Complexity can facilitate visual and auditory perception

Cameron T. Ellis^{1,*} & Nicholas B. Turk-Browne¹

¹Department of Psychology, Yale University, New Haven, CT, USA

Article type: Research Report

Word count: 6616 (main text incl. tables, captions and references)

*Please address correspondence to:

Cameron Ellis

Department of Psychology

Yale University

2 Hillhouse Avenue

New Haven, CT 06520

cameron.ellis@yale.edu

Draft version, 10/3/18. This paper has not been peer reviewed and is not the authoritative version of the manuscript. Please do not copy or cite without author's permission.

Abstract

Visual and auditory inputs vary in complexity. For example, driving in a city versus the country or listening to the radio versus not. How does such complexity impact perception? One possibility is that complex stimuli demand resources that exceed attentional or working memory capacities, reducing sensitivity to perceptual changes. Alternatively, complexity may allow for richer and more distinctive representations, increasing such sensitivity. We performed five experiments to test the nature of the relationship between complexity and perception during movie viewing. Experiment 1 revealed higher sensitivity to global changes in audio or video streams for clips with greater complexity, defined both subjectively (judgments by independent coders) and objectively (information-theoretic redundancy). Experiment 2 replicated this finding and further showed that this benefit does not result from complexity drawing attention. Experiment 3 provided a boundary condition by showing that change detection was unaffected by complexity when the changes were superimposed on, rather than dispersed throughout, the clips. Experiment 4 suggested that the effect of complexity, at least when defined objectively, does not depend on working memory demands. Experiment 5 suggested that complexity led to richer representations of the clips, as reflected in enhanced long-term memory. Collectively, these findings show that, despite increasing informational load, complexity can serve to ground and facilitate perception.

Keywords: information theory, multisensory perception, compression, change detection, load

Statement of Public Significance

It is commonly assumed that the perceptual systems struggle to cope with an overabundance of sensory information. To the contrary, we report certain cases in which naturalistic stimulus complexity can have an opposite, beneficial effect on perception.

Complexity can facilitate visual and auditory perception

In the study of perception it is typically assumed that sensory inputs are extremely complex and that this complexity impairs function. Researchers invoke the idea of “sensory overload”, the perceptual system diminishing in efficacy under the stress of complex stimulation, as a symptom of some mental disorders (Scheydt et al., 2017), a mistake made in marketing (Malhotra, 1984), and a cause of danger in some environments (Lindenmuth, Breu, & Malooley, 1980). That said, whether and how complexity diminishes perception has received surprisingly little direct investigation. Our goal in this study is thus to characterise the relationship between complexity and perception for naturalistic, multisensory stimuli across different task demands.

There is some suggestive evidence that complexity impairs perception and performance. In multiple object tracking, for example, every additional moving object diminishes a participant’s ability to track objects (Pylyshyn & Storm, 1988). In the temporal domain, stimuli presented close together in a sequence can mask one another, revealing limits on how much processing is possible within a fixed timespan (Marcel, 1983; Raymond, Shapiro, & Arnell, 1992). Moreover, orthographic complexity affects vowel detection time, slowing reading of more complex languages like Arabic (Abdelhadi, Ibrahim, & Eviatar, 2011). Finally, working memory capacity is lower for objects with greater complexity (Alvarez & Cavanagh, 2004; Brady, Konkle, & Alvarez, 2009) and memory recall is worse for lists of numbers with greater complexity (Mathy & Feldman, 2012). Common across these tasks is the implication that complexity is detrimental to the functioning of the perceptual system.

On the other hand, increased complexity might improve perception under some task demands. For instance, when complexity is measured by the number of speakers and the duration of speech, more complex auditory clips are better remembered than less complex clips (Potter & Choi, 2006). Perceptual load theory argues that complex scenes drain surplus cognitive resources and thus reduce task-irrelevant behavior (Lavie, 1995). Relatedly, additional complexity could help by providing scaffolding (e.g., schematic or hierarchical structure) that constrains and supports sensory processing. Rather than always helping or hurting, it is also possible that complexity has non-monotonic effects on perception. For example, infants are most engaged by stimuli that are neither too complex nor too simplistic (Kidd, Piantadosi, & Aslin, 2012). If this inverted-U shaped function exists in perception too, complexity may be beneficial up to a point for some tasks.

In addition to directly assessing the relationship between complexity and perception, we also sought to build upon methods that have been used previously in three ways. First, we employ movie stimuli whose range and variance of complexity may better approximate naturalistic experience, compared to distilled laboratory tasks such as visual short-term memory or multiple object tracking. Second, these movies contain both visual and auditory information, allowing us to characterize joint effects of complexity in multisensory stimuli, rather than by considering each modality separately. Third, complexity is a multifaceted concept (Forsythe, 2009), and here we compare results for two very different definitions — one psychological (coder ratings) and one computational (Kolmogorov complexity).

Five experiments were conducted in order to understand the relationship between complexity and perception in both the visual and auditory modality. Experiment 1 establishes the basic effect that video and audio complexity benefits visual and auditory sensitivity to global changes in a movie clip. Experiment 2 extends and replicates these findings but also suggests that selective attention may not contribute to this relationship. Experiment 3 further calls into question the role of attention in this effect by showing that detection of transient stimuli superimposed on a clip is less affected by complexity. Experiment 4 suggests the effect of complexity can be found for online processing that does not place demands on working memory. Experiment 5 indicates that complex movies are better remembered, suggesting that they afford richer, more distinctive representations. In most cases, these relationships hold for both subjective and objective measures of complexity. Overall, this research suggests that perception can be enhanced by complex stimuli.

