Scientific Reports

This document describes the structure for scientific reports and describes what is expected in each section of a scientific manuscript. You may vary slightly from this template, but most papers in the sciences are structured in this manner. Engineering papers have a slightly different structure, but in this class we will use the template that is generally used by the majority of publications in the sciences. Writing is a skill that requires organization, clarity, and grammatical accuracy. Being able to write well can have a profound effect on your career and graduate school is likely to be the last time that people will try to teach you to write well. If you feel uncomfortable writing, free one-on-one appointments are available at the SDSU writing center: https://writingcenter.sdsu.edu/ .

Title: A title that describes the work being done. For example: A study of phoneme recognition using recurrent neural networks

Author: You

Abstract

This is a paragraph that describes your work and is usually limited to no more than 200-250 words. It should present the study's goals, general ideas without equations, and a brief summary of the results. Here is a concrete example from a study that our lab did on North Atlantic right whales (Shiu *et al.*, 2020):

> Deep neural networks have advanced the field of detection and classification and allowed for effective identification of signals in challenging data sets. Numerous time-critical conservation needs may benefit from these methods. We developed and empirically studied a variety of deep neural networks to detect the vocalizations of endangered North Atlantic right whales (*Eubalaena glacialis*). We compared the performance of these deep architectures to that of traditional detection algorithms for the primary vocalization produced by this species, the upcall. We show that deep-learning architectures are capable of producing false-positive rates that are orders of magnitude lower than alternative algorithms while substantially increasing the ability to detect calls. We demonstrate that a deep neural network trained with recordings from a single geographic region recorded over a span of days is capable of generalizing well to data from multiple years and across the species' range, and that the low false positives make the output of the algorithm amenable to quality control for verification. The deep neural networks we developed are relatively easy to implement with existing software, and may provide new insights applicable to the conservation of endangered species.

Introduction

The introduction should provide an overview of the problem being solved and why the reader should care about it. In many cases, this is combined with a literature review where previous work in the area is cited and placed in relation to the current study. The literature review is sometimes split out into a separate section.

When the work is part of a class project as opposed to a research paper, students are usually demonstrating mastery of new concepts and this would be a good place to discuss the basic ideas behind these concepts.

Methods

The Methods section, sometimes called Methods and Materials, describes the techniques that you are using in your experiment as well as the data that were used in your study. For example, if you are describing a speech recognition experiment, you should have a section within methods that describes the data[1] that were used. How much data from how many speakers? Were the data collected in a quiet or noisy environment? What was the sample rate? For papers dealing with classification tasks, you should describe how the data were prepared, e.g. did you create a spectrogram? If so, what were the framing parameters, what window function did you use prior to the discrete Fourier transform? Were methods employed to discard or compensate for noise? Once you have described what will be presented to your machine learning algorithm (a neural network in this class), describe the network architectures that you examined in your study. If you have minor variations in your experiments such as testing different penalty weights for an L^2 regularizer, you can just state a parameter name (e.g., I used L^2 regularization with a penalty of λ) and then describe how this was varied in the experiments section.

A common mistake that computer scientists new to scientific writing make is to describe their code. Your description should be relatively independent of your implementation. For example, it is fine to say that the system was implemented in Python (or some other language and to cite libraries appropriately), but do not refer to specific subroutines in your architecture. For example, you might have a method called create_spectrogram, but instead you should describe the steps that your subroutine did: "I framed the speech signal using 20 ms windows that were advanced every 10 ms. These were windowed with a Hamming window and transformed to the frequency domain with a discrete Fourier transform. The spectra were converted to decibels…" Note that this is written in the past tense as the author is describing what they did.

Experiments

The experiments section describes how the study was conducted. For example, if you used N-fold cross validation, the number of folds would be described as well as any training/validation splits. If you varied a parameter such as the number of layers or

[1] Data are always plural. It is a common pattern of error to write sentences using data as if they were singular.

width of a network, the rationale for these tests and the values that you tested should be reported here. Example: "I used 10-fold cross validation to evaluate the performance of each network. I studied the effects of regularization by varying an L^2 penalty λ between 0.01 and 0.05 in steps of 0.02."

Results

This section shows the results of your study. It typically contains a description of the experiment outcomes without much interpretation other than stating trends or hypotheses that have been proven. Example: We found that regularization increased the ability of the network to generalize. Increasing the L^2 penalty improved results up to a value of …"

It is common to provide figures for this section. Think about how you want your data to be displayed. A common mistake for new report writers is to show too many figures. Scientific writing needs to tell a truthful story. You need to think about what the message of your study is and write in a way that convinces your reader of the points that you want to make. You should not "cherry pick" your data (show only the cases that tell the story you want), but if you have 10 plots that are all very similar, show one or two of them and note that the others are similar. The axes of your plots should always be labeled, and if appropriate, include a legend. Do not make your plots too small; plots whose text is unreadable are not useful.

