

Parameters estimation via unrolling for the PNP algorithm

Giuseppe Scarlato¹

Joint work with
Marco Donatelli¹

July 2, 2026

¹Università degli Studi dell'Insubria

MMIPU 2026 workshop
ICCSA 2026, University of Minho, Braga, Portugal



PNPD algorithm

Unrolling architecture

Training

Numerical results

PNPD algorithm

Consider a linear inverse problem of the form

$$b^\delta = Au + \eta_\delta = b + \eta_\delta, \quad (1)$$

where

- $b^\delta \in \mathbb{R}^s$ is the observed data corrupted by noise $\eta_\delta \in \mathbb{R}^s$
- $A \in \mathbb{R}^{s \times d}$ is a linear operator
- $u \in \mathbb{R}^d$ is the unknown signal we want to recover

A common variational approach to solve eq. (1) is to consider the following optimization problem

$$\begin{aligned} & \operatorname{argmin}_{u \in \mathbb{R}^d} \frac{1}{2} \|Au - b^\delta\|^2 + h(Wu) \\ & = \operatorname{argmin}_{u \in \mathbb{R}^d} f(u) + h(Wu), \end{aligned} \tag{2}$$

where

- $\|\cdot\|$ denotes the Euclidean norm
- $f(u)$ is the data fidelity term
- $h(Wu)$ is the regularization term
- $W \in \mathbb{R}^{d' \times d}$ is a linear operator
- $h : \mathbb{R}^{d'} \rightarrow \mathbb{R} \cup \{\infty\}$ is a convex and possibly non-smooth function.

Proximal gradient

Proximal gradient methods are first-order methods that alternate a gradient step on f and a proximal evaluation on $h \circ W$:

Proximal gradient

Proximal gradient methods are first-order methods that alternate a gradient step on f and a proximal evaluation on $h \circ W$:

$$\left\{ \begin{array}{l} u_{n+\frac{1}{2}} = u - \alpha_n \nabla f(u) \end{array} \right. \quad \leftarrow \text{gradient descent}$$

Proximal gradient

Proximal gradient methods are first-order methods that alternate a gradient step on f and a proximal evaluation on $h \circ W$:

$$\begin{cases} u_{n+\frac{1}{2}} = u - \alpha_n \nabla f(u) & \leftarrow \text{gradient descent} \\ u_{n+1} = \text{prox}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}) & \leftarrow \text{proximal evaluation} \end{cases} \quad (3)$$

Proximal gradient

Proximal gradient methods are first-order methods that alternate a gradient step on f and a proximal evaluation on $h \circ W$:

$$\begin{cases} u_{n+\frac{1}{2}} = u - \alpha_n \nabla f(u) & \leftarrow \text{gradient descent} \\ u_{n+1} = \text{prox}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}) & \leftarrow \text{proximal evaluation} \end{cases} \quad (3)$$

where

- $\alpha_n > 0$ is the step size,
- $\text{prox}_g(v) = \operatorname{argmin}_u \frac{1}{2} \|u - v\|^2 + g(u)$ is the proximal operator of g .

Proximal gradient

Proximal gradient methods are first-order methods that alternate a gradient step on f and a proximal evaluation on $h \circ W$:

$$\begin{cases} u_{n+\frac{1}{2}} = u - \alpha_n \nabla f(u) & \leftarrow \text{gradient descent} \\ u_{n+1} = \text{prox}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}) & \leftarrow \text{proximal evaluation} \end{cases} \quad (3)$$

where

- $\alpha_n > 0$ is the step size,
- $\text{prox}_g(v) = \operatorname{argmin}_u \frac{1}{2} \|u - v\|^2 + g(u)$ is the proximal operator of g .

A Nesterov-like extrapolation, like in FISTA, can be added to eq. (3) to accelerate convergence.

Proximal gradient methods are first-order methods that alternate a gradient step on f and a proximal evaluation on $h \circ W$:

$$\begin{cases} u_{n+\frac{1}{2}} = u - \alpha_n \nabla f(u) & \leftarrow \text{gradient descent} \\ u_{n+1} = \text{PROX}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}) & \leftarrow \text{proximal evaluation} \end{cases} \quad (3)$$

where

- $\alpha_n > 0$ is the step size,
- $\text{prox}_g(v) = \operatorname{argmin}_u \frac{1}{2} \|u - v\|^2 + g(u)$ is the proximal operator of g .

