Mapping varied mental representations: The case of representing illegalized immigrants

Joel E. Martinez^{1,2} & Alexander Todorov³

¹ Department of Psychology, Harvard University, Cambridge, MA, USA

² Data Science Initiative, Harvard University, Cambridge, MA, USA

³ Booth School of Business, University of Chicago, Chicago, IL

Note: This manuscript has not been peer-reviewed

Abstract

How do people represent social categories? To answer this question, aggregate mental representations would typically be estimated at the sample level or within theory-derived individual differences (e.g., conservatives vs. liberals). However, these approaches can fail to capture the heterogeneous contours of collective representations. We propose a fully data-driven approach that first clusters mental representations by similarity and then identifies which measures best differentiate between clusters. We apply this approach to understand mental representations of illegalized immigrants. Representations were estimated using face-based reverse correlation from border states, Texas and California ($N = 1002$), along with various measures thought to influence perceptions of immigrants: attitudes, demographics, ideologies, social geographic characteristics, and a manipulation of labels (i.e., “undocumented” vs “illegal” immigrant). Comparative analyses revealed how the aggregate approach hid representational clusters that differed on visualized facial phenotype and affective expressions. Some clustered representations were highly shared while others were differentiated by more idiosyncratic characteristics: age and local population size perceptions. Data-driven approaches therefore offer a useful tool for identifying unexpected sources of clustered representations by centering variation in investigations of mental representations.

Keywords: Disaggregation, Illegalized immigrants, Face Representations, Reverse Correlation, Machine Learning

Despite statistical averages operating as the analytic workhorse of social psychological science, recent scholarship has questioned the utility of over-relying on them (Martinez et al., 2020; Martinez & Paluck, 2020). The idea that group-level data does not necessarily faithfully reflect individual-level data has long been identified (Estes, 1956), yet the implications of this problem do not explicitly guide current theory-building and analytic practices (Fisher et al., 2018). For instance, judgment studies show how common assessments of agreement based on sample averages can hide disagreements arising from stable idiosyncratic interactions between individuals and stimuli (Martinez et al., 2020; Martinez & Paluck, 2020). The problem with averages is also present when trying to measure mental representations: computing representations at the individual-level can reveal patterns that are hidden when associations are instead computed from aggregated data (Martinez, Feldman, et al., 2021).

Previous work used face-based reverse correlation to capture the mental representation of various social categories (Brinkman et al., 2017; Dotsch et al., 2008; Dotsch & Todorov, 2012; Hinzman & Maddox, 2017; Imhoff et al., 2011; Krosch & Amodio, 2014; Lei & Bodenhausen, 2017). Main inferences from these studies relied on sample-averaged classification images that aggregate across individual-level images. Given the above critiques, an average classification image may not reflect distributions of representations across individuals. Aggregate misrepresentation has been identified as an issue for the inferences made from ratings given to average vs. individual classification images (Cone et al., 2020), but it even strikes at the conceptual validity of classification images themselves. For example, disaggregating the larger category of “undocumented immigrant” (Martinez, Oh, et al., 2021a) can change representations of illegalized immigrants (e.g., illegality is not always encoded as darker skin tones) (Martinez, Oh, et al., 2021b). However, partitioning the target category into subcategories is only one way of

disaggregating. Another way is to center representational variation as the target of investigation. In other words, what various understandings exist for a target category? This study explores the viability of this approach in the hopes of developing a more comprehensive approach to characterizing mental representations of social categories. We focus on the category of illegalized immigrants whose perceptions have been understudied despite its urgent political salience. Moreover, the evidence of a consistent dark-skinned threat representation of illegalized immigrants at the sample-average level that can shift depending on what additional information is presented suggests there exists important yet unexamined variation in understandings.

The importance and origins of representational variation

It is thought that the power of representations depends on consensus (Haslam et al., 1997; Stangor & Schaller, 2000). The more widely shared a belief or image is, the more that people can draw upon the same information and coordinate collective actions (Tappin et al., 2017). However, a focus on consensus at the expense of disagreement can blind researchers to the large influence of “noise” (e.g., stable patterns of idiosyncratic beliefs or experience) on judgment (Kahneman et al., 2021; Martinez et al., 2020). There is generally a larger share of idiosyncratic than shared contributions to most social evaluations (Hönekopp, 2006; Martinez et al., 2020), thus it is to be expected that meaningful variation also exists in representations of illegalized immigrants.

One initial source of variation is at the conceptual level. Stereotypes of immigrants have been shown to be ambivalent across dimensions (Reyna et al., 2013; Savaş et al., 2021), suggesting that no singular evaluation exists. Another source that we focus on in this study is from differences between people that are theorized to influence perceptions of illegalized

immigrants. These differences originate from attitudes, demographics, ideologies, social geography, language, and more. For attitudes, scales assessing immigration relevant evaluations have captured variation in attitudes towards illegalized immigrants (Ommundsen et al., 2014). For demographics, perceptions of immigrants have been influenced by education level (Abrajano & Hajnal, 2015), adherence to a Christianity infused with nationalism (Davis & Perry, 2020), and racial classification (Abrajano & Hajnal, 2015; Treitler, 2013). For ideologies, variation has also been identified from political orientations (Chavez, 2013), political party affiliations (Sides et al., 2018), and voting for or against specific nativist political actors like Donald Trump (Kunst et al., 2019; Sati, 2020).

For social geography, recent scholarship has called for attention to the importance of mapping geographic variation of psychological phenomena (Rentfrow & Jokela, 2016). This is essential for the psychological study of migration for its intimate ties to material and social space and accompanying demographic politics (Hopkins & Dixon, 2006; Rodríguez-Muñiz, 2021). Examples include the way that being near material and social boundaries (e.g., borders, social contact or segregation) shapes identity-based conflict (Enos, 2017), how perceptions that track the increasing size of demographic categories at local or national levels is associated with increased antipathy and hate crimes (Cikara et al., 2020; Craig et al., 2018), and how local or state legislation can shape immigrant sentiment for their constituents (Huo et al., 2018a, 2018b).

For language, variation in representations could arise from the various ways that illegalized immigrants get labeled in social discourse. A complex history of political criminalization, media agendas, and activist organizing have led to many names for illegalized immigrants (Ackerman, 2013, 2014; Plascencia, 2009; Rosa, 2019). Two of the more common labels are “illegal” and “undocumented”. These labels are thought to lead to different evaluative

and punitive consequences with illegal being worse presumably because it mentally conjures criminals while undocumented is seen as a more benign option (Ommundsen et al., 2014; Rucker et al., 2019). However, research focused on news media instead finds that labels have less of an influence than other types of frames that highlight narratives, negativity, or policy (Alamillo et al., 2019; Haynes et al., 2016). Thus, while labels might be important, it remains unknown what kind of immigrants they mentally conjure.

Theoretically, all these variables should describe meaningful variation in reverse correlation classification images.

Approaches to assessing representational variation

Theory-driven approach. While reverse correlation uses a data-driven process in the creation of classification face images, examination of representational variation has been typically guided by theory. Researchers decide on theory-derived individual difference variables by which to cleave the sample to examine representational differences (e.g., Dotsch et al., 2008; Lei & Bodenhausen, 2017; Young et al., 2014). For instance, one could decide that political orientation is an important source of variation based on past research and create an average classification image for liberal participants to compare to the average image for conservative participants. Any differences between the two images could elicit differential evaluations from naïve raters, which can then be interpreted as evidence for the importance of political orientation.

While the theory-driven approach can be useful, in practice it rests on untenable assumptions. First, analyzing psychology using a categorical lens often fails to capture the social contours of collective beliefs given that within-category disagreements can often be larger than between-category disagreements (Cikara, 2021; Martinez & Paluck, 2020). Second, it

essentializes people by assuming that individual difference variables can straightforwardly summarize people and their beliefs as if they experienced a static ontological existence (Birhane, 2021; Tucker, 2012). Third, it assumes that psychological phenomena are influenced by multiple independent and large main effects where the goal is to identify each one, even though the simultaneous contribution of many powerful main effects is highly improbable (Tosh et al., 2020). Representational variation is instead more likely to arise from complex interactions between dynamic forces and dynamic people. The theory-driven approach amounts to a static a priori guess about where representational variation resides, which can impede examination of other potentially important variables or complex interactions.

Proposed data-driven approach. Our approach advocates for a data-driven examination of representational variation which focuses on inter-individual similarity. This idea is not new. Since the introduction of correlating across people as a statistical technique for personality science (Stephenson, 1935), it has been adopted to identify how complex and socially contested concepts may be understood in varied ways across clusters of participants (Boutyline, 2017; Watts & Stenner, 2005). What *is* new is the development of advanced machine learning techniques that lessen previous computational barriers and that loosen restrictive assumptions imposed by confirmatory methods (Breiman, 2001b; Jones & Linder, 2015). The basic proposal is that we might learn more about how a social target is mentally represented if we first identify how representations are clustered by similarity and then describe the extent to which theory-derived individual differences track these clusters. This form of exploratory data-driven approach can help reveal unexpected features that can inform theory-building (Adolphs et al., 2016) and can remain dynamic as additional data may shift the clusters.

The analytic pipeline developed for this study occurs in four steps followed by a validation procedure (**Figure 1**). First, compute all the individual-level representations (i.e., the noise pattern before it is overlaid on the base face). Second, correlate these individual representations with each other to create a participant x participant similarity matrix. Third, use a clustering algorithm to identify structure in the similarity matrix – are there clusters of people who visualized similar faces? The last step is to compute an average classification image for each cluster to understand the content of each cluster’s average representation. Despite the data-driven spirit of this pipeline, many of the steps do require researchers to make multiple important analytic decisions (see method section).

This approach attempts to strike a balance between competing conceptualizations of the relationship between individual and average classification images. One perspective sees the individual images as noisy and the averaging as a process that reveals meaningful signal (Brinkman et al., 2017), while the other sees the individual images as a better reflection of existing representational variation while averages neglect to account for that variation (Cone et al., 2020). The former is limited by the issues with sample averages while the latter is limited by costs associated with analyzing a large collection of individual images and that some aggregation of information across people is eventually needed since shared beliefs cannot be quantified solely from individual representations. In our proposal, averages created from first clustering people with similar representations should be more representative of the averaged participants than previously used averaging approaches (e.g., full sample averages, sample-bootstrapped averages, or variable-derived averages).

A potential issue with our approach is that cluster algorithms may find clusters even within random noise, so an additional validation procedure can help identify whether meaningful

clusters have been found. We focus on two psychologically informed validations: 1) how the average cluster images are evaluated along various traits to help characterize the content of those representations and 2) whether theory-derived individual differences track differences in cluster membership as would be expected. Meaningful clusters can be said to be found to the extent that the individual differences match representational content in a theoretically consistent manner. To map variables to cluster memberships, we take advantage of fairly recent advances in machine learning: decision trees (Hothorn et al., 2006) and random forests which are an aggregation and summary of many decision trees (Breiman, 2001a). These combined approaches are promising for this application for the following reasons: decision trees can uncover complex interactions between variables on cluster membership in line with the idea that psychological phenomena are shaped by dynamic influences; random forests provide a robust way to cross-validate the important variables found in the decision tree; and innovations in interpretable machine learning have opened up the “black box” of random forest models to help characterize the relationships between the variables and cluster membership that the model learns (Molnar, 2020).

Figure 1. Analytic pipeline and validation procedures. The analytic pipeline consists of four steps: **a.** collect all the participants' individual noise representations, **b.** correlate the individual representations to create a participant x participant similarity matrix, **c.** find clusters using a hierarchical clustering algorithm, and **d.** compute an average face for each cluster. The validation pipeline consists of two main analyses to validate the choice of clusters: **e.** investigate whether evaluations of the average faces differ in ways that conceptually match the associated individual differences for each cluster and **f.** find individual difference measures that can predict cluster membership (top panel is a decision tree, bottom panel is variable-cluster mappings from a random forest).

