

A SUPERVISED LEARNING APPROACH TO PREDICTING MULTIGRID CONVERGENCE

NICOLAS NYTKO*

In collaboration with: Matthew West, Luke Olson, Scott MacLachlan

Abstract. Classical AMG solvers often require careful parameter tuning to achieve optimal convergence, and the way these parameters affect performance can be unpredictable in practice. Evidence is presented that supervised learning techniques are able to learn certain characteristics of two-level multigrid solvers, particularly the rate of convergence and optimal relaxation weight for a given coarse/fine mesh splitting. Random perturbations of C/F splittings are generated and evaluated in a multigrid solver to train a convolutional neural network (CNN) in order to predict convergence and a relaxation weight for the 1D variable coefficient Poisson equation, and to predict the convergence rate for a specific 2D convection-diffusion problem. Additionally for the 2D problem, the use of graph nets is explored for use on general finite-element meshes.

Key words. convergence, relaxation weight, machine learning, supervised learning, graph net, cnn, convolutional network

1. Introduction. Multigrid methods are some of the most widely used linear solvers for sparse systems, particularly ones that arise from the discretization of partial differential equations. By “downsampling” the original problem to a series of smaller and smaller problems, fast convergence can be obtained with minimal work [2]. Optimal “downsampling” is critical to the success of the method; classical Algebraic Multigrid (AMG) techniques attempt to create an interpolation operator by exploiting the structure of the original system [13]. However, these methods often require careful parameter tweaking to obtain a sensible operator and C/F splittings and often the best way to test the parameter space is to actually run the multigrid iterations.

In this paper, the following research question is addressed: can convergence rates of multigrid solvers be learned, and are neural networks an effective way of learning this attribute on the solver? The ability to predict this would aid in the use and study of classical AMG methods by allowing users to estimate efficacy ahead of time: for example being used to differentiate between different or the most optimal AMG methods for a specific problem. We address this question by introducing a method of generating data and training neural networks to predict convergence rates for a specific PDE or class of equations. Two problems are discussed here: (1) a 1D variable coefficient Poisson equation, and (2) a 2D recirculating flow convection-diffusion problem with specific parameters. Sections 2.1 and 2.3 describe how these random grid splittings are generated for the two problems and evaluated in a multigrid solver.

For the 1D poisson case, convolutional neural networks (CNNs) are trained to predict the convergence rate and optimal relaxation weight (Section 2.2). For the 2D convection-diffusion problem, a CNN was first trained to predict convergence (Section 2.4), and then two different graph nets were trained to extrapolate to differently sized inputs and non-grid meshes (Sections 2.5, 2.6). Numerical results consisting of the network error performance and sample predictions are shown in Section 3. Finally, further discussion and possible directions of future work are given in Section 4.

*University of Illinois at Urbana-Champaign, (nnytko2@illinois.edu).

43 **2. Methods.** The methods described here will be split into two subsections
 44 describing each problem that was explored. First, the Poisson case and the data
 45 generation is discussed (Section 2.1) followed by the respective convolutional neural
 46 network that was trained (Section 2.2). Afterwards, an overview of the data generation
 47 for the 2D convection-diffusion problem is given (Section 2.3) and then followed by
 48 a description of the convolutional network that was trained (Section 2.4). The use
 49 of two graph nets is also described: one that performs a simple edge convolution for
 50 each node (Section 2.5), and another that implements “message passing” and learns
 51 optimal edge weighting of the convolution (Section 2.6).

52 **2.1. 1D Poisson.** Formally, the problem being solved is the Poisson equation
 53 in one dimension with variable coefficients

$$54 \quad (2.1) \quad -\nabla \cdot (k(\mathbf{x}) \nabla \mathbf{u}) = f,$$

55 with the right-hand side $f(\mathbf{x})$ being arbitrarily chosen. Eqn. (2.1) is discretized using
 56 finite differences on a grid of $N = 31$ internal points on the domain $\Omega = [-1, 1]$ and
 57 Dirichlet boundary conditions, $\partial\Omega = 0$. To preserve the symmetric, positive definite
 58 properties of the resulting linear system, the $k(\mathbf{x})$ function is discretized on the grid
 59 *midpoints* [1].

