The forest, the trees, or both? Hierarchy and interactions between gist and object processing during perception of real-world scenes

Marcin Furtak¹, Liad Mudrik^{2,3}, Michał Bola¹

1. Laboratory of Brain Imaging, Nencki Institute of Experimental Biology of Polish Academy of Sciences, 3 Pasteur Street, 02-093 Warsaw, Poland
2. School of Psychological Science, Tel Aviv University, PO Box 39040, Tel Aviv, Israel
3. Sagol School of Neuroscience, Tel Aviv University, PO Box 39040, Tel Aviv, Israel

Corresponding author:

Michał Bola, PhD

Email: m.bola@nencki.edu.pl

Laboratory of Brain Imaging

Nencki Institute of Experimental Biology

Polish Academy of Science

3 Pasteur Street

02-093 Warsaw, Poland

Conflict of interest: The authors declare no competing interests.

Acknowledgements: This study was funded by a National Science Center Poland grant (grant number 2018/29/B/HS6/02152) and was conducted in the context of Action CA18106 funded by COST (European Cooperation in Science and Technology). LM is a CIFAR Tanenbaum Fellow in the Brain, Mind, and Consciousness program. We thank Florence Rémy for sharing the stimuli set, and Łucja Doradzińska and Natalia Rutkowska for assistance with data collection.

Author contributions: MB: conceived study, designed study, analyzed data, interpreted data, wrote manuscript; MF: programmed experimental procedure, collected data, analyzed data; LM: conceived study, designed study, interpreted data, revised manuscript;

Abstract

The global-to-local theories of perception assume that the gist of a scene is computed early and automatically, whereas recognition of objects occurs at a later processing stage, requires attentional resources, and is primed by the representation of gist. To test these theoretical predictions, we investigated the processing hierarchy of gist- and object-recognition and their interaction in two experiments (total N = 60). Backward-masked images of real-world scenes were presented for a range of brief durations - between 8 ms and 100 ms, and participants performed either an object or a background classification task, in separate blocks. We report three main findings. First, scenes' backgrounds were generally classified more accurately than foreground objects, but recognition of objects was boosted to the same level as backgrounds by cueing spatial attention to the exact object's location. Second, backgrounds influence objects' recognition, as objects presented within semantically incongruent backgrounds were classified less accurately. Third, objects influence background categorization, as backgrounds comprising incongruent objects were also classified less accurately. Therefore, the first two findings support the global-to-local theories, implying that gists are indeed more readily perceived than objects, probably at an earlier stage. Yet the latter finding that objects also influence gist recognition suggests a more parallel and interactive view of both processes than previously assumed.

Introduction

Do we first see the trees, and then understand that we are looking at a forest? Or rather realize that we are in a forest, and only then notice the trees? These two scenarios represent predictions of two different, seemingly incompatible theoretical models - one assuming perception is an analytic, local-to-global process; the other defining it as a holistic, global-to-local process. According to the analytic view, the local level is accessed early and automatically, whereas recognition of the global level occurs later and is based on subsequent integration of the elementary local features (Gibson, 1966; Riesenhuber and Poggio, 1999; Di Carlo et al., 2012). Findings from the early neurophysiological research on receptive fields' organization have been interpreted as uncovering a putative mechanism of such local-to-global integration, and thus providing support for the analytic models (Hubel and Wiesel, 1962). In contrast, the holistic model assumes that perception unfolds in a global-to-local manner. Accordingly, this approach emphasizes temporal precedence and general dominance of the global level, which is recognized early and in an automatic manner, and influences the subsequent recognition of local parts (Navon, 1977; Kimchi, 1992; Hochstein and Ahissar, 2002; Campana and Tallon-Baudry, 2013). While initially such models lacked a hypothesis regarding the neurophysiological mechanism, it is now clear that extensive feedback connectivity provides the neural scaffolding necessary for the global-to-local processing (Lamme and Roelfsema, 2000; Bullier, 2001).

Real-world scenes are a common example of hierarchical structures, with single objects representing the local level, and the *gist* and general meaning of a scene representing the global level. However, even though analytic and holistic theories differ fundamentally in their predictions and understanding of perception, it is still not known which one provides a better model to describe the scene recognition process. Evidence supporting the global-to-local view comes from studies showing that the gist of a complex scene can be recognized at a glance, with display times below 20 ms (Potter, 1975; Intraub, 1981; Rousselet et al., 2005; Fei-Fei et al., 2007; Joubert et al., 2007; Greene and Oliva, 2009). While this effect is indeed well established, these studies typically focused on the global level only and thus lacked direct comparisons between the global and local-level perception. Therefore, collected data does not allow drawing definitive conclusions regarding the hierarchy of visual processing and the contribution of local objects to the gist recognition, or vice versa. A parallel line of research demonstrates that classification of objects in scenes can be performed very accurately at comparable display times (e.g. Thorpe et al., 1996; Rousselet et al., 2002; Kirchner and Thorpe, 2006; Joubert et al., 2008; for review: Fabre-Thorpe, 2011). This might be considered to support the analytic view, but since these studies used images in which objects occurred in their typical context, the scenes' gist could have significantly contributed to the rapid and efficient categorization of objects. Moreover, the same caveat of the scene-processing studies applies here too: these studies also focused on one level only, this time the local one.

