

Trait knowledge forms a common structure across social cognition

Ryan M. Stolier^{1*}, Eric Hehman² and Jonathan B. Freeman^{3,4*}

Researchers have noted the resemblance across core models of social cognition, in which trait inferences centre on others' intentions and abilities (for example, warmth, competence). Current views posit that this common 'trait space' originates from the adaptive utility of the dimensions, predicting a relatively fixed and universal architecture. In contrast, we hypothesize that perceivers learn conceptual knowledge of how traits correlate, which shapes trait inferences similarly across domains (for example, faces, person knowledge, stereotypes), from which a common trait space emerges. Here we show substantial overlap between the structures of perceivers' conceptual and social perceptual trait spaces, across perceptual domains (studies 1–4) and that conceptual associations directly shape trait space (study 5). Furthermore, we find evidence that conceptual trait space is learned from social perception and actual personality structure (studies 6 and 7). Our findings suggest conceptual trait associations serve as a cornerstone in social perception, providing broad implications for the study of social behaviour.

To navigate an exceptionally complex social world, we ascribe countless traits to one another. Yet, this sea of trait inferences cohere into a small set of dimensions comprising a 'trait space' in social cognition: most often two dimensions, concerning others' intentions (for example, warmth, trustworthiness, communion) or capacity to enact those intentions (competence, dominance, agency; for reviews, see refs. ^{1,2}). This trait space seems conspicuously similar across a variety of distinct domains in social cognition, such as first impressions from faces³, knowledge of familiar others⁴ and group stereotypes⁵. Thus, it has been theorized that social cognition has a fixed architecture structured around a set of universal dimensions, often interpreted to reflect that humans track intention and capability traits given their utility in guiding adaptive social behaviour^{2,6}. While such a process may explain why certain traits are central to trait dimensions (for example, morality to the warmth dimension^{7,8}), recent research has found substantial variation in the dimensionality of trait space⁹, and it is still unclear why the countless traits (for example, sociability, humour, neuroticism, liberalism) correlate along these dimensions as they do. Moreover, it is unclear whether the organization of trait inferences along low-level dimensions is merely an emergent property of social perception (for example, tracking central traits of warmth and competence²), or plays a functional role in forming social perceptions and trait inferences in the first place.

Another possible explanation of a common trait space in social cognition is that perceivers may hold subjective conceptual knowledge of how personality traits correlate in others, which then guides trait inferences similarly across many social cognitive domains. For instance, perceivers may believe kind others are often intelligent, and thus judge a kind face, reputed other or social group to also be intelligent. This would cause trait inferences to correlate similarly across social perception, thus producing a common trait space. Classic research regarding such conceptual associations (that is, implicit personality theory¹⁰) has shown that they influence trait inferences and trait space during impression formation based on vignettes^{7,11}.

Decades of personality research indicates that actual personality traits are in fact correlated along a lower set of dimensions (for example, two-factor and five-factor solutions, such as the Big Five^{11,12}). It may be the case that perceivers learn how personality correlates through various means, such as cultural transmission, direct observations and interactions (for reviews, see refs. ^{10,13}), and apply this knowledge to infer others' traits whether in pursuit of accuracy or due to the inevitable effects of associative processing. This inferential process may be analogous to how perceivers' cognitive models of mental-state associations are applied to accurately predict a target's future mental states based on their current one^{14,15}. Thus, here we explore the possibility that perceivers conceptually learn actual trait associations and use those learned conceptual associations when perceiving others.

Here, we extend this conceptually driven stance on trait inferences to explain a commonality in trait space structure across social cognitive domains (including face impressions, familiar person knowledge and group stereotypes; studies 1–5; Fig. 1), which may be a by-product of applying learned trait knowledge to form initial inferences (studies 6 and 7). This process diverges from a universal account² in that the mechanism underlying trait space structure is not an evolved tendency for tracking functionally adaptive information, but rather reflects general conceptual knowledge about personality. Because such knowledge may differ across individuals depending on their experiences or learning, this perspective also provides a parsimonious account of emerging evidence that trait space changes across individual perceivers and social contexts^{9,16–18}. We find evidence in support of this account across several studies. All stimuli, data and analysis scripts (Python, R) are available on the Open Science Framework (OSF), from which results may be reproduced (<https://osf.io/2uzsx/>).

Results

Study 1. In study 1, we compared models of various social perceptual trait spaces to a model of conceptual trait space (Methods; Fig. 2). Distinct sets of participants reported their conceptual

¹Department of Psychology, Columbia University, New York, NY, USA. ²Department of Psychology, McGill University, Montreal, Québec, Canada.

³Department of Psychology, New York University, New York, NY, USA. ⁴Center for Neural Science, New York University, New York, NY, USA.

*e-mail: rms2262@columbia.edu; jon.freeman@nyu.edu

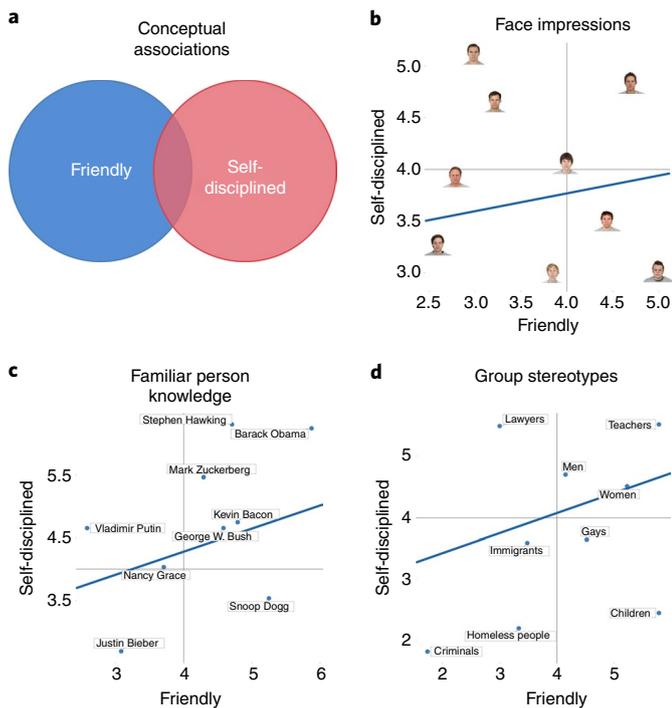


Fig. 1 | Conceptual illustration of theoretical and analytic approach. a–d, We theorize that conceptual trait associations (**a**) shape social perceptions across domains (**b–d**). Thus, our analytic strategy across studies was to evaluate whether the pairwise relationships between personality traits across these domains are similar. For instance, we observe that the conceptual association between two personality traits, ‘friendly’ and ‘self-disciplined’ (**a**), is mirrored by the correlation of ‘friendly’ and ‘self-disciplined’ perceptions of faces (**b**), familiar others (**c**) and social groups (**d**). Please note this figure is for illustrative purposes. It is also important to note that analyses reported do not test the similarity in magnitude of trait-pair correlations between domains, but rather whether the rank-ordering of associated trait pairs is similar between domains (Methods). Panel **a** is conceptually illustrative, whereas panels **b–d** depict a subset of data from study 1, nine data points per panel. Blue lines are predicted values via Pearson correlation of the two axis variables. Several trait spaces and stimulus examples are provided as data points in panels **b–d**. Face stimuli in panel **b** are adapted from ref. ⁵⁰ under a Creative Commons licence CC BY 4.0.

associations between traits ($n=116$; for example, ‘Are kind people often intelligent?’) and impressions towards unfamiliar faces (for a collection of social perceptual similarity matrices in studies 1 and 2, sample size reported is total participant raters, where subsets of this total reported sample rated each trait; $n=484$; for example, ‘How kind/intelligent is this face?’), familiar famous and historical people ($n=503$; for example, ‘How kind/intelligent is Barack Obama?’) and social groups ($n=488$; for example, ‘How kind/intelligent are teachers?’). We looked at a trait space of 15 trait terms across domains, made up of three subtraits per each of the Big Five factors of personality (Methods; Fig. 2¹⁹). From these data, we computed a similarity matrix for each of these four domain models (Fig. 2a). Each matrix is a collection of all pairwise ‘similarities’ in each domain, where similarity in the conceptual trait space matrix is the conceptual association between each trait pair (‘How likely are kind people to be intelligent?’; 1–7 Likert-type item), and in each perceptual trait space matrix is the pairwise Pearson correlation between each trait inference (for example, correlation of ‘kind’ and ‘intelligent’ face impressions). We then applied representational similarity analysis

(RSA), testing the Spearman correlation between unique values in each matrix pair (Spearman correlation used so as to not assume a strictly linear relationship between distances in any two spaces²⁰). In effect, this tests whether the relative degree of correlation between trait pairs across different trait spaces is the same. To account for structured dependency of matrices and allow inference towards our participants (given studies 1, 2 and 7 collapse across subjects), additional approaches for all such RSA analyses are reported in the Supplementary Results.

All pairwise trait space matrix analyses are depicted in Fig. 2b–d. We tested our hypothesis that a common social perceptual trait space may reflect a more domain-general conceptual trait space. Consistent with this hypothesis, we observed significant similarity between the conceptual trait space matrix and all social perceptual trait space matrices (Fig. 2b–d; conceptual matrix predicts: face trait space matrix, Spearman $\rho(103)=0.796$, $\rho^2(103)=0.634$, $P<0.0001$, 95% confidence interval (CI)=[0.713, 0.857]; familiar person trait space matrix, Spearman $\rho(103)=0.739$, $\rho^2(103)=0.545$, $P<0.0001$, 95% CI=[0.637, 0.815]; social group trait space matrix, Spearman $\rho(103)=0.779$, $\rho^2(103)=0.606$, $P<0.0001$, 95% CI=[0.690, 0.844]). Additional analyses confirmed a strong correspondence among the three social perceptual trait matrices (Supplementary Results). RSA using a valence similarity matrix (Supplementary Fig. 1) demonstrated that all reported effects occur above and beyond any effects due to valence alone. These findings provide evidence in support of our theoretical hypothesis that a domain-general conceptual trait space is reflected in a common social perceptual trait space seen across several domains.