Experiment 1

Experiment 1 aimed to establish the relationship between complexity and perception. Psychophysical metrics were used to compare sensitivity to global changes in movie clips, which differed in video and audio complexity.

Methods

Participants

Twenty-four people (11 females, aged 18 to 22), who reported normal or corrected-to-normal vision, participated in this experiment. As this was a novel task and analysis approach, we chose this sample size based on two pilot studies that each revealed

consistent but tentative effects in 16 participants. Three participants, not included in the total, were excluded from the experiment because they failed to complete the experiment in time (1) or misunderstood task instructions (2). All participants received course credit or \$12/hour as compensation and provided informed consent to a protocol designed in accordance with the principles expressed in the Declaration of Helsinki and approved by the Institutional Review Board of Princeton University.

Stimuli

An open source film, “Tears of Steel” (Blender Foundation, 2012) was partitioned into 24 5-s clips (24 frames per second) at 1080p uncompressed quality. These clips depicted a variety of scenes: indoors, outdoors, with people, without people, special effects, etc.

The subjective complexity of these clips was judged by separate groups of participants on Mechanical Turk, who were redirected to a Qualtrics survey hosting embedded movies from YouTube. To ensure their ratings were based on the correct modality, participants were exposed to either the video (N=56) or audio (N=51) stream. They were asked to “Evaluate the complexity of the clip” using a sliding scale, where complexity was previously defined to them as “how busy, unpredictable or elaborate the clip is”. Participants could re-play a clip if desired. Clips were all played with the same high quality (i.e., without lossy compression), so that load times were equated across clips. Ratings differed in reliability across modalities: intra-class correlation — a measure of inter-rater reliability (Shrout & Fleiss, 1979) — revealed poor reliability for audio (0.28) and fair for video (0.45) (Cicchetti, 1994). Overall subjective complexity scores for the audio and video stream of each clip were calculated by averaging ratings across

participants. These average audio and video scores were highly correlated across clips ($r(22)=0.71$, $p<.004$), suggesting a relationship between the contents in both modalities, given that participants were only ever presented with one stream.

The objective complexity of these clips was calculated with an information theoretic measure related to Kolmogorov complexity. This refers to the length of the shortest sequence sufficient to describe an object, and in this way reflects the amount of redundancy present in a data stream (Li & Vitanyi, 2009). To quantify redundancy in this way, we employed lossless compression algorithms and determined file size reduction in each modality. This specific metric was used previously as a proxy for complexity (Nagaraj, Kavalekalam, & Krishnan, 2011). For video complexity, we calculated the redundancy of each frame using the PNG algorithm and then averaged file size reduction across the clip. For audio complexity, file size reduction for the clip was based on the FLAC algorithm. We intentionally chose clips of variable and uncorrelated ($r(22)=0.03$, $p<.391$) audio and video complexity, in order to assess a broad range of complexities and to be able to quantify separate and joint effects of each modality. Only one objective measure was obtained for each stream and movie, from a deterministic algorithm, precluding analyses of reliability.

We measured perceptual sensitivity in visual and auditory modalities by presenting two clips on each given trial. One clip was the original version described above and the other clip was identical but perceptually degraded in one modality. Participants made a two-interval forced choice (2IFC) by discriminating between the clips and selecting which was degraded. To create a degraded version of each clip, the original audio and video streams were compressed using lossy algorithms. The jpeg2000

codec was used for the video stream, which reduced the file size by a given multiple: 16, 32, 64, 128 or 256. The mp3 codec was used for the audio stream, which reduced the bits per second encoding to a given rate: 128, 64, 32, 16 or 8. These degradation ranges were determined from two pilot experiments to produce comparable levels of 2IFC accuracy across modalities and ensure that we covered the full accuracy range of the psychometric function. The degraded clips were then up-sampled to a higher fidelity format (uncompressed AVI for video and FLAC for audio), preserving the degradation but equating file sizes and load times across levels.

Note that a potential source of confusion is the use of compression at two different points in the design. Specifically, lossless compression was used to estimate objective complexity, but was not applied to the stimuli used in the experiment, nor would have affected stimulus quality, by definition. On the other hand, lossy compression was applied to the stimuli for the degraded version of each clip in the 2IFC task, which was used to assess perceptual sensitivity. Note that this is one of many possible ways of introducing perceptual change (e.g., we could have blurred, changed colors or pitches, inserted blanks, etc.). We chose to degrade the stimuli because this introduce subtle, parametric, and global distortions that ensured attention to the entirety of the clips and allowed us to obtain a full range of accuracies. One concern we had in using lossy compression is that it might interact with complexity. For example, low complexity clips may provide more flexibility in terms of the amount of lossy compression possible for a given discrimination accuracy (e.g., at the extreme, heavy vs. minimal degradation of a solid black screen would be equally noticeable). We used the lossy compression approach anyway because similar concerns apply to other kinds of perceptual changes (e.g.,

blurring) and, regardless, these concerns do not apply as directly to the subjective measures of complexity used throughout.

