What does water look like?

Marta Kryven* William Cowan†

University of Waterloo

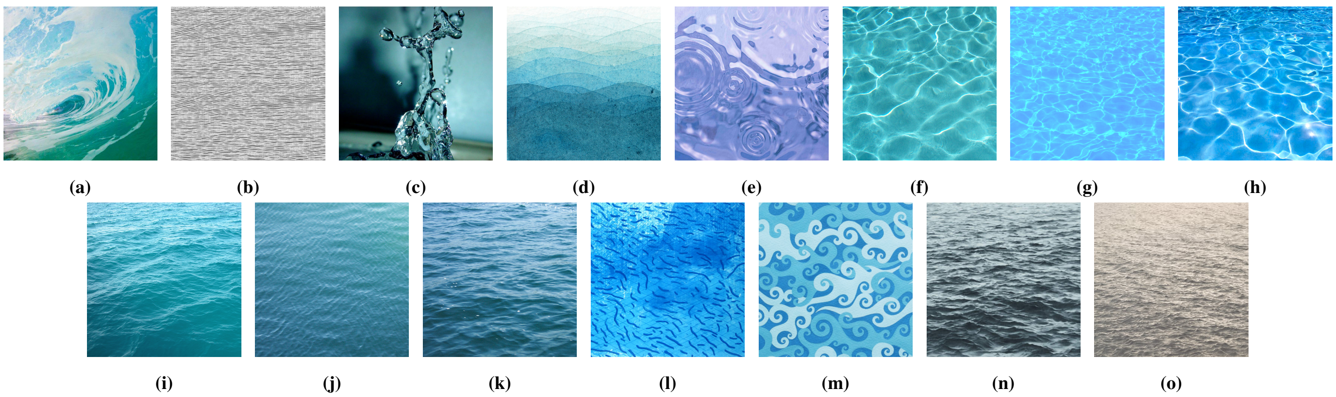

Figure 1: *A selection of test stimuli used during pilot experiments. Hard: 1a-1d; Easy: 1e -1k; Hard for some, easy for others: 1l - 1o.*

Abstract

What makes images of water look like water? We conducted four psychophysical experiments to isolate perceptual qualities that make water easy to recognize. Water recognition is facilitated by colour and by three patterns of waves. Low spatial frequencies (LSF) (<4.4 cpd) contribute more to recognition than high spatial frequencies (HSF). Here we describe the experimental methodology and results. Knowing which aspects of appearance identify water can inform perceptually inspired depiction of water, can create visual illusions and can reduce computation in realistic simulations.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation— [J.4]: Social and Behavioural Sciences—Psychology I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—Representations, data structures, and transforms I.4.8 [Image Processing and Computer Vision]: Scene analysis—;

Keywords: water simulation, spatial frequency scales, object recognition, non-photorealistic rendering, perception

1 Introduction

Water is interesting and ubiquitous. It can assume many forms and colours. Despite this diversity unattended recognition of water is common, essential for swimming, sailing and avoiding puddles. We would like to understand these visual skills and build on them when depicting water. Interactive graphics needs easy interpretation more than it does strict reality. In this paper we ask how does a substance as changeable as water should be depicted when it is not the focus of attention? Should the depiction be realistic? What colour should it be? Should it focus on local or global detail? To find out we adopted the experimental methods of perception to study early perceptual processes of identifying water.

2 Related Work

Most image recognition research supports computer vision, focusing on recognizing physical qualities of objects [Liu et al. 2010] [Perina et al. 2010] [Vogel and Schiele 2007]. In graphics the goal is human response and the focus on stimulating perception. Thus, unlike computer vision research we included illustrations of water in our experiments, which explored three aspects of visual processing. First, humans often respond to illustrations more readily than to realistic depictions [Mills 1985]. Line-drawings, for example, are easy to recognize since visual recognition is defined by the silhouette and by lines of self-occlusion [Hoffman and Singh 1997]. Although water has no fixed silhouette or shape, line-drawings of it are common.

Second, when colour is predictive, as when recognizing food but not when recognizing animals, it affects recognition accuracy [Delorme et al. 2000]. Does colour help water recognition? When water is shallow and pure it is transparent. Deep water is a saturated blue-green colour [Pope and Fry 1997]. Muddy water is brown. The variability of colours should make recognition of water insensitive to colour. In contrast, illustration often limits images of water to green-blue hues.

