tent: fully test-time adaptation by entropy minimization

ICLR'21 https://github.com/DequanWang/tent

⛺ helps a **model adapt itself** to changing conditions ☀️ 🌧️ ❄️
by updating on new and different data during testing
without altering training or requiring more supervision





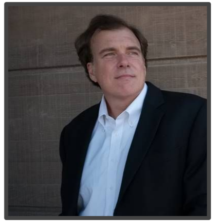

Dequan Wang*, Evan Shelhamer*, Shaoteng Liu, Bruno Olshausen, Trevor Darrell

Adobe BAIR BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

Fully Test-Time Adaptation: The Purpose

Adaptation to Natural Shifts How to reduce generalization error on shifted data corruptions, simulation-to-real discrepancies, and other shifts

Fully Test-Time? Adaptation during testing without relying on training data, offline optimization, … only the model and target data needed

Fully Test-Time Adaptation: The Need

how to generalize to new & different data **during testing?**

Hendrycks & Dietterich ICLR'19

1. Availability. A model might be distributed without source data for bandwidth, privacy, or profit.
2. Efficiency. It might not be computationally practical to (re-)process source data during testing.
3. Accuracy. A model might be too inaccurate without adaptation to serve its purpose.

ImageNet-C benchmark:

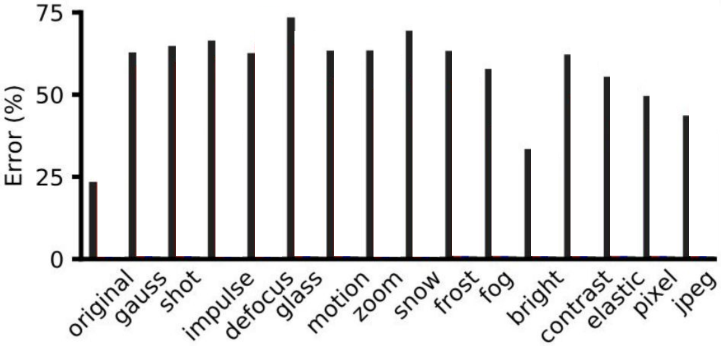

Fully Test-Time Adaptation: Our Method

in brief: **minimize entropy** of model predictions by **updating statistics and just a few parameters**
tent, for *test entropy* minimization

- training optimizes parameter θ
  to minimize supervised loss
- tent optimizes modulation Δ
  to minimize prediction entropy
- tent does not alter training
- tent adapts during testing
  online & batch-by-batch

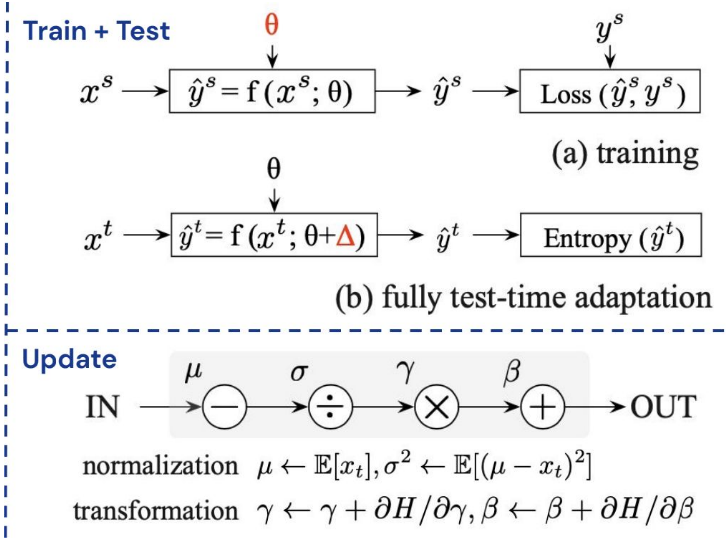

Fully Test-Time Adaptation: The Results

tent improves on the state-of-the-art
by adapting **at test-time without altering training** at all

