10. Proof of nesting properties [Fernandez2025, App. A.3] TCI.10

Below, we prove two important nesting properties that were stated in (TCI.8

> Pivots are nested w.r.t. T_ℓ if they are left-nested up to $\ell-1$ and right-nested up to $\ell+1$
> Then, the TCI form is exact on the one-dimensional slice T_ℓ
> $$\tilde{F}_{i_\ell \oplus (\sigma_\ell) \oplus j_{\ell+1}} \simeq [T_\ell]_{i_{\ell-1}\sigma_\ell j_{\ell+1}} = F_{i_\ell \oplus(\sigma_\ell)\oplus j_{\ell+1}} \qquad \mathcal{I}_0 < \cdots < \mathcal{I}_{\ell-1}\ \mathcal{S}_\ell\ \mathcal{J}_{\ell+1} > \ > \mathcal{J}_{\mathcal{L}+1} \tag{1}$$

> If pivots are fully nested, TCI form is exact on every T_ℓ and P_ℓ, i.e. on all slices used to construct it,
> thus it is an interpolation. (2)

For each ℓ define the matrices
$$A_\ell^\sigma = T_\ell^\sigma P_\ell^{-1}, \qquad B_\ell^\sigma = P_\ell^{-1} T_\ell^\sigma \tag{3}$$

$[A_\ell^{\sigma_\ell}]_{i_{\ell-1} i_\ell}$ = 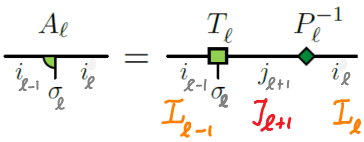 (4)

$[B_\ell^{\sigma_\ell}]_{j_\ell j_{\ell+1}}$ = 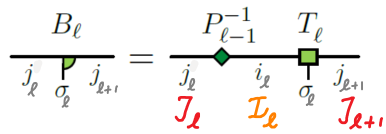

If left (row) indices of A or right (column) indices of B are restricted to pivots, they yield Kronecker symbols:

> If $i_{\ell-1}\oplus\sigma_\ell \in \mathcal{I}_\ell$ then $[A_\ell^{\sigma_\ell}]_{i_{\ell-1} i_\ell} = \delta_{i_{\ell-1}\oplus\sigma_\ell,\, i_\ell}$; if $\sigma_\ell \oplus j_{\ell+1} \in \mathcal{J}_\ell$ then $[B_\ell^{\sigma_\ell}]_{j'_\ell j_{\ell+1}} = \delta_{j_\ell,\, \sigma_\ell\oplus j_{\ell+1}}$
> left row index = pivot — right column index = pivot (5)

Reason: P_ℓ and $P_{\ell-1}$ are slices of T_ℓ :

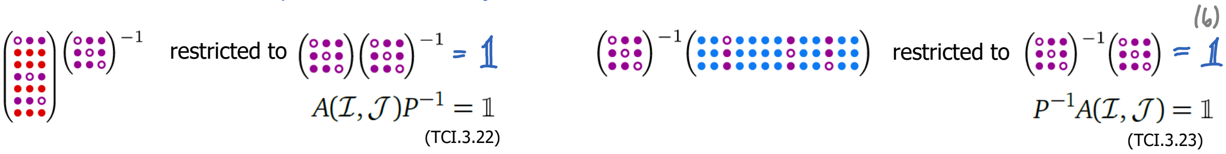 (6)

If pivots are left-nested up to ℓ, $\qquad \mathcal{I}_0 < \mathcal{I}_1 < \cdots < \mathcal{I}_{\ell'} < \cdots < \mathcal{I}_\ell$

and if i_ℓ is an index from a row pivot list, $\qquad$ if $i_\ell = (\underbrace{\sigma_1, \dots, \sigma_{\ell'}}, \dots, \sigma_\ell) \in \mathcal{I}_\ell$ (7)

then the same is true for any of its subindices, $i_{\ell'}$, for $\ell' < \ell$: $\qquad$ then $i_{\ell'} \in \mathcal{I}_{\ell'}$

(Because left-nested means: if you remove last index of an element of a row pivot list, $i_\ell \in \mathcal{I}_\ell$, you get an element of shorter pivot list, $i_{\ell-1} \in \mathcal{I}_{\ell-1}$.)