Yet to comprehensively describe the hierarchical processing of real-world scenes, and the interplay between local objects and global gist processing, a direct comparison of both levels is needed. To do so, several conditions should be met. First, both object- and gist-recognition should be probed using the same task and in the same group of participants. Second, the set of scene images should manipulate gists and objects in a full-factorial manner, so the two dimensions are orthogonal and not predictive of each other. Third, recognition should be tested for only one of the two levels at a time (hereby rendering the other one task-irrelevant), to make sure that the limit of performance for a given level is indeed reached, and minimize the potential effects of participants' biases or strategies to attend primarily to one of the levels (at the cost of the other one). The first two criteria were met only by two previous studies investigating perception of scenes (Davenport and Potter, 2004; Leroy et al., 2020), but both simultaneously probed object and gist recognition in the same trial, so their results might be affected by the aforementioned biases. Moreover, these studies yielded conflicting results, with one suggesting superior object processing (Davenport & Potter, 2004), while the other supporting preferable gist processing (Leroy et al., 2020). Thus, there is currently no direct comparison of gist and object processing during viewing of real-life scenes, which meets all three conditions, and existing findings are mixed with respect to this question (for a study with artificial stimuli meeting all criteria, see Campana et al., 2016). Additionally, both studies used fixed display times, either 80 ms in Davenport and Potter (2004) or 130 ms in Leroy et al (2020), which is a limitation when aiming to establish the hierarchy of processes since it is not known whether the observed differences in recognition of backgrounds and objects generalize across other display times, or is specific for the one selected. Finally, both studies used an open, basic-level recognition task, in which participants had to name

the displayed background or object. However, the difficulty of such an open recognition task might differ between backgrounds and objects, and results might be affected by participants' criteria or biases to a greater degree than in a 2-alternative forced choice task.

A second important point on which the analytic and holistic models differ is the direction of interaction between representations of local and the global levels. The global-to-local view assumes that the global gist is not only recognized earlier and dominates perception, but that it guides and affects subsequent exploration of the scene and recognition of objects (Bar, 2004). Indeed, objects that are semantically incongruent with the gist, namely those with a very low probability of occurring in a given context (Biederman et al., 1982), are identified slower and less accurately than congruent objects (Boyce et al., 1989; Davenport and Potter, 2004; Davenport, 2007; Joubert et al., 2008; Fize et al., 2011; Mack et al., 2017; Leroy et al., 2020), and their processing engages additional attentional resources (Vö and Henderson 2009; Mudrik et al., 2011). However, there is emerging evidence that representations of objects also affect representations of gist, as the presence of an incongruent object was found to impair gist recognition (Davenport and Potter, 2004; Davenport, 2007; Leroy et al. 2020). This observation is compatible with the analytic view, but challenges the global-to-local models, which assume gist- and object-recognition to unfold in a rather serial manner, with the former influencing the latter, but not vice versa. However, to the best of our knowledge none of the previous studies has convincingly shown that the influence of objects on gist is indeed automatic.

The present study was therefore designed to test predictions of the local-to-global and global-to-local models in the context of real-world scene perception. We specifically focused on recognition of backgrounds and objects in scenes, as reflecting the global and local levels, respectively, and investigated their hierarchy and interactions between their representations. In two experiments, participants were presented with briefly displayed and backward-masked images of real-world scenes, and asked to categorize either the scene's background or the foreground object as natural or artificial. Crucially, when designing the study we aimed to meet the three criteria needed to uncover the hierarchy of visual processing. First, we used a set of scenes in which backgrounds and objects were combined in a full-factorial way to create a set of congruent (e.g. a lion in a desert; a table in a kitchen) and incongruent scenes (a lion in a kitchen; a table in a desert). Thus, backgrounds were not predictive of objects, and vice versa. Second, we probed perception of both backgrounds and objects in the same group of participants and with the same 2-alternative forced-choice categorization task, which allowed testing recognition of both levels in relation to the same category (i.e. being natural or artificial). Third, we tested classification of each level in a separate block and defined the other level as task-irrelevant. This allowed testing recognition of a given level without participants' need to divide attentional resources, and assured that potential influence of the other, task-irrelevant level is indeed automatic. Additionally, in our study stimuli were presented across a whole range of display times - between 8 ms and 100 ms - to provide a finer-grained picture of the hierarchy of background and objects recognition and their interactions. Such a study design allowed us to test the following three hypotheses. First, that background will be classified more accurately than objects, across a range of presentation times. Here our

rationale was as follows: if the alternative is true, and object identification is needed to construct scene gist, then accuracy should be higher for object recognition, especially in shorter exposure duration (as objects would be recognized, but not to a sufficient extent for gist construction); but if gist is first extracted and used to guide object processing, backgrounds should be more readily detected. Second, we hypothesized that incongruent backgrounds will automatically impair classification of objects, in comparison to congruent ones. And third, that incongruent objects will also automatically impair classification of backgrounds, in comparison to congruent ones. While the first two hypotheses are in line with the global-to-local theories, the third one is not, as it suggests a more parallel and interactive view of visual processing.

