PERSPECTIVE/OPINION

On the Assumptions That We Make About the World Around Us: A Conceptual Framework for Feature Transformation Effects

Jan De Houwer*, Juliette Richetin†, Sean Hughes* and Marco Perugini†

Various phenomena such as halo effects, spontaneous trait inferences, and evaluative conditioning have in common that assumptions about object features (e.g., whether a person is intelligent or likeable) are influenced by other object features (e.g., whether that person is attractive or co-occurs with other liked persons). Surprisingly, these phenomena have rarely been related to each other, most likely because different phenomena are described using different terms. To overcome this barrier, we put forward a conceptual framework that can be used to describe a wide range of these phenomena. After introducing the four core concepts of the framework, we illustrate how it can be applied to various phenomena. Doing so helps to reveal similarities and differences between those phenomena, thus improving communication and promoting interactions between different areas of research. Finally, we illustrate the generative power of the framework by discussing some of the new research questions that it highlights.

Research on impression formation reveals that people constantly generate assumptions about the features of the objects surrounding them. For instance, when told that someone solved a difficult puzzle, we spontaneously infer that this person is intelligent (Wells, Skowronski, Crawford, Scherer, & Carlston, 2011). Likewise, studies on the halo-effect demonstrate that attractive persons are assumed to be more intelligent than unattractive persons (e.g., Landy & Sigall, 1974; but see Eagly, Ashmore, Makhijani, & Longo, 1991, for moderators of this effect). Similar effects occur for inanimate objects as is, for instance, the case when a product that is said to be expensive is assumed to be of high quality (Rao & Monroe, 1989). In these and many other cases, assumptions about one object feature (e.g., whether a person is intelligent; whether a product is of high quality) are influenced by other object features (e.g., whether a person can solve puzzles or is attractive; whether a product is expensive).[1]

Phenomena like these can be found not only in the literature on impression formation but also in learning research. For instance, studies on evaluative conditioning demonstrated that the mere spatio-temporal pairing of a picture showing an unknown person with a picture showing a liked person is enough to make people evaluate

the unknown person as being positive (see Hofmann, De Houwer, Perugini, Baeyens, Crombez, 2010, for a review). Also other features (e.g., the extent to which an unknown person is judged to be athletic) can be influenced by the act of pairing stimuli (e.g., presenting pictures of that person together with pictures of known athletes; Förderer & Unkelbach, 2015). In all of these conditioning studies, the features of one object (a liked person or athlete) influence judgements about the features of another object (an unknown person). If we take a broad view on the concept of an 'object' that encompasses not only stimuli but also responses, then other phenomena such as operant evaluative conditioning effects can also be conceived of in this way. Operant evaluative conditioning refers to a class of phenomena in which the liking of stimuli changes because they are related to valenced responses (De Houwer, 2007; Eder, Krishna, & Van Dessel, 2019). For instance, in line with the assumption that approach is a more positive response than avoidance, neutral faces that are repeatedly approached are, under certain conditions, liked more than neutral faces that are repeatedly avoided (Van Dessel, De Houwer, & Gast, 2016). Operant conditioning effects like these can be interpreted as the impact of the features of one object (whether a response is positive or negative) on judgements about the features of another object (whether the stimulus that is related to the response is judged to be positive or negative).

Although there is thus a large variety of psychological phenomena that involve the impact of one object feature

* Ghent University, BE

† University of Milan-Bicocca, IT

Corresponding author: Jan De Houwer (Jan.DeHouwer@UGent.be)

on assumptions about other object features, often these phenomena are studied in isolation, with different groups of researchers each busy documenting the moderators and mediating processes for the specific phenomenon they are interested in, but rarely interacting with each other. We believe that this lack of interaction is at least in part due to the fact that different terms are used in research on different feature transformation effects. For instance, evaluative conditioning researchers typically use technical terms such as conditioned stimulus (CS) and unconditioned stimulus (US) that are highly specific to the phenomenon that they study. In fact, these terms are so specific that they do not even apply to closely related effects such as operant evaluative conditioning, where other concepts such as stimulus reinforcer (Sr) and discriminative stimulus (Sd) are used. In the literature on the halo effect and spontaneous trait inferences, there seem to be few if any generally agreed upon concepts to refer to procedural elements and effects. Instead, effects are often described topographically, that is, in terms of specific features (e.g., attractiveness) and objects (e.g., a person called Bob).

