

Simen Selseng

Generating natural looking images with Generative Adversarial Networks

Specialization project, fall 2016

Data and Artificial Intelligence Group
Department of Computer and Information Science
Faculty of Information Technology, Mathematics and Electrical Engineering
Norwegian University of Science and Technology



Abstract

This paper provides a template for writing a Master's Thesis (parts of it can also be used when writing a fall project report). The template does not form a compulsory style that you are obliged to use, but rather provides a common starting point for all students. For a given project tuning of the template may still be required. Such tuning might involve moving a chapter to a section or vice versa due to the nature of the project.

The abstract is your sales pitch which encourages people to read your work, but unlike sales it should be realistic with respect to the contributions of the work. It should include:

- the field of research
- a brief motivation for the work
- what the research topic is and
- the research approach(es) applied.
- contributions

The abstract length should be roughly half a page of text — without lists, tables or figures.

Sammendrag

Husk at hvis du er en norsk student og skriver masteren din på engelsk, så *må* du lage et sammendrag på norsk.

(If you are a non-Norwegian student, it is not obligatory to include an abstract in Norwegian.)

Preface

The preface includes the facts: what type of project, where it is conducted, who supervised, and any acknowledgements you wish to give.

This Master Thesis template is based on one which was created for the 2016 “Experts in Team” course on Computational Creativity (TDT4853) at the Norwegian University of Science and Technology (NTNU), which in turn was heavily based on the AI Master Thesis template created by Anders Kofod-Petersen — and most of the explaining text comes from Anders’ original template.

Simen Selseng
Trondheim, 5th December 2016

Contents

1	Introduction	1
1.1	Background and Motivation	1
1.2	Goals and Research Questions	2
1.3	Research Method	2
1.4	Contributions	2
1.5	Thesis Structure	3
2	Background Theory	5
3	Related Work	9
3.1	Generative Adversarial Networks	9
3.2	State-of-the-art	10
3.2.1	Batch normalization	10
3.2.2	The All Convolutional net	10
3.2.3	Sparse connectivity	10
3.2.4	Feature matching	11
3.2.5	Minibatch Discrimination	11
3.3	Related work	11
3.3.1	Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks	11
3.3.2	Image Completion with Deep Learning in TensorFlow	11
3.3.3	Improved Techniques for Training GANs	12
3.3.4	Learning to Protect Communications with Adversarial Neural Cryp- tography	12
4	Architecture	13
5	Experiments and Results	15
5.1	Experimental Plan	15
5.2	Experimental Setup	15
5.3	Experimental Results	15
6	Evaluation and Conclusion	17
6.1	Evaluation	17
6.2	Discussion	17
6.3	Contributions	17
6.4	Future Work	17

Contents

Appendices	19
Bibliography	19

List of Figures

2.1	Boxes and arrows are nice	6
2.2	A missing figure	7

List of Tables

2.1	Example Table	7
-----	-------------------------	---

1 Introduction

All chapters should begin with an introduction before any sections. Further, each section should start with an introduction before its subsections begin. Chapters with just one section — or sections with just one sub-section — should be avoided. Think carefully about chapter and section titles as each title stands alone in the table of contents (without associated text) and should convey the meaning of the contents of the chapter or section.

In all chapters and sections it is important to write clearly and concisely. Avoid repetitions and if needed refer back to the original discussion or presentation. Each new section, subsection or paragraph should provide the reader with new information and be written in your own words. Avoid direct quotes. If you use direct quotes, unless the quote itself is very significant, you are conveying to the reader that you are unable to express this discussion or fact yourself. Such direct quotes also break the flow of the language (yours to someone else's).

1.1 Background and Motivation

Having a template to work from provides a starting point. However, for a given project, a slight variation in the template may be required due to the nature of the given project. Further, the order in which the various chapters and sections will be written will also vary from project to project, but will seldom start at the abstract and sequentially follow the chapters of the report. One critical reason for this is that you need to start writing as early as possible and that you will begin to write up where you are currently focusing. However, do not leave the abstract until the end. The abstract is the first thing anyone reads of an article or thesis — after the title; and thus it is important that it is very well written. Abstracts are hard to write, so create revisions throughout the course of your project.

