



DEEPER

Eleni Xynogala | XYN17515132 | 8.1.2018

Unit 1.1 | Theories & Technologies of Interaction Design | MA Interaction Design
Communication | London College of Communication | University of the Arts London

Contents

Introduction

Seeing the Invisible

Primary and Secondary Research

Background

Aestheticising the Invisible

Intervention in the Invisible

References

Introduction

“Deeper” is the project that created for the Theories & Technologies of Interaction Design brief. This project consisted of three parts. In the first part of the brief, I had to investigate an invisible system that we interact with through technology. Through the second part of the brief, the goal was to design a way to visualize the invisible system. Finally, in the third part, I had to develop a system that intervenes in the invisible system. The system that I chose to research is the artificial neural networks and specifically the art that is created by them. The final outcome of the project is a sculpture which contains a memory stick with a video.

Seeing the Invisible

To start this project, I had to choose a subject which I am interested in. So, I start research invisible systems and I inspired from a TEDx Talk from Mike Tyka, who uses artificial neural networks to create art. My goal was to answer the questions “what is this network”, “how does it work” and most importantly “why does this exist”.

Primary and Secondary Research

To do so, I started research using different methods so as to collect the knowledge that I needed to understand, analyze, visualize and finally intervene in the chosen subject. Here are some examples of my research methods that refer to all the journey from the initial research to the final part of the project:

Field research

I explored the machine learning field and especially the neural networks and the way they work. Specifically, I learned about neural networks’ architecture, image classification, data visualization and Google’s Deep Dream algorithm.

Online observation

Searching on the internet was the easiest way to explore how neural networks work and the audience's and experts' opinion about them. Some articles that influenced me was:

"Inceptionism: Going Deeper into Neural Networks" by Alexander Mordvintsev and Mike Tyka in Google Research Blog. In this post, they describe how neural networks are able to carry out difficult classification tasks and check what the network has learned during training, but also makes us wonder whether neural networks could become a tool for artists.

"Generative Models" by Andrej Karpathy, Pieter Abbeel, Greg Brockman, Peter Chen, Vicki Cheung, Rocky Duan, Ian Goodfellow, Durk Kingma, Jonathan Ho, Rein Houthoof, Tim Salimans, John Schulman, Ilya Sutskever and Wojciech Zaremba. This post describes four projects that share a common theme of enhancing or using generative models, a branch of unsupervised learning techniques in machine learning.

Generally, YouTube videos that explain the way that neural networks work were educated and helped me to clear my mind about this complicated issue.

Different blog posts which discuss the debatable side of this subject were also interested to read and influenced my progress in the progress. Some of them are listed below:

"Is Google's Deep Dream art?" By Marina Galperina.

"#Deepdream is blowing my mind" by Memo Akten.

"Sensual Machines" by Samim.

"Machine Bias" by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner

"Inside deep dreams: How google made its computers go crazy" by Steven Levy.

Reading articles and essays

The most interesting and useful papers that I rode about my topic was:

"Art by algorithm" by Ed Finn.

This essay discusses the fact that computation is changing aesthetics and debates the role of artists and appreciators of the arts in this trend.

"Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images" by Anh Nguyen, Jason Yosinski and Jeff Clune.

A really interesting article, which proves how neural networks can easily be fooled and imply a way that human can intervene in these systems.

"How the Machine 'Thinks': Understanding opacity in machine learning algorithms, Big Data & Society" by Jenna Burrell.

This article considers the issue of opacity as a problem in mechanisms of classification. It was useful especially for the third part of my project, that I wanted to criticize the network.

Ask experts in the field

To understand better the concept, I came in touch with two people that are close to the field in technological terms.

One of them is Will Gallia, who is currently in the prototyping lab of the Interaction Design Communication course and assist students with their projects. Will, except of the technician help that provided me, expressed also his valuable opinion about Deep Dream algorithm, as he was a practitioner in neural networks' visualization in the past.

Another programmer with experience in machine learning that helped me a lot in understanding the algorithm was Evangelos Michelioudakis, a PhD candidate in the Department of Informatics and Telecommunications of the National and Kapodistrian University of Athens and researcher at the Institute of Informatics and Telecommunications of the National Center for Scientific Research “Demokritos” in Athens.

Background

Neural networks are a form of computing which has its beginnings in the 40s when people started thinking about the brain as a computer and thinking about how it might work. One of those people was Alan Turing, who proposed a machine, which would be at first unorganized and then through training it would organize itself. People have built these machines in the 20th century and the last years have been an incredible revival of neural networks in the field of machine learning.

