



SINGD: Structured Inverse-Free Natural Gradient Descent

Wu Lin*, Felix Dangel*, Runa Eschenhagen, Kirill Neklyudov, Agustinus Kristiadi, Richard Turner, Alireza Makhzani



We present a **KFAC-style 2nd-order method** for DL that is **memory-efficient & numerically-stable in bfloat16** thanks to structured Kronecker factors and an inverse-free update rule.

SINGD = Inverse-free KFAC + Structures

Goal: Approximate natural gradient $\mathbf{F}^{-1}\mathbf{g}$ with gradient \mathbf{g} , inverse Fisher approximation \mathbf{F}^{-1} , and damping λ .

$$\mathbf{F}_{(\text{KFAC})}^{-1} = (\mathbf{S}_K + \lambda\mathbf{I})^{-1} \otimes (\mathbf{S}_C + \lambda\mathbf{I})^{-1}$$

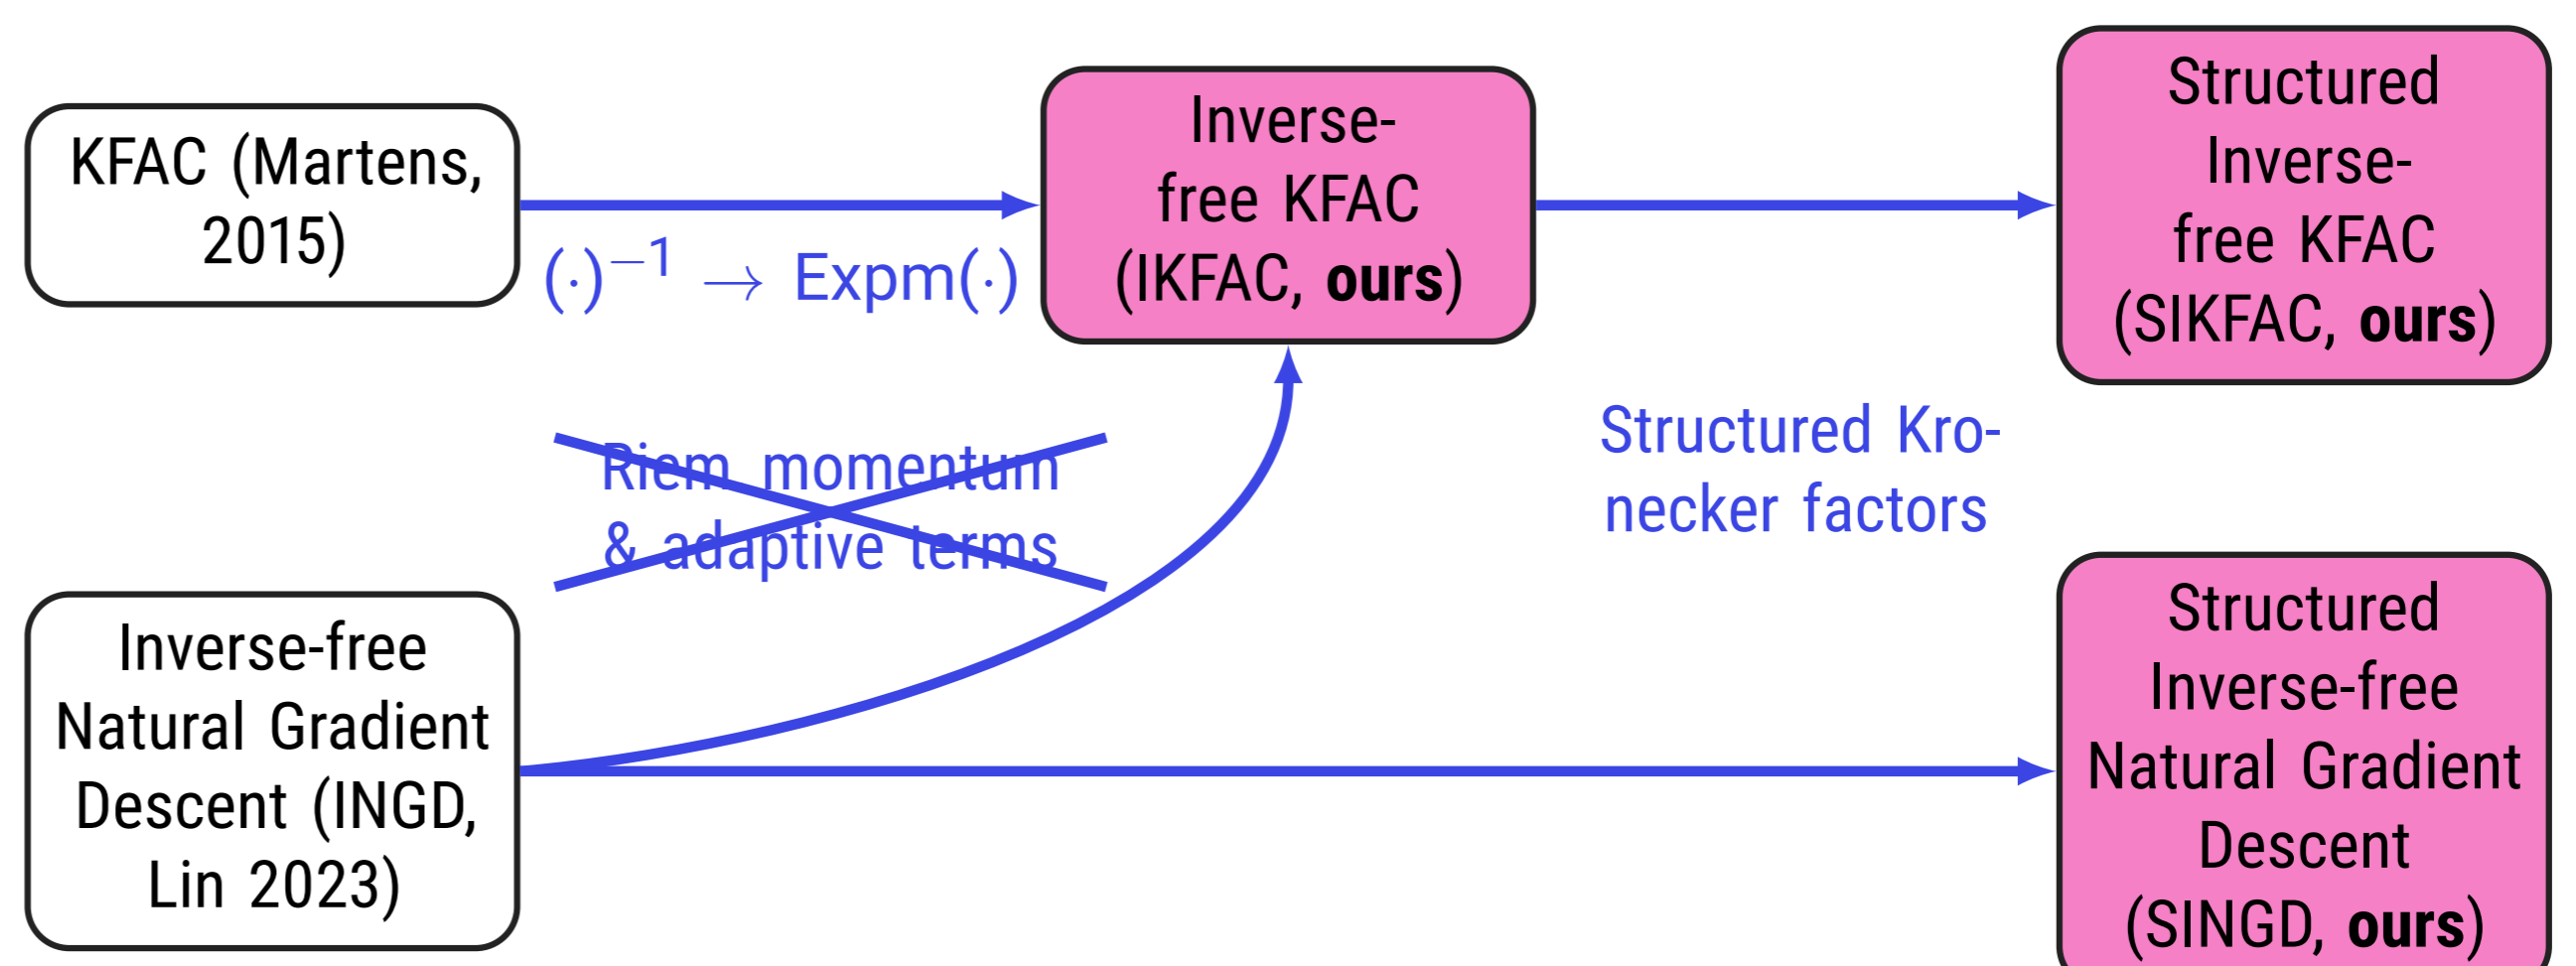
vs.

