

Research Article

The Self-Organization of Explicit Attitudes

Michael T. Wojnowicz,¹ Melissa J. Ferguson,¹ Rick Dale,² and Michael J. Spivey³

¹Department of Psychology, Cornell University; ²Department of Psychology, University of Memphis; and ³Department of Cognitive Science, University of California at Merced

ABSTRACT—How do minds produce explicit attitudes over several hundred milliseconds? Speeded evaluative measures have revealed implicit biases beyond cognitive control and subjective awareness, yet mental processing may culminate in an explicit attitude that feels personally endorsed and corroborates voluntary intentions. We argue that self-reported explicit attitudes derive from a continuous, temporally dynamic process, whereby multiple simultaneously conflicting sources of information self-organize into a meaningful mental representation. While our participants reported their explicit (like vs. dislike) attitudes toward White versus Black people by moving a cursor to a “like” or “dislike” response box, we recorded streaming x- and y-coordinates from their hand-movement trajectories. We found that participants’ hand-movement paths exhibited greater curvature toward the “dislike” response when they reported positive explicit attitudes toward Black people than when they reported positive explicit attitudes toward White people. Moreover, these trajectories were characterized by movement disorder and competitive velocity profiles that were predicted under the assumption that the deliberate attitudes emerged from continuous interactions between multiple simultaneously conflicting constraints.

Over the past two decades, much research in social psychology has suggested that the folk psychological concepts people use to predict and explain each other’s behavior, such as beliefs, fears, attitudes, and motives, can operate outside of cognitive control and perhaps introspective awareness (Bargh & Chartrand, 1999). For example, an implicit attitude toward a stimulus can be unintentionally activated by the mere presence of that stimulus. In many studies using personal interviews and self-

report questionnaires, very few undergraduate research participants reported preferring White people to Black people (Fazio, Jackson, Dunton, & Williams, 1995; Wittenbrink, Judd, & Park, 1997). Yet when these same participants were exposed to subliminal presentations of images of Black people, their recognition times during a linguistic test systematically sped up for subsequent negative words (e.g., “disaster,” “cancer”) and systematically slowed down for subsequent positive words (e.g., “sunshine,” “gift”), relative to their recognition times following subliminal presentations of images of White people. Given that many people demonstrate spontaneous initial biases toward traditionally stigmatized groups, how do they overcome these biases to explicitly report positive attitudes toward the same groups?

Researchers have proposed a variety of theoretical accounts to accommodate existing data on the formation of attitudes and choices. These accounts range from the broadly framed dual-attitude model (Wilson, Lindsey, & Schooler, 2000), to the more specific dual-process models (Devine, 1989; Smith & DeCoster, 2000), to dynamic-interaction models (e.g., Judd, Drake, Downing, & Krosnick, 1991; see also Roe, Busemeyer, & Townsend, 2001). What we find in common in all of these accounts is the coexistence of multiple attitudes and an emphasis on the temporal dynamics of how they influence evaluative responses. Rather than selecting among the specific theories, we invoked the encompassing theoretical framework of self-organization to guide an exploration of those temporal dynamics, and made specific predictions for what should result from multiple attitudes interacting over time.

Starting from the premise that mental representations in general are dynamically evolving states (Conrey & Smith, 2007), we suggest that explicitly reportable attitudes are merely the end result of a complex, nonlinear, time-dependent process of multiple less-explicit attitudes competing with one another over hundreds of milliseconds. As implemented in the brain, mental representations are distributed: Neural populations convey information through patterns of firing rates distributed across multiple neurons (Rogers & McClelland, 2004; Spivey, 2007),

Address correspondence to Michael T. Wojnowicz, G57 Uris Hall, Department of Psychology, Cornell University, Ithaca, NY 14850, e-mail: mtw28@cornell.edu.

even in higher-order decision-making regions (Bogacz & Gurney, 2007; Lapish, Durstewitz, Chandler, & Seamans, 2008). Contemporary researchers therefore model the decision-making process as a dynamic real-time evolution of a distributed pattern (Busemeyer & Townsend, 1993; Usher & McClelland, 2003). Preliminary mental representations provide rough sketches of information: In early moments of processing, distributed representations are partially consistent with multiple interpretations because of their proximity to multiple neural population codes. However, a continuous accrual of information causes the distributed pattern to dynamically “sharpen” into a confident (selected) interpretation, forcing other, partially activated, competing alternative representations, decisions, or actions to gradually die out.

Thus, in this self-organization framework, it is possible for one attitude (whose supporting biases rise quickly in activation) to be briefly prominent during early moments of forming an attitude choice, but for a different attitude (whose supporting biases are stronger but rise in activation more slowly) to take hold during later moments of forming that same attitude choice. The latter attitude will eventually activate other subsystems, such as language and memory, thus making the attitude seem explicit. What makes the first attitude implicit is not necessarily that it was generated in a different subsystem, but simply that it did not hold sway long enough to activate those language and memory subsystems.

