

---

# Loop Unrolled Shallow Equilibrium Regularizer (LUSER) - A Memory-Efficient Inverse Problem Solver

---

Peimeng Guan<sup>1</sup>, Jihui Jin<sup>1</sup>, Justin Romberg<sup>1</sup>, Mark A. Davenport<sup>1</sup>

<sup>1</sup>Department of Electrical and Computer Engineering  
Georgia Institute of Technology  
Atlanta, GA 30332  
guanpeimeng@gatech.edu

## Abstract

In inverse problems we aim to reconstruct some underlying signal of interest from potentially corrupted and often ill-posed measurements. Classical optimization-based techniques proceed by optimizing a data consistency metric together with a regularizer. Current state-of-the-art machine learning approaches draw inspiration from such techniques by unrolling the iterative updates for an optimization-based solver and then learning a regularizer from data. This *loop unrolling* (LU) method has shown tremendous success, but often requires a deep model for the best performance leading to high memory costs during training. Thus, to address the balance between computation cost and network expressiveness, we propose an LU algorithm with shallow equilibrium regularizers (LUSER). These implicit models are as expressive as deeper convolutional networks, but far more memory efficient during training. The proposed method is evaluated on image deblurring, computed tomography (CT), as well as single-coil Magnetic Resonance Imaging (MRI) tasks and shows similar, or even better, performance while requiring up to  $8\times$  less computational resources during training when compared against a more typical LU architecture with feedforward convolutional regularizers.

## 1 Introduction

In an inverse problems we face the task of reconstructing some data or parameters of an unknown signal from indirect observations. The forward process, or the mapping from the data to observations, is typically well known, but ill-posed or non-invertible. More formally, we consider the task of recovering some underlying signal  $\mathbf{x}$  from measurements  $\mathbf{y}$  taken via some forward operator  $\mathbf{A}$  according to

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}, \quad (1)$$

where  $\boldsymbol{\eta}$  represents noise. The forward operator can be nonlinear, but to simplify the notation, we illustrate the idea in linear form throughout this paper. A common approach to recover the signal is via an iterative method based on the least squares loss:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2. \quad (2)$$

For many problems of interest,  $\mathbf{A}$  is ill-posed and does not have full column rank. Thus, attempting to solve (2) does not yield a unique solution. To address this, we can extend (2) by including a regularizing term to bias the inversion towards solutions with favorable properties. Common examples of regularization include  $\ell_2$ ,  $\ell_1$ , and Total Variation (TV). Each regularizer encourages certain properties on the estimated signal  $\hat{\mathbf{x}}$  (e.g., smoothness, sparsity, piece-wise constant, etc.) and is often chosen based on task-specific prior knowledge.

Recent works [1] attempt to tackle inverse problems using more data-driven methods. Unlike typical supervised learning tasks that attempt to learn a mapping purely from examples, deep learning for inverse problems have access to the forward operator and thus should be able to guide the learning process for more accurate reconstructions. One popular approach to incorporating knowledge of the forward operator is termed *loop unrolling* (LU). These methods are heavily inspired by standard iterative inverse problem solvers, but rather than use a hand tuned regularizer, they instead learn the update with some parameterized model. They tend to have a fixed number of iterations (typically around 5-10) due to computational constraints. [2] proposes an interesting alternative that takes advantage of *deep equilibrium* (DEQ) models [3–6] that we refer to as DEQ4IP. Equilibrium models are designed to recursively iterate on their input until a “fixed point” is found (i.e., the input no longer changes after passing through the model). They extend this principle to the LU method, choosing to iterate until convergence rather than for a “fixed budget”.

**Our Contributions.** We propose an alternative novel architecture for solving inverse problems called *Loop Unrolled Shallow Equilibrium Regularizer* (LUSER). It incorporates knowledge of the forward model by adopting the principles of LU architectures while reducing its memory consumption by using a shallow (relative to existing feed-forward models) DEQ as the learned regularizer update. Unlike DEQ4IP that converts the entire LU architecture into a DEQ model, we only convert the learned regularizer at each stage. This has the advantage of simplifying the learning task for DEQ models, which can be unstable to train in practice. To our knowledge, this is the first use of multiple sequential DEQ models within a single architecture for solving inverse problems. Our proposed architecture (i) reduces the number of forward/adjoint operations compared to the work proposed by [2], (ii) reduces the memory footprint during training without loss of expressiveness as demonstrated by our experiments, and (iii) is more stable to train in practice than DEQ4IP. We empirically demonstrate better reconstruction across multiple tasks than state-of-the-art LU alternatives with a similar number of parameters, with the ability to reduce computational memory costs during training by a factor of up to  $8\times$ .