Procedure

Participants were seated in a dimly lit room in front of a large monitor (53 cm wide) and presented with movies using Psychtoolbox (Kleiner et al., 2007) at a comfortable audio level. Figure 1A outlines the structure of this experiment. A fixation dot blinked to indicate the start of a trial. Participants were shown the same clip twice (45° visual angle), separated by a 1000-ms pause. A 1500-ms response window began at the end of the second video. Participants were asked to judge whether the first or second clip was lower in quality. They were not instructed to focus on a particular modality and did not report to us whether they noticed which modality was degraded; instead there were asked to make a holistic judgment about quality. Whether the first or second clip was lower in quality was randomized. The accuracy of each response was the dependent variable. Upon response or at the end of this response window an inter-trial interval began, ranging from 350 to 850 ms. The identity of the movie clip, the modality reduced in quality, and the level of degradation were counterbalanced to produce 240 trials, with 10 breaks interspersed throughout.

Before the start of the experiment, each participant did 4 practice trials in which they were given feedback about the correct answer and what modality the change was in. These trials were controlled such that participants were shown some trials in which the change was obvious and some trials where the change was almost imperceptible.

Results and Discussion

The aim of Experiment 1 was to investigate how video and audio complexity relates to visual and auditory perceptual sensitivity. We initially expected that the subjective and objective measures of complexity would differ: the subjective measure has the potential to include a variety of components ranging from pixel/bit variability (inverse of redundancy), to the number of objects/sources, to narrative complexity; whereas the objective measure only reflects pixel/bit variability in the data stream. The correlation between subjective and objective audio complexity across movies was strong and positive ($r(22)=0.93, p<.001$), suggesting that subjective audio complexity is based on the temporal variability captured by objective complexity. However, the correlation between subjective and objective video complexity was not significant ($r(22)=-0.21, p=.245$), suggesting that subjective video complexity was largely based on different features of the stimulus.

Collapsing across the degree of degradation, participants were 74.2% correct on average when the change was visual and 74.0% when the change was auditory. To quantify the effect of complexity on accuracy, we used the Palamedes toolbox in MATLAB (Prins & Kingdom, 2009) to fit psychometric functions, with parameters for threshold, slope, and lapse rate (guessing rate was not fit because it was a forced choice task), over different levels of degradation for movies with high vs. low complexity (median split). The critical parameter was the threshold: the amount of change required for accuracy halfway between chance and the lapse rate. The parameters were fit to the average accuracy across participants at each level of degradation to improve the goodness of fit. The random-effects reliability of these estimates was assessed with bootstrapping

tests, by resampling participants with replacement 10,000 times prior to calculating the average accuracies and fitting the functions on each iteration. The resulting sampling distribution of the parameter estimates was used to generate 95% confidence intervals and for hypothesis testing. Note that sampling participants with replacement allows for non-parametric testing of reliability across participants in such pooled analyses (Efron & Tibshirani, 1986; Fan, Hutchinson, & Turk-Browne, 2016; Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014). The intuition for why this works is that to the extent that an effect is reliable across participants, their data should be substitutable with minimal impact on the results, leading to a narrow sampling distribution. Table 1 contains the descriptive statistics for these psychometric analyses.

Table 1.

Psychometric compression thresholds (video: proportion file size reduction [high is more compression]; audio: Kilobits per second [low is more compression]) calculated by averaging participants and movies. Confidence intervals are bootstrapped distributions of each condition.

Definition	Modality	Complexity	Mean	Low CI	High CI
Subjective	Video	High	64.7	49.6	74.1
		Low	79.4	66.2	94.4
		Diff	-14.7	-27.6	-1.9
	Audio	High	59.4	51.9	66.8
		Low	37.6	28.9	48.2
		Diff	21.7	10.4	31.3
Objective	Video	High	61.4	49.6	74.1
		Low	78.6	64.0	97.2
		Diff	-17.2	-34.3	-1.9
	Audio	High	55.8	49.1	63.7
		Low	37.4	25.1	52.7
		Diff	18.4	2.1	30.2

Figure 1. Methods and results for Experiment 1. (A) Example trial. Psychometric functions for high and low subjective (B) and objective (C) complexity in visual and auditory modalities. Above the plots are the resampled distributions of thresholds, significance of difference: * $p < .05$, *** $p < .001$.

The threshold for detecting visual changes was lower (better performance) for movies with high compared to low video subjective complexity ($p = .001$; Figure 1B). Similarly, movies with high versus low audio subjective complexity had lower auditory thresholds ($p < .001$). The same pattern of results was obtained for objective complexity,

video ($p=.002$; Figure 1C) and audio ($p=.002$). There was also a difference in lapse rate between high and low complexity for subjective and objective definitions, but our definition of threshold as halfway between chance and the lapse rate compensates for this.

To fit the psychometric functions above, we performed a median split on complexity and averaged across the movies in each bin to obtain generally smooth and monotonic data. In order to evaluate complexity as a continuous variable, we performed an additional analysis at the level of individual movies by calculating an average accuracy across quality levels of each movie. We used linear regression to predict these accuracies from measures of subjective and objective complexity and assessed the robustness of the slope with bootstrap resampling. As shown in Figure 2A-B, there was a positive relationship between movie accuracy and both subjective complexity (video: $b=0.08$, $CI=[-0.003, 0.338]$, $p=.053$; audio: $b=0.495$, $CI=[0.547, 0.732]$, $p<.001$) and objective complexity (video: $b=0.402$, $CI=[0.125, 0.506]$, $p<.001$; audio: $b=2.640$, $CI=[0.602, 0.770]$, $p<.001$). That is, as the complexity of movies increased, so did the ability of participants to detect perceptual changes.

