

Impact of learning on representation of parts and wholes in monkey inferotemporal cortex

Chris I. Baker¹, Marlene Behrmann^{1–3} and Carl R. Olson^{1–3}

¹ Center for the Neural Basis of Cognition, 115 Mellon Institute, 4400 Fifth Avenue, Pittsburgh, Pennsylvania 15213, USA

² Department of Neuroscience, University of Pittsburgh, 446 Crawford Hall, Pittsburgh, Pennsylvania 15260, USA

³ Department of Psychology, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, Pennsylvania 15213, USA

Correspondence should be addressed to C.I.B. (cbaker@cnbc.cmu.edu)

Published online 15 October 2002; doi:10.1038/nn960

Here we investigated the impact of visual discrimination training on neuronal responses to parts of images and to whole images in inferotemporal (IT) cortex. Monkeys were trained to discriminate among ‘baton’ stimuli consisting of discrete top and bottom parts joined by a vertical stem. With separate features at each end, we were able to manipulate the two parts of each baton independently. After training the monkeys, we used single-cell recording to compare neuronal responses to learned and unlearned batons. Responses to learned batons, though not enhanced in strength, were enhanced in selectivity for both individual parts and for whole batons. Whole-baton selectivity arose from a form of conjunctive encoding whereby two parts together exerted a greater influence on neuronal activity than predicted by the additive influence of each part considered individually. These results indicate a possible neural mechanism for holistic or configural effects in expert versus novice observers.

Visual object recognition is thought to depend on experience-induced changes in inferotemporal (IT) cortex, such that neurons become more selective for (or more responsive to) learned images^{1–4}. This view is consistent with evidence showing that lesions in IT interfere with pattern recognition^{5,6}, that neurons in IT are pattern-selective^{5,6} and that IT is a site of experience-dependent plasticity. Plasticity has been shown in IT by the use of three approaches: (i) repeated exposure to a stimulus over a short period of time leads to a decline in response strength^{7–10}, (ii) prolonged training on a visual paired associate task results in the emergence of neurons that are responsive to both members of the pair^{11–13} and (iii) discrimination training. Training monkeys to discriminate among images is thought to induce changes in the strength and selectivity of neuronal responses to those images, but studies to date have produced contradictory and inconclusive results. On the one hand, some neurons seem to become markedly selective for learned images. For example, in a study of monkeys trained to discriminate among wire objects^{14,15}, several units showed “a remarkable selectivity” for individual views that the monkey had learned to recognize. On the other hand, two studies involving a single day¹⁶ or several months¹⁷ of training do not show evidence of enhanced selectivity for images in the training set. In another study¹⁸, discrimination training resulted in a subtle enhancement of stimulus selectivity at the population level, but this effect was shown by comparing trained to untrained monkeys rather than by comparing responses to learned and unlearned stimuli in the same monkey. Thus, innate differences between monkeys may have contributed to the result. Finally, learned images have been reported to elicit higher firing rates than unlearned images^{18–20}.

An important question not addressed in previous studies is whether discrimination training enhances neuronal selectivity for whole images or for the parts of those images. This question is particularly relevant to the idea that experts process images in a qualitatively different way from novices, placing more weight on wholes and less on parts^{21,22}. It has been suggested that this ability arises from IT neurons that become selective for combinations of features contained in learned images—neurons for which “the whole is greater than the sum of the parts”^{23–25}. To show the existence of such neurons, the following are required: (i) the parts used in the experiment must be far enough apart so that juxtapositional features do not emerge where they abut, (ii) the parts must be manipulated independently and systematically, and (iii) the effect on neuronal activity of manipulating the parts together must be stronger than the summed effects of manipulating them individually. Experiments in which removal of part of a complex image causes a marked reduction in neuronal visual response strength^{26–30} do not meet these requirements.

To test for the occurrence of nonlinear part–part interactions, we used baton-shaped stimuli that consisted of two distinct elements joined by a vertical stem (Fig. 1a). The batons were organized into tetrads representing the four possible combinations of two top and two bottom parts. By monitoring neuronal responses elicited by the batons in these tetrads, we were able to determine whether the neuron was selective for the parts (the firing rate could be modeled as the sum of independent responses to the top and bottom elements) or for the wholes (the firing rate depended nonlinearly on the conjunction of top and bottom elements). To assess the impact of learning on selectivity for parts

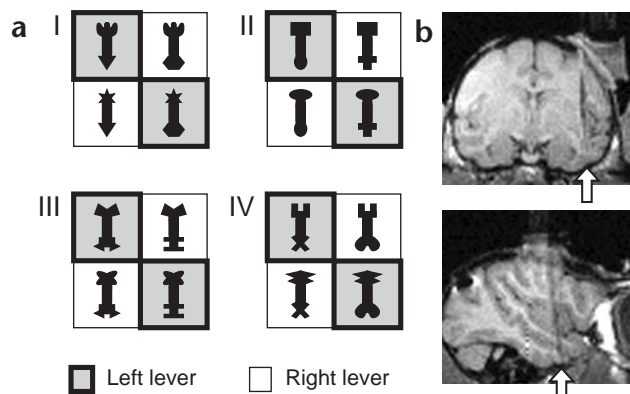


Fig. 1. Baton stimuli and recording location. **(a)** Four tetrads of batons were used in discrimination training. Monkey 1 was trained on tetrads I and II and monkey 2 on tetrads III and IV. The batons that were used for training for one monkey were also used as unlearned controls for the other monkey. Batons requiring right- and left-lever responses are indicated by white and gray backgrounds, respectively, although during experiments, the background was constant. **(b)** Coronal (top) and sagittal (bottom) magnetic resonance images showing recording locations in the right hemisphere of monkey I. The dark line running through the cortex is a shadow surrounding an electrode that was placed at the most medial recording site.

