# Connecting Weighted Automata and Recurrent Neural Networks through Spectral Learning (Supplementary Material) 

## A Proofs

## A. 1 Proof of Theorem 2

Theorem. Any function that can be computed by a vv-WFA with $n$ states can be computed by a linear 2-RNN with $n$ hidden units. Conversely, any function that can be computed by a linear 2-RNN with $n$ hidden units on sequences of one-hot vectors (i.e. canonical basis vectors) can be computed by a WFA with n states.

More precisely, the WFA $A=\left(\boldsymbol{\alpha},\left\{\mathbf{A}^{\sigma}\right\}_{\sigma \in \Sigma}, \boldsymbol{\Omega}\right)$ with $n$ states and the linear $2-R N N M=(\boldsymbol{\alpha}, \mathcal{A}, \boldsymbol{\Omega})$ with $n$ hidden units, where $\mathcal{A} \in \mathbb{R}^{n \times \Sigma \times n}$ is defined by $\mathcal{A}_{:, \sigma,:}=\mathbf{A}^{\sigma}$ for all $\sigma \in \Sigma$, are such that $f_{A}\left(\sigma_{1} \sigma_{2} \cdots \sigma_{k}\right)=$ $f_{M}\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{k}\right)$ for all sequences of input symbols $\sigma_{1}, \cdots, \sigma_{k} \in \Sigma$, where for each $i \in[k]$ the input vector $\mathbf{x}_{i} \in \mathbb{R}^{\Sigma}$ is the one-hot encoding of the symbol $\sigma_{i}$.

Proof. We first show by induction on $k$ that, for any sequence $\sigma_{1} \cdots \sigma_{k} \in \Sigma^{*}$, the hidden state $\mathbf{h}_{k}$ computed by $M$ (see Eq. (1)) on the corresponding one-hot encoded sequence $\mathbf{x}_{1}, \cdots, \mathbf{x}_{k} \in \mathbb{R}^{d}$ satisfies $\mathbf{h}_{k}=\left(\mathbf{A}^{\sigma_{1}} \cdots \mathbf{A}^{\sigma_{k}}\right)^{\top} \boldsymbol{\alpha}$. The case $k=0$ is immediate. Suppose the result true for sequences of length up to $k$. One can check easily check that $\mathcal{A} \bullet_{2} \mathbf{x}_{i}=\mathbf{A}^{\sigma_{i}}$ for any index $i$. Using the induction hypothesis it then follows that

$$
\begin{aligned}
\mathbf{h}_{k+1} & =\mathcal{A} \bullet_{1} \mathbf{h}_{k} \bullet_{2} \mathbf{x}_{k+1}=\mathbf{A}^{\sigma_{k+1}} \bullet_{1} \mathbf{h}_{k}=\left(\mathbf{A}^{\sigma_{k+1}}\right)^{\top} \mathbf{h}_{k} \\
& =\left(\mathbf{A}^{\sigma_{k+1}}\right)^{\top}\left(\mathbf{A}^{\sigma_{1}} \cdots \mathbf{A}^{\sigma_{k}}\right)^{\top} \boldsymbol{\alpha}=\left(\mathbf{A}^{\sigma_{1}} \cdots \mathbf{A}^{\sigma_{k+1}}\right)^{\top} \boldsymbol{\alpha} .
\end{aligned}
$$

To conclude, we thus have

$$
f_{M}\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \cdots, \mathbf{x}_{k}\right)=\boldsymbol{\Omega}_{k}=\boldsymbol{\Omega}\left(\mathbf{A}^{\sigma_{1}} \cdots \mathbf{A}^{\sigma_{k}}\right)^{\top} \boldsymbol{\alpha}=f_{A}\left(\sigma_{1} \sigma_{2} \cdots \sigma_{k}\right)
$$

## A. 2 Proof of Theorem 3

Theorem. Let $f:\left(\mathbb{R}^{d}\right)^{*} \rightarrow \mathbb{R}^{p}$ be a function computed by a minimal linear 2-RNN with $n$ hidden units and let $L$ be an integer such that $\operatorname{rank}\left(\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle\langle L, L+1\rangle\rangle}\right)=n$.