60 Before the neural network can be trained, a dataset of approximately 300,000
 61 random C/F grid splittings and their computed convergence rate and optimal relax-
 62 ation weight was generated. A set of 6 reference grids were first created according to
 63 various “coarsening” factors:

$$64 \quad (2.2) \quad r = \{2 \ 3 \ 4 \ 5 \ 6 \ 7 \ 8 \ 9\}.$$

65 Each of the values in 2.2, say $r_i = j$ refers to a grid in which each j th point is coarse,
 66 and the rest fine. So, a coarsening by 3 is a grid that has roughly $\frac{1}{3}$ of its distribution
 67 of points as coarse and all others fine. Each of these reference grids was randomly
 68 permuted such that each grid point had a random probability of being flipped to the
 69 opposite value: coarse point to fine and fine point to coarse. For each reference grid,
 70 approximately 1000 random trials of each following probability were run:

$$71 \quad (2.3) \quad p = \{0.01 \ 0.05 \ 0.1 \ 0.25 \ 0.5 \ 0.75\}$$

72 Additionally, each trial was run with a randomly-chosen function for the variable
 73 coefficients $k(\mathbf{x})$. The function is randomly chosen from one of four functional forms

$$74 \quad (2.4) \quad k(\mathbf{x}) = \begin{cases} \alpha & 0 < \alpha < 10 \\ \text{rand}() (\alpha + 1) & 0 < \alpha < 10 \\ \alpha \cos(\pi x \beta) + \gamma & 0 < \alpha < 10, 0 < \beta < 10, \alpha < \gamma < 10 \\ \left| \sum_{i=1}^5 \alpha_i x^i \right| + 0.01 & -10 < \alpha < 10, \end{cases}$$

75 with coefficient values specifically picked to prevent the function from being non-
 76 positive at any value in Ω .

77 The full set of C/F splittings was then given to a 2 level V-cycle Multigrid solver
 78 that was run for a maximum of 15 iterations. The solver is composed of one round of
 79 weighted Jacobi pre-smoothing, a coarse error correction, then another round of Jacobi
 80 post-smoothing. At each iteration, the absolute error between the approximation
 81 and the “exact” solution (pre-computed via sparse linear solve) was found and saved:

82 $e_i = \|\mathbf{A}^{-1}\mathbf{f} - \mathbf{u}_i\|$. This sequence of errors was then used to compute the average
 83 convergence rate. To acquire the optimal relaxation weight, the smoother was run
 84 through a bracketed numerical optimization method with the assumption that the
 85 convergence rate is unimodal as a function of relaxation weight. ¹

86 **2.2. Poisson CNN.** Two separate deep convolutional networks with residual
 87 connections were trained to separately predict the optimal relaxation weight and con-
 88 vergence rate when given a linear C/F splitting. The architecture for both networks
 89 is identical and a textual overview is given here:

- 90 1. 6 1D CNN layers of kernel size 7, input 2 channels output 7 channels. Padding
 91 by three on each side to keep dimensions static.
- 92 2. 6 CNN layers of kernel size 5, input/output 7 channels. Padding by two on
 93 each side.
- 94 3. 6 CNN layers of kernel size 3, input/output 7 channels. Padding by 1 on each
 95 side.
- 96 4. Max-pooling layer of kernel 2, stride 2. Effectively reduces input size by half.
- 97 5. 8 Fully-connected layers to gradually reduce output to a scalar describing the
 98 convergence or relaxation weight.

99 Each layer is followed by an implicit ReLU nonlinear activation function. Because
 100 the convolutional layers (except for the first) keep the input and output size static,
 101 we are able to push residual values and skip layers similarly to a ResNet [8, 6]. Even
 102 numbered layers n take as input both the output of layers $n - 1$ and $n - 2$, while
 103 odd-numbered layers only use the output from layer $n - 1$.

104 The C/F splitting is remapped such that a coarse point is given a value of 1
 105 and a fine point the value of -1 . The variable coefficients were also re-discretized
 106 to be defined on the nodal points instead of midpoints in order for the splitting
 107 and coefficients to be represented by vectors of same length. These two were then
 108 stacked into a two-channel tensor for input into the CNN. When training, input values
 109 (convergence rate and relaxation weight) are normalized to be within the range of
 110 $[0, 1]$. This normalization is undone when the output is displayed.

111 **2.3. 2D Convection-Diffusion.** This specific problem models what is called
 112 the *double glazing problem*, modeling the temperature distribution of a cavity with a
 113 single “hot” wall. This is given by the differential equation:

$$114 \quad (2.5) \quad -k\nabla^2 u + \mathbf{w} \cdot \nabla u = f.$$

115 The wind velocity function, $\mathbf{w}(x, y)$ is defined as $\mathbf{w}(x, y) =$
 116 $[2y(1 - x^2) \quad 2x(1 - y^2)]$. The domain is a square of side length two centered at the
 117 origin, $\Omega = [-1, 1] \times [-1, 1]$. Dirichlet boundary conditions are defined on $\partial\Omega$, with
 118 the one “hot” wall defined on $x = 1$ with value $\partial\Omega_H = 1$. The other boundaries are
 119 “cold” walls with $\partial\Omega_c = 0$. A diffusivity constant of $k = 0.1$ is used. This problem is
 120 derived from an example by Elman, Silvester, and Wathen [4].

