CS 8803 - The Uniform Illusion

Sahil Arora Haolun Hsu Kevin Shu Xiaofan Yuan

1 Stimulus and Response

The uniformity illusion is a class of optical illusions, which affects images which are 'close' to having a uniform pattern throughout the image. See www.uniformillusion.com for examples. Specifically, this class of images has the following characterization: in the center of the image, there is some 'uniform' pattern, such as a grid of rectangles or a collection of lines all facing the same direction, but on the boundary of the image, this pattern is altered slightly, (e.g. the rectangles in the grid are a slightly different shape, or the outer pattern is blurred).

When a person looks at this kind of image, they will initially perceive that the image consists of two distinct parts: the uniform pattern in the center and the altered pattern on the edges. As they stare at the image, they perceive the altered pattern on the edges fading, and the image is filled in with the uniform pattern.

2 Mathematical Model

When first presented with the image, people typically perceive two distinct parts: the nice uniform center, and the more chaotic boundary. Over time, their perception shifts to perceiving the illusion, which suggests that there is a computational process attempts to 'correct' this image. One way to model this phenomenon is to think of the brain as not just passively perceiving an image, but rather actively trying to find discrepancies in the image and its preconceptions about the world.

This idea is bolstered by the biological fact that peripheral vision in the human eye is significantly less sharp than the vision in the center. Given this fact, it makes sense that the brain would discount the sensory data that it obtains at the periphery and overemphasize the data that it receives at the center.

We can formalize this model using Bayes' rule. In some senses, this Bayesian model of the brain is a 'metamodel', in that different choices of priors and likelihoods can result in different models of perception.

We will suppose that there is some 'true' image, which is external to human perception, (in the case of the illusion, this true image is given by the true pattern that is presented by the researchers). What the human eye then perceives is similar to the true image, but will have some discrepancies due to the fact that human eye is not perfectly sharp. For example, the peripheral vision will be both blurrier and contain less color information than foveal vision.

We will let I be a random variable representing the true image, and we will let $\hat{I}$ be the image picked up by the eye when a person looks at image I. We will suppose that the random variable $\hat{I}$ only depends on I, and nothing else in the environment, so that we can describe what I is using a conditional probability distribution, say $\Pr(\hat{I}|I)$. We will refer to this distribution as the likelihood of $\hat{I}$ given I, and we will also suppose that in some way, the brain is able to compute this likelihood.

The other thing we will assume is that the brain is able to compute some **nonuniform** prior for how the image I is given. This may be because certain types of images are more common than others in daily perception, or it may simply be an unintentional effect of some other aspect of human perception.

Given this data, Bayes' rule dictates that

$$\Pr(I|\hat{I}) = \frac{\Pr(\hat{I}|I)\Pr(I)}{\Pr(\hat{I})}$$

What we observe is that the human brain is naturally biased towards producing images where I is 'simple', i.e. made of repeating patterns. So, when an image is presented which is not made of such repeating patterns, the brain might believe that the true image is in fact made of repeating patterns, and that there have simply been some errors in its reconstruction of the true image. If the brain wants to choose a maximum posterior estimator for the true image, then it might remove the nonrepeating parts of the pattern, and replace it with a simpler repeating image.

We will give two models which fall within this Bayesian metamodel. We will consider a simple 0-dimensional model of visual stimuli from a theoretical perspective, and we will also describe an error correcting autoencoder model for images that will experience a version of the uniformity illusion.

3 One Dimensional Ising Model

We will imagine that there are creatures which live in a 1 dimensional world, and which can only see in black and white, with $2n+1$ pixels of resolution, so that to them, an image can be expressed as a $2n+1$ dimensional vector of zeros and ones, $I = (I_{-n}, I_{-n+1} \dots, I_{n-1} I_n)$, where each $I_i \in \{0, 1\}$.

The creature does not have perfect vision, so there is some error probability for each pixel in its vision. Let the received image of the creature be $\hat{I} = (X_{-n}, \dots, X_n)$ and $\Pr(X_i = 1 | I_i = 1) = \Pr(X_i = 0 | I_i = 0) = p_i$.

For the purposes of modeling the uniformity illusion, we will suppose that there are only two distinct values for the p_i, say, for $|i| \leq k$, we have that $p_i = e^{-c}$, and for $|i| < k$, we have $p_i = e^{-b}$, where $e^{-c} > e^{-b}$.

We will consider an interesting model for what the prior distribution on images is, the so called Ising model distribution. Let $\sigma(x, y) = \begin{cases} 1 \text{ if } x = y \\ -1 \text{ if } x \neq y \end{cases}$. The probability that an image appears in the prior distribution is

$$\Pr(I) = k \exp\left(-\beta - \sum_{i=-n}^{n} \sigma(I_i, I_{i+1})\right)$$

where k is some constant independent of I, and $\beta > 0$. This distribution is very well studied in physics, and has historically been used for image correction in noisy models prior to the advent of neural networks [5].

