Journal of Experimental Social Psychology 62 (2016) 103–117

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

My friends are all alike — the relation between liking and perceived similarity in person perception

Hans Alves *, Alex Koch, Christian Unkelbach

University of Cologne, Germany

HIGHLIGHT

- People have more knowledge about liked others compared to disliked others.
- Positivity displays a smaller diversity than negativity.
- People perceive liked others as more similar to one another than disliked others.
- Thinking about people's positive traits makes them appear more similar to one another.

ARTICLE INFO

Article history:
Received 19 May 2015
Revised 20 October 2015
Accepted 29 October 2015
Available online 30 October 2015

Keywords:
Person perception
Impression formation
Valence asymmetries
Law of effect

ABSTRACT

Past research showed that people accumulate more knowledge about other people and objects they like compared to those they dislike. More knowledge is commonly assumed to lead to more differentiated mental representations; therefore, people should perceive others they like as less similar to one another than others they dislike. We predict the opposite outcome based on the density hypothesis (Unkelbach, Fiedler, Bayer, Stegmüller, & Danner, 2008); accordingly, positive impressions are less diverse than negative impressions as there are only a few ways to be liked but many ways to be disliked. Therefore, people should perceive liked others as more similar to one another than disliked others even though they have more knowledge about liked others. Seven experiments confirm this counterintuitive prediction and show a strong association between liking and perceived similarity in person perception. We discuss the implications of these results for different aspects of person perception.

1. Introduction

It seems evident that liking breeds differentiation. Wine lovers can differentiate between a Merlot and a Syrah, art enthusiasts see the differences between a Monet and a Renoir, and soccer fans can distinguish between the playing styles of Lionel Messi and Thomas Müller. In social psychological research, this notion is found in broad phenomena like the out-group homogeneity effect (Quattrone & Jones, 1980; Park & Rothbart, 1982) or the cross-race effect (Feingold, 1914; Young, Hugenberg, Bernstein, & Sacco, 2012) — people differentiate better between members of their usually preferred in-groups and between faces of their own ethnic identity compared to out-group members and faces from other ethnicities. Liking breeds differentiation because information sampling follows a hedonic principle (e.g., Fazio, Eiser, & Shook, 2004; Thorndike, 1898). People seek interactions with persons they like and avoid interacting with disliked persons (Denrell, 2005). As a result, people's mental representations of liked others are highly differentiated as opposed to the rather shallow representations they have of disliked persons (e.g., Smallman & Roese, 2008). A direct implication of this liking-breeds-differentiation principle is that social perceivers should see liked persons as more diverse while disliked persons should all seem alike.

Despite the intuitive appeal of a general dislike-homogeneity phenomenon, a different line of research suggests that liking might go along with increased perceived similarity. According to the "Density Hypothesis", positive information is less diverse and thus more densely clustered in spatial representations compared to negative information (Unkelbach, Fiedler, Bayer, Stegmüller, & Danner, 2008; Unkelbach, 2012). Similar to the principle observable in facial attractiveness, there are only a few possible ways to be liked but many different ways to be disliked (Potter, Corneille, Ruys, & Rhodes, 2007). Here, we apply this principle to person perception; based on the density hypothesis, we present a model that assumes perceived similarity among other people to be based on their matching and non-matching features. Because people should represent liked persons with predominantly positive features and because positive features are less diverse, a counterintuitive

* Corresponding author at: Social Cognition Center Cologne, Richard-Strauss-Str. 2, 50931 Köln, Germany.
E-mail address: hans.alves@uni-koeln.de (H. Alves).

http://dx.doi.org/10.1016/j.jesp.2015.10.011
0022-1031/

scenario follows: Social perceivers should see other people they like (e.g., their friends) as more similar to one another than other people they dislike even though they have collected more knowledge about liked others and therefore have a more differentiated representation of them. The goal of the present work is to examine the relation between liking and perceived similarity in person perception, and to test whether the liking-breeds-differentiation principle causes people to perceive liked others as more diverse, or whether the small diversity of positivity makes all liked others appear similar.

The nature of the relation between liking and perceived similarity has strong implications for social perception: Perceiving others as similar or different from one another is a determinant of many social cognitive processes including social categorization (Tajfel, Billig, Bundy, & Flament, 1971; Billig & Tajfel, 1973), generalization and stereotyping (Ames, 2004; Gawronski & Quinn, 2013; Linville, Salovey, & Fischer, 1986), social comparison (Mussweiler, 2003), as well as person memory (Earles, Kersten, Curtayne, & Perle, 2008; Heathcote, Freeman, Etherington, Tonkin, & Bora, 2009).

In the following, we first introduce the concept of differentiation and explain why liking should breed differentiation, and accordingly, why social perceivers might perceive liked others as more diverse compared to disliked others. We then delineate the density principle, namely why positive information should display low diversity, and accordingly, why social perceivers might see persons they like as more similar, despite having more knowledge. The following empirical part presents data from seven experiments that systematically investigated the relation between liking and perceived similarity in person perception. Finally, we discuss implications from our research for different aspects of person perception such as mood effects, and social comparison processes.

1.1. Differentiation and evaluation

Social perception has two fundamental characteristics: it is driven by the process of differentiation, and its outcome is typically evaluative. Differentiation is a core concept of human perception and cognition, and it is essential for any kind of categorization and plays a particularly prominent role in social psychological theorizing. Seeing the differences and the similarities among individuals determines many aspects of social perception and behavior, such as perceived group membership (Campbell, 1958; Zárate & Sanders, 1999), social comparison (Mussweiler, 2003), or interpersonal interaction (Tajfel, 1982). For example, differentiation was described to be “at the heart of the stereotype concept” (Linville et al., 1986, p. 165) as stereotypes arise from a lack of differentiation between individuals (Park & Hastie, 1987). Consequently, much social psychological research is devoted to the perception of variability among group members (e.g. Judd & Park, 1988; Judd, Ryan, & Park, 1991; see Rubin & Badea, 2012 for an overview), including fascinating phenomena, such as the mentioned out-group homogeneity effect and the cross-race effect (e.g., Quattrone & Jones, 1980; Ostrom & Sedikides, 1992). While social perception is driven by differentiation, its outcome is typically evaluative. In order to navigate complex social environments, humans have to distinguish good from bad people and decide whom they like and whom they dislike (Lewin, 1935). In the present work we want to address how differentiation and evaluation, the two concepts most central to social perception, are related in people's mental representation of their social world.

More specifically, we ask whether people perceive others they like or others they dislike as more similar to one another. As introduced above, we believe there are two competing scenarios: (1) Liking goes along with increased perceived diversity (i.e., decreased perceived similarity); (2) liking goes along with decreased perceived diversity (i.e., increased perceived similarity). We first review evidence for the first scenario building on the principle that liking breeds differentiation.

1.2. Liking-breeds-differentiation

There is substantial evidence indicating that people have highly differentiated representations of the things and individuals they like. For example, people divide liked objects into more categories than disliked objects and use finer evaluative distinctions when expressing attitudes about liked vs. disliked stimuli (e.g., Smallman & Roese, 2008; Smallman, Becker, & Roese, 2014). In addition, research on the perception of group variability shows that perceptions of homogeneity tend to have negative associations while heterogeneity has positive associations. Groups that are perceived as homogenous “are usually low status, low power, minority groups, whose members are perceived in less individualistic terms, receive less attention, and display less positive emotions” (Badea, Brauer, & Rubin, 2012, p. 1094; see also Brauer & Bourhis, 2006; Fiske, Haslam, & Fiske, 1991).

But why should people have a more differentiated representation of the objects and persons they like? One answer is that human information sampling follows a hedonic principle, first described by Thorndike (1898) as the law-of-effect. Accordingly, people are more likely to engage in exploratory behavior (e.g., interactions with another person) if the behavior is expected to have positive consequences (see also Chen & Bargh, 1999; Fazio et al., 2004; Hirt, Melton, McDonald, & Harackiewicz, 1996). This hedonic principle of information sampling also serves the purpose of maintaining internal cognitive consistency (Festinger, 1954). As a result, people collect larger information samples (i.e., more knowledge) and derive a more differentiated representation about liked others than about disliked others (Denrell, 2005). Smallman and Roese (2008) explicitly expressed this idea as “to cherish a loved one is to relish the fine nuances of his or her personality” while “the rejected and forsaken are construed on a relatively surface level” (p. 1228).

As more differentiated representations typically go along with decreased perceived similarity (e.g. Goldstone & Steyvers, 2001; Nosofsky, 1986; Shepard, 1987), liked persons should be perceived as less similar to one another than disliked others. In line with this idea, Linville, Fischer, and Salovey (1989) showed that when people become more familiar with a person or a group, they perceive them as more diverse. Likewise, familiarity leads to more differentiated categorical constructions and evokes less generalizations among objects (Medin, Lynch, Coley, & Atran, 1997; Rota & Zellner, 2007).

Taken together, a large body of empirical research suggests that liking a person invites repeated exposure, increases the amount of knowledge, and thereby leads to a more differentiated representation of liked compared to disliked others, which indicates less similarity. From the perspective of a hedonic principle of information sampling, it seems evident that people should perceive liked others as diverse and disliked others as similar, constituting a general dislike-homogeneity effect.

1.3. Diversity of positive and negative impressions

Despite the arguments reviewed so far, we believe that the seemingly obvious negative relation between liking and perceived similarity might not hold in most contexts. In fact, we suggest that liking often comes with increased perceived similarity. Clearly, information sampling follows a hedonic principle and people have rich representations of the people they like. However, these rich representations should entail predominantly positive information, while disliked others should be represented by negative information. Crucially, information valence is confounded with content diversity; negative information is more diverse than positive information. As a result, mental representations of different disliked people can be rather different while mental representations of different liked others should be rather similar.

The idea that positivity comes in small diversity was introduced as the “Density Hypothesis” by Unkelbach et al. (2008). The authors showed that positive stimuli display a higher density in spatial displays of mental representations than negative stimuli, an ecological phenomenon that is apparent across many different stimuli classes including evaluative words, self-generated nouns, trait words, IAPS-pictures, as well as facial features (Alves et al., 2015; Bruckmüller & Abele, 2013; Koch, Alves, Krüger, & Unkelbach, under review; Koch, Alves, & Unkelbach, under review; Potter et al., 2007; Unkelbach, 2012; Unkelbach et al., 2008).

The limited diversity of positivity is not the result of biased information processing but reflects a robust principle of the world humans live in. Positive states usually constitute a norm state which is characterized by the absence of multiple different negative norm deviations (Clark & Clark, 1977). On any given dimension, the positive norm states typically occupies a single range close to the midpoint of the dimension, and is opposed by two negative deviations towards the two ends of the dimension. The top part of Fig. 1 illustrates this principle. For example, humans are able to survive only within a narrow range of temperature, atmospheric oxygen concentration, and electromagnetic radiation given off by the sun. For each of these dimensions, there is a "not enough" as well as a "too much" leading to a larger diversity of unlivable compared to livable conditions. Likewise, the human body is constantly engaged in maintaining internal homeostatic conditions such as body temperature or the blood glucose level (Cannon, 1926). In general, life as we know it is possible only within tight boundaries. These examples might appear unrelated to psychological processes, but the same principle applies more generally to qualities humans prefer and therefore refer to as "positive". Facial beauty is an illustrative example of this principle as there are many ways to be ugly, but only one way to be beautiful (Potter et al., 2007; Rhodes, 2006).