121 This PDE is discretized using finite-elements on a structured grid of 25x25 internal
 122 points using the Firedrake software for FEM discretizations [3, 12, 7, 10]. Using a
 123 grid as a basis for the discretization allows use of both CNN and more sophisticated
 124 graph convolutional techniques.

125 Generating a dataset of mesh splittings and convergence rates was done in an
 126 overall similar way to the 1D Poisson equation with a few notable differences. Again,

¹Experimental testing generally asserts this to be true, however this will remain a conjecture for now.

127 a set of reference C/F splittings were generated that are later permuted. For the
 128 convection-diffusion case, the following reference splittings were used:

- 129 1. All fine points
- 130 2. All coarse points
- 131 3. Splitting as given by Ruge-Stüben AMG ($\theta = 0.25$) [13, 11]
- 132 4. Coarsening in each direction by 2
- 133 5. Coarsening in each direction by 3
- 134 6. Coarsening in each direction by 4
- 135 7. Coarsening in each direction by 5

136 The entries of each individual reference splitting were then randomly permuted
 137 according to a defined probability. The probability values used are the same as those
 138 in the Poisson case, repeated here for convenience:

$$139 \quad (2.6) \quad p = \{0.01 \quad 0.05 \quad 0.1 \quad 0.25 \quad 0.5 \quad 0.75\}.$$

140 Generated splittings that are unsolvable (i.e., consist of no coarse points) were re-
 141 jected and their trial re-run. This generated set of C/F splittings was passed along
 142 to another 2 level V-cycle multigrid solver run for a maximum of 50 iterations. The
 143 interpolation operator was formed by means of *direct interpolation*, the implementa-
 144 tion of which graciously taken from the PyAMG [11] library. This solver performs two
 145 rounds of Jacobi pre-and-post relaxation, with a coarse error correction between the
 146 relaxation steps. The absolute error between the approximation at each iteration and
 147 the “exact” solution was computed and saved, with the full sequence used to find the
 148 average convergence rate. Note the optimal Jacobi relaxation weight was not found
 149 here, as experimentation found that the optimal weight would nearly always be 1.

150 The convection-diffusion problem was additionally re-discretized at 4 different
 151 mesh sizes: 15×15 , 25×25 , 35×35 , 50×50 . The above process was re-run for
 152 each mesh size to generate a new dataset for use in the graph nets. To distinguish
 153 between the two datasets, they will hereby be referred to as the *statically-sized* (only
 154 containing 25×25 mesh and splittings) and the *variably-sized* datasets.

155 **2.4. Convection-Diffusion CNN.** A 2D convolutional network was trained
 156 on the statically-sized dataset to predict convergence when given a C/F splitting for
 157 the specific recirculating flow problem. Since the input parameters are slightly less
 158 complex, a less deep (*shallower*, if you will) network was trained:

- 159 1. 3 2D CNN layers of kernel 7, input two channels output 7 channels. Padding
 160 by three on each side to keep dimensions.
- 161 2. 3 CNN layers of kernel 5, input/output 7 channels. Padding by two on each
 162 side.
- 163 3. 3 CNN layers of kernel 3, input/output 7 channels. Padding by 1 on each
 164 side.
- 165 4. 2D Max-pooling layer of kernel 2, stride 2. Effectively reduces input size by
 166 half.
- 167 5. 1 Fully-connected layer to reduce output to a scalar predicting convergence
 168 rate.

169 Each layer is followed by a ReLU nonlinear activation function. Odd layers are fed
 170 the output of the previous two layers, while even layers are fed the input of only
 171 the previous layer, similarly to a ResNet [8, 6]. C/F splittings are mapped so that
 172 coarse points have value 1 and fine points have value -1 . The CNN thus has a
 173 $25 \times 25 = 625$ -length vector as input. Interestingly, normalization of the convergence
 174 rates is unneeded as the minimum and maximum recorded values are already close to

175 0 and 1, respectively.

176 **2.5. Convection-Diffusion Graph Convolutional Network (GCN).** The
 177 main downside of using traditional convolutional layers is that they are useful only on
 178 structured, grid-like inputs. Grid-based convolution is ineffective on the more complex
 179 meshes that may arise from finite-element discretizations. By treating the input mesh
 180 as a graph and using graph-based convolution techniques, we may get around this
 181 roadblock. The first graphnet that was tried uses the GCN layer introduced by Kipf
 182 and Welling [9] and is described here, the other uses an Edge-Conditioned Convolution
 183 (ECC)/message passing layer and is described in Section 2.6.