The remainder of the paper is organized as follows. Section 2 reviews related works in solving inverse problems. Section 3 introduces the proposed LUSER, which we compare with other baseline methods in image deblurring, CT, and MRI tasks in Section 4. We conclude in Section 5 with a brief discussion.

## 2 Related Work

### 2.1 Loop Unrolling

As noted above, a common approach to tackling an inverse problem is to cast it as an optimization problem consisting of the sum of a data consistency term and a regularization term

$$\min_x \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \gamma r(\mathbf{x}), \quad (3)$$

where  $r$  is a regularization function mapping from the domain of the parameters of interest to a real number and  $\gamma \geq 0$  is a well-tuned coefficient. The regularization function is chosen for specific classes of signals to exploit any potential structure, e.g.,  $\|\mathbf{x}\|_2$  for smooth signals and  $\|\mathbf{x}\|_0$  or  $\|\mathbf{x}\|_1$  for sparse signals. The total-variation (TV) norm is another popular example of a regularizer that promotes smoothness while preserving edges and is often used for imaging tasks.

When  $r$  is differentiable, the solution of (3) can be obtained in an iterative fashion via gradient descent. For some step size  $\lambda$  at iteration  $k = 1, 2, \dots, K$ , we apply the update:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \lambda \mathbf{A}^\top (\mathbf{y} - \mathbf{A}\mathbf{x}_k) - \lambda \gamma \nabla r(\mathbf{x}_k). \quad (4)$$

For non-differentiable  $r$ , the more generalized *proximal gradient* algorithm can be applied with the following update, where  $\tau$  is a well-tuned hyperparameter related to the proximal operator:

$$\mathbf{x}_{k+1} = \text{prox}_{\tau, r}(\mathbf{x}_k + \lambda \mathbf{A}^\top (\mathbf{y} - \mathbf{A}\mathbf{x}_k)). \quad (5)$$

The *loop unrolling* (LU) method considers running the update in (4) or (5), but replaces  $\lambda \gamma \nabla r$  or the proximal operator with a learned neural network instead. The overall architecture repeats the neural network based update for a pre-determined number of iterations, fixing the overall computational budget. Note that the network is only implicitly learning the regularizer. In practice, it is actually

learning an update step, which can be thought of as de-noising or a projection onto the manifold of the data. LU is typically trained end-to-end, i.e., when fed some initialization  $\mathbf{x}_0$ , the network will output the final estimate  $\mathbf{x}_K$ , compute a loss with respect to the ground truth, and then back-propagate across the entire computational graph to update network parameters. While end-to-end training is easier to perform and encourages faster convergence, it requires all intermediate activations to be stored in memory. Thus, the maximum number of iterations is always kept small compared to classical iterative inverse problem solvers.

Due to the limitation in memory, there is a trade-off between the depth of a LU and the richness of each regularization network. Intuitively, one can raise the network performance by increasing the number of loop unrolled iterations. For example, [2] extends the LU model to potentially infinite number of iterations using an implicit network, and [7] allows deeper unrolling iterations using invertible networks, while requiring recalculation of the intermediate results from the output in training phase. This approach can be computationally intensive for large-scale inverse problems or when the forward operator is nonlinear and computationally expensive to apply. For example, the forward operator may involve solving differential equations such as the wave equation for seismic wave propagation [8] and the Lorenz equations for atmospheric modeling[9].

Alternatively, one can design a richer regularization network. For example [10] uses a transformer as the regularization network and achieves extremely competitive results in the fastMRI challenge [11], but requires multiple 24GB GPU for training with batch size of 1, which is often impractical, especially for large systems. Our design strikes a balance in the expressiveness in regularization networks and memory efficiency during training. Our proposed work is an alternative method to achieve a rich regularization network without the additional computational memory costs during training.