Another common feature of water illustration is visual texture. What aspects of visual texture are important for recognition? Because vision has separate channels for HSF and LSF frequencies and natural scenes are weighted toward LSF [Field 1987], we expect LSF to be more important. In contrast, illustrations are weighted toward HSF.

In order to refine our intuitions we assembled a set of stereotypical depictions of water representing a wide variety of styles [Kryven and Cowan 2013] and in subsequent experiments narrowed down to comparing only images identified as easy. This led us to reject several usual representations of water. To our knowledge, this study

*e-mail:mkryven@uwaterloo.ca
†e-mail:wmcowan@uwaterloo.ca

CAe 2014, August 08 – 10, 2014, Vancouver, British Columbia, Canada.

is the first application of psychophysical experiments to defining visually salient properties of water relevant to computer graphics.

3 Methodology and Apparatus

We employed a method of brief exposure to visual stimuli where subjects' reaction times(RT) and percent correct are used as measures of accuracy. Visual processing of complex scenes requires 120ms [Holcombe 2009] of exposure, although some low-level information can be processed after only 39ms [Bar et al. 2006]. Humans can accurately respond to natural scenes in a categorization task in 400ms [Thorpe et al. 1996].

Method. Experiments took place in a dark room using a calibrated Apple laptop computer. 500x500-pixel images were displayed 50-cm from the subjects eyes subtending 10 degrees of visual angle. Responses were collected from the keyboard using custom built software. We conducted a series of two-alternative forced choice image recognition tasks. Figure 2 shows timing of each trial. First a fixation cross standardizes the direction of gaze between two uniform grey squares. After 500ms one square is replaced by an image of water, the other by an image unrelated to water. After 120ms the images are replaced by masks, which remain until the subject responds by pressing a key on the keyboard indicating the position of the water image. The response initiates the next trial. Subjects were asked to respond as quickly as possible without making mistakes. The average response time was 450ms, typical of early vision.

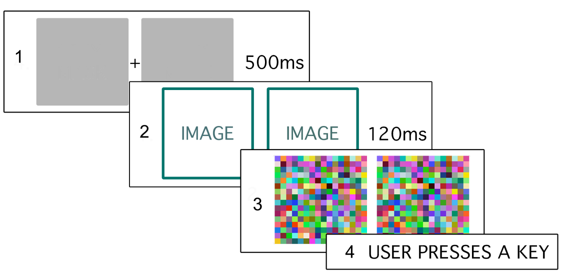

Figure 2: *Experiment trial sequence.*

Data analysis. Outliers more than two standard deviations away from the mean were recursively removed [Selst and Jolicoeur 1994] eliminating 5 % of the data. We regressed RT against trial number. Every subject's data showed a highly significant non-zero slope, usually negative. We retained the residuals to detrend the data. Side preference was similarly removed, data-set by data-set. Degrees of freedom consumed by data cleaning are negligible compared to amounts of data gathered. Distractors were contextually unrelated to water. Each target distractor pair occurred twice, once left right and once right left. Subjects always did 20 practise trials at the beginning the experiment.

4 Experiments

4.1 Pilot Experiment 1. Realism or Illustration?

Motivation. Water has many appearances. As with any image recognition experiment, assembling a large, inclusive and high quality data-set is hard. For example, while water in Turner's paintings is recognizable, when cropped out of context of the painting it becomes ambiguous. Thus, when choosing illustrations we selected unambiguous depictions, as, for example, line drawings of waves, realistic drawings by Vija Celmins and swimming pool paintings by David Hockney. We assembled a set of 30 images of water including different depiction styles: photographs, drawings, computer-generated images, vector art and paintings, representing surfaces, splashes and breaking waves. Figure 1 shows a subset for which statistically significant differences in RT were observed (the complete dataset is available at `http://www.cgl.uwaterloo.ca/~mkryven/`). Images were found on the Internet and cropped to remove context. Distractors were matched in photographic and illustrative styles.

Subjects. The subjects were 6 males and 6 females, graduate students and faculty, all had normal or corrected to normal vision. The same subjects repeated the experiments two to five times to generate a large corpus of data.