Current project

The goal of this study is to arrive at a more comprehensive understanding of how illegalized immigrants are represented. We attempt to do this by exploring representational variation using a data-driven approach, which should reveal various understandings that underlie the category of “undocumented immigrant”. First, we showcase the problems that arise when attempting to understand this question using a theory-driven aggregate approach. Then, we apply the proposed data-driven pipeline to the classification images, showcasing its advantages. Validation of the content of the resulting representations occurred through dangerous, American, and ethnoracial classification ratings, shown to be important for representations of immigrants (Martinez, Oh, et al., 2021a, 2021b). We additionally included dominance, which is considered a fundamental dimension of face perception (Todorov & Oosterhof, 2011).

Validation of the clusters used the individual difference variables discussed in the introduction. While the attitude, demographic, and ideological variables can easily be measured within the task, assessing contributions from geographic and label variables requires explicitly accounting for them in the study design. As such, we focused on collecting representations from two states at the Mexico-U.S. border: California and Texas. Within each state, we asked half of the participants to visualize “undocumented immigrants” and the other half “illegal immigrants”.

These design specifics were chosen carefully for they are related. Both Texas and California are border states, have large overall populations, and each have a population of over 400,000 illegalized immigrants (Pew Research Center, 2019b). However, they differ in the dominant immigrant sentiment of their political institutions and immigrant discourse. Texas had no sanctuary cities at the time of this study and “illegal immigrant” was a more common search

term on Google from its residents, while California had many sanctuary cities and its residents tended to search more for “undocumented immigrant” (Griffith, 2021) (see **Supplementary Figure S1**). This connection between place and language even occurs at the level of state legislation where California bills tend to use “undocumented” labeling while conservative states like Arizona mainly use “illegal” language in their legislation (Filindra & Kovács, 2012). Given that institutional signals of inclusion or exclusion can serve as norm cues (Tankard & Paluck, 2016), state-level political culture may become symbolically ingrained as a normative community value or accessible schema that shapes its residents’ representations of illegalized immigrants (Zaller, 1992). However, there is also considerable within-state variation in immigrant policy, sentiment, presence, and contact (Varsanyi et al., 2012), which could influence local demographic perceptions independent of state-level factors (Enos, 2017). These states are therefore useful cases for attempting to capture representational variation from labels and geographic differences at various geographic levels.

Methods

The data and analysis scripts are archived on the Open Science Framework: https://osf.io/z32fn/?view_only=0c4ef10f6bb643c58b257a3df9d6ad5f. The project was approved by the Princeton University IRB #7301. This study was not preregistered. We report how we determined our sample size, all manipulations, and all measures in the study.

Image generation task

Participants

Participants were sampled from Amazon Mechanical Turk in December 2019, with settings that restricted access to the task solely to participants who lived in California (N = 500)

or Texas (N = 502). Half of the participants in each state either visualized “undocumented” or “illegal”. Therefore, we collected ~250 participants per state x label condition as that is the sample size at which correlations stabilize (Schönbrodt & Perugini, 2013). The Texas illegal condition had 2 extra participants (N = 252) that were collected as pilot-testing participants. The average age across the full sample was 38 (SD = 12) and the gender distribution was slightly skewed towards women (53.9%). The ethnoracial distribution was as follows: white (57.2%), Latinx (16%), black or African (11%), East Asian (9.5%), South Asian-Indian (2.7%), race or ethnicity was not listed (2.4%), Hawaiian-Pacific Islander (.9%), Indigenous-Native American (.3%), Middle Eastern (.1%). In terms of citizenship status, 92.6% were born in the U.S. and 97.8% were U.S. citizens. A geospatial map of where in each state participants resided can be seen in **Supplementary Figure S2**.

Participant measures

The following individual difference measures were used to help characterize representational variation. The full descriptive table for participant characteristics can be found in **Supplementary Table S1**. Demographic variables included *age*, *gender*, *race*, *religious affiliation*, and *education level* (Master’s or Doctoral degree, Bachelor’s degree or similar, Professional qualification, High school or Baccalaureate or A-levels, I did not complete secondary/high school).

Attitudes were measured using the *Illegal Immigrant Scale* (IIS) (Ommundsen et al., 2014). The scale included 20 statements whereby participants responded with how strongly they agreed with each one (agree strongly, agree, undecided, disagree, disagree strongly). The original scale used the “illegal” label, so we changed the wording of the statements to “unauthorized

immigrants” so that the scale did not differentially match or mismatch across label conditions (“illegal” vs. “undocumented”). Statements ranged from economic (“Unauthorized immigrants should not receive food stamps”), to political (“Unauthorized immigrants who give birth to children in the United States should be made citizens”), to social (“Unauthorized immigrants should not be discriminated against”). Scores were computed by averaging the responses after reverse coding the relevant items. Scores ranged from 1 (positive) to 5 (negative) attitudes ($M = 3.00$, $SD = .30$).

There were four ideological variables: *political party affiliation* (Democrat, Independent, Republican, Green, Libertarian, Unaffiliated/Not political), *presidential vote in 2016* (Hillary Clinton & Tim Kaine: Democratic Party, Donald Trump & Mike Pence: Republican Party, Gary Johnson & Bill Weld: Libertarian Party, Jill Stein & Ajamu Baraka: Green party, other candidate, could not vote, chose not to vote), *support for building a wall* along the U.S.-Mexico border (agree strongly, agree, undecided, disagree, disagree strongly), and *political orientation*. Political orientation was measured using the Social and Economic Conservatism Scale (SECS) (Everett, 2013). Participants were presented with 12 political issues (e.g., “abortion”, “patriotism”, “traditional marriage”) and asked to designate on a thermometer scale from 1 (negative) to 100 (positive) how they feel about it. Scores were computed by reverse coding the relevant items and averaging the responses. Higher numbers indicate more conservative orientations ($M = 59.5$, $SD = 17.4$).

Social geographic variables included: *state lived in*, *how long they lived in that state* (Less than 1 month, Less than 1 year, Over 1 year, Over 10 years), estimates of *frequent contact with immigrants*, subjective *estimates of the size of the surrounding illegalized immigrant population*. While data collection was restricted to only Texas and California, 14 participants completed the

task from another state. We kept them in the analyses and labeled their state as “Other”. Most of the sample lived in their state for over 1 year (8.7%) or over 10 years (90%). Since immigrant status cannot always be readily perceived in fleeting daily social interactions, immigrant contact was assessed by having participants write down the initials of up to five immigrants they are in frequent contact with. Contact scores were calculated by the number of names they wrote down. Text input was processed to exclude non-name responses (e.g., statements that they do not know any immigrants). Population size estimates were calculated at three levels: their neighborhood, their city, and their state. The question asked was “Of all the people living in the location above, what percentage do you think are unauthorized immigrants?”. The order of the three locations was randomized and each location was estimated separately. The options were displayed as percentages from 0% to 50% in increments of 5 with the final option being “more than 50%”, however values in between could be chosen as the value of the slider was displayed above the scale and updated as the slider was moved ($M_{\text{neighborhood}} = 12.5\%$, Median = 5%, SD = 14.5; $M_{\text{city}} = 19\%$, Median = 15%, SD = 15.2; $M_{\text{state}} = 24.6\%$, Median = 21%, SD = 15.1). To put these estimates into perspective, the national illegalized population size in 2016 was 3.3% with Las Vegas having the largest population (8.2%) and Philadelphia having the lowest (2.6%) (Pew Research Center, 2019a). Keeping in mind the difficulty of enumerating this population and the uncertainty surrounding these statistics, the size of the surrounding illegalized immigrant populations tended to be overestimated, particularly for larger geographic zones.

A correlation matrix that describes the relationships between all continuous participant measures can be seen in **Supplementary Figure S3**. Most variables were weakly correlated except the three population size estimation measures.

Procedure

After providing consent, participants completed a standard reverse correlation task. They were asked to visualize either “illegal immigrants” or “undocumented immigrants” followed by 300 forced choice trials where they decided which of two presented face stimuli looked more like the visualized target. They pressed the “0” key for the right face and the “1” key for the left face. The stimuli were presented at 9 cm x 9 cm and side-by-side on screen. Trials were self-paced and a 1 second fixation cross appeared after each trial. After the task, participants completed the individual difference questionnaires.

Stimulus generation

The base face was a morph of all the male faces in the London Face Database (DeBruine & Jones, 2017). Overall, 770 stimuli were generated by repeatedly superimposing random sinusoid noise over the base face using the rcicr package (Dotsch, 2017). Out of 770 generated face stimulus pairs, only the first 300 were used for this study.

Classification image generation

Individual and average classification images were generated using the rcicr package. To generate individual classification images, the noise pattern of the chosen face stimuli was averaged. For the average classification images, the noise patterns of the individual-level classification images were averaged across the sample, across participants within different levels of variable, or across participants within a cluster.

Analyses

Analytic pipeline

Representational similarity. The first necessary result is a matrix that characterizes inter-individual representational similarity. This matrix was computed by first isolating the pixel noise pattern from each participants' classification image, estimating their idiosyncratic mental representation (**Figure 1a**). The base face is not used in this analysis, diverging from related analyses (Hong & Ratner, 2020), since it is typically added to the noise pattern to create all classification images and therefore adds no additional information and would inflate similarity between individuals. The noise pattern itself contains irrelevant pixel information that can also inflate similarity (e.g., the background around the face), so these pixels were removed by using an oval mask that isolated the area of the face. The masked noise pattern was then turned into a vector. All participants' masked noise vectors were then correlated and placed into a participant x participant similarity matrix (**Figure 1b**).

Cluster analysis. While there are many clustering methods, we decided to use hierarchical clustering for three reasons: it does not require an a priori decision about the number of clusters (various cluster numbers can be explored afterwards), the process is deterministic and thus arrives at the same solution, and the process of clustering is mapped on a dendrogram which provides a temporal visualization of the clustering decisions (**Figure 1c**). Despite these advantages, there are many algorithmic options and decisions to be made when running a hierarchical clustering algorithm. We explored a few options to find the most reasonable settings for our application.

The first decision is whether the hierarchical clustering algorithm should start with smaller clusters and iteratively build up to larger ones by finding the most similar clusters (i.e., agglomerative), or start with one giant cluster and break down to smaller clusters by finding the most distinct clusters (i.e., divisive). Whereas the former better identifies smaller clusters, the

latter better identifies larger clusters. Since the goal was to identify clusters of similar representations that may not necessarily be large, the choice for this analysis was agglomerative clustering. The next important decision is which metric (i.e., “linkage”) should be used to measure similarity between observations when iteratively deciding which observations to cluster together. We tested which of three algorithms led to a solution with clearly distinct clusters (**Supplementary Figure S4**). The participant similarity matrix was transformed into a Euclidean distance matrix and fed into a hierarchical algorithm with one of three linkage functions: average linkage, complete linkage, or Ward’s method. The average linkage measures the distance between the average similarity within two clusters. The complete linkage measures the distance between the furthest observations from each cluster. Ward’s method instead minimizes the total within-cluster variation. To identify which would be ideal for this study, we compared two metrics: the agglomerative coefficient (AC) and the sample sizes of the resulting clusters. The AC is a measure of how much cluster structure is found with a range from 0 (weak) to 1 (strong). A good cluster solution would show strong clustering structure. The ACs were as follows: average linkage AC = .56, complete linkage AC = .77, Ward AC = .96. A secondary metric we used was sample size (**Supplementary Figure S5**): the goal was to find an algorithm where the minimum sample size in any of the clusters would not be too low to obtain a reliable average image. Ward’s method was the only algorithm that stayed consistently above 20 minimum participants and provided a strong cluster structure. While we focus on Ward’s method, we also computed the same pipeline (e.g., clusters and average representations) for the complete linkage method with the goal of comparing the two. Therefore, the faces from the complete linkage pipeline are part of the rating task below. However, since clusters were not as differentiated using complete linkage, those results are not analyzed further.

The last critical choice is to decide how many clusters exist or are meaningful to examine. To aid in comparing cluster number solutions, we used two metrics: total within-cluster sum of squares (TWSS) and minimum cluster sample size (**Supplementary Figure S5 and S6**). Clusters with smaller TWSS have more similar observations, which is ideal and exactly what Ward's method optimizes. The TWSS and minimum sample size suggested that between 5 to 15 clusters would be reasonable solutions. Since there is no accurate a priori solution to compare to, we instead ran the same analytic pipeline for a 5, 10, and 15 cluster solution to see what insights could be gleaned from more or less cluster partitions. For brevity, we mainly report the 5 cluster solution, but the results from the 10 cluster (**Supplementary Figure S7, S8, S9**) and 15 cluster (**Supplementary Figure S10, S11, S12**) analyses can be found in the supplementary materials.