This basic framework places cognitive processes in the same domain as many other natural phenomena that evolve through self-organizing dynamics (Kelso, 1995; Van Orden, Holden, & Turvey, 2003). Self-organizing systems may change states autonomously over time, even under constant input, because continuous interactions between component parts (e.g., neurons or brain regions) drive such systems through a series of intermediate states toward a stable steady state. In the brain, these self-organizing dynamics are driven by the fundamental principles of mental processing. Mental processing generically involves recurrent processing loops (or cyclic feedback) between higher-order integrative regions and lower-level informational sources (Lamme & Roelfsema, 2000; O’Reilly, 1998; Spivey, 2007). These higher-order integrative regions enforce representational competition, in which increasing the activation of one particular interpretation inhibits alternatives. In this way, the brain dynamically morphs highly probabilistic mental states into nearly discrete symbolic representations. Many behavioral studies have supported the idea that higher-order mental states continuously evolve through the dynamic satisfaction of multiple simultaneously conflicting constraints, even in the case of seemingly categorical decisions in speech perception (McMurray, Tanenhaus, Aslin, & Spivey, 2003), syntactic rule construction (Farmer, Anderson, & Spivey, 2007), person construal (Freeman, Ambady, Rule, & Johnson, 2008), and semantic categorization (Dale, Kehoe, & Spivey, 2007). In the present work, we extended this framework to self-reported attitudes regarding social preferences.

A self-organizing, explicitly reported attitude requires a set of informational sources, including, for example, semantic features, evaluative conditioning, personal memories, motivational value, and response context. These informational sources should continuously cascade intermediate results of processing into integrative decision-making regions, such as the basal ganglia (Bogacz & Gurney, 2007) and cortical motor areas (Cisek & Kalaska, 2005). These informational sources send simultaneous probabilistic support for multiple candidate decisions. For example, during the early moments of processing, semantic knowledge might be 70% supportive of a “like” decision and 30% supportive of a “dislike” decision. However, higher-order integrative regions force the potential evaluative representations to compete, and these regions then send top-down recurrent feedback to the informational sources. Gradually, through multiple cycles of recurrent processing, the system self-organizes into a coherent response (Spivey, 2007). From this perspective, the research documenting pro-White implicit attitudes (e.g., Fazio et al., 1995) suggests that a stimulus referring to a Black person, compared with a stimulus referring to a White person, may evoke greater conflict distributed across probabilistic information sources as the positive, deliberate evaluation dynamically emerges. If this idea is correct, a temporally fine-grained analysis should reveal that people’s explicit liking judgments for Black people and White people evolve in real-time processing with qualitatively different dynamics.

How might such dynamic information be captured in real time? The unfolding cognitive dynamics may be revealed in continuous motor output. Because mental processing is recurrent, motor representations begin specifying movement parameters probabilistically, rather than waiting for a perfectly completed cognitive command (Erlhagen & Schoner, 2002). In fact, motor commands may initiate movement before specifying a unique target destination, because motor trajectory parameters can be continually updated midflight (Henis & Flash, 1995). For example, manual reaches toward a verbally named target object (e.g., “candy”) curve more toward a distractor that has a similar sounding name (e.g., “candle”) than toward a distractor with a dissimilar name (Spivey, Grosjean, & Knoblich, 2005). When participants provide taxonomic classifications for taxonomically equivocal animals (e.g., “mammal” for whales), their manual reaches curve more toward the distractor (e.g., “fish”) than when participants provide the same classifications for taxonomically unambiguous animals (e.g., apes; Dale et al., 2007).

The motor execution of explicitly reported attitudes toward different ethnic groups may exhibit similar nonlinear dynamics. To test whether explicit attitudes toward potentially conflicting stimuli show such competition during mental processing, we tracked participants’ motor trajectories toward “like” and “dislike” responses that represented their explicit attitudes. Given the findings concerning implicit attitudes toward Black versus White people, and assuming that explicit attitudes dynamically emerge through self-organization, we predicted that

hand trajectories would show greater motor curvature toward a “dislike” response while participants positively evaluated Black people than while they positively evaluated White people. This motor curvature would reveal a greater influence of a “dislike” decision during the process of settling into an explicit “like” decision about Black people than in the process of making the same decision about White people. In addition, there are two fine-grained predictions that would result exclusively from a competitive dynamics account of this phenomenon. If the phrase “Black people” evokes elevated dynamic competition between simultaneously active “like” and “dislike” representations, movement trajectories for “Black people” should exhibit evidence of nonlinear dynamics in their velocity profiles, as well as increased spatial disorder in the curviness of the trajectories.

