Brain & Language 135 (2014) 66–72

Contents lists available at ScienceDirect

Brain & Language

journal homepage: www.elsevier.com/locate/b&l

Short Communication

The role of language in multi-dimensional categorization: Evidence from transcranial direct current stimulation and exposure to verbal labels

Lynn K. Perry [a,b,*], Gary Lupyan [a]

[a] Department of Psychology, University of Wisconsin-Madison, United States
[b] DeLTA Center, University of Iowa, United States

ARTICLE INFO

Article history:
Accepted 26 May 2014
Available online 27 June 2014

Keywords:
Labeling
Categorization
Selective representation
Transcranial direct current stimulation (tDCS)
Wernicke's area

ABSTRACT

Human concepts differ in their dimensionality. Some, like GREEN-THINGS, require representing one dimension while abstracting over many others. Others, like BIRD, have higher dimensionality due to numerous category-relevant properties (feathers, two-legs). Converging evidence points to the importance of verbal labels for forming low-dimensional categories. We examined the role of verbal labels in categorization by (1) using transcranial direct current stimulation over Wernicke's area (2) providing explicit verbal labels during a category learning task. We trained participants on a novel perceptual categorization task in which categories could be distinguished by either a uni- or bi-dimensional criterion. Cathodal stimulation over Wernicke's area reduced reliance on single-dimensional solutions, while presenting informationally redundant novel labels reduced reliance on the dimension that is normally incidental in the real world. These results provide further evidence that implicit and explicit verbal labels support the process of human categorization.

1. Introduction

Because no two experiences are truly identical, using past knowledge to respond appropriately to present events requires forming categories of like things that can be treated equivalently (Murphy, 2002). To determine which things are alike we must selectively represent category-relevant properties and abstract across irrelevant ones. Importantly, the ratio of relevant to irrelevant properties differs for different categories. For example, consider the category GREEN THINGS which includes items like limes and grasshoppers while excluding closely related items like lemons and locusts. This category requires selectively representing color while excluding shape, taste, etc. The category BIRDS on the other hand requires simultaneously representing multiple features (e.g., feathers, wings)—none individually necessary or sufficient for membership. Human concepts can be placed on a continuum from low-dimensional (e.g., GREEN THINGS) to high-dimensional (e.g., BIRDS) (Lupyan, Mirman, Hamilton, & Thompson-Schill, 2012; Pothos, 2005).

It has been previously noted that high-dimensional categories (alternatively called information integration, Ashby & Maddox, 2011; or similarity-based, Sloutsky, 2010) are easier to learn for young children (Kloos & Sloutsky, 2008) and non-human primates to (Couchman, Coutinho, & Smith, 2010) than low-dimensional categories (alternatively called rule-based, Ashby & Maddox, 2011; or selection based, Sloutsky, 2010). When a stimulus space is structured ambiguously, children and non-human primates tend to partition it using multiple dimensions suggesting high-dimensional categorization is a kind of default (Couchman et al., 2010; Smith & Kemler, 1977). In contrast, adults, have little trouble forming low-dimensional categories.[1] Not only do humans overcome the apparent default of high-dimensional categorization, but given the choice, older children and adults show strong preferences for low-dimensional solutions (Couchman et al., 2010; Smith & Kemler, 1977).

1.1. Effects of language on categorization

What enables older children and adults to do easily what is so challenging to young children and non-human animals? One possibility is that low-dimensional categorization is aided by

---

* Corresponding author. Address: 1202 W. Johnson St., Madison, WI 53706, United States.

E-mail address: lkperry@wisc.edu (L.K. Perry).

[1] Although, when categorizing items into formal categories such as ODD NUMBERS and TRIANGLES—perhaps the most low-dimensional categories of all—adults never fully abstract from putatively irrelevant information (Lupyan, 2013).

http://dx.doi.org/10.1016/j.bandl.2014.05.005
0093-934X/

language. Indeed, Ashby and colleagues have noted that an effective strategy for learning low-dimensional categories is to verbalize a rule (e.g., green goes here, blue there). Such approaches are not feasible for high-dimensional categories if only because criteria for membership cannot be easily verbalized.

Additional support for the involvement of language in low-dimensional categorization comes from findings that children can learn low-dimensional categories at an earlier age if they are given category labels (Perry & Samuelson, 2013) or verbal instructions about category-relevant features (Kloos & Sloutsky, 2008). Conversely, disrupting language in adults through verbal interference (Lupyan, 2009), or more drastically, stroke-related aphasia (Lupyan & Mirman, 2013) impairs low- but not high-dimensional categorization.

A useful framework for understanding *why* labeling supports low-dimensional categorization is the Label Feedback Hypothesis (Lupyan, 2012): in associating a category name (i.e., a verbal label) with multiple exemplars, the label becomes most strongly associated with features that are most predictive/diagnostic of the category thereby facilitating selective activation of those features while simultaneously abstracting over irrelevant ones. Support for this hypothesis comes from findings that labels facilitate category learning (Lupyan, Rakison, & McClelland, 2007) and lead to faster object recognition (Lupyan & Thompson-Schill, 2012).

