Psychophysical Experiments on the Internet

Hendrik−Jan van Veen, Heinrich H. Bülthoff and Galia Givaty

Max−Planck−Institute for Biological Cybernetics • Tübingen • Germany

E−mail: hendrik−jan.veen@tuebingen.mpg.de

MOTIVATION

Psychologists have recently discovered the Internet for demonstrating visual illusions and for education. We have examined the feasibility of using the Internet for another purpose, namely large−scale data collection in visual psychophysics.

Web−experiments promise access to a huge number of anonymous subjects. The technique is therefore potentially suitable for experiments that need few data per subject from many subjects, and (i) need to cover a large parameter space, or (ii) are designed to catalogue possible behaviours, or (iii) demand subject diversity. Experiments that require precise control of timing, colour, contrast, size or resolution should in general not be considered for Internet implementation.

METHODS

Over the past year we have implemented five different web experiments using various Internet technologies (HTML, VRML, CGI, JAVA, etc.). We used a control group that performed each web−experiment under controlled conditions on one of our own computers to validate the results obtained from the anonymous Internet subjects (the main subject group). Members of the lab participating in web−experiments formed a second control group. Here we report on three of the five experiments we conducted. For one experiment (Exp. II) a directly comparable dataset from a traditional lab experiment is available.

RESULTS & CONCLUSIONS

In the first year more than 2000 subjects participated in one or more web−experiments. A comparison between the results of the anonymous Internet subjects, our laboratory subjects, labmembers as subjects, and subjects participating in classical laboratory experiments strongly confirms the general validity of our web−experiment data.

We conclude that web−experiments form a valuable and valid method for accessing large groups of subjects, provided careful thought is given to the limitations of using anonymous subjects on remote locations. The technique is especially well suited for performing quick pilot studies and for validating laboratory experiments using larger numbers of subjects.

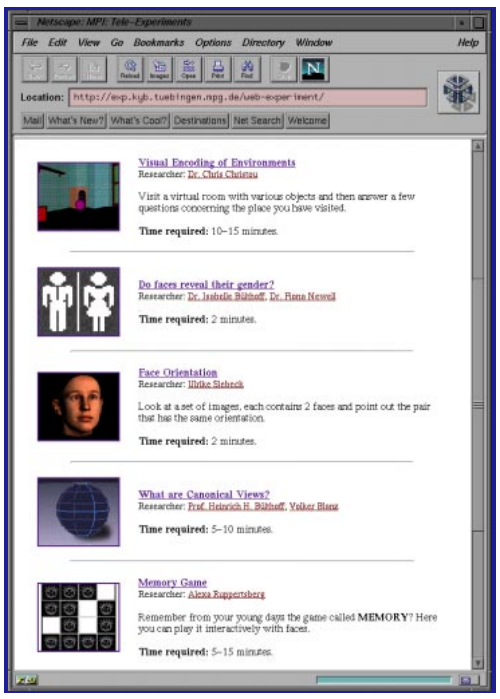

The figure shows part of the webpage containing the web−experiments (http://exp.kyb.tuebingen.mpg.de/web−experiment/). More specifically, it shows the list of experiments that are currently available. Apart from the three experiments described below we are running two more experiments on 'Canonical Views of 3D Objects' and 'Face Recognition'.

I: Gender Perception

collaborators: Isabelle Bülthoff & Fiona N. Newell

Is the texture or the shape of a face the better gender indicator?

STIMULI
Given a database of over two hundred 3D models of real faces (scanned), we computed the average male face, the average female face, and the overall average face (referred to as 'neutral'). We separated the shape of these faces from the textures and created 6 new faces by combining the shape of one face with the texture of another (see Vetter & Blanz for details). The stimuli consisted of colour pictures of these 9 faces in a fixed orientation and with a frontal illumination.

PARADIGM
Subjects were presented with 3 faces in sequence and were required to judge the gender of each face. The first face was always one of the combination faces or the neutral face, the second and third faces were the average male and female faces (presented in random order). Subjects that made mistakes in either of the last two faces were removed from the dataset.

