cse 493 Project Design

Outline

1. Project Expectations
   a. Does my project meet expectations?
   b. FAQs
2. Picking a Project idea
   a. Inspiration
   b. How to read a research paper
3. Proposal, milestones, and final report
   a. Due dates, expectation, logistics
   b. Support

Outline

1. Project Expectations
    a. Does my project meet expectations?
    b. FAQs
2. Picking a Project idea
    a. Inspiration
    b. How to read a research paper
3. Proposal, milestones, and final report
    a. Due dates, expectation, logistics
    b. Support

Project Expectation

Open ended. Anything related to deep learning!

Project Expectation

Open ended. Anything related to deep learning!

Completed in groups of 1,2, or 3 people

- More people = higher expectations

Project Expectation

Open ended. Anything related to deep learning!

Completed in groups of 1,2, or 3 people

- More people = higher expectations

Two project options:

-Applications: Pick a new problem, and apply a deep learning solution!

- Your own data, data from another scientific field

-Model: Pick a standard problem, and find a new solution.

- Kaggle challenges etc

Project Expectation

- Title, Author(s)
- Abstract
- Related Work
- Methods
- Experiments
- Discussion
- Supplementary Material (optional)

https://courses.cs.washington.edu/courses/cse493g1/23sp/project/

Does my project meet expectations?

Checklist

- ☐ I am using deep learning
- ☐ I am training a model
- ☐ I am not just download a git repository and running “train”
- ☐ I am trying to understand my results via analysis

Does my project meet expectations?

Strong projects might…

- Propose a novel variant of a technique (which takes a lot of effort)
- Adapt an existing technique to a totally new problem (which takes a lot of effort)

Weaker projects might…

- Spend several weeks collecting/cleaning data rather than testing hypotheses
- Clone an existing repo and do minimal stitching to make it work for a Kaggle competition

FAQs

Does my project need novelty?
No! Novelty is one way to fulfill the requirements but not the only way.

Do I need to get state of the art performance?
Not at all. Most research contributions do not lead to state of the art performance

How else can I show effort?
Compare and contrast different methods, show multiple design iterations leading to improved accuracy, show creative design choices to tackle a new dataset, and more

How do I show proper analysis?
Do your best to answer “why” in your discussion! What kinds of mistakes is your model making? Where is it improving? Why does the loss/accuracy curve look the way it does? etc

FAQs

Can I change my project after the proposal?
Yes! This is just a first idea for a direction

Can I change my project after the milestone?
We do not encourage this. If you feel you have to, come and speak to us about why.

Do I have to train a model? Can I just build on a deep learning API?
Yes, we expect that all projects will involve training a deep learning model. However, you can also incorporate other elements such as outputs from deep learning API