If the small diversity of positivity can be applied to person perception, positive impressions should be less diverse than negative impressions. This suggestion is supported by research showing that "positive" personality profiles tend to be those that show average scores on personality dimensions, which is why the correlation between item means of personality tests and item desirability typically exceeds $r = .80$ (e.g., Edwards, 1953; Leising, Ostrovski, & Zimmermann, 2013).

Fig. 1's lower half illustrates our application of the density hypothesis to person perception and why it should lead to greater perceived similarity among liked compared to disliked individuals. Let us assume three dimensions of social interactive behavior that one can observe during a social encounter (e.g., amount of talking, amount of laughing, amount of eye contact). Based on the reasoning above, the liked or

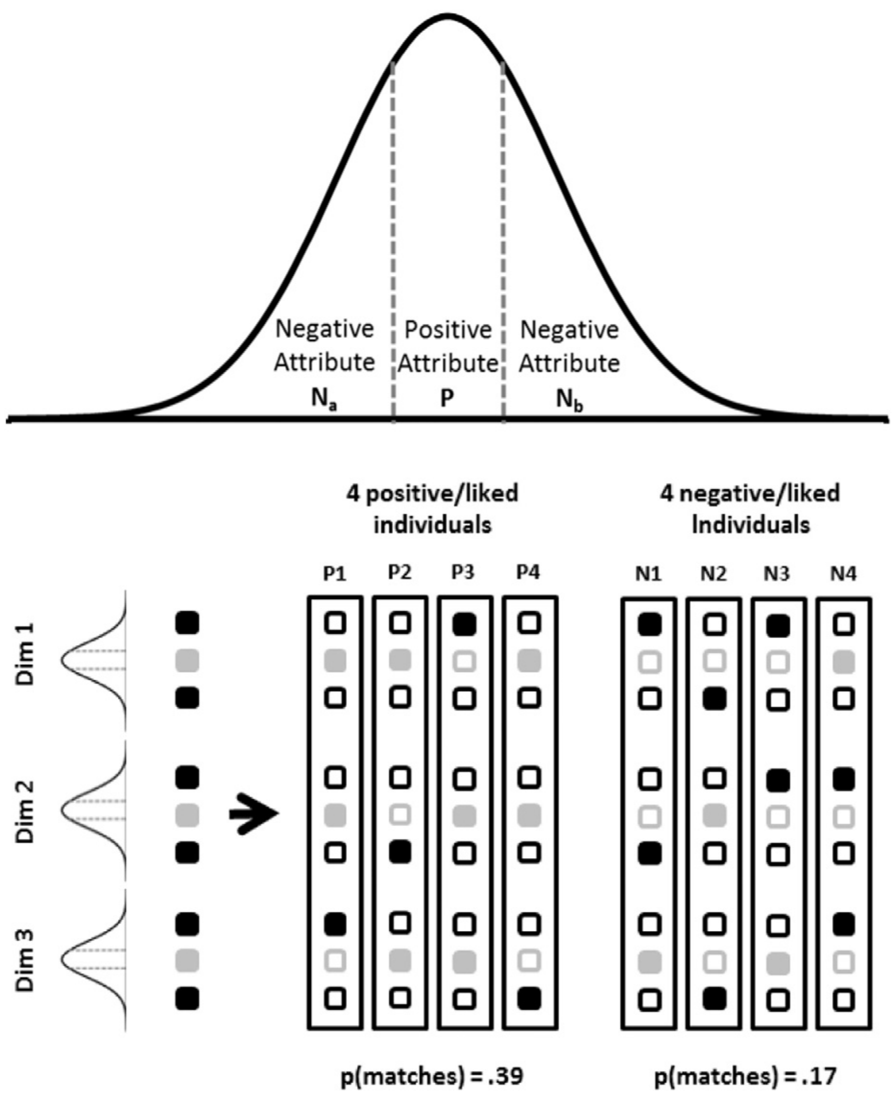

Fig. 1. The top part illustrates a normal distribution of values on a single dimension. Any positive quality is typically located in the center of the distribution and is opposed by two different negative qualities. The bottom part illustrates how this principle leads to higher similarity among positive/liked individuals compared to negative/disliked individuals when three attribute dimensions are combined. On each dimension, there is one positive attribute (grey square) and two negative attributes. Individuals that have mostly positive attributes (grey squares) vary along less attributes than individuals that have mostly negative attributes. In this example, the positive individuals on average have 1.16 out of 3 attributes in common, indicating a matching probability of .39, while the negative individuals have only 0.5 out of 3 attributes in common which equals a matching probability of .17. This principle can be applied to objectively present attributes as well as to subjective representations of other peoples' attributes.

"positive" range on each dimension is located around the midpoint of the dimension and constitutes a positive attribute, while there are two different negative ranges on each dimension that constitute two highly different attributes (e.g., barely talking vs. talking too much; not laughing vs. laughing all the time; no eye contact vs. constant eye contact). When combining the three dimensions, individuals can vary along three different positive attributes (gray squares) and along six different negative attributes (black squares). While a person can be perceived positively as talking, laughing, and making eye contact to an agreeable extent, he or she can be perceived negatively as incommunicative, or as a chatterbox; as humorless or clownish; and as making too little, or as making too much eye contact. Consequently, negative impressions come in greater diversity than positive impressions.

Fig. 1 further illustrates the mental representation of four liked individuals that each are represented by two positive and one negative attribute, and the mental representation of four disliked individuals that are represented by two negative and one positive attribute. Assuming that similarity is a function of matched and non-matched attributes/features (Tversky, 1977), the mental representations of liked individuals necessarily include more matching attributes than the representations of disliked individuals.

This principle should apply to "objective" attributes that target persons possess as well as to "subjective" representations of these attributes. In Garner's (1974) terminology, the principle is inherent in the intrinsic as well as in the extrinsic structure of the social world. That is, persons that display more likeable attributes than others should factually be more similar to one another, while at the same time persons that are perceived as more likeable than others by a given observer should also be perceived as more similar to one another. In the experiments that follow we investigate the latter case, that is, the subjective perception of similarities among liked and disliked others.

There is already empirical support for our model from research showing that the vocabulary for describing disliked persons is larger and more differentiated than the vocabulary for describing liked persons (Leising, Ostrovski, & Borkenau, 2012). Furthermore, research from Leising et al. (2013) suggests that people like others for similar reasons but dislike them for different reasons, a phenomenon that logically follows from our model (see also Leising, Erbs, & Fritz, 2010). When two people have a positive impression about another person, their impressions are likely to be highly similar, while negative impressions can be quite different.

In sum, the present model suggests a highly counterintuitive but theoretically predictable scenario: The people one likes and therefore knows the best will seem less diverse than the people one dislikes and therefore barely knows.

1.4. Overview of empirical investigation

We collected data from 7 experiments that investigated the relation between liking and perceived similarity in person perception. We employed different experimental designs (within, between), using different target persons (personally known to participants, celebrities), outcome measures (spatial arrangement, pairwise comparisons, trait overlap), and participants (students in Germany, Mturk participants in the US). All experiments applied a representative stimulus sampling approach (Brunswik, 1955), and asked participants to generate their own (liked and disliked) target persons, in order to increase external validity as well as construct validity (Wells & Windschitl, 1999; Judd, Westfall, & Kenny, 2012).

We used three different measures of perceived similarity. Experiments 1 and 2 used the spatial arrangement method (SpAM; Goldstone, 1994; Hout, Goldinger, & Ferguson, 2013), which builds on the intuitive and robust association between similarity and spatial distance (e.g. Casasanto, 2008; Lakoff & Johnson, 1980). SpAM asks participants to create their own mental map of a given set of stimuli by moving the stimuli around the computer screen so that distances between stimuli resemble their similarities (Goldstone, 1994; Hout et al., 2013; Koch et al., under review; Kriegeskorte & Mur, 2012).

We used pairwise similarity ratings as a second measure of perceived similarities in Experiments 3 to 7. This procedure presents participants with all possible pairwise combinations of stimuli and asks them to rate each pair's similarity. The similarity ratings serve either directly as a measure or the resulting similarity matrix can be analyzed using multidimensional scaling (MDS; Torgerson, 1965; Krumhansl, 1978). The MDS procedure estimates the coordinates of the stimuli in an n-dimensional space so that the Euclidean distances between stimuli are proportional to their similarities.

We employed a third measure of similarity in Experiments 4 and 7. Participants were asked to generate traits that applied to different target persons and we calculated the proportion of traits that applied to multiple target persons simultaneously (shared traits) as a proxy for perceived similarity.

In total, we conducted seven experiments to investigate the relation between liking and perceived similarity in person perception, and to test whether people perceive liked or disliked others as more similar to one another. All experiments we conducted are reported here, along with all experimental conditions and all collected variables.

2. Experiment 1

We started our empirical investigation with a simple design that asked participants to sample 4 personally-known others they liked and 4 personally-known others they disliked. First, we tested whether the hedonic principle of information sampling applied, namely that participants indeed collect more information about people they like than about people they dislike. Specifically, we expected participants to have spent more time with the liked target persons than with the disliked target persons, and that, accordingly, they have more knowledge about the liked compared to the disliked target persons. Second, we tested whether participants perceive liked or disliked target persons as more similar to one another.

2.1. Method

2.1.1. Participants and design

We had no specific prediction regarding the expected effect size in the first experiment; to achieve sufficient statistical power to detect small to medium effects (Cohen, 1988), we aimed for data from 70 participants. We factually collected data from 71 students of the University of Cologne (45 female, 26 male), who participated for a compensation of 3 Euros or for course credit. There was only one within-participants factor, as each participant provided ratings for the liked and disliked people they knew. Similarity assessment was realized using an adaptation of the spatial arrangement method (SpAM; e.g., Hout et al., 2013).

2.1.2. Procedure

Participants arrived in the laboratory and the experimenter seated them in front of a computer. If they agreed to participate after reading an informed consent, the experimenter started a Visual Basic program that presented instructions along with an adaption of the spatial arrangement method. First, participants were shown 8 text boxes and asked to provide forenames of four persons they knew and liked and forenames of four persons they knew but did not like. The order in which the program asked to provide names of liked and disliked persons was counterbalanced. The next screen showed participants 8 text labels displaying the provided names, arranged next to one another in the center of the screen. Instructions asked participants to create a similarity map by dragging and dropping the persons around the screen so that the distances between them indicated the similarities of their personalities (distant = dissimilar, close = similar).