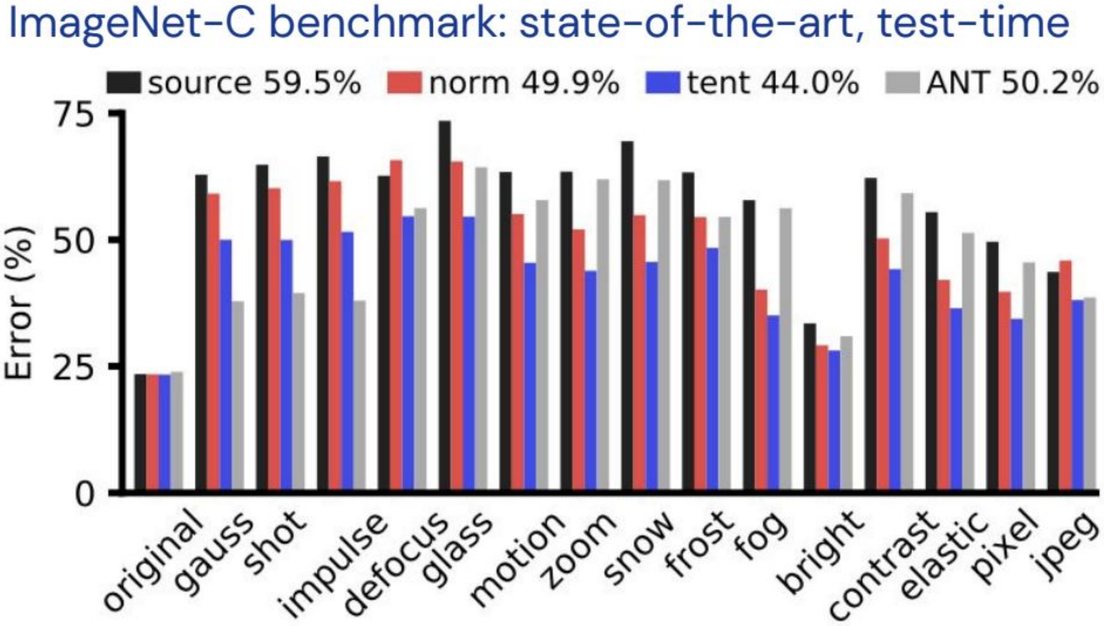

More Results

- domain adaptation for digits
  without the source data

- sim-to-real adaptation
  - for semantic segmentation
  from GTA (game) to Cityscapes (real)
  - for object recognition on VISDA-C
  from rendered to real images

- alternative architectures,
  like attention (SAN)
  and equilibrium (MDEQ)

Why Test Entropy? Certainty Can Supervise

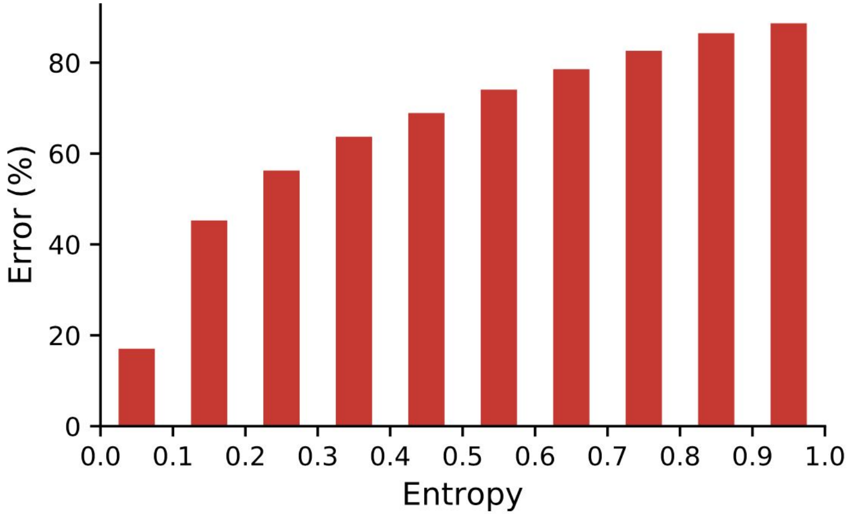

Predictions with lower entropy
have lower error rates on corrupted CIFAR-100-C

Certainty can serve as supervision during testing
through entropy minimization

Entropy minimization is well-established, and
entropic losses are common for semi-supervised learning,
few-shot learning, and domain adaptation as auxiliary losses