Successful Past Projects

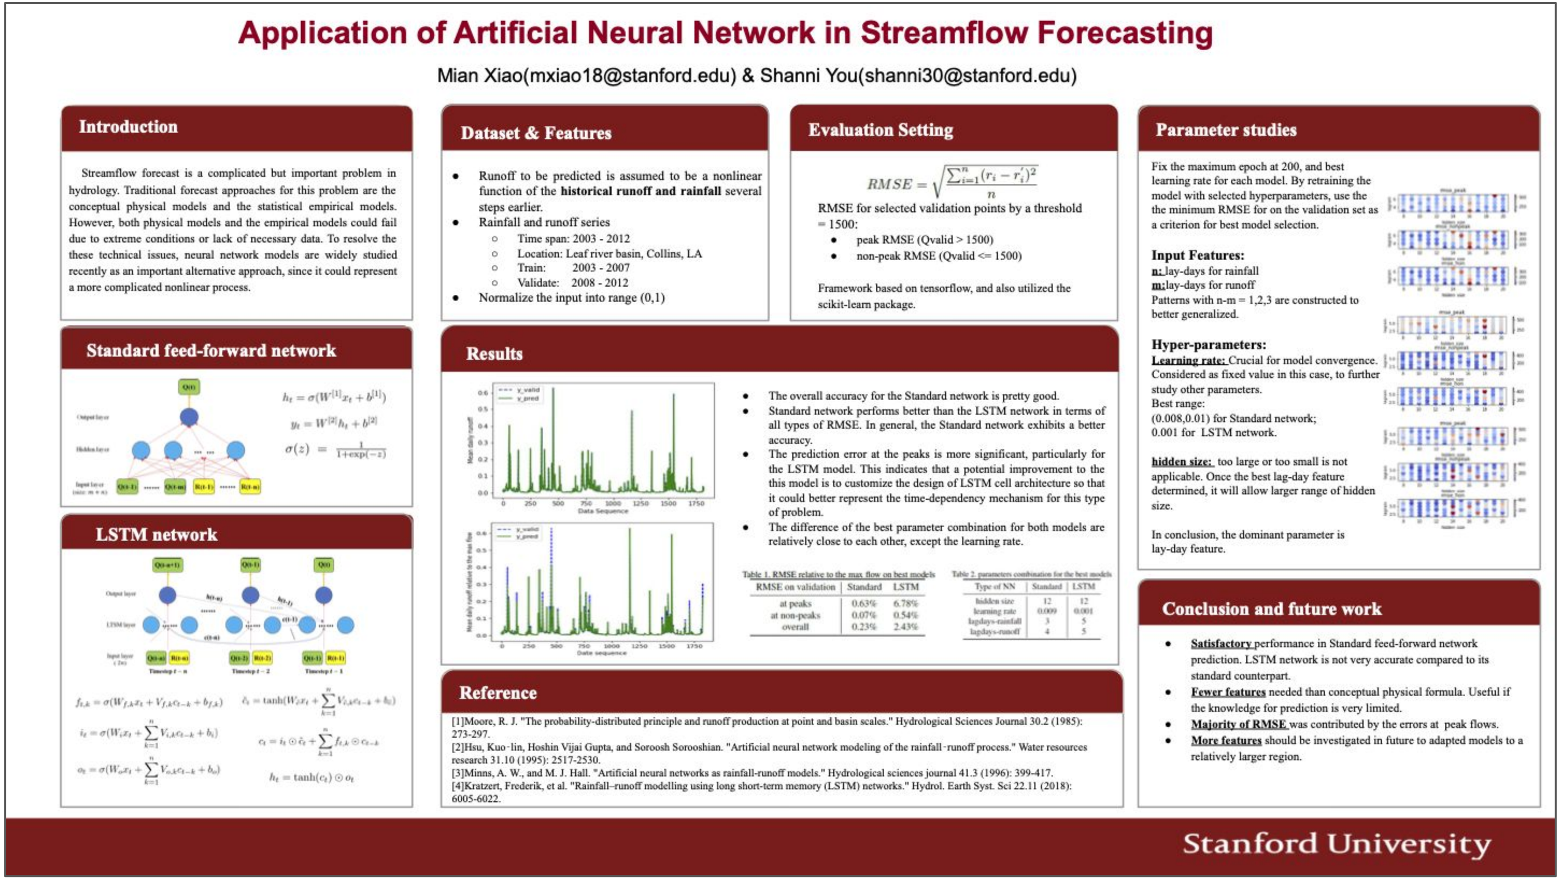

Application of Artificial Neural Network in Streamflow Forecasting

Mian Xiao(mxiao18@stanford.edu) & Shanni You(shanni30@stanford.edu)

Introduction

Streamflow forecast is a complicated but important problem in hydrology. Traditional forecast approaches for this problem are the conceptual physical models and the statistical empirical models. However, both physical models and the empirical models could fail due to extreme conditions or lack of necessary data. To resolve the these technical issues, neural network models are widely studied recently as an important alternative approach, since it could represent a more complicated nonlinear process.