Validation

Cluster validation occurred by characterizing the content of each cluster's average representation and examining which individual difference measures tracked cluster membership. A common validation metric for machine learning models is predictive accuracy on held out samples, however, for these data, hold out accuracy was not good. One reason may be that when clusters differ greatly in size the algorithm learns to optimize accuracy by guessing only the largest cluster. This means predictive accuracy in the case of imbalanced clusters reflects information about the distribution of cluster sizes instead of validity (Akosa, 2017). While there are ways to minimize this issue, predictive accuracy may also not be appropriate for this application. When a decision space is unknown, such as how many schemas of illegalized immigrants actually exist, it may make more sense to treat the problem as one of probabilistic prediction (i.e., the probability of an observation belonging to each cluster) than classification

(i.e., a forced decision that each observation absolutely belongs to a cluster) (Harrell, 2017). We use the probabilistic prediction approach for this study.

Cluster content. For each identified cluster, an average classification image was computed (**Figure 1d**) and rated on various traits (**Figure 1f**), see image rating task details below.

Cluster membership. Cluster membership was characterized using two approaches: decision trees and random forests. The decision trees used in this study were conditional inference trees with the partykit package version 1.2.11 (Hothorn et al., 2006). The algorithm conducts an exhaustive search through all variables and their constitutive levels to find the partitions that better distinguish cluster membership. At each step, it first conducts a general null hypothesis test that any of the variables predict cluster membership – the one with the lowest p value gets placed as the initial partition variable. Then the algorithm searches through all the levels of the chosen variable and conducts more null hypothesis tests searching for the lowest p value to create a split, suggesting the resulting probabilistic cluster memberships are more differentiated across the partition. The search continues through variables and variable levels until there are no more significant splits. The resulting decision tree provides a summary map of the search. The recursive nature of the search allows for an easy exploration and identification of complex interactions between variables. The model options were as follows: tests were corrected using Bonferroni correction, the alpha level was .10 to facilitate exploration, the split statistic was “quadratic”.

The descriptive advantage of decision trees’ exhaustive search is counteracted by the high possibility of overfitting. Random forests can mitigate this issue by combining multiple decision trees to provide a more robust idea of which variables are important and how (Breiman, 2001a).

The algorithm uses similar procedures for creating decision trees, however randomness and bootstrapping are embedded into the process. First, a random subset of the variables is chosen to search through, this creates more varied trees because the most important variables are not always the initial split. Second, a random subset of the data is held out of the search to also ensure tree variation and assess out-of-bag predictive accuracy. Third, for imbalanced clusters, there are procedures for weighting observations or under- and over- sampling them to try to create better cluster balance (Chen, 2004). Lastly, techniques for exploring variable interaction effects in random forests are less developed (Molnar, 2020), so we focus on confirming the importance of single variables.

Random forests were computed with the ranger package version 0.12.1 (Wright & Ziegler, 2017). To minimize cluster imbalance, observations were probability weighted (e.g., observations in larger clusters were proportionally less likely to be chosen for the subset training samples, as smaller clusters more likely to be chosen). Since random forests combine across trees, reporting single trees would not be representative of what the model ultimately learned. Instead, we report variable importance metrics, specifically a permutation p value that repeatedly shuffles the cluster labels and examines the predictive accuracy of each predictor to create a null distribution (Altmann et al., 2010). Actual predictive accuracies on the non-shuffled cluster variable are compared to the null distribution to obtain a p value. Variable importance only provides an idea of which variables are useful predictors; we also need to know how they map unto cluster membership. For understanding the variable-cluster mappings that the models learned, we report partial dependence plots (PDP) (Apley & Zhu, 2020; Molnar, 2020). These values are analogous to estimated marginal means in regression models - they plot how the average predicted probability of membership in each cluster varies across levels of a variable.

However, there exists no statistical inference techniques for these estimates, so they remain descriptive in this study.

As initial proof that our measures are valid and that random forest models can capture meaningful relationships, we tested which variables were related to support for building the Mexico-U.S. wall. The 2016 presidential vote, political orientation, and political party affiliation were identified as important variables for predicting support for building the wall in the expected directions (**Supplementary Figures S13 and S14**). These results build confidence that patterns found in the main analyses are meaningful.

Image rating task

Participants

There were four sets of raters for this portion of the study collected in May 2020. Previous simulation power analyses suggest 20 participants provided enough power for ratings (Martinez, Oh, et al., 2021a, 2021b). However, data from those studies also showed that test-retest reliability of ratings tends to be lower for classification images. To compensate for likely reliability exclusions (i.e., when test-retest reliability is negative), we collected more participants. The differences between the total number and the final number in the following sample sizes are due to reliability exclusions. Participants rated classification images on dangerousness ($N_{\text{total}} = 60$, $N_{\text{final}} = 46$, $M_{\text{age}} = 38.3$, $SD = 10.4$, 14 women, 32 men), Americanness ($N_{\text{total}} = 60$, $N_{\text{final}} = 46$, $M_{\text{age}} = 40.2$, $SD = 12.9$, 22 women, 24 men), dominance ($N_{\text{total}} = 59$, $N_{\text{final}} = 44$, $M_{\text{age}} = 36.9$, $SD = 12$, 19 women, 25 men), or categorized them on their perceived ethnoracial membership ($N_{\text{total}} = 97$, $N_{\text{final}} = 87$, $M_{\text{age}} = 38.1$, $SD = 12.4$, 33 women, 54 men). Overall, 96% of the raters were born in the U.S. and on average scored near the midpoint of the social and economic

conservatism scale ($M = 58.9$, $SD = 20.3$). Full sample characteristics can be found in **Supplementary Table S2**.

Procedure

Participants were sequentially presented with 61 average classification images derived from the various cluster algorithms (complete linkage vs. Ward's method) and number of clusters described above (5, 10, 15) and asked "How [dangerous, American, dominant] is this person?". Participants only evaluated one trait from 1 (not at all) to 9 (extremely). For ethnoracial categorization, participants saw each face separately and were asked "What race/ethnicity is this person?" with five options: Asian, black, Latino or Hispanic, Middle Eastern, or white. The set of face stimuli for all the ratings were randomized. Each face image was presented at 300 pixels x 300 pixels. All trials were self-paced and there was a 250 ms delay before the next trial. All images were repeated and rated a second time to assess test-retest reliability. Reliability was greatest for ethnoracial categorization (Mean $r = .47$, $SD = .23$, followed by dangerous (Mean $r = .35$, $SD = .21$, American (Mean $r = .37$, $SD = .20$, and dominant (Mean $r = .30$, $SD = .24$) ratings.

Analyses

Trait ratings were analyzed using mixed effect regression models in the lme4 package (Bates et al., 2015). Ratings were predicted by a face variable, random intercepts by participants and by the interaction between participants and faces to account for the repeated measures by participants and ratings given to each average face. Satterthwaite approximations were used to calculate degrees of freedom (Luke, 2016). Marginal means, two-sided pairwise comparisons of the average trait ratings given to each face, and effect sizes (Cohen's d) were estimated using the emmeans package version 1.5.3 (Lenth, 2020). Comparisons were corrected for false discovery

rate. Since we focused on the 5 cluster solution using Ward's method, we extracted and report those specific estimates from the full model. The ethnoracial classification ratings were descriptively summarized with a percentage of the sample that responded each category per face. The top two categories and their percentages are shown for each of the cluster faces in the figures.

Results

Theory-based approach

To illustrate the issues with the theory-based approach, **Figure 2** displays the average classification image for the full sample, the four main state-label conditions, and various levels of political orientation. The five score ranges on the SECS were chosen to reflect equidistant scores that could be interpreted as meaningfully different levels of political orientation (i.e., numbers closer to the ends are more politically extreme than numbers closer to 50). A visual examination could lead researchers to infer that a dark-skinned person is a shared representation held by most participants no matter the labels, state, or political orientation— a similar face was produced across all the variables. The one exception is the image produced by very liberal participants in the upper-right corner which might suggest liberals hold a different understanding of illegalized immigrants. This exception also averages across only 30 participants, so smaller sample sizes can reveal representations that get occluded when you average across too many participants. The average face produced by extreme liberals is not visible within the faces from the state-label conditions even though they are all derived from the same set of participants. Researchers would need to investigate all potentially relevant variables to find hidden representational variation, a very costly endeavor. Such an endeavor may even prove futile if single variables fail to capture

representational similarities by assuming cohesiveness. For instance, some extreme liberals may produce a similar face as some extreme conservatives which would go unnoticed if they are instead averaged with others who share a political orientation. This example situation *would* be accounted for by the data-driven approach which instead relies on probabilities to describe the relationships between individual differences and representations.

Figure 2. Average classification images using a theory-driven approach. The center face reflects the average classification images across the full sample (i.e., grand average). The left side shows how the grand average image partitions into average faces for the various state-label conditions, the right side partitions the grand average into various levels of the social and economic conservatism scale (higher numbers reflects more conservative participants). The number at the bottom represents the sample size of each average face.

Data-driven approach

Cluster content

The five-cluster solution led to imbalanced clusters. In terms of percentage of the participants, Cluster 1 was the largest (38%), followed by Cluster 2 (22%), Cluster 3 (17%), Cluster 4 (16%), and Cluster 5 was the smallest (7%) (**Figure 3a**). The average cluster classification images were phenotypically different across clusters (**Figure 3b**), however it's important to understand how they differ in terms of elicited social evaluations (**Figure 3c**).

Figure 3. Five cluster solution from the hierarchical clustering analysis. **a.** Dendrogram depicting the five clusters in various colors. **b.** The average faces and sample sizes from each cluster (number and color coded to match the dendrogram). **c.** Race and trait ratings of each cluster's average face. The percentages in the grey box reflect the top two ethnracial classifications. The plots show estimated marginal means from the trait ratings. The error bars depict 95% confidence intervals. The violins display the distributions of rating data while the dashed lines represent quartiles. The horizontal line is the neutral scale position (5) of each rating.

The following results will not report traits whose estimated means and confidence intervals crossed the neutral rating (5), suggesting that trait was less critical for characterizing that cluster's representation. Cluster 1 contained a representation perceived as somewhat dangerous ($M = 5.80$ CI[5.30, 6.31]), dominant ($M = 6.24$ CI[5.44, 7.04]), and mostly categorized as black. Cluster 2 instead contained a representation perceived as very American ($M = 6.98$ CI[6.13, 7.83]), not dangerous ($M = 3.80$ CI[3.30, 4.31]), and a mix of Latinx and black. Cluster 3 was perceived as mostly dominant ($M = 5.96$ CI[5.16, 6.75]) and a mix of black and Asian.

Cluster 4 was perceived as American ($M= 6.19$ CI[5.34, 7.04]), dominant ($M= 6.09$ CI[5.29, 6.89]), and a mix of black and Latinx. Cluster 5's representation was perceived as American ($M= 5.88$ CI[5.03, 6.73]), dominant ($M= 6.29$ CI[5.49, 7.09]), and mostly black.

Pairwise cluster comparisons of the ratings corroborated the cluster-specific content. Cluster 2 was the least dangerous representation ($bs_{range}= [.84, 2.0]$, $ds_{range}= [.41, .97]$, $ps_{range}= [<.0001, .003]$), Clusters 1 and 5 were both the most dangerous ($bs_{range}= [.73, 2.0]$, $ds_{range}= [.35, .97]$, $ps_{range}= [<.0001, .009]$). Clusters 1, 3, 4, and 5 were similarly dominant ($bs_{range}= [.05, .34]$, $ds_{range}= [.03, .19]$, $ps_{range}= [.290, .808]$) and all more dominant than Cluster 2 ($bs_{range}= [.72, 1.06]$, $ds_{range}= [.40, .59]$, $ps_{range}= [.0001, .006]$). Lastly, the American pattern was more complex. Cluster 2 was the most American representation ($bs_{range}= [.79, 1.95]$, $ds_{range}= [.41, 1.00]$, $ps_{range}= [<.0001, .005]$). Cluster 4 was more American than Cluster 1 ($b= .75$ CI[-.01, 1.51], $d=.38$, $p=.008$) and Cluster 3 ($b= 1.15$ CI[.40, 1.96], $d=.59$, $p=.0001$). Cluster 5 was more American than Cluster 3 ($b= .85$ CI[.09, 1.6], $d=.43$, $p=.003$). All other comparisons were not significant ($ps>.122$).