This introduction to background and motivation should state where your project is situated in the field and what the key driving forces motivating this research are. However, keep this section brief, as it is still part of the introduction. The motivation will be further extended in Chapter 3, presenting your complete state-of-the-art.

Note that this template uses italics to highlight where latin wording is inserted to represent text and the text of the template that we wish to draw your attention to. The italics themselves are not an indication that such sections should use italics.

1 Introduction

1.2 Goals and Research Questions

A research project needs to have one or several question(s) that should be answered. It is desirable to formulate such questions as early as possible as the formation of such questions provide both an important driving force for the project and provide clarity as to the goals sought. However, expect to refine the questions and thus the final path of the project as work progresses. Any refinements should be conducted with care, so as to avoid that the original aims and previous work are not lost. It is always good to have one (or max two) key questions and perhaps some sub questions.

Goal *Lorem ipsum dolor sit amet, consectetur adipiscing elit.*

Your goal/objective should be described in a single sentence. In the text under you can expand on this sentence to clarify what is meant by the short goal description. The goal of your work is what you are trying to achieve. This can either be the goal of your actual project or can be a broader goal that you have taken steps towards achieving. Such steps should be expressed in the research questions. Note that the goal is seldom to build a system. A system is built to to enable experiments to be conducted. The research question/goal would be the goal that the system is implemented to meet.

Research question 1 *Lorem ipsum dolor sit amet, consectetur adipiscing elit.*

Each research question provides a sub-goal and these should be precise and clearly stated enabling the reader to match your results to the original goals. They will also form the driving force for the experimental plan.

Research question 2 *Lorem ipsum dolor sit amet, consectetur adipiscing elit.*

1.3 Research Method

What methodology will you apply to address the goals: theoretic/analytic, model/abstraction or design/experiment? This section will describe the research methodology applied and the reason for this choice of research methodology.

1.4 Contributions

The main description of the contributions will come in Section 6.3 after the results are presented. This section just provides a brief summary of the main contributions of the work. This section can also be left out, leaving discussion to Section 6.3.

The format of this section will generally be as follows: *Donec non turpis nec neque egestas faucibus nec id neque. Etiam consectetur, odio vitae gravida tempus, diam velit sagittis turpis, a molestie ligula tellus at nunc. Nam convallis consequat vestibulum. Proin dolor neque, dapibus a pellentesque a, commodo a nibh.*

1. *Lorem ipsum dolor sit amet, consectetur adipiscing elit.*
2. *Lorem ipsum dolor sit amet, consectetur adipiscing elit.*
3. *Lorem ipsum dolor sit amet, consectetur adipiscing elit.*

1.5 Thesis Structure

This section provides the reader with an overview of what is coming in the next chapters. You want to say more than what is explicit in the chapter name, if possible, but still keep the description short and to the point.

2 Background Theory

The background theory depth and breadth depends on the depth needed to understand your project in the different disciplines that your project crosses. It is not a place to just write about everything you know that is vaguely connected to your project. The theory is here to help the reader that does not know the theoretical basis of your work so that he/she can gain sufficient understanding to understand your contributions. In particular, the theory section provides an opportunity to introduce terminology that can later be used without disturbing the text with a definition. In some cases it will be more appropriate to have a separate section for different theories. However, be careful so that you don't end up with too short sections. Subsections may also be used to separate different background theories.

When introducing techniques or results, always reference the source. Be careful to reference the original contributor of a technique and not just someone who happens to use the technique. For relevant results to your work, you would want to look particularly at newer results so that you have referenced the most up-to-date work in your area. If you don't have the source handy when writing, mark in the text that a reference is needed and add it later.

add reference

Web pages are not reliable sources — they might be there one day and removed the next; and thus should be avoided, if possible. A verbal discussion is not a source and should not be referenced or described in the text.

The bulk of citations in the report will appear in Chapter 3. However, you will often need to introduce some terminology and key citations already in this chapter.