Method

All experimental procedures were approved by the Research Ethics Committee at the Nicolaus Copernicus University (KB 447/2020). All participants taking part in the study declared normal or corrected-to-normal vision and no history of mental or neurological disorders, and all provided written informed consent and received monetary compensation for their time (50 PLN = c.a. 12 EUR).

Experiment 1

Participants

Data of 28 participants (21 females, mean age = 25.7 years; SD = 3.3 years; range: 21 - 33 years) was included in the analysis. Data of 3 additional subjects was collected, but excluded due to their in compliance to follow experimental instructions. The sample size was not pre-calculated because, first, we used a 3-way ANOVA design ($2 \times 2 \times 4$), for which there are no well-established methods to conduct a power analysis; and second, appropriate data based on which such analysis could have been conducted were lacking. Our aim was thus to collect 30 valid datasets, which is double the number collected in previous studies (e.g Davenport and Potter, 2004; Leroy et al., 2020).

Apparatus and stimuli

The experimental procedure was written in the Presentation software (Neurobehavioral Systems, Albany, CA) and presented on a Flex ScanEV-2450 computer monitor. The viewing distance was 60cm, maintained by a chinrest.

We used a stimulus set developed by Rémy et al. (2013, 2020). The original set included 400 scene images, each depicting either a natural (always outdoors) or an artificial (man-made, always indoors) background, with a single foreground object which was either an animal or a piece of furniture (**Fig. 1**). Thus, congruent scenes presented an animal within a natural background or furniture within an artificial background, whereas incongruent ones depicted an animal within an artificial background or furniture within a natural background. The stimuli were fully counterbalanced, so that each object appeared both in a congruent and an incongruent background, and each background appeared with a congruent or an incongruent object. Twenty images in which the background included both natural and artificial elements

(e.g. a balcony with a view of the sky; a room with a window depicting trees) were excluded, in all of their versions. Therefore, our final sample of stimuli consisted of 320 fully counterbalanced images.

Our initial analysis of physical features indicated differences in the red-green spectrum between natural and artificial backgrounds. Therefore, all images were converted to grayscale, as previous studies found that it does not affect performance in a classification task (Delorme et al., 2000). The luminance and contrast of the black and white images was equated using the histMatch function from the SHINE toolbox (Willenbockel et al., 2010). After this operation all images had a luminance of 128.2 and a contrast of 55.4 ALU (Arbitrary Linear Units) as estimated by the SHINE toolbox.

In the experiment we used 10 images as backward masks (in each trial 1 out of 10 was chosen randomly). Images consisted of 280 rectangles of the same size, filled randomly with one of seven different gray shades (thus, each shade occurred 40 times). The mask images were equated for luminance and contrast, again using SHINE, yielding luminance of 126.2 and contrast of 49.1, akin to the parameters of the stimuli.

Figure 1. A) Examples of scene images used in the study. Each scene depicted a natural (outdoors) or artificial (indoors) background, and a natural (animal) or artificial (furniture) foreground object. The set was created in a full-factorial way: each background and each object occurred in a congruent and an incongruent version. B) A schematic depiction of a trial sequence. Different display times were used in experiment 1 and experiment 2. The fixation cross was presented centrally, except for the object task in experiment 2, in which the cross position indicated the location of an object within the scene. Participants performed either a background classification task or an object classification task, in separate blocks. In both tasks they classified the respective element of the scene as natural or artificial, by pressing one of two response keys.

Procedure

The experiment consisted of two tasks (blocks): a gist (background) classification task and an object classification task. Within each block, half of the scenes were congruent and half incongruent. Finally, stimuli were displayed for four durations: 16, 32, 64 or 100 ms. Thus, the experiment had a 2 (background task, object task) x 2 (congruent scene; incongruent scene) x 4 (display times) design. Each task (block) included 640 trials (160 per display time; 80 congruent and 80 incongruent). Block order was

counterbalanced between subjects, and within each block the trial order was fully random. Three self-paced breaks were given during each block.

Each trial began with a 500 ms fixation cross (subtending $0.4^\circ \times 0.4^\circ$ of the visual angle) presented in the center of the screen. Then, the scene image was presented centrally (for one of the four durations) and followed by a backward mask presented for 50 ms. The scene and mask images subtended $13.8^\circ \times 9.2^\circ$ visual angle. Next, a question mark was presented centrally, which prompted participants to respond - by pressing one of two buttons using the index finger of their left or right hand - whether the background/object (depending on the block) was natural or artificial. Button-response assignment was fixed within the experimental session, but counterbalanced across participants. Subjects were asked to respond as quickly and accurately as possible. Response time was unlimited and the next trial started immediately after the manual response.

Experiment 2

In the second experiment data of 32 subjects (20 females, mean age = 28.4 years; SD = 4.3 years; range: 19 - 38 years) were analyzed. Data of 2 additional subjects was collected, but excluded due to lack of compliance with the task instructions. Similarly to Experiment 1, we aimed to collect 30 valid datasets.