The cluster representations can be interpreted in the following way. Clusters 1 and 5 are dark-skinned threat representations. Clusters 2 and 4 operate as non-threatening representations, although Cluster 2 is lighter skinned and less threatening than Cluster 4. Cluster 3 is a dark-skinned dominant representation, the most evaluatively neutral out of all the clusters. Cluster 4 and Cluster 3, to a lesser extent, paint a nuanced picture of collective schemas of illegality and racialization because they suggest that darker skin is not always represented or perceived as negative (e.g., Alter et al., 2016). These results provide evidence of representational variation in understandings of illegalized immigrants – some are positive, some are negative, and they vary in skin tone and phenotype, suggesting people might have different nationality or racial associations

in mind (Martinez, Oh, et al., 2021b). Interestingly, most of the representations were perceived as dominant (except Cluster 2), suggesting there are also various ways that dominance is being understood in relation to illegalized immigrants.

Cluster membership

Representational content alone is an insufficient procedure for validating clusters since even subtle differences between face images can lead to different evaluations. Validation here then further relies on identifying which individual differences track cluster membership and if those individual differences theoretically match the representational content of that cluster.

Decision tree. The decision tree provides a descriptive picture of which variables are important predictors of cluster membership and how they interact. The tree shows that age is the most important predictor, followed by estimates of the size of the illegalized population in one's neighborhood, followed by the state where participants live (**Figure 4**). The cluster distribution in the bottom figures suggests that the collection of variables do not optimally differentiate between clusters, so our analytic inferences focus on relative probabilities. Moreover, it should be noted that Cluster 1 will always show the largest probability as it was the largest cluster, so inferences will specifically focus on relative probabilities of the other clusters.

Starting from the age variable, the first important split is between individuals who are older vs. younger than 48 years old. Older participants show a higher probability of being in Clusters 3, 2, and 5 which are a mix of neutral, non-threatening, and threatening representations. For participants under 48, the next important split is between younger participants who perceive the illegalized population size in their neighborhood to be a more or less than 26%. Those who perceive a lot of illegalized immigrants in their neighborhood are very likely to be part of Cluster

2 and 4, the non-threatening representations. Lastly, for young participants who perceive few illegalized immigrants in their neighborhoods, the next tentative split is whether they live in Texas vs. California and the other states. The subtle difference here seems to be that Californians are more likely to be in Cluster 2, a non-threatening representation, while Texans are more likely to be in Cluster 5, a threatening representation.

The decision tree provided initial evidence that the clusters *are* meaningful. Threatening representations tended to occur for older participants or younger participants in Texas who do not perceive there to be many illegalized immigrants in their neighborhoods. Non-threatening representations tended to occur for younger participants who perceive there to be many illegalized immigrants in their neighborhoods, or who perceive fewer but live in California and other states. These results highlight the importance of geographic and specific demographic variables as sources of representational variation. However, since decision trees have overfitting issues, we test the robustness of these variables in random forests.

Figure 4. Conditional inference decision tree for five cluster membership. The ovals represent the most important predictor variables and their p values. The lines between variables and the bottom plots depict the level of the variables at which splits occur. The bottom plots depict the resulting probabilistic cluster membership distribution based on various tree partitions. The x-axes are the different clusters and the y-axis is the probability of belonging to each cluster. Sample sizes are provided above the plot for each partition. The top left faces are a legend to remind what each cluster represents. The “size.hood” variable is estimates of the undocumented immigrant population size in their neighborhood, “age” is the age of the participants, “state2” is the state they live in.

Random forest. The permutation-based p values for the importance of each variable can be seen in **Figure 5**. Only one of the variables identified by the decision tree remained significant: estimations of illegalized immigrant size in one’s neighborhood ($p = .030$). Age ($p = .079$), the 2016 vote ($p = .089$), estimations of illegalized immigrants in one’s city ($p = .139$), and lived-in state ($p = .188$) were the next highest important variables, though not significant.

Since neighborhood estimates and age were consistently identified as relatively more important, we explored how they are related to cluster membership more closely.

Figure 5. Permutation variable importance for five cluster solution. The x-axis is the permuted p value representing variable importance (significant alpha is .05). The exact p value is provided to the right of each bar. The y-axis lists all the individual difference measures.

Partial dependence plots suggested that the importance of neighborhood estimates occurred mainly for Cluster 2 and Cluster 3 (**Figure 6**). The more immigrants participants perceived in their neighborhood, the more likely they were to be a member of Cluster 2 (a non-threatening representation) and less a member of Clusters 1, 3, and 5 (more neutral and threatening representations). Similarly, the older the participants, the less likely they were to be

part of Cluster 2 and 4 (non-threatening representations) and more likely to be a member of Cluster 3 and 5 (more neutral and threatening representations).

Figure 6. Partial dependence plots for five cluster solution. a. Probabilistic cluster membership across undocumented population size in the neighborhood. **b.** Probabilistic cluster membership across age. The y-axis is the probability of cluster membership. The black tick marks above the x-axis are data points such that the less tick marks the less participants at that level of the x-axis variable.

Does the number of clusters matter?

When these five clusters are further partitioned into ten or fifteen smaller clusters, some results remained consistent while new ones emerged (**Supplementary Figures S7 to S12**). An

example of a pattern that remained consistent was the partitioning of Cluster 2 in the 10 and 15 cluster solutions. No matter how many subclusters Cluster 2 was partitioned into, the resulting representations tended to be lighter skinned and perceived as highly American and non-dangerous. Smaller clusters like Cluster 3 and Cluster 4 only partitioned into two separate representations in the 15 cluster solution and those representations were given similar ratings as the original cluster image. Cluster 5 remained intact in all the solutions. The largest cluster however, Cluster 1, was partitioned into many subclusters which revealed hidden images – some similar to the original cluster, others smiling and non-threatening, but all dark skinned. Moreover, while most of the ethnoracial classification ratings from the 5 cluster solution tended to fall along Latinx, black, and Asian categories, more subclusters resulted in some images also being perceived as white and Middle Eastern. Lastly, neighborhood size estimates and age remained important variables that predicted cluster membership for the 10 cluster solution and age for the 15 cluster solution in a way that mirrored the 5 cluster solution.

These results suggest that the hierarchical clustering uses different features for grouping different levels of hierarchy. It appears that larger clusters are mainly grouped by phenotypic information while the smaller clusters use affective facial expressions. This could occur because when a cluster of images share the same phenotype (e.g., skin tone) and thus overlap highly in visual information, what differentiates pairs of faces within this cluster are more fine-grained features, like facial expressions. These patterns suggest five clusters can provide a coarse summary of representational variation within the data, yet there may be more nuanced representations at higher cluster numbers.

Discussion

With the goal of mapping varied understandings of social categories, here illegalized immigrants, this project addressed two issues with standard reverse correlation practices that can restrict inferences about mental representations. First, averages can be misleading and, second, theory-based investigations of individual differences may not fully capture the social contours of collective representations. We proposed a data-driven approach that calculated inter-individual similarity in representations, identified clusters of similar representations, and characterized the extent to which individual differences tracked cluster membership. This approach found highly shared representations and less shared representations that were somewhat clustered by demographic (age) and geographic variables (size perceptions of illegalized population in one's neighborhood). We summarize the insights gleaned from this procedure followed by acknowledging limitations worth addressing in future research.

Varied schemas for illegalized immigrants

A limitation of data-driven approaches is that the resulting solutions may not always be easily categorizable or interpretable, the researcher must then reintroduce theoretical knowledge to aid in interpretation (Adolphs et al., 2016). With this in mind, we identified clusters whose average representations tended to fall along lines of threat vs. non-threat across varied skin tones. A shared thread among most of the representations was darker skin and facial features associated with dominance, suggesting these features are highly associated with illegality across more minds. Cluster 1's representation was the most shared and therefore more crosscutting and not predicted by any individual difference variables. This may be why the image from Cluster 1, a dark-skinned threat, appears similar to average classification images (**Figure 2**) and to sample-averaged representations from previous research (Martinez, Oh, et al., 2021a).

However, while that face is highly shared, meaningful pockets of individuals have a different person in mind. Representations from Cluster 5 and Cluster 3 portrayed a similar face to Cluster 1, except Cluster 5 was more threatening than Cluster 3. Thus, Cluster 1 and Cluster 5 could be interpreted as threat representations, while Cluster 3 could be interpreted as an evaluatively neutral representation that simply highlights beliefs that illegalized immigrants are darker skinned. Representations from Cluster 2 and Cluster 4 both tended to be non-threatening, yet the former was lighter skinned and the latter darker skinned. These two clusters are confirmation of the ambivalence that exists in portrayals of immigrants (Reyna et al., 2013): illegalized immigrants are not hated by everyone. It is important to note that these clusters should be seen as examples of existing representations, not a comprehensive typology of all possible representations. Collecting a broader sample across the U.S., in other countries, and across time could provide a wider typology of existing representations and reveal further sources that contribute to variation.

Sources of representational variation

While we tested many variables theorized to be relevant to immigrant representations, none fully differentiated between clusters, and many were not significant predictors of clustered variation. Attitudes measured from normative questions about economic, political, and social aid to illegalized immigrants did not differentiate between representational clusters (Ommundsen et al., 2014). Neither did political ideologies, party affiliations, or voting decisions (Chavez, 2013; Kunst et al., 2019; Sati, 2020; Sides et al., 2018), nor demographics like education level or religious affiliation (Abrajano & Hajnal, 2015; Davis & Perry, 2020). Race did not either, although was potentially expected given that a categorical race measure does not reflect that race is not an explanatory variable but an illusory multidimensional product of racist practices (Fields

& Fields, 2012; Helms et al., 2005; Sen & Wasow, 2016). Lastly, the label used to name illegalized immigrants also did not matter here, in line with the assertion that “undocumented” and “illegal” are too mired in criminalization that they both conjure similar representations (Martinez, Oh, et al., 2021a; Plascencia, 2009).

Our models instead revealed that Cluster 1’s crosscutting racialized representation could not be explained by any of the measures, suggestive of a common source like media (Martinez, Feldman, et al., 2021), while the rest of the clusters were differentiated by more idiosyncratic sources: age and geographic characteristics. Specifically, older participants tended to be more associated with clusters that contained either neutral or threatening and darker-skinned representations of illegalized immigrants. Younger participants instead tended to be associated with more non-threatening representations that varied in skin tone. Similarly, participants who perceived there to be many illegalized immigrants in their neighborhood were more associated with non-threatening representations, those who perceived less were more likely to hold threatening representations. The same patterns did not occur for perceptions at the level of the city or the state, suggesting that perceptions are shaped by a very localized space. Although we did not assess interactions between these variables in the random forests, the decision tree did suggest that age and neighborhood estimates interacted such that the cluster membership of younger participants with high neighborhood immigrant estimates tended to hold the most non-threatening representations (i.e., Cluster 2 and 4). Likewise, there seemed to be a tentative interaction with state such that those younger participants with smaller neighborhood immigrant estimates in Texas held more threatening representations than those in California, in line with each state’s general political environment against or for immigrants. These results converge with policy surveys of white respondents suggesting that those who lived around more Latinxs tended

to favor less restrictions for illegalized immigrants, whereas those who were embedded in older-aged social networks favored more restrictions against illegalized immigrants (Berg, 2009).

