Image statistics of American Sign Language: comparison with faces and natural scenes

Rain G. Bosworth

Department of Psychology, University of California, San Diego, La Jolla, California 92093

Marian Stewart Bartlett

Institute for Neural Computation, University of California, San Diego, La Jolla, California 92093

Karen R. Dobkins

Department of Psychology, University of California, San Diego, La Jolla, California 92093

Received January 4, 2006; accepted February 17, 2006; posted March 24, 2006 (Doc. ID 66719)

Several lines of evidence suggest that the image statistics of the environment shape visual abilities. To date, the image statistics of natural scenes and faces have been well characterized using Fourier analysis. We employed Fourier analysis to characterize images of signs in American Sign Language (ASL). These images are highly relevant to signers who rely on ASL for communication, and thus the image statistics of ASL might influence signers' visual abilities. Fourier analysis was conducted on 105 static images of signs, and these images were compared with analyses of 100 natural scene images and 100 face images. We obtained two metrics from our Fourier analysis: mean amplitude and entropy of the amplitude across the image set (which is a measure from information theory) as a function of spatial frequency and orientation. The results of our analyses revealed interesting differences in image statistics across the three different image sets, setting up the possibility that ASL experience may alter visual perception in predictable ways. In addition, for all image sets, the mean amplitude results were markedly different from the entropy results, which raises the interesting question of which aspect of an image set (mean amplitude or entropy of the amplitude) is better able to account for known visual abilities.

OCIS codes: 070.2590, 100.2960.

1. INTRODUCTION

Several lines of evidence suggest that visual experience with the environment plays a role in shaping visual abilities. One of the best known examples is in the domain of orientation processing. Early studies have shown that raising animals in environments with only horizontal or vertical contours results in enhanced sensitivities to those orientations.[1–3] (Similar results have been reported in children who grew up with uncorrected astigmatism.[4,5]) Humans raised under normal (and unrestricted) conditions also show anisotropies for orientations, which are thought to arise from a natural orientation bias in the environment. Specifically, human subjects exhibit an oblique effect in which acuity and contrast sensitivity are better for cardinal (vertical and horizontal), than for oblique, orientations[6–8] (and see Refs. 9–11 for evidence of an oblique effect in infants). It had long been speculated that the oblique effect is driven by a preponderance of cardinal orientations in the environment. In support of this notion, the results from Fourier analyses of natural scenes have shown that cardinal orientations contain more contrast (Fourier amplitude) than do oblique orientations.[12–16] As might be expected, the cardinal bias in the environment measured with Fourier analysis is stronger for scenes that contain man-made structures, referred to as carpentered environments, than for those that do not.[14,17,18] This difference between carpentered versus noncarpentered environments can be used to explain why people who live in less carpentered environments, such as the Cree Indians, exhibit a smaller oblique effect than do people who live in highly carpentered environments.[19] In sum, the results from a variety of studies investigating orientation processing suggest that visual sensitivity to different orientations is influenced by the prevalence of different orientations in the environment.

In addition to orientation, spatial frequency content of natural scenes has been statistically described using Fourier analysis (i.e., amplitude of contrast across spatial frequencies). A robust regularity that has been consistently observed across many studies is that amplitude declines with spatial frequency, where the amplitude at a given spatial frequency f is, approximately $1/f$.[20–24] (Note that many studies report power, which is the square of the amplitude, and the falloff is $1/f^2$.) Recent studies have shown that visual psychophysical performance in humans is well matched to the $1/f$ function typically encountered in natural scenes. For example, object discrimination sensitivity is best for stimuli that have natural power spectra slopes compared with those that have slopes that deviate from $1/f$.[25–28] Also, response properties of neurons in the retina, lateral geniculate nucleus, and visual cortex have

been suggested to be optimally organized to process stimuli with natural image statistics[12,16,21,22,29–33] (and see Simoncelli and Olshausen[34] for a review).

Statistics of other image properties such as contours and color content in natural scenes have also been shown to have a close link to perception. For example, Geisler *et al.*[35] showed that human performance for contour detection matched very closely to predicted ideal observer performance based on the co-occurrence of contours in natural images. Regan *et al.*[36] found that the optimal placement of M and L cones in the retina corresponds well to the observed wavelength distribution in the environment of New World monkeys (see also Refs. 37–39 for similar findings). Similar links between the environmental statistics of color and the cone spectral sensitivities in humans have also been reported, as well as relationships between environmental statistics and measures of color appearance, identification, and labeling.[40–44]

Another way researchers have investigated the role of the environment in shaping visual perception has been to ask which spatial frequencies or orientations in a set of images are most critical for discriminating one image in that set from another. Typically, such studies focus on a confined set of images for which people have ample experience. A good example of such an image set is printed letters. Several psychophysical studies have measured discrimination of letters filtered or masked along a band of spatial frequencies. The general consensus from these studies is that the most important band for letter identification peaks at an object spatial frequency of 3 cycles per letter (cyc/letter).[45–48] In recent Fourier analyses of letters, it has been shown that the highest amplitudes are also found near 3 cyc/letter.[49] Thus, these findings indicate that the task of letter discrimination relies most heavily on spatial frequencies that are most prevalent in letters, and it is reasonable to assume that this link arises as a result of experience with reading. Presumably, this tuning would be specific for the task of letter discrimination; for other image sets (such as faces), the spatial frequencies that contain the highest amplitudes are likely to differ from that of letters, and, accordingly, these spatial frequencies might be most important for discrimination of that set.

In the current study, we focused on images of signers producing American Sign Language (ASL), and we compared these images with those from established databases of natural scenes and faces. Our interest in ASL is based on the fact that many deaf people have lifelong experience with ASL, relying on it for communication (and there also exist many hearing people who converse fluently or nonfluently in ASL). Several studies have shown that native signers exposed to ASL since birth, by virtue of being raised by signing parents, do show altered or enhanced perceptual abilities.[50–56] Like the case for images of written letters (described above), it is possible that, for ASL images, the spatial frequencies (or orientations) that contain the highest amplitudes are also those most important in discriminating one sign from another. To address this possibility, the image statistics of ASL must first be described. To this end, we conducted Fourier analyses on a set of images containing a representative sample of the different hand–arm positions and spatial configurations seen in ASL. The Fourier analyses of static sign images in the current study focus only on the position and spatial configuration of hands and arms. Motion of the hands is the other integral piece of information that has been studied by linguists,[57–59] and we address motion statistics of ASL in a separate forthcoming paper.

From our Fourier analysis of static images (signs, faces, and natural scenes), we obtained two metrics. First, we quantified the mean amplitude of the Fourier component as a function of spatial frequency and orientation, which is the typically reported output from Fourier analysis. Second, we quantified the entropy of the amplitude as a function of spatial frequency and orientation, a metric that has yet to be described for natural image sets. Entropy describes the spread of a distribution and is related to, although not identical to, the variance of this distribution (see Cover and Thomas[60] or Atick[61] for a review). Here we looked at the distribution of amplitudes across the image set for a given spatial frequency or orientation. The entropy metric comes from information theory, with the notion that a spatial frequency (or orientation) whose amplitude distribution is broadly spread across the image set contains more information for discriminating one image from another than a spatial frequency (or orientation) that almost always has the same amplitude across images. It is important to point out that entropy refers to information in the image set and is not to be confused with psychophysical measurements of information that test (through filtering and masking) which spatial frequencies (or orientations) are most important to the visual system for discriminating one image in a set from another. Whether the two are related (information in the image set and critical information used by the visual system) is one of the questions of the current study.

The results of our Fourier analyses revealed two main findings. First, substantial differences were observed between signs and the two other image sets, faces and natural scenes. In particular, sign images contained more amplitude for vertical, than for horizontal, contours, whereas faces and natural scenes showed an opposite pattern. Given the above-mentioned relationship between the prevalence of different orientations in the environment and visual sensitivity, these results suggest that different anisotropies in orientation may exist for signers versus nonsigners. In addition, for sign images only, entropy versus spatial frequency curves were clearly bandpass, with a peak around 0.75 cyc/cm [which is roughly 0.19 cyc/deg at a viewing distance of 5 ft (1 ft = 0.3048 m)]. Interestingly, this peak in entropy maps on roughly to the spatial frequency that has been shown to be important for identifying signs in ASL (based on psychophysical studies by Riedl and Sperling,[62] discussed further, below), which suggests a potential link between entropy and visual discrimination.