and wholes, we trained monkeys to discriminate among batons within tetrads, setting up the stimulus-response associations so that the monkeys had to take into account the conjunction of parts in each baton in order to perform above chance (Fig. 1a). Each monkey was trained on two tetrads of batons. Tetrads learned by one monkey were used as unlearned controls for the other monkey so that effects that were due to the intrinsic properties of the stimuli could be dissociated from the effects of training. Once the monkeys were able to perform the task (Table 1), we carried out single-neuron recording in IT (Fig. 1b). In each neuron, we directly compared the visual responses elicited by learned and unlearned batons. The results indicate that IT neurons are more selective for learned stimuli, both at the level of individual parts and at the level of whole batons.

RESULTS

Magnitude and selectivity of neuronal responses

To compare neuronal responses to learned and unlearned batons under identical conditions, we collected data while monkeys maintained steady fixation without generating lever responses. First, we briefly assessed responses to all 16 batons; then we used the learned and unlearned tetrads that elicited the strongest responses for the first session of data collection. If possible, we followed up with a second session involving the less effective learned and unlearned tetrads. (Note that the term ‘session’ as used here and throughout the paper denotes collecting data on neuronal responses to batons in one learned and one unlearned tetrad). During a session, batons from the learned and unlearned tetrad were presented foveally in pseudorandom, interleaved sequence until 16 trials had been completed for each baton, for a total of 128 trials. In each of 502 sessions involving a total of 360 neurons, at least one of the eight batons elicited a significant visual response (Table 2).

We first asked whether IT neurons respond more strongly to learned than to unlearned batons. We computed, for each session, the mean firing rates elicited by the best learned baton and

the best unlearned baton (‘best’ denotes the baton that elicited the strongest response of the four in the tetrad). The distribution of values obtained across all sessions (Fig. 2) shows that the average response elicited by the best learned baton was not significantly different from the average response elicited by the best unlearned baton (monkey 1: unlearned, 13.36 spikes/s; learned, 13.24; $P > 0.7$ paired *t*-test; monkey 2: unlearned, 10.89 spikes/s; learned, 11.07; $P > 0.5$ paired *t*-test). This finding stands in contrast to previous reports of enhanced response strength for learned stimuli^{18–20}.

We next asked whether selectivity for batons in learned tetrads was enhanced relative to selectivity for batons in unlearned tetrads. To address this issue, we examined how sharply response strength fell off from the best baton to other batons in a tetrad. For each session, we ranked batons in the learned tetrad according to the strength of the elicited visual response and did the same for batons in the unlearned tetrad. Then, to eliminate any effect of absolute response strength, we normalized the response elicited by each of the three low-ranked batons to the response elicited by the best baton in the tetrad. On combining the results from all sessions (Fig. 3a and b), we found that low-ranked batons in learned tetrads elicited weaker normalized responses than low-ranked batons in unlearned tetrads. Thus, firing fell off more sharply from the best baton to other batons in learned compared to unlearned tetrads, indicating that neurons were more sharply tuned for batons in learned tetrads. To assess the significance

of this effect, we carried out a repeated measures ANOVA across all sessions, with normalized response strength as the dependent variable, and with training status (learned or unlearned) and rank (2, 3 or 4) as factors. Rank-1 batons were excluded from this analysis because there was no variance in rank-1 values after normalization. Rank was included as a factor to assess whether the effects of training varied across rank. In each monkey, there was a significant main effect of training status: learned batons below rank 1 elicited weaker normalized responses than unlearned batons below rank 1 (monkey 1, $P < 0.00001$; monkey 2, $P < 0.035$). *Post-hoc* analyses revealed that the effect was significant at all ranks in monkey 1 and at two of three ranks in monkey 2 (asterisks in Fig. 3a and b). The mean difference in responses elicited by rank-1 and rank-4 learned batons was 5.5 spikes/s (versus a 4.5 spikes/s mean difference for unlearned batons). The finding that neurons were more selective for batons in learned tetrads cannot be explained in terms of the visual attributes of the batons because the tetrads that were learned by monkey 1 were unlearned for monkey 2, and vice-versa. The difference between subjects in the strength of this effect, however, could well be explained by the batons’ visual properties. We therefore place little weight on inter-monkey differences in effect

Table 1. Discrimination task performance.

Monkey	1		2	
	I	II	III	IV
Reaction time (ms): mean (s.d.)	366 (29)	368 (29)	325 (22)	314 (22)
Percentage correct: mean (s.d.)	86 (11)	91 (9)	94 (9)	96 (7)

Performance of monkey 1 was assessed during physiological data collection sessions. Data for monkey 2 are from behavioral training sessions on the three days immediately prior to recording. Number of sessions for tetrads I, II, III and IV were 26, 28, 42 and 42, respectively.

strength. We conclude that IT neurons, considered as a population, responded more selectively—but not necessarily more strongly—to learned batons.

