Propositions accompanying the thesis

Making and breaking decision boundaries

Improving accuracy and assessing robustness
of deep learning for medical image analysis

1. Ensembling several best configurations of the method discovered during its tuning is an effective and practical way to improve accuracy and stability of method's predictions. (*Chapter 2*)
2. In deep-learning-based segmentation or detection, using prior information on what parts of the input cannot contain the target structure(s) can significantly improve accuracy, and this can often be accomplished by very small, simple modifications of the algorithm. (*Chapter 2 and 3*)
3. Training machine learning algorithms to be consistent under appropriately chosen types of transformations can improve accuracy, especially when unlabeled data can be used for the consistency training. (*Chapter 4*)
4. Off-the-shelf machine learning models can sometimes be very easily manipulated by black-box adversarial attacks. (*Chapter 5*)
5. Medical image analysis systems that are trained on private data and are close-sourced, with no extensive details on their underlying algorithm publicly available, may be naturally well-protected against adversarial attacks by external parties. (*Chapter 5*)
6. Complex AI models that are truly robust to adversarial attacks may never be possible with the current AI paradigm.
7. The ability of an algorithm to reproduce the type of label it was trained with does not guarantee its utility in practice.
8. Focus on novelty and performance leads to bad and boring science, as well as stressed researchers.
9. Reporting and evaluation standards in method-focused research in medical image analysis need to improve for the field to become truly scientific.
10. Before concluding one algorithm is better than another, alternative explanations for the performance difference should be considered, including trivial ones, such as differences in optimization or training data sampling.
11. In a parallel universe where 90% novelty-focused technical works in medical image analysis were never conducted (assuming 90% is sampled randomly), we have algorithms that are as performant, robust, and interpretable as they are in our universe, while having much fewer papers to read and review.