These results suggest that reverse correlation is not simply a measure of attitudes but can reflect participants' context-specific perceptions about people's appearances. People's local face diet develops or constrains their representations (Dotsch et al., 2016). Fittingly, it is specifically estimates from local environments (e.g., neighborhood) that are salient in perceptions of demographic population sizes (Wong, 2007). Unfortunately, we did not collect measures of how long participants resided in the same *neighborhood*. However, most of the participants lived in their *state* for over 10 years indicating how age plays a role in population size estimates. Chronological age as a variable confounds many age-related influences (North, 2019), yet within our analyses it may function more as a proxy measure for the amount of experience with local geography as opposed to clear-cut life stages or generational cohort effects. Those who lived in their state longer were on average older, their longer presence in a specific geo-political region had greater chance to shape their perceptions of their local social environment.

It is suggested that the diversity of local environments tunes mental representations such that more diversity leads to less differentiated evaluative representations of social groups (Bai et al., 2020). Congruently, participants who estimated a more illegalized immigrants in their neighborhoods tended to represent illegalized immigrants as phenotypically diverse yet similarly non-threatening. The majority of these participants also lived in their state for over 10 years, suggesting they potentially adjusted to local diversity (Ramos et al., 2019). Conversely, participants who estimated less local illegalized immigrants instead held more varied representations both phenotypically and evaluatively, indicative of more differentiated stereotypes. Yet, if estimates of local illegalized immigrants are low, how might their

representations have developed? Although visibility can be dangerous for illegalized immigrants (Asad, 2020) and interaction with illegalized immigrants can be perceived as rare (Martinez, Oh, et al., 2021a), people may draw from available social representations (e.g., media portrayals) to profile specific people in their local environments as illegalized immigrants and update their mental representations accordingly (Romero, 2006).

Limitations

While theory-based approaches have critical limitations, so do data-driven approaches. One disadvantage is that there is no baseline answer about how many representational schemas of illegalized immigrants exist. Inferences must instead be made from exploring a range of cluster numbers, similarity metrics, clustering algorithms – from which we examined a restricted range. Therefore, while this approach was used to extend our knowledge of variation in representations of illegalized immigrants, the procedure itself could be better validated using social categories with clear and known visual prototypes (e.g., age groups).

Another limitation is how less important variables may be influential through indirect means. For example, while political orientation contributed little to representational variation, agenda-laden media consumption could hypothetically influence estimates of immigrant population sizes and subsequent visual experiences. Collecting a more expansive set of measures could improve our descriptions of the relevant features that distinguish between clusters.

Lastly, big data approaches cannot easily provide a deep understanding of why variables relate the way they do or of the content of clusters without the combined aid of immersive insights gleaned from additional sources of data (e.g., Wang & Liu, 2021). For example, while this study focused on classifying clusters of representations, the goal is not to essentialize people

with certain traits or experiences as necessarily possessing specific representations or beliefs. It is critical to further examine the processes by which some younger or older individuals and those who perceive various amounts of immigrants in their neighborhoods come to develop different understandings of illegalized immigrants that shape their visual expectations. The current classification analyses therefore only provide an initial clue: a map of shared and idiosyncratic understandings, even if only temporary or contingent. A contextual and processual account of the development of these representational clusters is additionally required for understanding why these relationships exist at all.

Conclusions

We have shown that data-driven approaches are a viable way to map the diversity in representations of illegalized immigrants and how they pattern across people. Machine learning analyses revealed sources of representational variation that can go unnoticed by standard practices: representations of illegalized immigrants are not necessarily all dark-skinned threats, there exists variation in encoded valence and phenotype that is clustered by demographic and social geographic differences. This was an initial attempt at mapping variance and can likely be improved, but the results highlight a critical need to place representational variation at the forefront of investigations of mental representations. Variation is foundational, ubiquitous, and its mapping can uncover important insights (Kahneman et al., 2021; Martinez et al., 2020), yet it can be neglected by disciplinary norms that uphold averages as the inferential standard. If averages are not handled carefully by simultaneously exploring existing and meaningful heterogeneity, researchers themselves can establish reified caricatures of the public and their beliefs (Martinez & Paluck, 2020).

Context

This study is part of a larger project attempting to map mental representations of immigrants when they are criminalized, identify the sources of those representations, and develop ways to counteract them, for example with narratives (Martinez, Feldman, et al., 2021). In the process, it became clear that developing theoretical inferences from sample-averaged or measure-averaged reverse correlated face images was insufficient to answer these questions, they produced caricatures. We needed to build a methodology that allowed investigating variation within complex multidimensional data like face images, and that could serve as a necessary analytic framework for future research on representations. Concurrently, this project facilitates more use of machine learning techniques within psychology to answer important social and political questions at the level of cognition. It also pushes psychology's theoretical territory to examine previously neglected populations and see the value in investigating their conditions. In our case, illegalized immigrants are an interesting test case for this methodology. This population often avoids the danger of system and public visibility, suggesting some people's representations would develop from idiosyncratic contact, but most would develop from shared sources - which is what we found.

References

- Abrajano, M., & Hajnal, Z. (2015). *White backlash: Immigration, race, and American politics*. Princeton University Press.
- Ackerman, E. F. (2013). The Rise of the “Illegal Alien.” *Contexts*, 12(3), 72–74.
<https://doi.org/10.1177/1536504213499883>
- Ackerman, E. F. (2014). ‘What part of illegal don’t you understand?’: Bureaucracy and civil society in the shaping of illegality. *Ethnic and Racial Studies*, 37(2), 181–203.
<https://doi.org/10.1080/01419870.2012.705008>
- Adolphs, R., Nummenmaa, L., Todorov, A., & Haxby, J. V. (2016). Data-driven approaches in the investigation of social perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1693), 20150367–20150367.
<https://doi.org/10.1098/rstb.2015.0367>
- Akosa, J. (2017). *Predictive Accuracy: A Misleading Performance Measure for Highly Imbalanced Data*. 12.
- Alamillo, R., Haynes, C., & Madrid, R. (2019). Framing and immigration through the trump era. *Sociology Compass*, 13(5), e12676. <https://doi.org/10.1111/soc4.12676>
- Alter, A. L., Stern, C., Granot, Y., & Balcetis, E. (2016). The “Bad Is Black” Effect: Why People Believe Evildoers Have Darker Skin Than Do-Gooders. *Personality and Social Psychology Bulletin*, 42(12), 1653–1665. <https://doi.org/10.1177/0146167216669123>
- Altmann, A., Tološi, L., Sander, O., & Lengauer, T. (2010). Permutation importance: A corrected feature importance measure. *Bioinformatics*, 26(10), 1340–1347.
<https://doi.org/10.1093/bioinformatics/btq134>

- Apley, D. W., & Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. <https://doi.org/10.1111/rssb.12377>
- Asad, A. L. (2020). On the Radar: System Embeddedness and Latin American Immigrants' Perceived Risk of Deportation. *Law & Society Review*, 54(1), 133–167. <https://doi.org/10.1111/lasr.12460>
- Bai, X., Ramos, M. R., & Fiske, S. T. (2020). As diversity increases, people paradoxically perceive social groups as more similar. *Proceedings of the National Academy of Sciences*, 202000333. <https://doi.org/10.1073/pnas.2000333117>
- Bates, D. M., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Berg, J. A. (2009). White Public Opinion toward Undocumented Immigrants: Threat and Interpersonal Environment. *Sociological Perspectives*, 52(1), 39–58. <https://doi.org/10.1525/sop.2009.52.1.39>
- Birhane, A. (2021). The Impossibility of Automating Ambiguity. *Artificial Life*, 1–18. https://doi.org/10.1162/artl_a_00336
- Boutyline, A. (2017). Improving the Measurement of Shared Cultural Schemas with Correlational Class Analysis: Theory and Method. *Sociological Science*, 4, 353–393. <https://doi.org/10.15195/v4.a15>
- Breiman, L. (2001a). Random Forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Breiman, L. (2001b). Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3), 199–231. <https://doi.org/10.1214/ss/1009213726>

- Brinkman, L., Todorov, A., & Dotsch, R. (2017). Visualising mental representations: A primer on noise-based reverse correlation in social psychology. *European Review of Social Psychology*, 28(1), 333–361. <https://doi.org/10.1080/10463283.2017.1381469>
- Chavez, L. R. (2013). *The Latino threat: Constructing immigrants, citizens, and the nation* (Second edition). Stanford University Press.
- Chen, C. (2004). *Using Random Forest to Learn Imbalanced Data*. 12.
- Cikara, M. (2021). Causes and Consequences of Coalitional Cognition. *Advances in Experimental Social Psychology*, 64, 65–128. <https://doi.org/10.31219/osf.io/ktpf7>
- Cikara, M., Fouka, V., & Tabellini, M. (2020). *Hate crime increases with minoritized group rank* [Preprint]. Open Science Framework. <https://doi.org/10.31219/osf.io/2z3kw>
- Cone, J., Brown-Iannuzzi, J. L., Lei, R., & Dotsch, R. (2020). Type I Error Is Inflated in the Two-Phase Reverse Correlation Procedure. *Social Psychological and Personality Science*, 194855062093861. <https://doi.org/10.1177/1948550620938616>
- Craig, M. A., Rucker, J. M., & Richeson, J. A. (2018). Racial and political Dynamics of an Approaching “Majority-Minority” United States. *Annals of the American Academy of Political and Social Science*, 677(1), 204–214. <https://doi.org/10.1177/0002716218766269>
- Davis, J. T., & Perry, S. L. (2020). White Christian Nationalism and Relative Political Tolerance for Racists. *Social Problems*, spaa002. <https://doi.org/10.1093/socpro/spaa002>
- DeBruine, L., & Jones, B. (2017). *Face Research Lab London Set* (p. 281699312 Bytes) [Data set]. figshare. <https://doi.org/10.6084/M9.FIGSHARE.5047666.V3>
- Dotsch, R. (2017). *rcicr: Reverse correlation image classification toolbox*. (R package version 0.4.0) [Computer software].

- Dotsch, R., Hassin, R. R., & Todorov, A. (2016). Statistical learning shapes face evaluation. *Nature Human Behaviour*, 1(1), 1–6. <https://doi.org/10.1038/s41562-016-0001>
- Dotsch, R., & Todorov, A. (2012). Reverse correlating social face perception. *Social Psychological and Personality Science*, 3(5), 562–571. <https://doi.org/10.1177/1948550611430272>
- Dotsch, R., Wigboldus, D. H. J., Langner, O., & van Knippenberg, A. (2008). Ethnic Out-Group Faces Are Biased in the Prejudiced Mind. *Psychological Science*, 19(10), 978–980. <https://doi.org/10.1111/j.1467-9280.2008.02186.x>
- Enos, R. D. (2017). *The space between us: Social geography and politics*. Cambridge University Press.
- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, 53(2), 134–140. <https://doi.org/10.1037/h0045156>
- Everett, J. A. C. (2013). The 12 Item Social and Economic Conservatism Scale (SECS). *PLoS ONE*, 8(12). <https://doi.org/10.1371/journal.pone.0082131>
- Fields, K. E., & Fields, B. J. (2012). *Racecraft: The soul of inequality in American life*. Verso.
- Filindra, A., & Kovács, M. (2012). Analysing US State Legislative Resolutions on Immigrants and Immigration: The Role of Immigration Federalism. *International Migration*, 50(4), 33–50. <https://doi.org/10.1111/j.1468-2435.2010.00658.x>
- Fisher, A. J., Medaglia, J. D., & Jeronimus, B. F. (2018). Lack of group-to-individual generalizability is a threat to human subjects research. *Proceedings of the National Academy of Sciences*, 201711978–201711978. <https://doi.org/10.1073/pnas.1711978115>
- Griffith, B. (2021, March 22). *Map: Sanctuary Cities, Counties, and States*. CIS.Org. <https://cis.org/Map-Sanctuary-Cities-Counties-and-States>

