

Optical illusions and why neural networks can't seem to figure them out

[Optical illusions] are interesting because they provide insight into the nature of the visual system and perception. So ways of finding new illusions that explore these limits would be hugely useful.

...

[However,] Current machine-learning systems cannot generate their own optical illusions—at least not yet. Why not?

...

“The number of static optical illusion images is in the low thousands, and the number of unique kinds of illusions is certainly very low, perhaps only a few dozen,” [researchers Robert Williams and Roman Yampolskiy] say.

That represents a challenge for current machine-learning systems. “Creating a model capable of learning from such a small and limited dataset would represent a huge leap in generative models and understanding of human vision,” they say.

So Williams and Yampolskiy compiled a database of over 6,000 images of optical illusions and then trained a neural network to recognize them. Then they built a generative adversarial network to create optical illusions for itself.

The results were disappointing. “Nothing of value was created after 7 hours of training on an Nvidia Tesla K80,” say the researchers, who have made their database available for others to use.

Nevertheless, this is an interesting result. “The only optical illusions known to humans have been created by evolution (for instance, eye patterns in butterfly wings) or by human artists,” they point out.

In both cases, humans play a crucial role by providing valuable feedback—humans can see the illusion.

Read full, original post: [Neural networks don't understand what optical illusions are](#)