184 For the graphnets, the sparse FEM system was re-interpreted as a graph by
 185 assigning each row/column to be a node and defining connectivity between nodes as
 186 nonzero matrix entries. Edge weights were taken to be the entry values themselves,
 187 although normalized to be within the range of $[0, 1]$. This network was trained on the
 188 variably-sized dataset; as will be seen, the lack of any fully-connected layers allows
 189 for any arbitrary input.

190 The Graph Convolutional Network (GCN) layer is defined as an operator on the
 191 graph Laplacian, using the following propagation rule:

192 (2.7)
$$\mathbf{H}^{(i)} = \sigma \left(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(i-1)} \mathbf{W}^{(i)} + \mathbf{b}^{(i)} \right).$$

193 With $\mathbf{H}^{(i)}$ being the i 'th hidden layer, $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ the adjacency matrix with added
 194 self loops, $\tilde{\mathbf{d}}_{ii} = \sum_j \tilde{a}_{ij}$ a diagonal matrix consisting of the sum of outgoing edge
 195 weights, $\mathbf{W}^{(i)}$ the weight matrix at layer i , $\mathbf{b}^{(i)}$ a learned bias vector with same shape
 196 as $\mathbf{H}^{(i)}$, and $\sigma(\cdot)$ some nonlinear activation function. Hidden layers have dimensions
 197 $n \times f$, with n the number of graph nodes and f the number of features at each node.
 198 Thus, using only GCN layers it is not possible to reduce the number of rows in hidden
 199 layers (or the input layer for that matter).

200 Assume $\sigma(\cdot) = \text{ReLU}(\cdot) = \max\{\cdot, 0\}$ for all activation functions. Let \mathbf{X} be the
 201 $n \times 1$ vector containing C/F splitting values for each node. The network that was
 202 trained for this application consisted of the following architecture:

$$\begin{aligned} \mathbf{H}^{(1)} &= \sigma \left(\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \mathbf{X} \mathbf{W}^{(1)} + \mathbf{b}^{(1)} \right) \in \mathbb{R}^{n \times 2} \\ \mathbf{H}^{(2)} &= \sigma \left(\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \mathbf{H}^{(1)} \mathbf{W}^{(2)} + \mathbf{b}^{(2)} \right) \in \mathbb{R}^{n \times 3} \\ \mathbf{H}^{(3)} &= \sigma \left(\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \mathbf{H}^{(2)} \mathbf{W}^{(3)} + \mathbf{b}^{(3)} \right) \in \mathbb{R}^{n \times 2} \\ \mathbf{H}^{(4)} &= \sigma \left(\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \mathbf{H}^{(3)} \mathbf{W}^{(4)} + \mathbf{b}^{(4)} \right) \in \mathbb{R}^{n \times 1} \end{aligned}$$

$$\mathbf{R} = \begin{bmatrix} \dots & \mathbf{X}^T & \dots \\ \dots & \left(\mathbf{H}^{(1)} \right)^T & \dots \\ \dots & \left(\mathbf{H}^{(2)} \right)^T & \dots \\ \dots & \left(\mathbf{H}^{(3)} \right)^T & \dots \\ \dots & \left(\mathbf{H}^{(4)} \right)^T & \dots \end{bmatrix} \in \mathbb{R}^{9 \times n}$$

203
$$y = \sum_{j=1}^n \sigma \left(\sigma \left(r_j \mathbf{W}^{(5)} + \mathbf{b}^{(5)} \right) \mathbf{W}^{(6)} + \mathbf{b}^{(6)} \right) \in \mathbb{R}$$

204

Model	Metric	Dataset	Value
Jacobi Weight	MSE	Training	1.8331×10^{-3}
Jacobi Weight	MSE	Testing	1.8396×10^{-3}
Jacobi Weight	L1	Training	2.9254×10^{-2}
Jacobi Weight	L1	Testing	2.9271×10^{-2}
Convergence Factor	MSE	Training	1.4839×10^{-3}
Convergence Factor	MSE	Testing	1.5171×10^{-3}
Convergence Factor	L1	Training	2.3908×10^{-2}
Convergence Factor	L1	Testing	2.3865×10^{-2}

Table 3.1: Final training and testing Mean Squared Error (MSE)/ L^1 -norm loss values for the two Poisson models. Lower values correspond to higher model accuracy.

