CrossMark

BRIEF REPORT

Garner interference is not solely driven by stimulus uncertainty

Devin M. Burns[1]

Abstract The speeded classification tasks popularized by Garner (1974), with their accompanying labels of *separable* and *integral* dimensions, have become the dominant paradigm for characterizing perceptual interactions between stimulus dimensions. Separable dimensions like color and shape can be selectively attended to at will, but integral dimensions like hue and brightness cannot. When classifying stimuli with respect to a single dimension, integral stimuli lead to faster performance in the baseline block, where only that single dimension varies, than in the filtering block, in which a second dimension is allowed to vary irrelevantly. No such difference is predicted for separable dimensions lead to no difference. The comparison between baseline and filtering confounds a change in the number of stimuli with a change in the number of variant dimensions. A new experimental condition is proposed utilizing three stimulus dimensions to hold the number of stimuli constant while introducing variance along a second irrelevant dimension. Data indicate that interference increases, suggesting that changes in the number of potential stimuli within a block are not necessary for Garner interference, in contrast to the accounts provided by extant quantitative models.

Keywords Garner interference · Separable · Integral · EBRW · Tectonic

✉ Devin M. Burns
  burnsde@mst.edu

[1] Missouri University of Science and Technology, Rolla MO, USA

Introduction

One of the most popular ways for determining the perceptual interactions between stimulus dimensions, especially with regard to attentional concerns, is with the separable/integral dichotomy established by Garner (1974). Although much work has been done to classify various dimensions in this way, there is relatively little agreement about the internal processing characteristics that give rise to this distinctions.

Separability is often taken to signify that the stimulus dimensions are being processed independently (e.g., Bartlett et al. 2003; Ganel & Goshen-Gottstein, 2002; Kaufmann and Schweinberger, 2004), but definitions of processing independence are themselves often contested, and their relationship to Garner interference is unclear (Fitousi & Wenger, 2013). A variety of cognitive models have been applied to the Garner paradigm, and most attribute the crucial reaction time differences to changes in the number of potential stimuli used in each condition and the change in uncertainty that results. This work suggests that these changes are not necessary for interference to occur.

The Garner paradigm

The Garner paradigm uses two stimulus dimensions having two levels each, which are factorially combined to produce four stimuli, as in Fig. 1a (for now ignore the third dimension, saturation). In this example, brightness and hue are crossed to yield light-purple, dark-purple, light-blue, and dark-blue. The traditional paradigm compares performance in three conditions: *baseline*, *filtering*, and *correlated*. All
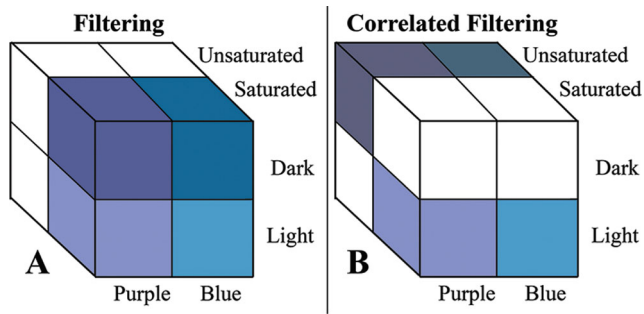
Fig. 1 Stimuli for the standard Garner filtering task (**a**) and the newly proposed correlated filtering task (**b**). The three stimulus dimensions are hue (*horizontal*), brightness (*vertical*), and saturation (*depth*). *White cubes* represent stimuli not used in a particular block of trials

conditions have the same task: the observer is shown one stimulus in each trial and asked to rapidly classify it with respect to a single dimension, for example to decide if it is blue or purple. Here, hue is the relevant dimension, but what distinguishes the three conditions are changes in the other dimension, brightness.