Tables are very useful for providing compact summaries of variation. While one could not reasonably show 10 confusion matrices for different parameters of a speaker identification experiment, 10 rows in a table are much easier to show. Figures with trends of experimental results can also be very helpful.

Regardless of what section they are in, figures and tables should always have meaningful captions. Many experienced scientists first look at a paper's abstract and the figures. They use the caption information to decide what the main message of the study is and whether or not it is worth spending time reading. As such, descriptive captions are important. Here's an example of a caption by a former student of mine that appeared in Madhusudhana *et al.* (2021):

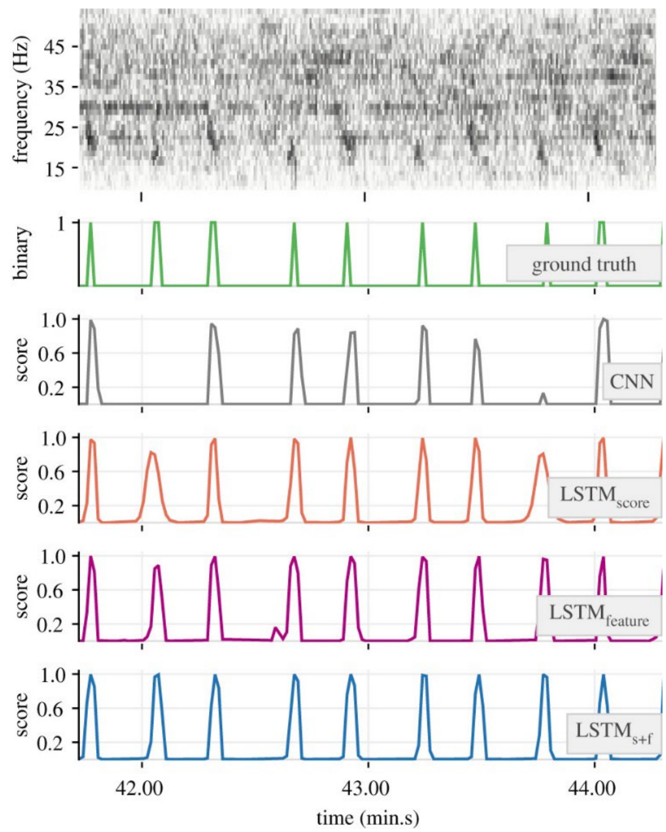

Figure 1 Demonstration of the trained detectors (from the experiment with 3 s segment overlap and PP = 50%) applied to a selected section of a recording. The top row shows the unsegmented spectrogram of the recording. Ground truth values associated with successive segments and the corresponding outputs from each model are shown as connected lines in subsequent rows.

<u>Discussion</u>

The Results section only discusses the measured results from your study. Speculation on why certain things worked well (or did not) belong in the discussion section. As an example, suppose you held the width of neural network layers fixed while increasing the rate of dropout for that layer. If one found that performance decreased as the dropout rate increased, one might speculate that part of this was due to a decreased capacity of the thinned network. It might be reasonable to go back and test this. If time did not permit one to do so, the conjecture could be voiced in this section.

This section frequently ends with the conclusions of the study, repeating the “take home” message along with he results. Some authors split the conclusion out into a separate section, but this is frequently not done.

Hopefully, your experiments will have gone well. In the classroom setting, you may occasionally find that you had problems completing the experiment. The discussion is a good place to clarify problems that you may have had that may have prevented obtaining the type of performance you had expected.

References

There are many citation styles. In this class, I would like you to use an alpha citation style, where the authors' name(s) and year are listed in parentheses. For example, "The King corpus consists of 51 male speakers … (Higgins and Vermilyea, 1995)." Many readers prefer this to the more concise number style as they do not need to flip back and forth to find the references if they know the literature well. If you do not already use a citation manager, it is worth learning one. LaTeX users typically use bibtex. Word and LibreOffice have a variety of managers, ranging from EndNote to free managers such as Zotero and Mendeley (these free managers can also be used with LaTeX). The university has a cite license for EndNote if you wish to use this product.

REFERENCES

Higgins, A., and Vermilyea, D. (1995). "King Speaker Verification," edited by International Telephone and Telegraph Corporation (Linguistic Data Consortium, 3600 Market Street, Suite 810, Philadelphia, PA 19104-2653), doi:10.35111/j0af-qf40.

Madhusudhana, S., Shiu, Y., Klinck, H., Fleishman, E., Liu, X., Nosal, E. M., Helble, T., Cholewiak, D., Gillespie, D., Sirovic, A., and Roch, M. A. (2021). "Improve automatic detection of animal call sequences with temporal context," J R Soc Interface 18(180), pp. 20210297, doi:10.1098/rsif.2021.0297.

Shiu, Y., Palmer, K. J., Roch, M. A., Fleishman, E., Liu, X., Nosal, E.-M., Helble, T., Cholewiak, D., Gillespie, D., and Klinck, H. (2020). "Deep neural networks for automated detection of marine mammal species," Sci Rep 10(1), doi:10.1038/s41598-020-57549-y.