If pivots are right-nested up to ℓ, $\qquad \mathcal{J}_\ell > \cdot > \mathcal{J}_{\ell'} > \cdots > \mathcal{J}_\mathcal{L} > \mathcal{J}_{\mathcal{L}+1}$

and if j_ℓ is an index from a column pivot list, $\qquad$ if $j_\ell = (\sigma_\ell, \dots, \underbrace{\sigma_{\ell'}, \dots, \sigma_\mathcal{L}}) \in \mathcal{J}_\ell$ (8)

then the same is true for any of its subindices, $j_{\ell'}$, for $\ell' > \ell$: $\qquad$ then $j_{\ell'} \in \mathcal{J}_{\ell'}$

(Because right-nested means: if you remove first index of an element of a column pivot list, $j_\ell \in \mathcal{J}_\ell$, you get an element of shorter pivot list, $j_{\ell+1} \in \mathcal{J}_{\ell+1}$.)

Iterative use of (5), starting from $A_1 A_2$ or $B_{\mathcal{L}-1} B_\mathcal{L}$, then yields a telescope collapse:

If $\mathcal{I}_0 < \mathcal{I}_1 < \cdots < \mathcal{I}_\ell$ and $i_\ell = (\sigma_1, \dots, \sigma_\ell) \in \mathcal{I}_\ell$, then: $\qquad$ If $\mathcal{J}_\ell > \cdots > \mathcal{J}_\mathcal{L} > \mathcal{J}_{\mathcal{L}+1}$ and $j_\ell = (\sigma_\ell, \dots, \sigma_\mathcal{L}) \in \mathcal{J}_\ell$, then:

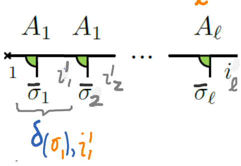 $= [A_1^{\sigma_1} \cdots A_\ell^{\sigma_\ell}]_{1 i_\ell} = \delta_{i'_\ell, i_\ell}$, $\qquad \delta_{j_\ell, j'_\ell} = [B_\ell^{\sigma_\ell} \cdots B_\mathcal{L}^{\sigma_\mathcal{L}}]_{j_\ell 1} =$ 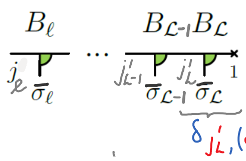 (9)

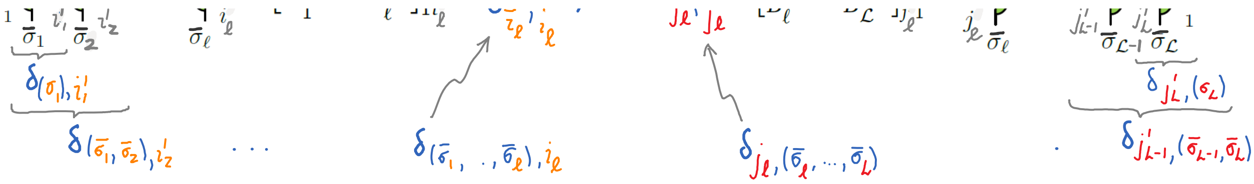

Important: such collapses do not apply for all configurations, only for pivots from left- or right-nested lists.

Thus, As and Bs are not isometries, $\sum_\sigma [A_\ell^{\sigma\dagger} A_\ell^{\sigma}]_{\bar i i'} \neq \delta_{\bar i i'}$, $\sum_\sigma [B_\ell^{\sigma} B_\ell^{\sigma\dagger}]_{j'\bar j} \neq \delta_{j'\bar j}$ (10)

because $\sum_\sigma$ -sum involve non-pivot configurations.