## 2.2 Deep Equilibrium Models for Inverse Problems (DEQ4IP)

*Deep equilibrium* (DEQ) models introduce an alternative to traditional feed-forward networks [3–6]. Rather than feed an input through a fixed (relatively small) number of layers, DEQ models solve for the “fixed-point” given some input. More formally, given a network  $f_\theta$  and some input  $\mathbf{x}^{(0)}$  and  $\mathbf{y}$ , we recursively apply the network via

$$\mathbf{x}^{(k+1)} = f_\theta(\mathbf{x}^{(k)}, \mathbf{y}), \quad (6)$$

until convergence.<sup>1</sup> In this instance,  $\mathbf{y}$  acts as an input injection that determines the final output. This is termed the forward process. The weights  $\theta$  of the model can be trained via implicit differentiation, removing the need to store all intermediate activations from recursively applying the network. This allows for deeper, more expressive models without the associated memory footprint to train such models.

[2] demonstrates an application of one such model, applying similar principles to a single iteration of an LU architecture. Such an idea is a natural extension as it allows the model to “iterate until convergence” rather than rely on a “fixed budget”. In another word, looping over (4) and (5) many times (in practice, usually around 50 iterations) until  $\mathbf{x}_k$  converges. However, such a model can be unstable to train and often performs best with pre-training of the learned portion of the model (typically acting as a learned regularizer/de-noiser). It is also important to note is that such a model would have to apply the forward operator (and potentially the adjoint) many times during the forward process. Although this can be accelerated to reduce the number of applications, it is still often more than the number of applications for an LU equivalent. This can be an issue if the forward operator is computationally expensive to apply, an issue LU methods somewhat mitigate by fixing the total number of iterations.

## 2.3 Alternative Approaches to Tackle Memory Issues

I-RIM [7] is a deep invertible network that address the memory issue by recalculating the intermediate results from the output. However it is not ideal when forward model is computationally expensive. Gradient checkpointing [12] is another practical technique to reduce memory costs for deep neural

<sup>1</sup>Note that, since our approach will ultimately use both methods, to aid in a clearer presentation we use subscript, i.e.,  $\mathbf{x}_k$ , to denote the LU iterations, and superscript with parenthesis, i.e.,  $\mathbf{r}^{(i)}$ , to denote the iterations in the deep equilibrium model.

networks. It saves intermediate activations of some checkpoint nodes, and recomputes the forward pass between two checkpoints for backpropagation. However, it is not an easy and efficient technique to implement for a weight-tied neural network.

### 3 Methodology

LU methods have been the state-of-the-art approach for solving inverse problems due to their stability when training, inclusion of the forward model, and their near instantaneous inference times. However, there is a noticeable trade-off in terms of the memory requirements when training these models and the accuracy for even medium scale problems, making them nearly impossible to use for much larger scale inverse problems. DEQ4IP offers an interesting alternative, drastically reducing the memory required during training and allowing the flexibility to adjust accuracy during reconstruction when performing inference. However, these models can potentially suffer when the forward/adjoint operators are computationally expensive, particularly when it takes more iterations to converge than a standard fixed LU method.

To address these concerns, we propose a novel architecture for solving inverse problems called *Loop Unrolled Shallow Equilibrium Regularizer* (LUSER) as an approach that limits the number of forward/adjoint calls for particularly complex inverse problems while drastically reducing the memory requirements for training, allowing us to scale up to much larger inverse problems (or require less GPU memory for existing problems). LUSER achieves this by adopting DEQ models as the learned regularizer update in a standard LU architecture. The implicit DEQ models are smaller in size but just as expressive as typical convolutional models used as the learned regularizer update allowing for an accurate reconstruction with less computational memory costs. Furthermore, learning a proximal update is a far simpler task compared to solving the inverse problem as a whole, reducing the instability during training commonly faced by DEQ models.

We adopt a ‘‘proximal gradient descent’’ styled LU architecture. The network takes in measurements  $\mathbf{y}$  and some initial estimate  $\mathbf{x}_0$ . The architecture consists of  $K$  stages, alternating between a gradient descent step  $\mathbf{d}_k = \mathbf{x}_k + \lambda \mathbf{A}^\top (\mathbf{y} - \mathbf{A}\mathbf{x}_k)$  for  $k = 1, 2, \dots, K$ , followed by a feed-forward pass of the shallow equilibrium model acting as a proximal update block.