1.2. Rationale and predictions

Insofar as implicit and explicit labeling supports adults' low-dimensional categorization abilities, decreasing the extent to which labels are activated may decrease the likelihood that people form low-dimensional categories. One way to study the role of labels in categorization is to manipulate the ease with which participants can use labels and observe the outcome of this manipulation on categorization. For example, if the word "green" supports selective representation of a grasshopper's color, then interfering with activation of the label should disrupt the speed or accuracy with which, for example, people group grasshoppers with limes. The method often used for *down-regulating* the labeling process—verbal interference—has a number of shortcomings, (see Perry & Lupyan, 2013) some of which can be overcome through use of noninvasive cortical stimulation.

In a previous study, Lupyan et al. (2012) examined how transcranial direct current stimulation (tDCS) applied over left inferior frontal gyrus (subsuming Broca's area) affects on categorization. They found down-regulating activity over Broca's area decreased accuracy in low-dimensional, but not high-dimensional categorization. However, because Broca's area has been associated with both linguistic processes such as speech production (Gernsbacher & Kaschak, 2003) and domain-general cognitive control (Kan & Thompson-Schill, 2004), it is difficult to draw conclusions about the role of language in categorization.

To assess more directly the relationship between labeling and categorization, here we stimulate BA 22—posterior superior temporal gyrus (subsuming Wernicke's area).[2] The involvement of Wernicke's area in lexical and phonological processes is well known (e.g., Binder et al., 1997; Geschwind, 1970; Price, 2000). Modulation of Wernicke's area using tDCS has been previously shown to affect name-learning (Flöel, Rösser, Michka, Knecht, & Breitenstein, 2008) and picture–word verification (Lupyan, in preparation). This cortical region, however, has not been previously implicated in domain-general cognitive control (Cole & Schneider, 2007). Finding that tDCS over Wernicke's area can affect nonverbal categorization—specifically low-dimensional categorization—would support the hypothesis that language is involved in the ability to form object representations that emphasize task-relevant dimensions.

We predicted that stimulating over Wernicke's area should, by down-regulating the labeling process, nudge people to represent stimuli in a higher-dimensional way than they would normally. We also attempt to up-regulate the labeling process through a behavioral manipulation by providing learners with novel redundant category labels (see Lupyan et al., 2007) with the expectation that these should nudge people to represent stimuli in a lower-dimensional way than they would otherwise.

2. Experiment 1: Modulating labeling processes in categorization

To examine the relationship between labeling and selective representation of category-relevant features we trained participants to discriminate between two types of "minerals"—some nutritious and some poisonous. The minerals comprised gabor patches varying in orientation and spatial frequency. The categories were structured such that using a uni-dimensional (either orientation or frequency) or bi-dimensional boundary (co-occurrence of *both* orientation and frequency) would lead to approximately equal accuracy (Fig. 1). This configuration allowed us to distinguish effects on overall accuracy from effects on the dimensionality of learned categories. Participants were assigned to one of four conditions: (1) Cathodal stimulation over Wernicke's area, (2) Control cathodal stimulation over the vertex, (3) No-stimulation group receiving redundant labels following each categorization trial (see Section 7), and (4) a no-stimulation baseline group.

3. Results and discussion

We first assessed performance by comparing accuracy and response times (RT) for the four conditions (Wernicke's-cathodal stimulation, vertex-cathodal stimulation, label, baseline). Next,

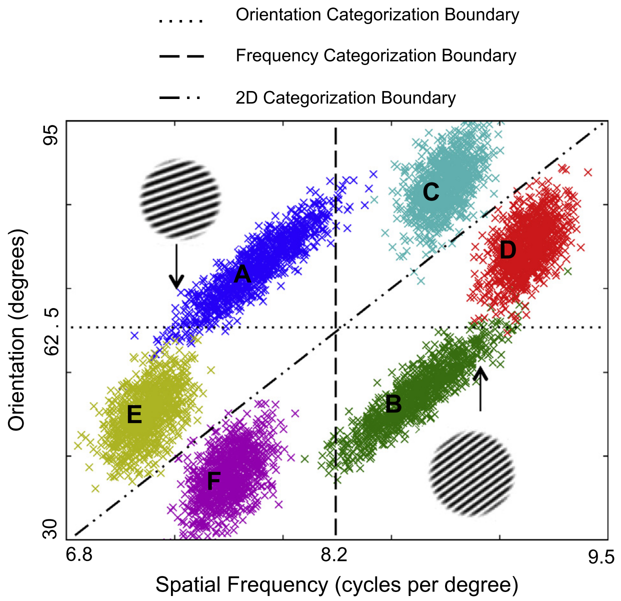

Fig. 1. Distribution of stimuli. Training stimuli was drawn from (A and B); two sample gabor patches from each distribution are shown. Generalization stimuli were drawn from (C–F). Lines denote potential category boundaries. Colors are used for visualization purposes only. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

---

[2] Much remains unknown about the functional extent of tDCS-induced changes to cortical excitability. Phrasing such as "stimulation over Wernicke's area" should therefore be taken to mean that we stimulated over the anatomical region corresponding to Wernicke's area (pSTG), not that the functional effects of stimulation were circumscribed strictly to Wernicke's area.

we quantified dimensionality of categorization solutions by determining for each subject to what extent their categorization was predicted by stimulus orientation, frequency, or both. Analyses were conducted using the lme4 package in R. Significance tests were calculated using chi-square tests that compared fit of mixed-effect models with and without the factor of interest on improvement in model fit.

