social contexts. Capturing basic mechanisms of social learning lays the groundwork for understanding why we are so successful — and why we sometimes fail — at interacting with others in dynamically evolving social environments.

Acknowledgments

The research was supported by institutional funding from the University of Pennsylvania awarded to R.W. and postdoctoral fellowships provided by MindCORE at the University of Pennsylvania awarded to R.C.P and E.R.P. M.Z. was supported by an individual postdoctoral fellowship grant from the Eunice Kennedy Shriver National Institute Of Child Health & Human Development of the National Institutes of Health under Award Number F32HD110174. We thank Michaela Flum and Yuheiry Rodriguez for their assistance with data collection.

References

- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest*, 20(1), 1–68.
- Blanco, N. J., & Sloutsky, V. M. (2020). Attentional mechanisms drive systematic exploration in young children. *Cognition*, 202, 104327.
- Chalnick, A., & Billman, D. (1988). Unsupervised learning of correlational structure. *Proceedings of the 10th Annual*

- Conference of the Cognitive Science Society* (pp. 510-516). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Clément, F., & Dukes, D. (2017). Social appraisal and social referencing: Two components of affective social learning. *Emotion Review*, 9(3), 253–261.
- Conley, M. I., Dellarco, D. V., Rubien-Thomas, E., Cohen, A. O., Cervera, A., Tottenham, N., & Casey, B. (2018). The racially diverse affective expression (RADIATE) face stimulus set. *Psychiatry Research*, 270, 1059–1067.
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433-442.
- McCrae, R. R., & Costa Jr, P. T. (1994). The stability of personality: Observations and evaluations. *Current Directions in Psychological Science*, 3(6), 173-175.
- Plate, R. C., Schapiro, A. C., & Waller, R. (2022). Emotional faces facilitate statistical learning. *Affective Science*, 3(3), 662–672.
- Plate, R. C., Woodard, K., & Pollak, S.D. (2022). Statistical learning in an emotional world. In D. Dukes, A.C. Samson, & E.A. Walle (Eds.), *The Oxford Handbook of Emotional Development* (pp. 78-92). Oxford University Press.
- Plate, R. C., Woodard, K., & Pollak, S. D. (2023). Category flexibility in emotion learning. *Affective Science*, 4(4), 722–730.
- Ruba, A. L., Pollak, S. D., & Saffran, J. R. (2022). Acquiring complex communicative systems: Statistical learning of language and emotion. *Topics in Cognitive Science*, 14(3), 432-450.
- Srivastava, S., Guglielmo, S., & Beer, J. S. (2010). Perceiving others' personalities: Examining the dimensionality, assumed similarity to the self, and stability of perceiver effects. *Journal of Personality and Social Psychology*, 98(3), 520–534.
- Sumner, E. S., Steyvers, M., & Sarnecka, B. W. (2019). It's not the treasure, it's the hunt: Children are more explorative on an explore/exploit task than adults. *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 2891-2897).
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., Marcus, D. J., Westerlund, A., Casey, B., & Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, 168(3), 242–249.
- Walle, E. A., Reschke, P. J., & Knothe, J. M. (2017). Social referencing: Defining and delineating a basic process of emotion. *Emotion Review*, 9(3), 245–252.
- Wu, Y., Schulz, L. E., Frank, M. C., & Gweon, H. (2021). Emotion as information in early social learning. *Current Directions in Psychological Science*, 30(6), 468-475.