The increase in selectivity measured across the entire population might reflect a moderate shift in many neurons or a dramatic shift in a few. To distinguish between these possibilities, we computed, for each session, an index of selectivity, $(b - w)/(b + w)$. The variables b and w are firing rates elicited by the best and worst batons in a tetrad. We found that the general pattern of the frequency distribution was similar for learned and unlearned batons (Fig. 3c). However, the learned distribution (median, 0.285) was shifted to the right (in the direction of greater selectivity) relative to the unlearned distribution (median, 0.229). This effect was significant in both monkeys (Wilcoxon matched pairs test: monkey 1, $P < 0.000001$; monkey 2, $P < 0.03$). Critically, we did not observe a second mode at the high end of the scale, as would have been created by a few neurons with a markedly enhanced selectivity for learned images. This finding stands in contrast to previous reports of neurons that are highly selective for particular learned images^{14,15}. We conclude that training exerted a moderate effect on the stimulus selectivity of many IT neurons, not a dramatic effect on the selectivity of a few.

To determine whether responses to learned batons were affected by task context, we compared responses elicited by a learned tetrad during both fixation and discrimination task performance. This analysis was based on 47 cases involving monkey 1 in which responses of the same neuron to the same stimuli were characterized in both tasks. To take into account the briefer presentation of the baton in the discrimination task, we confined the period of analysis to the epoch 130–230 ms after stimulus onset. This epoch began at the onset of the population response and continued for 100 ms (the duration of baton presentation in the discrimination task). To determine whether response magnitude differed systematically between tasks, we compared discharge rates elicited in fixation and discrimination contexts by the best baton in the tetrad. The mean firing rate during fixation was 17.71 spikes/s as compared to 17.84 spikes/s during task performance. The difference was not significant (paired t -test, $P > 0.89$). To determine whether the degree of selectivity depended on task context, we compared the selectivity index (defined above) between tasks. The mean values were 0.389 in fixation and 0.342 in discrimination—not significantly different (Wilcoxon matched pairs, $P > 0.21$). We conclude that the magnitude of the response and the degree of selectivity were hardly affected by task context, if at all. This result is consistent with previous reports indicating

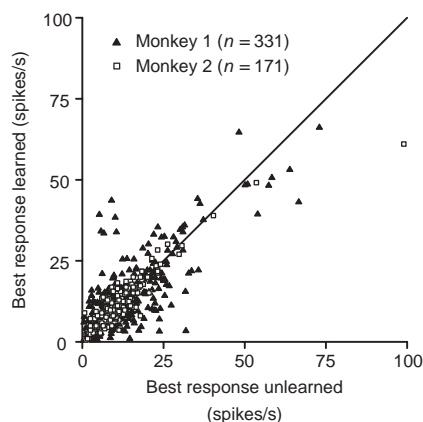


Table 2. Neuron and session counts.

		Total	Visually responsive
Monkey 1	Neurons	330	243
	Sessions	507	331
Monkey 2	Neurons	202	117
	Sessions	301	171

The number of data collection sessions was greater than the number of neurons studied because some neurons were used in two sessions that involved different pairs of learned and unlearned batons. Sessions that met the criterion of significant visual responsiveness were included in subsequent steps of data analysis.

that the effects of attention become prominent only when multiple stimuli compete to control neuronal activity³¹.

Selectivity for parts or wholes?

Enhanced selectivity for learned batons could arise from either a part-based or a whole-based mechanism. If a neuron's activity was affected consistently by the identity of a part at a given location on the baton, regardless of the feature at the other end of the baton, this indicated part-based selectivity. Some neurons, although responsive, were not selective (Fig. 4a). Other neurons were selective either for a part at just one location (Fig. 4b), or for the parts at both locations (Fig. 4c). To determine how frequently part-based selectivity occurred, we performed separate ANOVAs on responses to the learned and unlearned tetrads in each session, with identity of the top part and identity of the bottom part as factors and with the square-root transformed firing rate as the dependent variable (Fig. 4d–f). Out of 1,004 cases in which the identity of a part at a given location on a learned baton could have affected neuronal activity (502 sessions \times 2 locations), there were 366 cases (36%) in which a significant ($P < 0.05$) main effect was present. For unlearned batons, a significant main effect was present in only 293 cases (29%). In monkey 1, the higher incidence of selectivity for learned as compared to unlearned parts (43% versus 35%) was significant (χ^2 test, $P < 0.0024$). In monkey 2, the effect was of slightly greater magnitude, considered as a ratio of percentages (23% versus 18%), but did not attain significance because the number of observations was smaller (χ^2 test, $P < 0.075$). Because the ratio between the learned and the unlearned percentages was not significantly different between the monkeys (χ^2 test, $P > 0.9$), we carried out a test on the combined data, which revealed a highly significant effect (χ^2 test, $P < 0.0005$). We conclude that neuronal selectivity for individual parts was modestly enhanced for learned as compared to unlearned tetrads.

Having established that IT neurons were more selective for parts within learned tetrads, we asked whether they were more selective for batons considered as wholes. We took as a measure of whole-based selectivity the number of sessions in which the firing rate depended significantly on the interaction between top-part identity and bottom-part identity, as revealed by the two-

Fig. 2. Response to the best learned baton plotted against response to the best unlearned baton. Each point represents data from a session assessing the responses of one neuron to four batons from a learned tetrad and four batons from an unlearned tetrad. The 'best' baton in each tetrad was defined as the one eliciting the strongest response. There was no significant tendency in either monkey for responses elicited by the best learned baton to exceed those elicited by the best unlearned baton (paired t -test: monkey 1, $P > 0.7$; monkey 2, $P > 0.5$).