In the *baseline* (sometimes called *control*) condition, the second dimension is held constant. For a given relevant dimension, there are two such blocks of trials, with participants distinguishing a light-purple from a light-blue in one block, and then deciding between dark-purple and dark-blue in the other. In the *filtering* (sometimes called *orthogonal*) block, the second dimension is instead allowed to vary from trial to trial, with all four stimuli appearing equally often. This variation is irrelevant, in that stimuli will sometimes be light or dark, but participants are instructed to respond only with respect to hue. A third, *correlated* (sometimes called *redundant*) condition was originally included where values of the two dimensions are perfectly correlated within a block of trials, and then correlated the other way in another. In one block participants would distinguish light-blue from dark-purple, and in another dark-blue from light-purple. Because this condition encourages violations of selective attention, it is often excluded in recent applications (e.g., Atkinson et al. 2005; Ashby & Maddox, 1994) and will not be discussed further in this paper.

Once data is collected, the two dimensions (hue and brightness) are declared separable if reaction times for the conditions are equal. If the observer is capable of selectively attending to hue, variation in brightness should have no impact. If, however, selective attention is compromised in some way, variation in the irrelevant dimension could serve only to distract participants, leading to slower responses. Thus, the mean reaction time for filtering is predicted to be greater than baseline, showing what has come to be known

as *Garner interference*. This is interpreted to signal dimensions to which the observer cannot selectively attend, and are labeled integral.

Modeling interference

There are two key differences between the filtering and baseline blocks: the number of variant dimensions and the number of stimuli. The baseline condition is considered to be a single-dimensional condition, since only one dimension varies within the block of trials. In contrast, the filtering condition has two dimensions that vary, although only one of them is necessary for selecting a response. Allowing variation within a block of trials along this irrelevant dimension is frequently interpreted as the cause of slower performance in the filtering condition (Amishav & Kimchi, 2010; Ganel et al., 2005; Garner, 1976).

In sharp contrast, many of the formal models of Garner interference blame the slowdown on the number of stimuli. One of the earliest attempts at eliciting Garner interference from a processing model was by Ashby and Maddox (1994), who supplemented General Recognition Theory (Ashby & Townsend, 1986) with their RT-distance hypothesis to enable reaction time predictions from a previously accuracy-only model. In order to predict slower responses in the filtering condition, they offered several suggestions (Ashby & Maddox, 1994, p. 452). Their primary one was that the greater number of stimuli used in the filtering block would increase stimulus uncertainty, which would then increase variances for the perceptual distributions. The intuition behind this assumption is that if your expectations are less precise about the upcoming stimulus, you will perceive it less accurately. This increase in variance leads to slower average performance according to their model.

Another formal language that has been used for describing Garner effects, called *tectonic theory*, was laid out by Melara and Algom (2003). The primary application of this theory was to explain Stroop effects, but its application to Garner interference was also detailed. In this theory, the success of selective attention depends on two constructs: dimensional imbalance and dimensional uncertainty.

Dimensional imbalance deals with the relative salience between the two dimensions, and can be operationalized as the difference in mean RT between the baseline conditions for the two dimensions (e.g., participants are faster at categorizing hue than brightness when the other dimension is constant). Earlier work has demonstrated that Garner interference is sensitive to dimensional imbalance: when one dimension is more salient than the other the measure is unreliable (Melara & Mounts, 1993). Intuitively, if change in one

dimension is much more difficult to detect, the data may appear as though participants can pay perfect attention to the more salient dimension, regardless of whether the two dimensions are truly integral or separable.

Dimensional uncertainty is a combination of two separate effects: average uncertainty, which is simply a measure of how often a certain stimulus appears, and conditional uncertainty, which measures how the two dimensions are correlated. In the Garnerian context, the former depends only on the number of stimuli, and is therefore greater for filtering. Conditional uncertainty is constant for filtering and baseline, as dimensional values are uncorrelated.