Second, for all image sets (signs, natural scenes, and faces), the shape of the curve relating mean amplitude versus spatial frequency (and orientation) was markedly different from that relating entropy versus spatial frequency (and orientation). These differences raise an interesting question: which aspect of an image set (mean amplitude or entropy of the amplitude) is better able to account for known perceptual abilities? In particular, in

our orientation analyses of natural scenes, mean amplitudes were largest for cardinal orientations, in line with results from previous Fourier analyses and in line with the psychophysical oblique effect (see above). By contrast, entropy was found to be lowest for cardinal orientations, particularly for horizontal contours. Interestingly, these results are in line with recent psychophysical data showing that the oblique effect (i.e., poorest contrast sensitivity for oblique contours) that is known to exist for stimuli of a single spatial frequency turns into a horizontal effect (i.e., poorest contrast sensitivity for horizontal contours) for stimuli of broadband spatial frequencies that are more naturalistic.[17,63] Together, these current Fourier analyses and previous psychophysical results suggest a potential link between entropy and visual sensitivity.

2. METHODS

A. Photographs of American Sign Language

To obtain a database of ASL images, we took photographs of two female signers (RB and DH) producing signs from ASL. The photographs were taken with a Canon Rebel, 35 mm single-lens reflex camera, with a 200 speed setting and with a flash, at a distance of 63.5 in. The signers were asked to produce 40 different signs, which were selected with the goal of photographing a range of hand shapes and arm positions commonly seen in ASL (~15 different hand shapes). Note that it was not our goal to have each image associable with a word but rather to capture positional information that makes up the structure of signs. For this reason, we refer to our ASL database as sign images. In addition to providing a representative sample of the different hand–arm positions in ASL, the 40 signs were also selected such that the proportion of signs that were one handed versus two handed (30% versus 70%, respectively) and the proportion of two-handed signs that were symmetric versus asymmetric (50%) roughly matched that which has been observed in ASL.[64]

To make our sign images as naturalistic as possible, signers were asked to sign a given word but to momentarily freeze at a certain point so that the photograph of the prototypical hand–arm position for that sign could be taken. For example, in the sign CANADA [see Fig. 1(a)] a photograph was taken with the signer's fist on her chest. For signs that involved uniform, continuous circular motion (e.g., BICYCLE, ENJOY, GESTURE, WASH), we asked signers to freeze halfway through the motion cycle. These sign images are segments of whole words and can be thought of as making up constituents within ASL, much the way spoken language has a finite set of vowels and consonants that are combined sequentially to create spoken syllables and words. In total, our ASL database consisted of 59 and 46 photographs of sign images taken of signers DH and RB, respectively. In addition to obtaining photographs of sign images, for comparison purposes, we obtained a photograph of each signer with her arms resting, which we refer to as neutral pose [see Fig. 1(b)].

B. Sign Image Preparation

Each photograph was scanned using an Epson scanner set at 600 dots per inch, gamma 1.8, and with a window size 1800×2400 pixels. Mean interocular distance across all images was 138 pixels. For each image, we removed the face, the torso, and the background, leaving only the arms, from shoulders to fingers, as shown in Figs. 1(c) (sign image from the sign CANADA) and 1(d) (neutral pose). The purpose in doing so was to restrict our analyses to the signers' arms and hands because the arms and hands provide the lexical and grammatical information in ASL. Having the same face in each ASL image would add redundancy to the analysis. This is not to say that the face (specifically, facial expression) does not contribute to ASL comprehension; however, the lexical identity of any sign can be comprehended without the face. (Note that we conducted a separate analysis on a set of faces obtained from a published database, where each face in the database is a different identity; see below.) Finally, eliminating all but the arms in these photographs also served to eliminate noise from irrelevant portions of the image, for example, wrinkles in the clothing. (Partial results for full images are reported in Bosworth *et al.*[65] and are very similar.) Examples of several sign images used in our analyses are provided in Fig. 2.

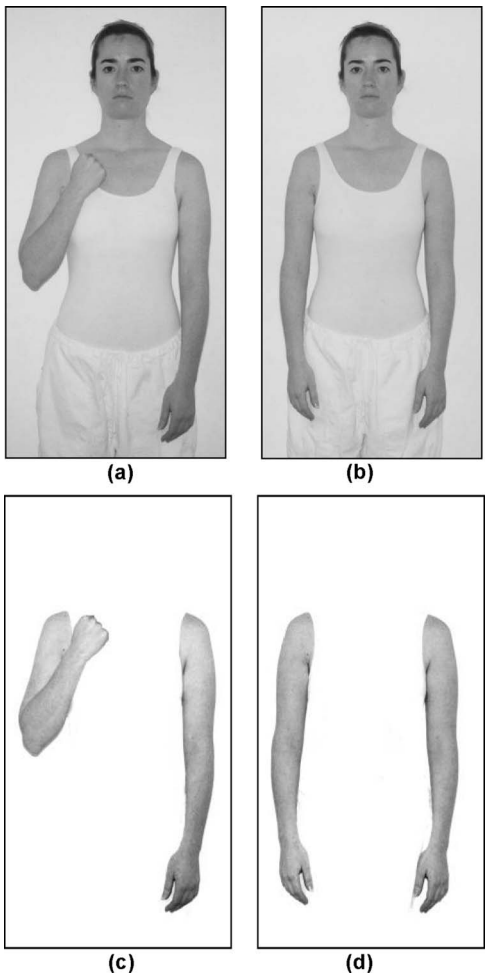

Fig. 1. Photography of ASL signs and image preparation. (a) An example of an original photograph image of DH signing CANADA and (b) original photograph image of DH in a neutral pose (arms at side). (c) and (d) The face, torso, and background in these two images have been removed, leaving only the arms, from shoulders to fingers. Fourier and entropy analyses were conducted on the arms-only images [(c) and (d)].

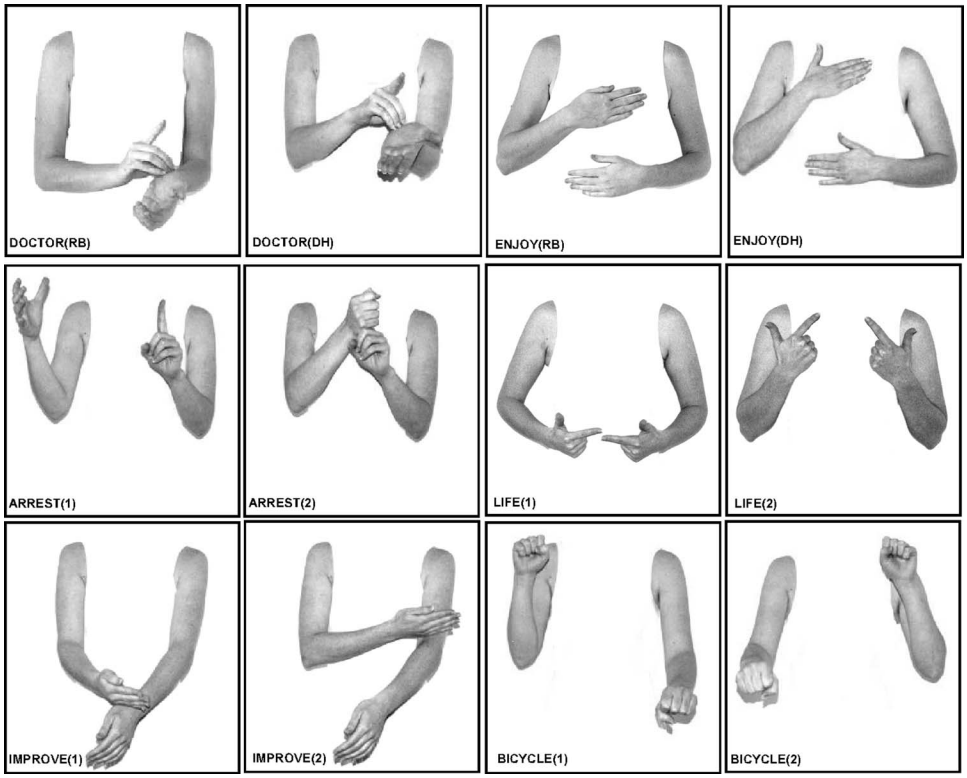

Fig. 2. Example sign images. Several examples are shown of sign images included in our ASL sample. The top row has images for the signs DOCTOR and ENJOY made by RB and DH. For some signs that involve critical changes in position (e.g., ENJOY, LIFE) or hand shape (e.g., ARREST), two consecutive images were obtained, and these are shown in the middle and last rows (only for DH).