In total, the total posterior distribution looks like

$$\Pr(I|\hat{I}) = k \exp\left(-\beta H(I) - \sum_{i=-k}^{k} c\sigma(X_i, I_i) - \sum_{|i|>k} b\sigma(X_i, I_i)\right)$$

Now that we have our model for how visual stimuli work, we will describe a simple version of the uniform illusion, where a strong pattern in the center will be extrapolated to the edges of the image.

Consider an image I so that for $|i| < k$, $I_i = 0$, and so that $I_i = 1$ for all of the remaining pixels. If we imagine this in 1 dimension, this is an image where the central pixels are all white, and the pixels on the sides are all black.

Theorem 1. *If $\beta + 2kc > 2(n-k)b$, then the maximum posterior image for $\hat{I} = I$ is all white.*

3.1 Posterior Sampling

It is well known that the Ising model can be sampled effectively using the Metropolis-Hastings algorithm, which is a Markov chain algorithm (see [4]). We simulated this process for the model described in this section, and confirmed that after not too many iterations, the sampled image converges to the completely white image. In particular, all that is needed to sample from this distribution is the ability to compute the prior and likelihood functions for this Ising model. One important aspect of the Ising model is the fact that the prior can be computed 'locally', where the only information needed to compute the change in potential when changing a pixel are the pixels immediately adjacent to it. Because the prior and posterior are very simple, it is plausible that the brain might be able to simulate this sort of algorithm.

4 Autoencoder Mechanism

We simulate the visual response to the blurriness uniformity illusion with an autoencoder, which is a popular model for image denoising. An autoencoder consists of three key features: an encoding function for compression, a decoding function for decompression, and a loss function that measures the distance between the amount of information loss between the compressed representation of your data and the decompressed representation. The encoder and decoder will be neural networks that are differentiable with respect to the loss function, so the parameters can be optimized via gradient descent.

The autoencoder is trained to take in an image, which has had some blur applied, and output the image without any blur. We train it by starting with a uniform pattern, applying a blur to it, and then feeding the blurred image into the autoencoder. Our loss function is a measurement of how different the output is from the starting image. When the loss on the validation dataset reaches a local optimum, the trained autoencoder emulates the behavior of filling in the blurred periphery with a uniform pattern across the image.

To better emulate the Ising model, we use a variational autoencoder, which learns a latent variable model for its input data. This constrains the encoding and decoding functions to learn the parameters of a probability distribution modeling the training data instead of any arbitrary function. First, an encoder network turns the input samples x into two parameters in a latent space, which we will note $\bar{z}$ and z_σ. Then, we randomly sample similar points z from the latent normal distribution that is assumed to generate the data, via $z = \bar{z} + e^{z_\sigma} * \epsilon$, where ϵ is a randomly sampled from the normal distribution. Finally, a decoder network maps these latent space points back to the original input data. The parameters of the model are trained via two loss functions: a reconstruction loss forcing the decoded samples to match the initial inputs like in the vanilla autoencoder, and the Kullback-Leibler (KL) divergence between the learned latent distribution and the prior distribution, acting as a regularization term.

4.1 Simulation Results

To quantify the effectiveness of our reconstruction, we use the structural similarity index measure (SSIM) which is commonly used to measure how similar two images are to each other (see [7]). We calculate SSIM across all the examples in training and testing and average this to get a global similarity measure. We found that predicted images from the autoencoder were more similar to the uniform image than they were to the image with a blurry border. This is in line with qualitative results and makes sense as the autoencoder was specifically trained to be as close to uniform as possible.

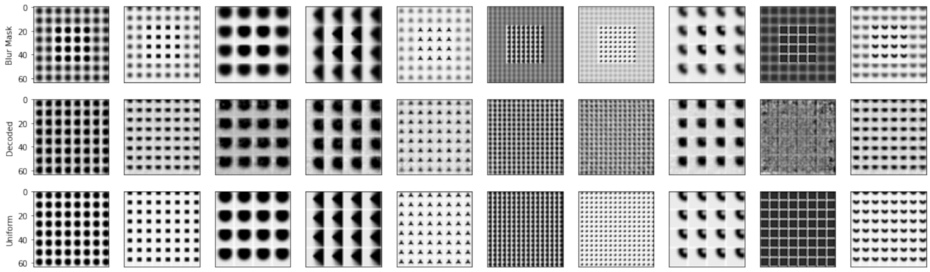

Figure 1: Patterns (Blur, Decoded, Uniform) from Variational Autoencoder

4.2 Similarity Thresholds

Another key aspect to emulate is that the uniformity illusion occurs less frequently and effectively when the central and peripheral stimuli become more different. So ideally our autoencoder would emulate this by reconstructing weaker blurred peripheries better than stronger blurs, while both would have better reconstruction than a uniform image with a noisy border. This is confirmed with SSIM results when evaluating decoded outputs from low blur, high blur, and noisy examples.

