
Projection onto Minkowski Sums with Application to Constrained Learning: Supplement

Joong-Ho Won¹ Jason Xu² Kenneth Lange³

A. Additional Application of Projection onto Minkowski Sums

Constraint relaxation If B_r is a ball $\{\mathbf{x} : \|\mathbf{x}\| \leq r\}$ with respect to a norm $\|\cdot\|$, the sum $C_i + B_r$ is a kind of halo around C_i , yielding a relaxation of the constraint C_i . Within any projection-based method for solving constrained problems, one may replace a given C_i by its Minkowski sum $C_i + B_r$ to account for a degree of error encoded by r . In Bayesian methods, priors sharply constrained to a support set C_i lead to computational issues. Recent work considers posterior projections (Patra & Dunson, 2018) and relaxing such constraints via “ d -expansion” (Duan et al., 2018), a concept that coincides exactly with the Minkowski sum described above.

B. Additional Definitions

Definition B.1 (Strongly convex and smooth functions). *Let $C \in \mathbb{R}^d$ is a convex set and $\|\cdot\|$ be a norm over the smallest vector space containing C . We say function $f : C \rightarrow \mathbb{R}$ is α -strongly convex with respect to $\|\cdot\|$ if*

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\alpha}{2} \|\mathbf{y} - \mathbf{x}\|^2,$$

and is β -smooth with respect to $\|\cdot\|$ if

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\beta}{2} \|\mathbf{y} - \mathbf{x}\|^2$$

for all $\mathbf{x}, \mathbf{y} \in C$.

Definition B.2 (Domain of an extended real-valued function). *Let $\psi : \mathbb{R}^d \mapsto \mathbb{R} \cup \{\infty\}$ be an extended real-valued function. The domain of ψ is defined by $\text{dom}(\psi) = \{\mathbf{x} \in \mathbb{R}^d : \psi(\mathbf{x}) < \infty\}$. If $\text{dom}(\psi) \neq \emptyset$, then ψ is called proper.*

C. Projection for Internal and Boundary Points

Though we are naturally interested in external points, for completeness here we discuss the case when \mathbf{x} is an internal or boundary point of the Minkowski sum. If $\mathbf{x} \in A + B$, then finding appropriate summands \mathbf{a} and \mathbf{b} also succumbs to block descent. As already noted in the main text, the decomposition $\mathbf{x} = \mathbf{a} + \mathbf{b}$ is not necessarily unique. In practical examples, convergence of block descent can be exceedingly slow for an internal or boundary point of $A + B$. For this reason, it is useful to explore alternatives. One possibility is to minimize the proximity function

$$f(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \|\mathbf{x} - \mathbf{a} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \text{dist}(\mathbf{a}, A)^2 + \frac{\rho}{2} \text{dist}(\mathbf{b}, B)^2,$$

whose minimal value is 0 for any $\mathbf{x} \in A + B$ and $\rho > 0$. The MM principle (Lange, 2016) summarized in the section below suggests minimizing the surrogate function

$$g(\mathbf{a}, \mathbf{b} \mid \mathbf{a}_n, \mathbf{b}_n) = \frac{1}{2} \|\mathbf{x} - \mathbf{a} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{a} - P_A(\mathbf{a}_n)\|_2^2 + \frac{\rho}{2} \|\mathbf{b} - P_B(\mathbf{b}_n)\|_2^2$$

to generate improved values \mathbf{a}_{n+1} and \mathbf{b}_{n+1} . The stationarity conditions for the surrogate read

$$\mathbf{0} = -(\mathbf{x} - \mathbf{a} - \mathbf{b}) + \rho[\mathbf{a} - P_A(\mathbf{a}_n)]$$

$$\mathbf{0} = -(\mathbf{x} - \mathbf{a} - \mathbf{b}) + \rho[\mathbf{b} - P_B(\mathbf{b}_n)].$$

One can readily verify that this linear system has the solution

$$\begin{aligned} \mathbf{a}_{n+1} &= \frac{1}{2 + \rho}[\mathbf{x} - P_B(\mathbf{b}_n)] + \frac{1 + \rho}{2 + \rho}P_A(\mathbf{a}_n) \\ \mathbf{b}_{n+1} &= \frac{1}{2 + \rho}[\mathbf{x} - P_A(\mathbf{a}_n)] + \frac{1 + \rho}{2 + \rho}P_B(\mathbf{b}_n). \end{aligned}$$