Standard feed-forward network

$$h_t = \sigma(W^{[1]}x_t + b^{[1]})$$

$$y_t = W^{[2]}h_t + b^{[2]}$$

$$\sigma(z) = \frac{1}{1+\exp(-z)}$$

LSTM network

$$f_{t,k} = \sigma(W_f x_t + V_{f,k}c_{t-k} + b_{f,k}) \qquad \tilde{c}_t = \tanh(W_c x_t + \sum_{k=1}^{n} V_{c,k}c_{t-k} + b_c)$$

$$i_t = \sigma(W_i x_t + \sum_{k=1}^{n} V_{i,k}c_{t-k} + b_i) \qquad c_t = i_t \odot \tilde{c}_t + \sum_{k=1}^{n} f_{t,k} \odot c_{t-k}$$

$$o_t = \sigma(W_o x_t + \sum_{k=1}^{n} V_{o,k}c_{t-k} + b_o) \qquad h_t = \tanh(c_t) \odot o_t$$

Dataset & Features

- Runoff to be predicted is assumed to be a nonlinear function of the **historical runoff and rainfall** several steps earlier.
- Rainfall and runoff series
  - Time span: 2003 - 2012
  - Location: Leaf river basin, Collins, LA
  - Train: 2003 - 2007
  - Validate: 2008 - 2012
- Normalize the input into range (0,1)

Evaluation Setting

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}(r_i - r_i')^2}{n}}$$

RMSE for selected validation points by a threshold = 1500:

- peak RMSE (Qvalid > 1500)
- non-peak RMSE (Qvalid <= 1500)

Framework based on tensorflow, and also utilized the scikit-learn package.

Results

- The overall accuracy for the Standard network is pretty good.
- Standard network performs better than the LSTM network in terms of all types of RMSE. In general, the Standard network exhibits a better accuracy.
- The prediction error at the peaks is more significant, particularly for the LSTM model. This indicates that a potential improvement to the this model is to customize the design of LSTM cell architecture so that it could better represent the time-dependency mechanism for this type of problem.
- The difference of the best parameter combination for both models are relatively close to each other, except the learning rate.

Table 1. RMSE relative to the max flow on best models

RMSE on validation	Standard	LSTM
at peaks	0.63%	6.78%
at non-peaks	0.07%	0.54%
overall	0.23%	2.43%

Table 2. parameters combination for the best models

Type of NN	Standard	LSTM
hidden size	12	12
learning rate	0.009	0.001
lagdays-rainfall	3	5
lagdays-runoff	4	5

Parameter studies

Fix the maximum epoch at 200, and best learning rate for each model. By retraining the model with selected hyperparameters, use the the minimum RMSE for on the validation set as a criterion for best model selection.

Input Features:
n: lay-days for rainfall
m: lay-days for runoff
Patterns with n-m = 1,2,3 are constructed to better generalized.

Hyper-parameters:
Learning rate: Crucial for model convergence. Considered as fixed value in this case, to further study other parameters.
Best range:
(0.008,0.01) for Standard network;
0.001 for LSTM network.

hidden size: too large or too small is not applicable. Once the best lag-day feature determined, it will allow larger range of hidden size.

In conclusion, the dominant parameter is lay-day feature.

Conclusion and future work

- Satisfactory performance in Standard feed-forward network prediction. LSTM network is not very accurate compared to its standard counterpart.
- Fewer features needed than conceptual physical formula. Useful if the knowledge for prediction is very limited.
- Majority of RMSE was contributed by the errors at peak flows.
- More features should be investigated in future to adapted models to a relatively larger region.

Reference

[1]Moore, R. J. "The probability-distributed principle and runoff production at point and basin scales." Hydrological Sciences Journal 30.2 (1985): 273-297.
[2]Hsu, Kuo-lin, Hoshin Vijai Gupta, and Soroosh Sorooshian. "Artificial neural network modeling of the rainfall-runoff process." Water resources research 31.10 (1995): 2517-2530.
[3]Minns, A. W., and M. J. Hall. "Artificial neural networks as rainfall-runoff models." Hydrological sciences journal 41.3 (1996): 399-417.
[4]Kratzert, Frederik, et al. "Rainfall–runoff modelling using long short-term memory (LSTM) networks." Hydrol. Earth Syst. Sci 22.11 (2018): 6005-6022.