205 Note that the matrix \mathbf{R} is formed whose columns are the propagation history
 206 of each node; this is an attempt to emulate traditional ResNet architectures. The
 207 final scalar output is computed by summing each column of \mathbf{R} through a two layer
 208 fully-connected neural network, with the first layer taking as input 9 features and
 209 outputting 5 features, and the second layer taking as input 5 features and outputting
 210 1 feature. This local transformation of each nodal value following by a global averaging
 211 removes any dependency for a fixed input size – instead any sized graph and splitting
 212 could (in theory) be used.

213 **2.6. Convection-Diffusion Message Passing Network (MPNN).** A down-
 214 side of the GCN layer, while being simple to implement and understand, is that it
 215 does not effectively learn edge features between nodes and thus is only an approxima-
 216 tion of a convolution on a graph. The Edge-Conditioned Convolution (ECC) layer,
 217 as defined by Simonovsky and Komodakis [14], attempts to combat this by learning
 218 optimal edge weights given an arbitrary node-edge neighborhood. This approach is
 219 also referred to as a “message-passing network”, as each connected neighbor attempts
 220 to learn a “message” that is passed to the original node [5]. The convolution operation
 221 is defined as

$$222 \quad (2.8) \quad \left(\mathbf{H}^{(i)}\right)_j = \frac{1}{|\mathcal{N}(j)|} \sum_{k \in \mathcal{N}(j)} F^{(i)}(e_{j,k}) \left(\mathbf{H}^{(i-1)}\right)_k + \mathbf{b}^{(i)}$$

223 where $\mathbf{H}^{(i)} \in \mathbb{R}^{n \times f_i}$ is the i 'th hidden layer, $\mathcal{N}(i)$ is a map returning the neighborhood
 224 of vertex i (including itself), $F^{(i)} : E \mapsto \mathbb{R}^{f_i \times f_{i-1}}$ is some function (in our case a
 225 small fully-connected neural network) that outputs a learned weight matrix given an
 226 edge, and $\mathbf{b}^{(i)}$ is a bias term for layer i . The notation $\left(\mathbf{H}^{(i)}\right)_j$ here denotes the j 'th
 227 row of the i 'th hidden layer.

228 Using the ECC layer, another graph net was trained to predict the convergence
 229 rate given some C/F mesh splitting. The architecture is identical to that in Section
 230 2.5 with the exception that the propagation layers $\mathbf{H}^{(i)} : i = 1 \dots 4$ are replaced
 231 with (2.8) and surrounded with a ReLU activation layer. Each $F^{(i)}$ is a two-layer
 232 fully-connected network taking some edge $e \in \mathbb{R}^{f_i}$ and outputting a weight matrix
 233 $\Theta \in \mathbb{R}^{f_i \times f_{i-1}}$.

234 3. Numerical Results.

235 **3.1. 1D Poisson.** A total of 336,000 C/F “grid” splittings were randomly gener-
 236 erated for the variable-coefficient Poisson model. Of these, a random 85% of data
 237 items were designated as a training set and the remaining 15% reserved for a testing
 238 set. The relaxation and convergence CNNs were trained for a total of 30 epochs each,
 239 where each epoch is a total pass over all training set entries. Models were trained
 240 using an MSE loss, and final loss values for both networks are detailed in Table 3.1.

241 Overall, the resulting data is encouraging and seems to indicate that both conver-
 242 gence factors and relaxation weights for the Poisson problem are predictable using a
 243 convolution-based network. Sample convergence rate predictions on the training and
 244 testing subsets are shown in Figure 3.1; these plots display predicted rates as a func-
 245 tion of the true rates, so error is indicated by deviation from the diagonal of the plots.
 246 It is interesting to visually see that the network has an easier time predicting more
 247 convergent grids ($c < 0.6$ or so) than more degenerate grids. A lack of predictions less
 248 than ≈ 0.1 is a deficiency of the input data not having grids with those conver-
 249 gence values.

250 Predictions for relaxation weights on both training and testing subsets are shown
 251 in Figure 3.2. Like the previous convergence plots, more accurate predictions are
 252 closer to the plot diagonals. There is a positive correlation between the predictions
 253 and the true values, though unlike the convergence rates there is no obvious “split” in
 254 the prediction data. The lack of predictions below 0.1 is likely to also be caused by
 255 the input data.