Study 2. An important aspect of our theoretical perspective is that trait conceptual knowledge drives inferences regarding social cognition, which are used for understanding other people and predicting their behaviours. We may predict a ‘kind’ person who behaves affectionately to be ‘extroverted’ and socialize frequently. One alternative interpretation of the results is that the correlation of any two trait inferences, such as ‘kind’ and ‘extroverted’, is due merely to how participants find any two words synonymous in semantic meaning (for reviews, see refs. ^{10,21}). To highlight the role of the trait concepts measured here as meaningful concepts that reflect perceivers’ differential predictions about human behaviour¹⁵, eliminate concerns regarding semantics and provide a conservative conceptual replication of study 1, in study 2 we designed a set of tasks emphasizing traits as distinct concepts used to predict distinct behaviours in a substantive manner¹⁵. Rather than asking participants about the same trait terms across domains (for example, ‘Is this face kind?’ and ‘Is this group kind?’), we used different items for each domain, which asked about the behavioural tendencies thought to underlie personality traits (for example, instead of ‘kind’: ‘Is this face likely to agree with others?’ and ‘Is this social group likely to compliment others?’). We gathered several items to correspond uniquely to each trait. Thus, we used behavioural tendency descriptions as proxies for traits for each of the different social perceptual domains, and compared the similarity matrices to the conceptual similarity matrix collected in study 1 that used direct trait terms. Items were chosen from the NEOPI (Neuroticism–Extraversion–Openness Personality Inventory)¹⁹, given both its use of behavioural tendency descriptions to collect information about people’s personalities, and prior validation of these items and their relation to actual personality traits.

We collected new matrices of face ($n=496$), familiar person knowledge ($n=478$) and social group trait space ($n=489$) using distinct trait descriptions between each task (Methods). The data of study 1 were used for the conceptual similarity matrix. Consistent with our hypothesis, we again observed a significant correlation between the conceptual trait space matrix and the three social

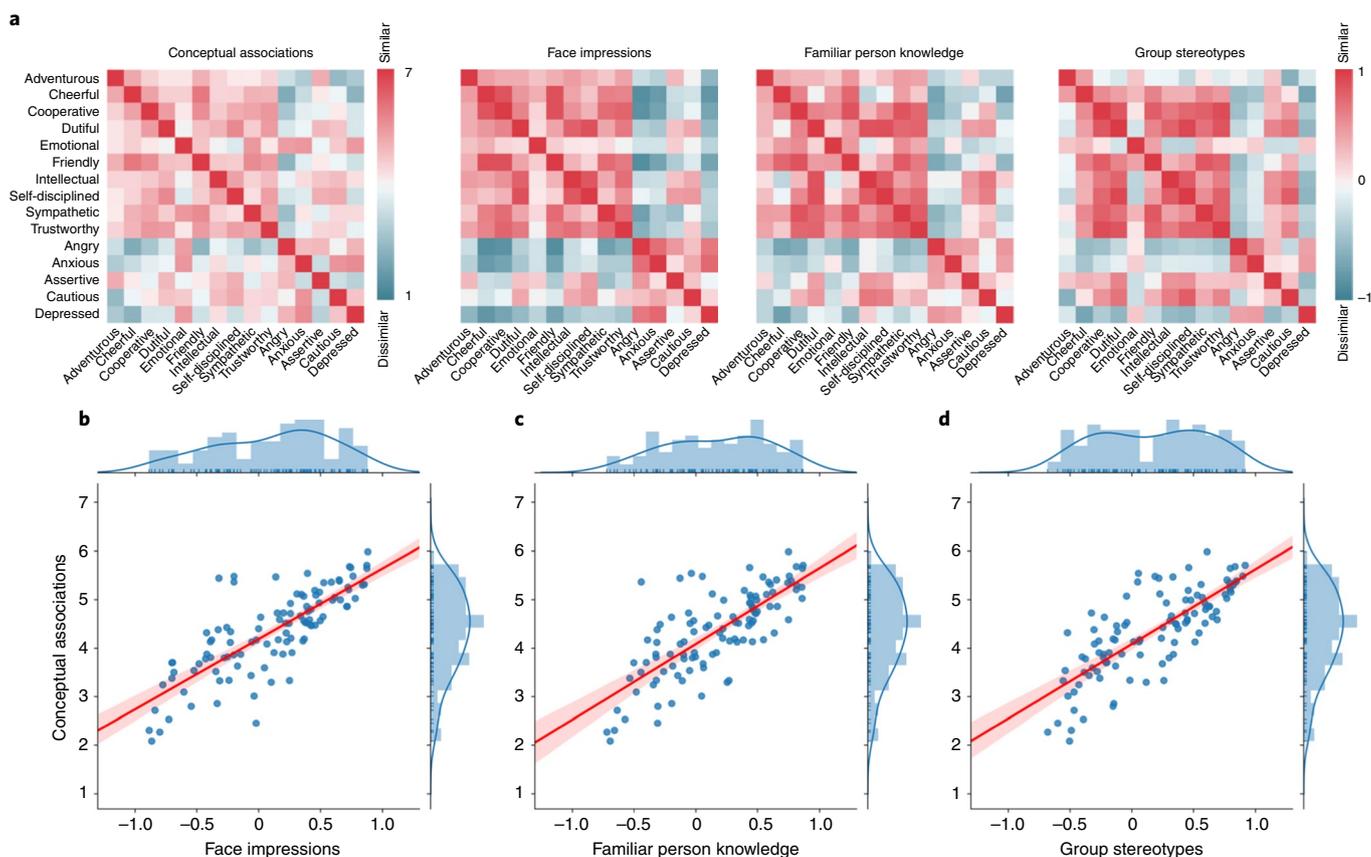


Fig. 2 | Trait inferences across social cognition mirror conceptual knowledge. **a**, All trait space similarity matrices from study 1, each made of the pairwise similarity values between each trait pair (plotted from dissimilar (blue) to similar (red)). Each matrix is sorted by the *k*-means cluster solution of the conceptual trait space matrix, to most intuitively depict their similar structure. Each matrix was collected from a distinct task, set of stimuli and set of participants. **b–d**, The results show that conceptual trait space ($n = 116$) is largely reflected in social perceptual trait spaces across domains (face impressions (**b**), $n = 484$, Spearman $\rho(103) = 0.796$, $\rho^2(103) = 0.634$, $P < 0.0001$, 95% CI = [0.713, 0.857]; person knowledge (**c**), $n = 503$, Spearman $\rho(103) = 0.739$, $\rho^2(103) = 0.545$, $P < 0.0001$, 95% CI = [0.637, 0.815]; group stereotypes (**d**), $n = 488$, Spearman $\rho(103) = 0.779$, $\rho^2(103) = 0.606$, $P < 0.0001$, 95% CI = [0.690, 0.844]). Error ribbons show the standard error of the estimate, and there are 105 trait pairs as data points per panel. Distributions of data points along each variable are plotted along each variable's axis as both a histogram and density plot. While Pearson correlations of original similarity data points are plotted for ease of interpretation, statistical analyses were of rank-ordered data points. In each plot (**b–d**), trait space matrices (**a**) are flattened into their unique pairwise similarity values and plotted against one another (conceptual on the y axis, left-most matrix; social perceptual matrices along the x axes, right three matrices). Each data point is a trait pair (for example, 'friendly'–'self-disciplined'). In each comparison, as two traits become more associated in conceptual knowledge (y axis), they become more correlated in trait inferences across domains (x axes). This pattern is found in study 1, in which trait terms were used in each task (for example, 'friendly', 'self-disciplined'), and in study 2, in which different trait descriptors were used in each task (for example, 'friendly' in the conceptual task, 'likely' to agree with others' in the face task; see study 2 results and Supplementary Fig. 2).

perceptual trait space matrices, despite the use of unique items to construct the different matrices (Supplementary Fig. 2; conceptual matrix predicts: face trait space matrix, Spearman $\rho(103) = 0.575$, $\rho^2(103) = 0.331$, $P < 0.0001$, 95% CI = [0.431, 0.691]; familiar person trait space matrix, Spearman $\rho(103) = 0.576$, $\rho^2(103) = 0.332$, $P < 0.0001$, 95% CI = [0.432, 0.691]; social group trait space matrix, Spearman $\rho(103) = 0.574$, $\rho^2(103) = 0.329$, $P < 0.0001$, 95% CI = [0.430, 0.690]). Similar results were obtained when controlling for valence (Supplementary Results). These findings again suggest that trait conceptual associations and inferences are correlated across domains of social cognition in a similar fashion, suggesting that domain-general conceptual associations may be applied across each domain, resulting in a common trait space. The results also suggest that the commonality in trait space is due to beliefs about personality traits as concepts used to predict meaningful social behaviour, rather than a mere artefact of semantic relatedness among trait terms.

Study 3. While studies 1 and 2 provide initial evidence that the structure of trait inferences reflect that of conceptual associations, these high-level assessments only ask whether, on average across perceivers, traits more correlated in conceptual associations are more correlated in trait inferences. An important component of our theoretical perspective is that conceptual trait associations may shift initial trait inferences, which would entail that variance between perceivers' conceptual knowledge should uniquely shape their idiosyncratic trait inferences across domains. People who believe two traits are more or less correlated (for example, 'kind people are/not intelligent') should make more or less tethered inferences of those two traits (for example, 'kind faces and groups are/not intelligent').

In study 3 ($n = 162$), we extended the methodology of studies 1 and 2. Focusing on face impressions, we collected both conceptual and face trait space matrices per subject along eight traits ('adventurous', 'assertive', 'cautious', 'depressed', 'emotional', 'friendly', 'self-disciplined', 'trustworthy'). We performed RSA within a linear