You can cite a paper in the following manner (and several other versions, see the natbib package documentation):

- when referring to authors:
? stated something rather nice.
- to cite indirectly:
Papers should be written nicely (?)
or
In ?, a less detailed template was presented.
- To just cite the authors:
? wrote a nice paper.
- Or just the year: ?.
- You can even cite specific pages: ?, p. 3.

2 Background Theory




Figure 2.1: Boxes and arrows are nice (adapted from ?)

You should obviously always cite your supervisor’s work (?), even if it is completely irrelevant (?) or very old (?). Digging up an even older book can also appear impressive (?).

Introducing figures

Remember that when you borrow figures you should always credit the original author — such as Figure 2.1 (adapted from ?). Also do not just put the figure in and leave it to the reader to try to understand what the figure is. The figure should be included to convey a message and you need to help the reader to understand the message intended by explaining the figure in the text.

It is good practice to add a note about a missing figure in the text, such as the completely amazing stuff that will appear in Figure 2.2.

Introducing tables in the report

As you can see from Table 2.1, tables are nice. However, again, you need to discuss the contents of the table in the text. You do not need to describe every entry, but draw the reader’s attention to what is important in the table.



Here we will add an amazing figure explaining it all

Figure 2.2: A missing figure

Table 2.1: Example Table

Langs	Source	Lang1	Lang2	Univ	NE	Mixed	Undef
EN-HI	FB+TW	54.22	22.00	19.70	4.00	0.05	0.03
	FB	75.61	4.17	18.00	2.19	0.02	0.01
	TW	22.24	48.48	22.42	6.71	0.08	0.07
	Vyas	54.67	45.27	0.06	-	-	-
	FIRE	45.57	39.87	14.52	-	0.04	-
EN-BN	TW	55.00	23.60	19.04	2.36	-	-
	FIRE	32.47	67.53	-	-	-	-
EN-GU	FIRE	5.01	94.99	-	-	-	-
DU-TR	Nguyen	41.50	36.98	21.52	-	-	-
EN-ES	EMNLP	54.79	23.50	19.35	2.08	0.04	0.24
EN-ZH		69.50	13.95	5.88	10.60	0.07	-
EN-NE		31.14	41.56	24.41	2.73	0.08	0.08
AR-AR		66.32	13.65	7.29	11.83	0.01	0.90

There is always some more stuff that you will need to add at some later point. Be sure to at least make a note about it somewhere.

3 Related Work

This chapter describes state-of-the-art techniques that can be applied to neural networks to better generate natural looking images. Section 3.1 gives an introduction to Generative Adversarial Networks (GAN). Section 3.2 will focus on state-of-the-art techniques used to bolster GAN. Section 3.3 will describe related work and how the methods from Section 3.2 are applied to different tasks.

3.1 Generative Adversarial Networks

Yann LeCun, Director of AI Research at Facebook and Professor at NYU, wrote in late July that the advances in deep learning he was most excited about, was adversarial training, or Generative Adversarial Networks (GAN) (LeCun and Bennett, 2016). The concept was introduced by Ian J. Goodfellow in a paper from 2014 (Goodfellow et al., 2014). The overall idea is to train two or more neural networks and let them work against each other as adversaries. On the most basic level, this approach consists of two neural networks, the generator G and the discriminator D . The task of generator G is to create an output of some kind, while the discriminator D 's task is to evaluate G 's output. Both models will use the output of the other model to improve their output. As Goodfellow puts it in his paper:

"The generative model can be thought of as analogous to a team of counterfeiters, trying to produce fake currency and use it without detection, while the discriminative model is analogous to the police, trying to detect the counterfeit currency. Competition in this game drives both teams to improve their methods until the counterfeits are indistinguishable from the genuine articles." [Goodfellow et al., 2014, p. 1]

GANs have for most part been used to generate natural looking images. The discriminator D 's task is to rate images on how natural they appear by output a scalar value given an input image. Ideally, an image from the training set, a real image, should get a good rating while an image that is generated from G should get a poor rating. During training, the discriminator network is shown different images from both the training set and samples from G , and told how it should rate them. This is then used to adjust the network and improve the discriminator by slowly modifying what features are used to spot fake images. The trick to GAN is that these adjustments are supplied back to the generator and in turned used to generate better samples. The generator G generates its samples given an input Z , which is generally an image consisting of randomly generated noise. The network will try to recognise patterns in the noise and apply elements of pre-

3 Related Work

viously seen images, resulting in an image that should look natural. A trained network can be fed any image and it will try to apply learned images to that image, morphing from the original image to something similar to what it was trained with.