These updates are guaranteed to reduce the objective $f(\mathbf{a}, \mathbf{b})$. It is straightforward to prove that the update map is nonexpansive when A and B are both convex. The objective $f(\mathbf{a}, \mathbf{b})$ is also convex in this setting, and stationary points and global minima coincide. Furthermore, any fixed point (\mathbf{a}, \mathbf{b}) with $\mathbf{a} \in A$ and $\mathbf{b} \in B$ satisfies $\mathbf{x} = \mathbf{a} + \mathbf{b}$. This algorithm is a special case of a class of algorithms called proximal distance algorithms (Xu et al., 2017), with the distinction that the tuning constant ρ need not be sent to ∞ when $\mathbf{x} \in A + B$.

D. Majorization-Minimization

The majorization-minimization (MM) principle provides a generic recipe for converting an optimization problem that is not immediately solvable (for instance, it may be non-convex or non-smooth) into a sequence of manageable problems. MM algorithms have become increasingly popular for large-scale optimization in statistics and machine learning (Lange, 2016), and includes expectation-maximization (EM) algorithms as a special case. MM algorithms operate by successively minimizing a sequence of *surrogate functions* $g(\mathbf{x} \mid \mathbf{x}_m)$ majorizing the objective function $f(\mathbf{x})$ at the current iterate \mathbf{x}_m . The notion of majorization requires two conditions: tangency $g(\mathbf{x}_m \mid \mathbf{x}_m) = f(\mathbf{x}_m)$ at the current iterate, and domination $g(\mathbf{x} \mid \mathbf{x}_m) \geq f(\mathbf{x})$ for all \mathbf{x} . The update rule

$$\mathbf{x}_{m+1} := \arg \min_{\mathbf{x}} g(\mathbf{x} \mid \mathbf{x}_m)$$

implies the descent property

$$f(\mathbf{x}_{m+1}) \leq g(\mathbf{x}_{m+1} \mid \mathbf{x}_m) \leq g(\mathbf{x}_m \mid \mathbf{x}_m) = f(\mathbf{x}_m).$$

Note that minimizing g is not strictly necessary: the weaker condition $g(\mathbf{x}_{m+1} \mid \mathbf{x}_m) \leq g(\mathbf{x}_m \mid \mathbf{x}_m)$ also decreases $f(\mathbf{x})$. Maximizing a function can be accomplished by an analogous combination of sequential minorization and maximization.

E. Additional Simulation Results for the $\ell_{1,p}$ -Overlapping Group Lasso

The full runtime comparison results of the simulation study of Sect. 5.1 of the main text is presented. Figure E.1 shows the runtime for varying dimension for fixed number of groups. The top right panel corresponds to the top left panel of Figure 1 in the main text. Figure E.1 illustrates the same information in a different format: runtime for varying number of groups for a fixed dimension. The bottom right panel corresponds to the top right panel of Figure 1 in the main text.