To overcome this fragmentation, we introduce a conceptual framework that can be applied to a wide range of these phenomena. Put differently, we propose a small set of concepts that can be used to describe a large set of phenomena. In what follows, we first describe the core concepts of our framework. We then illustrate the heuristic value of the framework by clarifying how it can be used (a) to describe various existing phenomena (which reveals the generality of the framework) and (b) to describe similarities and differences between those phenomena (which reveals the specificity allowed by the framework). Finally, we address the generative power of the framework by giving examples of how it highlights new questions and opportunities for research.

To prevent misunderstandings, we want to clarify from the outset that our framework is primarily a tool for improving description and communication (i.e., it is a *conceptual* framework). It is not (meant to be) a theoretical model that incorporates assumptions about the conditions under which certain phenomena can be found (i.e., it does not specify moderators) or about the processes that produce these phenomena (i.e., it does not specify mediators). Nevertheless, we believe that the framework is useful because, by promoting interactions between different areas of research, it helps discover gaps in the literature and stimulates new research.

The Conceptual Framework

As noted above, there are many phenomena in which information about certain object features influences assumptions about other object features. We propose to use the term "target feature" to refer to the object feature about which assumptions are being made. In experimental studies, the target feature can be easily identified on the basis of the dependent variable: it is the object feature that is assessed by the dependent variable. The term "source feature" can be used to refer to the object feature that influences assumptions about

the target feature. This second concept is related to the independent variable. More specifically, in experimental studies, the value of the source feature is varied across situations and it is examined whether this manipulation influences the value of the target feature. For instance, when people judge an attractive person to be more intelligent than an unattractive person, intelligence is the target feature whereas attractiveness is the source feature. Note that we use the term "feature" in a broad sense that encompasses any assumed state of an object that can have multiple values. These states can relate to many types of properties, such as physical (e.g., height), psychological (e.g., intelligence, valence)[2], and behavioral properties (e.g., the way in which an object responds to its environment).

When the source feature influences the assumptions that people make about the target feature, we say that feature transformation has taken place. With "assumptions", we mean the beliefs that people have about the value of target features as, for instance, indicated by their verbal reports or the choices that they make. We chose the term "transformation" because it can be used independently of the value of the source feature and the nature of the change in the target feature. Other terms like "transfer effect" are more restrictive in that they imply that the value of the target feature changes in the same direction as the value of the source feature (i.e., assimilative effects; e.g., judging attractive persons as more intelligent than unattractive persons). The term "transformation effect", on the other hand, encompasses any type of change in the value of the target feature, including changes in the opposite direction to the value of the source feature (i.e., contrast effects; e.g., judging attractive people as being less intelligent than unattractive people; also see Dymond & Rehfeldt, 2000).

Source and target features are typically part of objects that have also other features. Quite unimaginatively, we propose to refer to the object that holds the target feature as the "target feature object" and to use the concept "source feature object" to refer to the object that has the source feature. It is useful to distinguish target and source feature from target and source feature objects because the target and source features can either be part of the same object (e.g., when an attractive person is judged to be more intelligent than an unattractive person; Landy & Sigall, 1974) or they can be part of different objects (e.g., when the friend of an attractive person is judged to be more intelligent than the partner of an unattractive person; Sigall & Landy, 1973).

Illustrating the Heuristic Value of the Conceptual Framework

Applying the framework to existing phenomena

In this section, we make explicit how our conceptual framework can be applied to the various feature transformation effects that we alluded to in the introduction of this paper. More specifically, we clarify for a number of these effects which elements can be referred to as the source feature, the target feature, the source feature object, and the target feature object. Note, however, that

the framework can be applied to many other known and yet to be discovered feature transformation effects (e.g., to evaluative learning via intersecting regularities; see Hughes, De Houwer, & Perugini, 2016). Hence, this section is merely meant to illustrate the heuristic value of our framework.

