## Supplementary Material for: Learning Sparse Nonparametric DAGs

## A Proofs

In this Appendix, we prove Proposition 1. For completeness, note that

$$
\begin{gathered}
\mathcal{F}=\left\{f \mid f(u)=\operatorname{MLP}\left(u ; A^{(1)}, \ldots, A^{(h)}\right),\right. \\
\left.f \text { independent of } u_{k}\right\}
\end{gathered}
$$

and

$$
\begin{aligned}
\mathcal{F}_{0}=\{f \mid f(u) & =\operatorname{MLP}\left(u ; A^{(1)}, \ldots, A^{(h)}\right) \\
A_{b k}^{(1)} & \left.=0, \forall b=1, \ldots, m_{1}\right\}
\end{aligned}
$$

We omit the bias terms in each layer as it does not affect the statement.

Proof of Proposition 1. We will show that $\mathcal{F} \subseteq \mathcal{F}_{0}$ and $\mathcal{F}_{0} \subseteq \mathcal{F}$.
(1) $\mathcal{F}_{0} \subseteq \mathcal{F}:$ for any $f_{0} \in \mathcal{F}_{0}$, we have $f_{0}(u)=$ $\operatorname{MLP}\left(u ; A^{(1)}, \ldots, A^{(h)}\right)$, where $A_{b k}^{(1)}=0$ for all $b=$ $1, \ldots, m_{1}$. Hence the linear function $A^{(1)} u$ is independent of $u_{k}$. Therefore,

$$
\begin{aligned}
f_{0}(u) & =\operatorname{MLP}\left(u ; A^{(1)}, \ldots, A^{(h)}\right) \\
& =\sigma\left(A^{(h)} \sigma\left(\cdots A^{(2)} \sigma\left(A^{(1)} u\right)\right)\right)
\end{aligned}
$$

is also independent of $u_{k}$, which means $f_{0} \in \mathcal{F}$.
(2) $\mathcal{F} \subseteq \mathcal{F}_{0}$ : for any $f \in \mathcal{F}$, we have $f(u)=$ $\operatorname{MLP}\left(u ; A^{(1)}, \ldots, A^{(h)}\right)$ and $f$ is independent of $u_{k}$. We will show that $f \in \mathcal{F}_{0}$ by constructing a matrix $\tilde{A}^{(1)}$, such that

$$
\begin{equation*}
f(u)=\operatorname{MLP}\left(u ; \tilde{A}^{(1)}, A^{(2)}, \ldots, A^{(h)}\right) \tag{14}
\end{equation*}
$$

and $\tilde{A}_{b k}^{(1)}=0$ for all $b=1, \ldots, m_{1}$.
Let $\tilde{u}$ be the vector such that $\tilde{u}_{k}=0$ and $\tilde{u}_{k^{\prime}}=u_{k}$ for all $k^{\prime} \neq k$. Since $\tilde{u}$ and $u$ differ only on the $k$ th dimension, and $f$ is independent of $u_{k}$, we have

$$
\begin{equation*}
f(u)=f(\tilde{u})=\operatorname{MLP}\left(\tilde{u} ; A^{(1)}, \ldots, A^{(h)}\right) \tag{15}
\end{equation*}
$$

Now define $\tilde{A}^{(1)}$ be the matrix such that $\tilde{A}_{b k}^{(1)}=0$ and $\tilde{A}_{b k^{\prime}}^{(1)}=A_{b k}^{(1)}$ for all $k^{\prime} \neq k$. Then we have the following observation: for each entry $s \in\left\{1, \ldots, m_{1}\right\}$,

$$
\begin{aligned}
\left(\tilde{A}^{(1)} u\right)_{s} & =\sum_{k^{\prime}=1}^{d} \tilde{A}_{s k^{\prime}} u_{k^{\prime}}=\sum_{k^{\prime} \neq k} A_{s k^{\prime}} u_{k^{\prime}} \\
& =\sum_{k^{\prime}=1}^{d} A_{s k^{\prime}} \tilde{u}_{k^{\prime}}=\left(A^{(1)} \tilde{u}\right)_{s}
\end{aligned}
$$

Hence,

$$
\begin{equation*}
\tilde{A}^{(1)} u=A^{(1)} \tilde{u} \tag{16}
\end{equation*}
$$

Therefore, by (15)

$$
\begin{aligned}
f(u) & =f(\tilde{u}) \\
& =\operatorname{MLP}\left(\tilde{u} ; A^{(1)}, \ldots, A^{(h)}\right) \\
& =\sigma\left(A^{(h)} \sigma\left(\cdots A^{(2)} \sigma\left(A^{(1)} \tilde{u}\right)\right)\right) \\
& =\sigma\left(A^{(h)} \sigma\left(\cdots A^{(2)} \sigma\left(\tilde{A}^{(1)} u\right)\right)\right) \\
& =\operatorname{MLP}\left(u ; \tilde{A}^{(1)}, A^{(2)}, \ldots, A^{(h)}\right)
\end{aligned}
$$

By definition of $\mathcal{F}_{0}$, we know that $\operatorname{MLP}\left(u ; \tilde{A}^{(1)}, A^{(2)}, \ldots, A^{(h)}\right) \in \mathcal{F}_{0}$. Thus, $f \in \mathcal{F}_{0}$ and we have completed the proof.