3.1. Categorization performance

Classification accuracy across training blocks for the four tested conditions is shown in Fig. 2A: Overall accuracy, in decreasing order was $M_{\text{Wernicke's-cathodal-stimulation}} = .86$, $M_{\text{baseline}} = .82$, $M_{\text{label}} = .81$, $M_{\text{vertex-cathodal-stimulation}} = .73$. A comparison of logistic-regression predicting accuracy from block showed that although participants in all conditions learned over time, $p \ll .001$, condition was a significant predictor, $X^2(3) = 12.45$, $p = .006$, as was the interaction between block and condition, $X^2(3) = 7.68$, $p = .05$. Planned comparisons revealed this effect was due to lower performance by participants in the vertex-cathodal-simulation condition who had significantly lower accuracy than the remaining three groups, $.0004 < p < .05$, and learned more slowly than those in the baseline and label conditions, $.009 < p < .05$.

An analysis of RTs revealed a congruent pattern: $M_{\text{Wernicke's-cathodal-stimulation}} = 489$ ms, $M_{\text{baseline}} = 496$ ms, $M_{\text{label}} = 493$ ms, $M_{\text{vertex-cathodal-stimulation}} = 574$ ms. A comparison of mixed-effects models showed a marginal overall effect of condition on RTs, $X^2(3) = 6.24$, $p = .10$. Planned comparisons showed that participants in the vertex-cathodal-stimulation condition responded more slowly than those in the Wernicke's-cathodal-stimulation condition, $p = .03$, and baseline $p = .04$, and marginally more slowly than those in the label condition, $p = .08$. No other differences were significant.

Together, these analyses show that neither labeling nor cathodal stimulation over Wernicke's area affected overall accuracy or RTs relative to baseline. We discuss possible reasons for the effects of vertex stimulation on overall accuracy and RTs in Section 6.