EXPERIMENT 1

Method

Streaming x - and y -coordinates of mouse-cursor movements were recorded from 68 Cornell University undergraduates (43 female and 25 male) as they performed a simple explicit-attitude task. Trials began with 2 s for participants to view the evaluative response options (“like” and “dislike”), which were randomly

assigned to the upper corners of the screen. Participants then clicked on a small box at the bottom of the screen to reveal a stimulus word or phrase and dragged the mouse toward their selected evaluative response to that stimulus. The 40 stimuli included the target stimuli, “Black people” and “White people,” as well as 19 positively valenced distractors (e.g., “sunshine,” “babies”) and 19 negatively valenced distractors (e.g., “rats,” “murderers”). These 40 stimulus words and phrases were presented in a randomly assigned order within each of two blocks. Responses to the two stimulus repetitions were averaged together to yield a single measurement for each participant for all statistical analyses. (Responses to the stimulus repetitions were not averaged for distributional analyses.) We analyzed data only from the 61 participants who selected “like” for both “White people” and “Black people” on both stimulus repetitions.

Results and Discussion

Compared with the trajectories for “White people,” the trajectories for “Black people” curved significantly more toward the “dislike” response option, as shown in Figure 1 (upper half). The maximum perpendicular deviation from a hypothetical straight line connecting the trajectory’s starting point and endpoint was

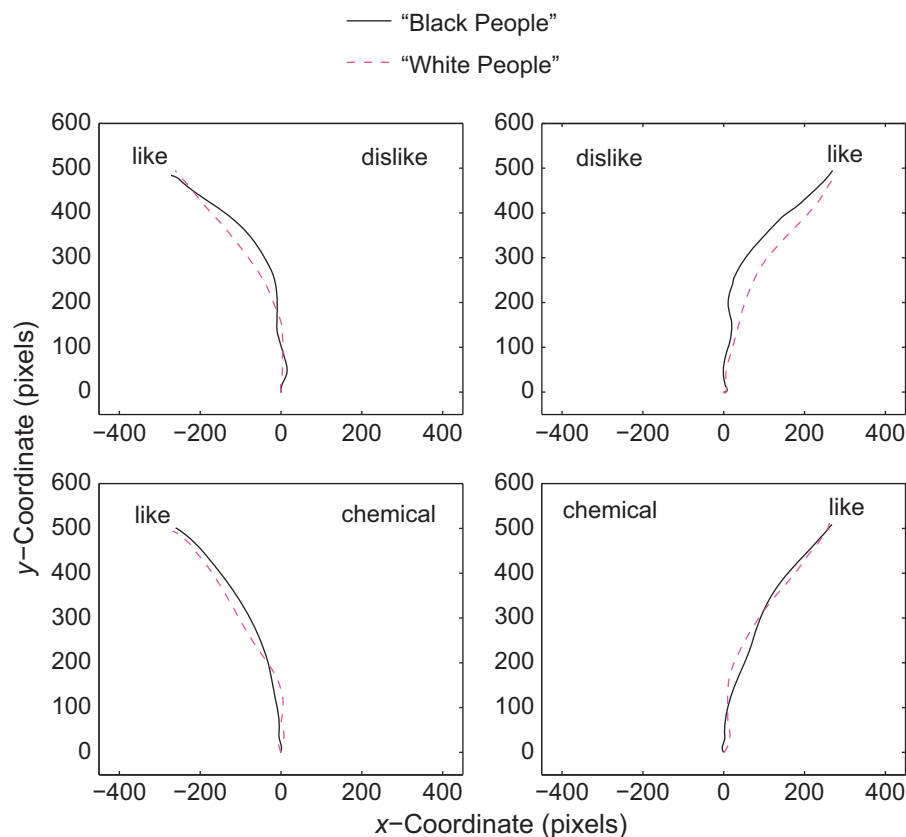


Fig. 1. Mean mouse-movement trajectories toward the evaluative response for the “Black people” and “White people” stimuli. In Experiment 1 (upper half), the participants guided the mouse from a starting point toward their choice of the “like” or “dislike” response box. In Experiment 2 (lower half), participants guided the mouse toward their choice of the “like” or “chemical” response box.

