

Guides for deep learning submissions

Aiming to strengthen scientific reproducibility among our peers, this document lays down some guidelines for submissions to IEEE T-IFS. Reproducibility is one aspect that should be assessed by reviewers and considered by the AE in the final acceptance decision.

The guidelines explicitly refer to papers dealing with steganalysis and multimedia forensics, however they are applicable to all papers submitted to T-IFS in which the proposed system relies on deep learning.

Thanks to recent advances in machine learning, forensic and steganalytic attacks have grown in sophistication and complexity. With more parameters, it becomes easier to omit key information. This is particularly so for Deep Learning (DL) methods, which have many variants and details. When some of these are omitted from the paper it weakens reproducibility, and often extends the review process.

Authors are encouraged to make both code and data sets available¹ (preferably in time for the review process) but this is not a formal requirement at the moment. In any case, supplementary code is not a substitute for a proper written description of the method. To help authors, the T-IFS EB has drawn up this checklist of details, that DL submissions on IFS topics should take particular care to include.

- The type of network (CNN, RNN, BDRNN, etc) and loss function.
- The topology:
 - type of each layer (pooling, conv, non-linearity, etc),
 - activation function of each layer (tanh, sigmoid, max- or avg-pooling, etc)
 - any parameters (strides, weights size, etc).
- Any pre-processing of network inputs, and interpretation of outputs. Note that there are different rescalings commonly called "normalization", so this needs to be specified.

Training phase

- How the network is initialized.
 - If randomly, give the exact parameters of the initializing distributions.
 - If there is pre-training, the procedure used (SAE, RBM, etc) and any parameters.
- The learning algorithm (SGD, Adagrad, AdaDelta, HF, etc).
 - Hyperparameters of the learning algorithm: (momentum, etc).
 - Mini-batches and (if appropriate) number of representatives of each class in learning batches. Where appropriate (e.g. steganalysis), specify if the data are stratified (i.e. paired cover and stego kept together).
 - If used, exactly what form of batch normalization and pooling.
 - Initial learning rate and its evolution during the learning, number of iterations or epochs, dropout or other regularization factors, etc.
 - The order that training data is visited during the learning process, and whether shuffled between epochs.

¹ From the Transactions website: "The Transactions encourages authors to make their publications reproducible by making all information needed to reproduce the presented results available online. This typically requires publishing the code and data used to produce the publication's figures and tables on a website; see the supplemental materials section of the information for authors. It gives other researchers easier access to the work, and facilitates fair comparisons."

- Stopping criterion (is there a fixed number of iterations or not?)
- Which data was used to determine the hyperparameters, learning rates, regularization parameters, and stopping criterion; the policy for hyperparameter search.

Test phase

- How data is broken down into training, validation, and testing sets.
- If cross-validation is used, how it is arranged and whether it is repeated.

Implementation

- If a GPU is used, which model.
- If you have used a standard deep learning tool, the exact version and any tool-specific parameters.
- If you propose techniques for scalability (kernel approximations, etc) then specify.

We encourage the inclusion of a “methodology” section in the paper to contain these details, which could be in an appendix, or as supplementary material providing that this is available for review. Algorithms and methods need not be described from scratch: a good citation is adequate, as long as the exact variety is clear and all parameters have been specified.

While we encourage the publication of negative results contradicting previously published research, in these cases a higher standard of reproducibility applies. Barring exceptional circumstances, data sets should be made available for the review process (if not, this should be explained to the AE). Finally, authors should demonstrate the statistical significance of results claimed.