In both of these models, slower performance in the filtering condition is ascribed not to the addition of variation along the second dimension per se, but rather the increased uncertainty that comes from having additional stimuli within a block. A third formal model used to describe Garner interference, the Exemplar Based Random Walk (EBRW) model (Nosofsky & Palmeri, 1997), also predicts that the larger number of stimuli used in the filtering block will directly impede performance by decreasing the strength of the memory trace for each particular exemplar. However, in contrast to the other two models, a second source of interference arises from the similarities between stimuli in a way that is not solely dependent on the number of stimuli.

In the EBRW model, performance is positively correlated with within-category similarity: participants do better when stimuli assigned to the same response are tightly clustered. In contrast, performance is negatively correlated with between-category similarity: participants fare worse when the two response categories are more confusable. The filtering condition has the same between-category similarity as the baseline condition (blue and purple are always equally confusable), but introduces within-category differences (changes in brightness) that cannot exist when only two stimuli are used.

Although this change in within-category similarity is clearly the result of the additional stimuli in the filtering condition, it is a separate effect that is not explicitly dependent on the number of stimuli. To see this clearly, we can imagine a filtering condition in which distances along the irrelevant dimension (brightness) are infinitesimally small. Although there are four separate stimuli, the EBRW model would regard this condition as being identical to the baseline condition, since within-category similarity approaches 100 %.

As already mentioned, the researchers aiming to use these tests to diagnose the relationship between two perceptual dimensions often believe it is variation along the second dimension itself that causes interference. If interference were simply caused by uncertainty, it would appear to tell us little about any particular pairing of dimensions! A fundamental weakness of the Garner paradigm is that the test for Garner interference confounds these two distinct changes: the filtering condition has both an additional variant dimension and two additional stimuli as compared to the baseline condition.

Stimuli or dimensions?

Researchers have tried several different tactics to empirically separate the effects of number of stimuli and number of variant dimensions. One approach has been to look at sequential trial-to-trial effects. In the filtering condition there are four different possible relations between a given stimulus and its predecessor: both dimensional values are repeated from the previous trial, both are changed, the relevant dimension is repeated but the irrelevant dimension changes, or the relevant dimension changes while the irrelevant dimension is repeated. In the baseline condition, however, the value of the irrelevant dimension is always repeated from trial to trial.

Dyson and Quinlan (2010) noticed that the traditional measure of Garner interference, the degree to which average performance in filtering is slower than average performance in the baseline condition, could be divided into two separate comparisons defined in terms of stimulus sequence. The first measures the effect of trial-to-trial irrelevant variation by comparing those filtering trials where the irrelevant dimension changed to those filtering trials in which it repeated. Slower performance for the former should directly implicate irrelevant variation instead of uncertainty, since both trial types are drawn from the filtering condition. Their second measure compares those filtering trials in which the value of the irrelevant dimension was repeated to trials from the baseline condition, where it is always repeated. Slower performance for the former here is said to implicate the change in the number of stimuli. The sum of these two measures is mathematically identical to the standard test for interference, and the authors argued that decomposing it in this way allows them to analytically separate the effects of irrelevant variation and stimulus uncertainty.

While this line of investigation is helpful in pointing out that trial-to-trial effects could be an important and often overlooked component of Garner interference, it does not necessarily succeed in the goal of separating the effects of stimulus uncertainty and irrelevant variation. Their measure of stimulus uncertainty, which compares the filtering and baseline conditions using only those trials where the irrelevant dimension was repeated, may still be influenced by

irrelevant variation. Just because the irrelevant dimension did not change from one trial to the next does not mean that it is having no effect. Many processing models predict that stimulus dimensions are only processed when they vary within a block, so the filtering condition engenders very different processing than the baseline condition, regardless of the stimulus sequence. Any effects of irrelevant variation that apply at this block level are thus being confounded with any effects due to the number of stimuli within a block.