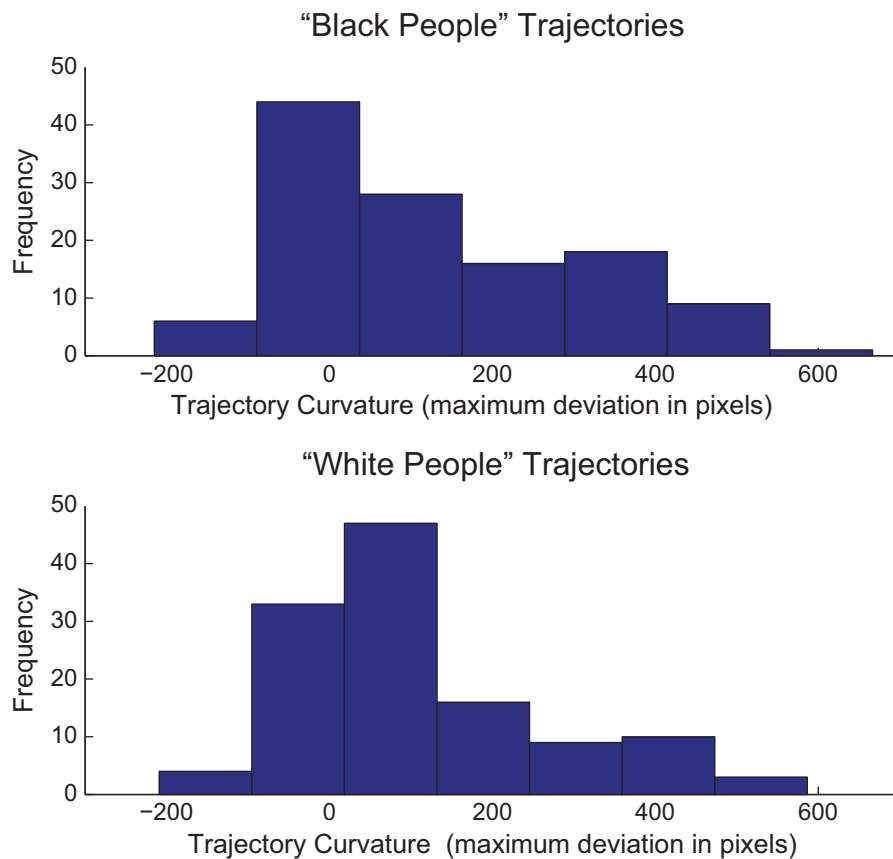


Fig. 2. Distributions of trajectory curvature in Experiment 1. The upper graph shows “Black people” trajectories, and the lower graph shows “White people” trajectories. Trajectory curvature was calculated as the maximum deviation between each trajectory and a straight line connecting its start and endpoint.

greater for the “Black people” trajectories than for the “White people” trajectories, $t(60) = 2.17$, $p_{\text{rep}} = .94$, $d = 0.29$. As a result, the mean distance traveled en route to the “like” response was also longer for “Black people” trajectories than for “White people” trajectories, $t(60) = 2.44$, $p_{\text{rep}} = .98$, $d = 0.32$. Responses to “Black people” and “White people” did not differ in total reaction time, $t(60) = 1.44$, $d = 0.18$.

In principle, the observed differential motor curvatures could have been generated by a stage-based sequence of decisional commands, rather than by continuous motor attraction to the “dislike” response. If motor execution required the complete prespecification of a unique target destination, rather than tracking of motor trajectory parameters that continuously evolved midflight, then a mean trajectory could look differentially curved because of the effect of averaging in replanned trajectories (i.e., some proportion of trials could have involved an initial motor command guiding movement directly toward “dislike” and then become aborted and replaced by a second motor command toward “like”). To accommodate the empirical mean trajectory, which initially moved upward rather than actually toward “dislike,” such an account would need to predict a bimodal distribution of curvatures that included some trajec-

tories that were very curved and others that were not curved. However, the distribution of trajectory curvatures shows no evidence of bimodality, as shown in Figure 2. The degree of bimodality can be quantified with a bimodality coefficient, which is capable of detecting bimodality in a mouse-tracking paradigm (Spivey et al., 2005). The bimodality statistic (b) is computed through the following formula (DeCarlo, 1997):

$$b = \frac{\text{skewness}^2 + 1}{\text{kurtosis} + \left[3 * \frac{(n-1)^2}{(n-2) * (n-3)} \right]},$$

where n is the number of observations in the distribution of interest. The standard cutoff for inferring bimodality in a distribution is $b > 0.55$. Neither the “Black people” nor the “White people” trajectories had distributions that met this cutoff, and in fact, the “Black people” trajectories formed a distribution of movement curvature that was closer to normal ($b = 0.24$, skewness = 0.613, kurtosis = 2.57) than the “White people” trajectories ($b = 0.301$, skewness = 0.98, kurtosis = 3.44).