Halo and horn effects

Let us start with the prototypical halo-effect: attractive persons are often assumed to be more intelligent than unattractive persons (Landy & Sigall, 1974). In terms of our framework, this is one instance of an assimilative feature transformation effect in which attractiveness is the source feature, intelligence is the target feature, and both features belong to the same object (i.e., the source and target feature object is one and the same object). Variants of the halo effect can also be accommodated within the framework. For instance, the term "horn effect" is sometimes used to refer to situations in which the source feature has a negative connotation and transforms the value of the target feature in an undesirable manner (e.g., baldness leads to attributions of less desirable personal and interpersonal characteristics, Cash, 1990). In other variants of the halo effect, the source and target feature are part of different objects (e.g., when examining the assumed intelligence of the partner of an attractive person; Sigall & Landy, 1973).

Spontaneous trait inference and transference

In a typical study on spontaneous trait inferences (Uleman, Saribay, & Gonzalez, 2008), researchers examine whether a source feature that refers to the behavior of the source feature object influences assumptions about a target feature of that same object. For instance, when told that the secretary solved the mystery halfway through the book, people will assume that the secretary is clever (Winter & Uleman, 1984). Solving the mystery is the source feature, intelligence is the target feature, and the source and target feature belong to the same object. In a related phenomenon called spontaneous trait transference, the source and target feature objects differ. For instance, people describing the behavior of others (e.g., solving a mystery) themselves become linked with the traits that are implied by that behavior (e.g., judged to be intelligent; see Skowronski, Carlston, Mae, Crawford, 1998, for evidence and Goren & Todorov, 2009, for boundary conditions).

Evaluative conditioning, attribute conditioning, and operant evaluative conditioning

Typical evaluative conditioning studies focus on valence as the source and target feature. They almost always involve separate target and source feature objects that are related in terms of their spatio-temporal presence. For instance, the target feature object can be an unknown person and the source feature object a liked person. Both objects are related in that they appear together in space and time. Results show that the valence of the liked person (source feature) influences the valence of the unknown person (target feature).[3]

The framework also applies to variants of evaluative conditioning. First, attribute conditioning refers to changes in a non-evaluative target feature (e.g., athleticism) as the result of pairing source and target features objects (e.g., a known athlete and an unknown person). Typically, the source feature (e.g., the fact that the athlete is known to be athletic) and target feature (e.g., the extent to which the unknown person is judged to be athletic) are identical (e.g., Förderer & Unkelbach, 2015). Second, in operant evaluative conditioning, the response is the source feature object (e.g., approaching), and the stimulus linked to the response is the target feature object (e.g., the stimulus that is repeatedly approached), the valence of the response is the source feature (e.g., the fact that approaching is positive), the valence of the stimulus linked with the response is the target feature (e.g., the liking of the stimulus that is repeatedly approached).

On the relation between various phenomena

Our conceptual framework not only provides ways of describing a wide range of phenomena with a common set of concepts but also ways of capturing some of the differences between different feature transformation effects. When described in terms of our framework, feature transformation effects can differ with regard to the nature of the four core elements (source feature, target feature, source feature object, target feature object) as well as the relations between those elements. More specifically, differences can be framed in terms of structure and content of the four core elements (see **Table 1**). At the structural level, the framework highlights two orthogonal dimensions on which feature transformation effects can differ: (1) whether the source and target feature are identical or different and (2) whether the source and target object are identical or different. At the content level, distinctions can be drawn with regard to the specific nature of the features, objects, and relations that are involved. Although the framework itself is not committed to specific ways of classifying features (e.g., evaluative, personality, behavioral, …), objects (e.g., persons, animals, inanimate objects, responses, …), or relations (e.g., equivalence, opposition, hierarchical, …), it does allow one to use various classifications to describe differences.[4]

In the remainder of this section, we illustrate how these ideas can help us clarify the differences between the different feature transformation effects that were described in the previous section. At the structural level,

Table 1: A taxonomy of feature transformation effects.