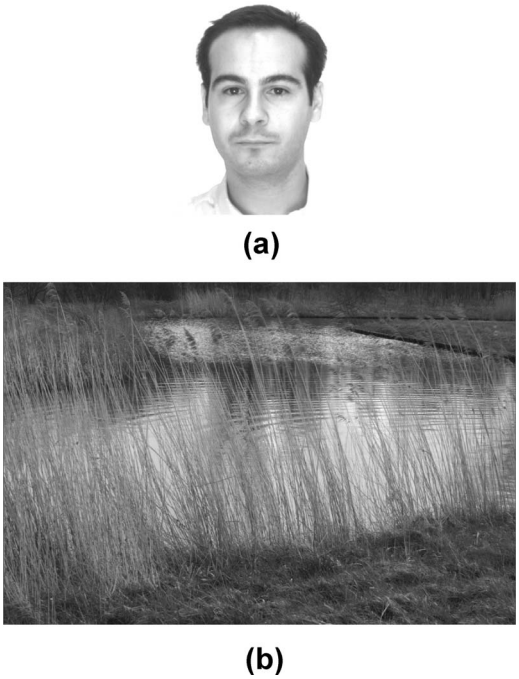

Fig. 3. (a) Example face image from the AR face database.[66] (b) Example of a natural scene image from the van Hateren and van der Schaaf database.[67]

C. Images of Faces and Natural Scenes

For comparison with sign images, we performed the same analyses on images of faces and natural scenes. The AR face database was collected by Martinez and Benavente[66] using a Sony 3 CDD camera with a Matrox Meteor red-green-blue frame grabber. Face images consisted of neutral expression images of 100 individuals of various races and genders from the AR face database [see example in Fig. 3(a)].[66] These images were 768×576 pixels with a mean interocular distance of 106 pixels. Natural scenes were taken from the Groningen natural image database[67] and were 1536×1024 pixels [see example in Fig. 3(b)]. The natural images were obtained from a Kodak DCS420 digital camera and were linearized in intensity, as described by van Hateren and van der Schaaf.[67] How the different image sets (signs, faces, and natural scenes) were aligned in scale is discussed further below. This database contains 4000 images of outdoor scenes, including buildings as well as natural landscapes and close-up foliage texture. We selected 100 images from the first 400 images of the database that contained no visible man-made structures. For all image sets (signs, faces, and natural scenes), images were normalized in luminance by setting the brightest and darkest gray values of each image to 255 and 0, respectively, and linearly rescaling the gray values in between.

D. Scale Alignment

The Fourier analysis (described below) outputs spatial frequency in terms of cycles per image. However, the three image databases contained images of different size. The images were not rescaled to contain the same number of pixels. Instead, cycles per degree was estimated for each image database, as well as cycles per centimeter in the world where applicable. (We chose this because down-

scaling causes high-frequency information to be lost while upscaling does not add high-frequency information and because cycles per image are not as relevant as cycles per degree or cyles per centimeter in the world.) For the ASL and face images, cycles per image was converted first into cycles per centimeter in the real world using the ratio of face width in pixels to actual face width in centimeters. Cycles per degree was then estimated by assuming a viewing distance of 5 ft (i.e., 2.54 cm = 1 deg). We chose this distance because it is typical for conversing signers. We also assumed that this distance would be typical for viewing faces during conversation between nonsigners. For the natural scenes database, cycles per centimeter in the world could not be calculated, since the images contain a range of scales in the foreground and background. Instead, cycles per image was converted directly into cycles per degree using the published angular resolution of the pictures, which is approximately 1 arc min per pixel, making a viewing angle of 17 by 25.5 deg.[67]

E. Fourier Analysis

Fourier analysis was conducted on each image by using a discrete fast Fourier transform function in MATLAB (by Math Works), using wraparound with no padding. Although wraparound can cause artificial edge responses, we chose this because the edge responses are generally much smaller than those caused by zero padding. Analyses were performed on the amplitude output (defined as the square root of the power, or energy). For each image, the Fourier transform can be displayed as a three-dimensional (3D) polar plot (r, θ, a), where the radial distance from the origin (r) is spatial frequency, the angle (θ) is orientation, and the brightness at each point in the plot is the amplitude (a) at a given spatial frequency and orientation. Examples of these Fourier transform plots are provided in Fig. 4 for two images: an image from the sign HEART [Fig. 4(a)] and an image of a natural scene [Fig. 4(b)]. For all analyses, the corners of the Fourier transform were excluded, since they contain high frequencies at oblique orientations that are beyond the resolution limit at the cardinal orientations. Comers were defined as radial distances that exceed the maximum radial distance for any one direction $[r > \min_\theta(\max(r_\theta))$ where r_θ is the radial distance in direction θ and $\min_\theta$ denotes the minimum over all orientations $\theta]$. In other words, the analysis was performed only for those radial distances r for which data were available at all orientations. Also, note that horizontal contours in an image have vertical Fourier energy, since orientation refers to the orientation of individual sine waves. In the results presented from here forward, orientation refers to the orientation in the image, which is 90 deg shifted from the Fourier orientation.

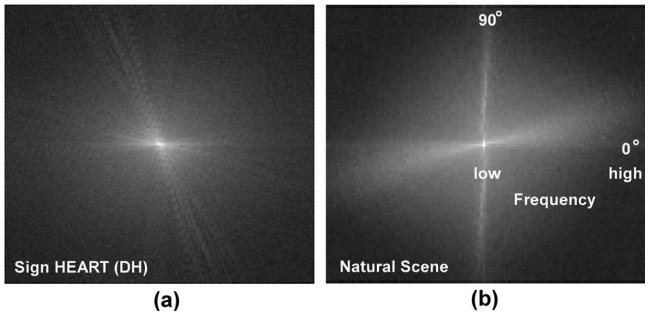

Fig. 4. Example Fourier plots for (a) sign image obtained from the sign HEART and (b) the natural scene shown in Fig. 3. For each image, the Fourier transform is displayed as a 3D polar plot, where log amplitude, a, is plotted as a function of both spatial frequency (distance from the origin, r) and orientation (angle, θ). Note that horizontal objects in an image have vertical Fourier energy, since orientation refers to the orientation of individual sine waves. The natural scenes plot shows high energy in the vertical orientation, which has been related to horizontal layering in distant natural scenes.[18] (In the results and figures, we refer to orientation as orientation in the image, which is 90 deg shifted from the Fourier orientation.)

For all image sets, Fourier data were analyzed at spatial frequencies ranging from 0.06 to 23.0 cyc/deg (0.02 to 9.0 cyc/cm). Fourier data for orientation were analyzed from 0 to 180 deg (with 0 and 90 deg denoting horizontal and vertical contours, respectively). From the Fourier analyses, we computed two metrics as a function of spatial frequency and orientation, mean amplitude and entropy of the amplitude, described in detail, below.

F. Mean Amplitude Analyses

For each image set, mean amplitudes were calculated across spatial frequency and across orientation. Two-dimensional (2D) plots of amplitude versus spatial frequency were obtained by averaging across orientation within each image and then averaging across images in a given image set. The data were then plotted in log–log coordinates, which is the conventional method of analyzing amplitude by spatial frequency data, thus allowing us to compare our data with those of previous studies. For the 2D plots of amplitude versus orientation, we first logged (base 10) the values before collapsing across spatial frequency because spatial frequency data are known to be highly nonlinear, whereas they are linear on a $\log_{10}$ scale. The 2D data were then averaged across the set of images.

G. Entropy of the Amplitude

A spatial frequency (or orientation) with high mean amplitude may nevertheless provide little information for differentiating images if it tends to frequently take the same high value, compared with a spatial frequency (or orientation) with high variability in amplitude across images. Therefore, in addition to measuring mean amplitude, we also measured the entropy of the amplitude distribution, In brief, entropy measures the spread of the probability distribution and is highest for broad distributions and lowest for narrow distributions. For Gaussian distributions, entropy is equivalent to variance. However, it is well known that many natural signals have distributions that are highly non-Gaussian, including measures of contrast in natural scenes.[68–72]

Entropy is a metric from information theory, which describes the amount of information in a signal for differentiating exemplars. (See Cover and Thomas[60] for an introductory text on information theory.) Information theory defines the amount of information (I) in an outcome x as follows:

$$I(x) = -\log_2(p(x)), \tag{1}$$

where $p(x)$ is the probability that the measure takes the value x across the set of images. [Note that the log in Eq.