We further analyzed these computer-mouse trajectories for evidence of nonlinear competitive dynamics, a signature of complex self-organizing systems. Velocity profiles were con-

structed by analyzing the temporal derivatives of motion toward the “like” response box along the x -coordinate. Because the mouse-movement’s starting location was equidistant between the two response boxes along the horizontal dimension, movement along the x -coordinate reflects relative confidence in deciding on one evaluation over the other. Our velocity predictions came from Usher and McClelland’s (2003) differential equations for modeling the dynamics of competition between mental representations:

$$\begin{aligned} dx_1 &= (I_1 - x_1 - \beta f_2)dt \\ dx_2 &= (I_2 - x_2 - \beta f_1)dt, \end{aligned}$$

where, in this case, x_1 and x_2 represent the activations of the mental representations for “like” and “dislike,” dx_1 and dx_2 represent the change in the activation of the two mental representations in a time step of size dt , I_1 and I_2 represent excitatory input to the representations from informational sources, βf_1 and βf_2 represent the inhibitory input from each mental representation to the other (*lateral inhibition*), and f_i (where $i = 1$ or 2) is equal to x_i if x_i is greater than zero.

According to these differential equations for competition dynamics, a strong evaluative competitor (dislike, x_2) sends intensified and prolonged lateral inhibition (βf_2) to the “like” evaluation (x_1). Thus, strong competition alters the velocity profile of the movement toward the evaluative attractor (dx_1/dt), reducing velocity toward the attractor early on in processing. However, as the more active alternative begins to win the competition, this lateral inhibition is gradually lifted, thus increasing velocity later in processing to produce greater acceleration. Therefore, strong competition predicts higher acceleration into the “like” response box (d^2x_1/dt^2) in normalized time. Moreover, this particular dynamic pattern (reduced early velocity and greater later acceleration) should lead to greater peak velocity, if jerk is minimized as the system achieves equivalent integral under the curve (where the integral represents net change in activation or location). Thus, dynamic conflict does not simply delay processing, but also changes its composition: Strong competition should lead to compressed, high-spiking derivative profiles toward the preferred interpretation, even in normalized time.

The observed mouse trajectories approached the “like” response boxes with precisely the temporal derivative profiles predicted by Usher and McClelland’s (2003) model of competition dynamics, as shown in Figure 3. The “Black people” trajectories had significantly greater maximum x -coordinate acceleration (shown as steeper velocity slope) into the “like” response box than the “White people” trajectories, $t(60) = 2.69$, $p_{\text{rep}} = .96$, $d = 0.41$. Moreover, the “Black people” trajectories had significantly greater peak velocity (shown as higher velocity curve peak), $t(60) = 2.65$, $p_{\text{rep}} = .95$, $d = 0.36$. These findings suggest that mental representations for both response alternatives, “like” and “dislike,” may be simulta-

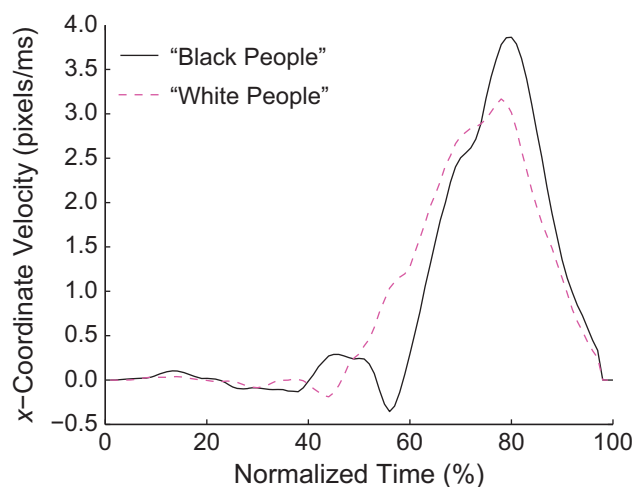


Fig. 3. Mean x -coordinate velocities for hand movements toward the evaluative response as a function of normalized time in Experiment 1. Velocities are shown separately for “Black people” and “White people” stimuli. The x -coordinate velocity reflects the evolution of confidence in the evaluative decision toward the “like” response.

neously active and competing over time, as in Usher and McClelland’s model.

The spatial-disorder analysis investigated the regularity of change in x -coordinate location over time. Our prediction about spatial disorder drew upon previous work on natural and physical self-organizing systems, which has established that increasingly conflicting constraints on a system’s state invoke dynamic state-space trajectories that show more disorder or irregularity in their pathways (Kauffman, 1993; see also Dale et al., 2007, and McKinstry, Dale, & Spivey, 2008). In the present study, a self-organizing framework predicted that the motor trajectories for “Black people” should have greater disorder than the trajectories for “White people,” even in the segments of the trajectories that had already committed to a “like” response. To investigate whether the “Black people” trajectories had more wiggles, blips, and other irregularities than the “White people” trajectories, we analyzed x -coordinate location over time, but only after the trajectory began moving in the positive x direction. A sigmoidal fit (which snugly fits curves that asymptotically approach “like” in an orderly and regular manner, as shown in Fig. 4) was then imposed on the obtained curve. The “Black people” trajectories had significantly greater deviation from the sigmoidal fit, as revealed in a significantly lower R^2 value, $t(60) = 2.29$, $p_{\text{rep}} = .92$, $d = 0.31$, which indicated disorderly variation around the x dimension in those trajectories.