Now we are ready to prove three important facts:

1-site nesting w.r.t. T_ℓ : $\boxed{\text{If pivots are nested w.r.t. } T_\ell \text{ , then } \tilde F \text{ is exact on the slice } T_\ell \text{ .}}$ (11)

Proof: let $\bar\sigma \in \mathcal I_{\ell-1} \times \mathcal S_\ell \times \mathcal J_{\ell+1}$ be any configuration from which T_ℓ is built: $T_\ell = F(\mathcal I_{\ell-1}, \mathcal S_\ell, \mathcal J_{\ell+1})$

$\mathcal I_0 < \cdots < \mathcal I_{\ell-1}\ \mathcal S_\ell\ \mathcal J_{\ell+1} > \cdots > \mathcal J_{\mathcal L+1}$

(TCI.7.12)

$$\boxed{\tilde F_{\bar\sigma}} = [T_1^{\bar\sigma_1} P_1^{-1} \cdots T_{\ell-1}^{\bar\sigma_{\ell-1}} P_{\ell-1}^{-1} T_\ell^{\bar\sigma_\ell} P_\ell^{-1} T_{\ell+1}^{\bar\sigma_{\ell+1}} \cdots P_{\mathcal L-1}^{-1} T_{\mathcal L}^{\bar\sigma_{\mathcal L}}]_{11} \qquad (12)$$

$$\overset{(4)}{=} [A_1^{\bar\sigma_1} \cdots A_{\ell-1}^{\bar\sigma_{\ell-1}} T_\ell^{\bar\sigma_\ell} B_{\ell+1}^{\bar\sigma_{\ell+1}} \cdots B_{\mathcal L}^{\bar\sigma_{\mathcal L}}]_{11} \qquad (13)$$

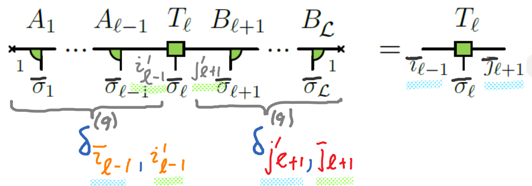

$$= [T_\ell^{\bar\sigma_\ell}]_{\bar\imath_{\ell-1}\bar\jmath_{\ell+1}} = \boxed{F_{\bar\sigma}} \qquad (14)$$

Through repeated use of (9), a sequence of telescope collapses pin all internal summations over primed indices to corresponding barred indices.

0-site nesting w.r.t. P_ℓ : $\boxed{\text{If pivots are nested w.r.t } P_\ell \text{ , then } \tilde F \text{ is exact on the slice } P_\ell \text{ .}}$ (14)

Proof: since $\mathcal I_{\ell-1} < \mathcal I_\ell$ and $\mathcal J_{\ell+1} > \mathcal J_{\ell+2}$, P_ℓ is a subslice of both T_ℓ and $T_{\ell+1}$. $P_\ell = F(\mathcal I_\ell, \mathcal J_{\ell+1})$

$\mathcal I_0 < \cdots < \mathcal I_\ell\ \ \mathcal J_{\ell+1} > \cdots > \mathcal J_{\mathcal L+1}$

But $\tilde F$ is exact on both, hence $\tilde F$ is exact on P_ℓ .

Moreover, $\tilde F(\mathcal I_\ell, \mathcal J_{\ell+1})$, viewed as a matrix with elements $[\tilde F]_{i_\ell, j_{\ell+1}} = [P_\ell]_{i_\ell, j_{\ell+1}}$, has

$$\boxed{\text{rank}[\ \tilde F(\mathcal I_\ell, \mathcal J_{\ell+1})] = \dim(P_\ell) = \chi_\ell} \qquad (15)$$

2-site nesting w.r.t. Π_ℓ : $\boxed{\text{If pivots are nested w.r.t } \Pi_\ell \text{ , then the local and global errors on that slice are equal.}}$

Proof: let $\bar\sigma \in \mathcal I_{\ell-1} \times \mathcal S_\ell \times \mathcal S_{\ell+1} \times \mathcal J_{\ell+2}$ be any configuration from which Π_ℓ is built: $\Pi_\ell = F(\mathcal I_\ell, \mathcal S_\ell, \mathcal S_{\ell+1}, \mathcal J_{\ell+2})$

$\mathcal I_0 < \cdots < \mathcal I_{\ell-1}\ \mathcal S_\ell\ \ \mathcal S_{\ell+1}\ \mathcal J_{\ell+2} > \cdots > \mathcal J_{\mathcal L+1}$

Then, by definition: (TCI.7.11) $F_{\bar\sigma} = [\Pi_\ell]_{\bar\sigma}$ (16)