We used the same procedure as in experiment 1, but introduced several changes. First, considering relatively good classification accuracy at 16 ms display time in experiment 1, in experiment 2 we included an 8 ms display time. To this end we used the DELL AW2518HF (24.5") monitor, as it is characterized by a 120Hz refresh rate. To keep the design of the experimental procedure and number of trials similar to experiment 1, the 64 ms presentation time was not included. Second, in experiment 2 masks were similar grids of 240 rectangles filled with 7 different shades of gray, but this time edges were located at random locations. This created an irregular pattern and thus introduced greater variability of spatial frequencies in mask images. Third, in experiment 2 in the object classification task the fixation cross preceding each image was always presented at the location of the to be classified object. Thus, attention was cued to the exact location of an object, to investigate the role of attentional resources in object recognition. Before the experiment, the location of the object was determined individually for every image and during the procedure the coordinates were used as the fixation cross location. Fourth, in each trial participants used a four-point scale to evaluate their confidence regarding the categorization task response. Confidence data are not reported here, as they address a different research question.

Analysis and statistics

The d' index was used as the measure of perceptual sensitivity. To calculate d' , which indicates to what extent an observer is able to distinguish signal from noise, one category (natural) was arbitrarily defined as signal, while the other category (artificial) as noise (importantly, a vice versa assignment would result in exactly the same estimates). Hits and false alarms equal to 0 or 1 for each subject were replaced using the log-linear rule, the least biased method of correcting extreme values (Stanislaw and Todorov 1999). Subjects were informed that response time is unlimited and thus trials were not excluded based on the response time.

Statistical analyses of d' were conducted using the JASP software and cross-checked with Statcheck (<http://statcheck.io/index.php>). Values are reported as Mean $\pm$ SEM. For all statistical tests, probability values were reported (p) and the standard 0.05 alpha level was used as a threshold for rejecting the null hypothesis.

Repeated-measures ANOVA models were used to analyze d' values. The within-subjects factors were: task (object, background), congruency (congruent, incongruent), and presentation time (16 ms, 32 ms, 64 ms, 100 ms in experiment 1; 8 ms, 16 ms, 32 ms, 100 ms in experiment 2). Results were reported as $F(df)$ and partial eta-squared, the indicator of the effect size, was reported as η_p^2 .

Additionally, to investigate whether participants' performance was better than chance-level at shortest display times, the d' index calculated for shortest presentation times was compared to 0 at the group level. Data distribution was first tested with the Shapiro-Wilk test. A two-tailed one-sample t-test or a Wilcoxon test (depending whether the distribution was Gaussian) was used and results were reported as $t(df)$ or W , respectively.

Data availability statement

Raw data in the form of logs, aggregated data used for statistical analysis, and scripts used to run the experimental procedure are available from OSF (<https://osf.io/cwnrv/>).

Results

We used a three-way repeated-measures ANOVA to investigate how the performed task, the semantic congruency of an image, and its presentation time, affected recognition abilities estimated by the d' index (**Fig. 2**). In the first experiment, we found the expected main effect of *presentation time* ($F(3,81) = 188.35$, $p < 0.001$, $\eta_p^2 = 0.875$), indicating that longer display times were associated with better performance. We also found an effect of a *task* ($F(1,27) = 17.33$, $p < 0.001$, $\eta_p^2 = 0.391$), suggesting that participants were better at recognizing backgrounds (scene gist; $M = 2.83 \pm 0.14$) than objects ($M = 2.44 \pm 0.11$), in line with our hypothesis. Finally, we found an effect of *congruency* ($F(1,27) = 28.76$, $p < 0.001$, $\eta_p^2 = 0.516$), so that recognition of both backgrounds and objects was better when the scene was congruent ($M = 2.84 \pm 0.14$) than incongruent ($M = 2.43 \pm 0.11$). None of the interactions reached the significance level (all $p > 0.05$).

In the second experiment we observed the main effect of *presentation time* ($F(3,93) = 297.27$, $p < 0.001$, $\eta_p^2 = 0.947$), *task* ($F(1,31) = 5.47$, $p < 0.026$, $\eta_p^2 = 0.150$), and *congruency* ($F(1,31) = 29.22$, $p < 0.001$, $\eta_p^2 = 0.485$) and, additionally, a significant interaction between *task* and *congruency* ($F(1,31) = 6.16$, $p = 0.019$, $\eta_p^2 = 0.166$). To explore the interaction we conducted analysis of simple main effects, which indicates that accuracy was similar in both tasks when presented scenes were congruent ($F(1) = 1.018$, $p = 0.28$), but recognition of backgrounds was better than objects when scenes were incongruent ($F(1) = 9.06$, $p = 0.005$). Thus, when spatial attention was cued to the object's location it was recognized as accurately as the background, but only when presented in a congruent context. The effect of congruency was present in both

the object- ($F(1) = 24.40$, $p < 0.001$) and the background-detection task ($F(1) = 8.05$, $p = 0.008$). None of the other interactions reached the significance level.