The curvature results (Fig. 1, upper half) clearly demonstrate a greater motor attraction toward the “dislike” response option for “Black people” than for “White people,” indicating some initial prominence of this negative evaluation in responses that, a fraction of a second later, manifest themselves as positive attitude choices. It is worth noting that this difference in curvature emerged in the absence of a difference in total reaction time. The findings in the velocity and spatial-disorder analyses further

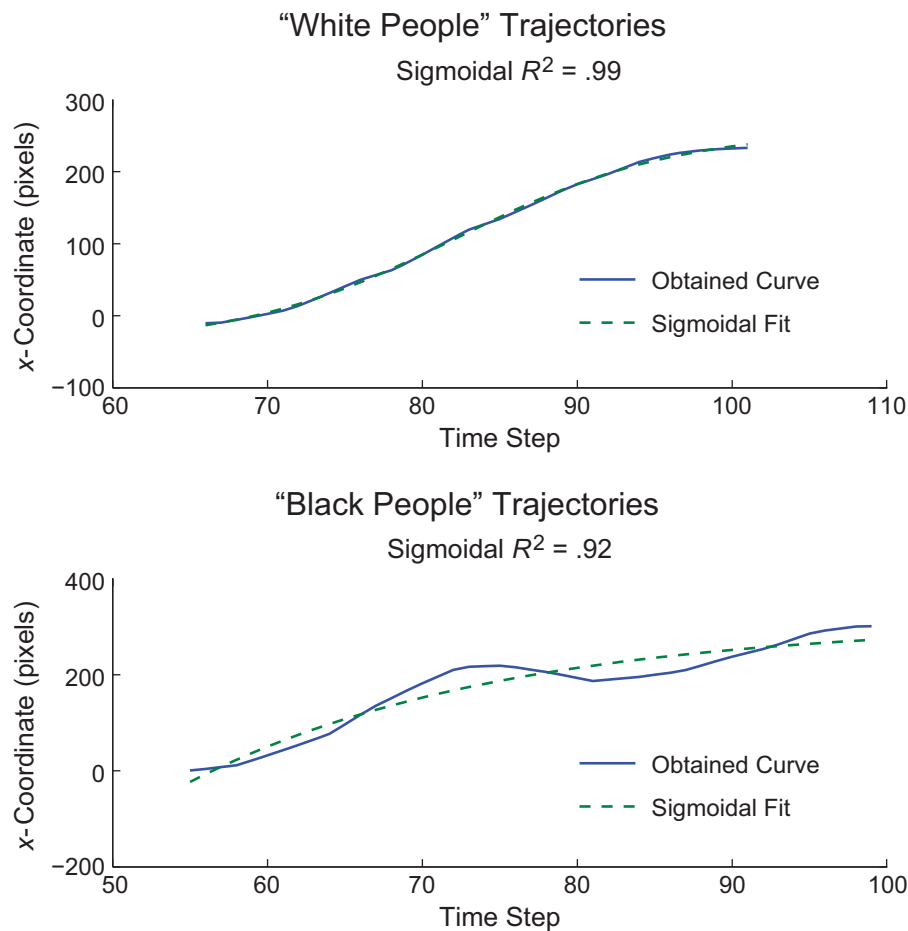


Fig. 4. Spatial disorder in hand trajectories' x -coordinates over time during evaluations of "Black people" and "White people" in Experiment 1. Each graph shows the observed x -coordinate location of a participant's hand as a function of time for a representative trial. Sigmoidal fit curves are also shown, and the R^2 values for the fit between the observed and sigmoidal curves are given. A lower R^2 value indicates more disorder.

suggest that this initial prominence of the negative evaluation may be part of a dynamic process of parallel competition between partially active positive and negative implicit evaluations; the winner of these evaluations becomes the explicit attitude choice.

EXPERIMENT 2

Because our claim is that multiple, partially active mental representations compete for the privilege of driving evaluative responses, imposing a set of response options that are not particularly competitive should change the motor dynamics. If the response box opposite the "like" box does not provide any semantic match to the content of the self-organizing evaluative response, then "White people" and "Black people" trajectories should lose their differential curvature. In particular, the targets "Black people" and "White people" should evoke much stronger support for interpretations as positive entities than as chemical elements.

Method

Sixty-six Cornell University undergraduates (40 female and 26 male) were asked to classify words (e.g., "ice cream," "sunshine," "boron") as something they liked ("like") or as the name of a chemical element ("chemical"). We analyzed data only from the 63 participants who consistently chose the "like" response for both "Black people" and "White people" on both repetitions of these trials, and who reported in a poststudy questionnaire that they were not forced into selecting "like" by the paradigm.