Source and Target Feature	Source and Target Feature Object	
	Identical	Different
Identical	(a)	(b)
Different	(c)	(d)

Note: Within each cell, further distinctions are possible on the basis of content, that is, the nature of the feature, object, and/or relation.

clear differences can be observed between typical halo and spontaneous trait inference effects, on the one hand, and typical (evaluative, attribute, or operant evaluative) conditioning effects, on the other hand. Whereas the former involve different source and target features that belong to the same object (e.g., attractiveness of a person influences assumptions about the intelligence of that same person; Cell c in **Table 1**), the latter incorporate identical source and target features that belong to different objects (e.g., pairing a liked and a neutral person influences the liking of the neutral person; Cell b in **Table 1**). Some instances of the halo effect and spontaneous trait inferences are structurally different from prototypical instances in that the source and target feature object are distinct rather than identical (e.g., attractiveness of a person influences assumptions about the intelligence of the partner of that person; Cell d in **Table 1**).

At the level of content, a number of interesting observations can be made. First, different feature transformation effects involve different types of source and target features. For instance, whereas halo effects typically involve an impact of physical or psychological source features on (physical, psychological, and behavioral) target features, spontaneous trait inferences involve effects of behavioral source features on (predominantly psychological) target features.

Second, target and source objects can be all kinds of living and non-living entities in the environment, including persons, animals, or products, and activities. Many existing feature transformation studies (e.g., on the halo effect and spontaneous trait inferences) focus on changes in assumptions about person features (Zebrowitz & Montepare, 2010). However, there is also substantial research on the transformation of features of products (e.g., Rao & Monroe, 1989).

Third, when the source and target features differ, the nature of the relation between both can vary. For instance, whereas in studies on the halo effect, the source feature (e.g., attractiveness) provides little diagnostic information about the target feature (e.g., intelligence), in studies on spontaneous trait inferences, the source feature (e.g., solving a mystery) provides potentially valid information for judging the target feature (e.g., intelligence).

Fourth, when the target and source feature objects differ, the nature of the relation between both objects can vary. For instance, when examining the effect of the attractiveness of a source feature person on assumptions about the intelligence of a different target feature person, one can vary information about the relation between the two persons (e.g., partner, friend, enemy). From this perspective, it is interesting to note that evaluative conditioning studies focus on one very specific type of relation between the target and source feature objects: their spatio-temporal relation (i.e., the fact that they occur together in space and time). Nevertheless, there have also been evaluative conditioning studies that added information about other relations between the two objects (e.g., whether they are friends or enemies; Fiedler & Unkelbach, 2011). Finally, when distinct source and target feature objects are used, one can also manipulate the temporal relation between the two objects. For instance, in studies on spontaneous trait inferences and halo effects, judgements about the target feature are typically made at a time at which also the source feature is present (often because both features are part of the same object). In conditioning studies, on the other hand, judgements about the target feature are made at a time when the source feature is no longer present (e.g., judging the valence of an unknown person after that person has been paired with a liked person).

Illustrating the Generative Power of the Framework

Our conceptual framework not only provides new ways of describing and organizing past research but also highlights new avenues for future research. At the structural level, it is striking to see that different research traditions limit themselves to a subset of the available options (see **Table 1**). For instance, in (evaluative, attribute, and operant evaluative) conditioning research, the source and target features are always identical (Cell b in **Table 1**). From the perspective of the framework, there is no reason why conditioning researchers cannot look also at the impact of source features on other target features (Cell d in **Table 1**). For instance, after pairing a first neutral face with a liked face and a second neutral face with a disliked face, one could assess changes in perceived intelligence of the neutral faces rather than changes in liking. These studies require only one change in the procedures that are used, namely a change in the target feature, that is, the feature assessed by the dependent variable (e.g., intelligence rather than liking). Given how minimal this procedural change is, it is striking to see that such studies have not yet been conducted. Vice versa, in the limited set of studies on the halo effect with distinct source and target feature objects, the source and target features have always been different (Cell d in **Table 1**). Except for the boundaries that research traditions place on creativity, we see no reasons why the source and target features could not be identical in those studies (Cell b in **Table 1**). For instance, instead of assessing how the attractiveness of a person influences judgments about the intelligence of the partner of this person (e.g., Sigall & Landy, 1973) one could assess how it influences the perceived attractiveness of the partner.