Further, to investigate the limits of real-world scenes categorization we tested whether participants showed above-chance performance in the shortest display-time conditions. In the first experiment, where backward-masked scenes were presented for 16 ms, we indeed found above chance-level performance in both the background recognition (congruent trials: $M = 1.55 \pm 0.15$, $t(27) = 10.2$, $p < 0.001$; incongruent trials: $M = 1.31 \pm 0.14$, $W = 406$, $p < 0.001$) and the object recognition task (congruent trials: $M = 1.37 \pm 0.11$, $t(27) = 11.7$, $p < 0.001$; incongruent trials: $M = 0.95 \pm 0.10$, $t(27) = 9.0$, $p < 0.001$). Crucially, in the second experiment, in which backward-masked scenes were displayed for only 8 ms, we also found above chance performance in the background recognition task (congruent trials: $M = 0.68 \pm 0.10$, $t(31) = 6.7$, $p < 0.001$; incongruent trials: $M = 0.49 \pm 0.07$, $W = 493$, $p < 0.001$) and in the object recognition task (congruent trials: $M = 0.55 \pm 0.09$, $t(31) = 5.6$, $p < 0.001$; incongruent trials: $M = 0.30 \pm 0.08$, $t(31) = 3.4$, $p = 0.001$). Thus, even with 8 ms display time, participants were able to classify not only backgrounds but also the objects, and not only when the scenes were congruent, but also when they were incongruent.

Figure 2. Classification performance, as indicated by the d' index, in experiment 1 (two upper panels) and experiment 2 (two lower panels). Upper and lower panels present the same data, in the upper panel data averaged across participants are presented, in the lower panel individual data are plotted. The line within the gray box indicates the median, while boxes' lower and upper borders indicate 1st and 3rd quartiles.

Discussion

How does visual perception unfold over time? The analytic theories of vision assume that we perceive the local level of objects and details first, and that representation of the global level of meaning is created by subsequent integration of these simpler elements (Hubel and Wiesel, 1962; Gibson, 1966; Riesenhuber and Poggio, 1999; Di Carlo et al., 2012). In stark contrast, the holistic, global-to-local theories assume that our initial visual impression is dominated by the meaning and global properties of a scene. According to these models, it is the global level that is recognized early and automatically, whereas recognition of local objects occurs at a later processing stage, requires attentional resources, and is primed by the global-level information (Navon, 1977; Kimchi, 1992; Hochstein and Ahissar, 2002; Bar, 2004; Oliva and Torralba, 2006; Campana and Tallon-Baudry, 2013). Real-world scenes are a prime example of complex and highly hierarchical naturalistic stimuli, and thus provide a yet untapped opportunity to directly test the fundamentally different predictions of both models.

Gist is classified more accurately than objects

The main finding of the present study is that the global level of a scene, defined as the scenes' background, was classified more accurately than the local level, defined as foreground objects. This effect was observed across a whole range of display times used in the study, including the shortest exposure durations of 16 ms and 8ms. While perception of both the global-level gist and local-level objects has already been extensively investigated (review: Bar, 2004, Oliva and Torralba, 2006; Fabre-Thorpe, 2011), what prevented drawing conclusions regarding hierarchical processing was that each aspect was typically tested separately. Thus, the current study goes beyond previous ones in the direct comparison of both levels - using data collected from the same participants, with the same task, and the same stimuli in which both levels are orthogonal - to uncover the hierarchical nature of the scene perception process. Importantly, two previous studies that are similar to our work in this respect provided conflicting findings. Davenport and Potter (2004, experiment 4) asked participants to recognize backgrounds and foreground objects in scenes presented for 80 ms and backward-masked. They found that objects were in fact recognized more accurately than backgrounds, which is in disagreement with our results and with predictions of the global-to-local theories. In contrast, Leroy and colleagues (2020) found more accurate recognition of backgrounds than objects across five experiments, in which scenes were displayed for 130 ms and backward masked. However, what we consider a limitation of both studies is that participants were asked to recognize both background and object in every trial (i.e. both aspects were task-relevant), and thus results might have been affected by the participants' need to divide attention between levels or their preference to focus primarily on one of them (which was not controlled for). Conversely, in our study participants were exclusively focused on either the global or the local level, which allowed probing the true limits of each level's recognition. Our results accordingly strengthen the results of Leroy's and colleagues (2020), while excluding potential confounding factors that might have explained their findings. Our study is also compatible with work of Campana and colleagues (2016), who - by using artificial stimuli similar to the Navon letters and a similar block design, with only one level being task-relevant - found that participants reacted faster to the global-

than to the local-level of such stimuli, and the global-level was decoded from the EEG response earlier than the local-level. Taken together, these studies join our results in supporting the global-to-local accounts, suggesting that gist is processed more readily, and earlier, than objects.