Figure 2. Regression results for Experiment 1. Relationship of (A) subjective complexity and (B) objective complexity with proportion correct in detecting visual and auditory changes. Shaded area represents the 95% confidence interval of the slope values from bootstrap resampling.

We began this experiment uncertain whether complexity would overwhelm perception and hurt performance or allow for richer representations that help performance. Our results from both psychometric and regression analyses supported the latter possibility. The consistency of results between subjective and objective measures of complexity is striking, especially for the visual modality, given that these two definitions were uncorrelated. This suggests the existence of two forms of complexity that impact perception.

Experiment 2

Experiment 1 was an exploratory study, in the sense that we anticipated that either direction of effect was possible. Therefore, in Experiment 2 we sought to replicate these findings and at the same time test a potential explanation based on attention. Specifically, more complex video and audio streams may have attracted attention, which in turn has been shown to affect perceptual sensitivity (Pestilli & Carrasco, 2005; Störmer, Winther, Li, & Andersen, 2013). Here we explicitly manipulate attention to the video or audio stream using a cuing paradigm (Posner, Nissen, & Ogden, 1978). As a sanity check, we expect improved detection of global changes that are validly vs. invalidly cued. However, the key prediction of the attention account involves an interaction between complexity and cuing: If complex movies inherently attract or hold attention, we would expect a diminished impact of the cue. We tested this prediction by calculating the relationship across movies between complexity with the validity effect. Moreover, insofar as the positive relationship between complexity and sensitivity from Experiment 1 depends on attention, it should be diminished when analysis is restricted to invalid trials.

Methods

Participants

The compensation and consent for the 24 participants (13 female, age 18-21) in this experiment were the same as Experiment 1.

Procedure

The room, computer, and movie clips were the same as Experiment 1. Before the two clips appeared on each trial, the word “Auditory” or “Visual” was presented for 750ms to indicate with 75% probability whether the audio or video stream contained the

change in quality between clips (Figure 3A). Participants were told that these cues would be accurate “most of the time.” The degree of audio or video quality change for each movie was determined using the 75% accuracy threshold identified in Experiment 1 (average 71:1 reduction for video, 40 Kb/s for audio). In other words, whereas we used multiple quality change levels in Experiment 1, here we picked one sensitive level and looked at the impact of the attentional manipulation on accuracy at that level. Each of the 24 movies was shown eight times: four with video change, four with audio change, and within each modality three valid trials and one invalid trial (192 total trials). Whether the lower quality clip appeared first or second was randomized. Akin to Experiment 1 participants completed 12 practice trials before the experiment began, also with 75% cue validity.

For this and subsequent experiments, we made use of an eye-tracker (Tobii X120, sampling rate = 60Hz) available in the testing room to monitor participant alertness and compliance. We did not have any hypotheses about eye movements related to the task design, especially since the cuing task was sequential, and thus do not report any data or analyses.

Results and Discussion

Overall participants were correct 70.2% of the time when there was a visual change and 80.3% when there was an audio change. Hence, the amount of quality change based on the 75% threshold from Experiment 1 was approximately correct.

To attempt to replicate Experiment 1, we first calculated the correlation across clips between complexity and sensitivity without consideration to cuing, and used the same resampling procedure to assess reliability. The relationship between complexity and

sensitivity was positive for subjective complexity in both modalities (video: $b=0.213$, $CI=[0.153, 0.439]$, $p=.001$; audio: $b=0.647$, $CI=[0.601, 0.784]$, $p<.001$) and objective complexity in both modalities (video: $b=1.022$, $CI=[0.416, 0.673]$, $p<.001$; audio: $b=2.842$, $CI=[0.508, 0.711]$, $p<.001$). Hence, as in Experiment 1, both subjective and objective complexity seem to enable improved performance in perceptual tasks.

As verification that the attention cue was effective, accuracy was higher on valid vs. invalid trials, for quality changes in both modalities (video: 71.9 vs. 63.3%, $t(23)=2.86$, $p=.009$; audio: 82.1 vs. 74.7%, $t(23)=3.17$, $p=.004$).

Insofar as complexity attracts attention, this validity effect may be attenuated for more complex clips because attention is drawn to them regardless of the cue. We calculated the validity effect for each clip as the average across participants of the difference between valid and invalid trials containing that clip. We then predicted this effect across clips with the subjective and objective complexity measures. To assess the reliability of the slopes, we resampled participants used in calculating the per-clip validity effects. There was no relationship across clips between the validity effect and subjective complexity (video: $b=-0.057$, $CI=[-0.356, 0.274]$, $p=.857$; audio: $b=0.024$, $CI=[-0.339, 0.422]$, $p=.974$) or objective complexity (video: $b=-0.062$, $CI=[-0.796, 0.574]$, $p=.794$; audio: $b=0.048$, $CI=[-1.551, 2.056]$, $p=.916$). These null effects fail to support an attentional explanation of the relationship between complexity and perceptual sensitivity, especially in light of the robust overall validity effect.