At the level of content, the conceptual framework highlights potential moderators that have not or hardly been examined in any of the existing areas of research on feature transformation effects. First, little or no research is available on the moderating impact of the nature of the relation between features and objects. Typically, researchers focus on the features that objects are assumed to have or display (e.g., to be attractive or to be able to solve complex puzzles). However, objects *having* features is just one of many ways in which objects and features can be related. For instance, objects can be said *not* to have a feature, to cause or predict other objects to have a feature, to be part of a class of objects that have that feature, to have more or less of a feature, or to simply be unrelated to a feature. For instance, rather than saying that a person can solve complex puzzles, one could say that the person cannot solve complex puzzles, helps other people solve complex puzzles, often interacts with other people who

solve complex puzzles, is to a certain extent better at solving complex puzzles than the average person, or that the ability to solve complex puzzles has not been assessed in that person. In each case, one could examine how this information influences the perceived intelligence of that person. A comparison of these different cases will reveal how variations in the type of relation between the object (e.g., the person) and the object feature (e.g., solving complex puzzles) moderates feature transformation (see Hughes, Ye, Van Dessel, & De Houwer, 2019, for an example of this type of research in the context of evaluative conditioning and Goren & Todorov, 2009, for an example in the context of spontaneous trait transference). Note that the nature of the object-feature relation can be varied not only for source feature objects (e.g., how a person is related to solving complex puzzles; see above) but also for target feature objects (e.g., how a person is related to intelligence, such as being intelligent, the likelihood to interact with other intelligent persons, having the potential to become intelligent …).

Second, there is little systematic research on the moderating impact of the relation between source and target features. When these features differ, there are different ways in which they can be related (e.g., semantically, affectively, with regard to the extent to which they co-occur in real life). For instance, whereas attractiveness and intelligence both have a positive valence, attractiveness and vanity are typically considered to be opposite in valence. Such differences are likely to have a profound effect on the direction and magnitude of feature transformation.

Finally, when the source and target feature object differ, it would be interesting to examine systematically how information about the nature of the relation between both objects (mere spatio-temporal relation, friends, enemies, members of a same group, …) moderates different types of feature transformation. One possibility is that people assume coherence in the way in which objects are related. For instance, people might think that if two objects are similar in one respect that they are similar also in other respects. In line with this shared features idea, recent studies showed that if a valenced word (e.g., HAPPY; source feature object) and a novel neutral word (target feature object) share an arbitrary feature like color, this results in a change of the valence of the originally neutral word (target feature) into the direction of the valence of the valenced word (source feature; Hughes, De Houwer, Mattavelli, & Hussey, 2019). In other words, if two objects are known to be similar in color, they are assumed to be similar in valence too. On the other hand, if two objects are opposite in one respect, people might assume that they are opposite also in other respects. Recent studies indeed showed that merely asking participants to press a left key upon presentation of a valenced word (source feature object) and a right key upon presentation of a novel neutral word (target feature object) changes the valence of the originally neutral word (target feature) in the direction opposite to the valence of the valenced word (source feature; Hussey & De Houwer, 2019). In other words, if two objects are known to be opposite in

terms of the responses to which they are assigned, they are assumed to be opposite in valence too. Formulated in terms of the conceptual framework put forward in this paper, these findings suggest that feature transformation involves not only a source and target feature (e.g., valence of the known word and the novel word, respectively) but also other contextual features (e.g., color of words, responses to which words are assigned) that moderate the way in which the source feature influences the target feature. Exploring these and other possibilities is bound to extend our understanding of feature transformation in general.

Final considerations and limitations
Our conceptual framework offers a unified way to describe and analyze feature transformation effects. Because it allows researchers to relate various phenomena using a common language, we hope that the framework will contribute to a broader and deeper understanding of the way in which people arrive at assumptions about the features of objects in their environment. Although the framework does not specify the mental processes that mediate feature transformation, it does provide an encompassing and theory-neutral set of concepts that can be used to describe feature transformation and its moderators. As such, the framework is bound to stimulate new research that will eventually lead to new insights in the underlying mental processes.