There are two more methodological aspects which reflect the robustness of conclusions drawn from the discussed studies. First, both Davenport and Potter (2004) and Leroy and colleagues (2020) used a basic-level recognition task, in which participants had to name the displayed background or object. This raises the concern that the difficulty of such an open recognition task differed between backgrounds and objects, and that results might be affected by criteria or biases (e.g. when participants decide they did not see clearly enough to try to name a stimulus). We therefore used a 2-alternative forced choice task in which participants had to judge both backgrounds and objects on the very same dimension (i.e. being natural or artificial). Because participants judged stimuli on the super-ordinal level, held to be recognized before the basic level (Loschky and Larson, 2010), and were forced to respond in every trial (even if it was a guess), our study tested the true limits of perception, allowing less room for participants' criteria or biases to affect results. Second, both Davenport and Potter (2004) and Leroy et al (2020) used fixed display times, which were 80 ms and 130 ms, respectively. We consider this a limitation when drawing conclusions regarding the hierarchy of processes, since it is not known whether the observed difference in recognition accuracy between backgrounds and objects generalizes across other display times, or is specific for the one selected. Therefore, our study included a systematic manipulation of display times, which allowed us to demonstrate if the effects are found across a range of presentation durations (i.e. congruency or task did not interact with display time in ANOVA models). Additionally, we were also able to show that performance in both background and object task plateaus at 64 ms exposure, with d' values > 3 (in line with Bacon-Mace et al., 2005), and that presentation durations as short as 16 ms and 8 ms, resulted in above-chance classification in both tasks (Van Rullen and Koch, 2003). Such accuracy levels observed for briefly displayed and backward masked stimuli support the previously demonstrated ability to analyze complex real-world scenes very rapidly and efficiently (Thorpe et al., 1996; Greene and Oliva, 2009).

Our conclusions strongly aligns with the Reverse Hierarchy Theory (Hochstein and Ahissar, 2002), according to which our initial visual impression is dominated by global-level features, which are established by a rapid and automatic feedforward sweep of activity along the ventral stream, activating the top ventral stream areas. Conversely, recognizing local features requires feedback activity, which gradually activates lower cortical areas and consequently adds more specific details to the visual representation. Thus, because of this reverse hierarchy, the lower the feature is represented in the cortical hierarchy, the more time is needed to activate it and add to the percept. We were able to test this crucial prediction of RHT by using backward masking, which acts by disturbing feedback activity (Fahrenfort et al., 2007), and systematically varying the backward-mask onset time. Therefore, our data support RHT by showing that recognition of objects is more affected by impairing reentrant activity than recognition of high-level backgrounds.

An important finding from the second experiment is that the object-classification accuracy can be at the same level as background-classification accuracy, provided that spatial attention is cued to the object's location. This finding does not necessarily challenge global-to-local models, as RHT strongly emphasizes that attention, the neural basis of which is ventral stream feedback activity, is required for object recognition (Hochstein and Ahissar, 2002). Importantly however, in none of the display times attention cueing boosted object-classification accuracy to exceed the background-classification accuracy, as would have been predicted from the analytic perspective. Also, classification accuracy was increased by attention only for objects presented in congruent backgrounds, and not when objects were presented in an incongruent context (i.e. as indicated by a significant *task x congruency* interaction; see Fig. 2). This supports the crucial role of the global level in priming and recognition of the local level, in line with the global-to-local view. Future studies will investigate the interaction between attention and congruence in more detail.

Background affects object recognition, and vice versa

What constitutes a scene is that both the elements of background and discrete objects are arranged according to certain spatial and semantic rules (Henderson and Hollingworth, 1999). The importance of such rules has been emphasized by global-to-local models, which assume that gist facilitates recognition of objects congruent with the context and thus expected in a given location, but might hamper recognition of the incongruent and unexpected ones (Bar, 2004). Here, we show that objects located within an incongruent background are indeed categorized less accurately, which confirms the theoretical prediction and replicates results of several previous studies (e.g. Davenport and Potter, 2004). Importantly, we also found less accurate classification of backgrounds in the presence of an incongruent object. This observation is particularly important, as it seems to be at odds with the discussed global-to-local theoretical models, and suggests that representations of objects do contribute to the gist recognition process, and thus the influence is bidirectional, rather than unidirectional. Importantly, while several previous studies found a similar effect (Davenport and Potter, 2004; Davenport, 2007; Leroy et al., 2020), our study is the first to show it is automatic, as it was caused by objects that are task-irrelevant and observed across the entire range of presentation times, including the shortest ones. Such an automatic computation of objects and their influence on gist classification is particularly problematic for the global-to-local theories, which assume that recognition of objects and details is an optional processing step, performed only when required by the task (Hochstein and Ahissar, 2002; Camapan and Tallon-Baudry, 2013).

However, an important limitation of our work is that it does not address at which processing stage objects are influenced by the gist, and vice versa. There has been an ongoing discussion whether such interactive effects occur at the perceptual stage, as proposed by the perceptual schema model (e.g. Biderman et al., 1982) and the object selection model (e.g. Mudrik et al., 2010), or at the post-perceptual, decision-related stage, as assumed by the functional isolation model (e.g. Hollingworth, 1998). Importantly, under the latter interpretation our results regarding influence between levels do not speak to theories of perception, as then the effect of incongruency stems from response-conflict, rather than a processing difficulty. While our study

was not designed to directly address this question, we consider the functional isolation interpretation less likely. The functional isolation model has been initially proposed based on the observation that many early studies did not correct for a bias and guessing in a recognition task (i.e. performed an analysis of accuracy rates without accounting for false alarms; see Hollingworth, 1998). However, more recent studies, which controlled and accounted for sources of bias, provide support for the perceptual locus of the effect (Leroy et al., 2020). Similarly, studies using electrophysiology already demonstrated that context affects objects recognition early, at the perceptual stage (Mudrik et al., 2010; Brandman and Peelen, 2017). Thus, though we cannot rule it out, we find the interpretation assuming a post-perceptual locus of the observed effects less plausible. What is more, this potential limitation only relates to the interactions between objects and gist, and accordingly do not pertain to our main conclusions regarding the hierarchy of backgrounds and objects perception. Notably however, the locus of interaction has so far been discussed only in the context of gist influencing objects recognition, but it is not known whether objects affect gist at the same processing stage (see Brandman and Peelen, 2019). Therefore, future studies should shed more light on mechanisms of bidirectional interactions between gist and objects.