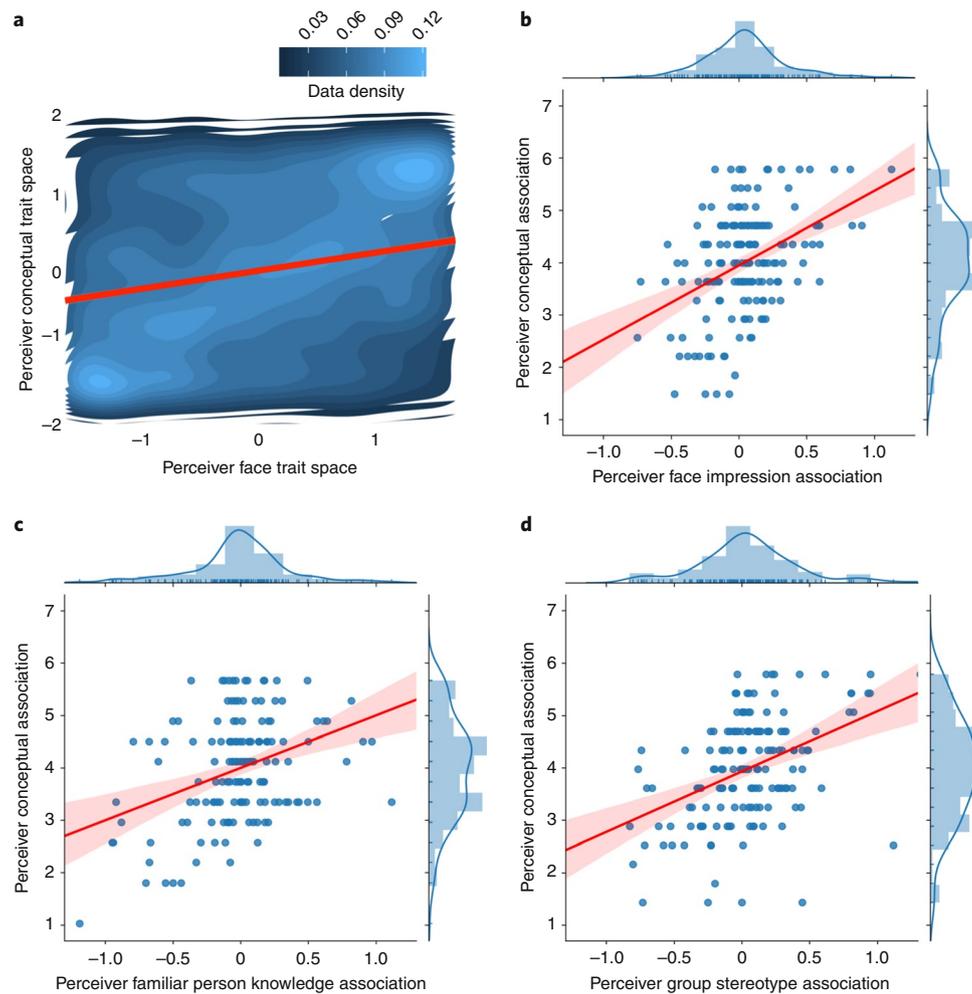


Fig. 3 | Individual differences in conceptual knowledge predict social perception. **a**, Study 3 tested whether the subjective conceptual trait space of a perceiver uniquely predicts their face trait space. A linear mixed-effects model was fit to effectively perform RSA clustered per subject (Methods), and participant subjective conceptual trait space matrices (y axis) predicted their face trait space matrices (x axis), over and above the group-average conceptual trait space matrix (to isolate the effect of subjective associations; estimate of this fixed effect is plotted; $b = 0.145$, $s.e. = 0.020$, $t(141.5) = 7.432$, $P < 0.0001$, $95\% \text{ CI} = [0.11, 0.19]$). Each data point is a trait pair (28 pairs), unique to each subject ($n = 162$; total of 4,563 data points; a contour plot is provided due to the quantity of data, where the colour lightness of the density function represents the probability of each value given the range of values, and the red line is predicted values estimated through multiple regression). **b–d**, In study 4, we see that, across domains, perceivers who believe two traits are more correlated (y axis) also see those traits more similarly in targets (x axes; face impressions (**b**), $n = 167$, Spearman $\rho(165) = 0.331$, $\rho^2(165) = 0.110$, $P < 0.0001$, $95\% \text{ CI} = [0.189, 0.460]$; familiar person knowledge (**c**), $n = 155$, Spearman $\rho(153) = 0.308$, $\rho^2(153) = 0.095$, $P < 0.0001$, $95\% \text{ CI} = [0.158, 0.444]$; group stereotypes (**d**), $n = 162$, Spearman $\rho(160) = 0.435$, $\rho^2(160) = 0.189$, $P < 0.0001$, $95\% \text{ CI} = [0.301, 0.552]$). Error ribbons display standard error of the estimate and data points are each participant per study. While Pearson correlations are plotted for ease of interpretation, statistical analyses were of rank-ordered data points. These results suggest that subjects' idiosyncratic conceptual knowledge and trait inferences are related.

mixed-effects model (Methods), predicting participants' face trait space matrices via their subjective conceptual trait space matrices. Importantly, we allowed for random effects of subject and controlled for the group-average conceptual trait space matrix, therefore testing whether there is a unique contribution of subjective conceptual knowledge to face impressions. Subjective conceptual trait space significantly predicted subjective face trait space over and above group-average conceptual trait space (Fig. 3a; standardized beta (b) = 0.145, $s.e. = 0.020$, $t(141.5) = 7.432$, $P < 0.0001$, $95\% \text{ CI} = [0.11, 0.19]$; similar results were obtained when controlling for valence; Supplementary Results).

Study 4. To further explore the role of perceiver's idiosyncratic conceptual structure and individual differences, as well as survey each domain of trait inferences, in study 4 we tested whether individual

differences in conceptual trait associations correspond to individual differences in specific trait inferences in each domain: face impressions ($n = 167$), familiar person knowledge ($n = 155$) and social groups stereotypes ($n = 162$). In this task, each participant first rated target stimuli along a pair of two randomly assigned traits, then reported their conceptual association between the assigned trait pair. We then tested whether individual differences in conceptual associations correlated with individual differences in trait inference associations. In support of our account, we found a consistent relationship between perceiver conceptual and trait inference associations. The more perceivers conceptually associated trait pairs the more they saw those traits similarly in targets (Fig. 3b–d; conceptual associations correlate with: face impressions, Spearman $\rho(165) = 0.331$, $\rho^2(165) = 0.110$, $P < 0.0001$, $95\% \text{ CI} = [0.189, 0.460]$; familiar person knowledge, Spearman $\rho(153) = 0.308$,

$\rho^2(153)=0.095$, $P<0.0001$, 95% CI=[0.158, 0.444]; social group stereotypes, Spearman $\rho(160)=0.435$, $\rho^2(160)=0.189$, $P<0.0001$, 95% CI=[0.301, 0.552]). In addition to study 3, these results demonstrate that perceivers' subjective trait inferences reflect their unique conceptual associations. Importantly, these results suggest a common trait space is observed within perceivers in line with their own subjective conceptual knowledge, and a common yet divergent structure of trait space between perceivers may emerge to the extent perceivers share or diverge in their conceptual trait knowledge.

Study 5. A key premise of our perspective is that conceptual associations between traits are used in the trait inference process, shaping their initial formation and consequently their intercorrelations, from which emerges a conceptually bound trait space across social perception. So far, while we have found evidence for the relationship between conceptual associations and trait inferences, this evidence has been correlational in nature. In study 5, we manipulated perceiver conceptual associations to more directly examine their directional influence on trait inferences. Participants ($n=141$) were randomly assigned to one of two between-subjects conditions, in which they were led to believe two traits were either positively or negatively correlated. At the beginning of the study, participants were randomly allocated a trait pair from the six unique pairings of 'friendly', 'depressed', 'intellectual'. To manipulate the direction of conceptual associations, participants read a faux science article about personality, which described research finding that the two traits assigned to that participant were either positively or negatively correlated (Methods). Participants then completed a face rating task.

Our analysis tested whether the associations of participants' face impressions were affected by their assigned conceptual association. For instance, we predicted that participants led to believe 'friendly' people are more versus less often 'depressed' would rate friendly appearing faces as more versus less 'depressed'. As repeated face ratings were nested within participant, we examined our hypothesis in a multilevel model. We regressed average ratings of the faces along one trait dimension (for example, 'friendly'; average ratings taken from independent raters) on our participants' ratings of the faces along the other trait dimension (for example, 'intellectual'), their assigned association condition and interaction of these two variables (Methods and Supplementary Methods for details). Participant ratings were group centred. In the model, intercepts were random and slopes were fixed. Consistent with our hypothesis, the strength of association between the two trait dimensions varied by association condition (Fig. 4; unstandardized beta (B)=0.023, s.e.=0.003, $P<0.0001$, 95% CI=[0.017, 0.030]). Simple slopes revealed a more negative trait-pair association in the negative condition ($B=-0.084$, s.e.=0.005, $P<0.0001$, 95% CI=[-0.094, -0.074]) than the positive condition ($B=-0.037$, s.e.=0.003, $P<0.0001$, 95% CI=[-0.047, -0.029]). Note that in both conditions, the trait-pair association should reflect not only the effect of the manipulation but also priors or a 'baseline' level of association between the particular traits. Thus, although the regression coefficient is negative in both conditions, what is critical is whether its magnitude differs across conditions. These results provide causal evidence in support of our theoretical account that conceptual trait associations structure initial inferences and their correlations, from which trait space emerges.

Study 6. While conceptual associations may guide trait inferences, it is also certainly plausible that this influence is bidirectional: the origins of conceptual trait associations may derive from inferences about the social world. Study 6 provided an initial test of this possibility. Rather than manipulate a conceptual trait association and measure its effect on the correlation of face judgements (as in study 5), here we test whether the reverse also holds true. In study 6, participants learned that two traits ('friendly' and 'cautious') were

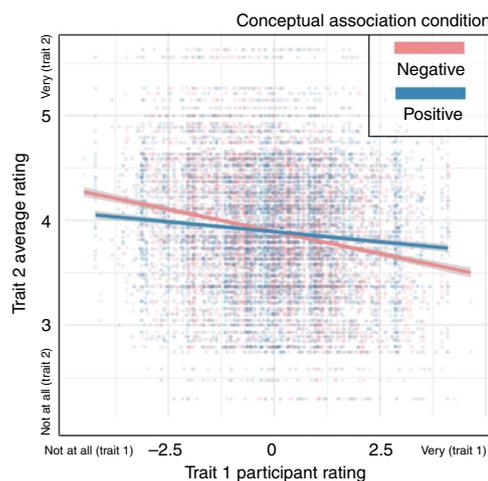


Fig. 4 | Conceptual associations directly shape trait inferences and space. In study 5, we manipulated perceiver conceptual associations with faux science articles and found that perceivers in the negative conceptual association condition see traits less similarly in targets compared with participants in the positive association condition ($n=141$, 90 faces rated, with 12,690 data points plotted; $B=0.023$, s.e.=0.003, $P<0.0001$, 95% CI=[0.017, 0.030]). Red versus blue data points and lines distinguish the negative versus positive association conditions. Error ribbons display standard error of the estimates. Participant ratings of faces (each data point is a face) along one trait (x axis; 'trait 1') correlated with average ratings of the face (from study 1) along another trait (y axis; 'trait 2') more negatively in the negative association condition ($B=-0.084$, s.e.=0.005, $P<0.0001$, 95% CI=[-0.094, -0.074]) than the positive association condition ($B=-0.037$, s.e.=0.003, $P<0.0001$, 95% CI=[-0.047, -0.029]). (Note that in both conditions, priors likely set baseline associations between traits and tether the manipulation down. Therefore, the magnitude and direction of a slope is not relevant to the effects. What is critical is whether its magnitude differs between conditions.) These findings demonstrate that conceptual associations may directly influence the initial trait inference process, from which the structure of trait space may emerge.

correlated positively or negatively depending on their between-subjects condition. Participants were tasked with assessing the 'cautiousness' of different individuals via their faces. After they rated each face, they were given feedback after the trial as to whether they were correct or incorrect regarding the face's 'cautiousness'. These faces varied in how friendly they appeared. Friendliness judgements are highly consistent across perceivers²², and we used faces rated low versus high in perceived friendliness from study 1 for the present study. In the 'positive association' condition, faces above average in how 'friendly' they looked were labelled as 'more cautious' and faces below average in how 'friendly' they looked were labelled as 'less cautious'. In the 'negative association' condition, this pattern was reversed. Therefore, as participants judged the 'cautiousness' of different faces, they received feedback that reinforced either the positive or negative association of targets' 'cautiousness' with the targets' 'friendliness'. Afterwards, participants reported their conceptual associations between the two traits as in previous studies.