Stanford University

https://cs229.stanford.edu/proj2020spr/poster/Xiao_You.pdf

Successful Past Projects

Open-Ended Generative Commonsense Question Answering with Knowledge Graph-enhanced Language Models
Final Project Report for CS229, Spring 2021

Hanson Lu
Stanford University
hansonlu@stanford.edu

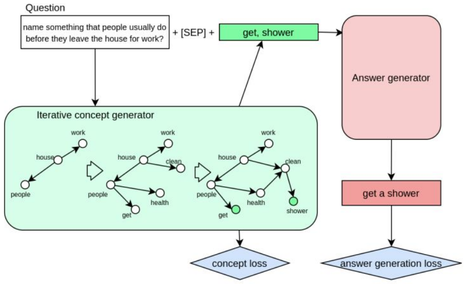

Figure 1: Planned Architecture of our model.

https://cs229.stanford.edu/proj2021spr/report2/81976767.pdf

Successful Past Projects

For Reports: https://cs229.stanford.edu/proj2021spr/

For Posters: https://cs229.stanford.edu/proj2020spr/

Outline

1. Project Expectations
    a. Does my project meet expectations?
    b. FAQs
2. Picking a Project idea
    a. Inspiration
    b. How to read a research paper
3. Proposal, milestones, and final report
    a. Due dates, expectation, logistics
    b. Support

Picking a project idea

One way to have novelty in your project is to take inspiration from things you care about outside of the course.

- Interested in healthcare? Robotics? Animals? Finance? Sports? Solve a problem that you are uniquely positioned to solve!

Picking a project idea

One way to have novelty in your project is to take inspiration from things you care about outside of the course.

- Interested in healthcare? Robotics? Animals? Finance? Sports? Solve a problem that you are uniquely positioned to solve!

Practical considerations:

1. Data - does data exist for your problem
2. Code - does an implementation exists?
3. Compute - do not train a 100 Billion parameter language model

Picking a project idea

Conferences:

CVPR: IEEE Conference on Computer Vision and Pattern Recognition

ICCV: International Conference on Computer Vision

ECCV: European Conference on Computer Vision

NeurIPS: Neural Information Processing Systems

ICLR: International Conference on Learning Representations

ICML: International Conference on Machine Learning

(I personally like looking at best paper awards)