Only after each name label was moved at least once, an "OK" button was activated, which enabled participants to submit their similarity maps. Fig. 2 shows an example of the spatial arrangement method. On the next screen, participants used sliders, ranging from 0 to 100, to indicate their liking for each of the 8 persons ("not at all" to "extremely"). This item served as a manipulation check. Participants also indicated the amount of knowledge they had about each person ("no knowledge at all" to "very much knowledge"), and the amount of time they had spent with each person ("no time at all" to "very much time") using the same type of sliders. At the end of the experiment, participants were thanked, paid, and debriefed about the purpose of the experiment. Experimental sessions lasted about 5 min.

2.2. Results

2.2.1. Liking, knowledge, and time spent together

We computed mean liking ratings for the liked and the disliked persons for each participant as a manipulation check. A paired-samples t-test confirmed that participants adhered to the instructions, as liking substantially differed between liked and disliked persons ($M_{liked} = 93.54$, $SD_{liked} = 6.54$ vs. $M_{disliked} = 18.97$, $SD_{disliked} = 9.96$), $t(70) = 46.70$, $p < .001$. We then computed mean values for knowledge and time spent together for the liked and disliked persons within each participant. As shown in Table 1, participants reported to have more knowledge about liked compared to disliked persons, $t(70) = 22.89$, $p < .001$, $d = 3.43$; and to have spent more time with liked compared to disliked persons in the past, $t(70) = 19.47$, $p < .001$, $d = 2.98$, confirming the expected pattern.

2.2.2. Similarity

The critical test regards the similarity (i.e., distance) between liked and disliked persons. The spatial arrangement method provides distances in percentage of the screen size. We calculated the mean distances between the four liked persons (i.e., six distances) and the mean distances between the four disliked persons for each participant. As shown in Table 1, the mean distance between the liked persons was smaller than the distance between the disliked persons, $t(70) = -5.33$, $p < .001$, $d = 0.75$, as predicted by the density hypothesis and the present model (cf. Fig. 1).

2.3. Discussion

Experiment 1 provides first support for the predicted counterintuitive scenario in person perception. First, the hedonic principle of information sampling clearly applies to participants' mental representations of other people. Participants indicated to have spent more time with liked than with disliked others and to have more knowledge

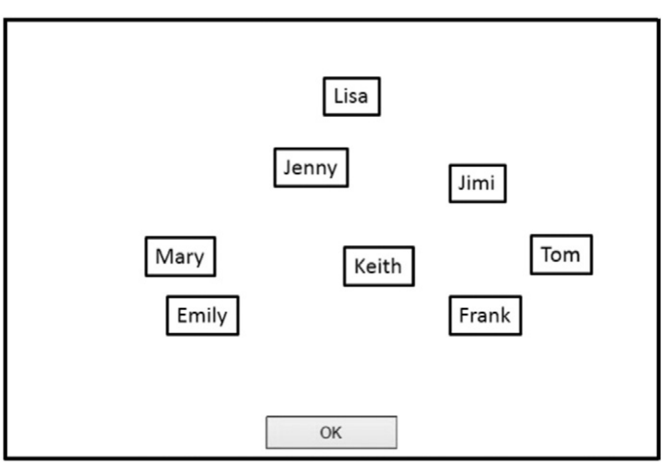

Fig. 2. Illustration of the spatial arrangement method (SpAM). Participants arranged liked and disliked others according to the similarities of their personalities.

Table 1
Mean and standard deviations of the dependent variables for liked and disliked target persons across Experiments 1 to 6.

	Liked	Disliked
Experiment 1a: spatial arrangement within		
Knowledge	90.93 (8.08)	47.55 (15.95)
Time spent together	88.80 (10.16)	40.07 (20.81)
Distance (screen size %)	11.08 (7.17)	18.81 (12.64)
Experiment 2: spatial arrangement between		
Knowledge	85.94 (7.51)	51.73 (18.56)
Time spent together	84.36 (7.47)	49.27 (17.81)
Distance	13.83 (8.67)	20.61 (12.34)
Experiment 3: pairwise similarity ratings		
Knowledge	90.54 (8.27)	42.98 (17.97)
Time spent together	89.69 (8.43)	39.84 (17.80)
(Mean Euclidean) distance	1.91 (1.06)	2.64 (1.06)
Experiment 4: trait generation		
Knowledge	90.01 (10.07)	42.41 (17.41)
Time spent together	88.67 (9.72)	37.88 (17.70)
(Mean Euclidean) distance	1.44 (0.86)	2.40 (0.93)
Experiment 5: celebrities		
Knowledge	56.25 (21.41)	40.63 (16.55)
(Mean Euclidean) distance	2.25 (1.02)	2.90 (1.05)
Experiment 6: celebrities sampled		
(Mean Euclidean) distance	2.31 (1.14)	2.64 (1.22)

about liked others. Second, despite this knowledge asymmetry, participants arranged liked others closer together indicating that they perceive them as more similar to one another regarding their personalities.

However, Experiment 1 allows a number of alternative explanations. First, the observed similarity asymmetry might hinge on the simultaneous arrangement method which requires participants to locate liked and disliked others within the same similarity space. A second concern regarding the spatial arrangement method is that participant's similarity dragging solutions might be influenced by other factors than perceived similarity. For example, they might be motivated to group liked others closely together to indicate that they stand together as friends. Likewise, participants might express that their liked others are "close with one another", meaning that they hold relationships with one another. A third concern relates to possible differences in the processing of positive and negative information (e.g. Taylor, 1991). It is a well-known phenomenon that negative stimuli trigger deeper and more elaborate processing which could increase differentiation and lead to the perception of dissimilarity. Fourth, we argue that despite the fact that participants have more knowledge about liked others, this knowledge indicates greater similarity. However, we do not know whether participants use this knowledge when judging between-person similarities. It is possible that they actually retrieve more knowledge about disliked others from their memory as would also be implied by a processing depth explanation. A final concern is that personally known liked others are likely to include persons that are related to one another, (i.e., family members), or who come from the same group, which increases their likelihood of actually possessing more similar personalities.

The following experiments will address these concerns. Experiment 2 first tests the possibility that our results hinge on the simultaneous arrangement of liked and disliked targets in the same similarity space. Simultaneous arrangement implies that similarities of liked and disliked targets should be judged on the basis of the same trait dimensions using equal resolutions of these dimensions. One might argue that people usually compare liked others on different and more specific trait dimensions than disliked others and perceive subtle differences on these dimensions as magnified.

Thus, liked targets might appear more similar to one another only in the context of disliked targets while the reverse might be true if both liked and disliked others are compared separately. To test this possibility, Experiment 2 asked for liked and disliked target persons between conditions. This allows participants in each condition to base their similarity judgments on different trait dimensions and to use different resolutions of these dimensions.

3. Experiment 2

Experiment 2 repeated Experiment 1 with a between-participants design. Instead of comparing both liked and disliked others, participants compared either liked or disliked others.

3.1. Method

3.1.1. Participants and design

Based on the data from Experiment 1, we expected a medium to large effect size for the similarity asymmetry. Therefore, we again aimed for data from 70 participants which provided us with sufficient statistical power for a between-participants comparison (Cohen, 1988). We factually collected data from 71 participants (47 females, 24 males) who were students from the University of Cologne and participated for a compensation of 3 Euros or for course credit. Procedure and materials were similar to Experiment 1, except that participants were randomly assigned to a "like" or "dislike" condition. Participants in the "like" condition were asked to provide forenames of four persons they knew and liked while participants in the "dislike" condition were asked to provide forenames of four persons they knew but did not like. Participants than performed the spatial arrangement task before they indicated liking, amount of knowledge, and amount of time spent together for each of the four persons.

3.2. Results

3.2.1. Liking, knowledge, and time spent together

As a manipulation check, we compared the liking ratings in the "like" condition with the ratings in the "dislike" condition. This difference was highly significant ($M_{liked} = 91.50$, $SD_{liked} = 6.99$ vs. $M_{disliked} = 26.51$, $SD_{disliked} = 14.13$), $t(69) = 24.46$, $p < .001$. As shown in Table 1, participants in the like condition also indicated to have more knowledge about the target persons than participants in the dislike condition, $t(69) = 10.13$, $p < .001$, $d = 2.44$. Participants in the like condition also reported to have spent more time with the persons than participants in the disliked condition, $t(69) = 10.77$, $p < .001$, $d = 2.59$, replicating the finding that liking goes along with more knowledge and exposure in a between-participants comparison.

3.2.2. Similarity

We calculated the mean distances between the four persons in the like condition and the mean distances between the four persons in the dislike condition. Table 1 shows that distances between target persons in the like condition were smaller than in the disliked condition, $t(69) = 2.67$, $p = .009$, $d = .64$. Liked others were again perceived as being more similar regarding their personalities than disliked others.

3.3. Discussion

Experiment 2 showed that the effects found in Experiment 1 did not hinge on a simultaneous comparison of liked and disliked others. When participants compared liked or disliked others in two conditions, they still perceived liked others as more similar to one another. Participants were not restricted to compare liked and disliked others within the same similarity space. Instead, participants in the liked condition had the possibility to base their similarity judgments on completely different trait dimensions and to use finer resolutions than participants in the disliked condition. In addition, participants still indicated to have more knowledge about liked others, and to have spent more time together with liked others.

Having ruled out a possible comparison set artifact, we now address other limitations related to the spatial arrangement method. Even though SpAM offers an elegant tool for simultaneous similarity assessment, this method leaves room for some alternative explanations regarding the similarity asymmetry we found.

First, participants might be motivated to spatially arrange liked others close together because they want them to "stand together" as friends. Second, the dense spatial arrangement of liked others might reflect participants' experience that their friends often appear together or even live close together. Further, participants might simply indicate that their liked others are "close with one another", or form a "circle of friends". To address these concerns, we conducted another experiment using pairwise similarity ratings instead of SpAM.

Yet another concern is that results might reflect affect-induced differences in the processing of positive and negative stimuli (e.g. Taylor, 1991; Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001). Accordingly, negative affect triggers deeper and more elaborate processing than positive affect. Assuming that the comparison of liked and disliked others triggers affective reactions in participants, it is possible that participants engage in deeper processing when comparing disliked others. As a result, they might even retrieve more knowledge about disliked others from memory and might conclude that they are dissimilar, while the more shallow processing when comparing liked others might give rise to a heuristic judgment that they are similar. To test such a processing explanation, we measured response latencies for each pairwise similarity ratings as a proxy for processing (e.g., Craik & Tulving, 1975). Assuming that response latencies are a measure for processing depth, this allows to test whether processing asymmetries can account for the differential similarity effect.

4. Experiment 3

Experiment 3 asked participants to compare the personalities of four liked and four disliked others using pairwise similarity ratings while we recorded response latencies for each comparison. We expected to replicate the similarity asymmetry from the previous two experiments and expected this asymmetry to be unrelated to possible processing depth differences as indexed by response latencies.

4.1. Method

4.1.1. Participants and design

Assuming a similar effect size, we again collected data from 70 participants (54 females, 16 males). All participants were students from the University of Cologne and participated for a bag of gummy bears or for course credit. There was only one within-participants factor; each participant rated similarities of liked as well as disliked persons.