Then, for any $\mathbf{P} \in \mathbb{R}^{d^{L} \times n}$ and $\mathbf{S} \in \mathbb{R}^{n \times d^{L} p}$ such that $\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle\langle L, L+1\rangle\rangle}=\mathbf{P S}$, the linear 2-RNN $M=(\boldsymbol{\alpha}, \mathcal{A}, \boldsymbol{\Omega})$ defined by

$$
\boldsymbol{\alpha}=\left(\mathbf{S}^{\dagger}\right)^{\top}\left(\mathcal{H}_{f}^{(L)}\right)_{\langle\langle L+1\rangle\rangle}, \quad \mathcal{A}=\left(\left(\mathcal{H}_{f}^{(2 L+1)}\right)_{\langle\langle L, 1, L+1\rangle\rangle}\right) \times_{1} \mathbf{P}^{\dagger} \times_{3}\left(\mathbf{S}^{\dagger}\right)^{\top}, \quad \boldsymbol{\Omega}^{\top}=\mathbf{P}^{\dagger}\left(\boldsymbol{\mathcal { H }}_{f}^{(L)}\right)_{\langle\langle L, 1\rangle\rangle}
$$

is a minimal linear $2-R N N$ computing $f$.
Proof. Let $\mathbf{P} \in \mathbb{R}^{d^{L} \times n}$ and $\mathbf{S} \in \mathbb{R}^{n \times d^{L} p}$ be such that $\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle\langle L, L+1\rangle\rangle}=\mathbf{P S}$ Define the tensors

$$
\mathcal{P}^{*}=\llbracket \mathcal{A}^{\star} \bullet_{1} \boldsymbol{\alpha}^{\star}, \underbrace{\mathcal{A}^{\star}, \cdots, \mathcal{A}^{\star}}_{L-1 \text { times }}, \mathbf{I}_{n} \rrbracket \in \mathbb{R}^{d \times \cdots \times d \times n} \quad \text { and } \quad \mathcal{S}^{*}=\llbracket \mathbf{I}_{n}, \underbrace{\mathcal{A}^{\star}, \cdots, \mathcal{A}^{\star}}_{L \text { times }}, \boldsymbol{\Omega}^{\star} \rrbracket \in \mathbb{R}^{n \times d \times \cdots \times d \times p}
$$

of order $L+1$ and $L+2$ respectively, and let $\mathbf{P}^{\star}=\left(\mathcal{P}^{*}\right)_{\langle\langle l, 1\rangle\rangle} \in \mathbb{R}^{d^{l} \times n}$ and $\mathbf{S}=\left(\mathcal{S}^{*}\right)_{\langle 1, L+1\rangle\rangle} \in \mathbb{R}^{n \times d^{l} p}$. Using the identity $\mathcal{H}_{f}^{(j)}=\llbracket \mathcal{A} \bullet{ }_{1} \boldsymbol{\alpha}, \underbrace{\mathcal{A}, \cdots, \mathcal{A}}_{j-1 \text { times }}, \boldsymbol{\Omega}^{\top} \rrbracket$ for any $j$, one can easily check the following identities:

$$
\begin{gathered}
\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle\langle L, L+1\rangle\rangle}=\mathbf{P}^{\star} \mathbf{S}^{\star}, \quad\left(\mathcal{H}_{f}^{(2 L+1)}\right)_{\langle L L, 1, L+1\rangle\rangle}=\mathcal{A}^{\star} \times_{1} \mathbf{P}^{\star} \times_{3}\left(\mathbf{S}^{\star}\right)^{\top}, \\
\left(\mathcal{H}_{f}^{(L)}\right)_{\langle L, 1\rangle\rangle}=\mathbf{P}^{\star}\left(\mathbf{\Omega}^{\star}\right)^{\top}, \quad\left(\mathcal{H}_{f}^{(L)}\right)_{\langle\langle L+1\rangle\rangle}=\left(\mathbf{S}^{\star}\right)^{\top} \boldsymbol{\alpha} .
\end{gathered}
$$