As predicted, we found that perceivers manipulated to believe friendliness and cautiousness were positively associated conceptually associated the two traits to a stronger degree (mean, $M=4.508$, s.d.=1.593) than perceivers led to believe the traits were negatively associated ($M=4.019$, s.d.=1.416; mean difference=0.4888, independent t -test, $t(144)=2.127$, $P=0.035$, $r^2=0.03$, mean difference 95% CI=[0.035, 0.943]; Fig. 5a). These findings support the

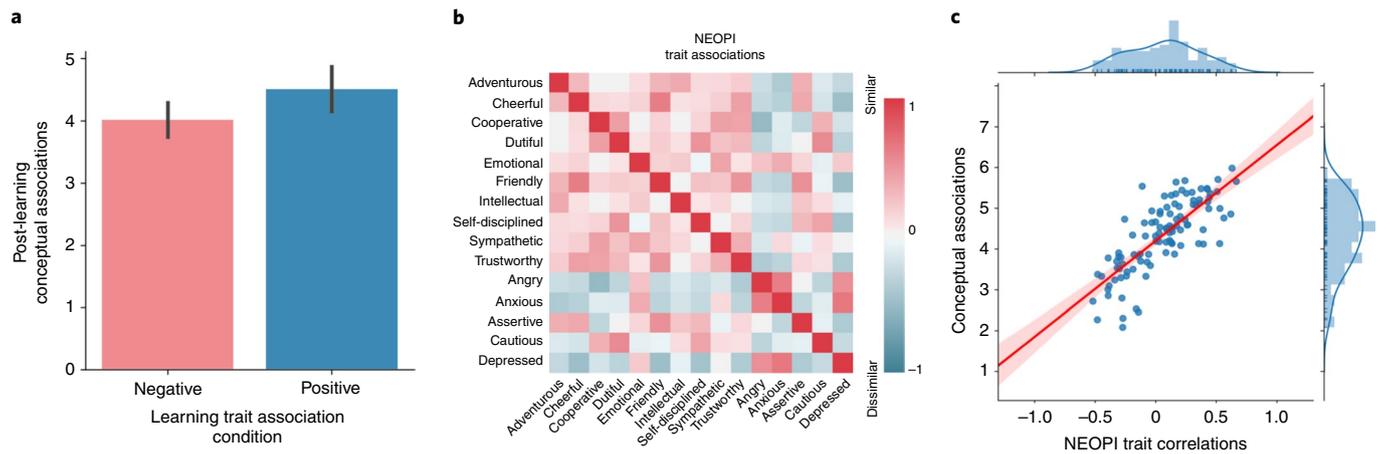


Fig. 5 | Social perception and trait inferences influence conceptual trait space. In studies 6 and 7, we tested whether the relationship between conceptual knowledge and trait inferences is bidirectional. **a**, In study 6 ($n = 146$), we found that conceptual associations between two traits ('cautious', 'friendly'; y axis) were stronger for participants assigned to observe those two traits positively correlating in target faces ($M = 4.508$, $s.d. = 1.593$), compared with participants assigned to perceive their negative correlation ($M = 4.019$, $s.d. = 1.416$; mean difference = 0.4888, independent t -test, $t(144) = 2.127$, $P = 0.035$, $r^2 = 0.03$, mean difference 95% CI = [0.035, 0.943]). Mean (bar height) and standard error (error bars) of participants' conceptual associations are plotted (negative association condition in pink, positive in blue). **b, c**, In study 7, we tested whether perceivers' conceptual knowledge is learned to some extent from the actual structure of human personality. **b**, We collected a trait space matrix of actual personality trait correlations of those traits used in previous studies (via the NEOPI, $n = 307,313$; personality trait correlations plotted from negative (blue) to positive (red)). **c**, RSA found that the NEOPI trait space matrix and conceptual trait space matrices explain a sizeable proportion of variance in one another (105 trait pairs as data points, Spearman $\rho(103) = 0.77$, $\rho^2(103) = 0.60$, $P < 0.0001$; 95% CI = [0.684, 0.841]). Error ribbons display standard error of the estimate, and there are 105 trait pairs as data points per panel. While Pearson correlations are plotted for ease of interpretation, statistical analyses were of rank-ordered data points. These findings suggest conceptual trait space is also shaped through social perception and that one potential source is direct observational or indirect social learning of the actual correlation of personality traits in others.

hypothesis that perceivers not only apply conceptual knowledge to social perception but also learn about trait concepts from social perception. However, an important limitation of these experiments is the strong priors that individuals hold through a lifetime of learning that would precede this experiment. More critically, while these findings demonstrate learning of conceptual associations are possible, the results are agnostic to the actual source of the associations that perceivers bring to the table in trait inference (such as those found in studies 1–4).

Study 7. One of the candidate sources of conceptual trait space is learning of the actual structure of others' personality traits. Much like conceptual trait space, people's actual personality traits are highly correlated along a relatively small set of dimensions (for example, the Big Five factors of personality^{23,24}). If actual personality traits were in fact correlated, a simple strategy to optimize trait inference for perceivers would be to learn this structure and make predictions accordingly. Not all of our personality traits are worn on our sleeve²⁵, so perceivers may take trait information at hand to surmise the whole of a target. For example, perceivers may use 'talkativeness', a more visible trait, to infer a target's 'anxiety', a less visible trait, based on their conceptual association between 'talkative' and 'anxious'²⁵. If perceivers learn the actual structure of personality traits they believe are more similar conceptually would also be more similar in actual human personality structure. This, of course, would be only one among many candidate sources for trait knowledge²⁶.

To test this possibility, we compared conceptual trait space (as measured in study 1) to an estimate of actual personality trait space via the NEO personality inventory (henceforth referred to as 'NEOPI trait space'^{19,24}). The NEOPI is a canonical and empirically validated model of personality structure, ideal for the current research as participants whose personality traits are measured do not explicitly evaluate whether they possess the traits of 'trustworthiness' or 'anxiety'. This greatly reduces the potential confound that

our NEOPI trait space matrix could reflect perceivers' social cognitive trait spaces merely due to semantic similarities in measurement (for example, reporting of 'warmth' trait in self and in others). The Big Five factor model of personality is composed of a larger set of personality traits underlying each of the Big Five factors, known as its 'facets', which were the trait adjectives used in our research above taken from the facet subscales of the NEOPI^{19,27}. Therefore, we were able to calculate a NEOPI trait space matrix comparable to the social cognition trait space matrices used in studies 1 and 2, as the same 15 traits are measured in all domains. The NEOPI trait space (Fig. 5b) was calculated via data acquired from a large open source dataset ($n = 307,313$ participants; retrieved from <https://osf.io/tbmh5>)²⁷. This allowed us to effectively test whether trait pairs associated in conceptual knowledge are also associated in a ground-truth model of personality.

Strikingly, perceiver social conceptual knowledge (via study 1) closely tracked the NEOPI trait structure, where trait pairs that perceivers conceptually relate are also more correlated in personality as measured in the NEOPI (Fig. 5c; Spearman $\rho(103) = 0.77$, $\rho^2(103) = 0.60$, $P < 0.0001$; 95% CI = [0.684, 0.841]). Supplementary analyses confirmed that social perceptual trait spaces also strongly resembled NEOPI trait space, as would be expected through transitivity given our hypothesis that social perceptual trait spaces reflect conceptual trait space (Supplementary Fig. 3 and Supplementary Results). These findings show that perceiver trait conceptual knowledge reflects actual personality structure, suggesting the possibility that its structure may be learned through some mechanism, such as cultural transmission or accurate observation.

Discussion

Taken together, our results broadly demonstrate that conceptual trait knowledge shapes trait inferences across distinct domains of social perception, including face impressions, familiar person knowledge and group stereotypes. The similarity structures of social

perceptual trait inferences were all highly correlated with that of conceptual trait space (studies 1 and 2). Participants' idiosyncratic conceptual knowledge was reflected in their inferences and social perceptual trait spaces (studies 3 and 4), and manipulation of perceiver conceptual associations influenced trait inferences accordingly (study 5). To probe the source of trait concept knowledge, we found evidence suggesting that conceptual knowledge may be learned through social perception, demonstrating the bidirectional nature of this process, through direct observation (study 6) or learning about the actual structure of personality traits (study 7).

These findings provide quantitative evidence for a common trait space across social cognition², which may emerge as trait inferences are similarly shaped by learned conceptual trait space across the many domains of social perception. A prominent perspective is that a common trait space arises due to the adaptive utility of its core dimensions—namely that, across social domains, perceivers track those traits significant to our function and survival (intentions and capabilities; for example, 'competence' and 'warmth'^{2,6}). Evidence suggests that this is the case, as traits with adaptive utility play a central role in dimensions of social cognition. Yet there is much additional covariation in the true expanse of trait space that is less easily explained by this functional perspective, such as the perceived relationships between humour, sociability, risk aversion or neuroticism. The findings reported here support a parsimonious explanation by a more proximal mechanism to perceptions, that perceivers' conceptual knowledge about how traits correlate in others shape how correlated they are in social perceptions regardless of their domain^{10,13}. This perspective provides a unifying framework through which we may understand trait space as part of a dynamic cognitive process, from which we may generate broad and general hypotheses about social perception based on context-varying models of social conceptual structure. This perspective also fits trait inferences generally, especially outside of face perception, into an emerging picture of the conceptual nature of social perception^{28,29}.

This flexible account may be indispensable for accommodating emerging findings of dynamic shifts in social cognitive models. Variation in social trait spaces has been increasingly well documented, suggesting trait space may in fact be dynamic rather than fixed in its structure, both shifting in its core dimensions and their relations depending on social factors^{18,30–35}. While trait space generally tends to be consistent across perceivers, various perceiver factors (for example, stereotypes, motivations, emotions) and social context may shape trait space (for a review, see ref. ⁹), as much of the variance in trait inferences is due to perceiver characteristics²². A trait space shaped by perceiver conceptual knowledge could, in theory, underlie these various findings. For instance, competence and warmth inferences come to correlate positively towards liked groups³⁶ and negatively towards disliked groups or groups with specific stereotypes (for example, outgroups and women^{30,37}). Perhaps conceptual associations between personality traits vary across these contexts in systematic ways. Future research should investigate how conceptual associations shift across social contexts, and whether these shifts are reliably reflected in different social trait spaces.