4.1.2. Procedure

Similar to Experiment 1, participants first provided forenames of four liked and four disliked persons they personally knew. The next screen informed participants about the study's purpose to measure how similar these persons were regarding their personalities. Participants were further instructed that they would be asked to compare the persons with one another sequentially. Thus, each participant performed a total of 28 comparisons. Comparison orders were randomized for each participant. Each comparison trial showed a text label on the screen stating "How similar are [Person X] and [Person Y] to each other regarding their personalities?" Participants then provided their rating on a 9 point scale with endpoints labeled "not at all similar" (0) and "extremely similar" (9). For each comparison, the computer recorded the time it took participants to respond. After completing the 28 comparisons, participants indicated their liking for each of the 8 persons, the amount of knowledge they had about that person, and the amount of time they had spent with that person. Then experimenters thanked, paid, and debriefed participants about the purpose of the experiment. Experimental sessions lasted about 5 min.

4.2. Results

4.2.1. Liking, knowledge, and time spent together

Participants adhered to the instruction, as liking substantially differed between liked and disliked persons ($M_{liked} = 93.98$, $SD_{liked} = 6.33$ vs. $M_{disliked} = 19.50$, $SD_{disliked} = 12.05$), $t(69) = 42.62$, $p < .001$. Table 1 shows that participants again indicated to have more knowledge about liked compared to disliked persons, $t(69) = 19.59$, $p < .001$, $d = 3.40$, and to have spent more time with liked compared to disliked persons, $t(69) = 20.28$, $p < .001$, $d = 3.58$.

4.2.2. Similarity

To compare similarities among liked and disliked persons, we used the MDS procedure (Krumhansl, 1978) provided by the SAS system. We assumed an ordinal structure of the similarity ratings and submitted participants' ratings to separate MDS procedures. Analogous to the spatial arrangement method, we defined a two-dimensional similarity space and the MDS procedure estimated coordinates for each person in this space by minimizing a loss function. Based on each participant's MDS solution, we then calculated the mean Euclidean distances from each person to all other persons of the same valence (i.e., a "density" index for liked vs. disliked others; see Unkelbach et al., 2008). A high index indicates low density; the person is dissimilar to others. A low index indicates high density; the person is similar to others. For each participant, we then computed the mean Euclidean distance among the four liked persons and the mean distance among the four disliked persons. As shown in Table 1, the mean distance between the liked persons was substantially smaller than the distance between the disliked persons, $t(69) = -4.33$, $p < .001$, $d = .69$.

Next, we wanted to assure that the observed density difference did not hinge on the two-dimensionality of the MDS procedure. Yet, the differential density was also present when we estimated distances based on three-dimensional ($t(69) = -4.26$, $p < .001$), and four-dimensional scaling solutions ($t(69) = -4.27$, $p < .001$).

4.2.3. Response latencies

Each participant performed six comparisons of two liked persons and another six comparisons of two disliked persons. A paired-samples t-test found that participants performed the liked–liked comparisons faster than the disliked–disliked comparisons ($M_{liked} = 5.25$ s, $SD_{liked} = 1.96$ s vs. $M_{disliked} = 6.03$ s, $SD_{disliked} = 2.33$ s), $t(69) = -3.76$, $p < .001$, $d = .36$. This pattern is in line with affect-based processing theories which suggest that people engage in deeper processing when confronted with negative stimuli. Thus, we investigated if these latency differences accounted for the observed similarity differences.

4.2.4. Regression coefficient analysis

We conducted the reaction time analysis on the level of the individual pairwise similarity ratings. We analyzed whether response latency of a given comparison was related to its respective (similarity) rating and if so, whether this could account for the larger perceived similarity among liked vs. disliked targets. In other words, if processing asymmetries explain the observed pattern, longer latencies should indicate lower similarity. Conversely, if our model holds, similarity should simply be predicted by the valence of the target persons.

We tested whether response latency and valence predicted similarity ratings using regression coefficient analysis (RCS; Lorch & Myers, 1990; Thompson, 2008; see also Baayen, Davidson, & Bates, 2008, for a discussion). This procedure also enabled us to test whether the density asymmetry was evident at the level of individual similarity ratings that did not undergo a multidimensional scaling procedure. For each participant, there were twelve similarity ratings of interest (i.e., six like–like and six dislike–dislike comparisons), along with their respective response latencies. In addition, we included a dichotomous valence variable that contrasted these six liked–liked and the six disliked–disliked comparisons. For each participant we then specified a simultaneous regression model predicting the 12 similarity ratings by their respective response latencies and target valences. Regression could not be performed for one participant whose similarity ratings were always 1. We excluded this participant and then calculated mean standardized regression coefficients for reaction time and for valence across participants. Contradicting a processing depth explanation, response latencies did not predict similarity ratings (*Mean Beta* $= -.04$, $SD = 0.31$), $t(68) = -0.95$, $p = .34$. Thus, while it took participants longer to compare the personalities of disliked persons, this had no influence on their similarity ratings. In accordance with the density hypothesis, only target valence significantly predicted similarity ratings (*Mean Beta* $= .20$, $SD = 0.49$), $t(68) = 3.46$, $p = .001$, $d = .42$.

4.3. Discussion

Experiment 3 replicated the previous findings using pairwise similarity ratings instead of the spatial arrangement method. The data suggest that the differential similarity of liked and disliked persons is no artifact of spatial arrangement.

The experiment also provided some deeper insights. Participants took longer to make disliked–disliked comparisons than liked–liked comparisons, indicating a possible processing difference. While the observed latency difference mirrors the phenomenon that positive information is processed faster than negative information (e.g. Unkelbach, 2012; Unkelbach et al., 2008), regression coefficient analysis showed that response latencies did not predict similarity ratings, ruling out a processing depth explanation. In line with the model presented in Fig. 1, the only significant predictor for participant's raw similarity ratings was whether the target persons were liked or disliked.

The experiments so far showed that participants report more knowledge about liked others, but they seem to perceive them as more similar than disliked others regarding their personalities. However, it is unclear if participants use their knowledge when assessing the similarities, either pairwise or by spatial arrangement. The next experiment therefore aimed to activate participants' knowledge structures prior to similarity judgments by including a task that asked participants to generate for each of the persons as many traits as they could come up with.

This procedure also provides us with a more direct measure of participants' knowledge structures and it allows us to test our predictions in several new ways. First, we could test whether participants are able to generate more traits for liked than for disliked others as predicted by the hedonic principle of information sampling. Second, we could test whether these larger information samples are actually more similar; that is, we would expect participants to describe their liked others with relatively more of the same traits and their disliked others with different traits.

Third, analyzing the trait structure allows for another critical test of our model that can discount two other alternative explanation for the results we obtained thus far. That is, our model predicts that positive impressions are more similar than negative impression even across different perceivers. This means that we expect a participant's liked targets to show a stronger trait overlap with other participants' liked targets as well. Such a larger between participants trait overlap would discount the possibility that our effects are due to systematic differences between liked and disliked targets that are specific for a given participants. For example, participants might know more liked than disliked targets, which enables them to sample liked targets from the same context (e.g. family, group of friends) but forces them to sample disliked targets from different contexts. Further, as known from research on interpersonal attraction, people like others who are similar to themselves (Byrne, 1971; Berscheid, 1985; Montoya, Horton, & Kirchner, 2008). From this perspective one could argue that liked others might be perceived as more similar to one another because they are more similar to the perceiver.

These two alternative explanations predict a larger trait overlap among liked targets only within participants, but not across different participants. The present density model on the other hand predicts a larger trait overlap among liked targets both within participants as well as across different participants.

5. Experiment 4

Experiment 4 was similar to Experiment 3 except that participants were asked to generate for each target person as many traits as they could come up with prior to providing similarity judgments. We thereby tested whether knowledge activation influenced perceived similarity and whether the target persons' trait structure was in line with the present density model.

5.1. Method

5.1.1. Participants and design

Based on the previous effects, we collected data from 70 participants (48 females, 22 males). All 70 participants were students from the University of Cologne and participated for a bag of gummy bears or for course credit. There was only one within-participants factor, as each participant generated traits for and rated similarities of liked as well as disliked persons they knew.

5.1.2. Procedure

The procedure was almost identical to the previous experiments. Yet, participants were asked to provide for each person as many traits as they could come up with. Participants used text boxes to provide the traits while the order of the eight target persons was randomized. As in Experiment 3, participants then provided pairwise similarity ratings as well as liking, knowledge, and time spent together ratings. Finally, participants were thanked, paid, and debriefed about the purpose of the experiment. Experimental sessions lasted about 10 min.

5.2. Results

5.2.1. Liking, knowledge, and time spent together

Again, liking ratings differed substantially between liked and disliked target persons ($M_{liked} = 95.26$, $SD_{liked} = 5.07$ vs. $M_{disliked} = 18.50$, $SD_{disliked} = 13.85$), $t(69) = 41.00$, $p < .001$. Table 1 shows that participants indicated to have more knowledge about liked compared to disliked persons, $t(69) = 17.83$, $p < .001$, $d = 3.35$, and to have spent more time with liked than with disliked persons, $t(69) = 19.81$, $p < .001$, $d = 3.56$.

5.2.2. Similarity

We submitted each participant's pairwise similarity ratings to the same density analyses as in Experiment 3. As shown in Table 1, the mean Euclidean distance between liked persons was again smaller than the mean distance between disliked persons, $t(69) = -6.80$, $p < .001, d = 1.07$. Liked persons clustered more densely then disliked persons. The differential density was also present when we submitted similarity ratings to three-dimensional ($t(69) = -6.15$, $p < .001, d = 1.12$), and four-dimensional solutions ($t(69) = -6.53$, $p < .001, d = 1.12$).

5.2.3. Response latencies

We again computed the mean response latencies for the liked–liked and disliked–disliked comparisons. A paired-samples t-test found that participants again performed liked–liked comparisons faster than disliked–disliked comparisons ($M_{liked} = 4.74$ s, $SD_{liked} = 1.64$ s vs. $M_{disliked} = 5.53$ s, $SD_{disliked} = 1.85$ s), $t(69) = -3.49$, $p = .001; d = .45$.

5.2.4. Regression coefficient analysis

We again performed RCA to test if response latencies of the individual comparisons predicted the respective similarity ratings. As in Experiment 3, response latencies did not predict similarity ratings (*Mean Beta* $= -.02$, $SD = 0.35$), $t(69) = -0.37$, $p = .71$. However, as in the previous analyses, valence of the comparison targets did predict similarity ratings (*Mean Beta* $= .27$, $SD = 0.48$), $t(69) = 5.06$, $p < .001$, $d = .63$.

5.2.5. Trait overlap within participants

Experiment 4's trait generation task affords a new test of the knowledge asymmetry and of the density asymmetry. We excluded one participant from the following trait analysis as this participant failed to generate any traits for three of the four disliked persons. However, including this participant did not alter any of the following results.