Architecture

Neural networks are composed of simple units. Each unit usually does a simple calculation like an addition. It takes input from many other neurons and process that data that comes in and send them to other neurons. All these are connected in a kind of network. The architecture, that is often used, is in layers where each layer only connects to the next layer. The connections that connect the neurons have weights associated with them and those weights are random at first (unorganized machine-it doesn't work).

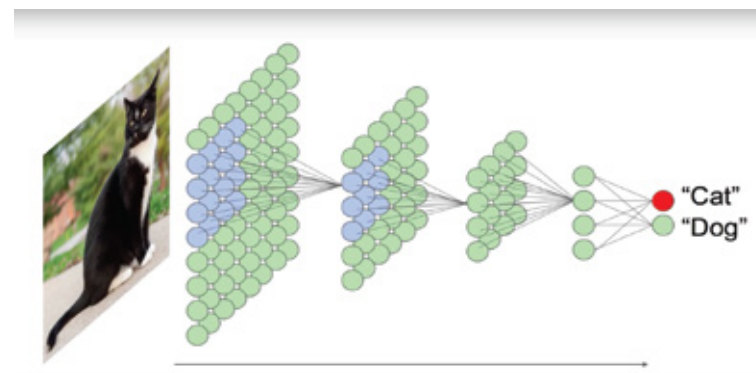


Image classification

In this characteristic example of image classification, we need to train the network, showing thousands and thousands of pictures of cats and dogs in this case, in order to distinguish objects. The layers that are close to the input data have neurons that become reactive to simple features (edges, corners) and then as you move through the layers the neurons respond to complex features (eyes, cat).

Visualisation

It is also interesting to run this program backwards (Generative neural networks). If you have a system that knows everything about what a cat is like, it should be able to produce new pictures that look like cats or dogs.

This started from experiments from Alexander Mordvintsev at Google. He was interested to create visualizations of the internal knowledge stored in the neural network and he came up with an algorithm. In this algorithm, you show a network that's trained on whatever (like a

picture), you run the network forward and then you say okay whatever you saw, we are going to adjust the pixels in the image towards that interpretation just a little bit. If you do this repeatedly you get these kinds of images.



How We End Up At The End Of Life—Mike Tyka

Art

Mike Tyka, who is both artist and computer scientist at Google, used it to make art. He started with random noise and then he kept zooming into the picture. A fractal world appeared, where the network keeps over interpreting the image. As it seems in the video, it can last forever essentially.

Ethical issues

So, to continue with my research and the next parts of the projects I set the following questions:

Is it real art?

Can a machine outdo human?

Who is the artist? Is the creator of the algorithm? Is the person who inputs the image in the system? Is it the computer? Or is it a combination of all three?

Is the machine biased against blacks?

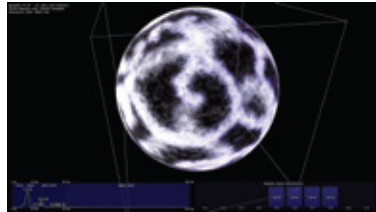
Aestheticising the Invisible

Through this part of the brief, the goal was to demonstrate the operation and structure of the neural networks. I experimented with different approaches, through many iterations and tests until to find my final visualization. I used sketches to clarify my ideas and I built two prototypes before my final outcome for this part of the project.

Inspiration




Neural Network 3D Simulation - Denis Dmitriev



Spiking Neural Network Visualisation with SpikeFun - Ivan Dimkovic



Uncertain Journey - Chiharu Shiota



Amazon Futures - La Loma & Tactical Tech - The Glass Room exhibit



Shades - Robert Rauschenberg

Ideas in Sketches



Light Installation / 3D simulation - Rejected

These ideas rejected because I didn't feel comfortable with some technologies like Arduino or 3ds Max.



Installation in a room - Rejected

After testing this idea, I realized that it is difficult to succeed it on time and it is not necessary to present it on such a large scale.



Artwork/Synthesis - Approved

So, I reconstructed my idea on a smaller scale and started a new iteration.

Prototype



I created a prototype using simple elements that I had at home, so as to see how my idea can be implemented. Testing my idea, I came up with corrections that would improve my visualization and I took decisions for the materials and visuals that I used.




Inside the box, I created a web using black thread, that depicts the machine's neurons. It is a complex network, therefore there are numerous connections that you cannot really count. I used patterns engraved in transparent acrylic using the laser-cutter of the 3D workshop.

These symbolize the motives that the network is trained to recognize, like this owl.


Final Visualization



I created a sculpture, like an open black box, which represents my machine. I used MDF, which technicians helped me to cut in the required dimensions in the 3D workshop. Then I painted with the black colour which shows that it is dark and hidden. I stacked the pieces to create the shape that I wanted.