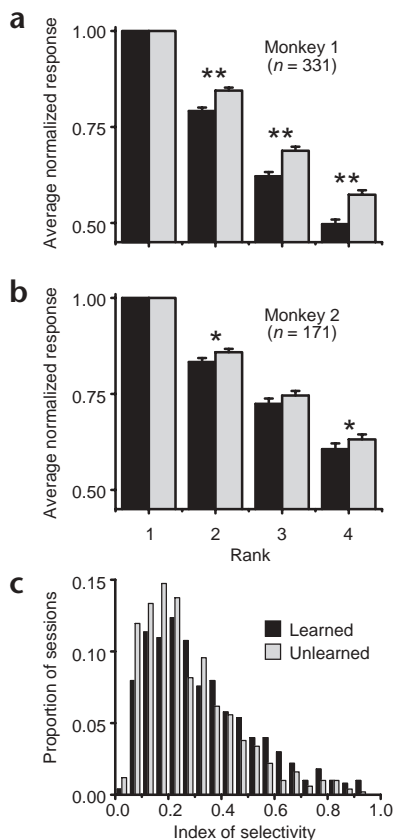


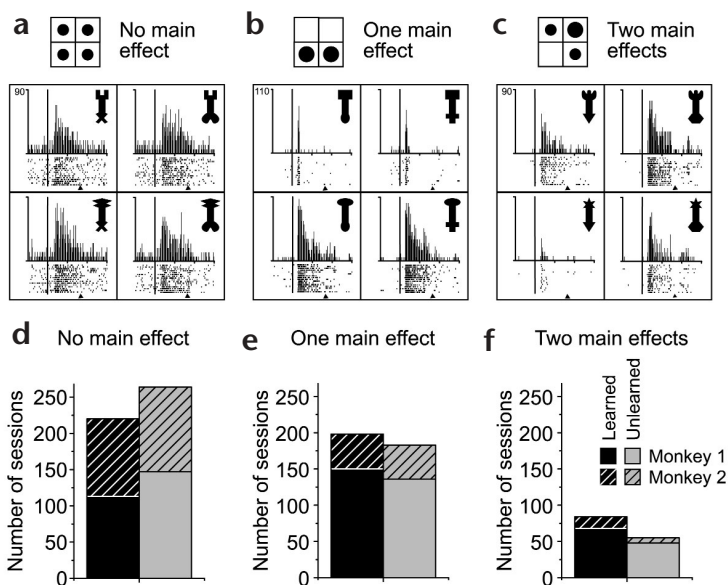
Fig. 3. Selectivity was enhanced for learned as compared to unlearned batons. (a, b) Bars represent mean normalized strengths, across all sessions, of neuronal responses to learned batons (black) and unlearned batons (gray) for monkey 1 (a) and monkey 2 (b). For each session, batons in the learned tetrad were ranked from the most effective (1) to the least effective (4). The same was done for the unlearned tetrad. Firing rates were then normalized to the firing rate elicited by the baton in rank 1 of the corresponding tetrad. The mean firing rate at each rank of each tetrad was then computed across all sessions. Error bars represent standard error of the mean (s.e.m.). Asterisks indicate the level of significance of the difference between learned and unlearned firing rates at each rank in each monkey, as determined by a *post hoc* analysis (Tukey HSD, ** $P < 0.0001$, * $P < 0.02$). (c) Enhancement of selectivity was accomplished by a subtle shift affecting the entire population rather than by the emergence of a few highly selective neurons. Graph shows distribution across all sessions of index values representing selectivity within learned (black) and unlearned (gray) tetrads. Index of selectivity = $(b - w)/(b + w)$ where b and w are firing rates elicited by best and worst batons in the tetrad. The rightward shift for learned batons was significant in both monkeys (Wilcoxon matched pairs test: monkey 1, $P < 0.00001$; monkey 2, $P < 0.03$).

factor ANOVA described above. An interaction effect indicated that the neuron was sensitive to the particular conjunction of parts in a baton. Interaction effects were twice as common for learned (18%) than for unlearned (9%) tetrads (Fig. 5a). This effect was significant in both monkeys (χ^2 test; monkey 1, $P < 0.0004$; monkey 2, $P < 0.026$), did not differ between monkeys (χ^2 test, $P > 0.9$) and became highly significant when data from the monkeys were combined (χ^2 test, $P < 0.0001$). We conclude that training led to enhanced selectivity for whole batons and not just for individual parts.

Interaction effects could take a variety of forms, including selectivity for a single baton (object-type

pattern, Fig. 5b) and selectivity for batons that were associated with the same response although they shared no parts (response-type pattern, Fig. 5c). To characterize the type of interaction in each neuron showing a significant interaction effect, we computed a normalized measure of the difference between the firing rates elicited by the most effective baton and the baton sharing no parts with it (Methods). This pattern-of-interaction index assumed a value of 4.00 in the case of a pure object-type pattern (one baton eliciting a stronger response than the other three) and of 0.00 in the case of a pure response-type pattern (two batons with no parts in common eliciting equally strong responses). Values for learned batons (median, 2.96) were centered closer to 4.00, whereas those for unlearned batons (median, 1.63) were centered closer to 0.00 (Fig. 5d). The offset between the two distributions was significant (Kolmogorov-Smirnov test, $P < 0.01$). Thus, the increase in interaction effects induced by training was disproportionately great for object-type effects. We conclude that learning enhanced the tendency for neurons to respond selectively to just one of the batons in a learned tetrad.