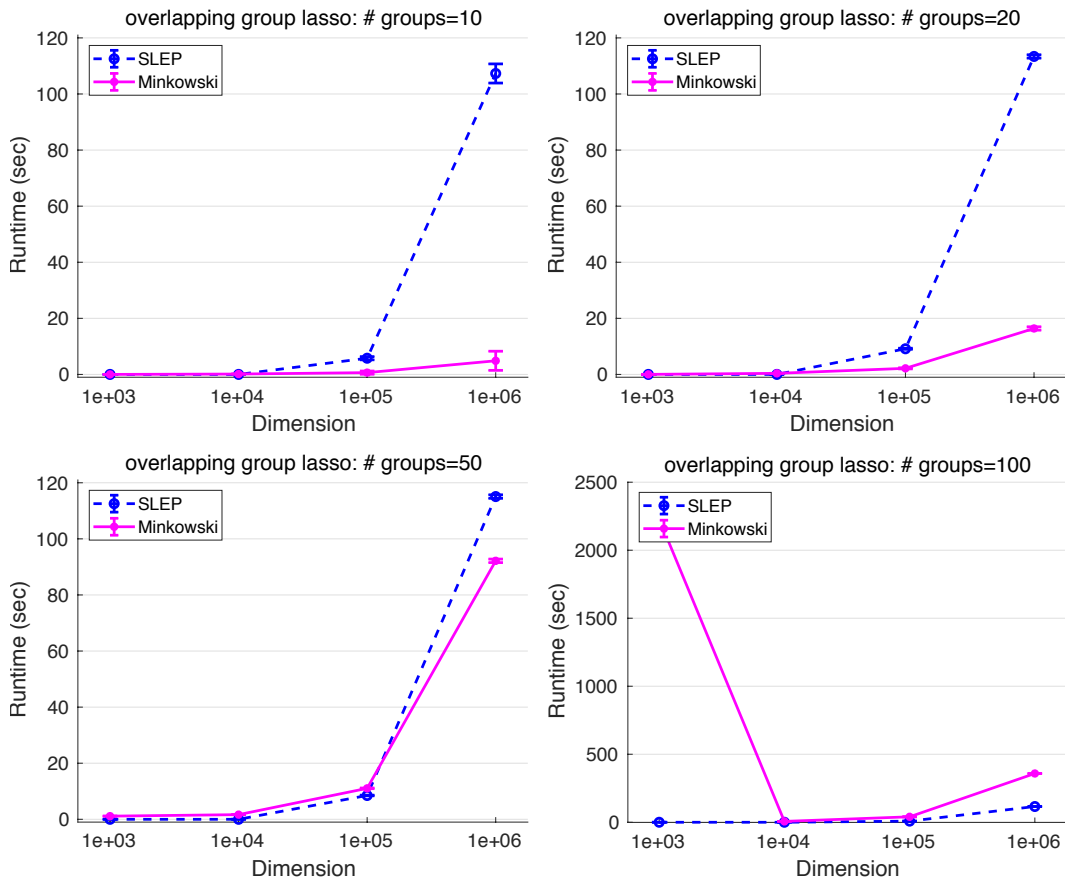


Figure E.1. Runtime-by-dimension comparison of the proposed Minkowski method and the dual projected gradient method by Yuan et al. (2011) for overlapping group lasso.

F. Additional Simulation Results for the Constrained Lasso

F.1. Additional Timing Results

The average runtime of path-following algorithm, Gurobi, ADMM, and the Minkowski projection-based proximal gradient descent methods are shown in Figure F.1 for $\lambda/\lambda_{\max} = 0.4$ and 0.8 . The results were obtained from the simulation as described in Sect. 5.2 of the main text, but omitted due to the space limit.

F.2. Accuracy

Following Gaines et al. (2018), we compare the objective value error of the path algorithm, ADMM, and the Minkowski methods relative to the final objective value of the Gurobi-solved quadratic program. Results for the zero-sum constrained lasso are shown in Figure F.2 for problem sizes $(n, d) = (100, 500)$, $(500, 1000)$, $(1000, 2000)$, $(2000, 4000)$, and $(4000, 8000)$ for which all four methods could be terminated. Likewise, results for the nonnegative lasso are presented in Figure F.3 for $(n, d) = (100, 500)$, $(500, 1000)$, and $(1000, 2000)$.

The presented results are consistent with those reported by Gaines et al. (2018) in case of the zero-sum lasso: the path algorithm is more accurate than the first-order methods; the accuracy of ADMM decreases as λ increases. It is notable that the Minkowski method tends to be more accurate than ADMM and is not sensitive to the sparsity level λ . Nevertheless, the accuracy of both first-order methods were less than 0.0001% , and this amount of error will not be significant practically.

On the other hand, in the nonnegative lasso the path-following algorithm was not as accurate as the other methods for $(n, d) = (100, 500)$, and at most similar for the larger problem sizes. Except for this outlier, all methods maintained high accuracy of less than $10^{-5}\%$ of relative error; ADMM was not sensitive to λ in this example.