Trying to make clear the connection between the human and artificial neural networks I curved the outside of the box with sketches of networks look like our brains' neural networks.



In front of the machine, I placed another acrylic with the photo of a black woman, which represents the input image that the system has never seen before.

After the network sees the input image, it produces new images that look like with the patterns that it is trained. So early layers care about simple features, like corners and edges. But as we move through the network the slightest features turn into eyes, birds, monkeys etc. To show the final visualization that the system produces, I input the image of the black image in the Deep Dream algorithm and I received the outcome, which I edited in photoshop, so as to take the different layers. All these images are printed on transparent paper and when are placed parallel one in front of the other the final outcome is presented.




The use of mirror has a double meaning. First of all, it is used to make threads look denser, depicting the complexity of the network. At the same time, they are like the eyes of the network that see the patterns, the input image, but also me, you and everyone that enters the system.

The use of the photo of the African-American woman is not irrelevant. These systems have criticized for the bias against black people. And here is obvious that the black woman is transformed into a gorilla.

Intervention in the Invisible

In this final part of the project, I had to find a way to intervene in the neural networks. This part of the brief was the hardest, as I had to choose one aspect of my subject that I wanted to focus more and help people learn new information, interacting with it.

Inspiration



The Kitty AI - Pinar Yoldas



Material Speculation: ISIS - Morehshin Allahyari



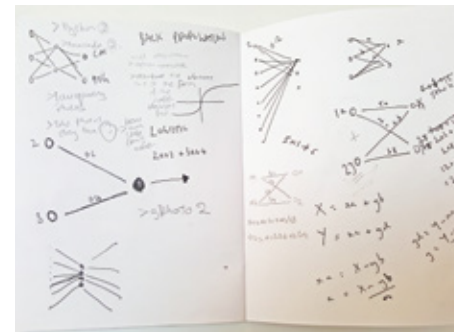
How Not to be Seen: A Fucking Didactic Educational .MOV File - Hito Steyerl

Ideas in Sketches




Installation - Rejected

This installation was based on the idea of taking real-time photos of people and edit them immediately with Deep Dream algorithm. It was rejected because it is not a real intervention and it doesn't have so impact on people's perception for the system.



Processing Representation - Rejected

This idea implemented, but the result wasn't compatible with what a real network would do. But a fake representation of the way it works. At the same time, it was like a visualization of the network's operation and not an intervention.




Video - Approved

Combining the knowledge of the coding with the information that I gained through research, I decided that a video that debate and criticize the system is what I want to create.


Final Outcome – Video

I created a video which I called Deeper. The name came in my mind because of all the names that are given to the algorithms related to the neural network, like Deep Dream, Deep Blue and DeepMind, but also because it implies that we explore the system deeper. The goal of this video is to inform people for some useful implementations of neural networks, like face and voice recognition, and at the same time to intervene in the system, showing the negative impact that they have in people's life, such as creating prejudice against black people, or creating fake porn.



The video created in Adobe Premiere Pro and saved on a memory stick which placed in the artwork which demonstrates the visualization of the neural network system of the previous part of the project. This symbolizes also an intervention in the system and its purpose is to raise audience's interest to discover what this artwork is, watching the video.

In each chapter, there are two sides. On the one hand, there are some videos and photos, found online, which show what neural networks can do.



On the other hand, there is an intervention which is visualized as repeatedly changing images, through which the system is trained.

It has the voice from a text to speech service, so as to sound less natural and to create a dialogue between the machine and its implication.

The video has the following structure:

Introduction

In the introduction, I show photos of the machine that I visualized in the second part of the brief and I set the following questions to the audience: "Does the machine know things?" and "Does it create art or racism?".

First chapter – The machine can think



All chapters' titles are created with a processing sketch which seems like networks.

The first video is a combination of a video, which shows the robot Sophia talking, and a voice from another video which talks about computers to outsmart humans by 2029.

Then, there are is a video about Deep Blue, a chess-playing computer developed by IBM, which defeated the world champion G. Kasparov in 1997. But also, another video which demonstrates Google's AI AlphaGo which is beating humanity at its own games.



The changes between the specific videos, which are from news broadcasts, are with an old TV changing channel noise effect created in Adobe After Effects.

The machine's voice claims that "Algorithms are simple mathematical formulas that nobody understands. What I cannot create I do not understand". In this way, I try to start a conversation with the audience.

Second chapter – The machine can create

In this chapter, the machine starts saying “I see animals. I dream animals. What do you dream?”

Then I show Mike Tyka’s video of zooming in a noise image, which is created from Deep Dream algorithm. And a part from his TEDx talk about the art that the neural networks create.