Fig. 4. Selectivity for the individual parts of learned batons was enhanced relative to selectivity for individual parts of unlearned batons. The two-way ANOVA with top part and bottom part as factors could yield any of three outcomes as shown here: selectivity (a) for neither part, (b) for one part or (c) for both parts of the batons in a tetrad. Each set of four histograms represents the responses of one neuron to batons from one tetrad. Traces are aligned on the onset of the 500-ms stimulus (vertical line). The duration represented by the entire horizontal axis is 2,000 ms. The triangle at the base of the raster indicates time of reward delivery. The icon above each set of histograms summarizes the pattern of significant selectivity. (d–f) Counts of sessions in which neurons showed no main effect (d), one main effect (e) or two main effects (f) of part identity for batons belonging to learned (black) or unlearned (gray) tetrads in monkey 1 (uniform texture) or monkey 2 (hatched).



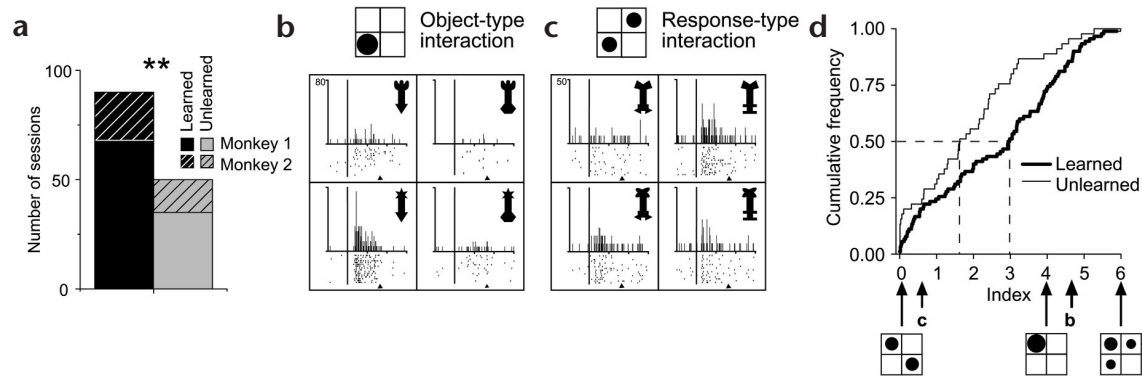


Fig. 5. Learning enhanced the tendency of neurons to respond selectively to just one baton within a tetrad. **(a)** Counts of sessions in which a two-way ANOVA with top part and bottom part as factors yielded evidence of a significant nonlinear interaction between the influences of top and bottom parts. For learned tetrads (black) as compared to unlearned tetrads (gray), significantly more interaction effects occurred (** $P < 0.0001$, χ^2 test). **(b, c)** Neurons showing significant interaction effects occupied a continuum extending from 'object-type' cases (**b**, one object elicited a particularly strong discharge) to 'response-type' cases (**c**, batons sharing no parts, but associated with the same behavioral response, elicited equal responses). **(d)** Cumulative frequency, with respect to a pattern-of-interaction index, of all cases in which firing rate depended significantly on a nonlinear interaction between the identity of the top part and the identity of the bottom part. Thick curve, learned tetrads; thin curve, unlearned tetrads. Index = $(x_1 - x_4)^2/V$, where x_1 and x_4 were the firing rates elicited by the best baton and the baton sharing no parts with it, respectively; V was the variance across the firing rates elicited by the four batons (Methods). Note that the curve for learned tetrads (thick) was shifted, relative to the curve for unlearned tetrads (thin), away from 0.0 (the value associated with a pure response-type pattern) and toward 4.0 (the value associated with a pure object-type pattern). This effect was significant ($P < 0.01$, Kolmogorov-Smirnov test). Index values for neurons in **(b)** and **(c)** are indicated by arrows.

The above estimates of part-based and whole-based selectivity may be low. If a neuron's receptive field (RF) did not encompass both ends of a baton, then the neuron could fall only into the nonselective category or into the category of showing a main effect for one part. To estimate how often this occurred, we took advantage of the principle that if an end of a baton fell outside the RF, then it would do so both for the learned and for the unlearned tetrad. Reasoning from the degree of concordance between results obtained with learned and unlearned tetrads (Methods), we estimated that 78% of the neurons in our sample had RFs that encompassed both ends of the baton. This estimate is commensurate with results from a previous study indicating that 80% of IT neurons have RFs of 5° or larger³². We conclude that the true incidence of selectivity was greater than the measured incidence by up to 20%. It should be noted that this was equally true for both learned and unlearned batons, so it cannot account for the observed learning-related effects.

DISCUSSION

To characterize the impact of discrimination training on neuronal selectivity for parts and wholes of images, we analyzed visual responses to learned and unlearned baton stimuli. The results point to four main conclusions: (i) neurons respond equally strongly on average to their preferred learned and unlearned batons, but (ii) they respond more selectively to learned batons, and this effect is evident (iii) in a modest enhancement of selectivity for the individual parts and (iv) in a marked enhancement of selectivity for combinations of parts. These effects constitute a possible neural mechanism for visual expertise, which is often attributed to greater reliance on configurations or conjunctions of elements as contrasted to individual elements.

The finding that learned and unlearned batons elicit responses of equivalent magnitude stands in contrast to previous reports that discrimination training enhances response strength. Some of these results^{19,20}, however, are from neurons selected on the basis of their responsiveness to learned stimuli. The selection procedure may

have biased the outcome. Others¹⁸ are based on a comparison between trained and untrained monkeys. Thus the findings could have arisen from inter-individual differences rather than from the effects of training. They may also have been specific to recording under anesthesia. However, it remains possible that the training procedures differed from those used here in some way that favored the development of stronger responses to learned stimuli.