Figure 3. Methods and results for Experiment 2. (A) Example trial. Relationship of (A) subjective complexity and (B) objective complexity with proportion correct in detecting visual and auditory changes. Shaded area represents the 95% confidence interval of the slope values from bootstrap resampling. Green refers to valid trials when the cued modality contained the change (75%); red refers to invalid trials (25%).

Indeed, the relationship between clip complexity and accuracy was immune to the attention cue. For valid trials (Figure 3B-C, green), there was a positive relationship for both subjective (video: $b=0.326$, $CI=[0.093, 0.358]$, $p<.001$; audio: $b=0.703$, $CI=[0.542, 0.758]$, $p<.001$) and objective complexity (video: $b=0.516$, $CI=[0.614, 1.349]$, $p<.001$; audio: $b=0.630$, $CI=[2.350, 3.444]$, $p<.001$). Critically, despite having only a third as much data, the same relationships — if anything, numerically stronger — were found on

invalid trials (Figure 3B-C, red), for both subjective (video: $b=0.363$, $CI=[0.017, 0.506]$, $p=.035$; audio: $b=0.728$, $CI=[0.245, 0.986]$, $p=.002$) and objective complexity (video: $b=0.552$, $CI=[0.532, 1.614]$, $p<.001$; audio: $b=0.635$, $CI=[0.899, 4.489]$, $p=.005$). Thus, even when attention is diverted to another modality, changes in complex clips are more easily noticed than changes in simple clips.

This study replicated the positive relationship between complexity and perceptual sensitivity, but also called into question the role of attention in this relationship. Although attention to the video or audio stream of a clip improved discrimination of that stream relative to the uncued stream, the complexity of the clip did not affect the benefit of this cue. In fact, even when attention was directed to the other stream, the complexity of the uncued stream still affected sensitivity to changes in that stream.

Experiment 3

One interpretation of Experiment 2 is that attention does not mediate the effect of complexity on perceptual sensitivity. In particular, the obtained relationship for the uncued stream suggests no role for goal-directed attention. However, this finding is also consistent with the possibility that complex stimuli capture attention in an involuntary, stimulus-driven manner, which in turn improves sensitivity. Here we test this possibility by removing the attention cue and embedding a visual or auditory oddball within movie clips varying in complexity. Insofar as complexity drives attention, the detection of these oddballs should be enhanced for more complex video and audio streams.

Methods

Participants

A new cohort of 24 participants (15 female, age 18-24) were recruited for this experiment. One participant was excluded because they failed to finish in time.

Procedure

On each trial, a single clip was presented uncompressed (Figure 4A). An oddball appeared on half of the trials, consisting of an unexpected transient superimposed in one modality on a single movie frame (42 ms). The oddball could not appear within the first or last five frames of the clip but was otherwise timed randomly. In the visual modality, the transient consisted of a circular patch of scrambled pixels 8.3° in diameter on either the left or right side of the screen. In the auditory modality, a high-frequency sine wave was played in either the left or right ear of the participant. After each clip, participants were asked to respond with a key press whether they detected a transient in the audio or video stream, or they could omit their response to indicate that no transient was detected. The identity of the movie clip, the presence of a transient, and, if present, the modality and side of the transient were counterbalanced, resulting in 384 trials shown in random order. Akin to Experiments 1 and 2, participants completed 4 practice trials before the experiment began.

Results and Discussion

Participants were accurate on average in reporting the presence of a transient in the visual (69.9%) and auditory modalities (75.4%). On trials with a visual transient, they rarely reported an auditory transient (1.0%), and vice versa (1.4%). On trials with no

transient, the false alarm rates were very low for both visual (1.6%) and auditory (1.9%) transients.

The critical analyses concerned whether transient detection accuracy improved with complexity. Inconsistent with this prediction, the relationship between subjective video complexity and visual sensitivity was negative ($b=-0.275$, $CI=[-0.359, -0.191]$, $p<.001$; Figure 4B). Moreover, there was no reliable relationship for subjective audio complexity ($b=-0.048$, $CI=[-0.197, 0.115]$, $p=.541$), objective video complexity ($b=0.1982$, $CI=[-0.055, 0.473]$, $p=.120$), and objective audio complexity ($b=-0.170$, $CI=[-0.920, 0.640]$, $p<.674$) (Figure 4C).

Figure 4. Methods and results for Experiment 3. (A) Example trial. The video stream has a transient on the last frame. Relationship of subjective (B) and objective (C) complexity to proportion correct in detecting the transient in visual and auditory modalities. Shaded lines represent the confidence interval of the slope from resampling.

Overall, we failed to extend the effect found with global change detection in Experiments 1 and 2 to the detection of local transients. In other words, the enhancement of perceptual sensitivity by complexity is selective to the complex stimulus itself rather than a general benefit for any stimuli appearing at the same location. This is inconsistent with the hypothesis that complexity captures attention in a stimulus-driven manner, and thus does not support an explanation of improved sensitivity mediated by attentional modulation. However, it is possible that complexity modulates attention in ways not examined here, such as in local vs. global scope.

Experiment 4

We observed a benefit of complexity for perceptual sensitivity in Experiments 1 and 2 but not Experiment 3. One important difference between these studies was that the discrimination in Experiments 1 and 2 was between two clips presented sequentially, a task that engages short-term, working memory to make the comparison. The task in Experiment 3 placed no such demands. Thus, it is possible that working memory mediates the relationship between complexity and perceptual sensitivity. Here we test this possibility by returning to the design of Experiments 1 and 2 with two movies per trial, but present them simultaneously, allowing participants to make immediate perceptual comparisons between movie frames. Insofar as the benefit of complexity depends on

working memory, there should no longer be a relationship between complexity and perceptual sensitivity.