Picking a project idea

Labs at UW

Vision + Graphics (GRAIL) - https://grail.cs.washington.edu/

Vision (RAIVN) - https://raivn.cs.washington.edu/

NLP (H2lab) - https://h2lab.cs.washington.edu/

NLP (Noah Ark) - https://noahs-ark.github.io/

Robotics (RSE lab) - http://rse-lab.cs.washington.edu/

Picking a project idea

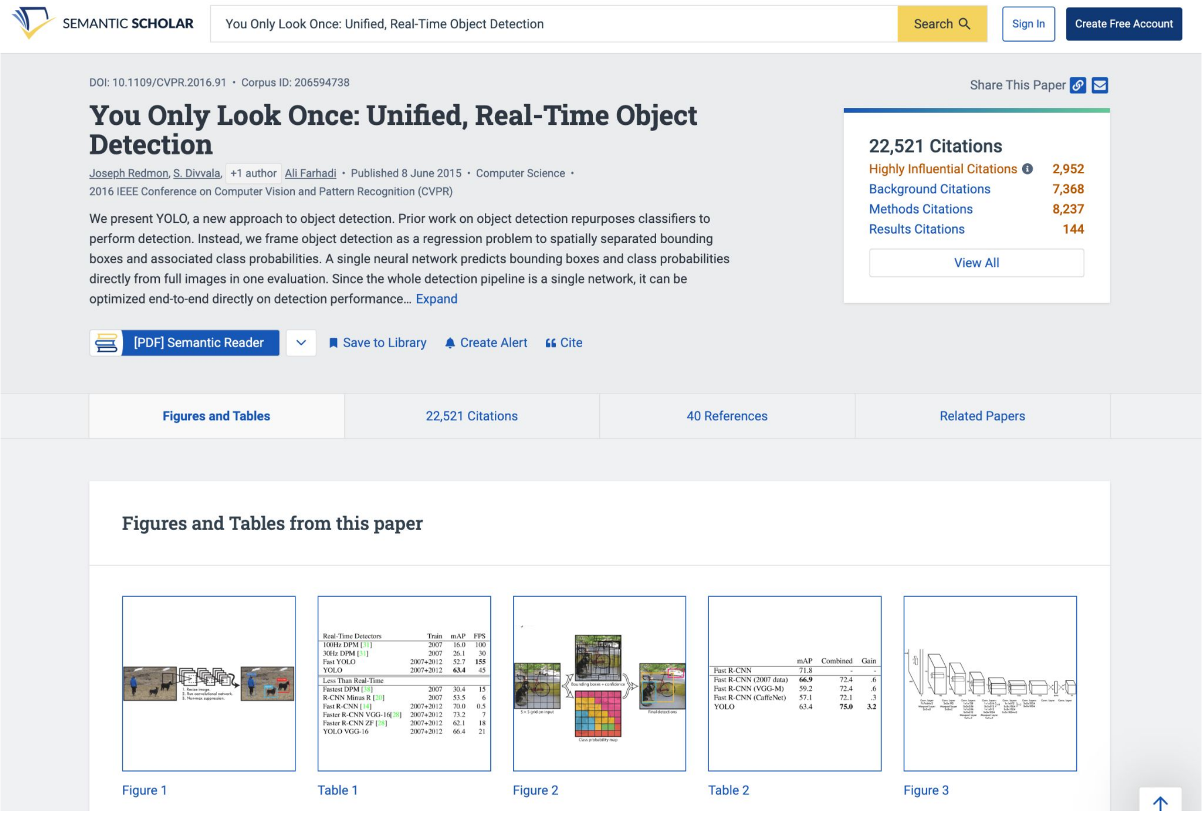

Reading papers

Do not read a paper linearly on your first pass

- First, read the abstract (word for word) as well as the figures & captions
- Does the paper still seem relevant? If so, read the methods + results
- Finally, read the entire paper linearly (if the additional detail seems useful)

Papers are not always the most efficient way to digest an idea. Also try looking for:

- Talks, videos, or blog posts on the topics
- Github repos, containing actual code for the idea

Reading papers

You Only Look Once: Unified, Real-Time Object Detection

Joseph Redmon[*], Santosh Divvala[*†], Ross Girshick[¶], Ali Farhadi[*†]
University of Washington[*], Allen Institute for AI[†], Facebook AI Research[¶]
http://pjreddie.com/yolo/

Tip to make your life easier

A	B	C	D
	Paper	**Summary**	**year**
	Prompt engineering (NLP)		
1	Chain of Thought Prompting Elicits Reasoning in Large Language Models	Chain of thought helps	2022
1	PromptSource: An Integrated Development Environment and Repository...	Collaborativly creating prompts	2022
1	Show Your Work: Scratchpads for Intermediate Computation with Language Models	same as chain of thought	2021
1	Large Language Models are Zero-Shot Reasoners	"Lets think step by step"	2022
1	Exploiting Cloze-Questions for Few-Shot Text Classification and Natural Language Inference		2020
	Making pre-trained language models better few-shot learners		

Outline

1. Project Expectations
    a. Does my project meet expectations?
    b. FAQs
2. Picking a Project idea
    a. Inspiration
    b. How to read a research paper
3. Proposal, milestones, and final report
    a. Due dates, expectation, logistics
    b. Support

Due Dates

Project Proposal (1 Page) – Due April 24th 11:59pm

Project Milestone (3-4 Pages) – Due May 12th 11:59pm

Project Final Report (5-6 Pages) – Due June 3rd 11:59pm

Project Poster Session – Final Exam Week

Project Proposal

- Describe the state of related work,
- Explain a problem that is unsolved given that statement,
- Introduce your ideas as an unique insight to tackle the problem or research question,
- Articulate the technical challenges you are likely to encounter,
- Plan out the experiments that justify the utility of the insight or answers the question,
- Your expected outcome

Milestone (Due May 12th 11:59pm)

~3-page progress report, more or less containing:

1. Literature review (3+ sources)
2. Indication that code is up and running
3. Data source explained correctly
4. What Github repo or other code you're basing your work off of
5. Ran baseline model have results
    a. Yes, points are taken off for no model running & no preliminary results
6. Data pipeline should be in place
7. Brief discussion of your preliminary results

Support

Ranjay - Thursdays 11:30am - 12:30pm

Aditya - Thursdays 11:30am - 1:00pm

Efficient Neural Networks, Retrieval

Sarah - Fridays 1:00pm - 3:00pm

Zero shot image classification, RNNs, object detection, ~ RL, ~ LLMs