The finding that neurons discriminated more effectively between learned than between unlearned batons is in accord with several previous reports based on prolonged pattern-discrimination training^{14,15,18}. This finding can be reconciled with negative results obtained in a few other studies by reference to fundamental differences in behavioral methodology. The training regimen used in one study¹⁶ was very brief, spanning roughly a day; in another study¹⁷, monkeys were trained on a task requiring the discrimination of grating orientation, a skill to which IT may not contribute. With respect to the nature of the changes that mediated the enhancement of selectivity in our study, there are two notable points. First, the enhancement, although modest at the level of individual neurons, was widespread. This observation is consistent with results from an earlier systematic study¹⁸ and is fundamentally at variance with the widely held view^{14,15} that a few neurons become highly selective for learned images. Second, the learning-induced enhancement of selectivity may have arisen from a subtle reduction of each neuron's responses to its nonpreferred images. This would be consistent with previous speculation that response reduction mediates repetition-induced increases in selectivity⁹. Some uncertainty remains on this point because the reduction in responses to nonpreferred stimuli, although significant in both monkeys with response magnitude normalized to the firing rate elicited by the best baton, was significant in only one monkey without normalization.

The finding that selectivity for the parts of learned images was modestly enhanced is without precedent because previous studies have not used visual stimuli allowing selectivity for parts to be analyzed separately from selectivity for wholes. In monkeys trained to

categorize images on the basis of parts at certain internal locations³³, IT neurons are especially selective for parts at those locations. However, this result can be explained, without recourse to a mechanism based on plasticity in IT, by supposing that monkeys simply learned to attend to the relevant locations. That neuronal activity in IT is controlled by attended elements of a display is well known³⁴. Furthermore, the parts at the task-relevant locations (which were the same in all monkeys) may have been inherently more distinctive. The current results can be accounted for neither by supposing that monkeys learned to attend to the ends of the batons (this would have had an identical impact on responses to learned and unlearned batons) nor by supposing that the parts of learned batons were innately more distinctive (we counterbalanced across monkeys which batons were learned and unlearned).

Selectivity for whole images, based on nonlinear interactions between parts, has not been examined in previous studies either of learned or of unlearned stimuli. It is well known that removing even a small part from a neuron's preferred image can cause a drastic reduction in responsiveness^{26–30}. One might suppose, given that the residual image does not elicit a strong response, that the part in isolation does not elicit a strong response, and that the two together do elicit a strong response, that their influences combine nonlinearly in driving neuronal activity. However, in all such cases, the possibility exists that the neuron is selective for a juxtapositional feature in the region where the part and the residual image abut. An example of a juxtapositional feature is the T-junction where two orthogonal contours come together. Only by manipulating discrete and distant features, such as those at opposite ends of a baton, is it possible to circumvent this problem.

Our findings on whole-baton selectivity offer an incidental insight into the representation of category membership. IT neurons selective for specific combinations of parts might have preferred a single baton within a tetrad or, conversely, two batons sharing no parts but associated with a common lever-response (batons diagonally opposed in Fig. 1a). We found that learning enhanced the tendency for neurons to prefer one object out of the four, but not to prefer objects in a common response category. Previous studies have also shown that neurons in IT are not often selective for category membership of disparate stimuli as determined by their arbitrary association with motor responses^{35,36}. For category-based selectivity to appear in IT, it may be necessary that stimuli in a category be visually similar, as are images of trees³⁷ or dogs^{38,39}.

That discrimination training enhanced the selectivity of neurons in IT for whole batons is compatible with the theory that perceptual learning and visual discrimination depend on 'unitization'—the formation of a unitary representation of the collection of features in a learned image^{40,41}. The hypothesis that unitization is dependent on IT neurons that become selectively tuned for combinations of features in learned images, neurons for which 'the whole is greater than the sum of the parts'^{23–25}, is supported by our finding of enhanced selectivity for whole batons. Using stimuli in which distinctive features are segregated in discrete regions was necessary to demonstrate this phenomenon. It is reasonable to assume that this training effect also occurs with natural images, in which features are intermingled. If so, it could account for the ability of experts to discriminate among objects in their domain of expertise on the basis of 'configural' or 'holistic' cues.

METHODS

Tasks. In the discrimination task, the monkey depressed two levers with the right and left hands while maintaining fixation within 1.5° of a spot centered on the monitor. After 500 ms of fixation, the spot was replaced

for 100 ms by a centrally placed baton approximately 5° tall and 2° wide. The monkey had to maintain central fixation until releasing one of the levers within 800 ms in order to receive liquid reward. Training to 80% criterion required ~5,000 trials on each baton for monkey 1 and ~7,000 trials for monkey 2. Each monkey continued to perform the task throughout the neuronal data collection period.

In the fixation task, the monkey maintained central fixation during a 300-ms period before the stimulus, then during a 500-ms period when a single baton was presented centrally, and finally during a 300-ms period after the stimulus. This was followed immediately by reward. Before neuronal data collection began, monkeys 1 and 2 had passively viewed each of the 16 batons approximately 500 and 300 times, respectively.

Recording methods. Recordings were made with varnish-coated tungsten electrodes introduced into the cortex through a guide tube penetrating the dura. Eye movements were monitored with implanted ocular search coils. All procedures were approved by the Carnegie Mellon University Institutional Animal Care and Use Committee and were in accordance with the guidelines set forth in the NIH Guide for the Care and Use of Laboratory Animals. Recording sites, localized by magnetic resonance imaging, occupied the ventral aspect of the temporal lobe lateral to the anterior medial temporal sulcus, and thus were in visual area TE of the inferior temporal lobe, as distinct from perirhinal cortex (Fig. 1b). Monkey 1: 16–19 mm anterior, 17–20 mm lateral. Monkey 2: 16–20 mm anterior, 17–21 mm lateral (Horsley-Clarke coordinates).