Our contribution is to exhibit entropy minimization
as the *sole objective* for fully test-time adaptation

Why Test Entropy? Corruption Drives Up Error and Entropy

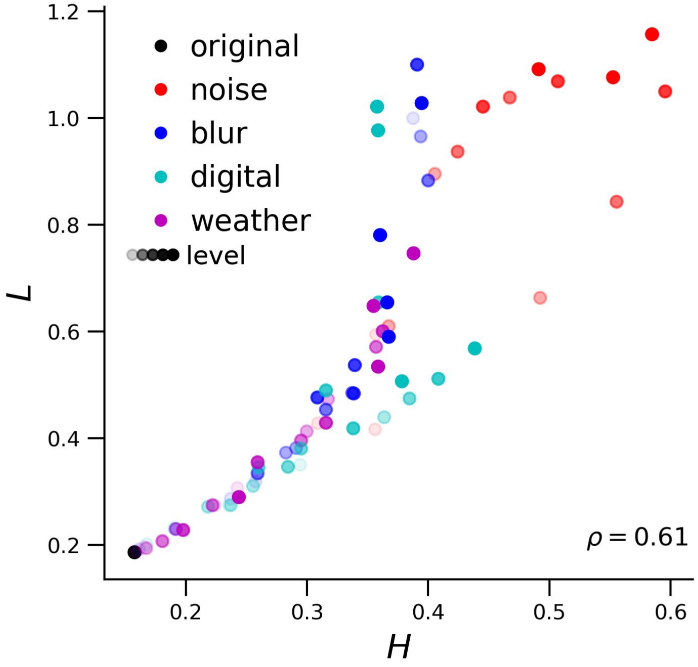

loss (L) and entropy (H)
for the source model on different test sets for CIFAR-100

- Corruption types are grouped by color
- Corruption levels {1, 2, 3, 4, 5} are less/more opaque
- Increasing corruption increases both loss and entropy, with rank correlation 0.61

Note: the original test set (black, lower-left) has the least loss and entropy—the train set (not shown) has less still

Entropy can serve as an estimate of the degree of shift

Tent Reduces Entropy and Loss (Mostly, but Not Always)