An important question concerns the origins of perceivers' conceptual trait associations, which we argue may lie at the foundation of a common social trait space. Human personality traits are in fact intercorrelated^{23,24}, and this similarity structure is tied to patterns of behaviour³⁸. Thus, it is possible people may come to learn actual personality structure to predict others' behaviour¹⁰. Here we found that conceptual trait space reflects that of actual personality traits. Previous research has found similar associations¹³, and that perceivers can use accurate knowledge of one personality trait to accurately predict other traits of which they are not explicitly informed³⁹. While our findings suggest perceivers learn actual personality structure, this is an area ripe for future research. Trait knowledge may be acquired through direct observation, such as social and statistical

learning of the social environment⁴⁰, or indirect sources, such as cultural learning and gossip⁴¹. Such knowledge is also probably shaped and biased by the host of processes and biases long known to influence trait inferences²⁶. One interesting question is the relative degree that semantic knowledge compared with cognitive biases contributes to trait space structure. Furthermore, our findings do not speak to which domains of social perception provide information about actual personality. We would speculate different domains must contribute differently. For instance, perception of more or less familiar individuals may provide more or less signal towards actual trait correlations^{25,42}, yet perceived trait correlations in faces and stereotypes may suffer from limited signal as these sources are often biased^{43–45} (cf. ref. ⁴⁶) and thus contribute less significantly to such accurate trait correlation learning. This line of research may therefore be of interest to the accuracy literature more broadly⁴⁷, and join other recent findings exploring accuracy through the perspective of trait space models¹⁵. Future research should quantify the contributions of different information sources to the development of conceptual trait knowledge.

It is crucial to note that, although trait space structure may be learned from actual personality structure, this should not imply perceivers' persistent accuracy in trait inferences themselves across domains, especially when perceivers begin with inaccurate and biased inferences (for example, in the case of face impressions⁴⁴). Rather, an accurately learned trait space may just as often lead to broad inaccuracies. Humans often begin with inaccurate and biased trait inferences. When initial inferences are inaccurate, other trait inferences made through what are accurate associations may increase in likely inaccuracy. For instance, if friendliness and sociability personality traits are in fact correlated, and perceivers understand this, an erroneous 'unfriendly' inference of a friendly target would lead to an 'unsociable' inference, although the genuinely friendly target is more likely to be sociable. Thus, an accurate trait space structure is easily misapplied by inaccurate inference content, and the structure of trait space can be an accurate reflection of reality while the content of inferences is far from it. It will be important for future work to identify when trait space may lead to accuracy or error in judgement.

We believe our perspective and findings here suggest a reorienting in the study of trait space is needed. These and other recent findings^{14,17} suggest trait space as a key process in forming initial perceptions⁷, in which trait space as measured in the context of target evaluations is merely an emergent property of this process. Some of the most interesting questions may be how and when conceptual trait space is used, for better or worse, and what unique predictions this framework affords models of social perception. One salient prediction, much like in the case of individuation in stereotyping⁴⁸, is that trait space is most used when other trait information is scarce. This is akin to saying perceivers may have 'trait stereotypes', or make further trait generalizations based on those they initially infer. In recent years, a resurgence in the study of social perceptual dimensions has occurred, with scientific interest in what dimensions best describe trait or mental-state space^{9,15}. Should trait space be dynamic and context dependent, attempts to identify and refine any 'true' universal dimensions may be misguided, as trait space is destined to vary when its conceptual basis and its application does. Future research might benefit from explicating the precise role and shape of trait space in the context at hand, from which we may develop nuanced models that predict the structure of trait inferences in particular domains.

There are several limitations of this work. First, although we manipulate conceptual associations or face trait covariations in studies 5 and 6, more thorough designs should be developed to test the bidirectional and mutually causal relationship between conceptual trait associations and social perception. Another limitation is the use of verbal stimuli to measure trait space, as it may

influence participant responses to conform across these tasks, especially between actual and perceived trait spaces (for a review of related research, see ref. ¹⁰). Lastly, studies 6 and 7 are demonstrations of possible sources of information shaping conceptual associations and should not be taken to present an exhaustive argument for the origins of conceptual trait knowledge.

In short, the present research provides evidence of a common trait space across social cognition, structured by perceivers' learned conceptions of how personality traits correlate. This account of trait space not only explains its homogeneity across social cognition, but also highlights that trait space may be dynamic rather than fixed in its architecture to the extent perceiver conceptual knowledge about personality shifts due to myriad social and contextual factors. We hope this work provides a parsimonious framework to understand trait space, and importantly, allows us to move beyond its measurement to questions of its foundational role in social perception.

Methods

All the studies conducted comply with ethical regulations for research on human subjects and all participants gave informed consent, as approved by the University Committee on Activities Involving Human Subjects at New York University. Participants were financially compensated US\$0.10 per minute for their participation. Statistical tests are two-tailed. Data distributions were assumed to be normal but this was not formally tested. No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those reported in previous publications^{17,29}. Randomization was applied where possible in all studies, and is described explicitly in each study methods section. Data collection and analysis were not performed blind to the conditions of the experiments.

All data, stimuli names, and data preparation and analysis code are available on the OSF (<https://osf.io/2uzsx/>). Analyses were performed in Python and R. Additional details on task instructions and approach to data preparation and analysis are provided in the Supplementary Methods.

Study 1. Participants. We aimed to recruit ample raters (participants) to acquire stable and reliable estimates of the trait ratings per each exemplar. Our target sample was 30 participants per trait rated in each of the social perception rating tasks below (face, familiar person and social group trait tasks), as trait ratings across traits stabilize at approximately this number of participant raters⁴⁹. Across traits and tasks, this totalled a target sample of $n = 450$ per social perception model. For the conceptual trait space model, involving conceptual ratings of traits with other traits, we based target sample size on previous work estimating a similar model²⁹, seeking a target sample of $n = 100$.

Conceptual trait task. We collected conceptual trait association data from 116 participants via Amazon Mechanical Turk (demographic data missing for 1 subject; all US residents; all primary English speakers; $M_{\text{age}} = 35.4$ yr, $s.d._{\text{age}} = 10.5$ yr; 58 female, 55 male, 2 other; 113 White, 2 other).

Face trait task. We collected face impression data from 484 participants via Amazon Mechanical Turk (demographic data missing for 2 subjects; all US residents; all primary English speakers; $M_{\text{age}} = 35.5$ yr, $s.d._{\text{age}} = 12.3$ yr; 281 female, 199 male, 2 other; 372 White, 44 Black, 31 Asian, 37 other). Participants were randomly assigned to evaluate one personality trait in all face stimuli, and were therefore divided roughly equally between all 15 personality trait conditions (32 participants per trait condition on average).

Familiar person trait task. We collected familiar person knowledge data from 503 participants via Amazon Mechanical Turk (demographic data missing for 4 subjects; all US residents; all primary English speakers; $M_{\text{age}} = 30.7$ yr, $s.d._{\text{age}} = 7.1$ yr; 308 female, 175 male, 16 other; 368 White, 44 Black, 42 Asian, 49 other). Participants were randomly assigned to evaluate one personality trait in all familiar person stimuli, and were therefore divided roughly equally between all 15 personality trait conditions (~34 participants per trait condition).

Group trait task. We collected group stereotype content data from 488 participants via Amazon Mechanical Turk (demographic data missing for 3 subjects; all US residents; all primary English speakers; $M_{\text{age}} = 30.4$ yr, $s.d._{\text{age}} = 6.9$ yr; 297 female, 183 male, 5 other; 368 White, 44 Black, 39 Asian, 37 other). Participants were randomly assigned to evaluate one personality trait in all group stimuli, and were therefore divided roughly equally between all 15 personality trait conditions (~33 participants per trait condition).

Valence task. We collected valence ratings of each personality trait adjective used in the above tasks from 69 subjects via Amazon Mechanical Turk ($n = 69$; $M_{\text{age}} = 31.4$ yr, $s.d._{\text{age}} = 6.6$ yr; 28 female, 40 male, 1 other; 52 White, 11 Black, 4 Asian, 2 other).

Stimuli. Personality trait stimuli. We chose personality trait stimuli that corresponded with many of the facet subscales of the NEOPI^{19,27}. We chose 15 facet subscale traits, including three from each of the Big Five personality factors to maintain a balance with the comparison of actual personality trait space. These were subtraits of the primary five factors: 'agreeableness', 'conscientiousness', 'extroversion', 'neuroticism' and 'openness'. The three chosen per primary factor were selected to most easily translate into adjectives participants could engage comfortably in each task. These traits included: 'adventurous', 'angry', 'anxious', 'assertive', 'cautious', 'cheerful', 'cooperative', 'depressed', 'dutiful', 'emotional', 'friendly', 'intellectual', 'self-disciplined', 'sympathetic' and 'trustworthy'.

Face stimuli. All stimuli were taken from the Chicago Face Database⁵⁰. Face stimuli included 90 portrait photographs of young White male individuals with neutral facial expressions. Exact stimulus identification numbers are provided on the OSF page (<https://osf.io/2uzsx/>). Example stimuli are presented in Fig. 1b.

Familiar person stimuli. All familiar person stimuli were chosen from recent work that used data-driven methods to identify individuals highest in familiarity to a similar online sample demographic, and maximize diversity in traits of the stimuli to guarantee a wide and generalizable sampling of trait space⁵¹. We used all 60 familiar person stimuli identified in Thornton and Mitchell⁵¹. Stimuli are provided in the OSF page (<https://osf.io/2uzsx/>). Example stimuli are presented in Fig. 1c.

Group stimuli. To obtain a diverse set of social group stimuli, we chose the 80 most frequently named social groups in the United States, as named in recent work by an online participant demographic similar to our own⁵³. Stimuli are provided in the OSF page (<https://osf.io/2uzsx/>). Example stimuli are presented in Fig. 1d.

Protocol. Conceptual trait task. Participants were informed that they would partake in a study on how different personality traits correlate in the world. After several examples and practice trials, participants began the task. Each trial item asked, "Given that an individual possesses one trait, how likely is it that they possess the other?" then presented the two trait stimuli for that trial separated by a hyphen (for example, 'friendly-self-disciplined'). Participants evaluated the conceptual relationship of each trait pair in the 15 trait stimuli (1–7 Likert-type scale, 1—'Not at all likely' to 7—'Very likely'), presented in both orders given the wording of the item question (for example, 'friendly-self-disciplined' and 'self-disciplined-friendly'). Therefore, there were a total of 210 trials for each participant (total possible permutations from all pairs of 15 trait stimuli). Participants then completed a general demographics survey.