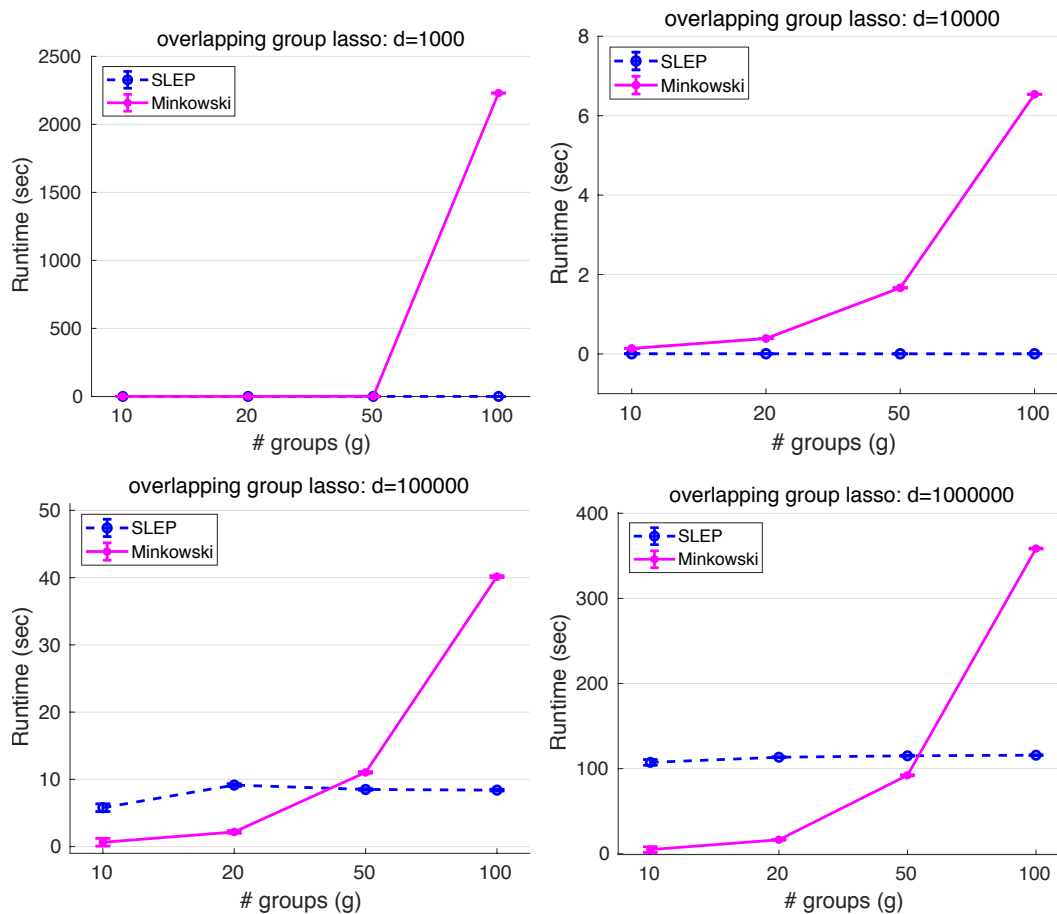


Figure E.2. Runtime-by-number-of-groups comparison of the proposed Minkowski method and the dual projected gradient method by Yuan et al. (2011) for overlapping group lasso.

References

- Duan, L. L., Young, A. L., Nishimura, A., and Dunson, D. B. Bayesian constraint relaxation. *arXiv preprint arXiv:1801.01525*, 2018.
- Gaines, B. R., Kim, J., and Zhou, H. Algorithms for fitting the constrained lasso. *Journal of Computational and Graphical Statistics*, 27(4):861–871, 2018.
- Lange, K. *MM Optimization Algorithms*. SIAM, 2016.
- Patra, S. and Dunson, D. B. Constrained bayesian inference through posterior projections. *arXiv preprint arXiv:1812.05741*, 2018.
- Xu, J., Chi, E., and Lange, K. Generalized linear model regression under distance-to-set penalties. In *Advances in Neural Information Processing Systems*, pp. 1385–1395, 2017.
- Yuan, L., Liu, J., and Ye, J. Efficient methods for overlapping group lasso. In *Advances in Neural Information Processing Systems*, pp. 352–360, 2011.

