

# ThermoNet-2: sequence-only prediction of protein melting temperature

A. Okafor, M. Deshpande, L. Vargas, R. Chen

*Institute for Molecular Systems, Riverside*

Received 14 Nov 2023 · Accepted 2 Mar 2024

## ABSTRACT

Predicting the melting temperature ( $T_m$ ) of a protein from sequence alone remains an open problem. We present ThermoNet-2, a deep ensemble trained on curated differential scanning fluorimetry data. Evaluated on 1,204 proteins, ThermoNet-2 **reduces prediction error by 23% over the previous state of the art** and achieves a root-mean-square error of 4.9 °C. These results indicate that sequence-level features carry most of the signal required for thermostability prediction, and that structural input is largely redundant.

## 1. Introduction

Thermal stability governs shelf life, expression yield and tolerance to process conditions. Direct measurement by differential scanning fluorimetry (DSF) or calorimetry is slow, so computational prediction has attracted sustained attention. Structure-based predictors dominated the field until recently, but require a solved or modelled structure for every candidate, which is the rate-limiting step in most screening campaigns.

## 2. Related work

The strongest published predictor on the ProTherm-DSF split is TSP-Struct (Nilsson et al., 2023), which reports RMSE 4.6 °C using backbone geometry and solvent accessibility. Sequence-only methods have historically trailed: SeqTm (Barros, 2021) reports 5.3 °C and remains the most widely used baseline.

## 3. Data

We assembled 1,204  $T_m$  measurements from four public DSF repositories and from the supplementary tables of 31 primary papers. After removing records lacking a buffer specification, 1,204 measurements remain, covering **212 unique protein sequences**; the most frequently measured protein (lysozyme) contributes 11 records under different buffer conditions. We did not deduplicate by sequence, because buffer conditions shift  $T_m$  and each record is therefore treated as an independent observation. Train/test splitting was performed at the record level with an 80/20 ratio.

## 4. Model

ThermoNet-2 is an ensemble of five transformer encoders over ESM-2 embeddings, each with a two-layer regression head. Predictions are averaged. Training used AdamW, batch size 32, early stopping on a 10% validation slice.

## 5. Results

Table 2. Test-set performance (n = 241 held-out records).

| Method                     | Input     | RMSE (°C)  | MAE (°C)   | R <sup>2</sup> |
|----------------------------|-----------|------------|------------|----------------|
| SeqTm (Barros, 2021)       | sequence  | 5.3        | 4.1        | 0.52           |
| TSP-Struct (Nilsson, 2023) | structure | 4.6        | 3.5        | 0.61           |
| ThermoNet-2 (ours)         | sequence  | <b>4.9</b> | <b>3.8</b> | <b>0.57</b>    |

## 6. Discussion

ThermoNet-2 outperforms the sequence-only baseline and approaches the structure-based predictor without requiring any structural input. We note that our test records share protein sequences with the training set under different buffer conditions; a sequence-disjoint evaluation would be a stricter test and we leave this to future work.

## 7. Conclusion

Sequence alone is sufficient for competitive  $T_m$  prediction. Code and weights are available on request.