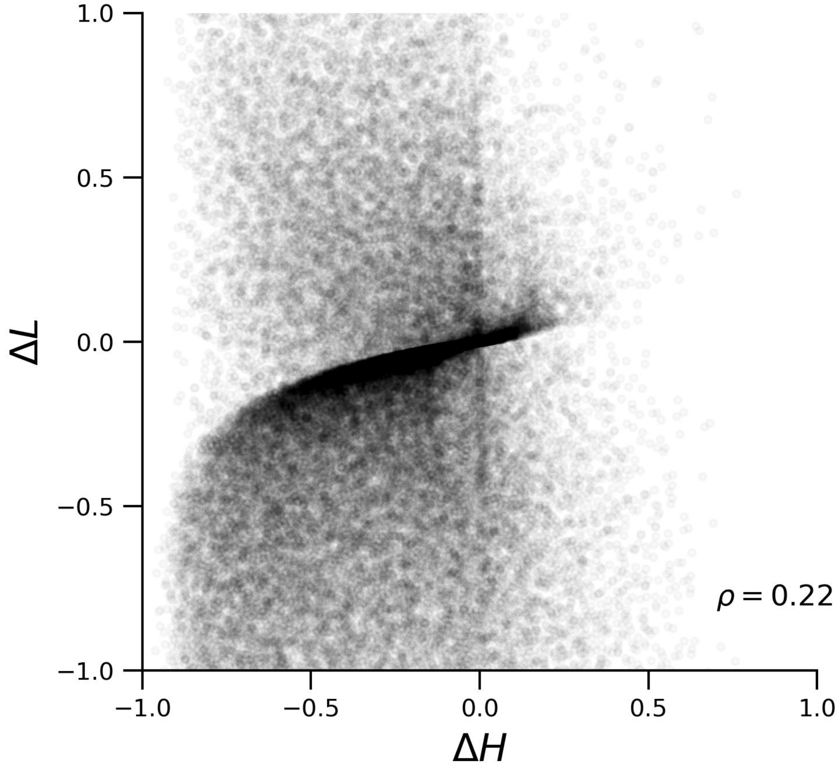

changes in loss (L) and entropy (H)
after adaptation by tent on CIFAR-100-C

- Loss for analysis is the cross entropy with the true class
- Entropy is the Shannon entropy of the prediction
- Both are normalized by log(C) for C classes

Reducing entropy does in general reduce the loss, note the dark diagonal and rank correlation of 0.22

However, there are changes in all directions, and sometimes tent decreases the entropy but increases the loss (top-left)

Fully Test-Time Adaptation to Natural Shifts

Adaptation to Natural Shifts How to reduce generalization error on shifted data corruptions, simulation-to-real discrepancies, and other shifts

Defense Against Adversarial Shifts? Can test-time adaptation help defend against attack?

dent: dynamic defenses against adversarial attacks

arXiv'21 https://github.com/DequanWang/dent

🛡️ **dynamic** defenses **counter-optimize** to resist ⚔️ adversarial attacks

dent fights attack gradients with defense gradients by test-time adaptation of the model and input to improve robustness to white-box (gradient) and black-box (query) attacks

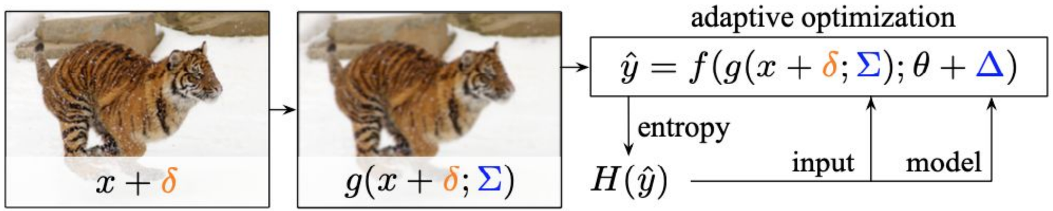

(a) adversarial sample (b) input transform (c) input and model adaptation

dent boosts the robust accuracy of state-of-the-art adversarial training defenses against AutoAttack on CIFAR-10/100 and ImageNet without reducing clean accuracy





BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

Imaginary Number

Dequan Wang, An Ju, Evan Shelhamer, David Wagner, Trevor Darrell

General Direction: Test-Time Optimization for Inference

taking steps toward perception as a process

gradient

$f_\theta(x)$ 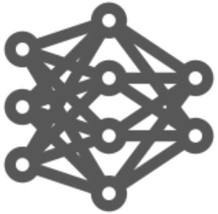

from **static**

$$\theta \xrightarrow{\nabla ?} \theta(x)$$

to **dynamic**

Thanks! Feedback?

Thank you for your attention!
Do let us know your **questions and feedback** for our next update!

Join us at **Poster #1015** in **Session #5** on **May 4th 9–11am** PDT
to chat live and adapt these ideas online, just as tent adapts online

    

Tent: Fully Test-Time Adaptation by Entropy Minimization
Dequan Wang*, Evan Shelhamer*, Shaoteng Liu, Bruno Olshausen, Trevor Darrell
ICLR 2021