(1) is base 2, since the definition of information comes from bits.[60]] In the current study, x is an amplitude value (at a given spatial frequency or orientation). Entropy (H) is the expected value of the information across the full distribution of values x, as follows:

$$H = E(-\log_2(p(x))), \tag{2}$$

where E is the expected value. For each image database, entropy was calculated across images for each spatial frequency and orientation. Entropy was calculated using Eq. (2), where $x = \log(a_{f\theta})$, which is the set of log amplitudes at frequency f and orientation θ across all images in the database. Probability distributions were estimated by histogramming x into 20 equally spaced bins. Identical bins were employed at all spatial frequencies and orientations. The bin range for each image database was defined by $\min(x)$ and $\max(x)$ over all spatial frequencies, orientations, and images. The entropy analysis produced a single 3D entropy plot (frequency × orientation × entropy) for each image database. Like the amplitude data (described above), the 3D surface was then reduced to two dimensions by collapsing across spatial frequency and orientation separately, producing plots of entropy versus spatial frequency and entropy versus orientation. For our analysis, we compare the entropy peaks and shape of entropy functions for each image database rather than compare the absolute entropy values across databases.

3. RESULTS

A. Spatial Frequency: Mean Amplitude Analysis

Mean amplitudes are plotted as a function of spatial frequency in Fig. 5 on a log–log scale. Data obtained for sign images (thick black curves) are presented separately for the two signers, RB (top) and DH (bottom). For comparison, data are also presented for a neutral pose (with arms at rest) for each signer (gray curves). Also, for comparison, data are shown for the set of faces taken from the AR database (dotted curves) and natural scenes taken from the Groningen database (thin curves). Note that the faces and natural scenes data are plotted redundantly in the top and bottom figures, to allow comparison with each signer's data.

For all image sets, amplitudes were found to be largest at low spatial frequencies, declining approximately according to $1/f$, where f is spatial frequency, and were approximately linear on a log–log scale. This $1/f$ pattern has been widely reported for natural scenes.[20–24] More precisely, the amplitude spectrum of natural scenes has been found to take the form $1/f^{\text{exp}}$, with the exponent being equivalent to the slope of log amplitude against log frequency. In log–log coordinates, the slopes are found to be near −1, ranging from −0.7 to −1.6 across studies.[21,22,24] (Note that some studies report results for the power spectrum,[23] which is the square of the amplitude spectrum, and hence the exponents and slopes from these studies will be double those for the amplitude spectrum.) In the current study, where we measure the amplitude spectrum in log–log coordinates, we find a slope value for natural scenes of −1.17, which is within the previously reported range. For the face database, the slope was −2.06. For sign images, the slope value was approximately

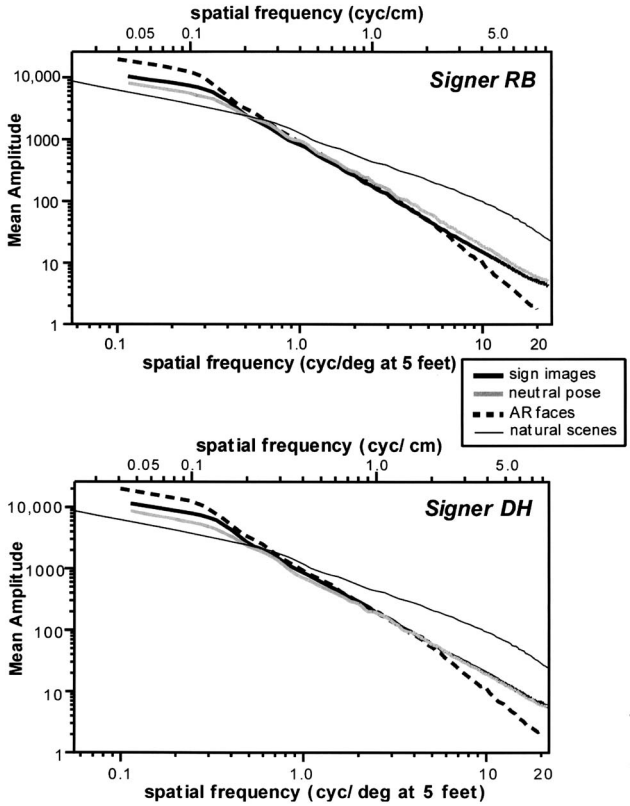

Fig. 5. Amplitude versus spatial frequency. Mean log amplitude is plotted as a function of log spatial frequency (data collapsed across orientation) for signers (top) RB ($n = 46$ images) and (bottom) DH ($n = 59$ images). Data are shown for sign images (black curves) and the neutral pose (gray curves). For comparison, data are shown for a set of 100 faces taken from the AR database (dotted curves) and a set of 100 natural scene images taken from the van Hateren and van der Schaaf database (thin curves). Note that the faces and natural scenes data are plotted redundantly in both figures, to allow comparison with each signer's data.

−1.66, which was nearly identical for the two signers (RB = −1.68, DH = −1.64). The slope value for the neutral pose was similar to that of the sign images (RB = −1.64, DH = −1.61), with a trend to be slightly shallower. In sum, all image sets conformed to a $1/f^{\text{exp}}$ pattern, with the steepest slopes found for faces, followed by sign images and then natural scenes. However, the fact that signing arms and the neutral pose yielded quite similar curves suggests that signing arms are unlikely to contain spatial frequency information much different from that of nonsigning arms, an issue we return to in Section 4.

B. Spatial Frequency: Entropy Analysis

Entropy of the amplitude distribution is plotted as a function of spatial frequency in Fig. 6, separately for sign images of signer RB (thick black curve), and signer DH (gray curve), faces (dotted curve), and natural scenes (thin curve). Note that entropy cannot be calculated for the neutral pose condition, since entropy was calculated across images and the neutral pose was a single image. For sign images from both signers, the entropy curves appeared bandpass, with a peak in entropy at 2.0 and 1.8 cyc/deg for RB and DH, respectively (0.8 and 0.7 cyc/cm, respectively). These peaks in entropy occur at

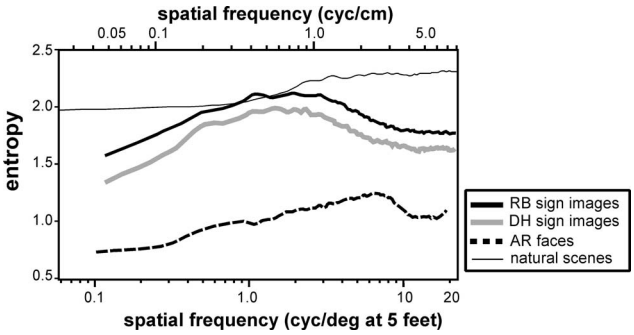

Fig. 6. Entropy versus spatial frequency. Entropy is plotted as a function of log spatial frequency (data collapsed across orientation) for signers RB (thick black curve) and DH (gray curve), faces (dotted curve), and natural scenes (thin curve). Entropy could not be calculated for the neutral pose, which is a single image, since in this study entropy was calculated for a distribution of many images.

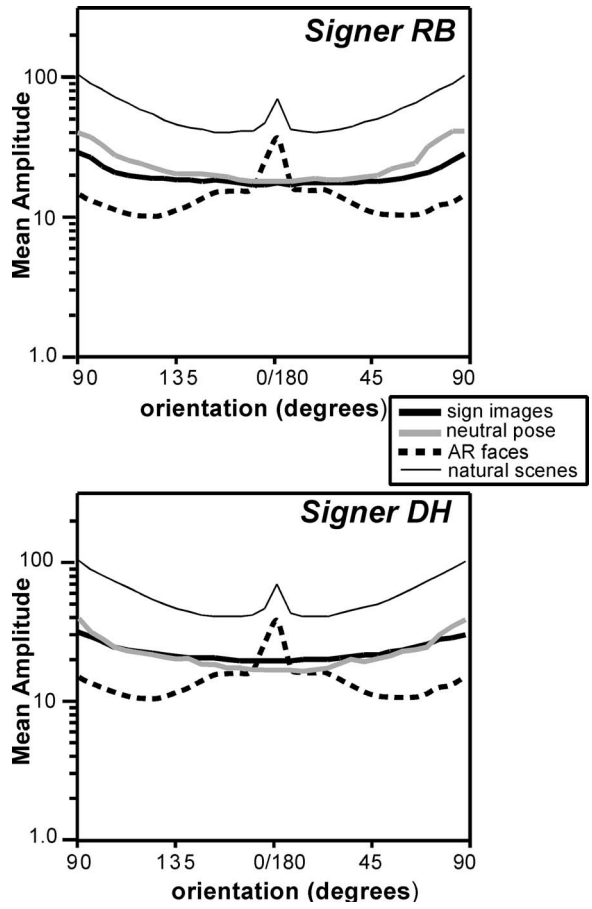

Fig. 7. Amplitude versus orientation. Mean log amplitude is plotted as a function of orientation (data collapsed across spatial frequency) for signers (top) RB ($n=46$ images) and (bottom) DH ($n=59$ images). Data are shown for sign images (black curves) and the neutral pose (gray curves). For comparison, data are shown for the set of 100 faces taken from the AR database (dotted curves) and the set of 100 natural scene images taken from the van Hateren and der Schaaf database (thin curves). Note that the faces and natural scenes data are plotted redundantly in both figures, to allow comparison with each signer's data. Also note that the X axis is centered at 0 deg (horizontal) to emphasize the sharpest peak in the data.