Telescope like (12)-(14) yields: (TCI.7.12)

$$\tilde F_{\bar\sigma} = [\underbrace{T_1^{\bar\sigma_1} P_1^{-1} \cdots T_{\ell-1}^{\bar\sigma_{\ell-1}} P_{\ell-1}^{-1}}_{\text{telescope collapse}} T_\ell^{\bar\sigma_\ell} P_\ell^{-1} T_{\ell+1}^{\bar\sigma_{\ell+1}} \underbrace{\cdots P_{\mathcal L-1}^{-1} T_{\mathcal L}^{\bar\sigma_{\mathcal L}}}_{\text{telescope collapse}}]_{11} \qquad (17)$$

$$= [T_\ell^{\bar\sigma_\ell} P_\ell^{-1} T_{\ell+1}^{\bar\sigma_{\ell+1}}]_{\bar\imath_{\ell-1}, \bar\jmath_{\ell+2}} \overset{(\text{TCI.9.3})}{=} [\tilde\Pi_\ell]_{\bar\sigma} \qquad (18)$$

(16) - (18): $$\boxed{[\Pi_\ell - \tilde\Pi_\ell]_{\bar\sigma} = [F - \tilde F]_{\bar\sigma}} \qquad (19)$$

local error global error

$\Longrightarrow$ $\boxed{\text{Local update reducing local error will also reduce the global error!}}$ [cf. (TCI.9.6)] (20)

11. CI-canonicalization [Fernandez2025, Sec. 4.5.1]

TCI.11

Any tensor train can be transformed exactly into TCI form at costs $\mathcal{O}(\chi^3)$, described uniquely in terms of pivot lists and corresponding slices of the tensor train. The TCI form corresponds to a particular choice of gauge.

Starting point: an arbitrary MPS

$$F_{\boldsymbol{\sigma}} = [M_1^{\sigma_1}]_{1a_1}[M_2^{\sigma_2}]_{a_1a_2}\cdots[M_{\mathcal{L}}^{\sigma_{\mathcal{L}}}]_{a_{\mathcal{L}-1}1} = \qquad (1)$$

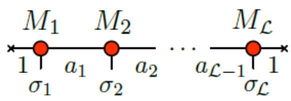

ordinary MPS indices, not multi-indices $a_\ell \neq (\sigma_1, \ldots, \sigma_\ell)$

First forward sweep: (swallow up a_ℓ indices, generated left-nested row pivot lists)

Initialize: do exact CI-decomposition of M_1 : (2)

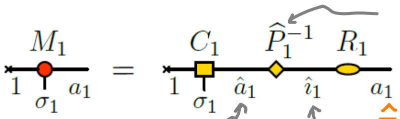

not a slice of F, since $\hat{a}_1 \neq$ multi-index

not multi-index $\quad \in \hat{\mathcal{I}}_1 = \{\sigma_1\}$ = multi-index!

Insert (2) into (1): $F_{\boldsymbol{\sigma}} =$ (3)

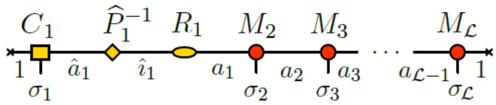

Iterate for $\ell \geq 2$:

Reshape, do exact CI-decomposition, define A_ℓ tensors (which 'swallow up non-multi-indices via internal summations): (4)

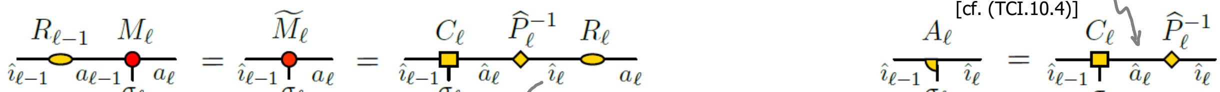

$\hat{\mathcal{I}}_{\ell-1}$ $\quad \hat{\mathcal{I}}_\ell = (\sigma_1, \ldots, \sigma_\ell)$

[cf. (TCI.10.4)]

= multi-indices! $\quad \hat{\mathcal{I}}_{\ell-1} \quad \hat{\mathcal{I}}_\ell$

row indices $\hat{\mathcal{I}}_\ell \subseteq \hat{\mathcal{I}}_{\ell-1} \times \mathcal{S}_\ell$ are nested by construction: $\hat{\mathcal{I}}_{\ell-1} < \hat{\mathcal{I}}_\ell$ column indices: $\hat{a}_\ell$, a_ℓ = not multi-indices! (5)