Figure 1 illustrates a single block in LUSER as a learned proximal operator. The input  $\mathbf{d}_k$  from the previous stage is processed only once by a set of input injection layers to avoid redundant computation. The input injection will determine the final fixed point output [5]. The recursive portion of the proximal block consists of two sets of layers: the data layer and the mixing layer. The current estimate of the fixed point is first passed through the data layer before being concatenated with the input injection and processed by the mixing layer. This process repeats until a fixed point is found. Although Figure 1 shows the simplest way of finding the fixed point, in practice acceleration or other fixed point solvers are applied to solve for the fixed point. We apply Anderson acceleration [13] for all of our models.




Figure 1: A proximal block in LUSER

At each stage, the shallow equilibrium regularizer attempts to output its best estimate of the ground truth  $\mathbf{x}^*$ . We introduce a skip connection between the input  $\mathbf{d}_k$  and the final output so that the regularizer is learning the residual  $\mathbf{r}$  between the input  $\mathbf{d}_k$  and the ground truth  $\mathbf{x}^*$  instead. When  $\mathbf{x}_k$  converges to  $\mathbf{x}^*$ ,  $\mathbf{d}_k$  will also converge to  $\mathbf{x}^*$  and we expect the residual to be closer and closer to the zero vector. Therefore, we initialize the input of data layer with a zero vector with the same dimensions as the input  $\mathbf{d}_k$  in the hopes that fewer iterations will be needed at later stages as the current estimate  $\mathbf{x}_k$  converges towards the ground truth. Let  $\mathbf{r}^{(i)}$  denote the  $i^{th}$  update of the residual,

where  $\mathbf{r}^{(0)} = \mathbf{0}$  and  $\mathbf{r}^{(\infty)}$  be the fixed point solution of the residual. Let  $\oplus$  denote the concatenation operator. The process in one loop unrolled block can be formulated as the following:

$$\mathbf{x}_{k+1} = \mathbf{d}_k + \text{MixLayer} \left( \text{DataLayer}(\mathbf{r}^{(\infty)}) \oplus \text{InjectLayer}(\mathbf{d}_k) \right). \quad (7)$$

Because of the fixed point solver within each regularization block, LUSER uses fewer layers to achieve a similar level of performance, thus we call LUSER a shallow equilibrium regularizer.

We also explore two variants of LUSER, dubbed LUSER-SW and LUSER-DW. LUSER-SW refers to a shared-weight version of the proposed algorithm, where the proximal operator in all loop unrolled iterations are weight-tied and thus identical. Theoretically, the regularizer should be able to handle any input regardless of the iteration step. However, in practice, the distribution of intermediate reconstructions can be vastly different. Thus, training a single model to handle all these instances can be a daunting task, leading to poor generalization across the different stages. Since the total computational budget is fixed, one approach can be to use different weights (DW) for the learned proximal operator at each stage to handle the potentially different distributions. This will increase the number of parameters that need to be trained and stored, but since LUSER already has so few parameters to begin with, expanding to the different weight variant is still feasible. This paper will compare two variants of LUSER with other architectures in different tasks.

DEQ4IP relies on the learned proximal operator to have the same weights, thus we cannot include a different weight variant for comparison. On the other hand, the same principles can be applied to the LU variant of DEQ4IP. However, we aim to only compare models of similar number of parameters (or less in the case of LUSER-SW), and thus restrict our attentions to the shared weight variant of LU only.

## 4 Experiments and Results

In this section, we compare our proposed networks to LU with DnCNN as proximal operator (LU-DnCNN) and DEQ4IP on three different tasks: image deblurring, computed tomography (CT), and single-coil accelerated Magnetic Resonance Imaging (MRI). The datasets we use are RGB CelebA [14], LoDoPaB-CT [15], and single-coil knee data from fastMRI [11] respectively. We also experiment with incorporating an auxiliary loss of including an MSE loss on intermediate reconstructions with the ground-truth instead of just the final output. This is done for the LU models (LU-DnCNN, LUSER).