Rather than looking to sequential analyses of data obtained from the typical Garner experimental design, other researchers have attempted to separate these two effects by using different stimulus sets. At first blush, it may seem that one could solve this problem simply by increasing the number of stimuli while holding the number of dimensions constant, but this gives rise to new confounds. Melara and Mounts (1994) created additional stimuli for the filtering condition by maintaining the two levels of the relevant dimension (hue in our example) and picking additional levels for the irrelevant dimension (brightness) that were evenly spaced between the former levels (i.e., if filtering used brightness values of 3 and 6, a new condition would also have stimuli with values 4 and 5). They showed that incorporating such stimuli caused a substantial decrease in Garner interference, meaning that filtering performance improved back toward the level of the baseline condition.

This result may seem counterintuitive, since all three models under consideration predicted a decrease in performance with increasing numbers of stimuli, but consider that adding additional levels along the irrelevant dimension increases within-category similarity by decreasing the average trial-to-trial change in that dimension, reducing its salience as compared to the relevant dimension. One would expect that if we instead chose new levels of brightness outside the range of our previous stimuli (e.g., brightness values of 1 and 8), this would increase the amount of interference by increasing the average trial-to-trial change in brightness, thus lowering the within-category similarity.

A different way to create new stimuli without increasing the number of dimensions would be to pick different levels of the relevant dimension rather than the irrelevant dimension. Even more clearly than the previous example, however, this introduces additional confounds due to the variable distances from the decision boundary that these new stimuli would have: those that are close to the bound would be more confusable and therefore slower, while those far from the boundary would be classified quickly and easily. Thus, new stimuli cannot be created by choosing new levels of either dimension without introducing new confounds. Creating new stimuli via any other change, such as altering saturation, would then be introducing another dimension and thus fail to achieve our goal of deconfounding these two changes.

If we cannot increase the number of stimuli without increasing the number of variant dimensions (or introducing confounds in terms of stimulus discriminability), what about increasing the number of dimensions without changing the number of stimuli? Expanding the stimulus space to include a third dimension opens up this possibility, as seen in Fig. 1b.

By correlating the values of two irrelevant dimensions, a new condition called *correlated filtering* can be created, which has the same number of stimuli as the standard filtering condition, but yet has irrelevant variation along two dimensions rather than one. If Garner interference is caused solely by the number of stimuli used in a given block, as models have represented, then correlated filtering should be no slower than standard filtering. If, however, Garner interference depends on the number of dimensions that vary within a given block of trials, as many applications have interpreted, irrelevant variation along a second dimension should produce additional interference.

Experiment

An experiment was conducted with the goal of using the correlated filtering condition to determine whether increasing the number of variant dimensions without changing the number of stimuli can produce an interference effect. Color stimuli were chosen due to their relatively balanced three-dimensional nature and because of their position as the canonical example of integral dimensions (Melara et al. 1993). A between-participants design was chosen in which a given subject always makes classification decisions according to the same relevant dimension and only participates for a single, one-hour session. This was to eliminate day-to-day variation within a participant, negate attentional concerns if participants had to switch relevant dimensions, and due to the concerns of Ashby and Maddox (1994) that well practiced observers may implement more complicated decision boundaries to optimize performance.

Method

Participants Participants were college students with ages ranging from 18 to 25. They were compensated for their participation with a coupon redeemable for a bookstore item valued at $10. Data was initially collected from 30 participants. One participant's data were thrown out for having chance accuracy (only 44 % correct), and an additional participant was recruited to replace the data.

Stimuli The stimuli for this experiment were square patches of color, which were chosen using the Munsell

color system, which attempts to equate the perceptual discriminability between changes in the three dimensions: saturation (referred to as *chroma*) brightness (*value*), and hue. In Munsell notation, the eight stimuli had a chroma of either 4 or 8, a value of either 4 or 6, and a hue of 10B or 7.5PB. These stimuli can be seen in Fig. 1.

Materials All trials took place in a dark room, with stimuli shown against a uniform grey background, $rgb =$ (200, 200, 200), on a 16" Dell Trinitron CRT monitor set to 1024×768 pixel resolution with a refresh rate of 75 Hz. Participants were seated 70 cm away from the monitor. Data was collected using the freely available PsychoPy experimental software (Peirce, 2007). Stimuli were displayed in the center of the screen and were 150×150 pixels in size, which equated to 3.8 degrees of visual angle. Auditory feedback was given on all trials through noise-attenuating headphones, with different tones denoting correct, incorrect, or slow responses (those longer than 2 s).