Results

According to statistical analyses on maximum deviation and distance traveled, the "Black people" and "White people" trajectories no longer differed in their curvature toward the competing response, as shown in Figure 1 (lower half), $t(62) = -0.10$, $p_{\text{rep}} = .16$, $d = -0.01$. Thus, the results of Experiment 1 are not attributable merely to responses to "Black people" involving a longer latency to settle on a positive evaluation, and

thereby drifting for longer in empty regions of movement space before curving toward the “like” response box. Rather, the “dislike” response option in Experiment 1 was actively pulling movement trajectories toward it, in a way that the “chemical” response option in Experiment 2 did not.

The trials with “Black people” and “White people” did not differ significantly in maximum acceleration, $t(62) = -1.06$, $p_{\text{rep}} = .64$, $d = 0.17$, or in peak velocity, $t(62) = -1.39$, $p_{\text{rep}} = .74$, $d = 0.22$. Likewise, the trials with “White people” and “Black people” did not differ significantly in spatial disorder, $t(62) = -0.13$, $p_{\text{rep}} = .19$, $d = -0.02$.

EXPERIMENT 3

Whereas we have framed our results with respect to explicit attitudes toward people of different races or ethnicities, the mouse-cursor response trajectories to “Black people” and “White people” in Experiment 1 may have diverged because of subtle confounds that do not refer to people at all. For example, perhaps these differences reflected different evaluations of the color terms “Black” and “White” that preceded the term “people.”

Method

Seventy-one Cornell University undergraduates (37 female and 34 male) were asked to classify stimuli as something they liked (“like”) or disliked (“dislike”). The crucial stimuli in this experiment were “African Americans” and “Caucasians.” We analyzed data only from the 64 participants who consistently chose the “like” response for both “African Americans” and “Caucasians” on both stimulus repetitions.

Results

The trajectories for “African Americans” curved significantly more toward the “dislike” response than the trajectories for “Caucasians,” $t(63) = 3.65$, $p_{\text{rep}} = .99$, $d = 0.56$ (Fig. 5). The motor trajectories evolved over time in accordance with the competitive velocity predictions, as reported in Experiment 1. The “African Americans” trajectories, compared with the “Caucasians” trajectories, had significantly greater maximum x -coordinate acceleration, $t(62) = 3.55$, $p_{\text{rep}} = .99$, $d = 0.47$. “African Americans” trajectories also obtained higher peak velocity than “Caucasians” trajectories, $t(62) = 4.54$, $p_{\text{rep}} = .99$, $d = 0.63$. Moreover, as we found for “Black people” trajectories in Experiment 1, the “African Americans” trajectories exhibited greater spatial disorder than the “Caucasians” trajectories, even after moving toward the “like” response, as indicated by significantly greater mean deviation from the sigmoidal fit, $t(62) = 2.49$, $p_{\text{rep}} = .94$, $d = 0.44$. In tandem, these results demonstrate that the same general constellation of findings was observed with the labels “African Americans” and “Caucasians” as was observed with the labels “Black people” and “White people.”

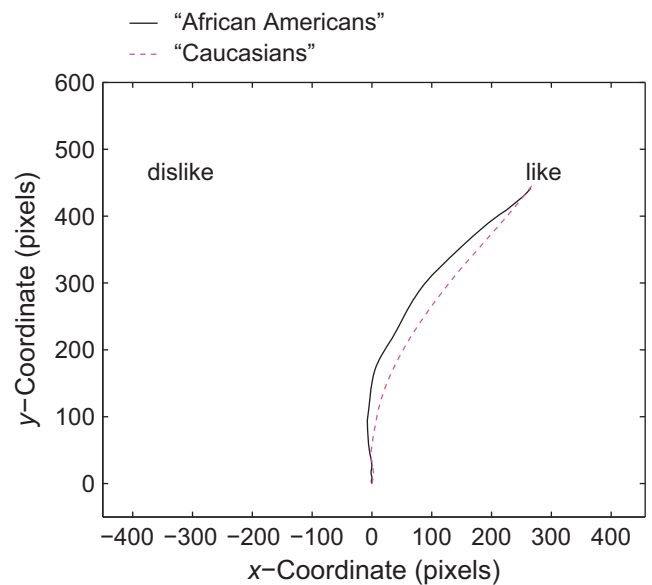


Fig. 5. Mean mouse-movement trajectories toward the evaluative response for the “African Americans” and “Caucasians” stimuli in Experiment 3. Because the “like” response box could appear in either the right or the left corner of the screen, both rightward and leftward trajectories were obtained. For this graph, leftward trajectories were flipped before being averaged with rightward trajectories.