After full forward sweep: $F_{\boldsymbol{\sigma}} =$ (6)

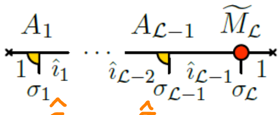

$\hat{\mathcal{I}}_0 < \cdots < \hat{\mathcal{I}}_{\mathcal{L}-1}$

all non-multi-indices have been 'swallowed'

row pivot lists are fully left-nested

Telescope collapse property [cf. (TCI.10.9)]:

If $\hat{\mathcal{I}}_0 < \hat{\mathcal{I}}_1 < \cdots < \hat{\mathcal{I}}_\ell$ and $\bar{i}_\ell = (\bar{\sigma}_1, \ldots, \bar{\sigma}_\ell) \in \mathcal{I}_\ell$, then:

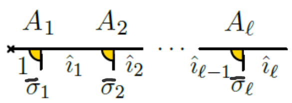

$= \delta_{\bar{i}_\ell \hat{i}_\ell}$ if $\hat{i}_\ell \in \mathcal{I}_\ell$ (7)

Therefore, $\widetilde{M}_\mathcal{L}$ is a slice of F : $\quad \widetilde{M}_\mathcal{L} = F(\hat{\mathcal{I}}_{\mathcal{L}-1}, \mathcal{S}_\mathcal{L})$ [cf. (TCI.10.14)] (8)

All C_ℓ and P_ℓ have full rank when viewed as matrices $[C_\ell]_{(\hat{i}_{\ell-1}\oplus\sigma_\ell),\hat{a}_\ell}$ and $[P_\ell]_{\hat{i}_\ell,\hat{a}_\ell}$ (9)

However, C_ℓ and $\widetilde{M}_\mathcal{L}$ may still be rank-deficient when viewed as matrices $[C_\ell]_{\hat{i}_{\ell-1},(\sigma_\ell\oplus\hat{a}_\ell)}$ or $[\widetilde{M}_\ell]_{\hat{i}_{\mathcal{L}-1},\sigma_\mathcal{L}}$ (10)

Backward sweep: (generate right-nested column pivot lists)

Initialize: do exact CI-decomposition of $\widetilde{M}_\mathcal{L}$: (11)

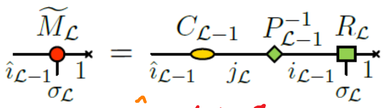

$\hat{\mathcal{I}}_{\mathcal{L}-1}\{\sigma_\mathcal{L}\} = \mathcal{J}_\mathcal{L} \quad \mathcal{I}_{\mathcal{L}-1} \subseteq \hat{\mathcal{I}}_{\mathcal{L}-1}$ all multi-indices!

$R_\mathcal{L}$ and $P_\mathcal{L}$, being subslices of $\widetilde{M}_\mathcal{L}$, are slices of F, namely: $R_\mathcal{L} = F(\mathcal{I}_{\mathcal{L}-1}, \mathcal{S}_\mathcal{L})$ and $P_\mathcal{L} = F(\mathcal{I}_{\mathcal{L}-1}, \mathcal{J}_\mathcal{L})$ (12)

Make identification $T_\mathcal{L} = R_\mathcal{L}$ (13)

Iterate for $\ell \le \mathcal{L}-1$:

Reshape, do exact CI-decomposition, define $T_\ell = R_\ell$ tensors (which are slices of F) and B_ℓ :

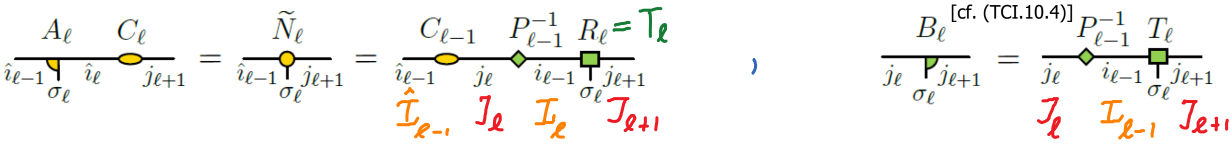

(14)

row indices: $\mathcal{I}_{\ell-1} \subseteq \hat{\mathcal{I}}_{\ell-1}$ column indices $\mathcal{J}_\ell \subseteq \mathbb{S}_\ell \times \mathcal{J}_{\ell+1}$ are nested by construction: $\mathcal{J}_\ell > \mathcal{J}_{\ell+1}$ (15)

left-nesting may be broken if pivots are discarded

R_ℓ is a slice of F (will be demonstrated below), hence we rename it $T_\ell = R_\ell$ (16)