Data analysis. In analyzing data from each session, we first determined whether at least one baton out of the eight elicited a significant visual response (*t*-test comparing firing rates during epochs 0–250 ms before fixation spot onset and 50–550 ms after stimulus onset, $P < 0.01$). Subsequent steps of analysis focused on 502 sessions meeting this criterion and on the firing rate 50–550 ms after stimulus onset. The 502 sessions comprised 142 cases in which a neuron contributed data on two pairs of learned and unlearned tetrads and 218 cases in which a neuron contributed data on only one pair. In population analyses, the results of the 502 sessions were treated as independent observations. Follow-up analysis showed that the observed learning effects (i) did not differ significantly in magnitude between cases in which a neuron was tested in one versus two sessions and (ii) persisted in the reduced data set that contained just the first session for each neuron.

To determine how frequently part- and whole-based selectivity occurred, two-way ANOVAs were performed on the responses of neurons to learned and unlearned tetrads. Because the number of observations per condition was small ($n = 16$) and because neural counts tend to follow a Poisson distribution with variances proportional to the mean, we transformed firing rates before the ANOVA. The square root transformation $X' = (X + 0.5)^{0.5}$ was used, where X is raw firing rate and X' is transformed firing rate. This transformation stabilizes variances when samples are taken from a Poisson distribution⁴².

A pattern-of-interaction index was computed as $(x_1 - x_4)^2/V$, where the x variables were the mean firing rates elicited by the most effective baton (x_1), the two batons sharing a part with it (x_2, x_3), and the baton sharing no parts with it (x_4), and V was the variance of x_n . Each value in the range 4.00–6.00 could, in principle, have arisen from either of two firing rate patterns differentiated by whether $(x_1 - x_2 - x_3 + x_4)$ was positive or negative. However, there was no case in which this term had a negative value.

A neuron might be unaffected by the identity of a part at a given location either (a) because it was nonselective for the part or (b) because the part fell outside its classic or its object-centered⁴³ receptive field (RF). To estimate the frequency of such cases, we carried out an analysis based on the principle that, for a given neuron, type-a effects should vary independently across tetrads, whereas type-b effects should be consistent across tetrads. Considering 1,004 cases (502 sessions \times 2 baton-based locations), we counted cases in which neuronal activity was significantly affected by the identity of a part at a given location. Cases were categorized by whether selectivity was present in both the learned and the unlearned tetrad ($n_b = 134$), in the learned tetrad alone ($n_l = 226$), in the unlearned tetrad alone ($n_u = 159$) and in neither ($n_n = 485$). From these counts, we estimated three probabilities: the probabilities that a neuron discriminated

between two learned parts (P_l) or between two unlearned parts (P_u) at a given location when that location fell in its RF, and the probability that a given location fell in the neuron's RF (P_r). These probabilities were computed from the counts n_b , n_l , n_u and n_n , and the identities: $n_b = P_r^*P_l^*P_u^*N$, $n_l = P_r^*P_l^*[1 - P_u]^*N$, $n_u = P_r^*[1 - P_l]^*P_u^*N$ and $n_n = [1 - P_r]^*N + P_r^*[1 - P_l]^*[1 - P_u]^*N$, where $N = n_b + n_l + n_u + n_n$. Resulting estimates of P_l , P_u and P_r were 0.457, 0.372 and 0.783, respectively.

Acknowledgments

This work was supported by National Institute of Health grant RO1 EY11831. We thank K. Medler and K. McCracken for technical assistance.

Competing interests statement

The authors declare that they have no competing financial interests.