Let $\mathbf{M}=\mathbf{P}^{\dagger} \mathbf{P}^{\star}$. We will show that $\boldsymbol{\alpha}=\mathbf{M}^{-\top} \boldsymbol{\alpha}^{\star}, \mathcal{A}=\mathcal{A}^{\star} \times_{1} \mathbf{M} \times{ }_{3} \mathbf{M}^{-\top}$ and $\boldsymbol{\Omega}=\mathbf{M} \boldsymbol{\Omega}^{\star}$, which will entail the results since linear 2-RNN are invariant under change of basis (see Section 2). First observe that $\mathbf{M}^{-1}=\mathbf{S}^{\star} \mathbf{S}^{\dagger}$. Indeed, we have $\mathbf{P}^{\dagger} \mathbf{P}^{\star} \mathbf{S}^{\star} \mathbf{S}^{\dagger}=\mathbf{P}^{\dagger}\left(\mathcal{H}_{f}^{(2 l)}\right)_{\langle l, l+1\rangle\rangle} \mathbf{S}^{\dagger}=\mathbf{P}^{\dagger} \mathbf{P S S}{ }^{\dagger}=\mathbf{I}$ where we used the fact that $\mathbf{P}$ (resp. S) is of full column rank (resp. row rank) for the last equality.

The following derivations then follow from basic tensor algebra:

$$
\begin{aligned}
\boldsymbol{\alpha} & =\left(\mathbf{S}^{\dagger}\right)^{\top}\left(\mathcal{H}_{f}^{(L)}\right)_{\langle L L+1\rangle\rangle}=\left(\mathbf{S}^{\dagger}\right)^{\top}\left(\mathbf{S}^{\star}\right)^{\top} \boldsymbol{\alpha}=\left(\mathbf{S}^{\star} \mathbf{S}^{\dagger}\right)^{\top}=\mathbf{M}^{-\top} \boldsymbol{\alpha}^{\star} \\
\boldsymbol{\mathcal { A }} & =\left(\left(\boldsymbol{\mathcal { H }}_{f}^{(2 L+1)}\right)_{\langle L L, 1, L+1\rangle\rangle}\right) \times_{1} \mathbf{P}^{\dagger} \times_{3}\left(\mathbf{S}^{\dagger}\right)^{\top} \\
& =\left(\boldsymbol{\mathcal { A }}^{\star} \times_{1} \mathbf{P}^{\star} \times_{3}\left(\mathbf{S}^{\star}\right)^{\top}\right) \times_{1} \mathbf{P}^{\dagger} \times_{3}\left(\mathbf{S}^{\dagger}\right)^{\top} \\
& =\mathcal{A}^{\star} \times_{1} \mathbf{P}^{\dagger} \mathbf{P}^{\star} \times_{3}\left(\mathbf{S}^{\star} \mathbf{S}^{\dagger}\right)^{\top}=\mathcal{A}^{\star} \times_{1} \mathbf{M} \times_{3} \mathbf{M}^{-\top} \\
\mathbf{\Omega}^{\top} & =\mathbf{P}^{\dagger}\left(\boldsymbol{\mathcal { H }}_{f}^{(L)}\right)_{\langle\langle L, 1\rangle\rangle}=\mathbf{P}^{\dagger} \mathbf{P}^{\star}\left(\mathbf{\Omega}^{\star}\right)^{\top}=\mathbf{M} \boldsymbol{\Omega}^{\star}
\end{aligned}
$$

which concludes the proof.

## A. 3 Proof of Theorem 4

Theorem. Let $\left(\mathbf{h}_{0}, \mathcal{A}, \boldsymbol{\Omega}\right)$ be a minimal linear 2-RNN with $n$ hidden units computing a function $f:\left(\mathbb{R}^{d}\right)^{*} \rightarrow \mathbb{R}^{p}$, and let $L$ be an integer ${ }^{7}$ such that $\operatorname{rank}\left(\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle L, L+1\rangle}\right)=n$.

Suppose we have access to 3 datasets $D_{l}=\left\{\left(\left(\mathbf{x}_{1}^{(i)}, \mathbf{x}_{2}^{(i)}, \cdots, \mathbf{x}_{l}^{(i)}\right), \mathbf{y}^{(i)}\right)\right\}_{i=1}^{N_{l}} \subset\left(\mathbb{R}^{d}\right)^{l} \times \mathbb{R}^{p}$ for $l \in\{L, 2 L, 2 L+1\}$ where the entries of each $\mathbf{x}_{j}^{(i)}$ are drawn independently from the standard normal distribution and where each $\mathbf{y}^{(i)}=f\left(\mathbf{x}_{1}^{(i)}, \mathbf{x}_{2}^{(i)}, \cdots, \mathbf{x}_{l}^{(i)}\right)$.