After backward sweep up to site ℓ, and further all the way to site $\mathcal{L}$:

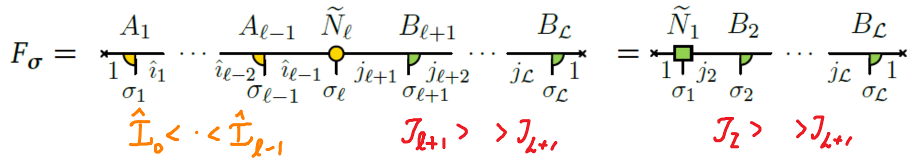

(17)

column pivot lists are fully right-nested

Telescope collapse property [cf. (TCI.10.9)]:

If $\mathcal{J}_{\ell+1} > \; > \mathcal{J}_{\mathcal{L}+1}$ and $\bar{j}_\ell = (\bar{\sigma}_\ell, .., \bar{\sigma}_\mathcal{L})$, then:

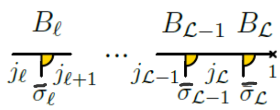

$= \delta_{j_\ell \bar{j}_\ell}$ if $j_\ell \in \mathcal{J}_\ell$ (18)

Therefore, $\tilde{N}_\ell$ is a slice of F : $\tilde{N}_\ell = F(\hat{\mathcal{I}}_{\ell-1}, \mathbb{S}_\ell, \mathcal{J}_{\ell+1})$ [cf. (TCI.10.14)] (19)

If $\bar{\sigma} \in \hat{\mathcal{I}}_{\ell-1} \times \mathbb{S}_\ell \times \mathcal{J}_{\ell+1}$ then

telescope collapses:

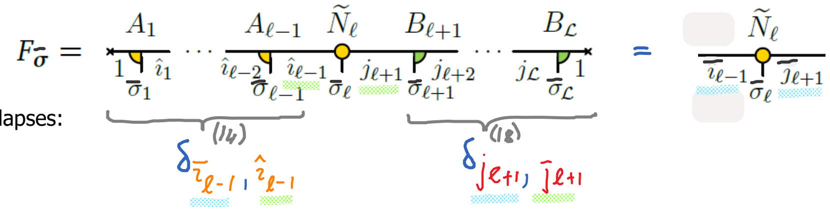

(20)

Similarly, R_ℓ and $P_{\ell-1}$, being subslices of $\tilde{N}_\ell$, are slices of F: $T_\ell = R_\ell = F(\mathcal{I}_{\ell-1}, \mathbb{S}_\ell, \mathcal{J}_{\ell+1})$ and $P_{\ell-1} = F(\mathcal{I}_{\ell-1}, \mathcal{J}_\ell)$ (21)

Thus, CI-decomposition of $\tilde{N}_\ell$ reveals bond dimension of F for bond $\ell-1$, namely $\boxed{\chi_{\ell-1} = |\mathcal{I}_{\ell-1}| = |\mathcal{J}_\ell|}$ (22)

Finally, telescope collapse shows that $\tilde{N}_1 = F(\mathbb{S}_1, \mathcal{J}_2)$ is a slice of F , too, so we identify $T_1 = \hat{N}_1$ (23)

Using $B_\ell \overset{(14)}{=} P^{-1}_{\ell-1} T_\ell$ in (17), we arrive at TCI form: $F_\sigma = [T_1^{\sigma_1} P_1^{-1} T_2^{\sigma_2} \cdots P_{\mathcal{L}-1}^{-1} T_\mathcal{L}^{\sigma_\mathcal{L}}]_{11}$ (24)

Here, all ingredients are slices of F , labeled by multi-indices.