- Harrell, F. (2017, January 15). *Classification vs. Prediction*. Statistical Thinking.
<https://www.fharrell.com/post/classification/>
- Haslam, S. A., Turner, J. C., Oakes, P. J., McGarty, C., & Reynolds, K. J. (1997). The Group as a Basis for Emergent Stereotype Consensus. *European Review of Social Psychology*, 8(1), 203–239. <https://doi.org/10.1080/14792779643000128>
- Haynes, C., Merolla, J. L., & Ramakrishnan, S. K. (2016). *Framing immigrants: News coverage, public opinion, and policy*. Russell Sage Foundation.
- Helms, J. E., Jernigan, M., & Mascher, J. (2005). The meaning of race in psychology and how to change it: A methodological perspective. *American Psychologist*, 60(1), 27–36.
<https://doi.org/10.1037/0003-066X.60.1.27>
- Hinzman, L., & Maddox, K. B. (2017). Conceptual and visual representations of racial categories: Distinguishing subtypes from subgroups. *Journal of Experimental Social Psychology*, 70, 95–109. <https://doi.org/10.1016/j.jesp.2016.12.012>
- Hönekopp, J. (2006). Once more: Is beauty in the eye of the beholder? Relative contributions of private and shared taste to judgments of facial attractiveness. *Journal of Experimental Psychology. Human Perception and Performance*, 32(2), 199–209.
<https://doi.org/10.1037/0096-1523.32.2.199>
- Hong, Y., & Ratner, K. G. (2020). Minimal but not meaningless: Seemingly arbitrary category labels can imply more than group membership. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspa0000255>
- Hopkins, N., & Dixon, J. (2006). Space, Place, and Identity: Issues for Political Psychology. *Political Psychology*, 27(2), 173–185. <https://doi.org/10.1111/j.1467-9221.2006.00001.x>

- Hothorn, T., Hornik, K., & Zeileis, A. (2006). Unbiased Recursive Partitioning: A Conditional Inference Framework. *Journal of Computational and Graphical Statistics*, 15(3), 651–674. <https://doi.org/10.1198/106186006X133933>
- Huo, Y. J., Dovidio, J. F., Jiménez, T. R., & Schildkraut, D. J. (2018a). Local policy proposals can bridge Latino and (most) white Americans' response to immigration. *Proceedings of the National Academy of Sciences*, 115(5), 945–950. <https://doi.org/10.1073/pnas.1711293115>
- Huo, Y. J., Dovidio, J. F., Jiménez, T. R., & Schildkraut, D. J. (2018b). Not Just a National Issue: Effect of State-Level Reception of Immigrants and Population Changes on Intergroup Attitudes of Whites, Latinos, and Asians in the United States: Local Immigration Policy. *Journal of Social Issues*. <https://doi.org/10.1111/josi.12295>
- Imhoff, R., Dotsch, R., Bianchi, M., Banse, R., & Wigboldus, D. H. J. (2011). Facing Europe: Visualizing spontaneous in-group projection. *Psychological Science*, 22(12), 1583–1590. <https://doi.org/10.1177/0956797611419675>
- Jones, Z., & Linder, F. (2015). *Exploratory Data Analysis using Random Forests*. 31.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown and Company.
- Krosch, A. R., & Amodio, D. M. (2014). Economic scarcity alters the perception of race. *Proceedings of the National Academy of Sciences*, 111(25), 9079–9084. <https://doi.org/10.1073/pnas.1404448111>
- Kunst, J. R., Dovidio, J. F., & Thomsen, L. (2019). Fusion with political leaders predicts willingness to persecute immigrants and political opponents. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-019-0708-1>

- Lei, R. F., & Bodenhausen, G. V. (2017). Racial assumptions color the mental representation of social class. *Frontiers in Psychology*, 8(April), 1–7.
<https://doi.org/10.3389/fpsyg.2017.00519>
- Lenth, R. V. (2020). *emmeans: Estimated Marginal Means, aka Least-Squares Means*. (R package version 1.5.3) [Computer software]. <https://github.com/rvlenth/emmeans>
- Luke, S. G. (2016). Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods*, 1–9. <https://doi.org/10.3758/s13428-016-0809-y>
- Martinez, J. E., Feldman, L. A., Feldman, M. J., & Cikara, M. (2021). Narratives shape cognitive representations of immigrants and immigration-policy preferences. *Psychological Science*, 1–18. <https://doi.org/10.1177/0956797620963610>
- Martinez, J. E., Funk, F., & Todorov, A. (2020). Quantifying idiosyncratic and shared contributions to judgment. *Behavior Research Methods*, 52(4), 1428–1444.
<https://doi.org/10.3758/s13428-019-01323-0>
- Martinez, J. E., Oh, D., & Todorov, A. (2021a). Immigrant faceism: Documentation statuses evoke racialized mental representations. *PsyArxiv*. <https://doi.org/10.31234/osf.io/xb62n>
- Martinez, J. E., Oh, D., & Todorov, A. (2021b). Mental representations of immigrants encode racialized expectations of socio-structural positions. *PsyArxiv*.
<https://doi.org/10.31234/osf.io/cvhze>
- Martinez, J. E., & Paluck, E. L. (2020). Quantifying shared and idiosyncratic judgments of racism in social discourse. *PsyArxiv*. <https://doi.org/10.31234/osf.io/kfpjg>
- Molnar, C. (2020). *Interpretable machine learning ; a guide for making black box models explainable*.

- North, M. S. (2019). A GATE to Understanding “Older” Workers: Generation, Age, Tenure, Experience. *Academy of Management Annals*, 13(2), 414–443.
<https://doi.org/10.5465/annals.2017.0125>
- Ommundsen, R., Larsen, K. S., van der Veer, K., & Eilertsen, D.-E. (2014). Framing Unauthorized Immigrants: The Effects of Labels on Evaluations. *Psychological Reports*, 114(2), 461–478. <https://doi.org/10.2466/17.PR0.114k20w0>
- Pew Research Center. (2019a). Metro area U.S. unauthorized immigrant population estimates, 2016 and 2007. *Pew Research Center’s Hispanic Trends Project*.
<https://www.pewresearch.org/hispanic/interactives/unauthorized-immigrants-by-metro-area-table/>
- Pew Research Center. (2019b). U.S. unauthorized immigrant population estimates by state, 2016. *Pew Research Center’s Hispanic Trends Project*.
<https://www.pewresearch.org/hispanic/interactives/u-s-unauthorized-immigrants-by-state/>
- Plascencia, L. F. B. (2009). The “Undocumented” Mexican Migrant Question: Re-Examining The Framing Of Law And Illegalization In The United States. *Urban Anthropology and Studies of Cultural Systems and World Economic Development*, 38(2/3/4), 375–434.
- Ramos, M. R., Bennett, M. R., Massey, D. S., & Hewstone, M. (2019). Humans adapt to social diversity over time. *Proceedings of the National Academy of Sciences*, 116(25), 12244–12249. <https://doi.org/10.1073/pnas.1818884116>
- Rentfrow, P. J., & Jokela, M. (2016). Geographical Psychology: The Spatial Organization of Psychological Phenomena. *Current Directions in Psychological Science*, 25(6), 393–398.
<https://doi.org/10.1177/0963721416658446>
- Reyna, C., Dobria, O., & Wetherell, G. (2013). The complexity and ambivalence of immigration attitudes: Ambivalent stereotypes predict conflicting attitudes toward immigration

policies. *Cultural Diversity and Ethnic Minority Psychology*, 19(3), 342–356.

<https://doi.org/10.1037/a0032942>

Rodríguez-Muñiz, M. (2021). *Figures of the future: Latino civil rights and the politics of demographic change*. Princeton University Press.

Romero, M. (2006). Racial Profiling and Immigration Law Enforcement: Rounding Up of Usual Suspects in the Latino Community. *Critical Sociology*, 32(2–3), 447–473.

<https://doi.org/10.1163/156916306777835376>

Rosa, J. (2019). Contesting Representations of Migrant “Illegality” through the Drop the I-Word Campaign Rethinking Language Change and Social Change. In N. Avineri, L. R. Graham, E. J. Johnson, R. Conley Riner, & J. Rosa (Eds.), *Language and social justice in practice* (pp. 35–43). Routledge.

Rucker, J. M., Murphy, M. C., & Quintanilla, V. D. (2019). The immigrant labeling effect: The role of immigrant group labels in prejudice against noncitizens. *Group Processes & Intergroup Relations*, 136843021881874. <https://doi.org/10.1177/1368430218818744>

Sati, J. C. (2020). On Trumpism and Illegality. In O. K. Obasogie (Ed.), *Trumpism and its discontents* (pp. 115–129). Berkeley Public Policy Press.

Savaş, Ö., Greenwood, R. M., Blankenship, B. T., Stewart, A. J., & Deaux, K. (2021). All immigrants are not alike: Intersectionality matters in views of immigrant groups. *Journal of Social and Political Psychology*, 9(1), 86–104. <https://doi.org/10.5964/jspp.5575>

Schönbrodt, F. D., & Perugini, M. (2013). At what sample size do correlations stabilize? *Journal of Research in Personality*, 47(5), 609–612. <https://doi.org/10.1016/j.jrp.2013.05.009>

Sen, M., & Wasow, O. (2016). Race as a bundle of sticks: Designs that estimate effects of seemingly immutable characteristics. *Annual Review of Political Science*, 19(1), 499–522. <https://doi.org/10.1146/annurev-polisci-032015-010015>

- Sides, J., Tesler, M., & Vavreck, L. (2018). *Identity crisis: The 2016 presidential campaign and the battle for the meaning of America*. Princeton University Press.
- Stangor, C., & Schaller, M. (2000). Stereotypes as individual and collective representations. In C. Stangor (Ed.), *Stereotypes and prejudice: Essential readings* (pp. 64–82). Psychology Press.
- Stephenson, W. (1935). Correlating persons instead of tests. *Journal of Personality*, 4(1), 17–24.
<https://doi.org/10.1111/j.1467-6494.1935.tb02022.x>
- Tankard, M. E., & Paluck, E. L. (2016). Norm perception as a vehicle for social change. *Social Issues and Policy Review*, 10(1), 181–211. <https://doi.org/10.1111/sipr.12022>
- Tappin, B. M., McKay, R. T., & Abrams, D. (2017). Choosing the right level of analysis: Stereotypes shape social reality via collective action. *Behavioral and Brain Sciences*, 40, e13. <https://doi.org/10.1017/S0140525X1500240X>
- Todorov, A., & Oosterhof, N. N. (2011). Modeling social perception of faces. *Signal Processing Magazine, IEEE*, 28, 117–122. <https://doi.org/10.1109/MSP.2010.940006>
- Tosh, C., Greengard, P., Goodrich, B., Gelman, A., & Hsu, D. (2020). *The piranha problem: Large effects swimming in a small pond*.
- Treitler, V. B. (2013). *The ethnic project: Transforming racial fiction into ethnic factions*. Stanford University Press.
- Tucker, I. (2012). Deleuze, sense, and life: Marking the parameters of a psychology of individuation. *Theory & Psychology*, 22(6), 771–785.
<https://doi.org/10.1177/0959354312442787>
- Varsanyi, M. W., Lewis, P. G., Provine, D. M., & Decker, S. (2012). A Multilayered Jurisdictional Patchwork: Immigration Federalism in the United States: A

- MULTILAYERED JURISDICTIONAL PATCHWORK. *Law & Policy*, 34(2), 138–158.
<https://doi.org/10.1111/j.1467-9930.2011.00356.x>
- Wang, D., & Liu, S. (2021). Doing Ethnography on Social Media: A Methodological Reflection on the Study of Online Groups in China. *Qualitative Inquiry*, 107780042110146.
<https://doi.org/10.1177/10778004211014610>
- Watts, S., & Stenner, P. (2005). Doing Q methodology: Theory, method and interpretation. *Qualitative Research in Psychology*, 2(1), 67–91.
<https://doi.org/10.1191/1478088705qp022oa>
- Wong, C. J. (2007). “Little” and “Big” Pictures in Our Heads: Race, Local Context, and Innumeracy About Racial Groups in the United States. *Public Opinion Quarterly*, 71(3), 392–412. <https://doi.org/10.1093/poq/nfm023>
- Wright, M. N., & Ziegler, A. (2017). ranger: A Fast Implementation of Random Forests for High Dimensional Data in C++ and R. *Journal of Statistical Software*, 77(1).
<https://doi.org/10.18637/jss.v077.i01>
- Young, A. I., Ratner, K. G., & Fazio, R. H. (2014). Political Attitudes Bias the Mental Representation of a Presidential Candidate’s Face. *Psychological Science*, 25(2), 503–510. <https://doi.org/10.1177/0956797613510717>
- Zaller, J. (1992). *The nature and origins of mass opinion*. Cambridge University Press.