Face trait task. Participants were informed they would partake in a study examining how people perceive others. Each participant was randomly assigned to evaluate only one of the 15 personality trait stimuli in faces. In the task, participants rated each of the 90 face stimuli on the personality trait they were assigned (1–7 Likert-type scale; for example, 1—'Not at all friendly', 4—'neutral', 7—'Very friendly'). Following the face trait rating task, participants completed a general demographics survey.

Familiar person trait task. Participants were informed they would partake in a study examining how people perceive famous people. Each participant was randomly assigned to evaluate only one of the 15 personality trait stimuli in familiar person stimuli. In the task, participants rated each of the 60 familiar person stimuli on the personality trait they were assigned (1–7 Likert-type scale; for example, 1—'Not at all friendly', 4—'neutral', 7—'Very friendly'). Following the familiar person trait rating task, participants completed a general demographics survey.

Group trait task. Participants were informed they would partake in a study examining common societal inferences, rather than their own, towards common social groups. Instructions were intended to reduce the influence of social desirability on responses^{5,33}. Each participant was randomly assigned to evaluate only one of the 15 personality trait stimuli in the social group stimuli. In the task, participants rated each of the 80 social group stimuli on the personality trait they were assigned (1–7 Likert-type scale; for example, 1—'Not at all friendly' to 7—'Very friendly'). Following the social group trait rating task, participants completed a general demographics survey.

Valence task. Participants were instructed to rate personality traits on their valence, or how negative to positive each trait is. Participants then rated one stimulus at a time (1–7 Likert-type scale; for example, 1—'Very negative' to 7—'Very positive'). Trials were randomized per subject. Participants responded to a basic demographics survey after the task.

Study 2. Participants. Face trait task. We collected face impression data from 496 participants via Amazon Mechanical Turk (all US residents; all primary English speakers; $M_{\text{age}} = 30.3$ yr, $s.d._{\text{age}} = 6.3$ yr; 257 female, 237 male, 2 other; 320 White, 101 Asian, 30 Black, 45 other). Participants were randomly assigned to evaluate one personality trait in all face stimuli, and were therefore divided roughly equally between all 15 personality trait conditions (~33 participants per trait condition on average).

Familiar person trait task. We collected familiar person knowledge data from 478 participants via Amazon Mechanical Turk (demographic data missing for 2 subjects; all US residents; all primary English speakers; $M_{\text{age}} = 29.8$ yr, $s.d._{\text{age}} = 6.3$ yr; 239 female, 237 male, 2 other; 309 White, 89 Asian, 40 Black, 40 other). Participants were randomly assigned to evaluate one personality trait in all familiar person stimuli, and were therefore divided roughly equally between all 15 personality trait conditions (~32 participants per trait condition).

Group trait task. We collected group stereotype content data from 489 participants via Amazon Mechanical Turk (all US residents; all primary English speakers; $M_{\text{age}} = 30.4$ yr, $s.d._{\text{age}} = 6.7$ yr; 263 female, 223 male, 3 other; 315 White, 89 Asian, 43 Black, 42 other). Participants were randomly assigned to evaluate one personality trait in all group stimuli, and were therefore divided roughly equally between all 15 personality trait conditions (~33 participants per trait condition).

Stimuli. Personality trait stimuli. We replaced each of the original 15 trait items from study 1 with items that asked about the likely behaviour of the target. A different item for each trait was used in each domain (that is, for a given trait, item A would only be used in the face task, item B in the familiar person task and item C in the social group task). We chose the new trait stimuli to replace trait terms with from the NEOPI facet items. In the NEOPI, to measure personality, participants are not asked directly whether they are 'kind', but asked multiple items that describe behavioural tendencies that have been found to relate to kindness. Given the long history of validation of these items and their correspondence to behaviours that relate to underlying personality traits, we chose our new trait stimuli from these items. Specifically, for each of the 15 traits in our similarity matrices, we chose three NEOPI items, so that one unique item per trait could be used in each of the three social perception tasks (face, familiar person, social group). Specific items may be viewed via the OSF (<https://osf.io/2uzsx/>).

Target stimuli. For face, familiar person and social group target stimuli, study 2 used the same stimuli as study 1.

Protocol. Study 2 used an identical task design for each of the three tasks as study 1, where only the items were replaced (see above).

Study 3. Participants. All data and tasks were performed by and within each participant. We collected face impression and conceptual association data from 162 participants via Amazon Mechanical Turk (original $n = 168$; 6 subjects dropped due to failure to follow task instructions; all US residents; all primary English speakers; $M_{\text{age}} = 31.9$ yr, $s.d._{\text{age}} = 5.9$ yr; 54 female, 108 male; 113 White, 33 Black, 9 Asian, 7 other).

Stimuli. As each subject completed multiple face rating tasks and a conceptual association task (compared with studies 1 and 2, in which participants completed a task for only one trait), a subset of trait adjective and face stimuli were used in study 3 in consideration of time constraints and participant fatigue.

Personality trait stimuli. We chose a subset of trait stimuli from those used in studies 1 and 2. Trait stimuli included: 'adventurous', 'assertive', 'cautious', 'depressed', 'emotional', 'friendly', 'self-disciplined' and 'trustworthy'. We used each trait in its own single block face rating task, and each pairwise combination of these traits in the conceptual association task.

Face stimuli. For face stimuli, each participant was assigned to a random subset of 25 stimuli from the face stimulus set used in studies 1 and 2.

Protocol. Participants first completed one block of face ratings for each personality trait stimulus, with the trait block order randomized (one trait within each block, for a total of eight blocks). The same 25 face stimuli were rated within each block. Following the series of face rating tasks, participants completed a conceptual association task in which they rated their pairwise conceptual association of each trait pair presented in random order. Each task otherwise had an identical design to that of study 1. Following completion of the two task sets, participants completed a standard demographics survey.

Study 4. Participants. Face trait task. We collected face impression data from 167 participants via Amazon Mechanical Turk (original $n = 174$; 5 subjects dropped due to task incompleteness; 2 subjects dropped due to failure to follow task instructions; all US residents; all primary English speakers; $M_{\text{age}} = 31.44$ yr, $s.d._{\text{age}} = 5.50$ yr; 102 female, 64 male, 1 decline; 128 White, 23 Black, 5 Asian, 11 other).

Familiar person trait task. We collected familiar person knowledge data from 155 participants via Amazon Mechanical Turk (original $n = 167$; 9 subjects dropped due to task incompleteness; 3 subjects dropped due to failure to follow task instructions; all US residents; all primary English speakers; $M_{\text{age}} = 32.34$ yr, $s.d._{\text{age}} = 6.52$ yr; 70 female, 82 male, 3 decline; 120 White, 20 Black, 6 Asian, 19 other).

Social group trait task. We collected group stereotype content data from 162 participants via Amazon Mechanical Turk (original $n = 168$; 6 subjects dropped due to task incompleteness; all US residents; all primary English speakers; $M_{\text{age}} = 31.45$ yr, $s.d._{\text{age}} = 5.53$ yr; 72 female, 90 male; 126 White, 20 Black, 8 Asian, 8 other).

Stimuli. Personality trait stimuli. We chose a diverse set of trait stimuli somewhat deviating from those in study 1 to assess generalizability. Trait stimuli included: 'creative', 'dishonest', 'friendly', 'intelligent', 'sociable' and 'stubborn'. We used all pairwise combinations of these trait pairs (for a total of 15 unique possible trait pairs). Participants were randomly assigned to one of the 15 total trait-pair combinations.

Target stimuli. For face, familiar person and social group target stimuli, study 4 used the same stimuli as study 1.

Protocol. Both social perception trait and conceptual trait tasks were largely identical in design within themselves to those in previous studies (see study 1 methods). A major distinction is that in this study, each participant both provided target trait and conceptual trait data. Each participant was randomly assigned to one of 15 trait pairs (the unique combinations of 6 trait stimuli: 'creative', 'dishonest', 'friendly', 'intelligent', 'sociable' and 'stubborn'). First, participants evaluated all stimuli on both assigned traits (either face, familiar person or group stimuli depending on the task). They evaluated all stimuli on one trait first, followed by the other, the order of which trait came first was randomized. The order of which trait was first evaluated was randomly determined per subject. In total, participants therefore completed: 180 trials of face impressions, 120 trials of familiar person impressions, 160 trials of group inferences. From this data, we were able to measure the correlation of inferences within each subject. Second, participants provided conceptual trait association ratings for their assigned trait pair. As participants only evaluated the similarity of two traits to one another (compared with the many trait pairs in study 1), there were only two trials in the conceptual trait task, randomly ordered. Following these tasks, participants completed a general demographics survey.

Study 5. Participants. We collected face impression data from 141 participants via Amazon Mechanical Turk (original $n = 192$; 51 subjects dropped due to insensitivity to experiment manipulation; all US residents; all primary English speakers; $M_{\text{age}} = 32.69$ yr, $s.d._{\text{age}} = 6.44$ yr; 64 female, 77 male; 102 White, 18 Black, 9 Asian, 12 other).

Stimuli. Personality trait stimuli. We chose three trait terms from the facets of the Big Five factors of personality, corresponding to the 'agreeableness', 'neuroticism' and 'openness' factors: 'friendly', 'depressed' and 'intellectual' (for three combinations of trait pairs). These traits were chosen given their correspondence to both relatively independent trait concepts, and to traditional dimensions of social perception ('friendly' to 'warmth', 'intellectual' to 'competence')³ and core personality traits ('friendly' to 'agreeable', 'intellectual' to 'openness', and 'depressed' to 'neuroticism')¹⁹. Further, we chose traits whose conceptual associations could be realistically manipulated in perceivers (for example, given likely strong priors for associations of traits that load along the same factors, such as 'warmth' and 'sociability').

Trait association manipulation article. To manipulate participant conceptual associations between traits, participants read a fake scientific article about the actual correlation of personality traits. Participants read an adapted article from previous research that was successful in manipulating lay theories of gender³². The article explained a research study conducted by personality researchers, who find that on average individuals with one personality trait (for example, friendliness) are more or less likely to have another personality trait (for example, depression). Each participant was randomly assigned to one trait pair at the beginning of the study, which was inserted into the article. The manipulation articles are available on the OSF (<https://osf.io/2uzsx/>).

Face trait stimuli. Study 5 used the same face stimuli as study 1.

Manipulation check. After the experiment, participants completed a brief questionnaire to assess effectiveness of the manipulation. Modelled from previous research and our own measurement methods^{17,52}, participants were asked direct questions about their conceptual associations between their assigned trait pair to assess manipulation effectiveness. We asked participants how likely individuals with the first assigned personality trait are likely to have the second trait assigned to the participant, and vice versa (for example, 'How likely is a friendly person to be intellectual?', 'How likely is an intellectual person to be friendly?'; Likert-type scale, 1—Not at all likely, 2, 3, 4—Neutral, 5, 6, 7—Very likely).