Image	Category	RT diff. from mean (ms)	p
	surface photograph	-9.766	0.0001
	splash	+5.878	0.0056
1a	splash	+9.2	0.04
1b	illustration	+13.5	0.0035
1c	splash	+10	0.004
1d	illustration	+17.6	0.0003
1e	surface photograph	-5.91	0.09
1f	surface photograph	-12.5	0.0004

Table 1: *RT for images categories and individual images from Figure 1 show that photographs are on average recognized faster.*

Result. We explored the data using one-way ANOVA (factors: subject, target category [photograph, splash, illustration], distractor category [illustration, natural scene, city scene]) and using one-way ANOVA (factors: subject, target, distractor). The results in top half of Table 1 show that surface photographs were generally easier, while splashes, including photographs of splashes, were harder. Illustrations as a group appeared as statistically insignificant. Analyzing RT to individual images, as shown in the bottom half of Table 1, shows that two images 1e and 1f stand out. They are both photographs and have distinctly expressed circular and caustic waves. Both have a colour close to blue. Surprisingly, while caustic patterns are rare, they are recognized easily, possibly because caustic patterns occur only in water and thus eliminate ambiguity. This result contradicts our expectation that illustrations of water would offer easier recognition than photographs. Subsequently we assume that easy images are photographs of water surfaces and that colour cues and visible patterns of waves are used to recognize water.

4.2 Pilot Experiment 2: Colour

Motivation. To test the hypothesis that colour eases recognition of water we compared 8 photographs in a variety of colours including blue, blue-green and also beige, the colour of water on a cloudy day. In the previous experiment all images identified as easy were largely blue. To test the intuition that humans expect water to be blue, we included four blue images in the set. We also included the green-blue image of caustics (Figure 1f) previously identified as easy. Because easy images might have identifiable patterns of waves, we included five images with visible ripples at different scales. To test whether colour takes precedence over other visual features half of the chosen distractors were blue, including two images of sky.

Subjects. The subjects were 5 males and 5 females, graduate students and faculty, all had normal or corrected to normal vision. Some did more than one session. There were 128 trials per session.

Result: Using exploratory data analysis we discovered significant

Image	Colour	RT diff. from mean (ms)	p
	beige	+12.9	0.0001
	blue	-4.116	0.0643
	blue-green	-8.779	p <0.0001
1i	blue-green	-8.498	0.031
1k	blue-green	-11.56	0.0029
1o	beige	+24.8	p <0.0001
Distractor	Colour	RT diff. from mean (ms)	p
	black-and-white	-5.277	0.0403
	blue	+4.826	0.0651

Table 2: *Blue-green is easy, beige is hard.*

results with one-way ANOVA (3 factors: subject, target colour [blue-green, blue, beige], distractor category [blue scene, sky, coloured scene, black-and-white scene]). Table 2 shows the significant factors, with saturated blue-green images significantly faster. This included the green-blue image of caustics and images of ripples shown on Figures (1i-1k). Blue distractors somewhat slowed down recognition without affecting error rate, but images of sky used as distractors increased error rates without slowing recognition, possibly due to sky being seen as a reflection in water. Some subjects, but not all, recognized beige water easily, which suggests variability in the human concept of water. The results of the two pilot experiments support the hypothesis that easily recognized images of water are blue photographs of water surfaces and have visually prominent patterns, either circles, caustics or waves.

4.3 Experiment 3: Spatial Frequencies

Motivation. Knowing that colour and surface geometry are important, we ask whether water recognition depends on local or global geometric features? Much contemporary computer graphics reproduces accurate small-scale surface features, assuming that fine detail is required for easy recognition. If so HSF contributes more than LSF. We desaturated 9 images of easy surfaces and 9 distractor images and filtered them approximately at the peak of spatial sensitivity [De Valois and De Valois 1988]: LSF <4.4 cpd and HSF >4.4 cpd so that each image was represented in three versions: full spatial frequency (FSF), LSF and HSF. All images had the same mean luminance. Test images were combined with distractors in matching versions.

Subjects.The subjects were 7 males and 2 females, graduate students and staff who were paid, all had normal or corrected to normal vision. The experimental procedure was approved by University of Waterloo Research Ethics Committee. There were 486 (9x9x2x3) trials per session.

Result. We used exploratory data analysis. One-way ANOVA (factors: subject, target frequency, distractor frequency) shows that FSF images were recognized faster than filtered images (Table 3, top) and LSF images were easier than HSF, suggesting that water is recognized by its global structure. However, one-way ANOVA (factors: subject, target, distractor) (Table 3, bottom) shows that some types of water surfaces (image 4 and 8) rely on LSF and FSF components equally and become harder to recognize when some of the frequencies are missing. For such images fine detail and global structure are equally important.