RESULTS
Out of 1089 subjects in total, 374 made an error judging the gender of one or both of the last two faces (this includes 2 out of 16 control subjects but none of the 32 lab members). 22% of those subjects judged the male to be female, and 86% called the female a male. We restrict the rest of the analysis to those 715 subjects who judged those two faces correctly. The graph to the right plots the percentage of female judgements for the different combination faces and the neutral face. There is a clear tendency to answer 'male' (e.g., see NT−NS). Apart from this interesting bias, we do get significant fractions of 'female' answers only when the texture is female or when the texture is neutral but the shape is female. Apparently, texture is a much stronger indicator of gender than shape is. We found a slight but not significant tendency of female subjects to answer relatively more 'female' than male subjects do, but they are still strongly biased towards answering 'male'.

red (control), green (lab members), blue (Internet), yellow (all subjects together). Along the abscissa the faces are decoded as follows: F (female), M (male), N (neutral), S (shape), T (texture). So, FT−MS, means a female texture combined with a male shape.

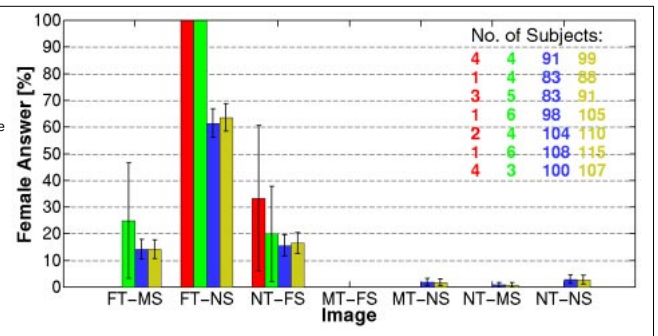

CONCLUSION
Without using the Internet it would have been practically impossible to get more than thousand subjects to perform this very short experiment which uses stimuli that can be used only once per subject. Note that we still need more subjects and face stimuli to decide whether the subject−gender effect is significant, and whether it depends on the particular stimulus properties or not (e.g., currently the heads are shown without hair).

Vetter, T. & Blanz, V. (1998). Estimating coloured 3D face models from single images: an example based approach. Proceedings of ECCV'98 Freiburg, Germany. Springer, LNCS 1407, 499−513.

II: Illusory Face Rotation

collaborator:Ulrike Siebeck

How does illumination direction influence the apparent orientation of a face?

STIMULI
We used colour pictures of five 3D face−models (scanned heads) in various orientations and illuminated from different directions (see below).

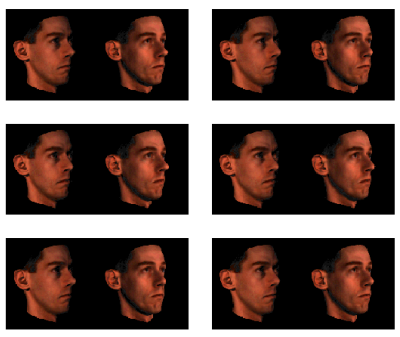

Example stimulus

PARADIGM
Subjects saw 6 pairs of images of the same head simultaneously. Within each pair one face was illuminated 30° from the left, the other 30° from the right. The mean orientation of the faces was either 15°, 30°, or 45° away from straight ahead (between subjects factor). The only difference between the 6 pairs was the physical orientation difference between the two depicted faces, which could be either −3°, 0°, 3°, 6°, 9°, or 12°. The subject's task was to select the pair in which the faces appeared to have the same orientation. Each subject did this with only one of the 5 faces (between subjects factor).

