Understanding and Simplifying Architecture Search in Spatio-Temporal Graph Neural Networks

Zhen Xu[1] **Quanming Yao**[2] **Yong Li**[2] **Qiang Yang**[3]

[1]4Paradigm
[2]Department of Electronic Engineering, Tsinghua University
[3]Department of Computer Science and Engineering, Hong Kong University of Science and Technology

Abstract Compiling together spatial and temporal modules via a unified framework, Spatio-Temporal Graph Neural Networks (STGNNs) have been popularly used in the multivariate spatio-temporal forecasting task, e.g. traffic prediction. After the numerous propositions of manually designed architectures, researchers show interest in the Neural Architecture Search (NAS) of STGNNs. Existing methods suffer from two issues: (1) hyperparameters like learning rate, channel size cannot be integrated into the NAS framework, which makes the model evaluation less accurate, potentially misleading the architecture search (2) the current search space, which basically mimics Darts-like methods, is too large for the search algorithm to find a sufficiently good candidate. In this work, we deal with both issues at the same time. We first re-examine the importance and transferability of the training hyperparameters to ensure a fair and fast comparison. Next, we set up a framework that disentangles architecture design into three disjoint angles according to how spatio-temporal representations flow and transform in architectures, which allows us to understand the behavior of architectures from a distributional perspective. This way, we can obtain good guidelines to reduce the STGNN search space and find state-of-the-art architectures by simple random search. As an illustrative example, we combine these principles with random search which already significantly outperforms both state-of-the-art hand-designed models and recently automatically searched ones. [1]

1 Broader Impact Statement

After careful reflection, the authors have determined that this work presents no notable negative impacts to society or the environment.

2 Submission Checklist

All authors must include a section with the AutoML **Submission Checklist** in their manuscripts, both at submission and camera-ready time. The submission checklist is a combination of the NeurIPS '22 checklist and the NAS checklist. For each question, change the default `\answerTODO{}` (typeset **[TODO]**) to `\answerYes{[justification]}` (typeset [Yes]), `\answerNo{[justification]}` (typeset [No]), or `\answerNA{[justification]}` (typeset [N/A]). **You must include a brief justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license of the code and datasets? [Yes] See Section **??**.
- Did you include all the code for running experiments? [No] We include the code we wrote, but it depends on proprietary libraries for executing on a compute cluster and as such will not be

[1]Our code is available at `https://github.com/AutoML-Research/SimpleSTG`.

runnable without modifications. We also include a runnable sequential version of the code that we also report experiments in the paper with.

- Did you include the license of the datasets? [N/A] Our experiments were conducted on publicly available datasets and we have not introduced new datasets.

Please note that if you answer a question with `\answerNo{}`, we expect that you compensate for it (e.g., if you cannot provide the full evaluation code, you should at least provide code for a minimal reproduction of the main insights of your paper).

Please do not modify the questions and only use the provided macros for your answers. Note that this section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] In section 3 and Section 4, we develop a framework to support the claims. In Section 5, our experiments show the effectiveness and efficiency as by the claims.
   (b) Did you describe the limitations of your work? [Yes] This has been discussed in Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] This has been discussed in the Section Broader Impact Statement.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? `https://automl.cc/ethics-accessibility/` [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results, including all requirements (e.g., `requirements.txt` with explicit version), an instructive `README` with installation, and execution commands (either in the supplemental material or as a URL)? [Yes] The environment, dataset, running command is included in `https://github.com/LARS-research/SimpleSTG/blob/main/README.md`
   (b) Did you include the raw results of running the given instructions on the given code and data? [No] However, all the results can be easily reproduced by the open sourced code.
   (c) Did you include scripts and commands that can be used to generate the figures and tables in your paper based on the raw results of the code, data, and instructions given? [Yes]
   (d) Did you ensure sufficient code quality such that your code can be safely executed and the code is properly documented? [Yes]
   (e) Did you specify all the training details (e.g., data splits, pre-processing, search spaces, fixed hyperparameter settings, and how they were chosen)? [Yes]
   (f) Did you ensure that you compared different methods (including your own) exactly on the same benchmarks, including the same datasets, search space, code for training and hyperparameters for that code? [Yes]

(g) Did you run ablation studies to assess the impact of different components of your approach? [Yes]

(h) Did you use the same evaluation protocol for the methods being compared? [Yes]

(i) Did you compare performance over time? [Yes]

(j) Did you perform multiple runs of your experiments and report random seeds? [Yes]

(k) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]

(l) Did you use tabular or surrogate benchmarks for in-depth evaluations? [No]

(m) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]

(n) Did you report how you tuned hyperparameters, and what time and resources this required (if they were not automatically tuned by your AutoML method, e.g. in a NAS approach; and also hyperparameters of your own method)? [Yes]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes]

(b) Did you mention the license of the assets? [Yes]

(c) Did you include any new assets either in the supplemental material or as a URL? [No]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]