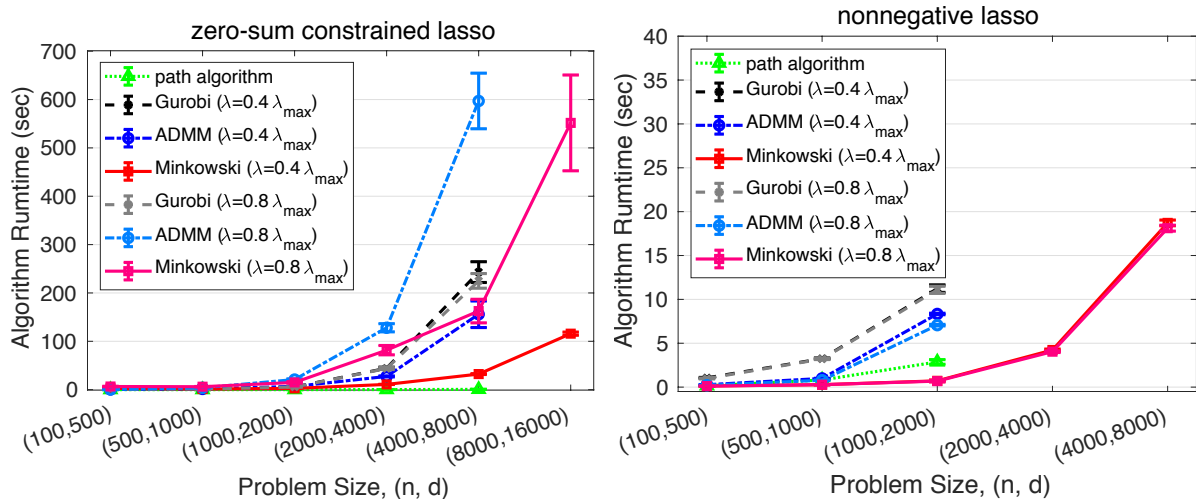


Figure F.1. Timing comparison of the proposed Minkowski method and the other methods by Gaines et al. (2018) for the constrained lasso. Left, runtime for the zero-sum constrained lasso. Right, runtime for the nonnegative lasso.

