

From: "jeffrey E." <jeevacation@gmail.com>

To: [REDACTED] <[REDACTED]>

Subject: Re: [Dewayne-Net] Google teaches "AIs" to invent their own crypto and avoid eavesdropping

Date: Sat, 29 Oct 2016 16:00:41 +0000

Paris

On Saturday, 29 October 2016, [REDACTED] <[REDACTED]> wrote:

Good. Between Raleigh & Seattle with an occasional side trip for an event here or there.

May need a jaw bone graft, but hoping to avoid.

Hating this election.

You?

Co-authored with iPhone auto-correct

On Oct 29, 2016, at 8:45 AM, jeffrey E. <jeevacation@gmail.com> wrote:

How are you

On Saturday, 29 October 2016, [REDACTED] <[REDACTED]> wrote:

Interesting

Co-authored with iPhone auto-correct

Begin forwarded message:

From: Hendricks Dewayne <[REDACTED]>

Date: October 29, 2016 at 1:50:41 AM PDT

To: Multiple recipients of Dewayne-Net <[REDACTED]>

Subject: [Dewayne-Net] Google teaches "AIs" to invent their own crypto and avoid eavesdropping

Reply-To: [REDACTED]

Google teaches "AIs" to invent their own crypto and avoid eavesdropping

Neural networks seem good at devising crypto methods; less good at codebreaking.

By SEBASTIAN ANTHONY (UK)

Oct 28 2016

<<http://arstechnica.com/information-technology/2016/10/google-ai-neural-network-cryptography/>>

Google Brain has created two artificial intelligences that evolved their own cryptographic algorithm to protect their messages from a third AI, which was trying to evolve its own method to crack the AI-generated crypto. The study was a success: the first two AIs learnt how to communicate securely from scratch.

The Google Brain team (which is based out in Mountain View and is separate from Deep Mind in London) started with three fairly vanilla neural networks called Alice, Bob, and Eve. Each neural network

was given a very specific goal: Alice had to send a secure message to Bob; Bob had to try and decrypt the message; and Eve had to try and eavesdrop on the message and try to decrypt it. Alice and Bob have one advantage over Eve: they start with a shared secret key (i.e. this is symmetric encryption).

Importantly, the AIs were not told how to encrypt stuff, or what crypto techniques to use: they were just given a loss function (a failure condition), and then they got on with it. In Eve's case, the loss function was very simple: the distance, measured in correct and incorrect bits, between Alice's original input plaintext and its guess. For Alice and Bob the loss function was a bit more complex: if Bob's guess (again measured in bits) was too far from the original input plaintext, it was a loss; for Alice, if Eve's guesses are better than random guessing, it's a loss. And thus an adversarial generative network (GAN) was created.

Alice, Bob, and Eve all shared the same "mix and transform" neural network architecture, but they were initialised independently and had no connection other than Alice and Bob's shared key. For Alice the key and plaintext are input into the first layer of the neural network; for Bob the key and the ciphertext were input; and for Eve, she got just the ciphertext. The first layer is fully-connected, so the text and key can mix about. Following the first layer there are a number of convolutional layers, which learn to apply a function to the bits that were handed to it by the previous layer. They don't know what that function might be; they just learn as they go along. For Alice, the final layer spits out some ciphertext; Bob and Eve output what they hope is the plaintext.

The results were... a mixed bag. Some runs were a complete flop, with Bob never able to reconstruct Alice's messages. Most of the time, Alice and Bob did manage to evolve a system where they could communicate with very few errors. In some tests, Eve showed an improvement over random guessing, but Alice and Bob then usually responded by improving their cryptography technique until Eve had no chance (see graph).

[snip]

Dewayne-Net RSS Feed: <<http://dewaynenet.wordpress.com/feed/>>

--

please note

The information contained in this communication is confidential, may be attorney-client privileged, may constitute inside information, and is intended only for the use of the addressee. It is the property of JEE

Unauthorized use, disclosure or copying of this communication or any part thereof is strictly prohibited and may be unlawful. If you have received this communication in error, please notify us immediately by return e-mail or by e-mail to jeevacation@gmail.com, and destroy this communication and all copies thereof, including all attachments. copyright -all rights reserved

--

please note

The information contained in this communication is confidential, may be attorney-client privileged, may constitute inside information, and is intended only for

the use of the addressee. It is the property of
JEE

Unauthorized use, disclosure or copying of this
communication or any part thereof is strictly prohibited
and may be unlawful. If you have received this
communication in error, please notify us immediately by
return e-mail or by e-mail to jeevacation@gmail.com, and
destroy this communication and all copies thereof,
including all attachments. copyright -all rights reserved