Then, whenever $N_{l} \geq d^{l}$ for each $l \in\{L, 2 L, 2 L+1\}$, the linear 2-RNN $M$ returned by Algorithm 1 with the least-squares method satisfies $f_{M}=f$ with probability one.

Proof. We just need to show for each $l \in\{L, 2 L, 2 L+1\}$ that, under the hypothesis of the Theorem, the Hankel tensors $\hat{\mathcal{H}}^{(l)}$ computed in line 4 of Algorithm 1 are equal to the true Hankel tensors $\boldsymbol{\mathcal { H }}^{(l)}$ with probability one. Recall that these tensors are computed by solving the least-squares problem

$$
\hat{\mathcal{H}}^{(l)}=\underset{T \in \mathbb{R}^{d \times \cdots \times d \times p}}{\arg \min }\left\|\mathbf{X}(\boldsymbol{\mathcal { T }})_{\langle l, 1\rangle\rangle}-\mathbf{Y}\right\|_{F}^{2}
$$

where $\mathbf{X} \in \mathbb{R}^{N_{l} \times d_{l}}$ is the matrix with rows $\mathbf{x}_{1}^{(i)} \otimes \mathbf{x}_{2}^{(i)} \otimes \cdots \otimes \mathbf{x}_{l}^{(i)}$ for each $i \in\left[N_{l}\right]$. Since $\mathbf{X}\left(\mathcal{H}^{(l)}\right)_{\langle l, 1\rangle\rangle}=\mathbf{Y}$ and since the solution of the least-squares problem is unique as soon as $\mathbf{X}$ is of full column rank, we just need to show that this is the case with probability one when the entries of the vectors $\mathbf{x}_{j}^{(i)}$ are drawn at random from a standard normal distribution. The result will then directly follow by applying Theorem 3.
We will show that the set

$$
\mathcal{S}=\left\{\left(\mathbf{x}_{1}^{(i)}, \cdots, \mathbf{x}_{l}^{(i)}\right) \mid i \in\left[N_{l}\right], \operatorname{dim}\left(\operatorname{span}\left(\left\{\mathbf{x}_{1}^{(i)} \otimes \mathbf{x}_{2}^{(i)} \otimes \cdots \otimes \mathbf{x}_{l}^{(i)}\right\}\right)\right)<d^{l}\right\}
$$

has Lebesgue measure 0 in $\left(\left(\mathbb{R}^{d}\right)^{l}\right)^{N_{l}} \simeq \mathbb{R}^{d l N_{l}}$ as soon as $N_{l} \geq d^{l}$, which will imply that it has probability 0 under any continuous probability, hence the result. For any $S=\left\{\left(\mathbf{x}_{1}^{(i)}, \cdots, \mathbf{x}_{l}^{(i)}\right)\right\}_{i=1}^{N_{l}}$, we denote by $\mathbf{X}_{S} \in \mathbb{R}^{N_{l} \times d^{l}}$ the matrix with rows $\mathbf{x}_{1}^{(i)} \otimes \mathbf{x}_{2}^{(i)} \otimes \cdots \otimes \mathbf{x}_{l}^{(i)}$. One can easily check that $S \in \mathcal{S}$ if and only if $\mathbf{X}_{S}$ is of rank strictly less than $d^{l}$, which is equivalent to the determinant of $\mathbf{X}_{S}^{\top} \mathbf{X}_{S}$ being equal to 0 . Since this determinant is a polynomial in the entries of the vectors $\mathbf{x}_{j}^{(i)}, \mathcal{S}$ is an algebraic subvariety of $\mathbb{R}^{d l N_{l}}$. It is then easy to check that the polynomial $\operatorname{det}\left(\mathbf{X}_{S}^{\top} \mathbf{X}_{S}\right)$ is not uniformly 0 when $N_{l} \geq d^{l}$. Indeed, it suffices to choose the vectors $\mathbf{x}_{j}^{(i)}$ such

[^0]that the family $\left(\mathbf{x}_{1}^{(i)} \otimes \mathbf{x}_{2}^{(i)} \otimes \cdots \otimes \mathbf{x}_{l}^{(i)}\right)_{n=1}^{N_{l}}$ spans the whole space $\mathbb{R}^{d^{l}}$ (which is possible since we can choose arbitrarily any of the $N_{l} \geq d^{l}$ elements of this family), hence the result. In conclusion, $\mathcal{S}$ is a proper algebraic subvariety of $\mathbb{R}^{d l N_{l}}$ and hence has Lebesgue measure zero [14, Section 2.6.5].