Methods

Participants

A new cohort of 24 participants (12 female, age 18-24) were recruited for this experiment.

Procedure

Two movies were displayed simultaneously on either side of fixation and different audio was played to each ear (Figure 5A). To fit on the screen, the videos were reduced in size (13° width, 0.125° separation). To help compensate for the increased difficulty of this task from stimulus conflict and size, the amount of quality degradation for the movie that had to be detected was increased. Specifically, the upper end of the quality changes from Experiment 1 were used: factor of 64, 128 or 256 for the video and bit rate of 32, 16 or 8 Kb/s for the audio. For data analysis, we collapsed across the three levels of degradation, as it was not possible to fit psychometric curves to this small number of steps. Participants had to respond whether the left or right movie was lower in quality, and the location of this degraded movie was randomized. The fixation and inter-trial interval were the same duration as Experiments 1-2. The identity of the movie clip, the modality reduced in quality, and the level of degradation were counterbalanced; each combination of parameters appeared twice, resulting in 288 trials. Participants completed 8 practice trials before the experiment began.

Results and Discussion

Participants were correct 65.0% of the time when there was a visual change and 78.3% when there was an audio change. We replicated the positive relationships from Experiments 1 and 2 between subjective audio complexity and auditory sensitivity ($b=0.341$, $[CI=0.416, 0.649]$, $p<.001$) (Figure 5B), objective video complexity and visual sensitivity ($b=0.685$, $CI=[0.360, 0.668]$, $p<.001$), and objective audio complexity and auditory sensitivity ($b=1.523$, $CI=[0.360, 0.607]$, $p<.001$) (Figure 5C). However, there was no reliable relationship between subjective video complexity and visual sensitivity ($b=0.025$, $CI=[-0.117, 0.233]$, $p=.560$). Note, however, that this relationship was also the weakest of the four in both Experiments 1 and 2, and was only marginally significant in Experiment 1.

Figure 5. Methods and results for Experiment 4. (A) Example trial. Relationship between subjective (B) and objective (C) complexity and proportion correct visual sensitivity and auditory sensitivity. Shaded lines represent the confidence interval of slope values from resampling.

Thus, the qualitative pattern of results was similar, but now without demands on working memory. This suggests that the benefit of complexity on perceptual sensitivity can manifest at the time of immediate perception, at least when the changes to be detected are integral to the complex stimulus (unlike Experiment 3).

Experiment 5

The fact that complexity benefits perception during online processing but does not seem to depend on attention means that a different explanation is needed. One possibility is that complex clips result in richer perceptual representations, which in turn allow global distortions to be better detected. This could occur because complex clips contain more elements, which enable participants to evaluate the quality of the clip over multiple locations, time points, and feature dimensions, improving the likelihood of noticing a distortion and thereby increasing sensitivity. Here we sought corroborating evidence for this account by examining long-term memory. Insofar as complexity enhances perceptual representations, a downstream consequence may be more detailed memories for complex clips, including as a result of the availability of additional cues at retrieval. Thus, we conducted a long-term memory task with individual clips that were encoded incidentally, followed after a delay by a surprise episodic memory test. We hypothesized that complex clips would be better remembered in the test.

Methods

Participants

A new cohort of 49 participants (34 female, age 18-22) were recruited for this experiment. Three participants, not included in this total, were excluded from the analysis because they failed to complete the experiment in time.

Stimuli

Twenty-four additional 5-s clips were included in this experiment for a total of 48 clips. They were taken from the same film but at previously unused time points intermingled with the original set. The complexity ratings for these new movies were

collected using the same Mechanical Turk procedure as Experiment 1 (N=51 for audio, N=51 for video). The ratings were as reliable as the first set: audio had poor interclass correlation (0.12) and video had fair (0.39). Like the first set of movie clips, the correlation between subjective and objective measures of audio complexity was strong and positive ($r(22)=0.67, p=.007$), whereas subjective and objective video complexity were not reliably correlated ($r(22)=0.27, p<.179$). The two sets of movie clips were counterbalanced across participants, with half of participants receiving the first set during encoding and as old items at test, and the second set serving as novel lures at test.

Procedure

As outlined in Figure 6A, participants completed three phases. In the Encoding phase, participants were shown 24 movies while their eyes were tracked. They were given no instruction about what to do while watching these other than to remain vigilant. In the Distractor phase, participants completed a distractor task for five minutes in which they summed two two-digit numbers at their own pace. In the Recognition phase, they were tested on their memory for the previously presented clips and 24 novel lure movies randomly intermixed. Because audio and video complexity vary with respect to each other and may impact memory differently, we tested the two modalities of each clip separately. Specifically, participants were presented with either just the video or just the audio from one of the 48 movies (96 total trials) and reported whether they had: high confidence it was previously presented (sure old), low confidence it was previously presented (unsure old), low confidence it was novel (unsure new), or high confidence it was novel (sure new). To assess episodic memory, “sure old” responses were coded as hits for previously presented clips and as false alarms for novel lures; all other responses

were coded as misses and correct rejections, respectively (Kim et al., 2014; Turk-Browne, Yi, & Chun, 2006). The pattern of results was nearly identical when hits/false alarms were calculated for “old” responses irrespective of confidence. We quantified sensitivity with A' , a non-parametric measure of sensitivity (Aaronson & Watts, 1987; Grier, 1971; Snodgrass & Corwin, 1988).