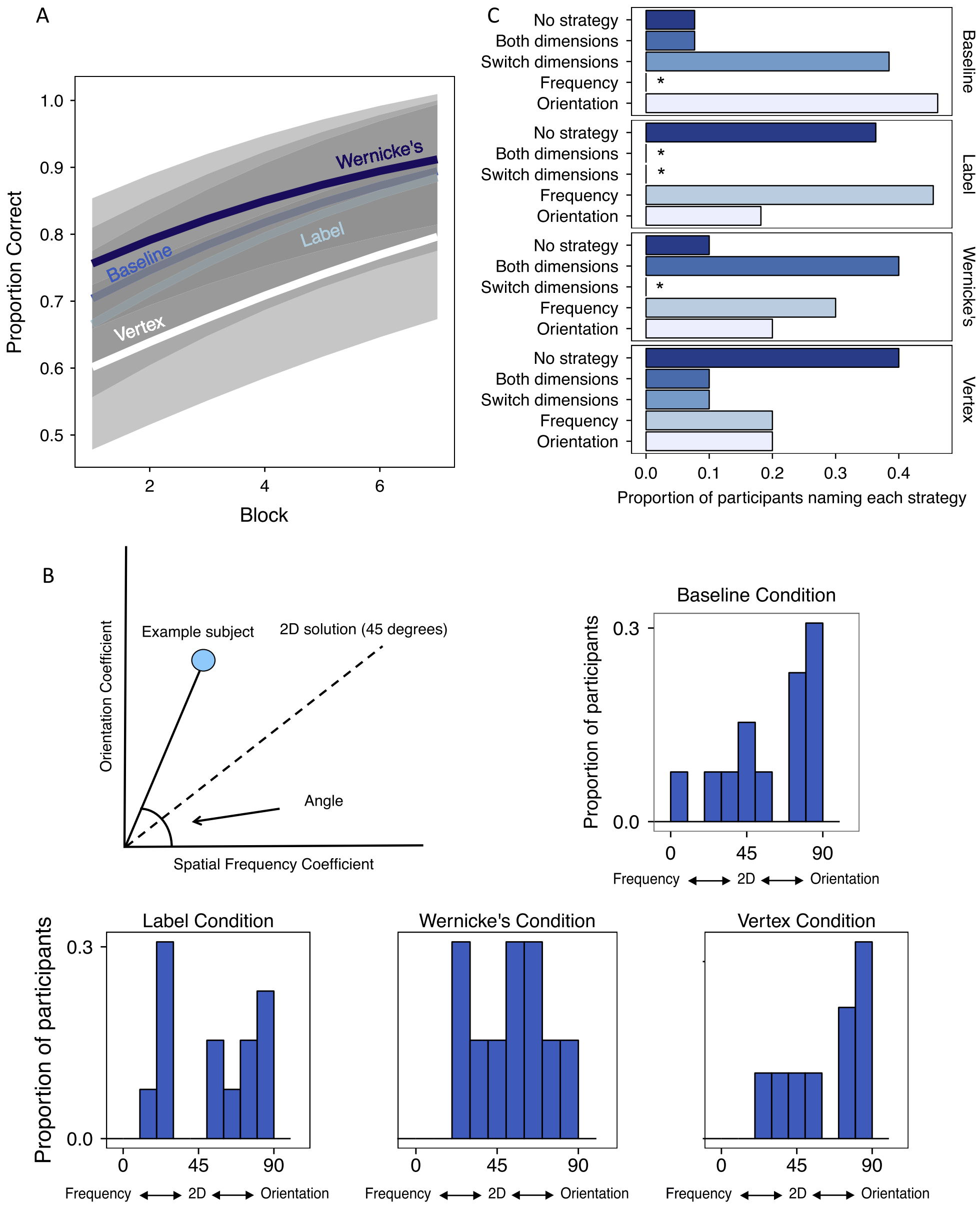

Fig. 2. (A) Predicted training accuracy for each condition based on logistic mixed regression (error bands represent standard error of the mean). (B) Schematic for computing categorization solution angle for each participant and histogram of categorization solutions. *X*-axis shows angle computed using method shown in schematic. (C) Reported categorization strategies for each condition. Asterisks (*) denote strategies never used by participants in a condition.

3.2. Categorization dimensionality

Our main prediction concerned effects of labeling on the dimensionality of the categorization solution. To quantify dimensionality, we fit each participant's responses to a logistic mixed regression model predicting classification (poisonous/nutritious) from the frequency and orientation of each gabor patch. These values were standardized (z-scored) to place them on the same scale. We then fit two models to each participant's responses: The uni-dimensional model included only the dimension most predictive for that participant. The bi-dimensional model included both dimensions. The two models were compared using Akaike Information Criterion (AIC). To determine if stimulation or labels affected whether participants' categorization decisions were better fit to the bi-dimensional model than uni-dimensional model, we compared linear regression models with/without condition, revealing a significant effect of condition on dimensionality, $F(3,41) = 3.05$, $p = .04$. Planned comparisons showed participants' decisions in the Wernicke's-cathodal-stimulation condition were significantly better fit to the bi-dimensional model than uni-dimensional model compared to participants in the vertex-cathodal-stimulation condition, $F(1,17) = 4.76$, $p = .04$; $b = -79.04$, 95% CI[-146.05, -12.03], label condition, $F(1,21) = 5.37$, $p = .03$; $b = -77.97$, 95% CI[-143.90, -12.04], or baseline, $F(1,21) = 5.12$, $p = .03$; $b = -77.09$, 95% CI[-138.44, -15.74]. There were no other significant differences.

To visualize differences in dimensionality we used coefficients of each predictor from each participant's best-fit model as coordinates in two-dimensional stimulus space (Fig. 2B). Someone relying only on frequency would fall on a 0° (horizontal) line. Someone relying only on orientation would fall on a 90° (vertical) line. Someone relying equally on both dimensions would fall on a 45° (diagonal) line. As shown in Fig. 2B, people in the baseline and vertex-cathodal-stimulation condition tended to rely on orientation while those in the label condition tended to rely on frequency. Participants in the Wernicke's-cathodal-stimulation condition were clustered around 45° indicating more bi-dimensional solutions.[3]

As evident from Fig. 2B, although individuals in the label condition relied on a single dimension to approximately the same extent as participants in the baseline and vertex-cathodal-stimulation conditions, the label participants appeared to rely more on frequency than orientation. A comparison of logistic mixed regression models with/without the interaction between condition (baseline, label) and the spatial-frequency value of each stimulus as a predictor revealed a significant effect on categorization, $X^2(1) = 11.48$, $p = .0007$, $b = -.15$, 95% CI[$-.23$, $-.07$].

We additionally queried participants about what strategy they used in performing the categorization task. The results are shown in Fig. 2C. Participants in the vertex-cathodal-stimulation condition were significantly more likely than baseline to report not using a strategy, $b = .32$, 95% CI[.01, .63], $z = 2.05$, $p < .05$. Participants in the Wernicke's-cathodal-stimulation condition and label condition demonstrated more consistent use of a single strategy. In particular, Wernicke's-cathodal-stimulation condition participants most frequently mentioned using both dimensions, and label condition participants were most likely to mention frequency. Logistic regression models revealed Wernicke's-cathodal-stimulation condition participants were more likely to mention using both dimensions simultaneously than baseline, $b = .32$, 95% CI[.05, .59], $z = 2.36$, $p < .05$; label condition participants were more likely to mention frequency than baseline, $b = .45$, 95% CI[.14, .76], $z = 2.77$, $p < .01$. These reports were consistent with our findings that people undergoing cathodal stimulation over Wernicke's area generally used both dimensions while those provided with labels used spatial-frequency.

Our results show that cathodal stimulation over Wernicke's area led to an increase in high-dimensional categorization. Importantly, it is not the case that participants in the Wernicke's-cathodal-stimulation condition were less likely to use orientation information than those in the baseline condition. Rather, the results of participants' performance on generalization trials (where stimuli crossed the training boundary) and the results of participants' self-report of strategy use both suggest that instead of using a single dimension, these participants were using the co-occurrence of *both* orientation and frequency information. We hypothesize this effect is due to stimulation disrupting the normally automatic labeling that participants engage in during the category-learning task.[4] Disrupting this process partially disrupts formation of low-dimensional categories. Including explicit category names did not lead to more uni-dimensionality, but given large reliance on a single dimension at baseline that we observed, this is not especially surprising. More surprising was the finding that including redundant category names caused participants to shift from orientation to frequency. We probe this result further in the next experiment.