## B Lifting the simplifying assumption

We now show how all our results still hold when there does not exist an $L$ such that $\operatorname{rank}\left(\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle L L, L+1\rangle\rangle}\right)=n$. Recall that this simplifying assumption followed from assuming that the sets $\mathcal{P}=\mathcal{S}=[d]^{L}$ form a complete basis for the function $\tilde{f}:[d]^{*} \rightarrow \mathbb{R}^{p}$ defined by $\tilde{f}\left(i_{1} i_{2} \cdots i_{k}\right)=f\left(\mathbf{e}_{i_{1}}, \mathbf{e}_{i_{2}}, \cdots, \mathbf{e}_{i_{k}}\right)$. We first show that there always exists an integer $L$ such that $\mathcal{P}=\mathcal{S}=\cup_{i \leq L}[d]^{i}$ forms a complete basis for $\tilde{f}$. Let $M=\left(\boldsymbol{\alpha}^{\star}, \mathcal{A}^{\star}, \boldsymbol{\Omega}^{\star}\right)$ be a linear $2-$ RNN with $n$ hidden units computing $f$ (i.e. such that $f_{M}=f$ ). It follows from Theorem 2 and from the discussion at the beginning of Section 4.1 that there exists a vv-WFA computing $\tilde{f}$ and it is easy to check that $\operatorname{rank}(\tilde{f})=n$. This implies $\operatorname{rank}\left(\left(\mathcal{H}_{f}\right)_{(1)}\right)=n$ by Theorem 1 . Since $\mathcal{P}=\mathcal{S}=\cup_{i \leq l}[d]^{i}$ converges to $[d]^{*}$ as $l$ grows to infinity, there exists an $L$ such that the finite sub-block $\tilde{\mathcal{H}}_{f} \in \mathbb{R}^{\mathcal{P} \times \mathcal{S} \times p}$ of $\mathcal{H}_{f} \in \mathbb{R}^{[d]^{*} \times[d]^{*} \times p}$ satisfies $\operatorname{rank}\left(\left(\tilde{\mathcal{H}}_{f}\right)_{(1)}\right)=n$, i.e. such that $\mathcal{P}=\mathcal{S}=\cup_{i \leq L}[d]^{i}$ forms a complete basis for $\tilde{f}$.
Now consider the finite sub-blocks $\tilde{\mathcal{H}}_{f}^{+} \in \mathbb{R}^{\mathcal{P} \times[d] \times \mathcal{S} \times p}$ and $\tilde{\mathbf{H}}_{f}^{-} \in \mathbb{R}^{\mathcal{P} \times p}$ of $\mathcal{H}_{f}$ defined by

$$
\left(\tilde{\mathcal{H}}_{f}^{+}\right)_{u, i, v,:}=\tilde{f}(u i v), \quad \operatorname{and}\left(\tilde{\mathbf{H}}_{f}^{-}\right)_{u,:}=f(u)
$$