In line with the proposed knowledge asymmetry, participants generated substantially more traits for liked persons ($M_{liked} = 6.67$, $SD_{liked} = 3.13$) than for disliked persons ($M_{disliked} = 3.91$, $SD_{disliked} = 1.76$), $t(68) = 11.59$, $p < .001$, $d = 1.09$.

To test for the proposed higher similarity of these larger information samples, we then computed the proportion of traits that participants assigned multiple times to different liked and different disliked others (i.e., how often participants used the same trait within a valence group). We calculated for each participant the mean proportion of traits that are shared among their liked targets and the proportion of traits shared among their disliked targets. Indeed, the proportion of shared traits was larger for the liked pairs ($M = .33$, $SD = .17$) than for the disliked pairs ($M = .22$, $SD = .17$), $t(68) = 5.06$, $p < .001, d = .65$. These numbers mean that 33% of the traits generated for a given liked target are on average shared with another liked target, while only 22% of a disliked target's traits are shared with another disliked target. Thus, participants generated more traits for the liked targets than for the disliked targets while they also generated a larger proportion of the same traits for their liked targets, supporting the prediction derived from our model.

There is one possible caveat arising from the fact that participants generated more traits for the liked than for the disliked others. We propose that this mirrors the knowledge asymmetry that participants also report on the explicit knowledge measure. However, this asymmetry might indicated that negative traits have a stronger impact on impressions than positive traits (see Baumeister et al., 2001 for a review). In order to be liked, many positive traits might have to be present while in order to be disliked, a few or even only one negative traits might be sufficient. This would in turn restrict the variation of liked individuals' traits as many positive traits have to simultaneously be present in them. There would simply be less degrees of freedom for variation among liked compared to disliked targets. If this alternative explanation holds, liked targets' larger trait overlap (proportion of shared traits) should be accounted for by the larger total number of traits that they are represented with.

We ran another regression coefficient analysis in order to test for this alternative explanation. For each of the 6 liked–liked and each of the 6 disliked–disliked target pairs, within each participant, we calculated the proportion of shared traits, serving as regression criterion, and we calculated the total number traits that each pair was described with, serving as predictor. Thus, for each of these twelve target pairs, we now had the proportion of shared traits, the number of traits and their valence (liked vs. disliked). We defined a regression model within each participant, predicting trait overlap (proportion of shared traits) by number of traits and valence. Ruling out the alternative explanation and in line with the density hypothesis, valence did predict trait overlap (*Mean Beta* $= .32$, $t(68) = 3.39$, $p = .001$, $d = .82$), but number of traits did not (*Mean Beta* $= -.03$, $t(68) = -.35$, $p = .72$). Thus, liked others show a stronger trait overlap independent of the number of generated traits.

5.2.6. Trait overlap across participants

As discussed above, a specific prediction of the density hypothesis and the present model is that trait overlap should also be larger for liked compared to disliked targets across different participants. We calculated the across-participants trait overlap separately for the eight target persons. For example, the traits of the first liked target that a given participant generated was compared to all first liked targets that the other participants generated, and so on. The resulting four liked and four disliked proportions were then averaged.[1] In line with the present density model, liked targets showed a larger trait overlap across participants compared to disliked targets ($M_{liked} = 0.07$, $SD_{liked} = 0.03$ vs. $M_{disliked} = 0.03$, $SD_{disliked} = 0.02$), $t(68) = 10.68$, $p < .001$; $d = 1.57$. These values mean that a liked target on average shared 7% of its traits with another participant's liked target, while a disliked target shares only 3% of its traits with another participant's disliked target. These differential proportions across participants only follow from the present model, but neither from the alternative that participants might sample liked others from the same and disliked others from different contexts (as this is not possible across participants), nor from the possibility that participants sample liked others that are similar to themselves and therefore similar to one other (as similarity to the self cannot play a role in trait overlap across participants).

5.3. Discussion

Experiment 4 replicated the findings from the previous experiments; participants rated liked persons as more similar than disliked persons. This effect occurred after participants initially engaged in a trait generation task which required them to thoroughly think about all persons. The trait generation task ensured that participants' knowledge structures were activated. This, however, did not weaken the similarity asymmetry; rather, the effect increased numerically ($d_{Exp.3} = 0.69$ vs. $d_{Exp.4} = 1.07$), $z = 1.53$, $p = .12$. Thus, even when participants carefully and explicitly considered what they knew about liked and disliked others, liked others appeared more similar to one another than disliked others. Participants again performed the liked–liked comparisons faster than the disliked–disliked comparisons; but again, response latencies did not predict similarity ratings while target valence did.

The trait generation task provided new insights regarding participants' knowledge structures. First, participants were able to generate almost twice as many traits for the liked targets compared to the disliked targets. Second, participants perceived liked targets' personalities as more similar to one another, as they assigned a larger proportion of the same traits to them. Thus, the asymmetry in participants' similarity judgments was in accordance with the traits they generated. This shows that even though participants have more knowledge about the personalities of people they like (i.e., larger information samples), this knowledge contains highly similar attributes (i.e., similar content); while knowledge about disliked people contains less but rather diverse attributes.

In addition, liked targets showed a larger trait overlap than disliked targets even across participants. That is, the liked targets of different participants were described with more of the same traits than the disliked targets. This is strong evidence for our model's prediction of a generally small positive diversity. In addition, this finding rules out the possibility that liked others have a larger trait overlap because they stem from the same context or because they are more similar to the perceiver. If this was true, liked others should be more similar than disliked others only within participants, but not across participants.

To summarize so far, Experiments 2–4 ruled out alternative explanations of the basic effect in Experiment 1 and supported the model put forth in Fig. 1. We now turn to a concern regarding the type of target persons used in the previous experiments. So far we showed the association between liking and similarity for the mental representation of persons that participants personally knew. The question remains to what extent the results are specific to these target persons. For example, participants might have selected multiple family members as personally known liked others. Liked others would then have more similar personalities due to the fact that they are relatives. In addition, liked others are more likely to know one another and to hold relationships with one another, while the same might not be true for personally known disliked others. Consequently, liked others might display stronger entitativity as a group than disliked others, which may cause participants to judge them as more similar (e.g., Lickel et al., 2000). Although the trait overlap data from Experiment 4 is not in line with this alternative, it seems prudent to rule out this alternative experimentally.

In addition, liked targets displayed a much smaller variability in terms of their likeability than disliked targets even within participants (Experiment 1: $Mean(SD_{liked}) = 6.16$ vs. $Mean(SD_{disliked}) = 13.62$; Experiment 2: $Mean(SD_{liked}) = 7.67$ vs. $Mean(SD_{disliked}) = 16.14$; Experiment 3: $Mean(SD_{liked}) = 6.01$ vs. $Mean(SD_{disliked}) = 14.21$; Experiment 4: $Mean(SD_{liked}) = 5.17$ vs. $Mean(SD_{disliked}) = 13.34$). This indicates that liked others are perceived as equally likeable while disliked others vary in terms of their likeability. Liked others might then be judged as more similar because they are similar in the sense that they are all liked to the same extent.

If any of these alternative explanations apply, results from the previous experiments would not be evidence for our model of limited positive diversity in person perception, but merely be an artifact of the target persons generated by participants. To rule out these alternative explanations and to put our model to another test, Experiment 5 used celebrities instead of personally known others as target persons.

6. Experiment 5

Experiment 5 tests the generality of our model by moving from the representation of personally known people to the representation of publicly known people. If liking goes along with decreased diversity, participants should perceive liked celebrities as more similar to one another than disliked celebrities.

6.1. Method

6.1.1. Participants and design

Based on the previous effects, we collected data from 71 participants (57 females, 14 males). All participants were students of the University of Cologne and participated for a bag of gummy bears or for course credit. There was only one within-participants factor, as each participant rated similarities of liked as well as disliked celebrities.

6.1.2. Procedure

The procedure followed Experiment 3 except that we asked participants to provide names of four liked and four disliked celebrities. Participants then provided pairwise similarity ratings while we recorded response latencies. Participants also provided liking and knowledge ratings for the celebrities, while time spent together ratings were omitted for obvious reasons. At the end, participants were thanked, paid, and debriefed about the purpose of the experiment. Experimental sessions lasted about five minutes.

6.2. Results

6.2.1. Liking and knowledge

Participants followed instructions, as liking substantially differed between liked celebrities and disliked celebrities ($M_{liked} = 80.40$, $SD_{liked} = 11.79$ vs. $M_{disliked} = 14.94$, $SD_{disliked} = 12.50$), $t(70) = 30.00$, $p < .001$. Table 1 shows that participants reported to have more knowledge

[1] We calculated the across-participants trait overlap within the eight different target persons to reduce the amount of necessary comparisons.

about liked celebrities compared to disliked celebrities, $t(70) = 8.47$, $p < .001$, $d = .82$.

6.2.2. Similarity

Table 1 shows the results of the same density analysis as in Experiments 3 and 4. The effect did not hinge on the dimensionality of the MDS solution as it was also present for three-dimensional ($t(70) = -3.19, p = .002, d = 0.55$), and four-dimensional solutions ($t(69) = -3.18, p = .002, d = .54$).

6.2.3. Response latencies

We again computed mean response latencies for the comparisons involving liked celebrity pairs and the comparisons involving disliked celebrity pairs for each participant. Different from the comparisons involving personally known others, participants performed liked–liked and disliked–disliked comparisons equally fast ($M_{liked} = 5.15$ s, $SD_{liked} = 2.08$ s vs. $M_{disliked} = 5.06$ s, $SD_{disliked} = 2.06$ s), $t(70) = 0.44$, $p = .664$.

6.3. Discussion

Experiment 5 asked participants to judge the similarities of celebrities. As Experiment 5 was identical to Experiment 3 in all other respects including sample size, comparing results from both experiments is of particular interest.

Participants indicated that they had more knowledge about liked celebrities than about disliked celebrities, but this effect was substantially smaller than for personally known others in Experiment 3 ($d_{Exp.5} = 0.82$ vs. $d_{Exp.3} = 3.40$), $z = 8.16$, $p < .001$. This suggests that the association between liking and knowledge is weaker when target information is acquired indirectly (e.g., via television) compared to when target information is acquired directly (via social interactions).

Importantly, participants also perceived liked celebrities as more similar to one another than disliked celebrities. This effect was of similar size as for personally known others in Experiment 3 ($d_{Exp.5} = 0.63$ vs. $d_{Exp.3} = 0.69$), $z = 0.25$, $p = .80$. As predicted, the density asymmetry is not limited to the representation of personally known people but applies to publicly known people as well.

This finding also excludes the possible alternative explanations that personally known others are perceived as more similar because they are related or because they display a larger entitativity due to personal relationships or group memberships they hold. Further, and in contrast to personally known target persons, liked and disliked celebrities were equal regarding their liking variability within participants ($Mean(SD_{liked}) = 11.13$ vs. $Mean(SD_{disliked}) = 11.52$), while the similarity difference was of similar size as in the previous experiments. This rules out the possibility that liked persons are judged as more similar because they are more similar in likeability. Different from the previous experiments, participants performed the liked–liked and disliked–disliked comparisons equally fast for celebrities. Thus, the density asymmetry occurred independently of the reaction time asymmetry, further ruling out processing depth explanations.