## B Experiment details

Baselines We consider the following baselines.

- Fast greedy equivalence search (FGS) ${ }^{2}$ (Ramsey et al., 2017) is based on greedy search and assumes linear dependency between variables.
- Greedy equivalence search with generalized scores (GSGES) ${ }^{3}$ (Huang et al., 2018) is also based on greedy search, but uses generalized scores without assuming a particular model class.
- DAG-GNN (GNN) ${ }^{4}$ (Yu et al., 2019) learns a (noisy) nonlinear transformation of a linear SEM using neural networks.
- NOTEARS (Linear) $)^{5}$ (Zheng et al., 2018) learns a linear SEM using continuous optimization.
- Causal additive model (CAM) ${ }^{6}$ (Bühlmann et al., 2014) learns an additive SEM by leveraging efficient nonparametric regression techniques and greedy search over edges.

For all experiments, default parameter settings are used, except for CAM where both preliminary neighborhood selection and pruning are applied.

Simulation Given the graph G, we simulate the SEM $X_{j}=f_{j}\left(X_{\mathrm{pa}(j)}\right)+z_{j}$ for all $j \in[d]$ in the topological order induced by G . We consider the following instances of $f_{j}$ :

[^0]

Figure 3: SHD (lower is better) with varying hidden layer size in NOTEARS-MLP.

- Additive GP: $f_{j}\left(X_{\mathrm{pa}(j)}\right)=\sum_{k \in \mathrm{pa}(j)} f_{j k}\left(X_{k}\right)$, where each $f_{j k}$ is a draw from Gaussian process with RBF kernel with length-scale one.
- Index model: $\quad f_{j}\left(X_{\mathrm{pa}(j)}\right) \quad=$ $\sum_{m=1}^{3} h_{m}\left(\sum_{k \in \operatorname{pa}(j)} \theta_{j m k} X_{k}\right)$, where $h_{1}=\tanh$, $h_{2}=\cos , h_{3}=\sin$, and each $\theta_{j m k}$ is drawn uniformly from range $[-2,-0.5] \cup[0.5,2]$.
- MLP: $f_{j}$ is a randomly initialized MLP with one hidden layer of size 100 and sigmoid activation.
- GP: $f_{j}$ is a draw from Gaussian process with RBF kernel with length-scale one.

In all settings, $z_{j}$ is i.i.d. standard Gaussian noise.

## C Sensitivity to number of hidden units

We also investigated the effect of number of hidden units in the NOTEARS-MLP estimate. It is well-known that as the size of the hidden layer increases, the functions representable by an MLP become more flexible. On the other hand, larger networks require more samples to estimate the parameters. Indeed, Figure 3 confirms this intuition. We plot the SHD with varying number of hidden units ranging from zero (i.e. linear function) to 100 units, using $n=1000$ and $n=200$ samples generated from the additive GP model on SF2 graph with $d=20$ nodes. One can first observe a sharp phase transition between zero and very few hidden units, which suggests the power of nonlinearity. Moreover, as the number of hidden units increases to 20 , the performance for both $n=1000$ and $n=200$ steadily improves, in which case the increased flexibility brings benefit. However, as we further increase the number of
hidden units, while SHD for $n=1000$ remains similar, the SHD for $n=200$ deteriorates, hinting at the lack of samples to take advantage of the increased flexibility.

## D Additional results

Full comparison We show $\{\mathrm{SHD}, \mathrm{FDR}, \mathrm{TPR}, \mathrm{FPR}\}$ results on all \{ER1, ER2, ER4, SF1, SF2, SF4\} graphs in Figure 4, 5, 6, 7 respectively. Similarly, see Figure $8,9,10,11$ for full comparison with CAM. As in Figure 1, each row is a random graph model, each column is a type of SEM. Overall NOTEARS-MLP has low FDR/FPR and high TPR, and same for NOTEARS-Sob on additive GP. Also observe that in most settings GNN has low FDR as well as low TPR, which is a consequence of only predicting a small number of edges.