The different level of blurs are generated by Gaussian filter with different standard deviations (SD). In Table 1, we show that the higher the SD will result in lower SSIM compared with the original the uniform images. We select SD as 1 for further comparison. Table 3 indicates that adding a mask should reduce the difference influenced via blur. As for the SSIM between uniform and decoded patterns is about 91%, implying that the decoded patterns are more similar to the uniform patterns than the masked patterns.

Table 1: SSIM between Uniform Patterns and the Patterns with Different Level of Blur

Standard Deviation of Gaussian Filter			
0.5	1	2	10
SSIM = 98.52%	SSIM = 85.36%	SSIM = 62.84%	SSIM = 35.46%

Table 2: SSIM among Different Patterns from Variational Autoencoders

	Mask	**Blurred**
Uniform	99.01%	85.36%
Decoded	87.94%	86.19%

5 Biological Plausibility and Cognitive Experiment Results

Cognitive neuroscience research suggests that the early perceptual processing in human brains works in a Bayesian way. For example, Knill and Pouget [3] "described psychophysical evidence that shows human observers to behave in a variety of ways like optimal Bayesian observers".

Moreover, researchers argue that “brains are essentially prediction machines. They are bundles of cells that support perception and action by constantly attempting to match incoming sensory inputs with top-down expectations or predictions.” As in [1], Clark concludes that “this ‘hierarchical prediction machine’ approach offers the best clue yet to the shape of a unified science of mind and action”.

Humans’ vision perception is not uniform. Our vision in the fovea is accurate and detailed, while visual resolution and color sensitivity are limited in the periphery (Anderson, Mullen, & Hess, 1991; Westheimer, 1982). On the other hand, vision seems rich and detailed for most of the visual field to us. M. Otten, Y. Pinto, C. Paffen, A. Seth, and R. Kanai [6] suggest that “Perhaps people’s actual experience is rich and detailed because the brain supplements the details and richness when bottom-up input is poor.”

They designed a series of experiments on this “uniform illusion” to show that that detailed peripheral visual experience is partially based on a reconstruction of reality. Those experiments involved 20 participants, and a wide range of visual features. The results showed that “fixating on centrally presented stimuli can reliably create an illusory perception of uniformity in peripheral stimuli and that this effect can occur for stimuli of different shape, orientation, luminance, shade, motion, identity, and pattern.” More precisely, participants reported the uniformity illusion in 33% to 96% of displays, depending on the type of stimulus and the degree of difference between the central and peripheral patches.

5.1 Comparing our mechanism with the reality

Our mechanism now is only tested to be working with blurriness, while uniformity illusions appear for images with very different characteristics, ranging from a display filled with objects moving at different speeds to a uniformly gray display with a difference in luminance.

Our mechanism works well with the low blur case, while it does not work that well anymore with high blur or noise. This coincides with the results in [6], in which “ all experiments showed that when central and peripheral stimuli were more dissimilar, participants less often reported seeing uniformity, and if they did, time to onset of the illusion increased.”

References

[1] A. Clark, Whatever next? Predictive brains, situated agents, and the future of cognitive science, *Behavioral & Brain Sciences*, 36, (2013) 181–204.

[2] Friston, Karl. ”The history of the future of the Bayesian brain.” NeuroImage 62.2 (2012): 1230-1233.

[3] D. Knill and A. Pouget, The Bayesian brain: the role of uncertainty in neural coding and computation, *TRENDS in Neurosciences*, 27(12), (2004), 712–719.

[4] Kotze, Jacques. ”Introduction to Monte Carlo methods for an Ising Model of a Ferromagnet.” arXiv preprint arXiv:0803.0217 (2008).

[5] Murphy, Kevin P. Machine learning: a probabilistic perspective. MIT press, 2012.

[6] M. Otten, Y. Pinto, C. Paffen, A. Seth, and R. Kanai, The uniformity illusion: Central stimuli can determine peripheral perception, *Psychological Science*, 28(1), (2017) 56–68.

[7] Wang, Z., Bovik, A. C., Sheikh, H. R., Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing.*

6 Appendix

Table 3: Contributions

Sahil	Image generation and autoencoder simulation
Haolun	Variational autoencoder simulation
Kevin	Mathematical modeling of uniformity illusion
Xiaofan	Biological plausibility of simulated model

Image generation and autoencoder simulations are available at github and Google Colab