### 4.1 Experimental Setup

In the image deblurring task, the blurry image is obtained by applying a  $(9 \times 9)$  Gaussian kernel with variance of 5 to an image with additive white Gaussian noise with a standard deviation of 0.01. If the image is RGB, the same kernel is applied to all channels. For accelerated MRI tasks, measurements in k-space (or frequency domain) are often subsampled due to the cost in measurement. The goal of MRI reconstruction is to recover the underlying physical structure from subsampled noisy measurements. We simulate the forward operator  $\mathbf{A}$  with a 2-dimensional Fourier transform with randomly selected rows. We consider two common subsampling scenarios: 4x and 8x acceleration, or subsampling the columns in full measurement by a factor of 4 and 8 respectively. For the CT task, the forward operator is a Radon transform, and we uniformly select 200 out of 1000 angles in measurements. The adjoint of measurement  $\mathbf{A}^\top \mathbf{y}$  is used as initialization for MRI and CT tasks, which brings the measurement back to the signal domain. However, in the deblurring task, since the measurement lies in a same domain as the underlying clean image,  $\mathbf{y}$  is used as the initial guess. Notice that some works use the filtered backprojection as the initialization for CT, such as [16], but we use the adjoint for the purpose of consistency.

We fix the budget of LU-DnCNN and LUSER to be a total of 8 iterations, while as we allow DEQ4IP to iterate until it reaches a fixed point. We use a DnCNN adopted from [17] with 17 convolutional layers with 64 channels, followed by BatchNorm and ReLU activations for the regularizer for LU-DnCNN and DEQ4IP. For the learned regularizer update in LUSER, we use 2 convolutional layers each for the input injection layer and data layer. The mixing layer contains 3 convolutional layers.

In order to stabilize training for DEQ inspired models, we wrap all convolutional layers in DEQ4IP and LUSER with Spectral Norm (SN) [18]. We list more details for the learned proximal network for

Table 1: Architecture of Proximal Network in LUSER

	Layer Details
Injection Layer	SN(conv( $C_{in}$ :1, $C_{out}$ :32, ks:3, pad:1)) + LeakyReLU SN(conv( $C_{in}$ :32, $C_{out}$ :32, ks:3, pad:1)) + LeakyReLU
Data Layer	SN(conv( $C_{in}$ :1, $C_{out}$ :32, ks:3, pad:1)) GN + LeakyReLU SN(conv( $C_{in}$ :1, $C_{out}$ :32, ks:3, pad:1)) GN + LeakyReLU
Mix Layer	SN(conv( $C_{in}$ :64, $C_{out}$ :64, ks:3, pad:1)) GN + LeakyReLU SN(conv( $C_{in}$ :64, $C_{out}$ :64, ks:3, pad:1)) GN + LeakyReLU SN(conv( $C_{in}$ :64, $C_{out}$ :1, ks:3, pad:1))

LUSER in Table 1 for the case when the input has a single channel.  $C_{in}$  and  $C_{out}$  denote the input and output channels, ks refers to the kernel size of a convolutional layer, pad denotes the padding in 2-dimension, and GN stands for GroupNorm.

We use two metrics to evaluate the quality of reconstruction: Peak Signal-to-Noise Ratio (PSNR) in dB and the Structural Similarity Index (SSIM). Note that although we use the same models as [2], we train our models from scratch and report lower values on the MRI task. We suspect that this is due to evaluating with a single channel only. When we include the imaginary channel (for a total of 2 channels), the metrics we recorded are more aligned with those reported in [2]. All models are trained with a single RTX6000 24GB GPU.

## 4.2 Main Results

Table 2 compares the average PSNR and SSIM. The different weight version of LUSER outperforms LU-DnCNN with a similar number of network parameters. LUSER-SW achieves similar level of performance in most tasks with only 5 layers, versus 17 layers in LU-DnCNN. DEQ4IP attains the best performance in image delurring task, but it requires repetitive computation of forward/adjoint operators. Training shared-weight architectures with auxillary losses improves the reconstruction quality in most tasks. Figure 2 shows representative reconstruction results. The overall performances are similar, but LUSER-DW attains better qualities in detailed structures, especially compared to LU-DnCNN. The areas with improvements are emphasized with red boxes in the ground truth images.

Table 2: Average PSNR and SSIM for test set, the best two performances are in bold.