for any $u \in \mathcal{P}=\mathcal{S}$ and any $i \in[d]$. One can check that Theorem 3 holds by replacing mutatis mutandi $\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle\langle L, L+1\rangle\rangle}$ by $\left(\tilde{\mathcal{H}}_{f}\right)_{(1)},\left(\mathcal{H}_{f}^{(2 L+1)}\right)_{\langle\langle L, 1, L+1\rangle\rangle}$ by $\tilde{\mathcal{H}}_{f}^{+},\left(\mathcal{H}_{f}^{(L)}\right)_{\langle L L, 1\rangle\rangle}$ by $\tilde{\mathbf{H}}_{f}^{-}$and $\left(\mathcal{H}_{f}^{(L)}\right)_{\langle\langle L+1\rangle\rangle}$ by vec $\left(\tilde{\mathbf{H}}_{f}^{-}\right)$.
To conclude, it suffices to observe that both $\tilde{\mathcal{H}}_{f}^{+}$and $\tilde{\mathbf{H}}_{f}^{-}$can be constructed from the entries for the tensors $\boldsymbol{\mathcal { H }}^{(l)}$ for $1 \leq l \leq 2 L+1$, which can be recovered (or estimated in the noisy setting) using the techniques described in Section 4.2 (corresponding to lines 2-12 of Algorithm 1).
We thus showed that linear 2 -RNNs can be provably learned even when there does not exist an $L$ such that $\operatorname{rank}\left(\left(\mathcal{H}_{f}^{(2 L)}\right)_{\langle L, L+1\rangle\rangle}\right)=n$. In this setting, one needs to estimate enough of the tensors $\boldsymbol{\mathcal { H }}^{(l)}$ to reconstruct a complete sub-block $\tilde{\mathcal{H}}_{f}$ of the Hankel tensor $\mathcal{H}$ (along with the corresponding tensor $\tilde{\mathcal{H}}_{f}^{+}$and matrix $\tilde{\mathbf{H}}_{f}^{-}$) and recover the linear 2-RNN by applying Theorem 3. In addition, one needs to have access to sufficiently large datasets $D_{l}$ for each $l \in[2 L+1]$ rather than only the three datasets mentioned in Theorem 4. However the data requirement remains the same in the case where we assume that each of the datasets $D_{l}$ is constructed from a unique training dataset $S=\left\{\left(\left(\mathbf{x}_{1}^{(i)}, \mathbf{x}_{2}^{(i)}, \cdots, \mathbf{x}_{T}^{(i)}\right),\left(\mathbf{y}_{1}^{(i)}, \mathbf{y}_{2}^{(i)}, \cdots, \mathbf{y}_{T}^{(i)}\right)\right)\right\}_{i=1}^{N}$ of input/output sequences.

## C Leveraging the tensor train structure for computational efficiency

The overall learning algorithm using the TIHT recovery method in TT format is summarized in Algorithm 2. The key ingredients to improve the complexity of Algorithm 1 are (i) to estimate the gradient using mini-batches of data and (ii) to directly use the TT format to represent and perform operations on the tensors $\mathcal{H}^{(l)}$ and the tensors $\boldsymbol{\mathcal { X }}^{(l)} \in \mathbb{R}^{M \times d \times \cdots \times d}$ defined by

$$
\begin{equation*}
\mathcal{X}_{i,,,, \cdots,:}=\mathbf{x}_{1}^{(i)} \otimes \mathbf{x}_{2}^{(i)} \otimes \cdots \otimes \mathbf{x}_{l}^{(i)} \quad \text { for } i \in[M] \tag{3}
\end{equation*}
$$

where $M$ is the size of a mini-batch of training data $\left(\mathcal{H}^{(l)}\right.$ is of TT-rank $R$ by design and it can easily be shown that $\boldsymbol{\mathcal { X }}^{(l)}$ is of TT-rank at most $M$, cf. Eq. (4)). Then, all the operations of the algorithm can be expressed in terms of these tensors and performed efficiently in TT format. More precisely, the products and sums needed to compute the gradient update on line 6 can be performed in $\mathcal{O}\left((R+M)^{2}(l d+p)+(R+M)^{3} d\right)$. After the gradient update, the tensor $\mathcal{H}^{(l)}$ has TT-rank at most $(M+R)$ but can be efficiently projected back to a tensor of TT-rank $R$ using the tensor train rounding operation [33] in $\mathcal{O}\left((R+M)^{3}(l d+p)\right)$ (which is the operation dominating the complexity of the whole algorithm). The subsequent operations on line 10 can be performed efficiently in the TT format in $\mathcal{O}\left(R^{3} d+R^{2} p\right)$ (using the method described in [26] to compute the pseudo-inverses of the matrices $\mathbf{P}$ and $\mathbf{S})$. The overall complexity of Algorithm 2 is thus in $\mathcal{O}\left(T(R+M)^{3}(L d+p)\right)$ where $T$ is the number of iterations of the inner loop.