GENERAL DISCUSSION

People’s hand-movement trajectories for explicitly evaluating “Black people” and “White people” were distinct as measured by three properties of movement dynamics: shape, time, and order. These findings suggest that explicit attitudes evolve through continuous temporal dynamics during real-time mental processing, with graded motor curvature revealing the influence of tendencies toward dislike. There was no evidence for cleanly separated (i.e., discrete, rather than continuous) explicit decisions, in which an initial response was executed solely toward the “dislike” response box and then a corrective response was executed midflight toward the “like” response box. Rather, the results suggest that a dynamic competition process may be what allows a single explicit attitude choice to emerge from multiple, potentially conflicting evaluative influences (e.g., Busemeyer & Townsend, 1993; Usher & McClelland, 2003). Thus, rather than switching from one singular (implicit) decision to a different singular (explicit) decision, the mind may host a continuously evolving blend of (implicit) evaluative decisions from which the eventual (explicit) behavioral choice emerges.

REFERENCES

- Bargh, J.A., & Chartrand, T.L. (1999). The unbearable automaticity of being. *American Psychologist*, *54*, 462–479.
- Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, *19*, 442–577.

- Busemeyer, J.R., & Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–462.
- Cisek, P., & Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, *45*, 801–814.
- Conrey, F.R., & Smith, E.R. (2007). Attitude representation: Attitudes as patterns in a distributed, connectionist representational system. *Social Cognition*, *25*, 718–735.
- Dale, R., Kehoe, C., & Spivey, M.J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Memory & Cognition*, *35*, 15–28.
- DeCarlo, L.T. (1997). On the meaning and use of kurtosis. *Psychological Methods*, *2*, 292–307.
- Devine, P.G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.
- Erlhagen, W., & Schoner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*, 545–572.
- Farmer, T.A., Anderson, S.E., & Spivey, M.J. (2007). Gradiency and visual context in syntactic garden-paths. *Journal of Memory and Language*, *57*, 570–595.
- Fazio, R.H., Jackson, J.R., Dunton, B.C., & Williams, C.J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes. *Journal of Personality and Social Psychology*, *69*, 1013–1027.
- Freeman, J.B., Ambady, N., Rule, N.O., & Johnson, K.L. (2008). Will a category cue attract you? Motor output reveals dynamic competition across person construal. *Journal of Experimental Psychology: General*, *137*, 673–690.
- Henis, E.A., & Flash, T. (1995). Mechanisms underlying the generation of averaged modified trajectories. *Biological Cybernetics*, *42*, 407–419.
- Judd, C.M., Drake, R.A., Downing, J.W., & Krosnick, J.A. (1991). Some dynamic properties of attitude structures: Context-induced response facilitation and polarization. *Journal of Personality and Social Psychology*, *60*, 193–202.
- Kauffman, S.A. (1993). *The origins of order*. Oxford, England: Oxford University Press.
- Kelso, J.A.S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Boston: MIT Press.
- Lamme, V.A.F., & Roelfsema, P.R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*, 571–579.
- Lapish, C.C., Durstewitz, D., Chandler, L.J., & Seamans, J.K. (2008). Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proceedings of the National Academy of Sciences, USA*, *105*, 11963–11968.
- McKinstry, C., Dale, R., & Spivey, M.J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science*, *19*, 22–24.
- McMurray, B., Tanenhaus, M.K., Aslin, R.N., & Spivey, M.J. (2003). Probabilistic constraint satisfaction at the lexical/phonetic interface: Evidence for gradient effects of within-category VOT on lexical access. *Journal of Psycholinguistic Research*, *32*, 77–97.
- O'Reilly, R.C. (1998). Six principles for biologically based computational models of cortical cognition. *Trends in Cognitive Sciences*, *2*, 455–462.
- Roe, R.M., Busemeyer, J.R., & Townsend, J.T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*, 370–392.
- Rogers, T.T., & McClelland, J.L. (2004). *Semantic cognition: A parallel distributed processing approach*. Boston: Bradford Books.
- Smith, E.R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, *4*, 108–132.
- Spivey, M.J. (2007). *The continuity of mind*. Oxford, England: Oxford University Press.
- Spivey, M.J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences, USA*, *102*, 10393–10398.
- Usher, M., & McClelland, J.L. (2003). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.
- Van Orden, G.C., Holden, J.C., & Turvey, M.T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, *132*, 331–350.
- Wilson, T.D., Lindsey, S., & Schooler, T.Y. (2000). A model of dual attitudes. *Psychological Review*, *107*, 101–126.
- Wittenbrink, B., Judd, C.M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, *72*, 262–274.

(RECEIVED 8/28/08; REVISION ACCEPTED 4/24/09)