A Nesterov-like extrapolation, like in FISTA, can be added to eq. (3) to accelerate convergence.

Equation (3) assumes that the proximal operator of $h \circ W$ can be computed in closed form. This is not the case for some common regularization terms, such as the **Total Variation (TV)**.

Nested Primal-Dual (NPD)

The **Nested Primal-Dual (NPD)**¹² algorithm is an inexact proximal gradient method summarized as

$$\begin{cases} \bar{u}_n = u_n + \gamma_n(u_n - u_{n-1}) & \leftarrow \text{Nesterov extrap.} \\ u_{n+\frac{1}{2}} = \bar{u}_n - \alpha_n \nabla f(\bar{u}_n) & \leftarrow \text{gradient descent} \\ u_{n+1} \approx \text{prox}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}) & \leftarrow \text{inexact proximal step} \end{cases}$$

where $\gamma_n \geq 0$ is the extrapolation parameter like in FISTA.

¹Chen and Loris 2019, “On starting and stopping criteria for nested primal-dual iterations”.

²Bonettini et al. 2023, “A nested primal-dual FISTA-like scheme for composite convex optimization problems”.

The **Nested Primal-Dual (NPD)**¹² algorithm is an inexact proximal gradient method summarized as

$$\begin{cases} \bar{u}_n = u_n + \gamma_n(u_n - u_{n-1}) & \leftarrow \text{Nesterov extrap.} \\ u_{n+\frac{1}{2}} = \bar{u}_n - \alpha_n \nabla f(\bar{u}_n) & \leftarrow \text{gradient descent} \\ u_{n+1} \approx \text{prox}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}) & \leftarrow \text{inexact proximal step} \end{cases}$$

where $\gamma_n \geq 0$ is the extrapolation parameter like in FISTA.

The inexact proximal step is computed by $k_{max} \in \mathbb{N}$ steps of a **dual sequence** v^k :

Let $u \in \mathbb{R}^d$, $v^0 \in \mathbb{R}^{d'}$, $\alpha > 0$, $0 < \beta < 2/\|W\|^2$, and define the sequence

$$v^{k+1} = \text{prox}_{\beta\alpha^{-1}h^*}(v^k + \beta\alpha^{-1}W(u - \alpha W^T v^k)),$$

having limit $\lim_{k \rightarrow \infty} v^k = \hat{v}$. Then we have

$$\text{prox}_{\alpha h \circ W}(u) = u - \alpha W^T \hat{v},$$

¹Chen and Loris 2019, "On starting and stopping criteria for nested primal-dual iterations".

²Bonettini et al. 2023, "A nested primal-dual FISTA-like scheme for composite convex optimization problems".

Applying **left preconditioning** to the least squares regularized problem, the norm of the data fidelity term changes. Given S positive definite, we have

$$\operatorname{argmin}_{u \in \mathbb{R}^d} \frac{1}{2} \|Au - b^\delta\|_{S^{-1}}^2 + h(Wu), \quad (4)$$

where

$$f_S(u) = \|Au - b^\delta\|_{S^{-1}}^2 = \|S^{-\frac{1}{2}}(Au - b^\delta)\|^2.$$

Applying **left preconditioning** to the least squares regularized problem, the norm of the data fidelity term changes. Given S positive definite, we have

$$\operatorname{argmin}_{u \in \mathbb{R}^d} \frac{1}{2} \|Au - b^\delta\|_{S^{-1}}^2 + h(Wu), \quad (4)$$

where

$$f_S(u) = \|Au - b^\delta\|_{S^{-1}}^2 = \|S^{-\frac{1}{2}}(Au - b^\delta)\|^2.$$