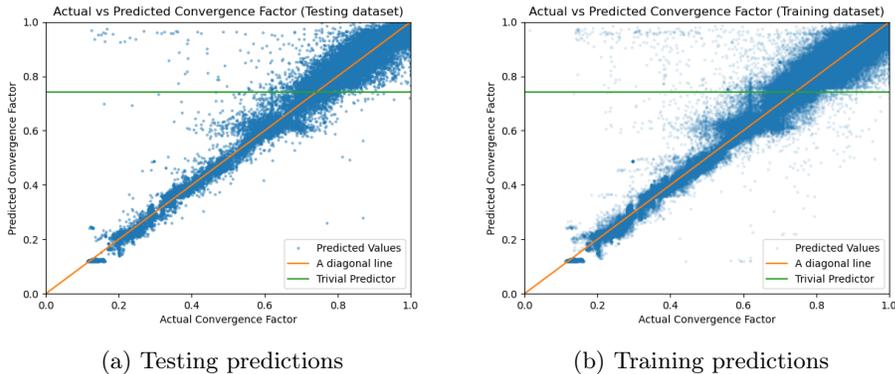


Fig. 3.1: Predicted convergence values vs. true convergence values on (3.1a) test-
 ing and (3.1b) training datasets for Poisson equation. Values closer to the diagonal
 represent more accurate predictions.

256 **3.2. 2D Convection-Diffusion.** Two datasets were created for the 2D recircu-
 257 lating flow problem: a set of *statically-sized* C/F splittings and a set of *variably-sized*
 258 splittings; both datasets consisted of 84,000 elements. The variably-sized dataset was
 259 evenly divided into the four mesh sizes: 15×15 , 25×25 , 35×35 , 50×50 . Each
 260 dataset was then split 85%-15% into training and testing subsets. All neural net-
 261 works were trained for 10 epochs each on an MSE loss, loss history for both datasets
 262 is displayed in Figure 3.3. The recirculating flow CNN was trained and tested on
 263 the static dataset, while the two graph nets were trained on the variable dataset and

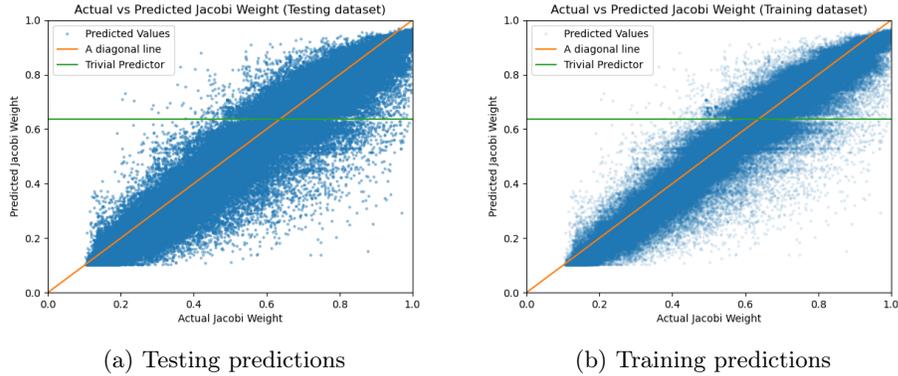


Fig. 3.2: Predicted relaxation weights vs. true relaxation weights on (3.2a) testing and (3.2b) training datasets for Poisson equation. Values closer to the diagonal represent more accurate predictions.

264 evaluated on both datasets.

265 Again, results show that networks are able to learn convergence values for the
 266 particular multigrid solver and recirculating-flow problem. Graph nets and convolu-
 267 tional networks seem to have more-or-less the same predictive power, with the graph
 268 nets having the obvious upside of being cable of handling non-rectangular grid inputs.
 269 Almost expectedly, the message-passing network (ECC) was able to more effectively
 270 learn the problem space than the GCN network, likely due to it learning edge weight-
 271 ings as well.

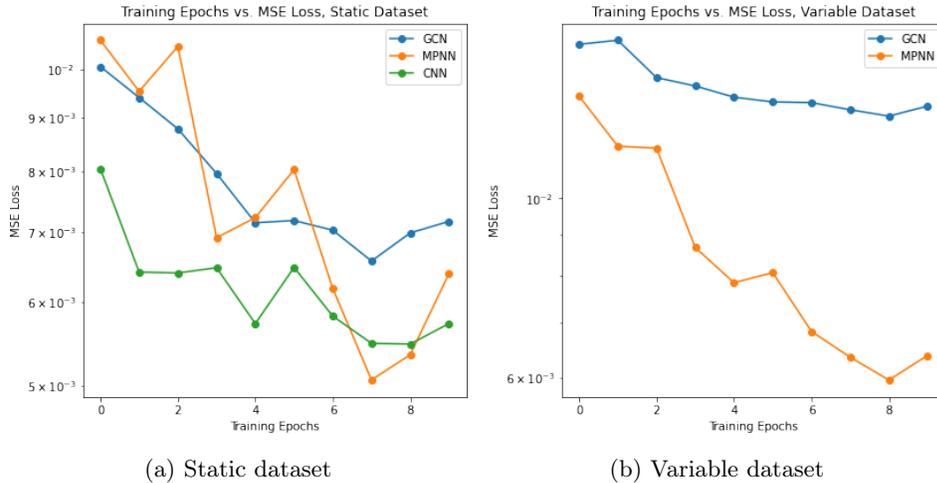


Fig. 3.3: MSE loss history on (3.3a) static and (3.3b) variable datasets. Network types are differentiated between each other and lower values correspond with lower prediction error against the entire dataset.