## Algorithm 2 2RNN-SL-TT: Spectral Learning of linear 2-RNNs in tensor train format

Input: Three training datasets $D_{L}, D_{2 L}, D_{2 L+1}$ with input sequences of length $L, 2 L$ and $2 L+1$ respectively, rank $R$, learning rate $\gamma$ and mini-batch size $M$.
for $l \in\{L, 2 L, 2 L+1\}$ do
Initialize all cores of the rank $R$ TT-decomposition $\mathcal{H}^{(l)}=\llbracket \mathcal{G}_{1}^{(l)}, \cdots, \mathcal{G}_{l+1}^{(l)} \rrbracket \in \mathbb{R}^{d \times \cdots \times d \times p}$ to $\mathbf{0}$.
// Note that all the updates of $\mathcal{H}^{(l)}$ stated below are in effect applied directly to the core tensors $\boldsymbol{\mathcal { G }}_{k}^{(l)}$, i.e. the tensor $\mathcal{H}^{(l)}$ is never explicitely constructed.
repeat
Subsample a minibatch

$$
\left\{\left(\left(\mathbf{x}_{1}^{(i)}, \mathbf{x}_{2}^{(i)}, \cdots, \mathbf{x}_{l}^{(i)}\right), \mathbf{y}^{(i)}\right)\right\}_{i=1}^{M} \subset\left(\mathbb{R}^{d}\right)^{l} \times \mathbb{R}^{p}
$$

of size $M$ from $D_{l}$.
5: $\quad$ Compute the rank $M$ TT-decomposition of the tensor $\boldsymbol{\mathcal { X }}=\boldsymbol{\mathcal { X }}^{(l)}$ (defined in Eq. (3)), which is given by

$$
\begin{equation*}
\boldsymbol{\mathcal { X }}=\llbracket \mathbf{I}_{M}, \mathcal{A}_{1}, \cdots, \mathcal{A}_{l} \rrbracket \text { where the cores are defined by }\left(\mathcal{A}_{k}\right)_{i,:, j}=\delta_{i j} \mathbf{x}_{k}^{(i)} \quad \text { and } \quad\left(\mathcal{A}_{l}\right)_{i,:}=\mathbf{x}_{k}^{(i)} \tag{4}
\end{equation*}
$$

for all $1 \leq k<l, i, j \in[M]$, where $\delta$ is the Kroencker symbol.
6: Perform the gradient update using efficient addition and product operations in TT format (see [33]):

$$
\left(\mathcal{H}^{(l)}\right)_{\langle\langle l, 1\rangle\rangle}=\left(\mathcal{H}^{(l)}\right)_{\langle\langle l, 1\rangle\rangle}+\gamma(\boldsymbol{\mathcal { X }})_{\langle\langle 1, l\rangle\rangle}^{\top}\left(\mathbf{Y}-(\boldsymbol{\mathcal { X }})_{\langle\langle 1, l\rangle\rangle}\left(\mathcal{H}^{(l)}\right)_{\langle\langle l, 1\rangle\rangle}\right)
$$

7: Project the Hankel tensor $\boldsymbol{\mathcal { H }}^{(l)}$ (which is now of rank at most $R+M$ ) back onto the manifold of tensor of TT-rank $R$ using the TT rounding operation (see again [33]):

$$
\mathcal{H}^{(l)}=\mathrm{TT} \text {-rounding }\left(\mathcal{H}^{(l)}, R\right)
$$

until convergence
9: Let $\mathbf{P}=\left(\llbracket \mathcal{G}_{1}^{(2 L)}, \cdots, \mathcal{G}_{L}^{(2 L)}, \mathbf{I}_{R} \rrbracket\right)_{\langle\langle L, 1\rangle\rangle}$ and $\mathbf{S}=\left(\llbracket \mathbf{I}_{R}, \mathcal{G}_{L+1}^{(2 L)}, \cdots, \mathcal{G}_{2 L+1}^{(2 L)} \rrbracket\right)_{\langle 1, L+1\rangle\rangle}$ (observe that $\left(\mathcal{H}^{(2 L)}\right)_{\langle\langle L, L+1\rangle\rangle}=\mathbf{P S}$ is a rank $R$ factorization).
10: Return the linear 2-RNN $\left(\mathbf{h}_{0}, \mathcal{A}, \boldsymbol{\Omega}\right)$ where