$$\mathbf{F}_{(\text{INGD})}^{-1} = \mathbf{K}\mathbf{K}^T \otimes \mathbf{C}\mathbf{C}^T$$

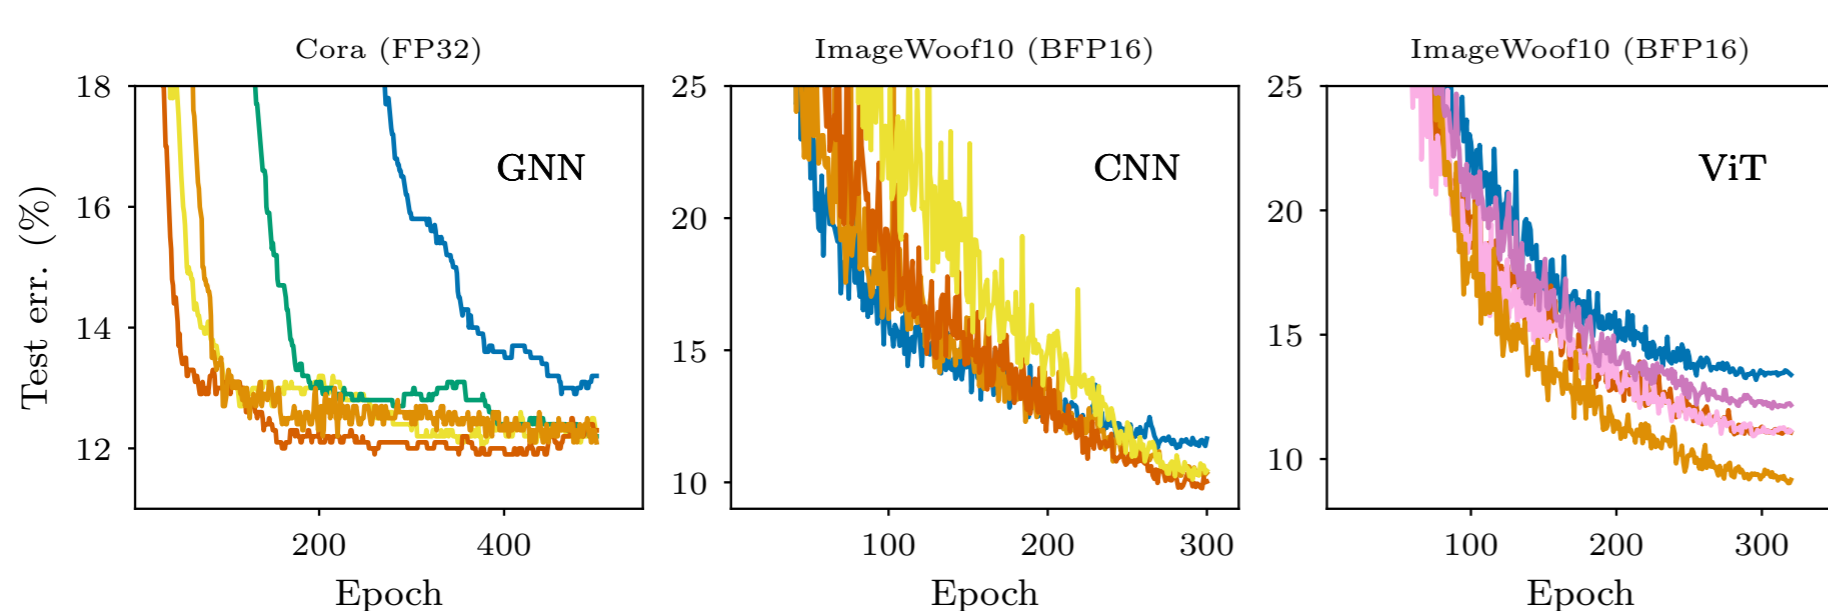
vs.