RESULTS
487 anonymous Internet subjects participated in the experiment. Plotted to the right is the illumination induced apparent orientation shift of the faces as given by the physical orientation difference compensating the illumination induced effects. We reproduce the results of Troje & Siebeck who performed a laboratory study using the same stimuli in a sequential 2AFC constant stimuli paradigm: Faces appear to rotate in the direction opposite to the rotation of the light source. The main difference with the Troje & Siebeck results is the increased effect size at 15° orientation. According to other results from Siebeck, this might be explained by the difference in visual size of the pictures between the two studies.

red (control), green (lab members), blue (Internet), yellow (all web experiment subjects together), black (experiment Troje & Siebeck)

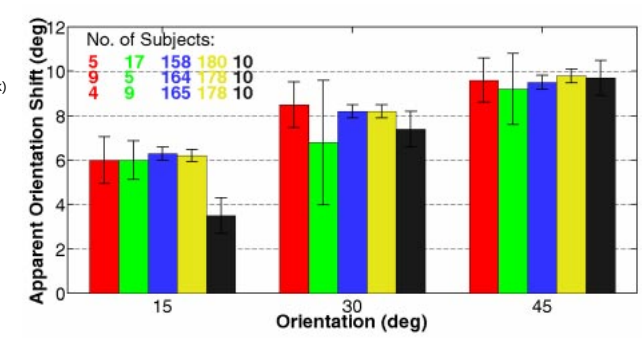

CONCLUSION
A visual illusion that strongly depends on fine contrast in the stimulus was reproduced quantitatively by anonymous subjects who used unknown equipment and participated over the Internet.

Ulrike Siebeck (1997). Orientierungseinschätzung bei menschlichen Gesichtern. Masters Thesis Fakultät für Biologie der Universität Tübingen.
Troje, N.F. & Siebeck, U. (1998). Illumination−induced apparent shift in orientation of human heads. *Perception, 27*, 671−680.

III: Visual encoding of Scenes

collaborator:Chris Christou

Are the position, shape, and colour of objects encoded independently or not?

STIMULI
We created a 3D interactive view of a small, irregularly shaped room containing 5 volumetric objects lying on the floor. Each of these objects had a unique colour and a unique shape.

View of the room

PARADIGM
Subjects were asked to explore the room as long as they wanted (up to 10 minutes). Using mouse & keyboard they could move through and around the room. To make sure that they had seen all 5 objects they had to click on each of them once. The exploration phase ended by clicking on the door of the room. Then a randomly oriented 2D map of the room appeared along with a set of 2D representations of the objects in all possible colours. The subject then had to recreate the interior of the room by selecting the appropriate objects from the set and putting them on the map of the room in one of the 5 given positions. They could correct any mistakes as often as they wanted.

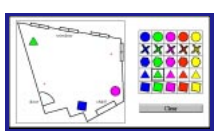

Room reconstruction

RESULTS
83 anonymous Internet subjects completed the experiment. The graph to the far right shows the average number of errors subjects made for different combinations of shape, colour and position. For instance, "position−colour" reflects the number of errors subjects made in combining the right colour with position, irrespective of shape. None of the combinations pos−col, pos−shape, and col−shape show decreased error rates, leading us to believe that those attributes are indeed encoded independently in the brain. More subjects are needed to reveal whether the tendency of decreased rates for the combination colour−shape proves to be a true effect. The lower error rates produced by lab members can be explained by their knowledge of the map creation task prior to the actual experiment.

red (control), green (lab members), blue (Internet), yellow (all web experiment subjects together).

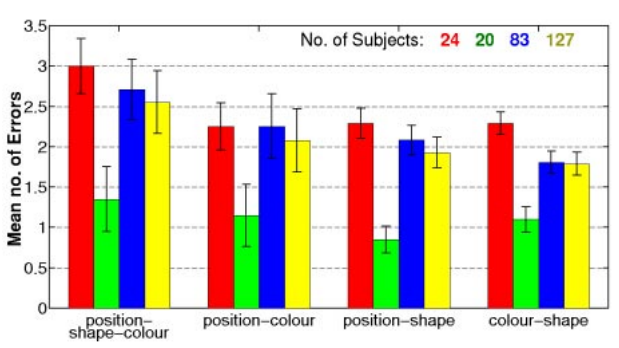

CONCLUSION
An experiment that uses a simple form of virtual reality on the Internet proved feasible, but attracted relatively few subjects due to the relatively long time it took to complete and because of browser−incompatibilitites (a VRML plug−in was required).