272 Predicted convergence values by the CNN over the static dataset are displayed
 273 in Figure 3.4. The subplots detail predicted convergence rates vs their true value for
 274 each C/F splitting. Almost opposite to the phenomenon observed in the Poisson CNN
 275 (Figure 3.1, the network predicts more poorly converging grids accurately. This could
 276 perhaps be explained by the data more densely populated with the poor convergence
 277 grids, and has a distinct lack of optimal grids.

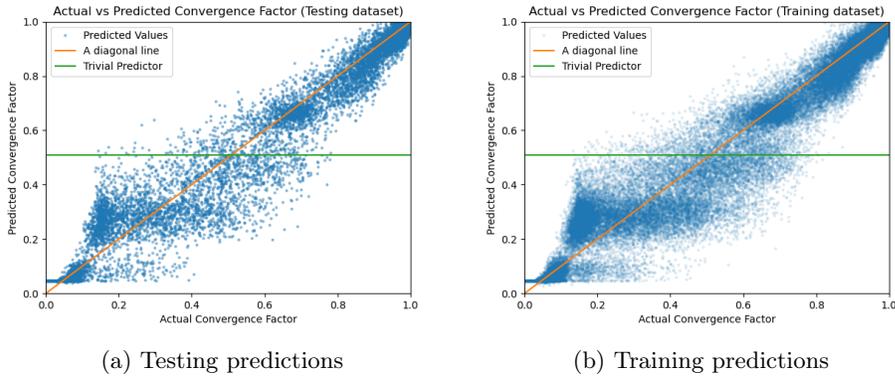


Fig. 3.4: Predicted convergence rates vs. true convergence rates on (3.4a) testing and (3.4b) training datasets for model convection-diffusion problem, using a convolutional network. Values closer to the diagonal represent more accurate predictions.

278 Convergence predictions by the ECC network over the variable-size dataset are
 279 given in Figure 3.5. Various horizontal trends of predictions are visible in the plots,
 280 with distinct subsets of convergence factors all being predicted as the same value.

281 **4. Discussion and Conclusions.** In this paper, a method for training convo-
 282 lutional and graph nets was presented to predict convergence of a two-level multigrid
 283 solver when given an input C/F mesh splitting. Results indicate that neural networks
 284 can indeed predict convergence for a multigrid solver. Convolutional networks are
 285 effective at learning attributes on structured grid-like 1D or 2D meshes, and graph
 286 nets are able to generalize on variably-sized structured meshes. The graph networks
 287 could be further developed and trained on more general classes of problems, instead of
 288 the specific convection-diffusion problem or the variable-coefficient Poisson equation.

289 Future works could build on the networks trained here and explore their appli-
 290 cations, using them for example to pick between different AMG aggregation methods
 291 to find the most optimal for a specific use case. Another more ambitious future work
 292 could be to use the neural networks in an optimization method to obtain the most
 293 convergent or most work efficient C/F splitting for a problem.

294

REFERENCES

- 295 [1] J. ADLER, H. D. STERCK, S. MACLACHLAN, AND L. OLSON, *Numerical partial differential*
 296 *equations*. Draft, 2020.
 297 [2] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, SIAM, 2000.
 298 [3] L. D. DALCIN, R. R. PAZ, P. A. KLER, AND A. COSIMO, *Parallel distributed computing using*
 299 *Python*, *Advances in Water Resources*, 34 (2011), pp. 1124–1139, <https://doi.org/http://>