Mapping varied mental representations of illegalized immigrants

Supplementary Material

Table of Contents:

- Table S1: Classification image generators' characteristics
- Table S2: Classification image raters' characteristics
- Figure S1: Google trends for labels across state and years
- Figure S2: Geographic map of sample based on IP addresses
- Figure S3: Correlation matrix of participant variables
- Figure S4: Agglomerative coefficient for different hierarchical clustering algorithms
- Figure S5: Maximum and minimum cluster size based on number of clusters and clustering algorithm
- Figure S6: Total within sums of squares (TWSS) across N cluster solutions
- Figure S7: Ten cluster solution from the hierarchical clustering analysis
- Figure S8: Permuted variable importance for ten cluster solution
- Figure S9: Partial dependence plots for ten cluster solution
- Figure S10: Fifteen cluster solution from the hierarchical clustering analysis
- Figure S11: Permuted variable importance for fifteen cluster solution
- Figure S12: Partial dependence plot for fifteen cluster solution
- Figure S13: Validation of random forest analysis and individual difference measures. The variable importance for individual difference predictors of support for building the Mexico-U.S. wall.
- Figure S14: Validation of random forest analysis and individual difference measures. Probabilistic responses of whether one agrees or disagrees strongly with building a Mexico-U.S. border wall across 2016 vote, political orientation, and political party.

Supplementary Table S1. Classification image generators' characteristics.

	California Illegal (N=250)	California Undocu (N=250)	Texas Illegal (N=252)	Texas Undocu (N=250)	Overall (N=1002)
Age					
Mean (SD)	37.4 (11.2)	38.7 (12.8)	38.3 (12.3)	37.7 (11.6)	38.0 (12.0)
Median [Min, Max]	35.0 [19.0, 72.0]	36.0 [18.0, 105]	35.0 [20.0, 73.0]	35.0 [19.0, 73.0]	35.0 [18.0, 105]
Gender Identity					
Agender	0 (0%)	0 (0%)	1 (0.4%)	0 (0%)	1 (0.1%)
Woman	132 (52.8%)	124 (49.6%)	144 (57.1%)	140 (56.0%)	540 (53.9%)
Man	114 (45.6%)	123 (49.2%)	106 (42.1%)	108 (43.2%)	451 (45.0%)
My gender is not listed	1 (0.4%)	0 (0%)	0 (0%)	0 (0%)	1 (0.1%)
Nonbinary	3 (1.2%)	3 (1.2%)	1 (0.4%)	2 (0.8%)	9 (0.9%)
Ethnoracial Background					
Black or African	27 (10.8%)	9 (3.6%)	43 (17.1%)	31 (12.4%)	110 (11.0%)
East Asian	32 (12.8%)	40 (16.0%)	9 (3.6%)	14 (5.6%)	95 (9.5%)
Hawaiian-Pacific Islander	4 (1.6%)	3 (1.2%)	0 (0%)	2 (0.8%)	9 (0.9%)
Indigenous-Native American	3 (1.2%)	0 (0%)	0 (0%)	0 (0%)	3 (0.3%)
Latinx-o-a or Hispanic	30 (12.0%)	46 (18.4%)	44 (17.5%)	40 (16.0%)	160 (16.0%)
Middle Eastern	0 (0%)	0 (0%)	1 (0.4%)	0 (0%)	1 (0.1%)
My race-ethnicity is not listed	12 (4.8%)	4 (1.6%)	2 (0.8%)	6 (2.4%)	24 (2.4%)
South Asian-Indian	8 (3.2%)	7 (2.8%)	5 (2.0%)	7 (2.8%)	27 (2.7%)
White or European	134 (53.6%)	141 (56.4%)	148 (58.7%)	150 (60.0%)	573 (57.2%)
Education Level					
Bachelors degree or similar	126 (50.4%)	117 (46.8%)	116 (46.0%)	104 (41.6%)	463 (46.2%)
High school or baccalaureate or A-levels	74 (29.6%)	67 (26.8%)	78 (31.0%)	91 (36.4%)	310 (30.9%)
I did not complete secondary- high school	0 (0%)	2 (0.8%)	1 (0.4%)	1 (0.4%)	4 (0.4%)
Masters or Doctoral degree	31 (12.4%)	40 (16.0%)	23 (9.1%)	31 (12.4%)	125 (12.5%)
Professional qualification	19 (7.6%)	24 (9.6%)	34 (13.5%)	23 (9.2%)	100 (10.0%)
Religious Affiliation					

	California Illegal (N=250)	California Undocu (N=250)	Texas Illegal (N=252)	Texas Undocu (N=250)	Overall (N=1002)
Agnostic	43 (17.2%)	42 (16.8%)	25 (9.9%)	33 (13.2%)	143 (14.3%)
Atheist	32 (12.8%)	32 (12.8%)	39 (15.5%)	24 (9.6%)	127 (12.7%)
Buddhist	4 (1.6%)	8 (3.2%)	3 (1.2%)	7 (2.8%)	22 (2.2%)
Hindu	2 (0.8%)	2 (0.8%)	2 (0.8%)	3 (1.2%)	9 (0.9%)
Jewish	2 (0.8%)	8 (3.2%)	3 (1.2%)	4 (1.6%)	17 (1.7%)
Mormon	3 (1.2%)	3 (1.2%)	1 (0.4%)	0 (0%)	7 (0.7%)
Muslim	4 (1.6%)	3 (1.2%)	4 (1.6%)	1 (0.4%)	12 (1.2%)
My religion is not listed	19 (7.6%)	11 (4.4%)	22 (8.7%)	31 (12.4%)	83 (8.3%)
Nothing in particular	53 (21.2%)	40 (16.0%)	41 (16.3%)	48 (19.2%)	182 (18.2%)
Orthodox such as Greek or Russian Orthodox	3 (1.2%)	0 (0%)	1 (0.4%)	0 (0%)	4 (0.4%)
Protestant	48 (19.2%)	49 (19.6%)	60 (23.8%)	55 (22.0%)	212 (21.2%)
Roman Catholic	37 (14.8%)	52 (20.8%)	51 (20.2%)	44 (17.6%)	184 (18.4%)
Born in U.S.					
No	19 (7.6%)	23 (9.2%)	18 (7.1%)	14 (5.6%)	74 (7.4%)
Yes	231 (92.4%)	227 (90.8%)	234 (92.9%)	236 (94.4%)	928 (92.6%)
U.S. Citizen					
No	2 (0.8%)	4 (1.6%)	6 (2.4%)	6 (2.4%)	18 (1.8%)
Prefer not to say	3 (1.2%)	1 (0.4%)	0 (0%)	0 (0%)	4 (0.4%)
Yes	245 (98.0%)	245 (98.0%)	246 (97.6%)	244 (97.6%)	980 (97.8%)
Immigrant Contact					
Mean (SD)	3.32 (1.88)	3.32 (1.85)	2.91 (1.89)	3.18 (1.95)	3.18 (1.90)
Median [Min, Max]	4.00 [0, 5.00]	4.00 [0, 5.00]	3.00 [0, 5.00]	4.00 [0, 5.00]	3.00 [0, 5.00]
Estimated immigrant pop. in neighborhood					
Mean (SD)	12.7 (14.0)	12.3 (13.6)	12.5 (15.1)	12.4 (15.3)	12.5 (14.5)
Median [Min, Max]	9.00 [0, 60.0]	6.00 [0, 60.0]	5.00 [0, 60.0]	5.00 [0, 60.0]	5.00 [0, 60.0]
Estimated immigrant pop. in city					
Mean (SD)	18.8 (14.8)	18.3 (14.3)	19.3 (16.3)	19.7 (15.5)	19.0 (15.2)

	California Illegal (N=250)	California Undocu (N=250)	Texas Illegal (N=252)	Texas Undocu (N=250)	Overall (N=1002)
Median [Min, Max]	15.0 [0, 60.0]	15.0 [1.00, 60.0]	15.0 [0, 60.0]	16.0 [0, 60.0]	15.0 [0, 60.0]
Estimated immigrant pop. in state					
Mean (SD)	24.3 (14.8)	23.4 (13.6)	24.4 (15.3)	26.3 (16.3)	24.6 (15.1)
Median [Min, Max]	22.0 [0, 60.0]	21.0 [2.00, 60.0]	21.0 [0, 60.0]	21.0 [2.00, 60.0]	21.0 [0, 60.0]
Illegal Immigrant Scale					
Mean (SD)	3.00 (0.297)	2.99 (0.307)	2.99 (0.315)	3.00 (0.300)	3.00 (0.304)
Median [Min, Max]	3.00 [2.05, 4.35]	3.00 [2.15, 4.20]	3.00 [1.80, 3.90]	3.05 [2.00, 3.85]	3.00 [1.80, 4.35]
Conservatism					
Mean (SD)	58.0 (17.5)	59.2 (17.2)	59.9 (17.7)	61.1 (17.0)	59.5 (17.4)
Median [Min, Max]	58.8 [14.8, 95.8]	60.0 [10.9, 98.3]	61.3 [8.25, 95.8]	63.2 [16.0, 95.4]	60.5 [8.25, 98.3]
Build the wall?					
Missing response	3 (1.2%)	2 (0.8%)	5 (2.0%)	1 (0.4%)	11 (1.1%)
Agree	24 (9.6%)	28 (11.2%)	33 (13.1%)	41 (16.4%)	126 (12.6%)
Agree Strongly	35 (14.0%)	41 (16.4%)	44 (17.5%)	41 (16.4%)	161 (16.1%)
Disagree	49 (19.6%)	42 (16.8%)	42 (16.7%)	44 (17.6%)	177 (17.7%)
Disagree Strongly	112 (44.8%)	106 (42.4%)	97 (38.5%)	90 (36.0%)	405 (40.4%)
Undecided	27 (10.8%)	31 (12.4%)	31 (12.3%)	33 (13.2%)	122 (12.2%)
Vote in 2016					
Chose not to vote	48 (19.2%)	45 (18.0%)	59 (23.4%)	63 (25.2%)	215 (21.5%)
Could not vote	9 (3.6%)	12 (4.8%)	17 (6.7%)	14 (5.6%)	52 (5.2%)
Donald Trump & Mike Pence, Republican Party	46 (18.4%)	54 (21.6%)	71 (28.2%)	71 (28.4%)	242 (24.2%)
Gary Johnson & Bill Weld, Libertarian Party	5 (2.0%)	18 (7.2%)	5 (2.0%)	7 (2.8%)	35 (3.5%)
Hillary Clinton & Tim Kaine, Democratic Party	118 (47.2%)	108 (43.2%)	89 (35.3%)	80 (32.0%)	395 (39.4%)
Jill Stein & Ajamu Baraka, Green Party	8 (3.2%)	5 (2.0%)	3 (1.2%)	5 (2.0%)	21 (2.1%)

