

# weight | watcher Cheat Sheet

weightwatcher.ai | Data-free Diagnostics for Deep Learning

## Install + minimal usage

```
pip install weightwatcher
```

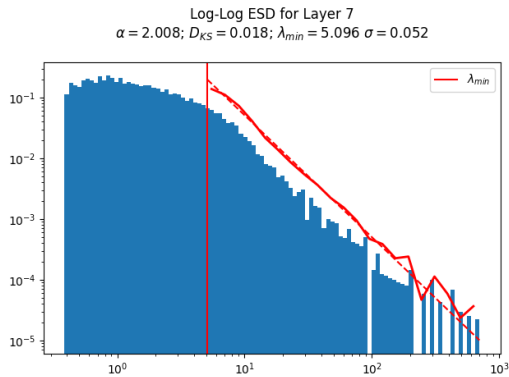
```
import weightwatcher as ww
watcher = ww.WeightWatcher(model=...)
df = watcher.analyze(...)
```

Model input: Keras or PyTorch model in memory, or a folder containing model files (.bin, .safetensors).

## Quick interpretation (HTSR $\alpha$ )

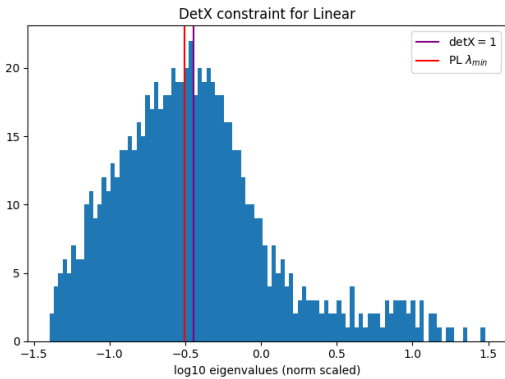
- $\alpha < 2$  **warning**: overfitting or memorization.
- $\alpha \approx 2$  **ideal**: best generalization.
- $\alpha \in (2, 4-6)$  **suboptimal**: correlated but not optimal
- $\alpha > 6$  **randomlike**: noise-dominated.

Rule: push  $\alpha$  toward 2 without inducing correlation traps. Improve layers by adjusting layer learning rates, regularization, and just training longer without overfitting



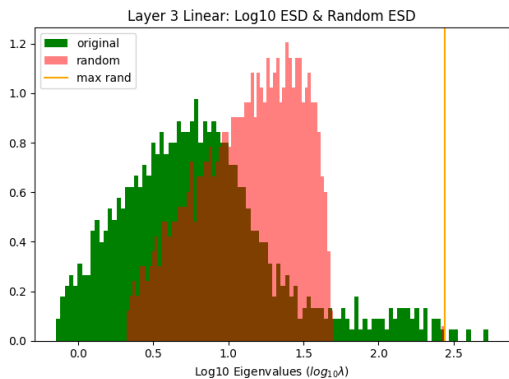
1) Standard ESD / Power-Law fit per-layer check  
Per-layer ESD + PL fit;  $\alpha, D_{KS}, \lambda_{min}$ .

```
df = watcher.analyze(plot=True)
df.alpha, df.D, df.xmin
```



2) SETOL ERG condition: check ideal performance  
Red and purple lines overlap when  $\alpha \approx 2$ .

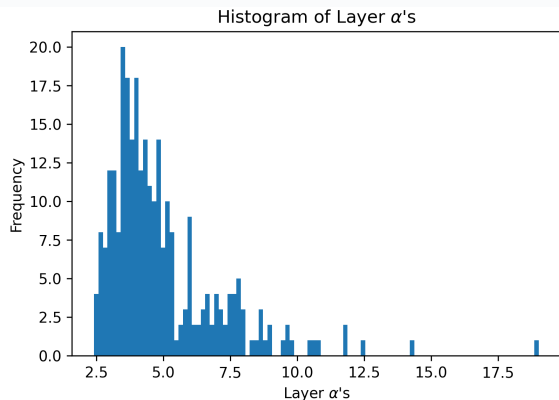
```
df = watcher.analyze(plot=True, detX=True)
df.erg_check = df.detX_num - df.num_pl_spikes
```



3) Correlation Traps

Randomized (red) ESD  $\approx$  MP bulk.  
Large outliers (orange) expose overfitting.

```
df = watcher.analyze(plot=True, randomize=True)
df.num_traps = df.rand_num_spikes
```



4) Large models: Histogram of  $\alpha$

Aggregate view when per-layer plots don't scale.  
Spot outliers, fat tails, and layer clusters fast.

```
df = watcher.analyze(...)
df.alpha.plot.hist(bins=100)
```

## More analyze() options:

Find Dragon Kings / Fix instabilities in  $\alpha$ :  
`watcher.analyze(fix_fingers='clip_xmax')`

Subtract Base model from Fine-Tuned model:  
`watcher.analyze(model="...", base_model="...")`

Eigenvector entropy & localization metrics:  
`watcher.analyze(vectors=True)`

Evaluate PEFT / LoRA adapter\_model.bin:  
`watcher.analyze(peft=True)`

Links: GitHub: [WeightWatcher](#) | Examples: [WeightWatcher-Examples](#) | WeightWatcher Pro: [weightwatcher-ai.com](#)

Papers: HTSR: [JMLR 2021](#) | WeightWatcher: [Nature 2021](#) | SETOL: [arXiv 2025](#) | Grokking: [ICML 2025](#)

Talks: [Silicon Valley ACM 2022](#) | [NeurIPS 2023](#) | [ICCF 24](#) | [TWIML 2024](#)

Author: [Charles H. Martin, PhD \(LinkedIn\)](#) | Community: [Discord](#)