While results from the present experiment using celebrities make it unlikely that liked targets hold relations with one another possibly causing the similarity asymmetry, there remains a final concern. Possibly, participants know more liked celebrities than disliked celebrities, which enables them to sample liked target from the same context (e.g. TV show) but forces them to sample disliked targets from different contexts. Experiment 6 therefore asked participants to draw a larger and more exhaustive sample of liked and disliked celebrities to increase context diversity of target persons.

7. Experiment 6

Experiment 6 tests whether previous results were due to participants retrieving liked and disliked targets from unequally diverse contexts. To account for this possible asymmetry, we instructed participants to sample 12 liked and 12 disliked celebrities before only four of each group were randomly determined as target persons. We expected the density asymmetry to hold even when target persons are drawn from a larger and thus more divers sample of celebrities.[2]

7.1. Method

7.1.1. Participants and design

We expected the random selection of target persons from a larger sample to introduce error variance as similarity ratings are likely to be sensitive to target persons' proximity in the generated list. That is, if targets appear close to one another in the initial list, they are likely to stem from the same context and are thus more likely to be similar to one another. We therefore increased the sample size and collected data from 121 participants (63 females, 58 males). All participants were recruited online via the Mturk platform; all were located in the US and were compensated with $ 0.70.

7.1.2. Procedure

We asked participants to provide names of twelve liked and twelve disliked celebrities. Four of the liked and four of the disliked celebrities were randomly determined as target persons for which participants then provided pairwise similarity ratings. We omitted liking and knowledge ratings in order to keep the online experimental sessions short and increase the likelihood of successful completion. At the end, participants were asked to rate their level of concentration during the experiment (1 = "extremely low" to 6 = "extremely high"), before they were thanked, and debriefed about the purpose of the experiment. Experimental sessions lasted about 15 min.

7.2. Results

7.2.1. Similarity

We removed data from one participant who rated his level of concentration as "extremely low". We then conducted a similar density analysis as in Experiments 3, 4, and 5. As Table 1 shows, the mean Euclidean distance between liked celebrities was again smaller than the distance between the disliked celebrities, $t(119) = -2.73$, $p = .007, d = 0.28$. The effect did not hinge on the dimensionality of the MDS solution as it was also present for three-dimensional ($t(119) = -2.54$, $p = .012$), and four-dimensional solutions ($t(119) = -2.53, p = .013$).

7.3. Discussion

As in the previous experiments, liked target persons were again perceived to be more similar to one another than disliked target persons. Crucially, the present experiment ruled out the possibility that this effect was due to differences in retrieving a small number of liked and disliked celebrity targets. That is, even when participants generated a larger sample of 24 celebrities, random sub-samples of 4 liked and 4 disliked celebrities showed the same asymmetry. The observed effect was smaller compared to the previous experiments, which might be due to the procedural differences, that is, the online data collection process and/or the sampling of 24 liked and disliked celebrities. Some participants reported that they found it difficult to come up with 24 celebrities they feel strongly about. Despite these variations, the predictions from the present model hold and the density asymmetry seems to be a general characteristic of people's mental representation of liked and disliked people.

[2] Experiment 7 was conducted following the comments of two anonymous reviewers.

8. Experiment 7

So far, we have tested our model's predictions by comparing the perceived similarities of liked and disliked target persons, and we have addressed a number of possible alternative explanations for the similarity asymmetry. However, we cannot completely rule out that the liked and disliked target persons that participants retrieve differ on some unassessed variables that create the observed similarity asymmetry, meaning that confounds remain a possibility. Experiment 7 therefore puts our model to a final critical test that does not rely on contrasting liked and disliked targets. From our model's assumption that positive traits are less diverse than negative traits it follows that different people have rather similar positive traits, but different negative traits; positive traits should have a higher matching probability across different people. In other words, positive traits can be expected to constitute people's similarities, while negative traits constitute their differences. If this is true, focusing on people's positive or negative attributes should make them appear more or less similar to one another.

Therefore, Experiment 7 asked participants to think about two persons they personally know, without specifying whether target persons had to be liked or disliked. Participants then generated positive or negative traits for the target persons before they were asked to judge the similarity of the target persons' personalities. Based on our model of small positive diversity, we predicted that participants focusing on positive traits generate a larger proportion of the same traits for both target persons (more shared traits), and that this leads participants to perceive them as more similar compared to participants focusing on negative traits. In addition, we predicted a larger overlap of positive traits across participants as well.

8.1. Method

8.1.1. Participants and design

We had no specific prediction regarding the expected effect size; for sufficient statistical power to detect small to medium effects in a between-participants design (Cohen, 1988), we aimed for data from 100 participants. We factually collected data from 101 participants (49 females, 52 males). All participants were recruited online via the Mturk platform; all were located in the US and were compensated with $ 0.70. We asked participants to either generate positive or negative traits for two target persons.

8.1.2. Materials and procedure

The study was conducted using Qualtrics. Participants first read and agreed with a consent form before answering demographical questions. On the next screen, participants used text boxes to provide the forenames of two people they personally knew. Depending on condition, participants then provided as many positive (positive condition) or negative (negative condition) traits as they could come up with for each person by typing them into two separate text boxes. The next screen asked participants, “What do you think, how similar are [Person 1] and [Person 2] regarding their personalities?” Participants used a slider ranging from 0 (very dissimilar) to 100 (very similar) to indicate the target persons' similarity. At the end, participants were asked to rate their level of concentration during the experiment (1 = “extremely low” to 6 = “extremely high”), before they were thanked, and debriefed about the purpose of the experiment. Experimental sessions lasted about six minutes.

8.2. Results

8.2.1. Similarity

As predicted, participants who generated positive traits judged the target persons as more similar ($M = 56.78$, $SD = 24.39$) than participants who generated negative traits ($M = 38.78$, $SD = 26.81$), $t(99) = 3.53$, $p = .001$; $d = .71$.

8.2.2. Trait overlap within participants

We excluded one participant from the trait analysis because she described target persons using whole sentences instead of individual trait words. We counted the number of traits participants generated and calculated the mean number of traits generated in the positive and negative conditions. Participants in the positive condition generated almost twice as many traits compared to participants in the negative condition ($M_{positive} = 6.42$, $SD_{positive} = 2.84$ vs. $M_{negative} = 3.78$, $SD_{negative} = 1.66$), $t(98) = 5.69$, $p < .001$, $d = 1.15$.

We then computed the proportion of traits that participants assigned to both target persons. For each participant, we counted the number of traits assigned to both persons simultaneously (shared traits) and divided this by the mean number of assigned traits. Indeed, the proportion of shared traits was larger in the positive condition compared to the negative condition ($M_{positive} = .21$, $SD = .18$ vs. $M_{negative} = .08$, $SD = .15$), $t(98) = 5.69$, $p < .001$, $d = .79$. Thus, participants in the positive condition were more likely to assign the same traits to both target persons compared to participants in the negative condition.

8.2.3. Regression analysis

We then conducted a regression analysis to test our prediction that participants in the positive condition perceived target persons as more similar because the traits they generated had a higher matching probability, that is, they were more likely to be shared by both target persons. In a simultaneous regression we predicted similarity ratings by experimental condition (positive traits vs. negative traits), and by the proportion of shared traits generated. The zero order correlation between the two predictors valence and proportion of shared traits was $r = .36$, $p < .001$; the criterion (similarity) was positively correlated with valence ($r = .33$, $p = .001$) and proportion of shared traits ($r = .54$, $p < .001$). In a simultaneous regression, valence was no longer a significant predictor of perceived similarity ($\beta = .15$, $p = .108$), while proportion of shared traits remained a significant predictor ($\beta = .49$, $p < .001$). Hence, the proportion of shared traits that participants generated accounted for the effect of trait valence on perceived similarity of target persons.

Given that participants in the positive condition also generated a larger total number of traits this might cause them to generate a larger proportion of shared traits. If the trait sample size increases sampling becomes more exhaustive and traits might be more likely to overlap. This possibility was ruled out by the fact that when we entered the total number of traits into the regression model as an additional predictor, it did not predict similarity ($\beta = -.03$, $p = .737$), while proportion of shared traits remained a significant predictor of similarity ($\beta = .50$, $p < .001$).

8.2.4. Trait overlap across participants

Similar to Experiment 4, we then tested whether trait overlap across participants was also larger for positive than for negative traits. We calculated the proportion of shared traits among different participants' first and among different participants' second target persons in the positive traits condition and in the negative traits condition.[3] As predicted by the present density model, the mean proportion of shared traits was larger among participants in the positive traits condition compared to those in the negative traits condition ($M_{positive} = .08$, $SD = .03$ vs. $M_{negative} = .01$, $SD = .01$), $t(98) = 14.28$, $p < .001$, $d = 2.88$. This means that different participants describe their target persons in a much more similar way when they name positive compared to negative traits.

[3] Note again that trait overlap across participants was calculated within the two different target persons as it was done in Experiment 5.

8.3. Discussion

Experiment 7 investigated the flip side of our prediction that liked persons should be more similar than disliked persons; namely, that generating positive attributes makes people more similar, and generating negative attributes makes them more dissimilar. In contrast to the previous experiments, these results preclude any possible confounds related to self-sampling liked and disliked targets.

When participants generated positive traits, they described target persons with more of the same traits compared to participants in the negative condition. Again in line with our model, positive traits had a higher matching probability across participants' target persons, leading to higher perceived similarity in the positive condition as shown by the regression results. In addition, positive compared to negative attributes had a higher matching probability across different participants, providing strong support for our model shown in Fig. 1.

9. General discussion

The present work examines the relationship between liking and perceived similarity in person perception. We introduced two principles that make opposite predictions regarding this relationship, namely liking-breeds-differentiation and the density hypothesis. The liking-breeds-differentiation principle builds on the assumption that information sampling follows a hedonic principle (Denrell, 2005; Fazio et al., 2004). People spend the majority of their social interactions with persons they like, their partners, friends, and family members. In addition, people usually avoid the unpleasant experience of interacting with someone they do not like. As a result, people acquire more knowledge about the persons they like and therefore have highly differentiated representations of them. Because differentiation typically reduces perceived similarity (e.g. Nosofsky, 1986; Goldstone & Steyvers, 2001; Linville et al., 1989; Shepard, 1987), liked others should appear less similar to one another in people's mental representation than disliked others.

On the other hand, liking is based on positive representations, which according to the density hypothesis, are less diverse compared to negative representations (Unkelbach, 2012; Unkelbach et al., 2008). That is, each positive quality is usually defined by the absence of several different negative qualities. While the ways in which an object or a person can meet certain "positivity" criteria are limited, the ways to diverge from them are numerous. We summarized this in the model displayed in Fig. 1. Regarding the mental representation of other people, knowledge about liked others should contain highly similar (i.e. matching) attributes and therefore indicate greater similarity, while knowledge about disliked others should contain rather diverse attributes.