4. Experiment 2: Effects of verbal labels on categorization biases

Why should labeling categories lead people to home in on frequency rather than orientation? Consider that for most real objects, orientation is incidental, i.e., irrelevant for category-membership: a cup is a cup regardless of orientation. If labels help to selectively represent dimensions most relevant for category-membership, one might expect that simply labeling a category causes learners to attend to dimensions that are *typically* diagnostic. In effect, labeling the "minerals" should lead them to inherit the diagnosticity biases present in the real world (cf. Brojde, Porter, & Colunga, 2011).

To test this hypothesis we presented participants with 8 gabor stimuli used in the training study, with the two mineral types grouped separately (Fig. 3A). In the label condition minerals were accompanied by labels (see Section 7). Participants were asked to describe how the groups differed, with the expectation that they would primarily list orientation and frequency. The critical question was which was listed first. We predicted that simply including labels would reduce reliance on orientation due to labels helping the ambiguous "minerals" to inherit real-world category biases. As a further test of this prediction, we ran a condition in which gabors were referred to as "flowers" and surrounded by schematic flower petals (Fig. 3B) reasoning that this context should reduce reliance on orientation even without labels. A rose at 45° is still a rose.

5. Results and discussion

When describing the differences between the two groups of "minerals", virtually all participants mentioned orientation as the primary distinguishing dimension ($M_{\text{mineral-no-label}} = .98$). Simply presenting novel labels alongside the groups decreased reliance on orientation ($M_{\text{mineral-label}} = .74$) and increased reliance on frequency, see Fig. 3C. Placing the very same gabor stimuli into a more explicit object context by surrounding them with "petals" and

---

[3] Fig. 2B is for visualization purposes only. There was insufficient power to perform a formal comparison of distributions.

[4] Flöel et al. (2008) found that compared to sham stimulation, anodal, but not cathodal stimulation over Wernicke's area affected learning of new names. In contrast, our findings demonstrate effects of cathodal stimulation on naming. Clearly, interpreting functional effects of tDCS is more complex than positing that the cathodal stimulation leads to down-regulation and anodal to up-regulation. Indeed, which polarity leads to behavioral change may be strongly task-dependent.

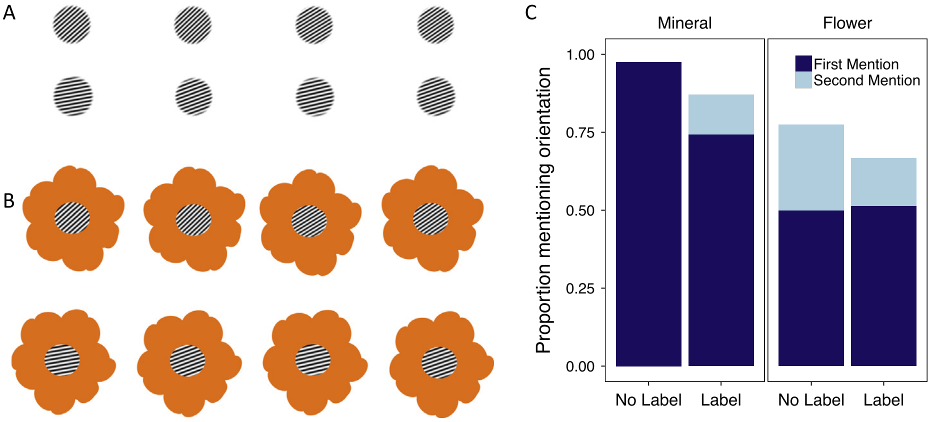

Fig. 3. Stimuli used in mineral (A) and flower (B) conditions of the incidental versus non-incidental dimension comparison study, and proportion of participants mentioning orientation in each condition (C).

calling them flowers led to a further drop in salience of orientation. In this more object-like context, using labels did not further decrease the likelihood of mentioning orientation ($M_{\text{flower-label}} = .51$; $M_{\text{flower-no-label}} = .50$). The significance of the above-mentioned effects was tested through logistic mixed regression including context (mineral, flower), label (label, no label), and time (first, second mention). Participants in the mineral condition were more likely to mention differences in orientation than those in the flower condition, $b = 8.95$, 95% CI[7.99, 9.91], $z = 4.01$, $p < .0001$, and the difference in mentioning orientation between label and no label mineral conditions was larger than in the flower conditions,[5] $b = -6.71$, 95% CI[-9.00, -4.41], $z = -2.66$, $p < .01$.

Even casually mentioned labels decreased the salience of orientation—a dimension typically irrelevant to category membership. Importantly, differences in frequency were quite subtle (see Fig. 3). That the mere presence of a label could de-emphasize the otherwise quite salient differences in orientation speaks to the power of this simple manipulation. To clarify: We believe that the decreased likelihood of mentioning orientation differences in the presence of labels is not due to something special about the relationship between labels and orientation. Rather, knowledge that contrastive labels tend not to be used for objects differing in orientation appears to carry over into novel category learning, leading the minerals, when accompanied by labels, to inherit biases of real-world categories. When this real-world bias is supported by visual context, as in the flower condition, reliance on orientation is further reduced.