PSNR SSIM	LU-DnCNN		DEQ4IP	LUSER-SW		LUSER-DW	
	Final loss	Aux loss	Final loss	Final loss	Aux loss	Final loss	Aux loss
CelebA	29.93	30.39	<b>31.57</b>	30.30	30.65	<b>31.40</b>	31.15
	0.862	0.871	<b>0.895</b>	0.869	0.878	<b>0.891</b>	0.888
CT	30.59	31.59	<b>31.79</b>	28.82	28.04	<b>31.83</b>	31.66
	0.844	0.859	<b>0.868</b>	0.801	0.797	<b>0.860</b>	0.859
4x MRI	29.01	29.02	29.01	28.82	29.18	<b>29.86</b>	<b>29.37</b>
	0.668	0.671	0.678	0.662	0.685	<b>0.740</b>	<b>0.713</b>
8x MRI	27.50	<b>27.65</b>	27.51	27.42	27.42	<b>28.06</b>	27.55
	0.576	0.572	0.570	0.562	0.560	<b>0.630</b>	<b>0.596</b>

Table 3 compares the network sizes (number of parameters) as well as the maximum training batch sizes for three commonly seen GPU RAM capacities: 8 GB, 10 GB and 24 GB. Batch sizes are recorded with maximum even numbers, except when it is 1 for stochastic gradient descent. Notice that in MRI, because the Fourier transformation is implemented in tensor form, the minimum batch size it can take is 2. We use the batch size as a proxy for the memory requirements during training. Since DEQ4IP is an extension of LU-DnCNN, their networks are of the same size, but implicit DEQ models support larger batch sizes making DEQ4IP far more memory-efficient during training. The advantages of using DEQ models for LUSER and DEQ4IP are particularly highlighted in the case of limited memory (smaller GPUs). LU-DnCNN is unable to even train for the CT and MRI task with limited memory constraints, while LUSER and DEQ4IP can. This pattern is expected to repeat for more large scale tasks where standard LU architectures will be unable to train at all due to memory requirements. It is important to note that memory requirements depend more on the depth of the




Figure 2: Representative reconstructions, where each row represents one task. The left-most column shows the initialization or input to the networks; the middle three columns show the reconstructions for LU-DnCNN, DEQ4IP and LUSER-DW; and the right-most column shows the underlying true images. Regions corresponding to qualitative improvements are emphasized in red boxes in the last column.

Table 3: Comparison of network sizes and maximum possible batch sizes during training. Entries with ‘-’ denote that the architecture with a batch size of 1 cannot fit into a particular GPU RAM capacity.

Max Batch Size	GPU RAM	LU-DnCNN	DEQ4IP	LUSER-SW	LUSER-DW
	Capacity	(17 layers)	(17 layers)	(5 layers)	(5 layers)
deblurring	#Params	558,580	558,580	96,503	770,073
	8 GB	1	16	10	10
	10 GB	1	16	14	14
	24 GB	4	68	34	34
CT	#Params	556,033	556,033	93,954	751,625
	8 GB	-	4	2	2
	10 GB	-	6	4	4
	24 GB	2	20	10	10
MRI	#Params	557,185	557,185	95,107	760,849
	8 GB	-	4	4	4
	10 GB	2	6	4	4
	24 GB	4	16	12	12

network than the number of parameters. For example, even though LUSER-DW and LUSER-SW have different numbers of parameters, they share the same architecture/depth and thus use roughly the same amount of memory during training.

We observe that LUSER takes longer to train in each epoch than LU-DnCNN, but the training converges in fewer iterations than LU-DnCNN. The total training time for all methods are roughly the same for the image deblurring and MRI tasks, but DEQ4IP takes longer to train on the CT problem as it is finding a fixed-point with a more complicated forward operator. During evaluation, LU-DnCNN is the fastest, and the DEQ4IP and LUSER are in similar speed. When the forward

operator is computationally intensive, LUSER evaluates faster than DEQ4IP. Table 4 compares the various properties among the three architectures.

