

---

## Appendix

---

### A. Multi-Airfoil Experiment

As an additional test of the generalization capabilities of our model, we performed an experiment in which we train the CFD-GCN on two airfoils and test on a previously unseen one. For this task, we use the same range of physical parameters as before on all airfoils. The training set is composed of simulations using the NACA4412 and the RAE2822 airfoils. The training set uses the same NACA0012 airfoil as the previous tasks.

Figure 11 presents the training curves for the CFD-GCN and the baselines. As before, we observe that the CFD-GCN model generalizes better to the test set, which contains the previously unseen airfoil. With an RMSE of 0.34 the upsampled coarse mesh baseline's results were too high to be displayed in the same plot.






Figure 11. Training curves for the multi-airfoil experiment. As in the previous generalization task, the CFD-GCN is able to generalize better to the new conditions present in the test set. The test loss is lower than the training loss here simply because the training and test set are composed of different airfoils, with meshes containing different number of nodes, and thus their losses are not directly comparable.

## B. Mesh Optimization


In Figure 12, we present changes to the mesh during the optimization process for the generalization task.




(a) The NACA0012 mesh after training.



(b) The original NACA0012 mesh before training.



(c) The meshes before and after training, superimposed.



(d) Farther view of before and after meshes, superimposed.

Figure 12. Comparison of the coarse mesh before and after being optimized during training. Changes are greater around the airfoil, where the gradients of the loss are large. Regions further away from the wing, which do not affect prediction strongly, are mostly unaltered.

### C. Interpolation Experiment

In Figure 13, we present the predictions and ground truth for the fields that were omitted in the main text for the interpolation task in Section 4.1.




Figure 13. CFD-GCN model prediction and ground truth for a test sample in the interpolation task. The y component of the velocity and the pressure output fields for the same sample as in Figure 6 are presented here.

## D. Generalization Experiment

In Figures 14 and 15, we present the predictions and ground truth for the fields that were omitted in the main text for the interpolation task in Section 4.2. In Figure 16, we present a prediction for the upsampled coarse mesh baseline.

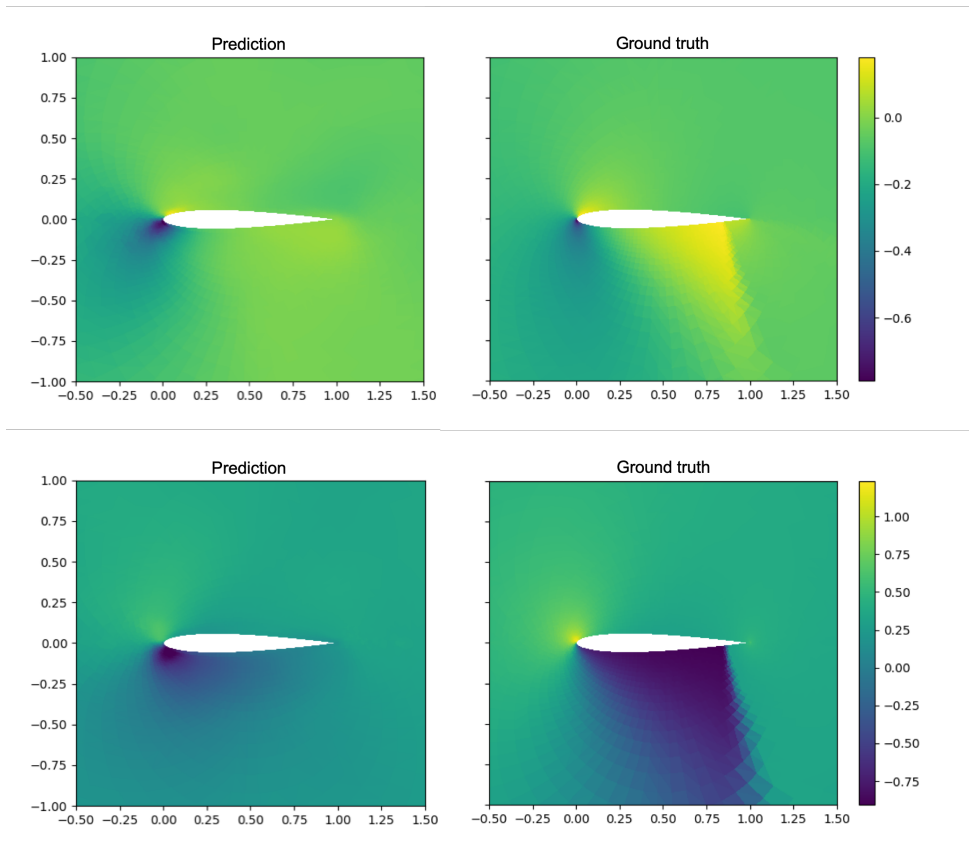


Figure 14. The GCN baseline prediction for a test sample with a large shock in the generalization task. The y component of the velocity and the pressure output fields for the same sample as in Figure 8 are presented here.