6. General discussion

Our main goal was to examine the relationship between labeling and selective representation of dimensions in category learning. We hypothesized that if labeling facilitates selective representation, cathodal stimulation over Wernicke's area should cause participants to form higher-dimensional (here, bi-dimensional) representations. Conversely, explicit labeling of the two categories may cause participants to rely more on a uni-dimensional solution. Our results confirmed this prediction: cathodal stimulation over Wernicke's area, hypothesized to down-regulate activation of labels, led to greater reliance on bi-dimensional solutions. This result speaks to the role of labels in categorization, suggesting that "normal" categorization, especially of low-dimensional categories, may be continuously augmented by linguistic processes (Lupyan, 2012).

An intriguing secondary finding was that although explicit labeling did not increase uni-dimensional categorization (which were already uni-dimensional at baseline), the inclusion of redundant labels changed *which* dimension participants relied on both in the category learning task and in the comparison task in which participants simply had to describe differences between the simultaneously-displayed categories. The simple addition of category names to the objects reduced the salience of orientation— a dimension normally incidental to named categories. Thus, although labeling supports selectively representing category-relevant information—necessary for low-dimensional categories—labeling comes with its own set of priors such that assigning even uninformative labels to novel categories may lead to abstracting over dimensions that are not generally diagnostic of real-world category distinctions.

Although low-dimensional categorization has been previously linked to easily verbalizable rules (Ashby & Maddox, 2011), it was previously unknown how labeling affected the number of and which particular dimensions were used in categorization. Thus, our work is an important step in understanding processes underlying people's ability to selectively represent task-relevant information and abstract over irrelevant information.

6.1. Dimensional biases

Despite orientation having minimal use for distinguishing real-world categories, we know that people are quite capable of learning to use it when required by the task (e.g., Helie & Ashby, 2012). In our particular stimuli, orientation turned out to be visually salient, a likely reason why so many participants defaulted to it in the category-learning experiment and the subsequent comparison experiment. Simply calling the gabors "minerals" is probably insufficient for the stimuli to inherit diagnosticity biases of real-world minerals (especially given participants' likely poor knowledge of what distinguishes minerals). Consequently, we believe placing the stimuli in a richer visual context—adding (category-irrelevant) flower petals, rather than calling them flowers, decreased reliance on orientation in this condition.

6.2. Specificity of stimulation

Although tDCS is not as spatially precise as noninvasive brain stimulation techniques such as TMS, it does show a surprising

---

[5] To rule out potential low-level confounds caused by viewing gabors on a colored background in the flower condition, we repeated the mineral condition with identically colored orange rectangles behind the gabors using 60 new participants recruited from Mechanical Turk. The results were similar to the original study: those in the mineral condition were more likely to mention orientation than those in the flower condition, $b = 3.00$, $z = 2.03$, $p < .05$, and the difference between the no label and label mineral conditions was larger than in the flower condition, $b = 2.52$, $z = 1.87$, $p = .06$.

degree of localization as confirmed by concurrent fMRI (Holland et al., 2011). Our findings that stimulating Wernicke's area and the vertex led to distinct patterns of results suggests our stimulation had some degree of localization, although we cannot be certain of exact precision. For example, our stimulation over BA 22 may have also stimulated regions of the inferior parietal lobule such as the supramarginal gyrus (BA 40) and the angular gyrus (BA 39), implicated in linguistic (word comprehension, reading) and nonlinguistic (spatial cognition) functions (e.g., Seghier, 2013).

Overall classification speed and accuracy was decreased by cathodal stimulation over the vertex (Cz). One possible explanation for this finding is that vertex stimulation perturbed activity in the supplementary motor area (SMA). Vollmann et al. (2013) found that anodal stimulation of SMA improved visuo-motor learning. It is conceivable that if our cathodal stimulation affected the SMA, it may have disrupted some aspect of associating visual forms with motor responses. Importantly, this effect was general and did not interact with our primary measure of interest—categorization dimensionality.

Although cathodal stimulation is often presumed to be inhibitory (e.g., Nitsche et al., 2003), and our behavioral data support this interpretation, this may not always be the case (Batsikadze, Moliadze, Paulus, Kuo, & Nitsche, 2013). The role of cathodal tDCS in perturbing neural activity and behavioral performance is still being explored (Nozari, Woodard, & Thompson-Schill, 2014).

6.3. Beyond categorization

A final consideration is whether in addition to supporting low-dimensional categorization labeling might support selective representation more generally. A growing body of evidence suggests there is a connection between cognitive control processes, such as inhibitory control and selective attention, and labeling. Verbal interference (thought to disrupt labeling) disrupts task switching—an ability dependent on cognitive control (Emerson & Miyake, 2003), while explicit labeling facilitates children's ability to switch rules in the dimensional change card sort task (Kirkham, Cruess, & Diamond, 2003), suggesting the ability to flexibly switch sorting dimensions may be partially driven by implicit labeling. However, the extent to which labeling plays a causal role in cognitive control remains unclear. To understand how people selectively represent task-relevant attributes to the exclusion of other information, future research will need to further assess the directionality of the relationship between labeling and cognitive control.

6.4. Conclusions

Our results demonstrate that labeling plays an important role in supporting selective representation category-relevant dimensions and in biasing the specific dimensions used to categorize. Language (particularly the process of verbal labeling) more than simply *describing* differences, appears to actively help to represent task-relevant distinctions.