Protocol. At the beginning of the study, participants were randomly assigned to one of two 'association direction' conditions, specifying whether the between-subjects manipulation would convince them their trait pair was negatively or positively correlated (for example, are 'friendly' people more likely to be 'depressed' or less

likely to be 'depressed'). Participants were randomly allocated one of the three trait pairs. Participants were informed that they would take part in a study on how people think of others. In the first part, we manipulated what they thought about personality by having them read an article about personality. Once participants began to read the article, we did not allow them to progress past the article for 2 min to further encourage their reading and engagement of the article given its length. The article explained research finding the participant's assigned trait pair (for example, 'depressed'–'friendly') was negatively or positively correlated, depending on the participant's association direction condition. After reading the article, participants were given a free response form to summarize the article and additionally provide their thoughts on the article and personality generally. Next, we informed participants that a new task would begin where they would make personality judgements of others based on only their face. This task and its instructions were identical to that of the face rating tasks in previous studies. Participants rated all 90 face stimuli on one trait to which they were assigned (trait randomly chosen; Likert-type scale, for example, 1—Not at all friendly, 4—Neutral, 7—Very friendly). Lastly, participants completed the manipulation check, reporting their conceptual association for their assigned trait pair. Instructions and item design were identical to those used in study 3.

Study 6. Participants. We collected data from 146 participants via Amazon Mechanical Turk ($n = 146$; 6 subjects dropped due to task incompletion; all US residents; all primary English speakers; $M_{\text{age}} = 37.2$ yr, $s.d._{\text{age}} = 12.7$ yr; 69 female, 77 male; 115 White, 15 Black, 9 Asian, 7 other).

Stimuli. Personality trait stimuli. Participants guessed the ostensible 'cautiousness' of face stimuli. We next measured their conceptual association between 'friendly' and 'cautious'. 'Friendliness' was chosen due to the spontaneity of 'friendly' face impressions^{33,34}, and we chose 'cautiousness' because of its relatively low conceptual association with 'friendliness' (where in study 1 it had the lowest relationship, closest to the 'neutral' response option, with a value of 4.13 on the 1–7 Likert-type scale).

Face stimuli. Face stimuli were a 56-face subset of those face stimuli used in previous experiments. The stimuli were split into two sets based on their above or below 'Neutral' 'friendly' ratings from study 1, allowing us to label responses to the more or less 'friendly' faces as 'more cautious' or 'less cautious' depending on the subjects' experimental conditions.

Protocol. The task was a two-part task, consisting of the learning phase, in which participants were manipulated to either positively or negatively associate 'friendliness' with 'cautiousness', followed by the conceptual trait association task, in which they reported their conceptual association between the personality trait stimuli. In the learning phase, participants guessed the ostensible 'cautiousness' of the faces, making a two-choice categorization as to whether each face was 'less cautious' or 'more cautious'. Based on their experimental condition, feedback to 'cautiousness' judgements of low versus high 'friendliness' faces indicated an incorrect or correct response. These data were not analysed as the task purpose was manipulation of an association, through feedback in which friendly or unfriendly faces were said to be more or less cautious. Following the feedback phase, they completed a conceptual trait association task identical to that in previous studies (but here it only included the 'friendly' and 'cautious' trait-pair ratings).

Study 7. Data utilized in study 7 analyses came from the study 1 conceptual trait task data, and a personality measurement dataset available from previously published research via the OSF²⁷. For reporting of methods and data selection from this outside dataset, see the Supplementary Methods.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Experiment materials information and all experiment de-identified data are publicly available at <https://osf.io/2uzsx/>. The materials used in this study are widely available.

Code availability

Data analysis script notebooks are publicly available at <https://osf.io/2uzsx/>.

Received: 2 May 2019; Accepted: 27 November 2019;
Published online: 13 January 2020

References

- Leary, T. *Interpersonal Diagnosis of Personality: A Functional Theory and Methodology for Personality Evaluation* (Wiley, 1957).
- Fiske, S. T., Cuddy, A. J. & Glick, P. Universal dimensions of social cognition: warmth and competence. *Trends Cogn. Sci.* **11**, 77–83 (2007).
- Oosterhof, N. N. & Todorov, A. The functional basis of face evaluation. *Proc. Natl Acad. Sci. USA* **105**, 11087–11092 (2008).
- Rosenberg, S., Nelson, C. & Vivekananthan, P. A multidimensional approach to the structure of personality impressions. *J. Pers. Soc. Psychol.* **9**, 283–294 (1968).
- Fiske, S. T., Cuddy, A. J., Glick, P. & Xu, J. A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *J. Pers. Soc. Psychol.* **82**, 878–902 (2002).
- Cuddy, A. J., Fiske, S. T. & Glick, P. The BIAS map: behaviors from intergroup affect and stereotypes. *J. Pers. Soc. Psychol.* **92**, 631–648 (2007).
- Asch, S. E. Forming impressions of personality. *J. Abnorm. Soc. Psychol.* **41**, 258–290 (1946).
- Wojciszke, B. Morality and competence in person-and self-perception. *Eur. Rev. Soc. Psychol.* **16**, 155–188 (2005).
- Stolier, R. M., Hehman, E. & Freeman, J. B. A dynamic structure of social trait space. *Trends Cogn. Sci.* **22**, 197–200 (2018).
- Schneider, D. J. Implicit personality theory: a review. *Psychological Bull.* **79**, 294–309 (1973).
- Peabody, D. & Goldberg, L. R. Some determinants of factor structures from personality-trait descriptors. *J. Pers. Soc. Psychol.* **57**, 552–567 (1989).
- Digman, J. M. Higher-order factors of the Big Five. *J. Pers. Soc. Psychol.* **73**, 1246–1256 (1997).
- Lay, C. H. & Jackson, D. N. Analysis of the generality of trait-inferential relationships. *J. Pers. Soc. Psychol.* **12**, 12–21 (1969).
- Thornton, M. A. & Tamir, D. I. Mental models accurately predict emotion transitions. *Proc. Natl Acad. Sci. USA* **114**, 5982–5987 (2017).
- Tamir, D. I. & Thornton, M. A. Modeling the predictive social mind. *Trends Cogn. Sci.* **22**, 201–212 (2018).
- Sutherland, C. A. et al. Facial first impressions across culture: data-driven modeling of Chinese and British perceivers' unconstrained facial impressions. *Pers. Soc. Psychol. Bull.* **44**, 521–537 (2018).
- Stolier, R. M., Hehman, E., Keller, M. D., Walker, M. & Freeman, J. B. The conceptual structure of face impressions. *Proc. Natl Acad. Sci. USA* **115**, 9210–9215 (2018).
- Xie, S. Y., Flake, J. K. & Hehman, E. Perceiver and target characteristics contribute to impression formation differently across race and gender. *J. Pers. Soc. Psychol.* **117**, 364–385 (2018).
- Goldberg, L. R. A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. *Pers. Psychol.* **7**, 7–28 (1999).
- Kriegeskorte, N., Mur, M. & Bandettini, P. Representational similarity analysis—connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* **2**, 4 (2008).
- Gara, M. A. & Rosenberg, S. Linguistic factors in implicit personality theory. *J. Pers. Soc. Psychol.* **41**, 450–457 (1981).
- Hehman, E., Sutherland, C. A., Flake, J. K. & Slepian, M. L. The unique contributions of perceiver and target characteristics in person perception. *J. Pers. Soc. Psychol.* **113**, 513–529 (2017).
- Goldberg, L. R. The structure of phenotypic personality traits. *Am. Psychol.* **48**, 26–34 (1993).
- Digman, J. M. Personality structure: emergence of the five-factor model. *Annu. Rev. Psychol.* **41**, 417–440 (1990).
- Vazire, S. Who knows what about a person? The self–other knowledge asymmetry (SOKA) model. *J. Pers. Soc. Psychol.* **98**, 281–300 (2010).
- Uleman, J. S. & Kressel, L. M. in *Oxford Handbook of Social Cognition* (ed. Carlston, D. E.) 53–73 (2013).
- Johnson, J. A. Measuring thirty facets of the Five Factor Model with a 120-item public domain inventory: development of the IPIP-NEO-120. *J. Res. Pers.* **51**, 78–89 (2014).
- Stolier, R. M. & Freeman, J. B. Neural pattern similarity reveals the inherent intersection of social categories. *Nat. Neurosci.* **19**, 795–797 (2016).
- Brooks, J. A. & Freeman, J. B. Conceptual knowledge predicts the representational structure of facial emotion perception. *Nat. Hum. Behav.* **2**, 581–591 (2018).
- Cuddy, A. J. C. et al. Stereotype content model across cultures: towards universal similarities and some differences. *Br. J. Soc. Psychol.* **48**, 1–33 (2009).
- Weisman, K., Dweck, C. S. & Markman, E. M. Rethinking people's conceptions of mental life. *Proc. Natl Acad. Sci. USA* **114**, 11374–11379 (2017).
- Sutherland, C. A. et al. Social inferences from faces: ambient images generate a three-dimensional model. *Cognition* **127**, 105–118 (2013).
- Koch, A., Imhoff, R., Dotsch, R., Unkelbach, C. & Alves, H. The ABC of stereotypes about groups: agency/socioeconomic success, conservative–progressive beliefs, and communion. *J. Pers. Soc. Psychol.* **110**, 675–709 (2016).
- Collava, J., Sutherland, C. & Rhodes, G. Testing the functional basis of first impressions: dimensions for children's faces are not the same as for adults' faces. *J. Pers. Soc. Psychol.* **117**, 900–924 (2019).