roughly the same spatial frequencies that have been shown to be important for identifying signs in ASL based on psychophysical studies,[62] an issue we return to in Section 4. There was also a trend for a bandpass shape in the face images, with a peak at 7.0 cyc/deg (2.8 cyc/cm). For natural scenes, entropy was roughly constant up to 1 cyc/deg, increased markedly between 1 and 3 cyc/deg, and then remained roughly constant as spatial frequency was increased further. Note that the shapes of the entropy versus spatial frequency curves for all image sets were very different from the shapes of the mean amplitude versus spatial frequency curves of Fig. 5: For all image sets, mean amplitude decreased with increasing spatial frequency in a $1/f$ fashion, whereas entropy curves for spatial frequency were largely bandpass or flat in nature.

C. Orientation: Mean Amplitude Analysis

Mean amplitudes are plotted as a function of orientation in Fig. 7 (0 deg denotes horizontal and 90 deg denotes vertical contours in the image). Data obtained for sign images (black curves) and the neutral pose (gray curves) are presented separately for the two signers, RB (top) and DH (bottom). For both sign images and the neutral pose, there was a clear peak in amplitude at 90 deg. This vertical bias (which was stronger for the neutral pose), is due, of course, to the fact that arms (specifically arms at rest, as in the neutral pose) are oriented vertically. The fact that signing arms and the neutral pose yielded quite similar curves suggests that signing arms are unlikely to contain orientation information much different from that of nonsigning arms, an issue we return to in Section 4.

For comparison, data are shown for the set of faces (dotted curves) and natural scenes (thin curves). Note that the faces and natural scenes data are plotted redundantly in the top and bottom figures, to allow comparison with each signer's data. For natural scenes, peaks in amplitude were observed at the cardinal orientations (0 and 90 deg), which is in agreement with previously reported results for amplitude analyses of natural scenes.[12,15,16,18,20] Similarly, for faces, peaks in amplitude were observed at the cardinal orientations; however, the peak was much stronger at 0 deg (horizontal).

D. Orientation: Entropy Analysis

Entropy of the amplitude distribution is plotted as a function of orientation in Fig. 8, separately for sign images of signer RB (thick black curve) and signer DH (gray curve), the set of faces (dotted curve), and natural scenes (thin curve). As was the case for the spatial frequency analyses, for the orientation analyses, the entropy curves and mean amplitude curves were quite different in shape. This is revealed in several interesting ways. First, for natural scenes, the clear cardinal bias seen in the mean amplitude data (i.e., peaks in mean amplitude near 0 and 90 deg; see Fig. 7) actually reversed in terms of entropy. That is, for entropy there was a clear nadir at 0 deg (and a smaller nadir at 90 deg). Second, for faces, the clear peak at 0 deg in the mean amplitude data was absent in the entropy data. Third, for sign images, the clear peak at 90 deg in the amplitude data (seen for both signers) was largely attenuated in the entropy data. In addition, the entropy data for signer RB showed a peak at 0 deg, which

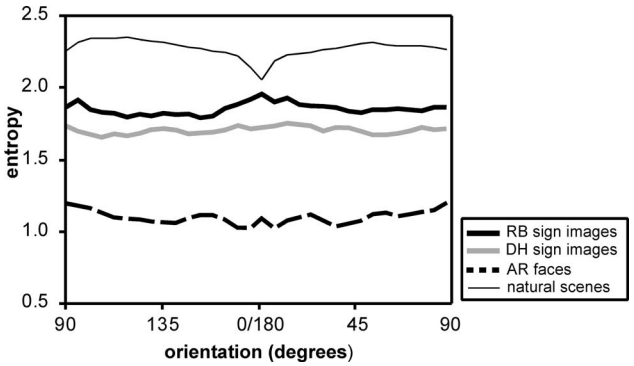

Fig. 8. Entropy versus orientation. Entropy is plotted as a function of orientation (data collapsed across spatial frequency) for signers RB (black curve) and DH (gray curve), faces (dotted curve), and natural scenes (thin curve). Entropy could not be calculated for the neutral pose, which is a single image, since in this study entropy was calculated for a distribution of many images. Note that the X axis is centered at 0 deg (horizontal) to emphasize the sharpest trough in the data.

was not observed in the amplitude data. We return to the relevance of these differences in Section 4.

4. DISCUSSION

The current study employed Fourier analysis to measure mean amplitude and entropy of the amplitude as a function of spatial frequency and orientation in three different sets of images: faces and natural scenes (from published databases) and ASL signs (from our own database). The analysis here addressed only the amplitude spectrum and not the phase. The amplitude spectrum contains information about the second-order statistics of the images, whereas the higher-order relationships are captured in the phase (e.g., see Bell and Sejnowski[73]). The natural image literature has focused primarily on amplitude as it is less clear how to characterize the phase structure of images. Progress on this front has been made by Thomson,[74] who showed that there are statistical consistencies in the phase spectra of natural scenes using the phase-only second spectrum, a fourth-order statistic that quantifies harmonic beat interactions in the data.

To our knowledge, this is the first study to quantify the image statistics of ASL. The novelty of this research is that it characterizes the low-level visual content of a visual language in the same way classic phonetics research characterizes the low-level auditory content of spoken languages. In addition, the current study introduces the concept of entropy across a set of images with the notion that, in addition to mean amplitude, this type of information in an image set may play a role in shaping visual perception.

Several studies have looked at entropy within images,[68,69,75,76] whereas the present study measured entropy across images. Kersten[75] showed that a given pixel value was 65% predictable by knowledge of the rest of the image. Ruderman[76] measured the mutual information between pairs of pixels as a function of separation distance and showed that the mutual information decreased nonlinearly with distance according to a power law. Daugman,[69] as well as Thomson,[74] compared the entropy of the original pixel values with the entropy of a wavelet distribution and showed that a wavelet representation reduced the entropy for natural images, i.e., created a sparse code. Brady and Field[68] showed that divisive normalization on a wavelet representation of natural images increases the entropy of the representation. These studies addressed the redundancy within images by measuring the joint entropy of pairs of image measures (or by measuring mutual information that is negatively related to entropy). The entropy across images measured in this study asked a different question, namely, which dimensions carry the most information for distinguishing images.

In this section, we discuss the results of the current (and previous) Fourier analyses conducted on different databases, with the goal of relating the findings to known (or predicted) visual psychophysical results. Before proceeding with the discussion, it is worthwhile pointing out that there are two main types of psychophysical tests that have been used to establish links between image statistics and visual perception. First, simple sensitivities (e.g., contrast sensitivity or acuity) are determined along a given stimulus dimension (like spatial frequency or orientation). Typically, the goal has been to link these psychophysical sensitivity data to biases in the environment as measured with the mean amplitude metric of the Fourier analysis (for example, as in the oblique effect). Second, with a set of real-world images (within a confined set, such as printed letters or faces), different spatial frequency or orientation bands are filtered or masked to determine which bands are most critical to the visual system for discriminating one image in that set from another. Here, the goal has been to ask whether the most critical band for visual discrimination is tied to that which contains the greatest mean amplitude (as has been suggested for printed letters; see Section 1). The novel contribution of the current paper, to our knowledge, is the suggestion that contrast sensitivity and/or visual discrimination may be tied to entropy of the amplitude in an image set rather than (or in addition to) the mean amplitude. Below, we begin with a discussion of natural scenes, followed by faces and then ASL signs.