4.4 Experiment 4: Is there a colour of water?

Motivation. Is there an optimal colour for recognition of water? We saw that images most easily identified as water in the previous experiments are blue-green, but other easy images also fall within

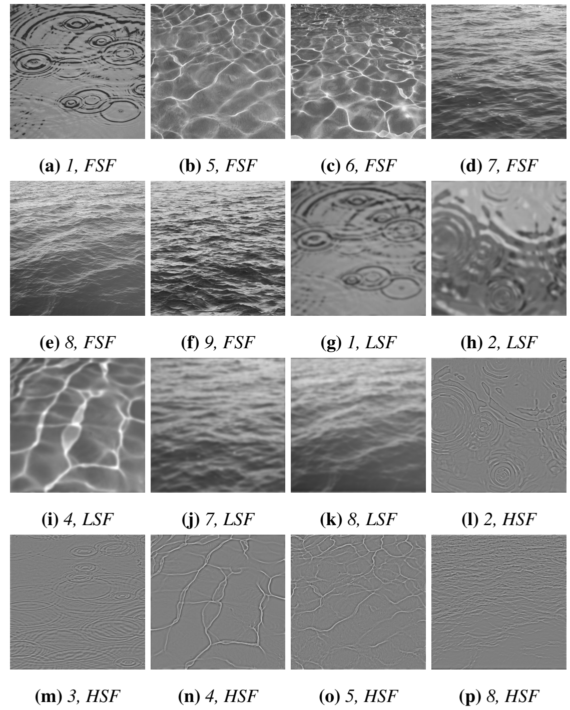

(a) *1, FSF* **(b)** *5, FSF* **(c)** *6, FSF* **(d)** *7, FSF*
(e) *8, FSF* **(f)** *9, FSF* **(g)** *1, LSF* **(h)** *2, LSF*
(i) *4, LSF* **(j)** *7, LSF* **(k)** *8, LSF* **(l)** *2, HSF*
(m) *3, HSF* **(n)** *4, HSF* **(o)** *5, HSF* **(p)** *8, HSF*

Figure 3: *FSF and LSF are easy, while HSF are hard to recognize.*

Image	Band	RT diff. from mean (ms)	p
	FSF	-9.627	0.0001
	HSF	+12.04	0.0001
1	FSF	-15.41	0.0004
5	FSF	-22.39	0.0001
6	FSF	-10.93	0.0099
7	FSF	-13.7	0.0021
8	FSF	-15.74	0.0003
9	FSF	-11.9	0.0063
1	LSF	-17.77	0.0001
4	LSF	+10.52	0.0178
8	LSF	+11.7	0.0089
2	HSF	+22.24	0.0001
3	HSF	+15.7	0.0008
4	HSF	+28.39	0.0001
8	HSF	+20.83	0.0001

Table 3: *Reaction times by frequency band and by image.*

a range of reddish blue and and blue hues. For example, image on Figure 1e is blue-red, 1f is blue-green, and 1h is blue. While luminance and saturation vary greatly across the easy images, hues vary little.

To find the optimal range of hue we selected 3 images of each surface type and 9 matching distractors and recoloured them to hues with xy near: blue(0.21, 0.23), blue-green (0.23, 0.31) and blue-red (0.25, 0.26) each group forming a clearly distinguishable cluster in xyz and in XYZ. We also included 3 grey images as controls, which we expected to be harder than the coloured images.

Subjects. The subjects were 7 males and 3 females, graduate students and staff who were paid, all had normal or corrected to normal vision. The experimental procedure was approved by University of Waterloo Research Ethics Committee. There were 288 (12x12x2) trials per session.

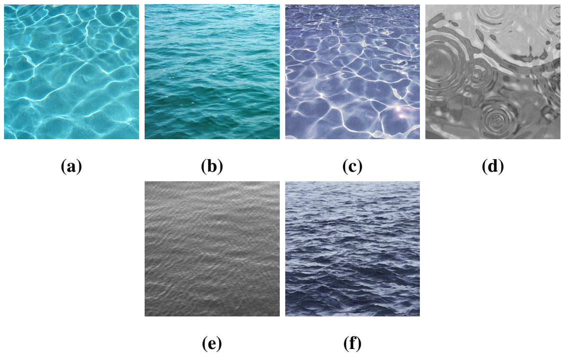

Figure 4: *Easy images: 4a - 4c; Hard images: 4d - 4f.*

Image	Colour	RT diff. from mean (ms)	p
	grey	+9.325	p <0.0001
	blue-green	-7.39	p <0.0001
4d	grey	+8.926	0.0076
4e	grey	+14.99	p <0.0001
4a	blue-green	-11.29	0.0003
4b	blue-green	-7.738	0.0160
4f	blue-red	+7.668	0.0201
4c	blue-red	-6.516	0.0450