Projection onto Minkowski Sums

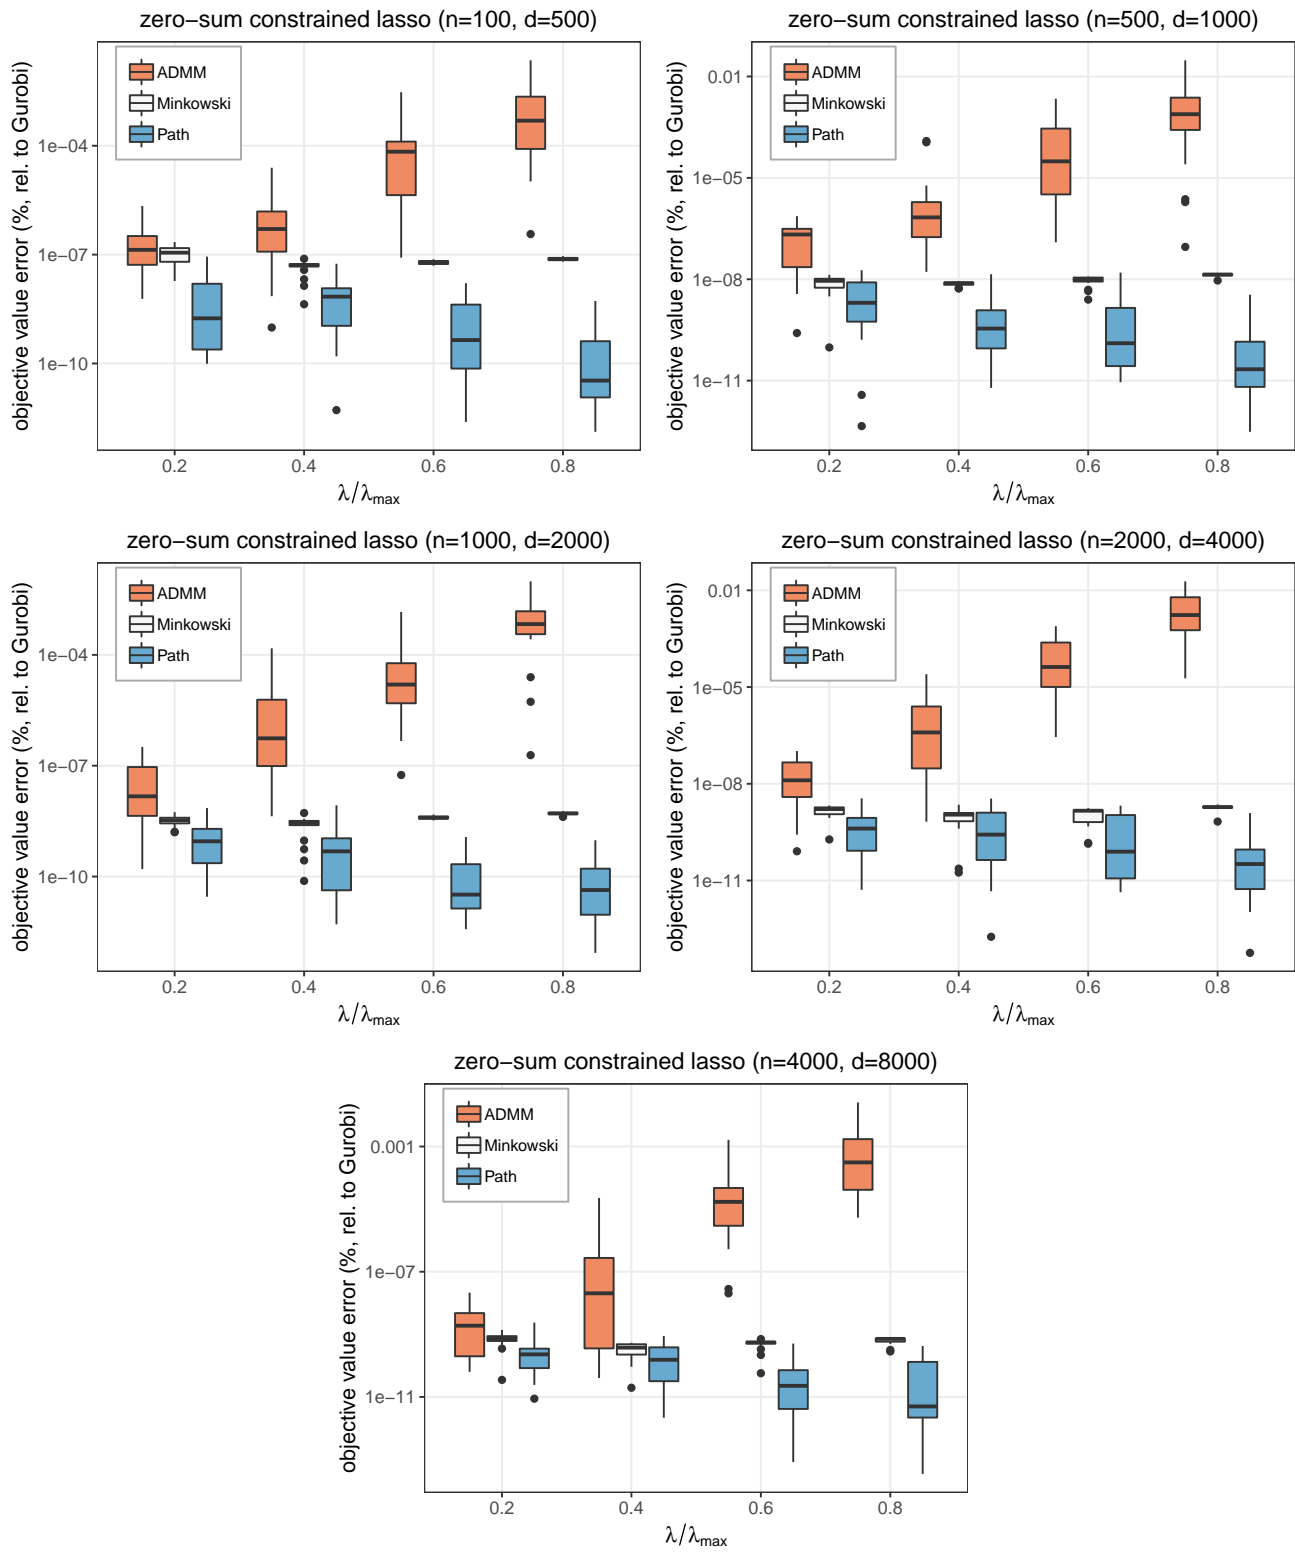


Figure F.2. Comparison of the proposed Minkowski method and the other methods by Gaines et al. (2018) for the zero-sum constrained lasso.