A. Natural Scenes

With regard to spatial frequency, the results of our amplitude analyses revealed a $1/f$ function, which is in line with previous findings for natural scenes.[20–24] The slope of -1.17 for amplitude versus spatial frequency on a log scale is similar to slopes reported previously in the literature (e.g., Tolhurst *et al.*[24] reported a mean slope of -1.2). See Ruderman[77] and Balboa *et al.*[78] for discussions of the origins of $1/f$ scaling in natural scenes, and see Atick and Redlich[29] for a theory of how the $1/f$ scaling relates to contrast sensitivity in the retina.

The results of our entropy analyses of natural scenes looked quite different from the mean amplitude analyses. Spatial frequencies over $\sim$2 cyc/deg had more entropy compared with spatial frequencies under $\sim$2 cyc/deg; however, above and below this spatial frequency, entropy was relatively constant. The difference between the amplitude and the entropy results raises the interesting question of which measure (if either) is more influential

in shaping perception. If entropy plays a role, psychophysical studies that employ bandpass filtering or masking of natural scenes might show that middle and high spatial frequencies are more critical than low spatial frequencies for discriminating one natural scene from another.

In the orientation domain, the lowest amplitudes in natural scenes were observed for oblique orientations, which is in agreement with findings from previous studies of natural scenes.[12–16] For many years it has been believed that this orientation bias in the environment is what drives the perceptual oblique effect; acuity and contrast sensitivity are worse at oblique orientations as compared with cardinal orientations[6–8] (see Section 1). However, the results of our entropy analysis were essentially opposite to the results of our amplitude analysis (compare thin curves in Fig. 7 versus Fig. 8). More specifically, we observed the lowest amount of entropy for horizontal contours (i.e., troughs in entropy at 0 deg).

At first glance, the fact that our mean amplitude results, but not our entropy results, mirror the perceptual oblique effect would seem to suggest that mean amplitude across an image set is more influential in shaping perception than is the entropy in that image set. However, recent psychophysical studies present a challenging point of view.[17,63] In these studies it was shown that the oblique effect (i.e., poorest contrast sensitivity for oblique contours) that is known to exist for stimuli of a single spatial frequency turns into a horizontal effect (i.e., poorest contrast sensitivity for horizontal contours) for stimuli of broadband spatial frequencies that are more naturalistic. Also, the horizontal effect was found to be strongest when the slope of the $1/f$ spatial frequency distribution in the broadband stimulus was between 0.75 and 1 (in log–log coordinates), which matches the statistics of natural scenes. The similarity between these psychophysical findings of a horizontal effect and the sharp decrease in entropy seen at horizontal contours in the current study suggests that entropy might, in fact, play a role in shaping perceptual sensitivity. One way in which this could occur is if the visual system allocates less dynamic range to orientations with lesser entropy in the environment.

Another explanation for the horizontal effect, suggested by Hansen and Essock,[17] is that the effect is due to contrast normalization in the visual cortex. Interestingly, these two hypotheses, contrast normalization (Hansen and Essock's study) and entropy in the environment (current study), are not incompatible. Both accounts involve relying on orientations with the least redundancy, either within images as in the contrast normalization model or across images as in the current entropy analysis. Brady and Field[68] showed that divisive contrast normalization increases the entropy of neural responses to natural scenes. Brady and Field's study was primarily a within-image analysis, whereas the present study looks at entropy across images, but it shows that the computational process involved in divisive normalization is related to information-carrying capacity, a finding that is supported by Simoncelli and Schwartz.[79] In any event, one way to test whether entropy plays a role in shaping visual perception would be to use the filtering or masking technique to determine psychophysically whether horizontal contours are least critical for discriminating one natural scene from another.

B. Faces

With regard to spatial frequency, the results of our amplitude analyses revealed a $1/f$ function, which is in line with previous findings for faces (e.g., Torralba and Oliva[18]). Interestingly, however, the slope in log–log coordinates was steeper than for natural scenes (see Fig. 5). In our entropy analysis, we observed a small peak at 2.8 cyc/cm (which converts into 7 cyc/deg, assuming a viewing distance of 5 ft). These entropy data for faces can be compared with the results of psychophysical studies using filtering or masking to determine which spatial, frequencies human observers rely on most for discriminating faces. Although there is some differences across studies, the general consensus is that 10 to 15 cyc/face contain most information used by human observers,[80–83] with one study[84] being slightly higher at 20 cyc/face width. The cycles per centimeter of the current study can be converted into cycles per face by multiplying by the mean face width of our two signers (15.7 cm). This shows the peak in the entropy data to be 43.2 cyc/face, which is much higher than the peak cycles per face determined psychophysically. However, note that, although entropy peaks at ~43.2 cyc/face in our analyses, there is a broad range in this peak, which includes the range of important frequencies nr reported in the psychophysical literature (10–15 cyc/face). Still, our entropy data seem to overestimate the contribution of high spatial frequencies.

The orientation results for faces revealed a cardinal bias in mean amplitude of faces (as was observed for natural scenes), although the effect was stronger for horizontal than vertical contours. This cardinal bias seen in the amplitude for faces, which was also reported by Torralba and Oliva,[18] was largely attenuated in the entropy data of faces (compare dotted curves Fig. 7 versus Fig. 8).

C. Signs in American Sign Language

The statistics of sign images were found to differ from those observed for natural scenes and faces in the following ways. First, although mean amplitude analyses of all image sets revealed a $1/f$ function, the slope for sign images was steeper than for natural scenes but shallower than for faces. Because the slope of the sign images falls midway between that of two other image sets for which signers should also have ample experience (faces and natural scenes), it is not entirely obvious whether spatial frequency sensitivity would be expected to differ between signers and nonsigners. To date, a single study has compared contrast sensitivity to different spatial frequencies in signers versus nonsigners.[85] For the three spatial frequencies tested in that study (0.5, 2, and 9 cyc/deg), no significant differences were found between groups, nor were any interactions between subject group and spatial frequency found. Such findings suggest that exposure to the spatial frequency distribution of sign images has no effect on contrast sensitivity, although the relatively small number of spatial frequencies tested in this study may not have been sufficient for revealing subject group differences.

Second, although all image sets revealed higher mean amplitudes for cardinal, than for oblique, orientations, sign images contained more amplitude for vertical, than for horizontal, contours, while faces and natural scenes showed an opposite pattern. This difference may not be surprising, since signing arms are often oriented vertically. (This is especially true for one-handed signs, where the nondominant arm is usually in a vertical position.) If there is a relationship between the prevalence of different orientations in the environment and visual sensitivity (see Section 1), these mean amplitude analyses for orientation suggest that anisotropies in orientation perception may differ between ASL signers and nonsigners. Future psychophysical studies will be required to test this hypothesis.

A third way in which the statistics of ASL images were found to differ from those observed for natural scenes and faces was in the entropy analyses for spatial frequency (see Fig. 6). Here, sign images revealed a clear bandpass pattern with a peak in entropy at approximately 0.75 cyc/cm (which converts into 1.9 cyc/deg, assuming a typical viewing distance of 5 ft). This result suggests that spatial frequencies near this peak contain the greatest amount of information for distinguishing one sign image from another. To investigate whether this entropy signal in the image set is actually used by human observers, one can ask whether the spatial frequencies that contain the highest entropy in an image set are those most critical for psychophysical discrimination of sign images. Psychophysical data from Riedl and Sperling[62] allow us to address this possibility. In their study, they presented spatial-frequency-filtered signs to determine which spatial frequencies are most critical for signers to be able to discriminate one sign from another. They used four spatial frequency filters with peaks of 0.05, 0.20, 0.35, and 0.75 cyc/cm and found that 0.75 cyc/cm was most important. Note that the investigators were actually attempting to make it so that the four different spatial frequency bands yielded equal performance (which they attempted to do by keeping the peak frequency constant but varying the bandwidth of the filters), yet they still found that the subjects' most accurate performance was seen at approximately 0.75 cyc/cm.

Although Riedl and Sperling[62] did not test high enough spatial frequency bands to determine whether the 0.75 cyc/cm band was the peak (since it was the highest band tested), their results are, at the very least, in line with those of the current study, which showed that the image statistics of signs contain the highest entropy (i.e., the greatest amount of information) at around 0.75 cyc/cm. This consistency between the entropy statistics of signs here and psychophysical discrimination of signs reported by Riedl and Sperling suggests that the mechanisms involved in comprehension of sign language might become tuned to rely on spatial frequencies in signs that contain the greatest entropy.