Procedure Each participant was instructed to classify stimuli according to a single dimension throughout the experiment, with ten participants randomly assigned to each (hue, saturation, and brightness). The experiment consisted of four blocks of each of three conditions: baseline, filtering, and correlated filtering.

The four baseline blocks differed in terms of the fixed values of the irrelevant dimensions: a participant instructed to classify stimuli according to hue distinguished the two dark-saturated colors in one block, the dark-unsaturated in another, then light-saturated, and finally light-unsaturated. The four filtering blocks differed in terms of which of the two irrelevant dimensions was allowed to vary within the block of trials, and the level at which the other irrelevant dimension was fixed (e.g., brightness varies and all stimuli are saturated). There are only two possible correlated filtering configurations, however, with the irrelevant dimensions either being correlated in one way (e.g., dark colors are always unsaturated while light are saturated) or the other, so each of these two blocks was repeated to maintain an equal number of trials for planned comparisons.

These 12 blocks were presented in a different randomized order for each participant. Within each block, each stimulus being used appeared 40 times in random order, with the caveat that two appearances of each stimulus were presented at the beginning of the block as practice and then excluded from analysis. This resulted in each baseline block consisting of 80 trials, while the other blocks (which used four stimuli) had 160 trials. Each participant thus received 1600 trials, 1520 of which were submitted for analysis. The experiment took approximately 50 min to complete.

Regardless of the particular condition, every block of trials unfolded in exactly the same manner. Participants were reminded which keys corresponded to the values of their assigned dimension (e.g. "F" for blue and "J" for purple), but were never informed as to which combination of stimuli would be appearing in a given block. They were allowed to rest as long as they wanted between blocks of trials. For all trials, there was a 500-ms blank screen followed by a fixation cross which was displayed for a random length of time uniformly distributed between 250 and 750 ms, so as to disrupt automatic responding. The stimulus was then presented, and remained on the screen until a response was given or the timeout value of 2 s was reached. A feedback tone was then played for 100 ms, and the next trial commenced. Each trial thus took around 1.6 s for a typical participant.

Results

Before any analysis was conducted, reaction times of 2 s or more (time-outs) and those less than 250 ms were thrown out. Out of a total of 45,600 trials, there were only 255 of the former and 305 of the later (less than 1 % of each). The remaining overall accuracy was 92 % with correct trials taking 506 ms on average. Individual participants ranged from 84 to 99 % accuracy, and 390 to 660 ms average RT for correct trials.

A mixed-effects model was implemented to model RT on correct trials using condition (baseline, filtering, or correlated filtering) and relevant dimension as fixed effects, with a random intercept for participant and a by-participant random slope for condition (allowing the effect of condition to differ for participants). There was a main effect of condition, $F(2, 28.86) = 48.24$, $p < 0.001$, but no effect for dimension $F(2, 26.94) = .19$, $p = 0.83$. Planned comparisons for condition revealed that baseline (454 ms) was faster than filtering (511 ms), $t(28.73) = 7.89$, $p < 0.001$, which in turn was faster than correlated filtering (530 ms), $t(28.97) = 2.60$, $p = 0.014$. A graph of the results is shown in Fig. 2.

To check for speed accuracy trade-offs, the same analysis was run with respect to accuracy. The results followed the same pattern, with a main effect of condition, $F(2, 28.89) = 21.25$, $p < 0.001$, but no effect of dimension, $F(2, 26.94) = 2.17$, $p = 0.13$. Mirroring the RT results, baseline (94.5 %) had better accuracy than filtering (92.2 %), $t(28.95) = 3.82$, $p < 0.001$, which in turn was more accurate than correlated filtering (91.2 %), $t(28.73) = 2.56$, $p = 0.016$.