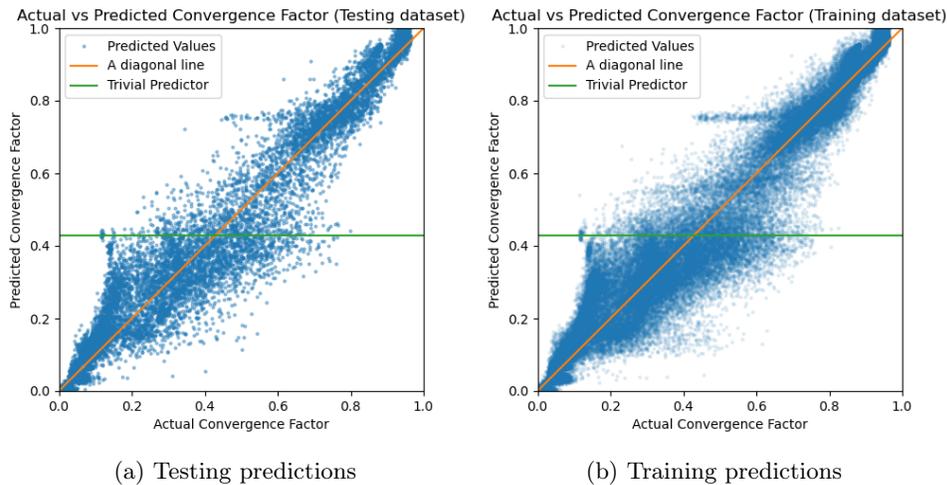


Fig. 3.5: Predicted convergence rates vs. true convergence rates on (3.5a) testing and (3.5b) training datasets for model convection-diffusion problem, using an Edge-Conditioned Convolution network. Values closer to the diagonal represent more accurate predictions.

- 300 [//dx.doi.org/10.1016/j.advwatres.2011.04.013](https://dx.doi.org/10.1016/j.advwatres.2011.04.013). New Computational Methods and Software
 301 Tools.
- 302 [4] H. ELMAN, D. SILVESTER, AND A. WATHEN, *Finite Elements and Fast Iterative Solvers With*
 303 *Applications in Incompressible Fluid Dynamics*, Oxford University Press, 2014.
- 304 [5] J. GILMER, S. S. SCHOENHOLZ, P. F. RILEY, O. VINYALS, AND G. E. DAHL, *Neural message*
 305 *passing for quantum chemistry*, 2017, <https://arxiv.org/abs/1704.01212>.
- 306 [6] K. HE, X. ZHANG, S. REN, AND J. SUN, *Deep residual learning for image recognition*, 2015,
 307 <https://arxiv.org/abs/1512.03385>.
- 308 [7] M. HOMOLYA AND D. A. HAM, *A parallel edge orientation algorithm for quadrilateral meshes*,
 309 *SIAM Journal on Scientific Computing*, 38 (2016), pp. S48–S61, [https://doi.org/10.1137/](https://doi.org/10.1137/15M1021325)
 310 [15M1021325](https://doi.org/10.1137/15M1021325), <http://arxiv.org/abs/1505.03357>, <https://arxiv.org/abs/1505.03357>.
- 311 [8] G. HUANG, Z. LIU, L. VAN DER MAATEN, AND K. Q. WEINBERGER, *Densely connected*
 312 *convolutional networks*, 2018, <https://arxiv.org/abs/1608.06993>.
- 313 [9] T. N. KIPF AND M. WELLING, *Semi-supervised classification with graph convolutional net-*
 314 *works*, 2017, <https://arxiv.org/abs/1609.02907>.
- 315 [10] A. T. T. McRAE, G.-T. BERCEA, L. MITCHELL, D. A. HAM, AND C. J. COTTER, *Automated*
 316 *generation and symbolic manipulation of tensor product finite elements*, *SIAM Journal*
 317 *on Scientific Computing*, 38 (2016), pp. S25–S47, <https://doi.org/10.1137/15M1021167>,
 318 <http://arxiv.org/abs/1411.2940>, <https://arxiv.org/abs/1411.2940>.
- 319 [11] L. N. OLSON AND J. B. SCHRODER, *Pyamg: Algebraic multigrid solvers in python v4.0*, 2018,
 320 <https://github.com/pyamg/pyamg>. Release 4.0.
- 321 [12] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. T. McRAE,
 322 G.-T. BERCEA, G. R. MARKALL, AND P. H. J. KELLY, *Firedrake: automating the*
 323 *finite element method by composing abstractions*, *ACM Trans. Math. Softw.*, 43 (2016),
 324 pp. 24:1–24:27, <https://doi.org/10.1145/2998441>, <http://arxiv.org/abs/1501.01809>, <https://arxiv.org/abs/1501.01809>.
- 325 [13] J. RUGE AND K. STEUBEN, *Multigrid Methods*, SIAM, 1987, ch. 4. Algebraic Multigrid.
- 326 [14] M. SIMONOVSKY AND N. KOMODAKIS, *Dynamic edge-conditioned filters in convolutional neural*
 327 *networks on graphs*, 2017, <https://arxiv.org/abs/1704.02901>.
- 328