Complexity and runtime Recall that numerical evaluation of matrix exponential involves solving linear systems, hence the time complexity is typically $O\left(d^{3}\right)$ for a dense $d \times d$ matrix. Taking NOTEARS-MLP with one hidden layer of $m$ units as an example, it takes $O\left(n d^{2} m+d^{2} m+d^{3}\right)$ time to evaluate the objective and the gradient. If $m / d=O(1)$, this is comparable to the linear case $O\left(n d^{2}+d^{3}\right)$, except for the inevitable extra cost from using a nonlinear function. This highlights the benefit of Proposition 1: the acyclicity constraint almost comes for free. Furthermore, we used a quasiNewton method to reduce the number of calls to evaluate the gradient, which involves computing the matrix exponential. Table 1 contains runtime comparison of different algorithms on ER2 graph with $n=1000$ samples. Recall that the kernel-based approach of GSGES comes with a $O\left(n^{3}\right)$ computational complexity, whereas NOTEARS-MLP and NOTEARS-Sob has $O(n)$ dependency on $n$. This can be confirmed from the table, which shows GSGES has a significantly longer runtime.

Comments on hyperparameter tuning The experiments presented in this paper were conducted under a fixed (and therefore suboptimal) value of $\lambda$ and weight threshold across all graph types, sparsity levels, and SEM types, despite the fact that each configuration may prefer different regularization strengths. Indeed, we observe substantially improved performance by choosing different values of hyperparameters in some settings. As our focus is not on attaining the best possible accuracy in all settings by carefully tuning the hyperparameters, we omit these results in the main text and only include here as a supplement. For instance, for ER4 graph with $d=40$ variables and $n=200$ samples, when the SEM is additive GP and MLP, setting $\lambda=0.03$ and threshold $=0.5$ gives results summarized in Table 2.

Learning Sparse Nonparametric DAGs

|  | NOTEARS-MLP | NOTEARS-Sob | FGS | Linear | GNN | GSGES |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $d=20$ | $92.12 \pm 22.51$ | $62.90 \pm 16.83$ | $0.55 \pm 0.43$ | $10.95 \pm 4.52$ | $498.32 \pm 43.72$ | $1547.42 \pm 109.83$ |
| $d=40$ | $282.64 \pm 67.46$ | $321.88 \pm 57.33$ | $0.59 \pm 0.17$ | $43.15 \pm 12.43$ | $706.35 \pm 64.49$ | $6379.98 \pm 359.67$ |

Table 1: Runtime (in seconds) of various algorithms on ER2 graph with $n=1000$ samples.

| SEM | Method | SHD | FDR | TPR | FPR | Predicted \# |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Additive-GP | NOTEARS-MLP | $124.3 \pm 6.65$ | $0.30 \pm 0.07$ | $0.35 \pm 0.04$ | $0.04 \pm 0.01$ | $81.70 \pm 10.49$ |
|  | GSGES | $121.3 \pm 5.02$ | $0.36 \pm 0.05$ | $0.28 \pm 0.03$ | $0.04 \pm 0.00$ | $69.30 \pm 5.01$ |
| MLP | NOTEARS-MLP | $88.40 \pm 11.29$ | $0.18 \pm 0.08$ | $0.57 \pm 0.06$ | $0.03 \pm 0.02$ | $111.70 \pm 15.97$ |
|  | GSGES | $121.60 \pm 11.95$ | $0.33 \pm 0.09$ | $0.33 \pm 0.06$ | $0.04 \pm 0.01$ | $77.10 \pm 7.13$ |

Table 2: ER4, $d=40, n=200$ with $\lambda=0.03$ and threshold $=0.5$.


Figure 4: Structure recovery measured by SHD (lower is better) to ground truth.


Figure 5: Structure recovery measured by FDR (lower is better) to ground truth.


Figure 6: Structure recovery measured by TPR (higher is better) to ground truth.


Figure 7: Structure recovery measured by FPR (lower is better) to ground truth.


Figure 8: Structure recovery measured by SHD (lower is better) to ground truth, compared with CAM.


Figure 9: Structure recovery measured by FDR (lower is better) to ground truth, compared with CAM.


Figure 10: Structure recovery measured by TPR (higher is better) to ground truth, compared with CAM.


Figure 11: Structure recovery measured by FPR (lower is better) to ground truth, compared with CAM.


[^0]:    ${ }^{2}$ https://github.com/bd2kccd/py-causal
    ${ }^{3}$ https://github.com/Biwei-Huang/
    Generalized-Score-Functions-for-Causal-Discovery/
    ${ }^{4}$ https://github.com/fishmoon1234/DAG-GNN
    ${ }^{5}$ https://github.com/xunzheng/notears
    ${ }^{6}$ https://cran.r-project.org/package=CAM