3.2 State-of-the-art

State-of-the-art techniques for artificial neural networks that can be applied to generative adversarial networks.

3.2.1 Batch normalization

Regular normalization is a step in training of neural networks that usually only occurs during the pre-processing of input data, before the training phase. It is common to scale the values of the dataset to be of a range between zero and one, or between minus one and positive one. This means that the largest value of the entire dataset is changed to one, while the lowest is set to zero or minus one. Batch Normalization (Ioffe and Szegedy, 2015) changes normalization from a step that is executed once during pre-training and on the whole dataset, to something that occurs at each iteration and only at a small batch at a time. This technique makes the network less susceptible to unoptimized parameters while also allowing for the use of much higher learning rates. The use of higher learning rates allows for fewer iterations as well as higher degree of accuracy than normal normalization. Regular normalization can still be applied during preprocessing. Salimans et al. (2016, p. 4) introduce virtual batch normalization which normalizes the images during preprocessing to be better suited for batch normalization later on. This technique makes the images less dependent on the other images within the same batch.

3.2.2 The All Convolutional net

An approach to convolutional neural networks that replaces the pooling layers, used to downsample input images, with a pure convolutional layer (Springenberg et al., 2016). Regular pooling layers are predefined to behave in a certain way, for instance max-pooling that is used to downsample an input by choosing its greatest value. An all convolutional network is able to learn its own kind of downsampling and not be restricted by the initial architecture of the network. This approach can be used by both the generator and discriminator network.

3.2.3 Sparse connectivity

Sparse connectivity is the concept of reducing the number of connections within the network (Goodfellow et al., 2016, p. 335-336), eliminating fully connected layers. This is also known as sparse connections, sparse weights and sparse interactions. The benefits to this approach is fewer number of operations as well as reduced memory consumption. The idea behind sparse connectivity is that not every pixel of an image is important on its own, but rather the shapes and patterns the pixels are a part of.

3.2.4 Feature matching

The generator in regular GANs trains by purely using the outputs from the discriminator. Feature matching (Salimans et al. 2016, p. 2) changes this by making the generator learn features directly from the training data, and not the discriminator. The discriminator is tasked with selecting the features that are most important. This technique makes the networks more stable during training and creates a more robust classifier, but not necessarily more natural images (Salimans et al. 2016, p. 6 - 8).

3.2.5 Minibatch Discrimination

This technique helps to stabilize the discriminator by letting it look at images in batches rather than a single sample at a time (Huszár, 2016). This forces the generator to improve its samples in a broader spectrum and result in better images. Salimans et al. (2016, p. 3) compare Minibatch Discrimination with Feature Matching and conclude that minibatch discrimination converges more quickly resulting in better images, but the resulting classifier is not as strong as one with Feature Matching. It should be noted that these methods are not compatible with each other because Minibatch Discrimination requires a regular discriminator.

3.3 Related work

This section will describe previous work done with GANs that are related to this project.

3.3.1 Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks

Further work has been done to improve GANs. Radford, Metz and Chintala (2015) introduce convolution to GANs, coining the term deep convolutional generative adversarial networks (DCGAN). Regular convolutional networks have mostly been used with supervised learning, the authors use DCGAN to train unsupervised on large imagesets with the goal of generating samples that appear natural. Their network employs the techniques discussed in the previous section; Batch normalization and all Convolutional network with Sparse connectivity. They trained their network on a containing over three million images of bedrooms [<http://lsun.cs.princeton.edu/2016/>]. The network is able to learn the general patterns of a bedroom and from that generates new images that are not the original dataset. The result after five epochs of training are shown in the below. Although low resolution, the images are distinctly looking like bedrooms. The project was open source and was ported to TensorFlow (Kim, 2016).