Third chapter – The machine can intervene

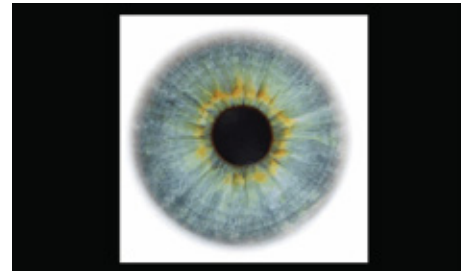
This chapter shows some videos edited with the use of Deep Dream algorithm, which creates a psychedelic effect in the videos. In this way, videos’ creators intervene in a movie like Fear and Loathing in Las Vegas, which is about LSD, so as to make a statement. In the same way, the other video is showing Donald’s Trump speech and the creator has transformed him to look like a monster.

In this case, the machine asks “Who is the artist? Is the engineer? Is the person who inputs the image? Is it me?”

Fourth chapter – The machine can NOT

The title implies that the machine cannot do everything, like doubting what said before. In this chapter, there are videos and pictures which show that these machines are biased against black criminals, or racists in face recognition, as they seem to have problems in recognizing black people or tag them as gorillas. Finally, I demonstrate a clip, showing Gal Gadot’s face on a porn star’s body, which created with a machine learning algorithm.

End



The eyes symbolize that machine observe as, as we train it with pictures that it sees.

The ending clip is created in Adobe After Effects and the two classical music pieces that are used in the video are created from Recurrent Neural Networks.



Video

Watch the video here:

<https://www.youtube.com/watch?v=-j9y7RI12-k>

For more information about the project, there are references below and all the process is documented in Medium blog in the following links:

<https://medium.com/@elen.xynogala/the-art-of-neural-networks-38483be2bb61>

<https://medium.com/@elen.xynogala/aestheticising-the-invisible-89d79f10ce1c>

<https://medium.com/@elen.xynogala/intervention-in-the-invisible-183a51109ac2>

References

- Akyen, M. (2015). *Deepdream is blowing my mind*. Medium. <https://medium.com/@memoakten/deepdream-is-blowing-my-mind-6a2c8669c698> (accessed 4 December 2017).
- Angwin, J., Larson, J., Mattu, S. and Kirchner, L. (2016). *Machine Bias*. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (accessed 4 December 2017).
- Auerbach, D. (2015). *Do androids dream of electric bananas?* Slate. http://www.slate.com/articles/technology/bitwise/2015/07/google_deepdream_it_s_dazzling_creepy_and_tells_us_a_lot_about_the_future.html (accessed 4 December 2017).
- Burrell, J. (2016). *How the Machine 'Thinks' Understanding opacity in machine learning algorithms*, Big Data & Society
- Finn, E. (2017). *Art by algorithm*. <https://aeon.co/essays/how-algorithms-are-transforming-artistic-creativity> (accessed 4 December 2017).
- Galperina, M. (2016). *Is Google's Deep Dream art?* Hopes&Fears. <http://www.hopesandfears.com/hopes/culture/is-this-art/215039-deep-dreamgoogle-art> (accessed 4 December 2017).
- Karpathy, A., Abbeel, P., Brockman, G., Chen, P., Cheung, V., Duan, R., Goodfellow, I., Kingma, D., Ho, J., Houthoof, R., Salimans, T., Schulman, T., Sutskever, I. and Zaremba, W. (2016). *Generative Models*. OpenAI. <https://blog.openai.com/generative-models/> (accessed 4 December 2017).
- Levy, S. (2015). *Inside deep dreams: How google made its computers go crazy*. Wired. <https://www.wired.com/2015/12/inside-deep-dreams-howgoogle-made-its-computers-go-crazy/#.i806t2o2x> (accessed 4 December 2017).
- Mordvintsev, A. and Tyka, M. (2015). *Inceptionism: Going Deeper into Neural Networks*. <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html> (accessed 4 December 2017).
- Nguyen, A., Yosinski, J. and Clune, J. (2015). *Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images*. <https://arxiv.org/pdf/1412.1897.pdf> (accessed 4 December 2017).
- Saffari, A. (2016). *An introduction to Deep Learning for Generative Models*. <http://www.ymer.org/amir/2016/11/21/an-introduction-to-deep-learningfor-generative-models/> (accessed 4 December 2017).
- Samim. (2015). *Sensual Machines*. Medium. <https://medium.com/@samim/sensual-machines-82858b32a4e5> (accessed 4 December 2017).
- Tyka, M. (2015). *The art of neural networks*. TEDx Talks. <https://www.youtube.com/watch?v=0qVOUD76JOg> (accessed 4 December 2017).