Results from seven experiments support the density hypothesis and our model. Even though participants had more knowledge about liked others they perceived liked others as more similar to one another than disliked others.

The hedonic principle of information sampling was also evident as participants consistently indicated to have more knowledge about liked others than about disliked others. In addition, Experiment 4 found that participants were able to describe liked others with almost twice as many character traits than disliked others. It seems to be the case that people are true experts regarding the people they like.

However, this knowledge asymmetry did not lead to the perception that liked others are less similar to one another than disliked others. In six out of the seven experiments we asked participants to compare the personalities of four liked and four disliked others, and liked others were clearly perceived as more similar to one another. Experiment 1 showed this asymmetry when participants simultaneously positioned liked and disliked others in a spatial arrangement procedure. Experiment 2 found the same asymmetry in a between-participants design which allowed participants in the liked and disliked conditions to base their similarity judgments on different dimensions with different resolutions. Experiment 3 replicated the effect using pairwise similarity ratings, thereby ruling out a number of alternative explanations related to the spatial arrangement method. Experiment 4 replicated the effect after participants engaged in a trait generation task, ensuring that their knowledge about the target persons was activated. This knowledge activation procedure did not weaken but strengthen the similarity asymmetry. Moreover, the greater perceived similarity among liked others was also visible among the traits that participants generated. Liked targets showed a larger trait overlap than disliked targets, both within and between participants. Experiment 5 found that the density asymmetry applies to personally as well as publicly known people ruling out additional alternative explanations regarding the relatedness of personally known others. Experiment 6 replicated the effect when participants generated a larger sample of liked and disliked celebrities from which targets were randomly chosen in order to increase variability among targets. Finally, Experiment 7 experimentally manipulated valence of participants' focus and found that focusing on target persons' positive traits compared to focusing on their negative traits increased perceived similarity.

In Experiments 3 to 5, we measured response latencies for the similarity judgments to test whether results were due to an affect-induced processing asymmetry (Taylor, 1991). It took participants indeed longer to compare personally known disliked others than liked others which might indicate deeper processing. Crucially, the response latencies did not influence the similarity ratings which were however predicted by the valence of the target persons. The reaction time difference was not present in Experiment 5 where participants compared celebrities, while the density asymmetry was still present.

To summarize, people adhered to a hedonic principle of information sampling and reported and showed more knowledge about liked others. Yet, the content of these larger samples consistently indicated higher similarity between liked compared to disliked others.

9.1. Implications

We believe the observed relationship between liking and perceived similarity affects many different aspects of person perception. First and foremost, our research shows that the liking-breeds-differentiation principle does not lead to the perception that disliked persons are generally more homogenous than liked persons. Given our data, the idea that people's preference to seek positive interactions and to avoid negative ones creates the perception that disliked members of their social world are all the same is not valid. Instead, people are surrounded by a social world in which negativity comes in various different forms and impressions about disliked persons contain little, but highly diverse information.

9.1.1. Positive similarities and negative differences

Another of our model's implication is that negative attributes tend to be different across people while positive attributes tend to be similar across people. Thus, while their negative traits make people unique, their positive traits make them similar. We directly tested this prediction in Experiment 7 and showed that focusing on the good (bad) aspects of other people makes them appear more (less) similar to each other. This notion has a powerful implication for apparent effects of mood on perceived variability. Accordingly, positive mood increases the perception of inter- and intragroup homogeneity. For example, happy participants viewing behavioral descriptions of highly variable groups perceived these groups as less diverse than participants in neutral moods exposed to the same descriptions (Queller, Mackie, & Stroessner, 1996). The authors suggested that happy perceivers focus on similarities rather than differences (Stroessner, Mackie, & Michalsen, 2005). When information that differentiates group members is made salient, happy participants no longer perceive them as more homogenous than their neutral mood counterparts. In line with this idea, Estes, Jones, and Golonka (2012) found that positive primes increased similarity ratings of unrelated social categories (e.g., musicians & dentists), and in an experiment by Isen and Daubman (1984), happy participants perceived exemplars as more similar to a given category than participants in neutral mood. The present research might provide a new explanation for these effects by assuming that participants in positive mood attend more strongly to positive and thus mood congruent information (e.g., Forgas & Bower, 1987). As our research shows that positive attributes are usually similar among objects and persons, representing stimuli (e.g., a group of people) primarily by their positive attributes will make them appear more similar to one another. The result that happy perceivers focus on similarities would therefore not be the explanation, but the outcome. Future research should test the possibility that the density principle together with the simple notion of mood-congruent processing can account for the well-documented finding that people in a good mood perceive the world in a more inclusive and integrative way (e.g., Bless & Fiedler, 2006).

9.1.2. Social comparison processes

The established positive relationship between liking and perceived similarity also has implications for social comparison processes. According to the selective accessibility model of comparison (Mussweiler, 2001, 2003), initial assessments of similarities between targets and standards determine whether judgments about the target are assimilated towards or contrasted away from the standard. If target and standard are perceived to be similar, target-congruent information is rendered accessible during the subsequent comparison process resulting in assimilation; when target and standard are perceived as dissimilar, target-incongruent information is more accessible resulting in contrast. Our results suggest that assimilation effects are more likely to occur when people compare others they like while contrast effects should more frequently occur when people compare others they dislike. At the same time, focusing on the similarities between other people might lead to more positive evaluations than focusing on the dissimilarities between people. This follows from the notion that people more strongly differ regarding their negative attributes than their positive attributes as a result of the larger diversity of negative attributes.

9.2. Reconciliation of two contradicting principles

Results from the present work follow the density hypothesis and contradict the liking-breeds-differentiation principle. However, we believe that both principles co-exist. To reconcile both theoretical accounts, we have to ask why on the one hand, liking is usually associated with greater differentiation and why on the other hand positive impressions are less diverse than negative impressions.

As outlined earlier, there is direct evidence for the influence of liking on differentiation. Smallman and Roese (2008) demonstrated that participants sort liked objects into more categories than disliked objects, indicating that liked objects are perceived as less similar. However, the authors used an evaluative conditioning procedure where neutral symbols (CS) were paired with either positive or negative pictures (US). Thus, the valence of the to-be-categorized target stimuli did not stem from the stimuli's own attributes. In this case, the density principle does not apply as it is based on an ecological diversity of attributes that positive and negative stimuli display.

A well-studied type of phenomena where liking also goes along with greater differentiation includes the out-group homogeneity effect and the cross-race effect (Feingold, 1914; Quattrone & Jones, 1980; Park & Rothbart, 1982; Young et al., 2012). Contrary to the density principle, people typically perceive members of their preferred in-group as less homogenous than members of an out-group. However, the preference for the in-group is most likely not based on certain positive and negative attributes of in- and out-group members (as proposed in our model), but on familiarity. People simply like others who they encounter frequently (Zajonc, 1968), which is also visible in people's general preference for objects, faces, or persons that are prototypical or "average" regarding a given environment (e.g., Langlois & Roggman, 1990). In other words, in-group members might be liked because social perceivers are repeatedly exposed to them, making them more familiar.

While familiarity leads to in-group preference it also decreases perceived similarity among in-group members (Linville et al., 1989). Given the simultaneous influence of familiarity on liking and perceived similarity, liking is sometimes negatively associated with perceived similarity. That is, on a group level analysis, a given in-group is perceived as more diverse than a given out-group. However, this magnifying effect does not even out the similarity asymmetry inherent in the evaluative information environment and thus does not lead to a general dislike-homogeneity effect. The fact that positivity exists within tight boundaries limits the possible diversity of positive impressions.

In our view, the density asymmetry and out-group homogeneity are not mutually exclusive but co-exist. For example, men should perceive female target persons as more similar to one another than male target persons (out-group homogeneity). This effect should be even larger when men compare liked women and disliked men (out-group homogeneity + density asymmetry). However, when men compare liked men and disliked women, the effect should vanish or even reverse (out-group homogeneity − density asymmetry). Research designs that demonstrate the existence of both, the density asymmetry and out-group homogeneity within the same paradigm would be a promising route for future research. Such a line of research could help to identify the boundary conditions of both effects and ultimately determine which effect dominates in which context.

One viable boundary condition that determines whether liking goes along with higher or lower perceived similarity is the type of attributes that similarity judgment are based on, more specifically, whether attributes are evaluative or non-evaluative. As argued above, people have more differentiated representations of other people they like as they accumulate a great amount of knowledge about these persons. This knowledge might often not be per se evaluative in nature. For example, one might learn about liked person's jobs, hobbies, family situation and eating habits. Liked others might therefore appear highly diverse regarding these non-evaluative attributes. Yet, in terms of evaluative attributes, and in particular their personality traits, the representation follows our model's prediction of a limited diversity of positive impressions.

9.3. Conclusion

In the opening of his novel "Anna Karenina", Leo Tolstoi stated, "Happy families are all alike; every unhappy family is unhappy in its own way." He thereby recognized the small diversity of positivity which we suggest to be a robust phenomenon that is inherent in the world humans live in, including their social world. Those people we represent with predominantly positive attributes and therefore like, appear highly similar. In the end, despite "relishing the nuances" of their personalities, our friends all seem alike.

Acknowledgments

The research was supported by a grant from the German Research Foundation (DFG; 273/1-1) awarded to the third author.

References

Alves, H., Unkelbach, C., Burghardt, J., Koch, A.S., Krüger, T., & Becker, V.D. (2015). A density explanation of valence asymmetries in recognition memory. *Memory & Cognition*, *43*, 896–909.

Ames, D.R. (2004). Strategies for social inference: A similarity contingency model of projection and stereotyping in attribute prevalence estimates. *Journal of Personality and Social Psychology*, *87*, 573–585.

Baayen, R.H., Davidson, D.J., & Bates, D.M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*, 390–412.