Each T_ℓ is full rank for both ways of viewing it as a matrix, $[T_\ell]_{(i_{\ell-1},\sigma_\ell), j_{\ell+1}}$ or $[T_\ell]_{i_{\ell-1},(\sigma_\ell, j_{\ell+1})}$ (25)

After full backward sweep, all column pivots are fully right-nested.

But row pivots may not be fully left-nested, since backward sweep may have discarded some row pivots.

To restore full left-nesting, do one more exact forward sweep, using 1-site TCI algorithm (explained in TCI.9).

This will not break right-nesting of columns, since no pivots will be discarded. Final result is a <u>fully nested</u> TCI form.

12. High-level TCI algorithms [Fernandez2025, Sec. 4.5.2] TCI.12

The various TCI algorithms (2-site, 1-site, 0-cite), decomposition options (CI, prrLU), and pivot search options (full, rook, block rook, addition of global pivots) can be combined in various ways.

- 2-site TCI in *accumulative* plus *rook pivoting* mode is the fastest technique. It requires the least pivot exploration and very often provides very good results on its own. The accuracy can be improved, if desired, by following this with a few (cheap) 1-site TCI sweeps to reset the pivots.
- 2-site TCI in *reset* plus *rook pivoting* mode is marginally more costly than the above but more stable. It is a good default. For small d, one should use the *full* search, which is even more stable and involves almost no additional cost if $d \le 2n_{\text{rook}}$.
- If good heuristics for proposing pivots are available or ergodicity issues arise, one should consider switching to *global pivot proposal* followed by 2-site TCI.
- To obtain the best final accuracy at fixed χ, one can build a TCI with a higher rank $\chi' > \chi$, then compress it using either SVD or CI recompression.
- For calculations of integrals or sums, we recommend the environment mode. [see (TCI.9.10)] In some calculations, we have observed it to increase the accuracy by two digits for the same computational cost.

Table 2: Computational cost of the main TCI algorithms in `xfac` / `tci.jl`.

action	variant		calls to F_σ	algebra cost
iterate	rook piv.	2-site	$\mathcal{O}(\chi^2 d n_{\text{rook}} \mathcal{L})$	$\mathcal{O}(\chi^3 d n_{\text{rook}} \mathcal{L})$
	full piv.	2-site	$\mathcal{O}(\chi^2 d^2 \mathcal{L})$	$\mathcal{O}(\chi^3 d^2 \mathcal{L})$
	full piv.	1-site	$\mathcal{O}(\chi^2 d \mathcal{L})$	$\mathcal{O}(\chi^3 d \mathcal{L})$
	full piv.	0-site	0	$\mathcal{O}(\chi^3 \mathcal{L})$
achieve full nesting			$\mathcal{O}(\chi^2 d \mathcal{L})$	$\mathcal{O}(\chi^3 d \mathcal{L})$
add n_p global pivots			$\mathcal{O}\big((2\chi + n_p) n_p \mathcal{L}\big)$	$\mathcal{O}\big((\chi + n_p)^3 \mathcal{L}\big)$
compress tensor train	SVD		0	$\mathcal{O}(\chi^3 d \mathcal{L})$
	LU			
	CI			

Operations on tensor trains

Function composition: $g(f(x))$ construct another TCI: $\tilde{G}_{\vec\sigma} = g(\tilde{F}_{\vec\sigma})$ (1)

Initialize $\tilde{G}_{\vec\sigma}$ using pivots of $\vec{F}_{\vec\sigma}$ then, applying g to each element of T_ℓ (slice of $\vec{F}_\sigma$)

Subsequently, optimize $\tilde{G}_{\vec\sigma}$ using 2-site TCI algorithm.

Element-wise tensor addition: given $\widetilde{F} = M_1 M_2 \cdots M_{\mathcal{L}}$ and $\widetilde{F}' = M_1' M_2' \cdots M_{\mathcal{L}}'$ (2)

element-wise sum: $\widetilde{F}''_\sigma = \widetilde{F}_\sigma + \widetilde{F}'_\sigma = \mathrm{Tr}(M_1''^{\sigma_1} M_2''^{\sigma_2} \cdots M_{\mathcal{L}}''^{\sigma_{\mathcal{L}}})$, $\quad M_\ell''^{\sigma_\ell} = \begin{pmatrix} M_\ell^{\sigma_\ell} & 0 \\ 0 & M_\ell'^{\sigma_\ell} \end{pmatrix}$ (3)

element-wise sum: $F_\sigma = F_\sigma + F'_\sigma = \Pi(M_1 \; M_2 \cdots M_{\mathcal{L}})$, $M_\ell = \begin{pmatrix} \cdot & \\ 0 & M'^{\sigma_\ell}_\ell \end{pmatrix}$ (3)