Sett in bilde her

3.3.2 Image Completion with Deep Learning in TensorFlow

Amos (2016) shows how a trained GAN can be used to fill in missing pieces of an image . The article is heavily based on the previously mentioned paper (Radford, Metz and

3 Related Work

Chintala, 2015) and the TensorFlow version (Kim, 2016) of the code. The author trains the network using a dataset consisting of faces. After the network has had sufficient training, the author inputs unseen faces where the middle part of the images are replaced by a grey box. The generator is able to fill in and complete the missing parts of the images in quite an impressive way. The project is implemented with TensorFlow and the code is publicly available (Amos, 2016).

3.3.3 Improved Techniques for Training GANs

Salimans et al. (2016) improve GANs further by presenting new architectural features that they apply to semi-supervised learning. The authors used the previously mentioned state-of-the-art techniques to generate even more natural images as well as a classifier. The most interesting aspect of this paper is the evaluation metric to score how well the model performed. Salimans et al. used the pre-trained image classifier network Inception, created by Google (Szegedy et al., 2015). Inception recognizes what objects are in the image and Salimans et al. used this to determine how natural the generated images appear.

3.3.4 Learning to Protect Communications with Adversarial Neural Cryptography

Researchers at Google applied GAN to the problem of protecting digital communication (Martín and Andersen, 2016). The networks were able to learn a basic form of encryption without any prescribed knowledge. One interesting aspect is that they actually use two neural networks as the generator, one as the sender and encrypter, and the other as the receiver and decrypter. The adversarial network acts as a man-in-the-middle and was tasked with intercepting and decrypting the encrypted messages between the generators. The project was implemented in TensorFlow using a convolutional network that is fully connected and trained using batch normalization. The adversarial network is given a slight advantage by training it for one additional minibatch for every batch that the generators were trained.

4 Experiments and Results

4.1 Experimental Plan

Trying and failing is a major part of research. However, to have a chance of success you need a plan driving the experimental research, just as you need a plan for your literature search. Further, plans are made to be revised and this revision ensures that any further decisions made are in line with the work already completed.

The plan should include what experiments or series of experiments are planned and what questions the individual or set of experiments aim to answer. Such questions should be connected to your research questions, so that in the evaluation of your results you can discuss the results wrt to the research questions.

4.2 Experimental Setup

The experimental setup should include all data — parameters, etc. — that would allow a person to repeat your experiments.

4.3 Experimental Results

Results should be clearly displayed and should provide a suitable representation of your results for the points you wish to make. Graphs should be labeled in a legible font. If more than one result is displayed in the same graph, then these should be clearly marked. Please choose carefully rather than presenting every result. Too much information is hard to read and often hides the key information you wish to present. Make use of statistical methods when presenting results, where possible to strengthen the results. Further, the format of the presentation of results should be chosen based on what issues in the results you wish to highlight. You may wish to present a subset in the experimental section and provide additional results in an appendix.

5 Evaluation and Conclusion

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Nam consequat pulvinar hendrerit. Praesent sit amet elementum ipsum. Praesent id suscipit est. Maecenas gravida pretium magna non interdum. Donec augue felis, rhoncus quis laoreet sed, gravida nec nisi. Fusce iaculis fermentum elit in suscipit.

5.1 Evaluation

When evaluating your results, avoid drawing grand conclusions, beyond that which your results can in fact support. Further, although you may have designed your experiments to answer certain questions, the results may raise other questions in the eyes of the reader. It is important that you study the graphs/tables to look for unusual features/entries, and discuss these as well as the main findings.

5.2 Discussion

In this section it is important to include a discussion of not just the merits of the work conducted but also the limitations.

5.3 Contributions

What are the main contributions made to the field and how significant these contributions are.

5.4 Future Work

Consider where you would like to extend or improve this work, or how somebody else could continue it. These extensions might either be continuing the ongoing direction or taking a side direction that became obvious during the work. Further, possible solutions to limitations in the work conducted, highlighted in Section 6.2 may be presented.

Appendices