Results and Discussion

To verify that participants were able to remember the clips, we computed an overall A' for each participant across movies and found that it was reliably above chance (0.50) for both video (mean=0.82, $t(48)=32.29$, $p<.001$) and audio (mean=0.74, $t(48)=22.38$, $p<.001$) tests.

To examine the relationship between memory and complexity, we computed an A' for each movie across participants (half of participants contributed to the hit rate and the other half to the false alarm rate, based on counterbalancing). We then attempted to predict these values from subjective and objective measures of complexity, and resampled participants to obtain confidence intervals for the regression slope. Consistent with our hypothesis, complexity positively predicted subsequent memory for most conditions: subjective video ($b=0.198$, $CI=[0.099, 0.310]$, $p<.001$), subjective audio ($b=0.662$, $CI=[0.388, 0.928]$, $p<.001$) (Figure 6B), and objective audio ($b=2.520$, $CI=[1.312, 3.698]$, $p<.001$) (Figure 6C); however, there was an unexpected negative relationship for objective video ($b=-0.530$, $CI=[-0.681, -0.402]$, $p<.001$).

This pattern of memory results is generally consistent with the perceptual benefits in Experiment 4, suggesting that they may be a downstream consequence. The exception was the objective definition of video complexity, which was positively related to

perceptual sensitivity and negatively related to mnemonic sensitivity. This shows that complexity, at least defined in certain ways, can have dissociable effects on perception and memory. Regardless, in general we interpret these findings as evidence that complex inputs afford richer representations that can provide more opportunities for finding changes during online processing and memory retrieval.

Figure 6. Methods and results for Experiment 5. (A) Schematic of the three phases. Participants incidentally encoded 24 clips, then did a distractor task, and finally had to judge whether they recognized the audio or video stream of a clip from earlier in the experiment. Relationship between subjective (B) and objective (C) complexity with A' for visual and auditory recognition memory. Shaded lines represent the confidence interval of slope values from resampling.

General Discussion

The goal of this work was to investigate how complexity affects perception of naturalistic, multisensory stimuli. Using human judgment and information theory, we evaluated the impact of subjective and objective complexity on the perception of movie clips. Across several experiments, perceptual sensitivity to global distortions was greater for more complex clips, regardless of whether they were presented sequentially (Experiment 1-2) or simultaneously (Experiment 4). Goal-directed attention enhanced sensitivity overall but not in a way that interacted with complexity (Experiment 2). We also did not find evidence that complexity captures attention in a stimulus-driven manner (Experiment 3). Instead, complex clips afforded better long-term memories, especially for auditory streams (Experiment 5).

Our tentative conclusion from these findings is that, rather than overwhelming perception, the additional elements and dynamics present in complex stimuli constrain and support perceptual analysis. As an analogy, consider evaluating whether the prescription of a pair of corrective eyeglasses is correct: looking at a rich and textured scene rather than a uniform surface provides more opportunities and variability over which to assess whether the scene appears correctly. It is unclear from the present work which features and dimensions in the complex clips provided this better scaffolding, though the failure to enhance detection of superimposed stimuli in Experiment 3 suggests that these features were integral to the stimulus.

We attempted to show that the findings reflect a general property of complexity by using two different definitions of complexity. However, the movie clips chosen resulted in a particular range of complexities of both types, and it is possible that the

findings may be different at higher or lower levels of complexity. Specifically, rather than the linear relationships observed here, the full relationship between complexity and perception/memory may be non-monotonic (Kidd et al., 2012). This range issue could also explain why certain types of complexity can be detrimental to perception (Abdelhadi et al., 2011; Pylyshyn & Storm, 1988). Relatedly, artificial stimuli (e.g., white noise) can be outside the range of complexity of naturalistic images, and for such stimuli it would be difficult to notice subtle changes in quality. The impact of complexity range may also differ for auditory vs. visual modalities, as well as for perceptual vs. mnemonic tasks. Finally, having more elements in complex clips against which to identify changes may come at the cost of having lower precision representations of individual elements. This could be examined in future studies that involve discrimination of local rather than global changes during perception and that employ highly similar lures in the memory test, both of which require greater precision.

The extent to which the two complexity measures captured different aspects of complexity differed by modality. In the auditory modality, subjective and objective complexity were highly correlated. Although participants were unconstrained in how they made subjective judgments, we know that objective audio complexity reflected the temporal diversity of the signal, suggesting that these judgments were also based on such features. In the visual modality, subjective and objective complexity were uncorrelated or, if anything, weakly negatively correlated. We think this reflects a difference in the nature of the objective measure used for video vs. audio complexity. In particular, whereas audio complexity tracked temporal diversity, video complexity was measured within each frame using lossless spatial compression and thus ignored temporal diversity

across frames. Despite this difference, objective video complexity behaved like objective audio complexity, with both leading to enhanced perceptual sensitivity. They differed however in their relationship to long-term memory, which could be explored in future studies by contrasting temporal vs. spatial redundancy in videos. Subjective video complexity showed the same relationship to long-term memory as objective and subjective audio complexity, suggesting that these subjective judgments may have been more sensitive to temporal redundancy.