Table 4: Method Comparisons

	LU	DEQ4IP	LUSER
Training Time	Fast	Slow	Moderate
Evaluation Time	Fast	Slow	Moderate
Training Stability	Stable	Unstable	Stable
Network Size	Large	Large	Small
Expressiveness	Low	Moderate	High
Training Memory Usage	Large	Small	Moderate

## 5 Conclusion

Loop unrolling architectures with deep convolutional layers as the learned regularizer update are a popular approach for solving inverse problems. Although its variants achieve state-of-the-art results across a variety of tasks, LU algorithms incur a huge memory cost when training due to the requirement of saving all intermediate activations, sometimes even requiring multiple GPUs to train on complex tasks. DEQ4IP offers an interesting alternative via extending LU to infinitely many layers by finding a fixed point solution, but can be impractical when the forward/adjoint operators are nonlinear or larger in scale. To address the memory issue, we proposed two variants of loop unrolling architectures with deep equilibrium models as the learned regularizer updates, LUSER-SW and LUSER-DW. We verify the memory savings (by comparing batch sizes) relative to loop unrolling algorithms with a DnCNN model. Across all tasks, LUSER-DW outperforms LU-DnCNN with a similar number of network parameters, while reducing the memory requirements by a factor of 5 or more. LUSER offers a path forward for large-scale, complex, non-linear inverse problems that are currently infeasible to train with existing methods.

## Acknowledgment

Guan and Davenport are supported by the Center for Energy and Geo Processing (CeGP) at Georgia Tech and the King Fahd University of Petroleum and Minerals. Jin and Romberg are supported in part by the Office of Naval Research Task Force Ocean under Grant No. N00014-19-1-2639.

## References

- [1] Gregory Ongie, Ajil Jalal, Christopher A Metzler, Richard G Baraniuk, Alexandros G Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020.
- [2] Davis Gilton, Gregory Ongie, and Rebecca Willett. Deep equilibrium architectures for inverse problems in imaging. *IEEE Transactions on Computational Imaging*, 7:1123–1133, 2021.
- [3] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. *Advances in Neural Information Processing Systems*, 32, 2019.
- [4] Shaojie Bai, Vladlen Koltun, and J Zico Kolter. Multiscale deep equilibrium models. *Advances in Neural Information Processing Systems*, 33:5238–5250, 2020.
- [5] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley J Osher, and Wotao Yin. Fixed point networks: Implicit depth models with jacobian-free backprop. 2021.
- [6] Laurent El Ghaoui, Fangda Gu, Bertrand Travacca, Armin Askari, and Alicia Tsai. Implicit deep learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021.
- [7] Patrick Putzky and Max Welling. Invert to learn to invert, 2019.



- [8] Chris Chapman. *Fundamentals of seismic wave propagation*. Cambridge university press, 2004.
- [9] Christian Oestreicher. A history of chaos theory. *Dialogues in clinical neuroscience*, 2022.
- [10] Zalan Fabian and Mahdi Soltanolkotabi. Humus-net: Hybrid unrolled multi-scale network architecture for accelerated mri reconstruction. *arXiv preprint arXiv:2203.08213*, 2022.
- [11] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdal, Adriana Romero, Michael Rabbat, Pascal Vincent, Nafissa Yakubova, James Pinkerton, Duo Wang, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastmri: An open dataset and benchmarks for accelerated mri, 2018.
- [12] Nimit Sharad Sohoni, Christopher Richard Aberger, Megan Leszczynski, Jian Zhang, and Christopher Ré. Low-memory neural network training: A technical report. *arXiv preprint arXiv:1904.10631*, 2019.
- [13] Homer F Walker and Peng Ni. Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis*, 49(4):1715–1735, 2011.
- [14] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [15] Johannes Leuschner, Maximilian Schmidt, Daniel Otero Baguer, and Peter Maass. Lodopabct, a benchmark dataset for low-dose computed tomography reconstruction. *Scientific Data*, 8(1):1–12, 2021.
- [16] AmirEhsan Khorashadizadeh, Konik Kothari, Leonardo Salsi, Ali Aghababaei Harandi, Maarten de Hoop, and Ivan Dokmani'c. Conditional injective flows for bayesian imaging. *arXiv preprint arXiv:2204.07664*, 2022.
- [17] Ernest Ryu, Jialin Liu, Sicheng Wang, Xiaohan Chen, Zhangyang Wang, and Wotao Yin. Plug-and-play methods provably converge with properly trained denoisers. In *International Conference on Machine Learning*, pages 5546–5557. PMLR, 2019.
- [18] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.