Table 4: *RT for statistically significant images shown on Figure 4.*

Result. As before we used exploratory data analysis. One-way ANOVA (factors: subject, target colour, distractor colour) shows (Table 4, top) that most subjects were fastest responding to blue-green images of water. The green-blue colour matches absorption spectrum of water. Blue-red and blue fell in between, presumably they have admissible, but less water-like hues. Removing colour cues completely makes water harder to recognize. This may be because in natural environments water usually reflects light. Results of one-way ANOVA(factors: subject, target, distractor) are shown at the bottom of Table 4. In addition, one blue-green distractor caused significant increase in error rate, which may be owing to being seen as a reflection of trees in water.

5 Discussion and Future Work

We tested the generalization of Mills [1985] for water finding no artistic depictions as easy to recognize as photographs. Possibly the illustrations we used were ill-chosen and the photographic depictions of water may be more familiar. We are now seeking artistic depictions of water that satisfy the colour and geometry constraints we discovered by isolating good features in photographs. If the new images are easily recognized two interesting opportunities for further research beckon. By automatic search for images of water in online databases we hope to find non-water images resembling water owing to misbinding of visual features under brief exposure [Treisman and Gelade 1980]. To enable the automated search we must detect linear and circular edge structures. We should also find better non-photorealistic depictions of water.

This study limited itself to static images, but water surfaces usually move. Motion is surely an effective recognition cue and must be taken into account in order to inform interactive graphics. Our research is continuing looking for motion patterns that define water.

Lastly, the LSF nature of water surfaces suggests ”good enough” computations for water in the background. The next step will be creating such algorithms and testing them.

6 Conclusion

In our experiments subjects recognized water using colour and surface geometry. The colour providing easiest recognition is a blue-green colour. In addition, three surface structures are important: either circular wavelets expanding around a common point, or an irregular pattern of caustics, or a texture of ripples. The information defining these surface structures is concentrated in LSF <4.4 cpd. The results provide guidance for effective depiction of water in artistic rendering.

References

BAR, M., NETA, M., AND LINZ, H. 2006. Very first impressions. *Emotion 6*, 269.

DE VALOIS, R., AND DE VALOIS, K. 1988. *Spatial Vision*. Oxford: Oxford University Press.

DELORME, A., RICHARD, G., AND FABRE-THORPE, M. 2000. Ultra-rapid categorisation of natural scenes does not rely on colour cues: a study in monkeys and humans. *Vision research 40*, 2187–2200.

FIELD, D. J. 1987. Relations between the statistics of natural images and the response properties of cortical cells. *JOSA A 4*, 2379–2394.

HOFFMAN, D. D., AND SINGH, M. 1997. Salience of visual parts. *Cognition 63*, 29–78.

HOLCOMBE, A. O. 2009. Seeing slow and seeing fast: two limits on perception. *Trends in cognitive sciences 13*, 216–221.

KRYVEN, M., AND COWAN, W. 2013. Modelling perceptually efficient aquatic environments. In *Proceedings of the ACM Symposium on Applied Perception*, ACM, 131–131.

LIU, C., SHARAN, L., ADELSON, E. H., AND ROSENHOLTZ, R. 2010. Exploring features in a bayesian framework for material recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, IEEE, 239–246.

MILLS, M. I. 1985. Image synthesis: optical identity or pictorial communication? In *Graphics Interface*, ACM, 303–306.

PERINA, A., CRISTANI, M., AND MURINO, V. 2010. Learning natural scene categories by selective multi-scale feature extraction. *Image and Vision Computing 28*, 6, 927–939.

POPE, R. M., AND FRY, E. S. 1997. Absorption spectrum (380–700 nm) of pure water. ii. integrating cavity measurements. *Applied optics 36*, 8710–8723.

SELST, M. V., AND JOLICOEUR, P. 1994. A solution to the effect of sample size on outlier elimination. *The Quarterly Journal of Experimental Psychology Section A 47*, 3, 631–650.

THORPE, S., FIZE, D., MARLOT, C., ET AL. 1996. Speed of processing in the human visual system. *Nature 381*, 520–522.

TREISMAN, A. M., AND GELADE, G. 1980. A feature-integration theory of attention. *Cognitive psychology 12*, 1, 97–136.

VOGEL, J., AND SCHIELE, B. 2007. Semantic modeling of natural scenes for content-based image retrieval. *International Journal of Computer Vision 72*, 2, 133–157.