35. Yzerbyt, V. Y., Kervyn, N. & Judd, C. M. Compensation versus halo: the unique relations between the fundamental dimensions of social judgment. *Pers. Soc. Psychol. Bull.* **34**, 1110–1123 (2008).
36. Kraft-Todd, G. T. et al. Empathic nonverbal behavior increases ratings of both warmth and competence in a medical context. *PLoS One* **12**, e0177758 (2017).
37. Sutherland, C. A., Young, A. W., Mootz, C. A. & Oldmeadow, J. A. Face gender and stereotypicality influence facial trait evaluation: Counter-stereotypical female faces are negatively evaluated. *Br. J. Psychol.* **106**, 186–208 (2015).
38. Paunonen, S. V. & Ashton, M. C. Big Five factors and facets and the prediction of behavior. *J. Pers. Soc. Psychol.* **81**, 524–539 (2001).
39. Burron, B. E., Carlson, K. A., Getty, G. R. & Jackson, D. N. The effects of informational characteristics on the perception of real and hypothetical target persons. *Psychon. Sci.* **23**, 145–147 (1971).
40. Jackson, D. N. A model for inferential accuracy. *Can. Psychol.* **13**, 185–195 (1972).
41. Baumeister, R. F., Zhang, L. & Vohs, K. D. Gossip as cultural learning. *Rev. Gen. Psychol.* **8**, 111–121 (2004).
42. Paunonen, S. V. Consensus in personality judgments: moderating effects of target-rater acquaintanceship and behavior observability. *J. Pers. Soc. Psychol.* **56**, 823–833 (1989).
43. Judd, C. M. & Park, B. Definition and assessment of accuracy in social stereotypes. *Psychol. Rev.* **100**, 109–128 (1993).
44. Todorov, A., Olivola, C. Y., Dotsch, R. & Mende-Siedlecki, P. Social attributions from faces: determinants, consequences, accuracy, and functional significance. *Annu. Rev. Psychol.* **66**, 519–545 (2015).
45. Tskhay, K. O. & Rule, N. O. Accuracy in categorizing perceptually ambiguous groups: a review and meta-analysis. *Pers. Soc. Psychol. Rev.* **17**, 72–86 (2013).
46. Jussim, L., Crawford, J. T. & Rubinstein, R. S. Stereotype (in) accuracy in perceptions of groups and individuals. *Curr. Dir. Psychol. Sci.* **24**, 490–497 (2015).
47. Funder, D. C. On the accuracy of personality judgment: a realistic approach. *Psychol. Rev.* **102**, 652–670 (1995).
48. Fiske, S. T. & Neuberg, S. L. A continuum model of impression formation from category-based to individuating processes: influences of information and motivation on attention and interpretation. *Adv. Exp. Soc. Psychol.* **23**, 1–74 (1990).
49. Hehman, E., Xie, S. Y., Ofosu, E. K. & Nespoli, G. A. Assessing the point at which averages are stable: a tool illustrated in the context of person perception. Preprint at *PsyArXiv* <https://psyarxiv.com/2n6jq> (2018).
50. Ma, D. S., Correll, J. & Wittenbrink, B. The Chicago Face Database: a free stimulus set of faces and norming data. *Behav. Res. Methods* **47**, 1122–1135 (2015).
51. Thornton, M. A. & Mitchell, J. P. Theories of person perception predict patterns of neural activity during mentalizing. *Cereb. Cortex* **28**, 3505–3520 (2017).
52. Coleman, J. M. & Hong, Y.-Y. Beyond nature and nurture: the influence of lay gender theories on self-stereotyping. *Self Identity* **7**, 34–53 (2008).
53. Freeman, J. B., Stolier, R. M., Ingbreten, Z. A. & Hehman, E. A. Amygdala responsivity to high-level social information from unseen faces. *J. Neurosci.* **34**, 10573–10581 (2014).
54. Willis, J. & Todorov, A. First impressions: making up your mind after a 100-ms exposure to a face. *Psychol. Sci.* **17**, 592–598 (2006).

Acknowledgements

We thank M. B. Meshar, J. A. Chwe, B. Barnett, H. Woo and C. Cogley for assistance in materials development and data collection. This work was supported in part by research grants NIH-F31-MH114505 (R.M.S.) and NSF-BCS-1654731 (J.B.F.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

R.M.S., E.H. and J.B.F. developed the theoretical perspective. R.M.S. developed the study concepts. All authors contributed to the study design. Testing and data collection were performed by R.M.S. R.M.S. and E.H. performed the data analysis and interpretation. R.M.S. drafted the manuscript, and all authors contributed to edits and revisions. All authors approved the final version of the manuscript for submission.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41562-019-0800-6>.

Correspondence and requests for materials should be addressed to R.M.S. or J.B.F.

Peer review information Primary Handling Editor: Aisha Bradshaw

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2020

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- | | | |
|-------------------------------------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | The statistical test(s) used AND whether they are one- or two-sided
<i>Only common tests should be described solely by name; describe more complex techniques in the Methods section.</i> |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | A description of all covariates tested |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> | A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals) |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
<i>Give P values as exact values whenever suitable.</i> |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> | For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> | For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes |
| <input type="checkbox"/> | <input checked="" type="checkbox"/> | Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated |

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

Data were collected using web-based surveys implemented on Mechanical Turk.

Data analysis

All data analysis code are available on the manuscript's associated Open Science Framework web address (<https://osf.io/2uzsx/>). Analysis were completed in Python 2.7 and R.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

All stimuli, data, and analysis scripts (Python, R) are available or directed towards on the OSF (including figure source data), from which results may be reproduced (<https://osf.io/2uzsx/>).

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Data are quantitative. Studies 1-4, and 7 are correlational, and Studies 5 and 6 are experimental.
Research sample	The research sample consisted of workers from Amazon's Mechanical Turk. While not a fully representative sample, the participants are more diverse in age, race, and socioeconomic status than typical undergraduate research samples (Buhrmester, Kwang, & Gosling, 2011; Perspectives on Psychological Science). See demographic information below.
Sampling strategy	In studies 1 and 2, our target sample was 30 participants per trait rated in each of the social perception rating tasks below (face, familiar person, and social group trait tasks, as trait ratings across traits stabilize at approximately this number of participant raters; see Hehman, Xie, Ofosu, & Nespoli, 2018, retrieved from osf.io/mwtuz). Across traits and tasks, this totaled a target sample of n = 450 per social perception model. For the conceptual trait space model, involving conceptual ratings of traits with other traits, we based target sample size on prior work estimating a similar model (Brooks & Freeman, 2018 - Nature Human Behavior), seeking a target sample of n = 100. In studies 3, 4, 5, and 6, we chose to target a sample size of n = 200 following the sample size of a similar task design implemented in prior research (Stolier, Hehman, Keller, Walker, & Freeman, 2018 - PNAS). Study 7 compared our prior data (Study 1) to one new open dataset provided by other researchers, and thus the sample size was restricted to the total number of participants in that dataset (n = 307,313; https://osf.io/tbmh5/).
Data collection	All data were collected online via Amazon Mechanical Turk via an application in the web browsers on participants personal computers, and researchers were not present during data collection (however we cannot guarantee others were not present given the collection method). Experimenters were blind to any assignments of stimuli or experimental conditions to subjects.
Timing	Study 1 was collected between 2/20/2018 and 2/22/2018. Study 2 data between 7/19/2018 and 8/7/2018. Study 4 between 9/24/2018 and 10/10/2018. Study 5 between 2/4/2019 and 2/7/2019. Studies 3, 6, and 7 between 7/4/19 and 8/15/19.
Data exclusions	Where applicable, exclusion criteria were established in before data collection. For Studies 1, 2, and 3 no data were excluded from analyses (all data by participants who completed the task were included in analyses). In Study 4, 20 subjects were lost due to task incompleteness, and 5 subjects due to failure to follow instructions by solely hitting a single response option. In Study 5, 51 participants who failed the manipulation check were excluded from analyses (those who did not report the conceptual association direction, positive vs. negative, of their condition assignment).
Non-participation	No participants dropped out or declined participation.
Randomization	Studies 1 to 4, and 7 were not experimental. However, stimuli varied between participants, and were randomly assigned. Studies 5 and 6 was experimental, and in each participants were randomly assigned to one of two between-subjects experimental conditions through a randomization function in the study code.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input type="checkbox"/>	<input checked="" type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics

Participants were all U.S. residents and primary English-speakers. Study 1 demographics were: for conceptual trait task, age mean = 35.4 years, age standard deviation = 10.5 years; 58 Female, 55 Male, 2 other; 113 White, 2 other; for the face trait task, age mean = 35.5 years, age standard deviation = 12.3 years; 281 Female, 199 Male, 2 other; 372 White, 44 Black, 31 Asian, 37 other; for the familiar person trait task, age mean = 30.7 years, age standard deviation = 7.1 years; 308 Female, 175 Male, 16 other; 368 White, 44 Black, 42 Asian, 49 other; for the group trait task, age mean = 30.4 years, age standard deviation = 6.9 years; 297 Female, 183 Male, 5 other; 368 White, 44 Black, 39 Asian, 37 other.

Study 2 demographics were: for face trait task, age mean = 30.3 years, age standard deviation = 6.3 years; 257 Female, 237 Male, 2 other; 320 White, 101 Asian, 30 Black, 45 other; for the familiar person trait task, age mean = 29.8 years, age standard deviation = 6.3 years; 239 Female, 237 Male, 2 other; 309 White, 89 Asian, 40 Black, 40 other; for the group trait task, age mean = 30.4 years, age standard deviation = 6.7 years; 263 Female, 223 Male, 3 other; 315 White, 89 Asian, 43 Black, 42 other.

Study 3 demographics were: all United States residents; all primary English-speakers; Mage = 31.9 years, SDage = 5.9 years; 54 Female, 108 Male; 113 White, 33 Black, 9 Asian, 7 other.

Study 4 demographics were: for the face trait task, age mean = 31.44 years, age standard deviation = 5.50 years; 102 Female, 64 Male, 1 decline; 128 White, 23 Black, 5 Asian, 11 other; for the familiar person trait task, age mean = 32.34 years, age standard deviation = 6.52 years; 70 Female, 82 Male, 3 decline; 120 White, 20 Black, 6 Asian, 19 other; for the group trait task, age mean = 31.45 years, age standard deviation = 5.53 years; 72 Female, 90 Male; 126 White, 20 Black, 8 Asian, 8 other.

Study 5 demographics were: age mean = 32.69 years, age standard deviation = 6.44 years; 64 Female, 77 Male; 102 White, 18 Black, 9 Asian, 12 other.

Study 6 demographics were: all United States residents; all primary English-speakers; Mage = 37.2 years, SDage = 12.7 years; 69 Female, 77 Male; 115 White, 15 Black, 9 Asian, 7 other.

Study 7 demographics were: Mage = 25.2 years, SDage = 10.0 years; 185,149 Female, 122,164 Male; race/ethnicity data unavailable.

Recruitment

Participants were recruited from Amazon's Mechanical Turk user base. The only restrictions placed on the sample were age (above 18), nationality (born and raised in the United States), and an "approval rate" (indicating that the participant pays attention and follows instructions correctly in tasks) of over 90%. While all participants are self-selected due to interest and motivation to participate in research studies, this is unlikely to introduce bias into the sample since the participants are more diverse in age, race, and socioeconomic status than typical undergraduate research samples (Buhrmester, Kwang, & Gosling, 2011; Perspectives on Psychological Science), and studies have shown that MTurk workers provide high-quality data that replicates many classic findings in experimental psychology (Piolacci & Chandler, 2014; Current Directions in Psychological Science).

Ethics oversight

The study protocol was approved by the University Committee on Activities Involving Human Subjects at New York University.

Note that full information on the approval of the study protocol must also be provided in the manuscript.