Badea, C., Brauer, M., & Rubin, M. (2012). The effects of winning and losing on perceived group variability. *Journal of Experimental Social Psychology*, *48*, 1094–1099.
Baumeister, R.F., Bratslavsky, E., Finkenauer, C., & Vohs, K.D. (2001). Bad is stronger than good. *Review of General Psychology*, *5*, 323–370.
Berscheid, E. (1985). Interpersonal attraction. In G. Lindzey, & E. Aronson (Eds.), (3rd ed.). *Handbook of social psychology*, *Vol. 2*. (pp. 413–484). New York: Random House.
Billig, M., & Tajfel, H. (1973). Social categorization and similarity in intergroup behaviour. *European Journal of Social Psychology*, *3*, 27–52.
Bless, H., & Fiedler, K. (2006). Mood and the regulation of information processing and behavior. In J.P. Forgas, K.D. Williams, & W. van Hippel (Eds.), *Hearts and minds: Affective influences on social cognition and behavior* (pp. 65–84). New York: Psychology Press.
Brauer, M., & Bourhis, R.Y. (2006). Social power. *European Journal of Social Psychology*, *36*, 601–616.
Bruckmüller, S., & Abele, A.E. (2013). The density of the big two: How are agency and communion structurally represented? *Social Psychology*, *44*, 63–74.
Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*, 193–217.
Byrne, D. (1971). *The attraction paradigm*. New York: Academic Press.
Campbell, D.T. (1958). Common fate, similarity, and other indices of the status of aggregates of person as social entities. *Behavioral Science*, *3*, 14–25.
Cannon, W. B. (1926). Physiological regulation of normal states: Some tentative postulates concerning biological homeostatics. *Jubilee volume to Charles Richet* (pp. 91–93). Paris: Editions Medicales.
Casasanto, D. (2008). Similarity and proximity: When does close in space mean close in mind? *Memory & Cognition*, *36*, 1047–1056.
Chen, M., & Bargh, J.A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin*, *25*, 215–224.
Clark, H.H., & Clark, E. (1977). *Psychology and language: An introduction to psycholinguistics*. New York: Harcourt Brace Jovanovich.
Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale: L. Erlbaum Associates.
Craik, F.M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*, 268–294.
Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review*, *112*, 951–978.
Earles, J.L., Kersten, A.W., Curtayne, E.S., & Perle, J.G. (2008). That's the man who did it, or was it a woman? Actor similarity and binding errors in event memory. *Psychonomic Bulletin & Review*, *15*, 1185–1189.
Edwards, A.L. (1953). The relationship between the judged desirability of a trait and the probability that the trait will be endorsed. *Journal of Applied Psychology*, *37*, 90–93.
Estes, Z., Jones, L.L., & Golonka, S. (2012). Emotion affects similarity via social projection. *Social Cognition*, *30*, 584–609.
Fazio, R.H., Eiser, J., & Shook, N.J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, *87*, 293–311.
Feingold, G.A. (1914). The influence of environment on identification of persons and things. *The Journal of Criminal Law and Criminology*, *5*, 39–51.
Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, *7*, 117–140.
Fiske, A.P., Haslam, N., & Fiske, S.T. (1991). Confusing one person with another: What errors reveal about the elementary forms of social relations. *Journal of Personality and Social Psychology*, *60*, 656–674.
Forgas, J.P., & Bower, G.H. (1987). Mood effects on person perception judgments. *Journal of Personality and Social Psychology*, *53*, 53–60.
Garner, W.R. (1974). *The processing of information and structure*. New York: Wiley.
Gawronski, B., & Quinn, K.A. (2013). Guilty by mere similarity: Assimilative effects of facial resemblance on automatic evaluation. *Journal of Experimental Social Psychology*, *49*, 120–125.
Goldstone, R. (1994). An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, & Computers*, *26*, 381–386.
Goldstone, R.L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, *130*, 116–139.
Heathcote, A., Freeman, E., Etherington, J., Tonkin, J., & Bora, B. (2009). A dissociation between similarity effects in episodic face recognition. *Psychonomic Bulletin & Review*, *16*, 824–831.
Hirt, E.R., Melton, R., McDonald, H.E., & Harackiewicz, J.M. (1996). Processing goals, task interest, and the mood–performance relationship: A mediational analysis. *Journal of Personality and Social Psychology*, *71*, 245–261.
Hout, M.C., Goldinger, S.D., & Ferguson, R.W. (2013). The versatility of SpAM: A fast, efficient, spatial method of data collection for multidimensional scaling. *Journal of Experimental Psychology: General*, *142*, 256–281.
Isen, A.M., & Daubman, K.A. (1984). The influence of affect on categorization. *Journal of Personality and Social Psychology*, *47*, 1206–1217.
Judd, C.M., & Park, B. (1988). Out-group homogeneity: Judgments of variability at the individual and group levels. *Journal of Personality and Social Psychology*, *54*, 778–788.
Judd, C.M., Ryan, C.S., & Park, B. (1991). Accuracy in the judgment of in-group and outgroup variability. *Journal of Personality and Social Psychology*, *61*, 366–379.
Judd, C.M., Westfall, J., & Kenny, D.A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, *103*, 54–69.
Koch, A., Alves, H., Krüger, T., & Unkelbach, C. (2015). *A general valence asymmetry in similarity: Good is more alike than bad*. (under review).
Koch, A., Alves, H., & Unkelbach, C. (2015). *Why good is more alike than bad: An ecological explanation of similarity asymmetries in evaluative information*. (under review).
Kriegeskorte, N., & Mur, M. (2012). Inverse MDS: Inferring dissimilarity structure from multiple item arrangements. *Frontiers in Psychology*, *3*.
Krumhansl, C.L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, *85*, 445–463.
Lakoff, G., & Johnson, M. (1980). The metaphorical structure of the human conceptual system. *Cognitive Science*, *4*, 195–208.
Langlois, J.H., & Roggman, L.A. (1990). Attractive faces are only average. *Psychological Science*, *1*, 115–121.
Leising, D., Erbs, J., & Fritz, U. (2010). The letter of recommendation effect in informant ratings of personality. *Journal of Personality and Social Psychology*, *98*, 668–682.
Leising, D., Ostrovski, O., & Borkenau, P. (2012). Vocabulary for describing disliked persons is more differentiated than vocabulary for describing liked persons. *Journal of Research in Personality*, *46*, 393–396.
Leising, D., Ostrovski, O., & Zimmermann, J. (2013). 'Are we talking about the same person here?': Interrater agreement in judgments of personality varies dramatically with how much the perceivers like the targets. *Social Psychological and Personality Science*, *4*, 468–474.
Lewin, K. (1935). *A dynamic theory of personality*. New York, NY, US: McGraw-Hill.
Lickel, B., Hamilton, D.L., Wieczorkowska, G., Lewis, A., Sherman, S.J., & Uhles, A.N. (2000). Varieties of groups and the perception of group entitativity. *Journal of Personality and Social Psychology*, *78*, 223–246.
Linville, P.W., Fischer, G.W., & Salovey, P. (1989). Perceived distributions of the characteristics of in-group and out-group members: Empirical evidence and a computer simulation. *Journal of Personality and Social Psychology*, *57*, 165–188.
Linville, P.W., Salovey, P., & Fischer, G.W. (1986). Stereotyping and perceived distributions of social characteristics: An application to ingroup–outgroup perception. In J.F. Dovidio, & S.L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 165–208). San Diego, CA, US: Academic Press.
Lorch, R.F., & Myers, J.L. (1990). Regression analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 149–157.
Medin, D.L., Lynch, E.B., Coley, J.D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*, *32*, 49–96.
Montoya, R.M., Horton, R., & Kirchner, J. (2008). Is actual similarity necessary for attraction? A meta-analysis of actual and perceived similarity. *Journal of Social and Personal Relationships*, *25*, 889–922.
Mussweiler, T. (2001). Focus of comparison as a determinant of assimilation versus contrast in social comparison. *Personality and Social Psychology Bulletin*, *27*, 38–47.
Mussweiler, T. (2003). Comparison processes in social judgment: Mechanisms and consequences. *Psychological Review*, *110*, 472–489.
Nosofsky, R.M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.
Ostrom, T.M., & Sedikides, C. (1992). Out-group homogeneity effects in natural and minimal groups. *Psychological Bulletin*, *112*, 536–552.
Park, B., & Hastie, R. (1987). Perception of variability in category development: Instance-versus abstraction-based stereotypes. *Journal of Personality and Social Psychology*, *53*, 621–635.
Park, B., & Rothbart, M. (1982). Perception of out-group homogeneity and levels of social categorization: Memory for the subordinate attributes of in-group and out-group members. *Journal of Personality and Social Psychology*, *42*, 1051–1068.
Potter, T., Corneille, O., Ruys, K.I., & Rhodes, G. (2007). 'Just another pretty face': A multidimensional scaling approach to face attractiveness and variability. *Psychonomic Bulletin & Review*, *14*, 368–372.
Quattrone, G.A., & Jones, E.E. (1980). The perception of variability within in-groups and out-groups: Implications for the law of small numbers. *Journal of Personality and Social Psychology*, *38*(1), 141–152.
Queller, S., Mackie, D.M., & Stroessner, S.J. (1996). Ameliorating some negative effects of positive mood: Encouraging happy people to perceive intragroup variability. *Journal of Experimental Social Psychology*, *32*, 361–386.
Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annual Review of Psychology*, 57199–57226.
Rota, L.M., & Zellner, D.A. (2007). The categorization effect in hedonic contrast: Experts differ from novices. *Psychonomic Bulletin & Review*, *14*, 179–183.
Rubin, M., & Badea, C. (2012). They're all the same!…but for several different reasons: A review of the multicausal nature of perceived group variability. *Current Directions in Psychological Science*, *21*, 367–372.
Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.
Smallman, R., & Roese, N.J. (2008). Preference invites categorization. *Psychological Science*, *19*, 1228–1232.
Smallman, R., Becker, B., & Roese, N.J. (2014). Preferences for expressing preferences: People prefer finer evaluative distinctions for liked than disliked objects. *Journal of Experimental Social Psychology*, 5225–5231.
Stroessner, S.J., Mackie, D.M., & Michalsen, V. (2005). Positive mood and the perception of variability within and between groups. *Group Processes & Intergroup Relations*, *8*, 5–25.
Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology*, 331–339.
Tajfel, H., Billig, M.G., Bundy, R.P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, *1*, 149–178.
Taylor, S.E. (1991). Asymmetrical effects of positive and negative events: The mobilization–minimization hypothesis. *Psychological Bulletin*, *110*, 67–85.

Thompson, G.L. (2008). Eliminating aggregation bias in experimental research: Random coefficient analysis as an alternative to performing a 'by-subjects' and/or 'by-items' ANOVA. *Tutorial in Quantitative Methods for Psychology*, *4*, 21–34.
Thorndike, E.L. (1898). *Animal Intelligence: An Experimental Study of the Associative Processes in Animals (Psychological Review, Monograph Supplements, No. 8)*. New York: Macmillan.
Torgerson, W.S. (1965). Multidimensional scaling of similarity. *Psychometrika*, *30*, 379–393.
Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.
Unkelbach, C. (2012). Positivity advantages in social information processing. *Social and Personality Psychology Compass*, *6*, 83–94.
Unkelbach, C., Fiedler, K., Bayer, M., Stegmüller, M., & Danner, D. (2008). Why positive information is processed faster: The density hypothesis. *Journal of Personality and Social Psychology*, *95*, 36–49.
Wells, G.L., & Windschitl, P.D. (1999). Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin*, *25*, 1115–1125.
Young, S.G., Hugenberg, K., Bernstein, M.J., & Sacco, D.F. (2012). Perception and motivation in face recognition: A critical review of theories of the cross-race effect. *Personality and Social Psychology Review*, *16*, 116–142.
Zajonc, R.B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, *9*, 1–27.
Zárate, M.A., & Sanders, J.D. (1999). Face categorization, graded priming, and the mediating influences of similarity. *Social Cognition*, *17*, 367–389.