$$\mathbf{F}_{(\text{SINGD})}^{-1} = \mathbf{K}\mathbf{K}^T \otimes \mathbf{C}\mathbf{C}^T,$$

with structured \mathbf{K}, \mathbf{C} .

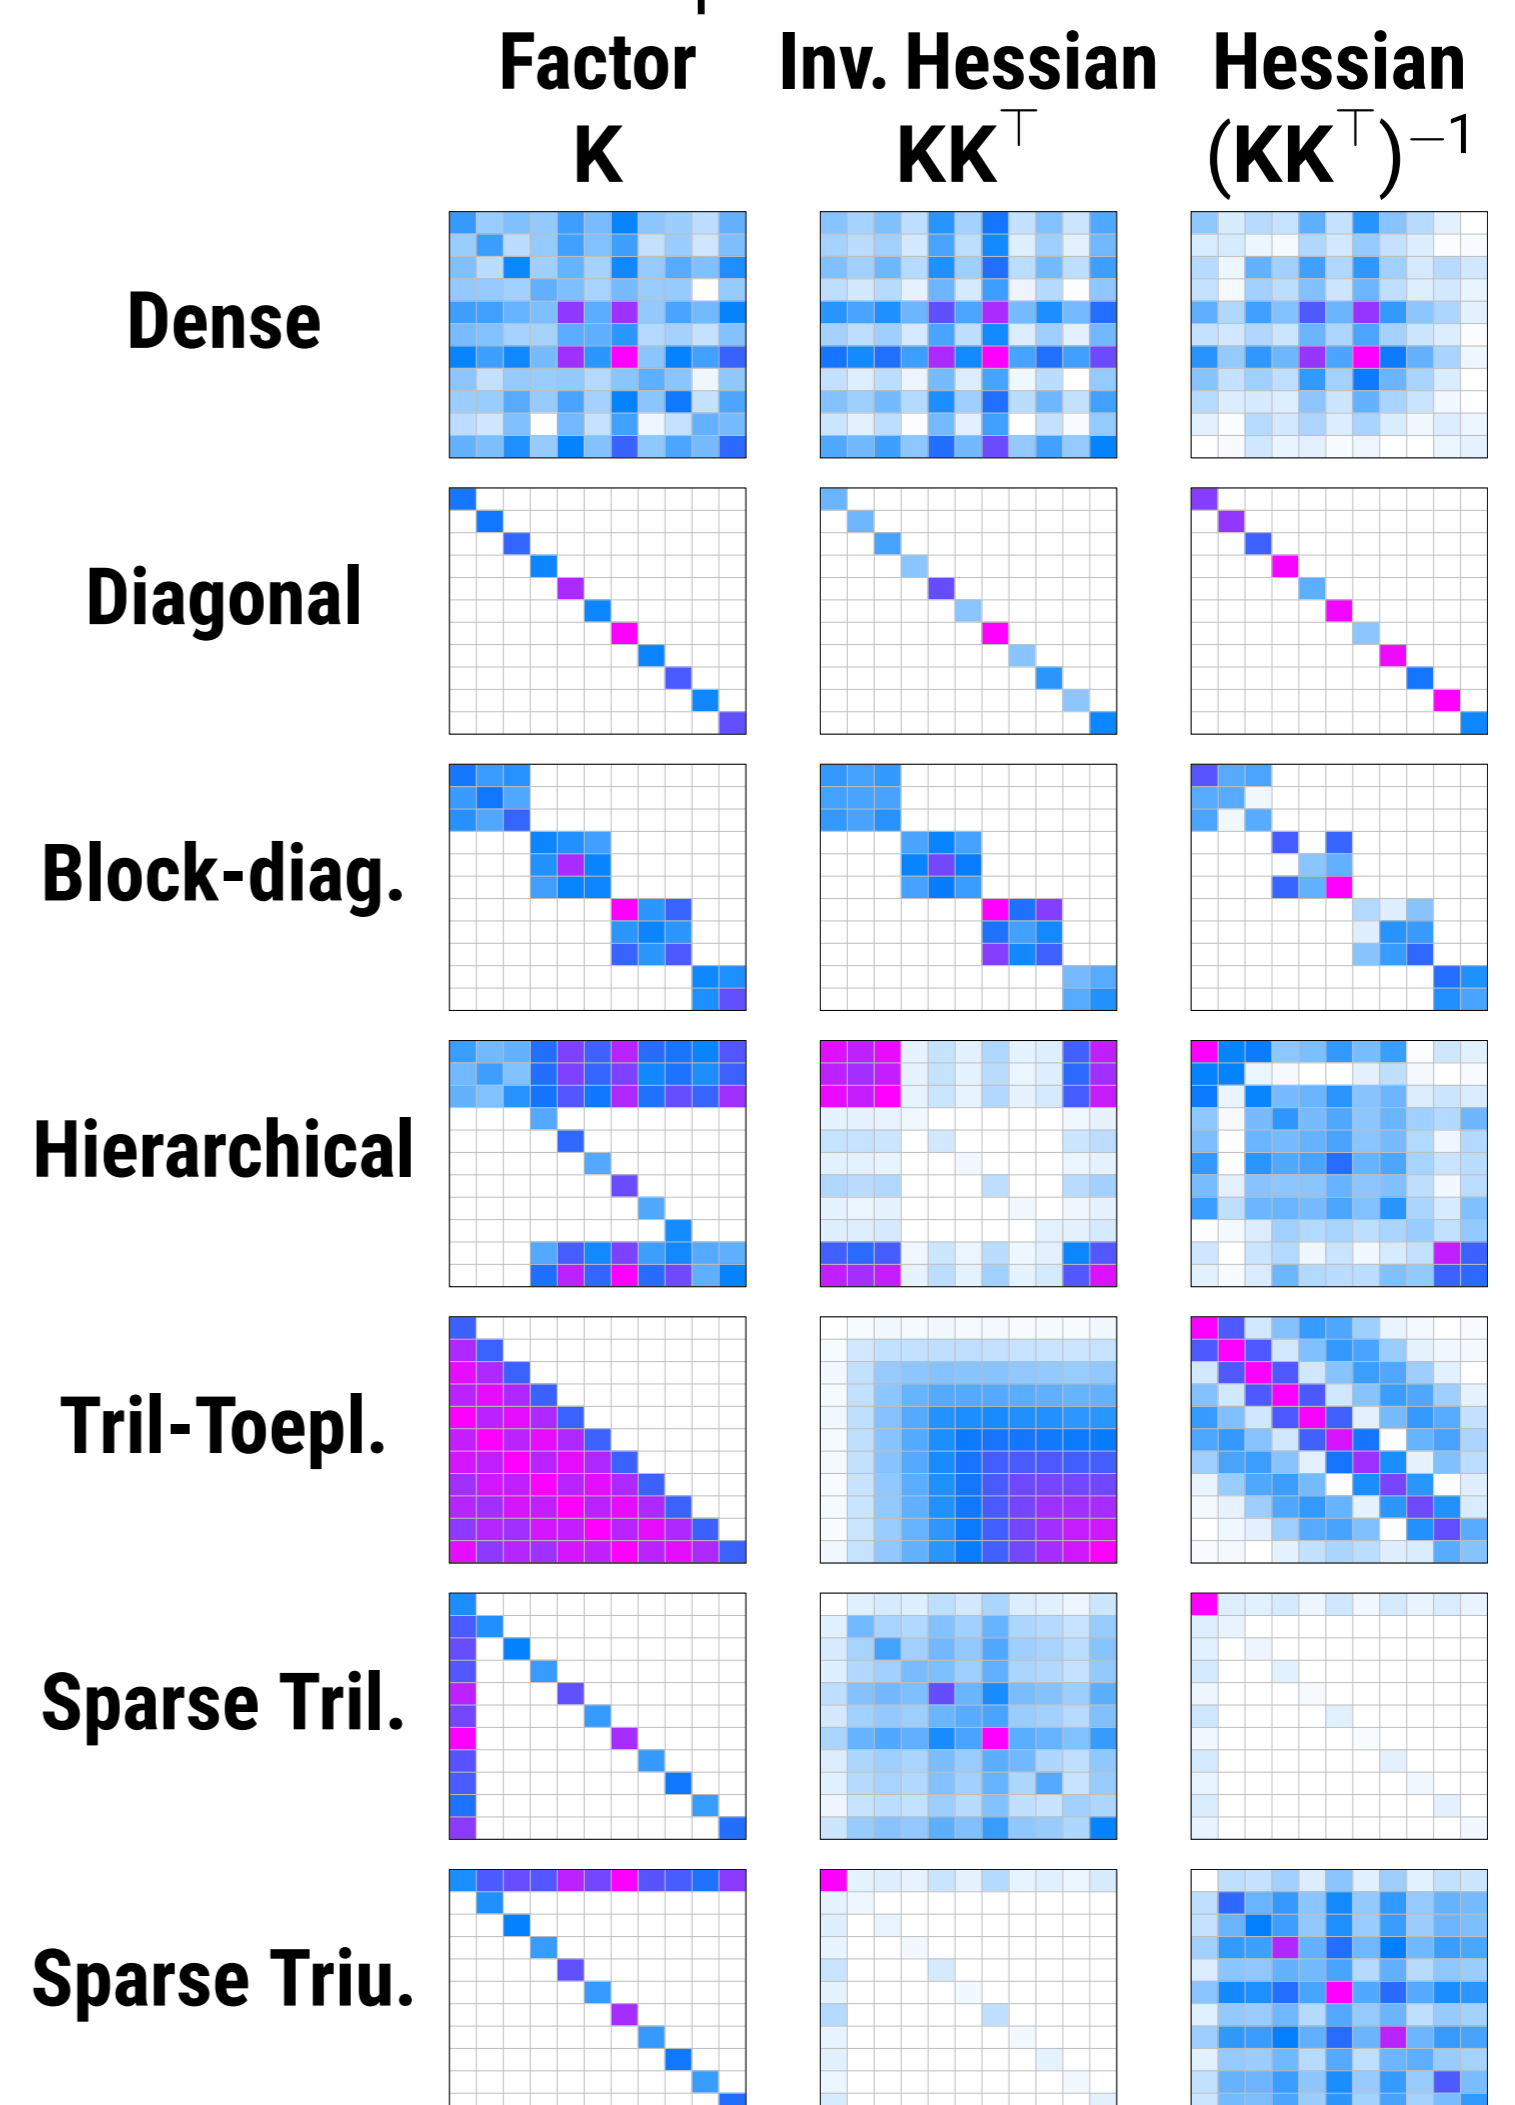


Performs Well on Many Problems & in bfloat16



Available Structures

Different structures capture different information.



Works like any Optimizer (pip install singd)

```

1 from singd.optim.optimizer import SINGD
2 from torch_utils import CrossEntropyLoss, alexnet, new_batch
3
4 # 1. Setup
5 model, loss_func = alexnet(), CrossEntropyLoss()
6 X, y = new_batch()
7
8 optim = SINGD( # Optimizer setup
9     model,
10    kfac_like=False, # Use True for IKFAC
11    structures=("dense", "diagonal"), # Structure of (K, C)
12 )
13
14 # 2. One training step
15 optim.zero_grad() # Clear previous gradients
16 loss_func(model(X), y).backward() # Forward + backward pass
17 optim.step() # Update pre-conditioner and parameters
    
```