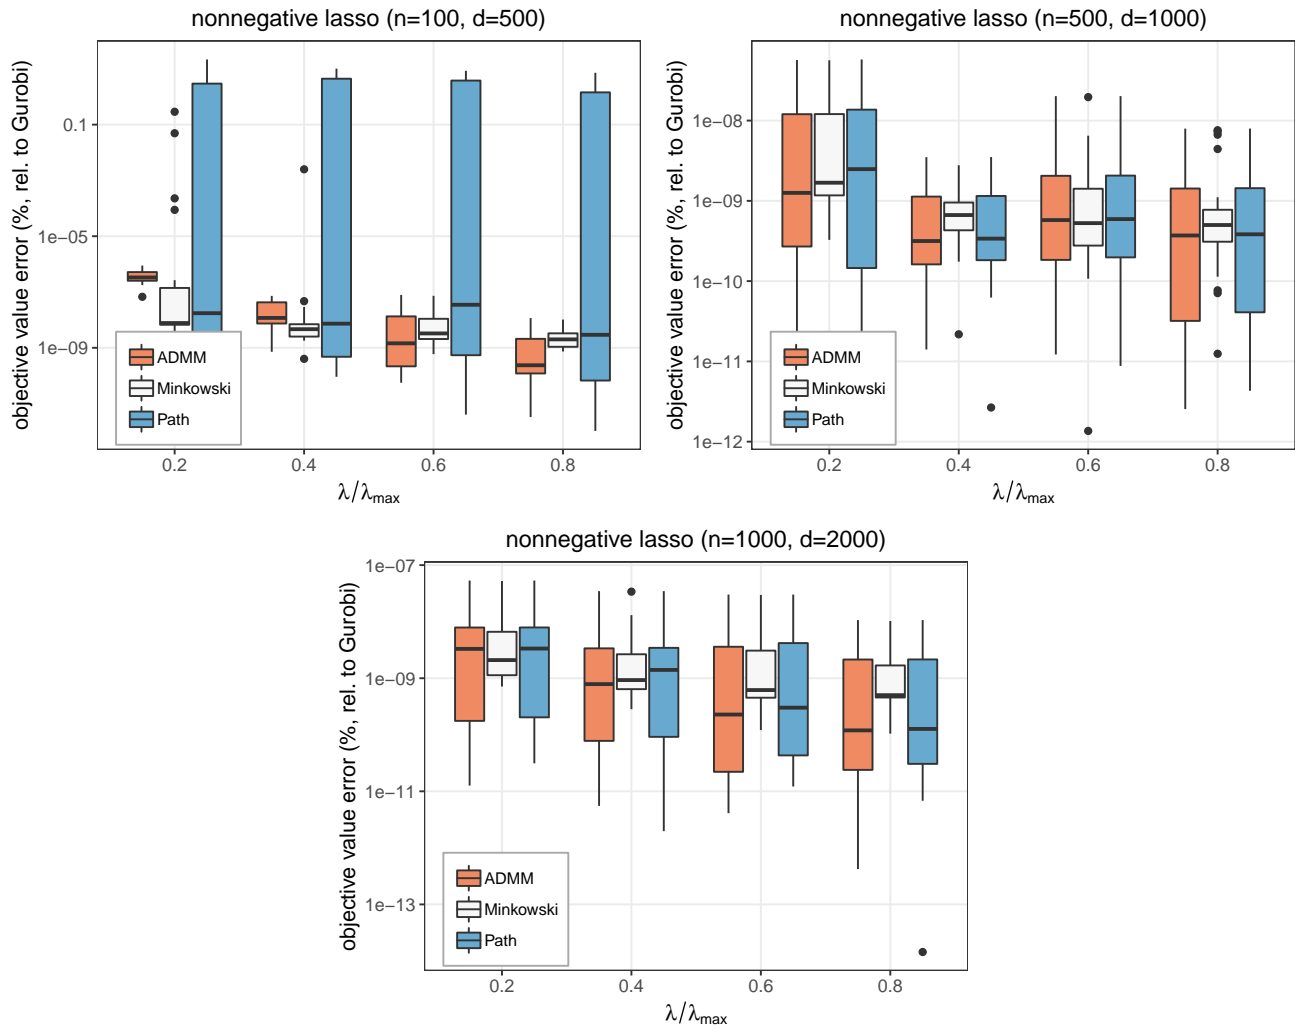


Figure F.3. Comparison of the proposed Minkowski method and the other methods by Gaines et al. (2018) for the nonnegative lasso.