	California Illegal (N=250)	California Undocu (N=250)	Texas Illegal (N=252)	Texas Undocu (N=250)	Overall (N=1002)
Other Candidate	16 (6.4%)	8 (3.2%)	8 (3.2%)	10 (4.0%)	42 (4.2%)
Political Party Affiliation					
Democrat	125 (50.0%)	109 (43.6%)	99 (39.3%)	90 (36.0%)	423 (42.2%)
Green	4 (1.6%)	3 (1.2%)	1 (0.4%)	0 (0%)	8 (0.8%)
Independent	54 (21.6%)	53 (21.2%)	53 (21.0%)	62 (24.8%)	222 (22.2%)
Libertarian	5 (2.0%)	9 (3.6%)	5 (2.0%)	6 (2.4%)	25 (2.5%)
Republican	39 (15.6%)	56 (22.4%)	67 (26.6%)	59 (23.6%)	221 (22.1%)
Unaffiliated/Not political	23 (9.2%)	20 (8.0%)	27 (10.7%)	33 (13.2%)	103 (10.3%)
Self-report state					
Arizona	0 (0%)	1 (0.4%)	3 (1.2%)	0 (0%)	4 (0.4%)
Arkansas	1 (0.4%)	0 (0%)	0 (0%)	0 (0%)	1 (0.1%)
California	247 (98.8%)	248 (99.2%)	1 (0.4%)	0 (0%)	496 (49.5%)
Colorado	1 (0.4%)	0 (0%)	0 (0%)	0 (0%)	1 (0.1%)
Illinois	0 (0%)	0 (0%)	1 (0.4%)	0 (0%)	1 (0.1%)
Iowa	1 (0.4%)	0 (0%)	0 (0%)	0 (0%)	1 (0.1%)
Louisiana	0 (0%)	0 (0%)	1 (0.4%)	0 (0%)	1 (0.1%)
Missouri	0 (0%)	0 (0%)	1 (0.4%)	1 (0.4%)	2 (0.2%)
New York	0 (0%)	1 (0.4%)	0 (0%)	0 (0%)	1 (0.1%)
Oklahoma	0 (0%)	0 (0%)	0 (0%)	1 (0.4%)	1 (0.1%)
Texas	0 (0%)	0 (0%)	244 (96.8%)	248 (99.2%)	492 (49.1%)
Washington D.C.	0 (0%)	0 (0%)	1 (0.4%)	0 (0%)	1 (0.1%)
State from IP address					
Arizona	0 (0%)	1 (0.4%)	0 (0%)	0 (0%)	1 (0.1%)
Arkansas	0 (0%)	0 (0%)	1 (0.4%)	0 (0%)	1 (0.1%)
California	248 (99.2%)	246 (98.4%)	0 (0%)	0 (0%)	494 (49.3%)
New Mexico	0 (0%)	0 (0%)	0 (0%)	1 (0.4%)	1 (0.1%)
North Carolina	0 (0%)	1 (0.4%)	0 (0%)	0 (0%)	1 (0.1%)
Texas	0 (0%)	0 (0%)	247 (98.0%)	247 (98.8%)	494 (49.3%)
Missing	2 (0.8%)	2 (0.8%)	4 (1.6%)	2 (0.8%)	10 (1.0%)

	California Illegal (N=250)	California Undocu (N=250)	Texas Illegal (N=252)	Texas Undocu (N=250)	Overall (N=1002)
How long lived in state?					
Less than 1 month	1 (0.4%)	0 (0%)	3 (1.2%)	2 (0.8%)	6 (0.6%)
Less than 1 year	0 (0%)	2 (0.8%)	4 (1.6%)	1 (0.4%)	7 (0.7%)
Over 1 year	20 (8.0%)	22 (8.8%)	24 (9.5%)	21 (8.4%)	87 (8.7%)
Over 10 years	229 (91.6%)	226 (90.4%)	221 (87.7%)	226 (90.4%)	902 (90.0%)

Table S2: Classification image raters' characteristics.

	Dangerous (N=46)	Dominant (N=44)	American (N=46)	Ethnoracial Category (N=87)	Overall (N=223)
<i>Age</i>					
Mean (SD)	38.3 (10.4)	36.9 (12.0)	40.2 (12.9)	38.1 (12.4)	38.3 (12.0)
Median [Min, Max]	36.0 [25.0, 69.0]	33.5 [20.0, 72.0]	36.5 [20.0, 73.0]	35.0 [21.0, 68.0]	35.0 [20.0, 73.0]
<i>Gender</i>					
Man	32 (69.6%)	25 (56.8%)	24 (52.2%)	54 (62.1%)	135 (60.5%)
Woman	14 (30.4%)	19 (43.2%)	22 (47.8%)	33 (37.9%)	88 (39.5%)
<i>Race</i>					
Black or African	4 (8.7%)	12 (27.3%)	5 (10.9%)	12 (13.8%)	33 (14.8%)
East Asian	3 (6.5%)	2 (4.5%)	3 (6.5%)	3 (3.4%)	11 (4.9%)
Latinx/o/a or Hispanic	3 (6.5%)	2 (4.5%)	1 (2.2%)	4 (4.6%)	10 (4.5%)
White or European	36 (78.3%)	27 (61.4%)	36 (78.3%)	63 (72.4%)	162 (72.6%)
Indigenous/Native American	0 (0%)	1 (2.3%)	0 (0%)	0 (0%)	1 (0.4%)
South Asian/Indian	0 (0%)	0 (0%)	1 (2.2%)	3 (3.4%)	4 (1.8%)
My race/ethnicity is not listed	0 (0%)	0 (0%)	0 (0%)	2 (2.3%)	2 (0.9%)
<i>Born in U.S.</i>					
No	1 (2.2%)	2 (4.5%)	1 (2.2%)	5 (5.7%)	9 (4.0%)
Yes	45 (97.8%)	42 (95.5%)	45 (97.8%)	82 (94.3%)	214 (96.0%)
<i>Social and Economic Conservatism Scale</i>					
Mean (SD)	58.4 (19.7)	60.3 (19.8)	58.4 (19.4)	58.8 (21.6)	58.9 (20.3)

	Dangerous (N=46)	Dominant (N=44)	American (N=46)	Ethnoracial Category (N=87)	Overall (N=223)
Median [Min, Max]	59.0 [19.1, 94.9]	64.5 [16.0, 90.4]	54.8 [15.4, 91.6]	59.0 [15.5, 100]	60.5 [15.4, 100]
<i>Test-Retest Reliability</i>					
Mean (SD)	0.352 (0.213)	0.296 (0.242)	0.373 (0.201)	0.471 (0.233)	0.391 (0.233)
Median [Min, Max]	0.380 [0.00673, 0.752]	0.250 [0.0108, 0.780]	0.363 [0.0548, 0.849]	0.501 [0.0135, 0.894]	0.407 [0.00673, 0.894]

Supplementary Figure S1. Google trends for labels across state and years. Each plot represents a different state that was under consideration in the design of this study before settling on Texas and California. States are color coded by zones considered to have different relationships to migration based on distance to a border or historical migration patterns. The solid line is for “illegal immigrant” searches, the dotted line is for “undocumented immigrant” searches. Missing lines suggest data were unavailable for those time periods or states. The x-axis is year (2006 to 2019) in which the label searches occurred and the y-axis is a normalized value that represents the number of searches as a proportion of all searches in the same time period and region. It can be interpreted as public interest.

Texas Sample

California Sample

Supplementary Figure S2. Geographic map of sample based on IP addresses. Left panel represents participants from Texas, right panel from California. The location of the dots represents various cities and the color of the dot is the number of participants in each city where darker colors represent more participants. Texas had more representation from metropolitan areas: Houston, Dallas-Fort Worth, Austin, Brownsville. California had more representation from the Bay Area (e.g., San Francisco) and southern Californian areas around Los Angeles and San Diego.

Supplementary Figure S3. Correlation matrix of participant variables. Cooler colors represent positively correlated variables, hotter colors represent negatively correlated variables. Variable order from top to bottom (or left to right): time on task, age, immigrant contact, social and economic conservatism, estimated immigrant population size in neighborhood, in city, in state, illegal immigrant scale average score. Note: “size” is the estimated size of the surrounding undocumented population, “hood” is neighborhood, “secs” is the political orientation scale, “iiasvg” is the illegal immigration attitude scale, and “totalTime” was the completion duration for the task.

Supplementary Figure S4. Agglomerative coefficient for different hierarchical clustering algorithms. The top panels and bottom left panel are the dendrograms for the various similarity linkage algorithms. The bottom right panel depicts the agglomerative coefficient for the various linkage algorithms. There is a clearer cluster structure in the Ward dendrogram than the other linkages, as reflected in the higher agglomerative coefficient.

Supplementary Figure S5. Maximum and minimum cluster size based on number of clusters and clustering algorithm. The left panel shows the minimum sample size across the clusters at various N cluster solutions. The colors represent the linkage algorithms. The y-axis is the sample size of the smallest cluster at each solution. The right panel depicts the largest or maximum sample size at each cluster solution. The 5-15 range which was chosen for the study is highlighted in grey.

Supplementary Figure S6. Total within sums of squares (TWSS) across N cluster solutions. The left panel depicts the TWSS values (y-axis) across different number of cluster solutions (x-axis). The right panel zooms in on the 5-15 cluster range. There is a leveling off at the 7 cluster area. After 15 clusters, there is a decrease. However, interpreting more than 15 face clusters is very difficult, so we focused on the restricted range of 5-15 clusters.

Supplementary Figure S7. Ten cluster solution from the hierarchical clustering analysis. A) Dendrogram depicting the ten clusters in the same colors as Figure 2 (e.g., cluster 1 is broken down into clusters 1, 6, 8, 9 all in red). B) The average faces and sample sizes from each cluster (color coded to match the dendrogram). C) The estimated marginal means from the trait ratings for each face. The error bars depict 95% confidence intervals. The horizontal line is the neutral scale position of each rating. The percentages reflect the top two ethnoracial classifications.

Supplementary Figure S8. Permuted variable importance for ten cluster solution. The x-axis is the permuted p value representing variable importance (significant alpha is .05). The exact p value is provided to the right of each bar. The y-axis lists all the individual difference measures.

Supplementary Figure S9. Partial dependence plots for ten cluster solution. A) Probabilistic cluster membership across undocumented population size in the neighborhood. B) Estimates across age. The y-axis is the probability of cluster membership. The black tick marks above the x-axis are data points such that the less tick marks the less participants at that level of the x-axis variable.

Supplementary Figure S10. Fifteen cluster solution from the hierarchical clustering analysis. A) Dendrogram depicting the fifteen clusters in the same colors as Figure 2 (e.g., Cluster 1 is broken down into clusters 1, 15, 7, 13, 9, 11 all in red). B) The average faces and sample sizes from each cluster (color coded to match the dendrogram). C) The estimated marginal means from the trait ratings for each face. The error bars depict 95% confidence intervals. The horizontal line is the neutral scale position of each rating. The percentages reflect the top two ethnoracial classifications.

Supplementary Figure S11. Permuted variable importance for fifteen cluster solution. The x-axis is the permuted p value representing variable importance (significant alpha is .05). The exact p value is provided to the right of each bar. The y-axis lists all the individual difference measures.

Supplementary Figure S12. Partial dependence plot for fifteen cluster solution. Probabilistic cluster membership across age. The y-axis is the probability of cluster membership. The black tick marks above the x-axis are data points such that the less tick marks the less participants at that level of the x-axis variable.

Supplementary Figure S13. Validation of random forest analysis and individual difference measures. As an initial step for validation, we calculated the variable importance for individual difference predictors of support for building the Mexico-U.S. wall. The x-axis is the permuted p value representing variable importance (significant alpha is .05). The exact p value is provided to the right of each bar. The y-axis lists all the individual difference measures. The results show that political variables (vote in 2016, social and economic conservatism, and political party affiliation) were the most important predictors of support for building the wall as could be expected given the politicized nature of Donald Trump’s anti-immigrant proposal. **Supplementary Figure S14** shows how these three variables relate to support.

Supplementary Figure S14. Validation of random forest analysis and individual difference measures. Probabilistic responses of whether one agrees or disagrees strongly with building a Mexico-U.S. border wall across 2016 vote (top left panel), political orientation (top right panel), and political party (bottom right panel). The y-axis is the response probability. The x-axes are various presidential candidates (top left panel), political orientation scores (top right panel), and political parties (bottom right panel). For the 2016 vote, those who voted for more liberal or “other” candidates mostly disagreed strongly with building the wall. More conservative or libertarian voters and those who did not vote were more split in their support, yet overall were more likely to support, especially Donald Trump voters. For the social and economic conservatism scale, those who scored higher than 50 (i.e., more conservative) were split in their support, but overall were more likely to agree strongly than individuals who scored below 50 (i.e., more liberal). The more liberal respondents predominantly disagreed strongly. For political party affiliation, respondents from most of the party options were split in their support, except for Democrats who mainly disagreed strongly. Overall, there was less support for the wall, however, higher levels of support occurred for respondents from the Republican and Libertarian parties and who were unaffiliated. These patterns make theoretical sense and show that these techniques can capture meaningful relationships.