To more closely examine the presence of traditional Garner interference for each possible pairing of dimensions, data were grouped by relevant dimension. Tests were performed comparing reaction times in the four baseline blocks to the two filtering blocks in which a particular second dimension was allowed to vary, and then repeated using
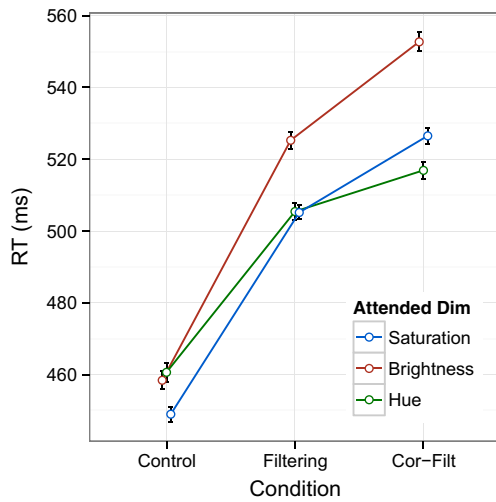
Fig. 2 Data showing the effects of condition and relevant dimension on reaction time

the other choice of irrelevant dimension. Note that all eight stimuli appear equally often on both sides of these comparisons. A linear model was used with condition as a within-participants variable and participant as a random effect. All effects were significant with $p < 0.001$, with full details shown in the top half of Table 1.

In a similar fashion, tests were conducted to compare correlated filtering with standard filtering in each of the six possible combinations: for a given relevant dimension,

Table 1 Tests of Garner interference (first six) and correlated filtering (last six) for each combination of relevant and irrelevant dimensions, with abbreviations for saturation, brightness, and hue

Rlv	Irlv	RT Diff	df	F-Val	*p* value	D
Sat	Brt	68.75	(1,5683)	414.40	<0.001	0.51
Sat	Hue	43.64	(1,5681)	182.99	<0.001	0.34
Brt	Hue	67.16	(1,5706)	319.81	<0.001	0.43
Brt	Sat	66.53	(1,5606)	337.36	<0.001	0.44
Hue	Sat	33.10	(1,5512)	79.36	<0.001	0.21
Hue	Brt	56.65	(1,5470)	183.67	<0.001	0.34
Sat	+	13.40	(1,5503)	10.93	0.001	0.08
Sat	−	20.30	(1,5525)	27.81	<0.001	0.13
Brt	+	15.40	(1,5536)	14.02	<0.001	0.09
Brt	−	18.35	(1,5492)	23.60	<0.001	0.12
Hue	+	12.09	(1,5504)	9.41	.002	0.07
Hue	−	21.61	(1,5524)	30.07	<0.001	0.13

For the correlated filtering comparisons, both dimensions vary irrelevantly, but can either be positively (+) or negatively (−) correlated. RT Diff represents filtering minus baseline for the top and correlated filtering minus standard filtering for the bottom. The final column is Cohen's measure of effect size

the two irrelevant dimensions could be either positively or negatively correlated. Positive correlation was defined for convenience as occurring when the values of saturated, dark, and/or blue were paired.

To control for potential salience differences between stimuli, tests were performed comparing trials from one of these correlated filtering conditions to those filtering trials that used the same relevant dimension and stimuli. This involved taking half of the trials from each of the four filtering conditions for a given relevant dimension. In this way, only half of the stimuli are used in each comparison, with an approximately equal number of trials from each condition. A linear model was used with condition as a within-participants variable and participant as a random effect. Almost all effects were significant with $p < 0.001$ (one had $p = 0.002$), and full details are shown in the bottom half of Table 1.