Assuming that there exist a positive definite matrix P such that

$$P^{-1}A^T = A^T S^{-1}, \quad (5)$$

Applying **left preconditioning** to the least squares regularized problem, the norm of the data fidelity term changes. Given S positive definite, we have

$$\operatorname{argmin}_{u \in \mathbb{R}^d} \frac{1}{2} \|Au - b^\delta\|_{S^{-1}}^2 + h(Wu), \quad (4)$$

where

$$f_S(u) = \|Au - b^\delta\|_{S^{-1}}^2 = \|S^{-\frac{1}{2}}(Au - b^\delta)\|^2.$$

Assuming that there exist a positive definite matrix P such that

$$P^{-1}A^T = A^T S^{-1}, \quad (5)$$

then NPD applied to eq. (4) gives the **Preconditioned NPD (PNPD)** algorithm:

$$\begin{cases} \bar{u}_n = u_n + \gamma_n(u_n - u_{n-1}), \\ u_{n+\frac{1}{2}} = \bar{u}_n - \alpha_n P^{-1} \nabla f(\bar{u}_n), \\ u_{n+1} \approx \operatorname{PROX}_{\alpha_n h \circ W}(u_{n+\frac{1}{2}}), \end{cases}$$

To satisfy the condition in eq. (5), we can choose S and P as follows:

$$P = A^T A + \nu I, \quad S = A A^T + \nu I,$$

with $\nu > 0$. With this choice, the gradient step is equivalent to the Levenberg-Marquardt method applied to linear inverse problems, or equivalently, to the Iterated Tikhonov method³.

³Hanke and Groetsch 1998, “Nonstationary Iterated Tikhonov Regularization”.

To satisfy the condition in eq. (5), we can choose S and P as follows:

$$P = A^T A + \nu I, \quad S = AA^T + \nu I,$$

with $\nu > 0$. With this choice, the gradient step is equivalent to the Levenberg-Marquardt method applied to linear inverse problems, or equivalently, to the Iterated Tikhonov method³.

In general, the identity in eq. (5) is satisfied whenever P is a polynomial of $A^T A$ and S is a corresponding polynomial of AA^T .

³Hanke and Groetsch 1998, “Nonstationary Iterated Tikhonov Regularization”.

Unrolling architecture

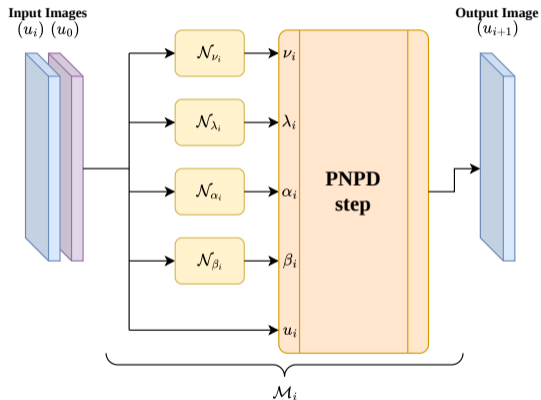


Figure 1: Single iteration of the unrolling architecture \mathcal{M}_i composed by the predictors for each parameter $\mathcal{N}_{\nu_i}, \mathcal{N}_{\lambda_i}, \mathcal{N}_{\alpha_i}, \mathcal{N}_{\beta_i}$ and by one step of the PNP algorithm.




Figure 2: Global unrolling architecture obtained by composing the base network \mathcal{M}_i N times.

Training

The training loss is defined as a weighted average of a base loss function ℓ computed on each intermediate iterate

$$\mathcal{L}(u_0, \dots, u_{N-1}, u) = \frac{\sum_{i=0}^{N-1} w_i \ell(u_i, u)}{\sum_{i=0}^{N-1} w_i}.$$

The training loss is defined as a weighted average of a base loss function ℓ computed on each intermediate iterate

$$\mathcal{L}(u_0, \dots, u_{N-1}, u) = \frac{\sum_{i=0}^{N-1} w_i \ell(u_i, u)}{\sum_{i=0}^{N-1} w_i}.$$

For the numerical experiments, we use the negative Structural Similarity Index (SSIM) as the base loss function

$$\ell(u_i, u) = -\text{SSIM}(u_i, u),$$

The training loss is defined as a weighted average of a base loss function ℓ computed on each intermediate iterate

$$\mathcal{L}(u_0, \dots, u_{N-1}, u) = \frac{\sum_{i=0}^{N-1} w_i \ell(u_i, u)}{\sum_{i=0}^{N-1} w_i}.$$