D. Are Signing Arms and Hands Unique?

On a final note, it is important to point out that if perceptual differences are found to exist between signers and nonsigners and if these differences appear to relate to the image statistics of the arms and hands in ASL, we suggest that it is more likely due to signers' reliance on ASL rather than to their mere exposure. This is because the case could easiy be made that both signers and nonsigners alike receive ample (and roughly equal) experience viewing arms and hands, whether those arms are signing (experienced by signers only), gesturing, or at rest (the latter two being experienced by signers and nonsigners). And in fact, the results of our mean amplitude analyses revealed only minor differences between sign images and the neutral pose (in terms of the shapes of the spatial frequency and orientation curves). Thus, differences between signers and nonsigners would most likely be attributable to the fact that signers pay much more attention to the hands and arms because this part of the body is crucial for ASL communication.

Finally, the current Fourier analysis focuses on the static information carried in ASL, i.e., in the hand–arm position that carries lexical information. However, the motion of the hands and arms also carries critical information in ASL, and we address the statistics of the motion in ASL in a forthcoming paper on the motion statistics of ASL. In any event, it is reasonable to suggest that visual perception may differ between signers and nonsigners in predictable ways based on the (static or moving) image statistics of signs. Many studies have indeed reported differential spatial or motion processing between signers and nonsigners, which is thought to be due to ASL experience.[50–56] However, because the spatial frequency–orientation makeup of the stimuli in these previous studies was not controlled, we cannot yet determine whether differences between signers and nonsigners in these studies might be linked to exposure to the visual statistics of signing arms and hands. Future experiments in our laboratory are currently underway to investigate these possibilities.

ACKNOWLEDGMENTS

This work was supported by a National Science Foundation (NSF) grant BCS-0241557 to K. R. Dobkins and an NSF grant IIS-0220141 to M. S. Bartlett. We thank Deborah M. Harlan for producing the sign images for the ASL database.

Corresponding author K. R. Dobkins can be reached at the Department of Psychology, 0109, University of California, San Diego, La Jolla, California 92093, or by phone, 858-534-5434; fax, 858-534-7190; or e-mail, kdobkins@ucsd.edu.

REFERENCES

1. C. Blakemore and G. F. Cooper, "Development of the brain depends on the visual environment," Nature **228**, 477–478 (1970).
2. C. Blakemore and G. F. Cooper, "Modification of the visual cortex by experience," Brain Res. **31**, 366 (1971).
3. M. P. Stryker, H. Sherk, A. G. Leventhal, and H. V. Hirsch, "Physiological consequences for the cat's visual cortex of effectively restricting early visual experience with oriented contours," J. Neurophysiol. **41**, 896–909 (1978).
4. J. Gwiazda, I. Mohindra, S. Brill, and R. Held, "Infant astigmatism and meridional amblyopia," Vision Res. **25**, 1269–1276 (1985).
5. D. E. Mitchell, R. D. Freeman, M. Millodot, and G.

Haegerstrom, "Meridional amblyopia: evidence for modification of the human visual system by early visual experience," Vision Res. **13**, 535–558 (1973).
6. S. Appelle, "Perception and discrimination as a function of stimulus orientation: the 'oblique effect' in man and animals," Psychol. Bull. **78**, 266–278 (1972).
7. F. W. Campbell, J. J. Kulikowski, and J. Levinson, "The effect of orientation on the visual resolution of gratings," J. Physiol. (London) **187**, 427–436 (1966).
8. D. E. Mitchell, R. D. Freeman, and G. Westheimer, "Effect of orientation on the modulation sensitivity for interference fringes on the retina," J. Opt. Soc. Am. **57**, 246–249 (1967).
9. S. Sokol, A. Moskowitz, and V. Hansen, "Electrophysiological evidence for the oblique effect in human infants," Invest. Ophthalmol. Visual Sci. **28**, 731–735 (1987).
10. S. Sokol, A. Moskowitz, and V. Hansen, "Evoked potential and preferential looking correlates of the oblique effect in 3-month-old infants," Doc. Ophthalmol. **71**, 321–328 (1989).
11. D. Y. Teller, R. Morse, R. Borton, and D. Regal, "Visual acuity for vertical and diagonal gratings in human infants," Vision Res. **14**, 1433–1439 (1974).
12. R. J. Baddeley and P. J. Hancock, "A statistical analysis of natural images matches psychophysically derived orientation tuning curves," Proc. R. Soc. London, Ser. B **246**, 219–223 (1991).
13. D. M. Coppola, H. R. Purves, A. N. McCoy, and D. Purves, "The distribution of oriented contours in the real world," Proc. Natl. Acad. Sci. U.S.A. **95**, 4002–4006 (1998).
14. M. S. Keil and G. Cristobal, "Separating the chaff from the wheat: possible origins of the oblique effect," J. Opt. Soc. Am. A **17**, 697–710 (2000).
15. E. Switkes, M. J. Mayer, and J. A. Sloan, "Spatial frequency analysis of the visual environment: anisotropy and the carpentered environment hypothesis," Vision Res. **18**, 1393–1399 (1978).
16. A. van der Schaaf and J. H. van Hateren, "Modelling the power spectra of natural images: statistics and information," Vision Res. **36**, 2759–2770 (1996).
17. B. C. Hansen and E. A. Essock, "A horizontal bias in human visual processing of orientation and its correspondence to the structural components of natural scenes," J. Vision **4**, 1044–1060 (2004).
18. A. Torralba and A. Oliva, "Statistics of natural image categories," Network **14**, 391–412 (2003).
19. R. C. Annis and B. Frost, "Human visual ecology and orientation anisotropies in acuity," Science **182**, 729–731 (1973).
20. R. M. Balboa and N. M. Grzywacz, "Power spectra and distribution of contrasts of natural images from different habitats," Vision Res. **43**, 2527–2537 (2003).
21. G. J. Burton and I. R. Moorhead, "Color and spatial structure in natural scenes," Appl. Opt. **26**, 157–170 (1987).
22. D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," J. Opt. Soc. Am. A **4**, 2379–2394 (1987).
23. D. L. Ruderman and W. Bialek, "Statistics of natural images: scaling in the woods," Phys. Rev. Lett. **73**, 814–817 (1994).
24. D. J. Tolhurst, Y. Tadmor, and T. Chao, "Amplitude spectra of natural images," Ophthalmic Physiol. Opt. **12**, 229–232 (1992).
25. D. C. Knill, D. Field, and D. Kersten, "Human discrimination of fractal images," J. Opt. Soc. Am. A **7**, 1113–1123 (1990).
26. C. A. Parraga, T. Troscianko, and D. J. Tolhurst, "The effects of amplitude-spectrum statistics on foveal and peripheral discrimination of changes in natural images, and a multi-resolution model," Vision Res. **45**, 3145–3168 (2005).
27. Y. Tadmor and D. J. Tolhurst, "Discrimination of changes in the second-order statistics of natural and synthetic images," Vision Res. **34**, 541–554 (1994).
28. D. J. Tolhurst and Y. Tadmor, "Discrimination of spectrally blended natural images: optimisation of the human visual system for encoding natural images," Perception **29**, 1087–1100 (2000).
29. J. Atick and A. Redlich, "What does the retina know about natural scenes?" Neural Comput. **4**, 196–210 (1992).
30. P. L. Clatworthy, M. Chirimuuta, J. S. Lauritzen, and D. J. Tolhurst, "Coding of the contrasts in natural images by populations of neurons in primary visual cortex (V1)," Vision Res. **43**, 1983–2001 (2003).
31. B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," Nature **381**, 607–609 (1996).
32. Y. Tadmor and D. J. Tolhurst, "Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes," Vision Res. **40**, 3145–3157 (2000).
33. M. V. Srinivasan, S. B. Laughlin, and A. Dubs, "Predictive coding: a fresh view of inhibition in the retina," Proc. R. Soc. London, Ser. B **216**, 427–459 (1982).
34. E. P. Simoncelli and B. A. Olshausen, "Natural image statistics and neural representation," Annu. Rev. Neurosci. **24**, 1193–1216 (2001).
35. W. S. Geisler, J. S. Perry, B. J. Super, and D. P. Gallogly, "Edge co-occurrence in natural images predicts contour grouping performance," Vision Res. **41**, 711–724 (2001).
36. B. C. Regan, C. Julliot, B. Simmen, F. Vienot, P. Charles–Dominique, and J. D. Mollon, "Frugivory and colour vision in Alouatta seniculus, a trichromatic platyrrhine monkey," Vision Res. **38**, 3321–3327 (1998).
37. C. C. Chiao, D. Osorio, M. Vorobyev, and T. W. Cronin, "Characterization of natural illuminants in forests and the use of digital video data to reconstruct illuminant spectra," J. Opt. Soc. Am. A **17**, 1713–1721 (2000).
38. J. N. Lythgoe and J. C. Partridge, "Visual pigments and the acquisition of visual information," J. Exp. Biol. **146**, 1–20 (1989).
39. D. Osorio and M. Vorobyev, "Colour vision as an adaptation to frugivory in primates," Proc. R. Soc. London, Ser. B **263**, 593–599 (1996).
40. J. Pokorny and V. C. Smith, "Evaluation of single-pigment shift model of anomalous trichromacy," J. Opt. Soc. Am. **67**, 1196–1209 (1977).
41. J. D. Mollon, "Color vision," Annu. Rev. Physiol. **33**, 41–85 (1982).
42. M. A. Webster, "Pattern selective adaptation in color and form perception," in *The Visual Neurosciences*, M. L. Chalupa and S. J. Werner, eds. (MIT, 2003), pp. 936–947.
43. M. A. Webster and J. D. Mollon, "Adaptation and the color statistics of natural images," Vision Res. **37**, 3283–3298 (1997).
44. N. Yendrikhovskij, "Computing color categories from statistics of natural images," J. Imaging Sci. Technol. **45**, 409–417 (2001).
45. S. T. Chung, "The effect of letter spacing on reading speed in central and peripheral vision," Invest. Ophthalmol. Visual Sci. **43**, 1270–1276 (2002).
46. N. J. Majaj, D. G. Pelli, P. Kurshan, and M. Palomares, "The role of spatial frequency channels in letter identification," Vision Res. **42**, 1165–1184 (2002).
47. D. H. Parish and G. Sperling, "Object spatial frequencies, retinal spatial frequencies, noise, and the efficiency of letter discrimination," Vision Res. **31**, 1399–1415 (1991).
48. J. A. Solomon and D. G. Pelli, "The visual filter mediating letter identification," Nature **369**, 395–397 (1994).
49. E. Poder, "Spatial-frequency spectra of printed characters and human visual perception," Vision Res. **43**, 1507–1511 (2003).
50. C. J. Brozinsky and D. Bavelier, "Motion velocity thresholds in deaf signers: changes in lateralization but not in overall sensitivity," Brain Res. Cognit. Brain Res. **21**, 1–10 (2004).
51. D. Bavelier, C. Brozinsky, A. Tomann, T. Mitchell, H. Neville, and G. Liu, "Impact of early deafness and early exposure to sign language on the cerebral organization for motion processing," J. Neurosci. **21**, 8931–8942 (2001).
52. D. Bavelier, A. Tomann, C. Hutton, T. Mitchell, D. Corina, G. Liu, and H. Neville, "Visual attention to the periphery is

