We thank the reviewers for their constructive comments. We provide our answers below.

Presentation: if the paper is accepted, we will improve our presentation, notably by better explaining "the big picture" and what our method does (R4); including detailed pseudo code for our algorithm (R1); emphasizing the difference between the introduced game setup and the standard bi-level formulation (R2); discussing data modalities without patch structure to which our approach is applicable, and distinguishing better our general methodology and its aspects specific to images (R1).

R1: "The authors consider both gradient l1 an l2 type regularization, as well as "standard" l1 and l2 methods [..]". In fact, our framework also handles a broader range of regularizers including non-local variants, sparse coding, and non-local group sparsity. "PSNR values by themselves don't really provide an overly compelling justification". We agree with the reviewer. If the paper is accepted we will include at least some qualitative examples in the main body of the paper and additional ones in the supplemental material. Finally, we also plan to include SSIM scores. See Table 1 below for an overview.

R2: " [..] all priors and problems within the context of this paper are convex". This is correct : we focus on convex games, i.e., each objective function h_j is convex wrt. $\mathbf{Z}$. "[..] the generalization to non-cooperative game seems a bit artificial to me". We can recover a classical optimization problem by summing convex regularizers. However an interesting insight from our experiments is that it is often beneficial to consider non-potential games instead. Please refer to Table 1 in our paper and Table A2 in the appendix for more details regarding this point. "Do you need fewer samples?". To answer this we have conducted additional experiments and report the results in Table 2 below. Interestingly, our method outperforms DnCNN for all numbers of training images used in the experiments, the gap in performance increasing significantly as the number of images decreases. This may be particularly interesting in applications where the amount of training data is very limited (for example in medical imaging). A more detailed study will be included if the paper is accepted. *Additional feedback*: we will correct typos and answer the rest of the questions below. Line 69: we do not have necessarily $p > m$ even though it is generally the case for sparse coding. "Is there multiple solutions". This could be the case (if h is not strongly convex). Line 113: $\mathbf{P}^\top$ denotes the operator which places patches on the reconstructed image. Line 119: Thank you, this will be corrected in the final version. Line 132-133:" this is fairly different from the TV from Chambolle". In the case of TV, we are in the pixel-level setting, hence $\mathbf{z}_j$ directly models the underlying value of the observed pixel $\mathbf{x}_j$. So our TVs are in fact similar (we consider here the anisotropic version). "What is $\mathcal{N}$? Is it the same as $\mathcal{N}_j$". Yes, the weights are shared. "I guess you mean Gradient". Yes we do.

R3: "[..] some of the extensions/generalizations seem relatively straightforward". We agree that our method makes the training of the studied priors relatively easy thanks to smoothing techniques and game encoding, and actually we believe that this is one of the strengths. To the best of our knowledge, no prior works proposed unrolled TV models and only one heuristic approach was proposed for group sparsity [28]. *Additional feedback*: "No θ dependence is explicitly written [..]". Thank you, we will clarify the notation. "[..] the update rule for $\mathbf{W}$ is never mentioned [..]". For patch-based models we employ a debiasing dictionary $\mathbf{W}$ to improve the quality of the reconstructions. Debiasing is commonly used when dealing with ℓ_1 penalty which is known to shrink the coefficients $\mathbf{Z}$ too much. In our method, $\mathbf{W}$ is learned through backpropagation. However, when dealing with pixel-based regularizers (including TV), we are in the setup where $q = 1 \times 1$ and $p = 1$ and the $\mathbf{W}$ matrix boils down to a single scalar coefficient. Empirically, this coefficient does not impact performance significantly so it was neglected in our implementation. We will clarify this point in the final version. "[..] it looks like $a_{j,k}$ depends on $\mathbf{W}$ through $\hat{\mathbf{y}}$". This is the case when the graph is non-local. We admit that the table can be confusing so it will be updated in the final version.

R4: "The weakness is the lack of theoretical guarantees, although its not a weakness of the paper itself, but of the family of methods in general". Our work does have theoretical in the case of a potential games. However, we did not manage to obtain convergence guarantees when considering a general convex game. Proving the convergence of the forward inference algorithm amounts to showing the monotonicity of the H operator, which turns out to be very challenging (very limited results on this type of problem are available in the literature). This is an interesting direction for future research. Thank you for reporting typos which will be corrected in the final version.

Table 1: Denoising results in terms of average PSNR(dB)/SSIM on BSD68.

Method	Noise level σ 15	25	50
DnCNN[70]	31.73/0.8907	29.23/0.8278	26.23/0.7189
TV	30.75/0.8614	28.24/0.7903	25.32/0.6619
SC+Var	31.49/0.8885	29.00/0.8234	26.08/0.7088
Group+Var	**31.75/0.8970**	**29.24/0.8341**	**26.34/0.7310**

Table 2: Denoising ($\sigma = 15$) with smaller training sets. Results in terms of average PSNR(dB) on BSD68.

Method	Params	Training images (BSD400) 400	200	100	50
DnCNN [70]	556k	31.73	31.65	31.47	31.23
TV	480	30.75	30.72	30.67	30.66
SC+Var	68k	31.49	31.49	31.46	31.40
Group+Var	68k	**31.75**	**31.66**	**31.62**	**31.54**