For the numerical experiments, we use the negative Structural Similarity Index (SSIM) as the base loss function

$$\ell(u_i, u) = -\text{SSIM}(u_i, u),$$

and we choose the weights w_i as

$$w_i = \frac{1}{N - i}.$$

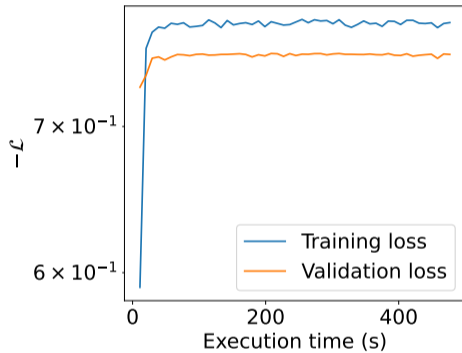



Figure 3: Evolution of the loss function as a function of epochs (left) and of time (right).

Numerical results

For the numerical experiments, we consider the image deblurring problem with Total Variation regularization, which can be formulated as

$$\operatorname{argmin}_{u \in \mathbb{R}^d} \underbrace{\frac{1}{2} \|Au - b^\delta\|^2}_{\text{data fidelity}} + \underbrace{\lambda \operatorname{TV}(u)}_{\text{regularization}},$$

which is a specific instance of the model in eq. (2) where

- $f(u) = \frac{1}{2} \|Au - b^\delta\|^2$,
- $h(Wu) = \lambda \operatorname{TV}(u)$,
- $A \in \mathbb{R}^{d \times d}$ is the blurring operator,
- $\lambda > 0$ is a regularization parameter
- $b^\delta \in \mathbb{R}^d$ is the observed image (blurred image with noise).

The TV is the Total Variation operator defined as

$$\text{TV}(x) = \sum_{i=1}^d \|\nabla_i u\|,$$

where $W \in \mathbb{R}^{2d \times d}$ is the discretization of the gradient operator, and $h : \mathbb{R}^{2d} \rightarrow \mathbb{R}$ is the sum of the pixel-wise euclidean norms.



Ground truth

Observed









Reconstructed



Ground truth

Observed

Reconstructed











Step sizes (α, β)

Reg. param. (λ)

Prec. param. (ν)

SSIM



Step sizes (α, β)

Reg. param. (λ)

Prec. param. (ν)

SSIM

- Use different designs for the predictor network \mathcal{N} .




- Use different designs for the predictor network \mathcal{N} .
- Input more information about the current iteration to the predictor, e.g. the dual variable v^0 of the inner loop or the residual.

- Use different designs for the predictor network \mathcal{N} .
- Input more information about the current iteration to the predictor, e.g. the dual variable v^0 of the inner loop or the residual.
- Test weighting strategies for the loss function.

- Use different designs for the predictor network \mathcal{N} .
- Input more information about the current iteration to the predictor, e.g. the dual variable v^0 of the inner loop or the residual.
- Test weighting strategies for the loss function.
- Refine training to avoid local minima and improve stability.

- Use different designs for the predictor network \mathcal{N} .
- Input more information about the current iteration to the predictor, e.g. the dual variable v^0 of the inner loop or the residual.
- Test weighting strategies for the loss function.
- Refine training to avoid local minima and improve stability.
- Test the unrolling architecture on different problems, algorithms, and regularization terms.

- Use different designs for the predictor network \mathcal{N} .
- Input more information about the current iteration to the predictor, e.g. the dual variable v^0 of the inner loop or the residual.
- Test weighting strategies for the loss function.
- Refine training to avoid local minima and improve stability.
- Test the unrolling architecture on different problems, algorithms, and regularization terms.

-  S. Aleotti, M. Donatelli, R. Krause, G. Scarlato
A Preconditioned Version of a Nested Primal-Dual Algorithm for Image Deblurring
J. Sci. Comput. 103, 85 (2025)
-  git repository for the PNPd unrolling codes
https://github.com/Giuseppe499/PNPd_unrolling
-  git repository for the PNPd codes
<https://github.com/Giuseppe499/PNPd>

Thank you for your attention!