Sequential analyses were also performed to examine trial-to-trial effects. Again, the critical comparison here is between correlated filtering and standard filtering. The purported test of stimulus uncertainty effects from Dyson and Quinlan (2010) can be used to directly compare these conditions: restricting our comparison to only those trials in which the irrelevant dimension did not change from the last trial, these two conditions should be equivalent, since they have the same number of stimuli. However, when these trials were compared using a mixed-effects model of RT as a function of condition with a random intercept for participant and a by-participant random slope for condition, the correlated filtering trials (514 ms) were significantly slower than filtering (499 ms), $t(28.94) = 2.15$, $p = 0.040$.

The sequential test for irrelevant variation, which compares trials in which the value of the irrelevant dimension was repeated to those in which it changed, is a within-condition test. Looking at the interaction between this factor and the condition, however, can indicate if the strength of this effect was greater for correlated filtering than for filtering. This was tested with a mixed-effects model of RT as a function of the product of condition with change in the irrelevant dimension, with a random intercept for participant and a by-participant random slope for condition. This revealed that in addition to overall slower performance when the irrelevant dimension changed (538 ms) than when it repeated (506 ms), $t(32994) = 14.10$, $p < 0.001$, an interaction shows that this effect was greater for the correlated filtering condition (36 ms) than the filtering condition (27 ms), $t(32994) = 2.40$, $p = 0.016$.

Discussion

These results suggest that stimulus uncertainty is not the only driving force behind Garner interference. Although

practitioners have sometimes assumed that Garner interference is the direct result of allowing a second dimension to vary irrelevantly, models of the effect often achieve slower performance in the filtering condition only due to the increase in stimulus uncertainty that happens when going from using two stimuli to four.

Although increasing stimulus uncertainty is likely a contributing factor, this experiment has shown that Garner interference can also be attributed directly to the increase in the number of variant dimensions, with these two effects confounded in the traditional comparison between baseline and filtering. Utilizing a third stimulus dimension allowed for the creation of a novel condition, correlated filtering, in which two irrelevant dimensions vary simultaneously in a correlated fashion. Comparing this condition to the standard filtering condition selectively manipulates the number of dimensions while controlling for the number of stimuli.

Data from this experiment showed that the correlated filtering condition was slower than the standard filtering condition across all six possible dimensional pairings. This comparison had an average effect size of $D = 0.10$, making it weaker than Garner interference, which averaged $D = 0.38$. A likely reason for this difference in strength is that the traditional test is a combination of both the stimulus uncertainty effect and the change in irrelevant variation.

This disambiguation of potential influences goes further than the sequential analyses of Dyson and Quinlan (2010), who attempted to improve the filtering test by separating it into two components that take account of the relation between a stimulus and its predecessor. While their measure of stimulus uncertainty is insulated from the effects of trial-to-trial variation in the irrelevant dimension (by using only trials where the value of this dimension is repeated), this does not mean that it is fully blind to the effects of irrelevant variation, which likely occur on a broader block-level as well.

These results support such a claim by showing a significant difference on this measure between filtering and correlated filtering, which were designed to be balanced in terms of stimulus uncertainty. This measured difference must therefore be attributed to the effects of irrelevant variation. The separate accounting of trial-to-trial effects from block-level differences is a useful tool provided by this research, but even these more detailed tests are unable to fully separate the effects of stimulus uncertainty from irrelevant variation.

One alternate explanation for these results comes from the Exemplar Based Random Walk (EBRW) model (Nosofsky & Palmeri 1997). This model can attribute interference effects both to the increase in stimulus uncertainty that results from having four stimuli, and to the decrease in within-category similarity. Thus, unlike models solely dependent on stimulus uncertainty, it would still be capable of modeling the results from correlated filtering since within-category similarity is lower here than for filtering.

It is important to note that this explanation differs from the theoretical narrative in that the EBRW model has no representation of which or how many dimensions vary, but rather is solely concerned with inter-stimulus similarity. From this point of view, the correlated filtering condition could be equivalent to a two-dimensional stretch filtering condition, as described by Nosofsky and Palmeri (1997), where distances along the irrelevant dimension are increased. Other models, however, take explicit account of the number of dimensions that vary within a condition, and therefore would treat the 3-D correlated filtering condition differently than the 2-D stretch filtering condition.