Assuming the effect does occur at the perceptual level, it can be explained by two potential mechanisms. First, the dual-process model, based on the observation that gist- and object- processing are separated at the neuronal level (Harel et al., 2013; Wischniewski and Peelen, 2021). Accordingly, scene- and object-selective regions compute semantic representations of the respective scene aspects, and only then they influence each other (Brandman and Peelen, 2017, 2019). However, an alternative explanation is that performance in the gist-categorization task, and possibly also in the object-categorization task, is based solely on global scene statistics - like spatial frequency (Oliva and Torralba, 2001; but see Loschky and Larson, 2008). Under this account, the influence of objects on gist recognition is not based on recognition of objects semantics, but rather on biasing the global scene statistics that are diagnostic for the categorization task (e.g. a roundy hippo adds more low spatial-frequency and affects the ensemble texture in different way than an edgy chair). This view is supported by computational work showing that influence of an incongruent object can be accounted for using scene statistics only (Mack and Palmeri, 2010). On the other hand, finding a congruence effect even when a scene and object are presented separately, in different screen locations, goes against this account and supports a semantic component (Leroy et al., 2020).

Conclusions

In conclusion, our study shows that scenes' backgrounds are categorized more accurately than foreground objects, and demonstrates the key role of attention in the object recognition process. Further, our data support global-to-local models of perception by demonstrating an automatic influence of context on the object categorization process, but also reveal that objects affect gist recognition, and such a bidirectional interaction has not been predicted by these models. Further research is needed to determine to what extent these findings generalize across different tasks and stimuli sets, and to investigate if these effects are perceptual, or rather post-perceptual and decision-related. An additional goal for future studies is elucidating

the neural mechanisms behind the behavioral findings, particularly the role of feedforward and feedback connectivity, and the pathways of interaction between gist and object representations.

References

1. Bacon-Macé, N., Macé, M. J. M., Fabre-Thorpe, M., & Thorpe, S. J. (2005). The time course of visual processing: Backward masking and natural scene categorisation. *Vision research*, 45(11), 1459-1469.
2. Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5(8), 617-629.
3. Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive psychology*, 14(2), 143-177.
4. Boyce, S. J., Pollatsek, A., & Rayner, K. (1989). Effect of background information on object identification. *Journal of Experimental Psychology: Human Perception and Performance*, 15(3), 556.
5. Brandman, T., & Peelen, M. V. (2017). Interaction between scene and object processing revealed by human fMRI and MEG decoding. *Journal of Neuroscience*, 37(32), 7700-7710.
6. Brandman, T., & Peelen, M. V. (2019). Signposts in the fog: Objects facilitate scene representations in left scene-selective cortex. *Journal of cognitive neuroscience*, 31(3), 390-400.
7. Bullier, J. (2001). Integrated model of visual processing. *Brain research reviews*, 36(2-3), 96-107.
8. Campana, F., & Tallon-Baudry, C. (2013). Anchoring visual subjective experience in a neural model: the coarse vividness hypothesis. *Neuropsychologia*, 51(6), 1050-1060.
9. Campana, F., Rebollo, I., Urai, A., Wyart, V., & Tallon-Baudry, C. (2016). Conscious vision proceeds from global to local content in goal-directed tasks and spontaneous vision. *Journal of Neuroscience*, 36(19), 5200-5213.
10. Davenport, J. L. (2007). Consistency effects between objects in scenes. *Memory & Cognition*, 35(3), 393-401.
11. Davenport, J. L., & Potter, M. C. (2004). Scene consistency in object and background perception. *Psychological science*, 15(8), 559-564.
12. Delorme, A., Richard, G., & Fabre-Thorpe, M. (2000). Ultra-rapid categorisation of natural scenes does not rely on colour cues: a study in monkeys and humans. *Vision research*, 40(16), 2187-2200.
13. DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How does the brain solve visual object recognition?. *Neuron*, 73(3), 415-434.
14. Fabre-Thorpe, M. (2011). The characteristics and limits of rapid visual categorization. *Frontiers in psychology*, 2, 243.
15. Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. (2007). Masking disrupts reentrant processing in human visual cortex. *Journal of cognitive neuroscience*, 19(9), 1488-1497.
16. Fei-Fei, L., Iyer, A., Koch, C., & Perona, P. (2007). What do we perceive in a glance of a real-world scene?. *Journal of vision*, 7(1), 10-10.
17. Fize, D., Cauchoux, M., & Fabre-Thorpe, M. (2011). Humans and monkeys share visual representations. *Proceedings of the National Academy of Sciences*, 108(18), 7635-7640.

18. Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.
19. Greene, M. R., & Oliva, A. (2009). The briefest of glances: The time course of natural scene understanding. *Psychological Science*, 20(4), 464-472.
20. Harel, A., Kravitz, D. J., & Baker, C. I. (2013). Deconstructing visual scenes in cortex: gradients of object and spatial layout information. *Cerebral Cortex*, 23(4), 947-957.
21. Henderson, J. M., & Hollingworth, A. (1999). High-level scene perception. *Annual review of psychology*, 50(1), 243-271.
22. Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36(5), 791-804.
23. Hollingworth, A. (1998). Does consistent scene context facilitate object perception?. *Journal of Experimental Psychology: General*, 127(4), 398.
24. Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1), 106-154.
25. Intraub, H. (1981). Rapid conceptual identification of sequentially presented pictures. *Journal of Experimental Psychology: Human Perception and Performance*, 7(3), 604.
26. Joubert, O. R., Rousselet, G. A., Fize, D., & Fabre-Thorpe, M. (2007). Processing scene context: Fast categorization and object interference. *Vision research*, 47(26), 3286-3297.
27. Joubert, O. R., Fize, D., Rousselet, G. A., & Fabre-Thorpe, M. (2008). Early interference of context congruence on object processing in rapid visual categorization of natural scenes. *Journal of Vision*, 8(13), 11-11.
28. Kimchi, R. (1992). Primacy of wholistic processing and global/local paradigm: a critical review. *Psychological bulletin*, 112(1), 24.
29. Kirchner, H., & Thorpe, S. J. (2006). Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. *Vision research*, 46(11), 1762-1776.
30. Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in neurosciences*, 23(11), 571-579.
31. Leroy, A., Faure, S., & Spotorno, S. (2020). Reciprocal semantic predictions drive categorization of scene contexts and objects even when they are separate. *Scientific reports*, 10(1), 1-12.
32. Loschky, L. C., & Larson, A. M. (2008). Localized information is necessary for scene categorization, including the natural/man-made distinction. *Journal of Vision*, 8(1), 4-4.
33. Loschky, L. C., & Larson, A. M. (2010). The natural/man-made distinction is made before basic-level distinctions in scene gist processing. *Visual Cognition*, 18(4), 513-536.
34. Mack, A., Clarke, J., Erol, M., & Bert, J. (2017). Scene incongruity and attention. *Consciousness and cognition*, 48, 87-103.
35. Mack, M. L., & Palmeri, T. J. (2010). Modeling categorization of scenes containing consistent versus inconsistent objects. *Journal of vision*, 10(3), 11-11.
36. Mudrik, L., Lamy, D., & Deouell, L. Y. (2010). ERP evidence for context congruity effects during simultaneous object–scene processing. *Neuropsychologia*, 48(2), 507-517.

37. Mudrik, L., Deouell, L. Y., & Lamy, D. (2011). Scene congruency biases binocular rivalry. *Consciousness and cognition*, 20(3), 756-767.
38. Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive psychology*, 9(3), 353-383.
39. Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International journal of computer vision*, 42(3), 145-175.
40. Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, 155, 23-36.
41. Potter, M. C. (1975). Meaning in visual search. *Science*, 187(4180), 965-966.
42. Rémy, F., Saint-Aubert, L., Bacon-Macé, N., Vayssière, N., Barbeau, E., & Fabre-Thorpe, M. (2013). Object recognition in congruent and incongruent natural scenes: A life-span study. *Vision research*, 91, 36-44.
43. Rémy, F., Vayssière, N., Saint-Aubert, L., Bacon-Macé, N., Pariente, J., Barbeau, E., & Fabre-Thorpe, M. (2020). Age effects on the neural processing of object-context associations in briefly flashed natural scenes. *Neuropsychologia*, 136, 107264.
44. Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11), 1019-1025.
45. Rousselet, G. A., Fabre-Thorpe, M., & Thorpe, S. J. (2002). Parallel processing in high-level categorization of natural images. *Nature neuroscience*, 5(7), 629-630.
46. Rousselet, G., Joubert, O., & Fabre-Thorpe, M. (2005). How long to get to the “gist” of real-world natural scenes?. *Visual cognition*, 12(6), 852-877.
47. Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior research methods, instruments, & computers*, 31(1), 137-149.
48. Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *nature*, 381(6582), 520-522.
49. VanRullen, R., & Koch, C. (2003). Visual selective behavior can be triggered by a feed-forward process. *Journal of Cognitive Neuroscience*, 15(2), 209-217.
50. Võ, M. L. H., & Henderson, J. M. (2009). Does gravity matter? Effects of semantic and syntactic inconsistencies on the allocation of attention during scene perception. *Journal of Vision*, 9(3), 24-24.
51. Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: the SHINE toolbox. *Behavior research methods*, 42(3), 671-684.
52. Wischniewski, M., & Peelen, M. V. (2021). Causal Evidence for a Double Dissociation between Object- and Scene-Selective Regions of Visual Cortex: A Preregistered TMS Replication Study. *Journal of Neuroscience*, 41(4), 751-756.