7. Methods

7.1. Categorization study

7.1.1. Participants

We randomly assigned 57 participants (35 females; average age: 19 years) to one of four conditions: baseline = 19; Wernicke's-cathodal-stimulation = 13; vertex-cathodal-stimulation = 10; label = 15. Exclusion criteria included history of neurologic or psychiatric disease, use of anti-convulsants, anti-psychotic, or sedative medications. The task proved unexpectedly difficult for some participants: 12 did not exceed chance performance (6, 3, 1, 2 from each condition, respectively) and were removed from analyses. 10 of these were nonnative English speakers, and potentially had difficulty comprehending task instructions. Participants received course credit for participating.

7.1.2. Stimuli

Stimuli were circular gabor patches varying in orientation and spatial frequency (Fig. 1). Values for each stimulus were selected by random sampling of a bivariate distribution (similar to Helie and Ashby, 2012) drawn from a space varying in orientation from 30° to 95° and frequency from 6.8 to 9.5 cycles/deg. Stimuli were centered on the screen and subtended ~1.5° of visual angle. The training category structures (A and B in Fig. 1) were constructed such that using either a fully uni-dimensional (vertical or horizontal boundary) or bi-dimensional (diagonal boundary) solution would lead to roughly equivalent accuracy. To determine which dimensions each individual was using, we included generalization stimuli at key positions (C–F in Fig. 1).

7.1.3. Experimental procedure

The procedure was similar to that used by Lupyan et al. (2007). Participants were told to imagine they were discovering minerals: some were edible and should be approached, some poisonous and should be avoided. On each trial, participants saw one mineral in the center of the screen and an explorer, either left, right, above, or below the mineral. After 400 ms the stimulus disappeared and participants responded by pressing an arrow key to approach/avoid the mineral. On training trials, participants heard feedback (bell/buzzer) after each response. On generalization trials there was no feedback. In the training trials of the label condition, participants heard novel labels ("leebish", "grecious"). The labels were redundant in that they were presented *after* feedback. Participants completed seven blocks of 40 training trials (20 per category) and eight generalization trials (two per generalization region; see Fig. 2) per block. After the experiment, participants were asked via questionnaire, "What strategy/rule did you use to complete the task? Did your strategy change over time?" These open-ended responses were coded as orientation, frequency, both, switching between strategies, or having no strategy, by a coder blind to condition.

7.1.4. tDCS procedure

tDCS was delivered by a battery-driven constant direct-current stimulator (Soterix 1 × 1 Low Intensity Stimulator). Rubber electrodes were inserted into saline-soaked 5 × 7 cm sponges. Placement of the cathodal electrode was made by reference to the 10–20 system: intersection of T5-C3 and T3-P3 for the Wernicke's-stimulation condition (posterior region of BA 22), selected due to belief that this area subsumes Wernicke's area and is implicated in comprehension and production of labels; and at Cz for vertex (Homan et al., 1987). The anodal electrode was attached to the right cheek. Current was increased over 30 s to 1.75 mA, and then the task began. Stimulation lasted 20 min, the approximate task length.

7.2. Incidental versus non-incidental dimension comparison study

7.2.1. Participants

120 adults (71 female) were recruited from Amazon Mechanical Turk and were assigned to one of four between-subjects conditions in a 2 × 2 design (label /no label; mineral/flower). Age information was not collected from these participants.

7.2.2. Stimuli

In all conditions, stimuli comprised eight patches, four from each of the 2 categories used in the categorization study. In the flower condition, each gabor was surrounded by schematic flower petals (see Fig. 3). Petals were rotated differently for each stimulus to avoid drawing extra attention to the orientation of gabor lines. As in the categorization study, the categories differed only in orientation and frequency.

7.2.3. Procedure

Participants saw stimuli arranged in 2 rows, each corresponding to a category. In the no-label condition each row was labeled by the line "Here is the [first/second] group of [minerals/flowers]." In the label condition the line labeling each row included the phrase "Let's call them the [leebish/grecious] group." Participants were asked to write the first and then second difference they noticed between the groups. The open-ended responses were coded as orientation, frequency, or miscellaneous by a coder blind to condition.

Acknowledgments

The study was partially supported by NSF #1331293 to G.L. We thank Anqi Fu, Ishaan Guptasarma, Shirley Hu, Martin Potter, and Jesse Sherman for help with data collection.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.bandl.2014.05.005.

References

Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annals of the New York Academy of Sciences, 1224*(1), 147–161.

Batsikadze, G., Moliadze, V., Paulus, W., Kuo, M.-F., & Nitsche, M. A. (2013). Partially non-linear stimulation intensity-dependent effects of direct current stimulation on motor cortex excitability in humans. *The Journal of Physiology, 591*(Pt 7), 1987–2000.

Binder, J. R., Frost, J. A., Hammeke, T. A., Cox, R. W., Rao, S. M., & Prieto, T. (1997). Human brain language areas identified by functional magnetic resonance imaging. *The Journal of Neuroscience, 17*(1), 353–362.

Brojde, C. L., Porter, C., & Colunga, E. (2011). Words can slow down category learning. *Psychonomic Bulletin & Review, 18*, 798–804.