Future work in this domain has much to answer. If Garner interference is not solely caused by stimulus uncertainty, as these results indicate, must our models explicitly account for the number of variant dimensions, or will inter-stimulus distance suffice? If both, or even all three of these potential factors contribute to interference with integral dimensions, then why not so with separable dimensions? Should this distinction even properly be considered as a dichotomy, or rather a continuum with varying degrees of attentional selectivity? If success at selectivity is driven solely by dimensional uncertainty and dimensional imbalance, as claimed by Melara and Algom (2003), then can we incorporate the number of variant dimensions into an updated conception of dimensional uncertainty?

This analysis gives partial support to the common presumption that Garner interference is based on the dimensional structure of the stimulus set rather than just an increase in the size of that set, but sets a challenge for how process models will account for these effects. The Garner paradigm has contributed greatly to our knowledge of how perceptual dimensions interact, presenting us with this mysterious dichotomy, but the psychological forces that drive its interference effects have yet to be fully brought to light.

References

Amishav, R., & Kimchi, R. (2010). Perceptual integrality of componential and configural information in faces. *Psychonomic Bulletin and Review*, *17*(5), 743–748.

Ashby, F. G., & Maddox, W. T. (1994). A response time theory of separability and integrality in speeded classification. *Journal of Mathematical Psychology*, *38*(4), 423–466.

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, *93*(2), 154–179.

Atkinson, A. P., Tipples, J., Burt, D. M., & Young, A. W. (2005). Asymmetric interference between sex and emotion in face perception. *Attention, Perception, & Psychophysics*, *67*(7), 1199–1213.

Bartlett, J. C., Searcy, J. H., & Abdi, H. (2003). What are the routes to face recognition. In Rhodes, M., & Peterson, G. (Eds.) *Perception of faces, objects and scenes: analytic and holistic processes (21–52)*. Oxford, UK: Oxford University Press.

Dyson, B. J., & Quinlan, P. T. (2010). Decomposing the Garner interference paradigm: Evidence for dissociations between macrolevel and microlevel performance. *Attention, Perception, & Psychophysics*, *72*(6), 1676–1691.

Fitousi, D., & Wenger, M. J. (2013). Variants of independence in the perception of facial identity and expression. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 133–155.

Ganel, T., & Goshen-Gottstein, Y. (2002). Perceptual integrality of sex and identity of faces: Further evidence for the single-route hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(4), 854–867.

Ganel, T., Goshen-Gottstein, Y., & Goodale, M. A. (2005). Interactions between the processing of gaze direction and facial expression. *Vision Research*, *45*(9), 1191–1200.

Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum Associates.

Garner, W. R. (1976). Interaction of stimulus dimensions in concept and choice processes. *Cognitive Psychology*, *8*(1), 98–123.

Kaufmann, J. M., & Schweinberger, S. R. (2004). Expression influences the recognition of familiar faces. *Perception*, *33*(4), 399–408.

Melara, R. D., & Algom, D. (2003). Driven by information: A tectonic theory of Stroop effects. *Psychological Review; Psychological Review*, *110*(3), 422–471.

Melara, R. D., Marks, L. E., & Potts, B. C. (1993). Primacy of dimensions in color perception. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(5), 1082–1104.

Melara, R. D., & Mounts, J. R. W. (1993). Selective attention to Stroop dimensions: Effects of baseline discriminability, response mode, and practice. *Memory and Cognition*, *21*(5), 627–645.

Melara, R. D., & Mounts, J. R. W. (1994). Contextual influences on interactive processing: Effects of discriminability, quantity, and uncertainty. *Attention, Perception, & Psychophysics*, *56*(1), 73–90.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*(2), 266–300.

Peirce, J. W. (2007). Psychopy—psychophysics software in Python. *Journal of Neuroscience Methods*, *162*, 8–13.