$$
\begin{aligned}
\boldsymbol{\alpha} & =\left(\mathbf{S}^{\dagger}\right)^{\top}\left(\mathcal{\mathcal { H }}_{f}^{(L)}\right)_{\langle L+1\rangle\rangle} \\
\boldsymbol{\mathcal { A }} & =\left(\left(\boldsymbol{\mathcal { H }}_{f}^{(2 L+1)}\right)_{\langle L L, 1, L+1\rangle\rangle}\right) \times_{1} \mathbf{P}^{\dagger} \times_{3}\left(\mathbf{S}^{\dagger}\right)^{\top} \\
\boldsymbol{\Omega}^{\top} & =\mathbf{P}^{\dagger}\left(\mathcal{H}_{f}^{(L)}\right)_{\langle L, 1\rangle\rangle}
\end{aligned}
$$

by performing efficient computations in TT format for the products [33] and pseudo-inverses (see e.g. [26]).

## D Real Data Experiment on Wind Speed Prediction

Besides the synthetic data experiments we showed in the paper, we have also conducted experiments on real data. The data that we use for the these experiments is from TUDelft ${ }^{8}$. Specifically, we use the data from Rijnhaven station as described in [30], which proposed a regression automata model and performed various experiments on the dataset we mentioned above. The data contains wind speed and related information at the Rijnhaven station from 2013-04-22 at 14:55:00 to 2018-10-20 at 11:40:00 and was collected every five minutes. To compare to the results in [30], we strictly followed the data preprocessing procedure described in the paper. We use the data from 2013-04-23 to 2015-10-12 as training data and the rest as our testing data. The paper uses SAX as a preprocessing method to discretize the data. However, as there is no need to discretize data for our algorithm, we did not perform this procedure. For our method, we set the length $L=3$ and we use the general algorithm described in Appendix B. We calculate hourly averages of the wind speed, and predict one/three/six hour(s) ahead, as in [30]. For our methods we use a linear 2-RNN with 10 states. Averages over 5 runs of this experiment for one-hour-ahead, three-hour-ahead, six-hour-ahead prediction error can be found in Table 1, 2 and 3. The results for RA, RNN and persistence are taken directly from [30]. We see that while TIHT+SGD performs slightly worse than ARIMA and RA for one-hour-ahead prediction, it outperforms all other methods for three-hours and six-hours ahead predictions (and the superiority w.r.t. other methods increases as the prediction horizon gets longer).

Table 1: One-hour-ahead Speed Prediction Performance Comparisons

| Method | TIHT | TIHT+SGD | Regression <br> Automata | ARIMA | RNN | Persistence |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| RMSE | 0.573 | 0.519 | 0.500 | $\mathbf{0 . 4 9 6}$ | 0.606 | 0.508 |
| MAPE | 21.35 | 18.79 | $\mathbf{1 8 . 5 8}$ | 18.74 | 24.48 | 18.61 |
| MAE | 0.412 | 0.376 | 0.363 | $\mathbf{0 . 3 6 1}$ | 0.471 | 0.367 |

Table 2: Three-hour-ahead Speed Prediction Performance Comparisons

| Method | TIHT | TIHT+SGD | Regression <br> Automata | ARIMA | RNN | Persistence |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| RMSE | 0.868 | $\mathbf{0 . 8 5 4}$ | 0.872 | 0.882 | 1.002 | 0.893 |
| MAPE | 33.98 | $\mathbf{3 1 . 7 0}$ | 32.52 | 33.165 | 37.24 | 33.29 |
| MAE | 0.632 | $\mathbf{0 . 6 2 4}$ | 0.632 | 0.642 | 0.764 | 0.649 |

Table 3: Six-hour-ahead Speed Prediction Performance Comparisons

| Method | TIHT | TIHT+SGD | Regression <br> Automata | ARIMA | RNN | Persistence |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| RMSE | 1.234 | $\mathbf{1 . 1 4 5}$ | 1.205 | 1.227 | 1.261 | 1.234 |
| MAPE | 49.08 | $\mathbf{4 4 . 8 8}$ | 46.809 | 48.02 | 47.03 | 48.11 |
| MAE | 0.940 | $\mathbf{0 . 8 6 5}$ | 0.898 | 0.919 | 0.944 | 0.923 |

[^1]
[^0]:    ${ }^{7}$ Note that the theorem can be adapted if such an integer $L$ does not exists (see supplementary material).

[^1]:    ${ }^{8}$ http://weather.tudelft.nl/csv/