The main aim of the present paper was to introduce the framework and to provide an initial exploration of its heuristic and generative value. Hence, our treatment of the issue is far from exhaustive. In all likelihood, there are still other known phenomena that could be conceptualized as instances of feature transformation effects. For instance, the effects of persuasive messages (see Bohner, Erb, & Siebler, 2008) could also be described using our conceptual framework. Consider an advertisement in which a new product is described being high quality. As a result of seeing the advertisement, people assume that the product is expensive. This effect can be described as a feature transformation effect in which the source feature is (information about) product quality, the target feature is price, and both features belong to the same object. Hence, at the structural level, this type of persuasion has much in common with halo effects (i.e., both are instances of Cell c in **Table 1**). Other instances of persuasion could be conceived of as involving identical source and target features of a single object (i.e., Cell a in **Table 1**). For instance, one could examine the persuasive impact of specific types of information about product quality (e.g., results of scientific tests, consumer experiences, …; source feature) on the assumed quality of the product (target feature), either immediately after providing the information or after a delay. Another important area of research that could be looked at from the perspective of feature transformation concerns social biases. For instance, the fact that Black persons are assumed to be aggressive can be interpreted as a halo effect in which a source feature (race) transforms a target feature (aggressiveness) of the same object (see also Forgas

& Laham, 2016). More generally, our framework could be used as a tool to screen the psychological literature in order to highlight and organize existing knowledge about feature transformation. As such, we hope it will foster a new area of research on this important phenomenon.

Notes
[1] Although we realize that for human beings, other human beings are special objects with unique properties, in this paper we do not focus on these differences but simply refer to objects in general. The fact that our conceptual framework encompasses both effects with social and non-social objects does not imply any *a priori* assumptions about similarities or differences between feature transformation effects with social and non-social objects. The framework is also not committed to any assumptions about the mechanisms via which one object feature has an impact on judgements about another object feature. When we use the terms "impact" or "influence", we are referring to functional causation, that is, to the idea that one feature is a function of (i.e., dependent on) the other, just like the impact of an independent variable on a dependent variable can be described without making any assumptions about the mechanism via which this impact occurs (Chiesa, 1992).

[2] One could argue that valence (good or bad) is not a property of an object but the consequence of the interaction between a person and an object (Lewin, Heider, & Heider, 1936). Although an object can have desirable and undesirable features, it is thus problematic to say that an object has a valence. However, this issue is not central to the aims of our paper. When we say that the valence of an object changes, this refers either to the fact that assumptions about the desirable and undesirable features of the objects have changed or, relatedly, that (assumptions about) the ability of the object to evoke desirable or undesirable feelings has changed. Many studies examine such changes in the valence of objects because it is assumed that object valence has a profound impact on behavior (e.g., whether an object is approached or avoided).

[3] In conditioning terms, the source feature object is the unconditioned stimulus (US), the target feature object is the conditioned stimulus (CS), the source feature is the unconditioned response (UR), and the target feature is the conditioned response (CR).

[4] Whether the features or objects are identical or different is considered to be a structural property of the feature transformation effect in the sense that it can be determined for all possible types of features and objects.

Acknowledgements

Funding Information

Competing Interests
The authors do not have competing interests.

Author Contributions
· Contributed to the conception: JDH, SH, MP
· Drafted and/or revised the article: JDH, JR, SH, MP
· Approved the submitted version for publication: JDH, JR, SH, MP

References
Bohner, G., Erb, H.-P., & Siebler, F. (2008). Information processing approaches to persuasion: Integrating assumptions from the dual- and single-processing perspectives. In W. D. Crano & R. Prislin (Eds.), *Attitudes and attitude change* (pp. 161–188). New York: Psychology Press.

Cash, T. F. (1990). The psychology of physical appearance: Aesthetics, attributes, and images. In T. F. Cash & T. Pruzinsky (Eds.), *Body images: Development, deviance, and change* (pp. 51–79). New York: Guilford Press.

De Houwer, J. (2007). A conceptual and theoretical analysis of evaluative conditioning. *The Spanish Journal of Psychology, 10*, 230–241. DOI: https://doi.org/10.1017/S1138741600006491

Dymond, S., & Rehfeldt, R. A. (2000). Understanding complex behavior: The transformation of stimulus functions. *The Behavior Analyst, 23,* 239–254. DOI: https://doi.org/10.1007/BF03392013

Eagly, A. H., Ashmore, R. D., Makhijani, M. G., & Longo, L. C. (1991). What is beautiful is good, but…: A meta-analytic review of research on the physical attractiveness stereotype. *Psychological Bulletin, 110*, 109–128. DOI: https://doi.org/10.1037/0033-2909.110.1.109