enhanced in congenitally deaf individuals," J. Neurosci. **20**, 1–6 (2000).
53. K. Emmorey, E. Klima, and G. Hickok, "Mental rotation within linguistic and non-linguistic domains in users of American sign language," Cognition **68**, 221–246 (1998).
54. K. Emmorey and S. M. Kosslyn, "Enhanced image generation abilities in deaf signers: a right hemisphere effect," Brain Cogn. **32**, 28–44 (1996).
55. R. G. Bosworth and K. R. Dobkins, "Left hemisphere dominance for motion processing in deaf signers," Psychol. Sci. **10**, 256–262 (1999).
56. R. G. Bosworth and K. R. Dobkins, "Visual field asymmetries for motion processing in deaf and hearing signers," Brain Cogn **49**, 170–181 (2002).
57. D. Brentari, *A Prosodic Model of American Sign Language Phonology* (MIT, 1998).
58. D. Perlmutter, "Sonority and syllable structure in American Sign Language," Linguist. Inquiry **23**, 407–442 (1992).
59. R. B. Wilbur and A. M. Martinez, "Physical correlates of prosodic structure in American Sign Language," presented at the Meeting of the Chicago Linguistics Society, April 25–27, 2002.
60. T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, 1991).
61. J. J. Atick, "Could information theory provide an ecological theory for sensory processing?" Network Comput. Neural Syst. **3**, 231–251 (1992).
62. T. R. Riedl and G. Sperling, "Spatial-frequency bands in complex visual stimuli: American Sign Language," J. Opt. Soc. Am. A **5**, 606–616 (1988).
63. E. A. Essock, J. K. DeFord, B. C. Hansen, and M. J. Sinai, "Oblique stimuli are seen best (not worst!) in naturalistic broad-band stimuli: a horizontal effect," Vision Res. **43**, 1329–1335 (2003).
64. P. Eccarius and D. Brentari, "Symmetry and dominance: a cross-linguistic study of signs and classifier constructions," Lingua (to be published); www.sciencedirect.com.
65. R. G. Bosworth, C. E. Wright, M. S. Bartlett, D. P. Corina, and K. R. Dobkins, "Characterization of the visual properties of spatial frequency and speed in ASL signs," in *Cross-Linguistic Perspectives in Sign Language Research. Selected Papers from TISLR 2000*, A. E. Baker, B. van den Bogaerde, and O. Crasborn, eds. (Signum, 2003), pp. 265–282.
66. A. M. Martinez and R. Benavente, "The AR face database," CVC Tech. Rep. 24 (Computer Vision Center, Universitat Autònama de Barcelona, 1998) Available at http://rvl1.ecn.purdue.edu/~aleix/ar.html.
67. J. H. van Hateren and A. van der Schaaf, "Independent component filters of natural images compared with simple cells in primary visual cortex," Proc. R. Soc. London, Ser. B **265**, 359–366 (1998).
68. N. Brady and D. J. Field, "Local contrast in natural images: normalisation and coding efficiency," Perception **29**, 1041–1055 (2000).
69. J. G. Daugman, "Entropy reduction and decorrelation in visual coding by oriented neural receptive fields," IEEE Trans. Biomed. Eng. **36**, 107–114 (1989).
70. D. Field, "What is the goal of sensory coding?" Neural Comput. **6**, 559–601 (1994).
71. J. Huang and D. Mumford, "Statistics of natural images and models," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Press, 1999), pp. 541–547.
72. M. S. Lewicki, "Efficient coding of natural sounds," Nat. Neurosci. **5**, 356–363 (2002).
73. A. J. Bell and T. J. Sejnowski, "The 'independent components' of natural scenes are edge filters," Vision Res. **37**, 3327–3338 (1997).
74. M. G. Thomson, "Beats, kurtosis and visual coding," Network **12**, 271–287 (2001).
75. D. Kersten, "Predictability and redundancy of natural images," J. Opt. Soc. Am. A **4**, 2395–2400 (1987).
76. D. L. Ruderman, "The statistics of natural images," Network Comput. Neural Syst. **5**, 517–548 (1994).
77. D. L. Ruderman, "Origins of scaling in natural images," Vision Res. **37**, 3385–3398 (1997).
78. R. M. Balboa, C. W. Tyler, and N. M. Grzywacz, "Occlusions contribute to scaling in natural images," Vision Res. **41**, 955–964 (2001).
79. E. P. Simoncelli and O. Schwartz, "Modeling surround suppression in V1 neurons with a statistically-derived normalization model," Adv. Neural Inf. Process. Syst. **11**, 153–159 (1999).
80. N. P. Costen, D. M. Parker, and I. Craw, "Effects of high-pass and low-pass spatial filtering on face identification," Percept. Psychophys. **58**, 602–612 (1996).
81. J. Gold, P. J. Bennett, and A. B. Sekuler, "Identification of band-pass filtered letters and faces by human and ideal observers," Vision Res. **39**, 3537–3560 (1999).
82. E. Peli, E. Lee, C. L. Trempe, and S. Buzney, "Image enhancement for the visually impaired: the effects of enhancement on face recognition," J. Opt. Soc. Am. A **11**, 1929–1939 (1994).
83. T. Tieger and L. Ganz, "Recognition of faces in the presence of two-dimensional sinusoidal masks," Percept. Psychophys. **26**, 163–167 (1979).
84. T. Hayes, M. C. Morrone, and D. C. Burr, "Recognition of positive and negative bandpass-filtered images," Perception **15**, 595–602 (1986).
85. E. M. Finney and K. R. Dobkins, "Visual contrast sensitivity in deaf versus hearing populations: exploring the perceptual consequences of auditory deprivation and experience with a visual language," Brain Res. Cognit. Brain Res. **11**, 171–183 (2001).