Figure 15. The CFD-GCN model prediction for a test sample with a large shock in the generalization task. The y component of the velocity and the pressure output fields for the same sample as in Figure 9 are presented here.

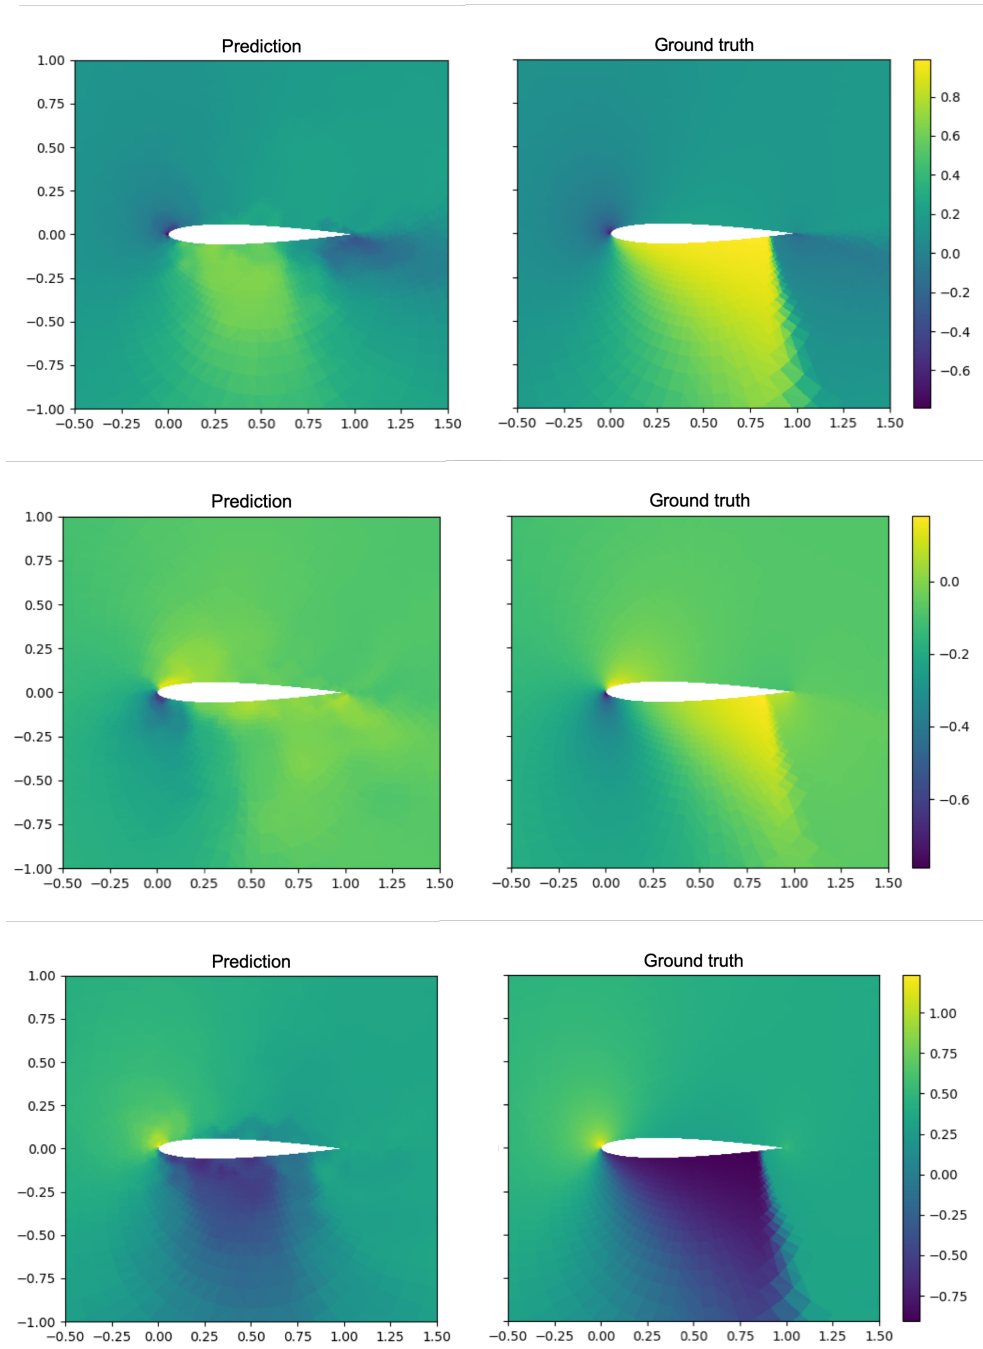


Figure 16. The upsampled coarse mesh baseline prediction for a test sample with a large shock in the generalization task. The x and y components of the velocity and the pressure output fields are presented here.