The current study could have practical implications for compression codecs. Specifically, our findings suggest that consumers will be less likely to notice quality decrements from lossy compression when audio and/or video epochs with higher complexity are encoded with higher fidelity. Lossy compression algorithms could thus dynamically vary the amount of compression based on an information-theoretic measure of complexity/redundancy, to efficiently store multimedia content with minimal impact on user experience.

In conclusion, cognitive systems for perception and memory can operate more effectively in the presence of complex input. We provide evidence for a tentative explanation of this phenomenon, namely that it occurs because complex clips contain more elements against which to detect perceptual changes and that can serve as retrieval cues. This stands in contrast to the conventional assumption that sensory systems are overwhelmed with information and that this deluge impairs performance by exceeding perceptual and mnemonic capacities (Alvarez & Cavanagh, 2004; Lavie, 2005). At least for some types of perceptual judgments, complexity can support perceptual sensitivity.

Indeed, given that naturalistic input is highly complex, the benefits reported here may be adaptive.

Acknowledgments

We thank Judy Fan and the Turk-Browne lab for helpful discussions on the project and Nour Malek for feedback on an early version of the manuscript.

Funding

This work was supported by NIH R01- EY021755.

References

- Aaronson, D., & Watts, B. (1987). Extensions of Grier's computational formulas for A' and B'' to below-chance performance. *Psychological Bulletin*, 102(3), 439-442.
- Abdelhadi, S., Ibrahim, R., & Eviatar, Z. (2011). Perceptual load in the reading of Arabic: Effects of orthographic visual complexity on detection. *Writing Systems Research*, 3(2), 117-127.
- Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, 15(2), 106-111.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4), 487-502.
- Blender Foundation (Ton Roosendaal). (2012). Tears of Steel. [Web video]. Available from <https://mango.blender.org/>
- Cicchetti, D. V. (1994). Guidelines, criteria, and rules of thumb for evaluating normed and standardized assessment instruments in psychology. *Psychological Assessment*, 6(4), 284-290.
- Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1(1), 54-75.
- Fan, J. E., Hutchinson, J. B., & Turk-Browne, N. B. (2016). When past is present: Substitutions of long-term memory for sensory evidence in perceptual judgments. *Journal of Vision*, 16(8), 1-12.

- Forsythe, A. (2009). Visual complexity: Is that all there is? Paper presented at the *International Conference on Engineering Psychology and Cognitive Ergonomics*.
- Grier, J. B. (1971). Nonparametric indexes for sensitivity and bias: computing formulas. *Psychological Bulletin*, 75(6), 424-429.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one*, 7(5), e36399.
- Kim, G., Lewis-Peacock, J. A., Norman, K. A., & Turk-Browne, N. B. (2014). Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences*, 201319438.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1-16.
- Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, 21(3), 451-468.
- Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences*, 9(2), 75-82.
- Li, M., & Vitanyi, P. M. (1997). *An introduction to Kolmogorov complexity and its applications* (Vol. 34, No. 10). Heidelberg: Springer.
- Lindenmuth, J. E., Breu, C. S., & Malooley, J. A. (1980). Sensory overload. *The American Journal of Nursing*, 80(8), 1456-1458.

- Malhotra, N. K. (1984). Information and sensory overload. Information and sensory overload in psychology and marketing. *Psychology & Marketing*, 1(3), 9-21.
- Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, 15(2), 197-237.
- Mathy, F., & Feldman, J. (2012). What's magic about magic numbers? Chunking and data compression in short-term memory. *Cognition*, 122(3), 346-362.
- Nagaraj, N., Kavalekalam, M. S., & Krishnan, N. (2011). Lossless Compression and Complexity of Chaotic Sequences. *arXiv preprint*: 1101.4341.
- Pestilli, F., & Carrasco, M. (2005). Attention enhances contrast sensitivity at cued and impairs it at uncued locations. *Vision Research*, 45(14), 1867-1875.
- Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. In H. L. Pick & I. J. Saltzman (Eds.), *Modes of perceiving and processing information* (pp. 137-157). Hillsdale, NJ: Erlbaum.
- Potter, R. F., & Choi, J. (2006). The effects of auditory structural complexity on attitudes, attention, arousal, and memory. *Media psychology*, 8(4), 395-419.
- Prins, N & Kingdom, F. A. A. (2018). Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes toolbox. *Frontiers in Psychology*, 9:1250
- Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3(3), 179-197.

- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 849-860.
- Scheydt, S., Müller Staub, M., Frauenfelder, F., Nielsen, G. H., Behrens, J., & Needham, I. (2017). Sensory overload: a concept analysis. *International Journal of Mental Health Nursing*, 26(2), 110-120.
- Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: uses in assessing rater reliability. *Psychological Bulletin*, 86(2), 420-428.
- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: applications to dementia and amnesia. *Journal of Experimental Psychology: General*, 117(1), 34-50.
- Störmer, V. S., Winther, G. N., Li, S.-C., & Andersen, S. K. (2013). Sustained multifocal attentional enhancement of stimulus processing in early visual areas predicts tracking performance. *Journal of Neuroscience*, 33(12), 5346-5351.
- Turk-Browne, N. B., Yi, D.-J., & Chun, M. M. (2006). Linking implicit and explicit memory: common encoding factors and shared representations. *Neuron*, 49(6), 917-927.