Then compress using CI-canonicalization algorithm (see TCI.11). Runtime costs: $\mathcal{O}((\chi + \chi')^3 d \mathcal{L})$ (4)

ranks of $\tilde{F}$, $\tilde{F}'$

Convolution = matrix-vector contraction: (5)

$$\int dx\, g(x', x) f(x) \simeq \sum_{\sigma} G_{\sigma'\sigma} F_{\sigma} \simeq \sum_{\sigma} \tilde{G}_{\sigma'\sigma} \tilde{F}_{\sigma} =$$

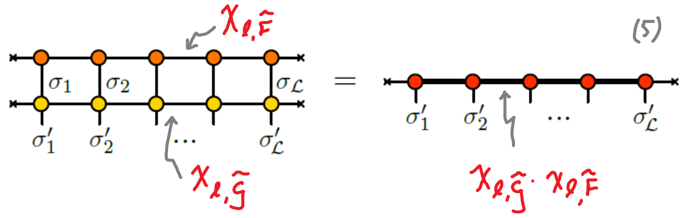

Standard tools for compressing the result:
(i) Fitting exact result to MPS with reduced bond dimensions;
(ii) Zip-up compression, where MPO-MPS contraction is performed one site at a time.

For both cases, one can use either SVD or CI/prrLU (favorable for large d, when one can use rook search)

Runtime costs of both: $\mathcal{O}(\chi^4 \mathcal{L} d^2)$ for $\chi_{\tilde{G}} = \chi_{\tilde{F}} = \chi_{\tilde{G}\tilde{F}} = \chi$

The χ^4 scaling is currently a dominant bottleneck! Mitigation attempts include

- parallelization: use different 'workers' to treat different parts of MPO-MPS contraction; combine results at the end [Stoudenmire2013].

- 'patching': divide domain of function into different patches, use different resolutions for different patches, according to needs. Adapt patch sizes dynamically while learning TCI decomposition. [Grosso2025]

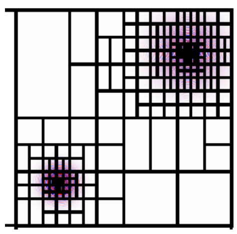

13. Relation to machine learning

[Fernandez2025, Sec. 4.8]

TCI.13

Traditional machine learning approach to 'learning' a compressed representation $\tilde{F}_{\vec{\sigma}}$ of $F_{\vec{\sigma}}$:

(i) draw 'training set' of configurations/values: $\{\vec{\sigma}, F_{\vec{\sigma}}\}$

(ii) design a 'model' $\tilde{F}_{\vec{\sigma}}$ (e.g. deep neural network);

(iii) fit model to training set by minimizing error $\|F - \tilde{F}\|$ w.r.t. some norm, typically using stochastic gradient descent.

(iv) use model to evaluate $\tilde{F}_{\vec{\sigma}}$ for new configurations.

TCI implements this program with some important differences / special features:

(i) TCI does not use a 'training set'; instead it actively requests configurations likely to bring most new information ('active learning').

(ii) The 'model' is not a neural network but a tensor train (highly structured model). <u>If</u> F is compressible, it can be approximated by low-rank $\tilde{F}$, with exponentially smaller memory footprint. Learning requires $\ll d^{L}$ samples.

(iii) The TCI learning algorithm used to minimize error is very different from steepest descent. It guarantees error is smaller than specified tolerance τ for all known samples.

(iv) Once $\tilde{F}$ has been found, its elements $\tilde{F}_{\vec{\sigma}}$ can be computed for <u>all</u> configurations $\vec{\sigma}$.
This is useful if function calls to $F_{\vec{\sigma}}$ are expensive. (Then learning $\tilde{F}$ is expensive, but calling $\tilde{F}_{\vec{\sigma}}$ is cheap.)
Moreover, subsequent operations on $\tilde{F}$ can be performed exponentially faster than on F.