Eder, A. B., Krishna, A., & Van Dessel, P. (2019). Operant evaluative conditioning. *Journal of Experimental Psychology: Animal Learning and Cognition, 45*, 102–110. DOI: https://doi.org/10.1037/xan0000189

Fiedler, K., & Unkelbach, C. (2011). Evaluative conditioning depends on higher order encoding processes. *Cognition and Emotion, 25*, 639–656. DOI: https://doi.org/10.1080/02699931.2010.513497

Förderer, S., & Unkelbach, C. (2015). Attribute conditioning: Changing attribute-assessments through mere pairings. *The Quarterly Journal of Experimental Psychology, 68*, 144–164. DOI: https://doi.org/10.1080/17470218.2014.939667

Forgas, J. P., & Laham, S. M. (2016). Halo effects. In R. F. Pohl (Eds.), *Cognitive illusions: Intriguing phenomena in judgement, thinking and memory* (2nd ed., pp. 276–290). New York: Psychology Press.

Goren, A., & Todorov, A. (2009). Two faces are better than one: Eliminating false trait associations with faces. *Social Cognition, 27*, 22–248. DOI: https://doi.org/10.1521/soco.2009.27.2.222

Hofmann, W., De Houwer, J., Perugini, M., Baeyens, F., & Crombez, G. (2010). Evaluative conditioning in humans: A meta-analysis. *Psychological Bulletin, 136*, 390–421. DOI: https://doi.org/10.1037/a0018916

Hughes, S., De Houwer, J., Mattavelli, S., & Hussey, I. (2019). The shared feature principle: If two objects share a feature, people assume those objects also share other features. *Manuscript under review*.

Hughes, S., Ye, Y., Van Dessel, P., & De Houwer, J. (2019). When people co-occur with good or bad events: Graded effects of relational qualifiers on evaluative conditioning. *Personality and Social Psychology Bulletin, 45(2)*, 196–208. DOI: https://doi.org/10.1177/0146167218781340

Hussey, I., & De Houwer, J. (2019). Discrimination leads to a reversed transfer of valence. Unpublished data.

Landy, D., & Sigall, H. (1974). Beauty is talent: Task evaluation as a function of the performer's physical attractiveness. *Journal of Personality and Social Psychology, 29*, 299–304. DOI: https://doi.org/10.1037/h0036018

Lewin, K., Heider, F., & Heider, G. M. (1936). *Principles of topological psychology*. New York: McGraw-Hill. DOI: https://doi.org/10.1037/10019-000

Rao, A. R., & Monroe, K. B. (1989). The Effect of Price, Brand Name, and Store Name on Buyers' Perceptions of Product Quality: An Integrative Review. *Journal of Marketing Research, 26*, 351–357. DOI: https://doi.org/10.2307/3172907

Sigall, H., & Landy, D. (1973). Radiating beauty: Effects of having a physically attractive partner on person perception. *Journal of Personality and Social Psychology, 28*, 218–224. DOI: https://doi.org/10.1037/h0035740

Uleman, J. S., Saribay, S. A., & Gonzalez, C. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual Review of Psychology, 59*, 329–360. DOI: https://doi.org/10.1146/annurev.psych.59.103006.093707

Van Dessel, P., De Houwer, J., & Gast, A. (2016). Approach-avoidance training effects are moderated by awareness of stimulus-action contingencies. *Personality and Social Psychology Bulletin, 42*, 81–93. DOI: https://doi.org/10.1177/0146167215615335

Wells, B. M., Skowronski, J. J., Crawford, M. T., Scherer, C. R., & Carlston, D. E. (2011). Inference making and linking both require thinking: Spontaneous trait inference and spontaneous trait transference both rely on working memory capacity. *Journal of Experimental Social Psychology, 47*, 1116–1126. DOI: https://doi.org/10.1016/j.jesp.2011.05.013

Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology, 47*, 237–252. DOI: https://doi.org/10.1037/0022-3514.47.2.237

Zebrowitz, L. A., & Montepare, J. M. (2010). Social Psychological Face Perception: Why Appearance Matters. *Social & Personality Psychology Compass, 2*, 1497–1523. DOI: https://doi.org/10.1111/j.1751-9004.2008.00109.x