RECEIVED 15 AUGUST; ACCEPTED 24 SEPTEMBER 2002

- Wallis, G. & Bülthoff, H. Learning to recognize objects. *Trends Cogn. Sci.* 3, 22–31 (1999).
- Tanaka, J. & Gauthier, I. in *Psychology of Learning and Motivation* Vol. 36 (eds. Goldstone, R. L., Schyns, P. G. & Medin, D. L.) 83–125 (Academic Press, New York, 1997).
- Sheinberg, D. L. & Logothetis, N. K. in *Perceptual Learning* (eds. Fahle, M. & Poggio, T.) 95–124 (MIT Press, Cambridge, MA, 2002).
- Hasegawa, I. & Miyashita, Y. Categorizing the world: expert neurons look into key features. *Nat. Neurosci.* 5, 90–91 (2002).
- Tanaka, K. Inferotemporal cortex and object vision. *Annu. Rev. Neurosci.* 19, 109–139 (1996).
- Logothetis, N. K. & Sheinberg, D. L. Visual object recognition. *Annu. Rev. Neurosci.* 19, 577–621 (1996).
- Baylis, G. C. & Rolls, E. T. Responses of neurons in the inferior temporal cortex in short-term and serial recognition memory tasks. *Exp. Brain Res.* 65, 614–622 (1987).
- Miller, E. K., Li, L. & Desimone, R. A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* 254, 1377–1379 (1991).
- Li, L., Miller, E. K. & Desimone, R. The representation of stimulus familiarity in anterior inferior temporal cortex. *J. Neurophysiol.* 69, 1918–1929 (1993).
- Xiang, J. -Z. & Brown, M. W. Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* 37, 657–676 (1998).
- Sakai, K. & Miyashita, Y. Neural organization for the long-term memory of paired associates. *Nature* 354, 152–155 (1991).
- Erickson, C. A. & Desimone, R. Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J. Neurosci.* 19, 10404–10416 (1999).
- Messinger, A., Squire, L. R., Zola, S. M. & Albright, T. D. Neuronal representations of stimulus associations develop in the temporal lobe during learning. *Proc. Natl. Acad. Sci. USA* 98, 12239–12244 (2001).
- Logothetis, N. K. & Pauls, J. Psychophysical and physiological evidence for viewer-centered object representations in the primate. *Cereb. Cortex* 3, 270–288 (1995).
- Logothetis, N. K., Pauls, J. & Poggio, T. Shape representation in the inferior temporal cortex of monkeys. *Curr. Biol.* 5, 552–563 (1995).
- Erickson, C. A., Jagadeesh, B. & Desimone, R. Clustering of perirhinal neurons with similar properties following visual experience in adult monkeys. *Nat. Neurosci.* 3, 1143–1148 (2000).
- Vogels, R. & Orban, G. A. Does practice in orientation discrimination lead to changes in the response properties of macaque inferior temporal neurons? *Eur. J. Neurosci.* 6, 1680–1690 (1994).
- Kobatake, E., Wang, G. & Tanaka, K. Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *J. Neurophysiol.* 80, 324–330 (1998).
- Sakai, K. & Miyashita, Y. Neuronal tuning to learned complex forms in vision. *Neuroreport* 5, 829–832 (1994).
- Miyashita, Y., Date, A. & Okuno, H. Configurational encoding of complex visual forms by single neurons of monkey temporal cortex. *Neuropsychologia* 31, 1119–1131 (1993).
- Gauthier, I. & Tarr, M. J. Unraveling mechanisms for expert object recognition: bridging brain activity and behavior. *J. Exp. Psychol. Hum. Percept. Perform.* 28, 431–446 (2002).
- Tanaka, J. W. & Farah, M. J. in *Analytic and Holistic Processes in Perception of Faces, Objects and Scenes* (eds. Peterson, M. A. & Rhodes, G.) (Oxford Univ. Press, New York, in press).
- Murray, E. A. & Bussey, T. J. Perceptual-mnemonic functions of the perirhinal cortex. *Trends Cogn. Sci.* 3, 142–151 (1999).
- Bussey, T. J. & Saksida, L. M. The organization of visual object representations: a connectionist model of effects of lesions in perirhinal cortex. *Eur. J. Neurosci.* 15, 355–364 (2002).
- Bussey, T. J., Saksida, L. M. & Murray, E. A. Perirhinal cortex resolves feature ambiguity in complex visual discriminations. *Eur. J. Neurosci.* 15, 365–374 (2002).
- Perrett, D. I., Rolls, E. T. & Caan, W. Visual neurons responsive to faces in the monkey temporal cortex. *Exp. Brain Res.* 47, 329–342 (1982).
- Desimone, R., Albright, T. D., Gross, C. G. & Bruce, C. Stimulus-selective properties of inferior temporal neurons in the macaque. *J. Neurosci.* 4, 2051–2062 (1984).
- Tanaka, K., Saito, H., Fukada, Y. & Moriya, M. Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *J. Neurophysiol.* 66, 170–189 (1991).
- Yamane, S., Kaji, S. & Kawano, K. What facial features activate face neurons in the inferotemporal cortex of the monkey. *Exp. Brain Res.* 73, 209–214 (1988).
- Tsunoda, T., Yamane, Y., Nishizaki, M. & Tanifuji, M. Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nat. Neurosci.* 4, 832–838 (2001).
- Desimone, R. Neural mechanisms for visual memory and their role in attention. *Proc. Natl. Acad. Sci. USA* 93, 13494–13499 (1996).
- Op De Beeck, H. & Vogels, R. Spatial sensitivity of macaque inferior temporal neurons. *J. Comp. Neurol.* 426, 505–518 (2000).
- Sigala, N. & Logothetis, N. K. Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature* 415, 318–320 (2002).
- Desimone, R. Visual attention mediated by biased competition in extrastriate visual cortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 353, 1245–1255 (1998).
- Op de Beeck, H., Wagemans, J. & Vogels, R. Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nat. Neurosci.* 4, 1244–1252 (2001).
- Xiang, J. -Z. & Brown, M. W. Differential neuronal responsiveness in primate perirhinal cortex and hippocampal formation during performance of a conditional visual discrimination task. *Eur. J. Neurosci.* 11, 3715–3724 (1999).
- Vogels, R. Categorization of complex visual images by rhesus monkeys. Part 2: single-cell study. *Eur. J. Neurosci.* 11, 1239–1255 (1999).
- Freedman, D. J., Riesenhuber, M., Poggio, T. & Miller, E. K. Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291, 312–316 (2001).
- Freedman, D. J., Riesenhuber, M., Poggio, T. & Miller, E. K. Visual categorization and the primate prefrontal cortex: neurophysiology and behavior. *J. Neurophysiol.* 88, 929–941 (2002).
- Goldstone, R. L. Unitization during category learning. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 86–112 (2000).
- Goldstone, R. L. Perceptual learning. *Annu. Rev. Psychol.* 49, 585–612 (1998).
- Zar, J. H. *Biostatistical Analysis* 4th edn. (Prentice Hall, Upper Saddle River, New Jersey, 1999).
- Pasupathy, A. & Connor, C. E. Shape representation in area V4: position-specific tuning for boundary conformation. *J. Neurophysiol.* 86, 2505–2519 (2001).