Cole, M. W., & Schneider, W. (2007). The cognitive control network: Integrated cortical regions with dissociable functions. *NeuroImage, 37*(1), 343–360.

Couchman, J. J., Coutinho, M. V. C., & Smith, J. D. (2010). Rules and resemblance: Their changing balance in the category learning of humans (*Homo sapiens*) and monkeys (*Macaca mulatta*). *Journal of Experimental Psychology: Animal Behavior Processes, 36*(2), 172–183.

Emerson, M. J., & Miyake, A. (2003). The role of inner speech in task switching: A dual-task investigation. *Journal of Memory and Language, 48*(1), 148–168.

Flöel, A., Rösser, N., Michka, O., Knecht, S., & Breitenstein, C. (2008). Noninvasive brain stimulation improves language learning. *Journal of Cognitive Neuroscience, 20*(8), 1415–1422.

Gernsbacher, M. A., & Kaschak, M. P. (2003). Neuroimaging studies of language production and comprehension. *Annual Review of Psychology, 54*, 91–114.

Geschwind, N. (1970). The organization of language and the brain. *Science, 170*, 940–944.

Helie, S., & Ashby, F. G. (2012). Learning and transfer of category knowledge in an indirect categorization task. *Psychological Research, 76*(3), 292–303.

Holland, R., Leff, A. P., Josephs, O., Galea, J. M., Desikan, M., Price, C. J., et al. (2011). Speech facilitation by left inferior frontal cortex stimulation. *Current Biology, 21*(16), 1403–1407.

Homan, R. W., Herman, J., & Purdy, P. (1987). Cerebral location of international 10–20 system electrode placement. *Electroencephalography and Clinical Neurophysiology, 66*(4), 376–382.

Kan, I. P., & Thompson-Schill, S. L. (2004). Selection from perceptual and conceptual representations. *Cognitive, Affective, & Behavioral Neuroscience, 4*(4), 466–482.

Kirkham, N. Z., Cruess, L., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science, 6*(5), 449–467.

Kloos, H., & Sloutsky, V. M. (2008). What's behind different kinds of kinds: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General, 137*(1), 52–72.

Lupyan, G. (2009). Extracommunicative functions of language: Verbal interference causes selective categorization impairments. *Psychonomic Bulletin & Review, 16*(4), 711–718.

Lupyan, G. (2012). Linguistically modulated perception and cognition: The label-feedback hypothesis. *Frontiers in Cognition, 3*, 54.

Lupyan, G. (2013). The difficulties of executing simple algorithms: Why brains make mistakes computers don't. *Cognition, 129*(3), 615–636.

Lupyan, G. (in preparation). Facilitation and impairment of lexical activation by noninvasive stimulation of two classical language areas.

Lupyan, G., & Mirman, D. (2013). Linking language and categorization: Evidence from aphasia. *Cortex, 49*(5), 1187–1194.

Lupyan, G., Mirman, D., Hamilton, R., & Thompson-Schill, S. L. (2012). Categorization is modulated by transcranial direct current stimulation over left prefrontal cortex. *Cognition, 124*(1), 36–49.

Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking redundant labels facilitate learning of novel categories. *Psychological Science, 18*(12), 1077–1083.

Lupyan, G., & Thompson-Schill, S. L. (2012). The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology: General, 141*(1), 170–186.

Murphy, G. L (2002). *Big book of concepts*. The MIT Press.

Nitsche, M. A., Nitsche, M. S., Klein, C. C., Tergau, F., Rothwell, J. C., & Paulus, W. (2003). Level of action of cathodal DC polarisation induced inhibition of the human motor cortex. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology, 114*(4), 600–604.

Nozari, N., Woodard, K., & Thompson-Schill, S. L. (2014). Consequences of cathodal stimulation for behavior: When does it help and when does it hurt performance? *PLoS ONE, 9*(1), e84338.

Perry, L. K., & Samuelson, L. K. (2013). The role of verbal labels in attention to dimensional similarity. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuch (Eds.), *Proceedings of the thirty-fifth annual conference of the Cognitive Science Society*.

Perry, L. K., & Lupyan, G. (2013). What the online manipulation of linguistic activity can tell us about language and thought. *Frontiers in Behavioral Neuroscience, 7*, 122.

Pothos, E. M. (2005). The rules versus similarity distinction. *Behavioral and Brain Sciences, 28*(01), 1–14.

Price, C. J. (2000). The anatomy of language: Contributions from functional neuroimaging. *Journal of Anatomy, 197*(Pt 3), 335–359.

Seghier, M. L. (2013). The angular gyrus: Multiple functions and multiple subdivisions. *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry, 19*(1), 43–61.

Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science, 34*(7), 1244–1286.

Smith, L. B., & Kemler, D. G. (1977). Developmental trends in free classification: Evidence for a new conceptualization of perceptual development. *Journal of Experimental Child Psychology, 24*(2), 279–298.

Vollmann, H., Conde, V., Sewerin, S., Taubert, M., Sehm, B., Witte, O. W., et al. (2013). Anodal transcranial direct current stimulation (tDCS) over supplementary motor area (SMA) but not pre-SMA promotes short-term visuomotor learning. *Brain Stimulation, 6*(2), 101–107.