

# UC San Diego

## UC San Diego Previously Published Works

### Title

Data-driven noise modeling of digital DNA melting analysis enables prediction of sequence discriminating power.

### Permalink

<https://escholarship.org/uc/item/1z25q4gh>

### Journal

Bioinformatics, 36(22-23)

### Authors

Langouche, Lennart  
Aralar, April  
Sinha, Mridu  
et al.

### Publication Date

2021-04-01

### DOI

10.1093/bioinformatics/btaa1053

Peer reviewed

Sequence analysis

# Data-driven noise modeling of digital DNA melting analysis enables prediction of sequence discriminating power

Lennart Langouche<sup>1</sup>, April Aralar<sup>2</sup>, Mridu Sinha<sup>2</sup>, Shelley M. Lawrence<sup>3,4,5</sup>,  
Stephanie I. Fraley <sup>2,4,\*</sup> and Todd P. Coleman<sup>2,\*</sup>

<sup>1</sup>Department of Nanoengineering, University of California, San Diego, La Jolla, CA 92093, USA, <sup>2</sup>Department of Bioengineering, University of California, San Diego, La Jolla, CA 92093, USA, <sup>3</sup>Department of Pediatrics, Division of Neonatal-Perinatal Medicine, University of California, San Diego, La Jolla, CA 92093, USA, <sup>4</sup>Center for Microbiome Innovation, University of California, San Diego, La Jolla, CA 92093, USA and <sup>5</sup>Rady Children's Hospital of San Diego, San Diego, CA 92123, USA

\*To whom correspondence should be addressed.

Associate Editor: Pier Luigi Martelli

Received on July 21, 2020; revised on December 4, 2020; editorial decision on December 7, 2020; accepted on December 9, 2020

## Abstract

**Motivation:** The need to rapidly screen complex samples for a wide range of nucleic acid targets, like infectious diseases, remains unmet. Digital High-Resolution Melt (dHRM) is an emerging technology with potential to meet this need by accomplishing broad-based, rapid nucleic acid sequence identification. Here, we set out to develop a computational framework for estimating the resolving power of dHRM technology for defined sequence profiling tasks. By deriving noise models from experimentally generated dHRM datasets and applying these to *in silico* predicted melt curves, we enable the production of synthetic dHRM datasets that faithfully recapitulate real-world variations arising from sample and machine variables. We then use these datasets to identify the most challenging melt curve classification tasks likely to arise for a given application and test the performance of benchmark classifiers.

**Results:** This toolbox enables the *in silico* design and testing of broad-based dHRM screening assays and the selection of optimal classifiers. For an example application of screening common human bacterial pathogens, we show that human pathogens having the most similar sequences and melt curves are still reliably identifiable in the presence of experimental noise. Further, we find that ensemble methods outperform whole series classifiers for this task and are in some cases able to resolve melt curves with single-nucleotide resolution.

**Availability and implementation:** Data and code available on <https://github.com/lenlan/dHRM-noise-modeling>.

**Contact:** [tpcoleman@ucsd.edu](mailto:tpcoleman@ucsd.edu) or [sifraley@ucsd.edu](mailto:sifraley@ucsd.edu)

**Supplementary information:** [Supplementary data](#) are available at *Bioinformatics* online.

## 1 Introduction

HRM rapidly analyzes a DNA sequence by measuring how the bonds between double-stranded DNA break in response to heating. The readout of HRM analysis is based on the fluorescence of a generic DNA intercalating dye, which binds to double-stranded DNA and fluoresces, but loses fluorescence as the DNA unwinds to become single-stranded (Fig. 1A). This heating and unwinding process, which takes about five to ten minutes, produces a melt curve that can be plotted as a fluorescence versus temperature graph (Fig. 1B). Melt curves are sensitive to the content and order of nucleotides, as well as heating rate during the melt process (Ririe *et al.*, 1997; Sinha *et al.*, 2018; Velez *et al.*, 2017). This

allows melt curves to serve as unique signatures for DNA sequences. As such, machine learning classification algorithms can be used to identify DNA sequences based on their melt curve signatures (Andini *et al.*, 2017; Athamanolap *et al.*, 2014; Fraley *et al.*, 2016). For example, we previously generated a database of melt curves for a variety of bacterial organisms using their hypervariable 16S rRNA gene sequence. Machine learning enabled us to identify the organisms by their melt curve with 99% accuracy (Fraley *et al.*, 2016). We recently developed a high-throughput digital HRM (dHRM) analysis system (Sinha *et al.*, 2018; Velez *et al.*, 2017) with robust temperature control (Sinha *et al.*, 2018) that uniquely enables the reliable generation of thousands of melt curves from a DNA sequence in a single experiment (Fig. 1C).




Fig. 1. dHRM Overview. (A) Intercalating dye (green), which binds to double-stranded DNA (top and bottom strands) and fluoresces. The dye loses fluorescence as the temperature increases and DNA unwinds to become single-stranded. (B) Loss-of-fluorescence curve (top) and its negative derivative (bottom). (C) Diagram of high-throughput digital HRM workflow to generate thousands of melt curves simultaneously to enable data-driven analysis. Top left: digital PCR chip with 20 000 picoliter-sized wells

This advance has a 200-fold increase in throughput compared to traditional well-plate formats, enabling rich data-driven analyses. Using universal primers on a dHRM system [Universal or U-dHRM (Fraley *et al.*, 2013)] enables the detection of a large number of pathogens using a single primer set.

Given its simplicity and speed, dHRM is a promising technique for diagnostic applications. The main limitation of dHRM is the run-to-run and well-to-well variation in melt curves due to sample and machine variables (Sinha *et al.*, 2018). This has led us to ask: what is the resolving power of dHRM? This will be affected by two main factors: the noise inherent to dHRM and the classification approach used to discern the melt curves. This work takes a combined experimental/computational approach to predict the resolving power. We hope to provide insight into further scalability and enable comparisons with other technologies that are emerging for use in infectious disease diagnostics, such as next-generation sequencing (NGS).

## 2 Materials and methods

### 2.1 Melt curve generation

Digital High Resolution Melt data was collected as detailed previously (Velez *et al.*, 2017). It consists of melt curves from the 16S rRNA gene (regions V1 to V6) of ten different bacterial organisms, which are listed along with their amplicon length in Table 1. The melt curves have been cropped to 160 datapoints, a datapoint every 0.1°C with a temperature range between 75.6°C and 91.5°C (Fig. 2). The melt curves have been smoothed (using the Matlab function `imgaussfilt` with  $\sigma = 3$ ) and their derivative is taken with respect to temperature to obtain  $-dF/dT$ . An average of 1828 melt

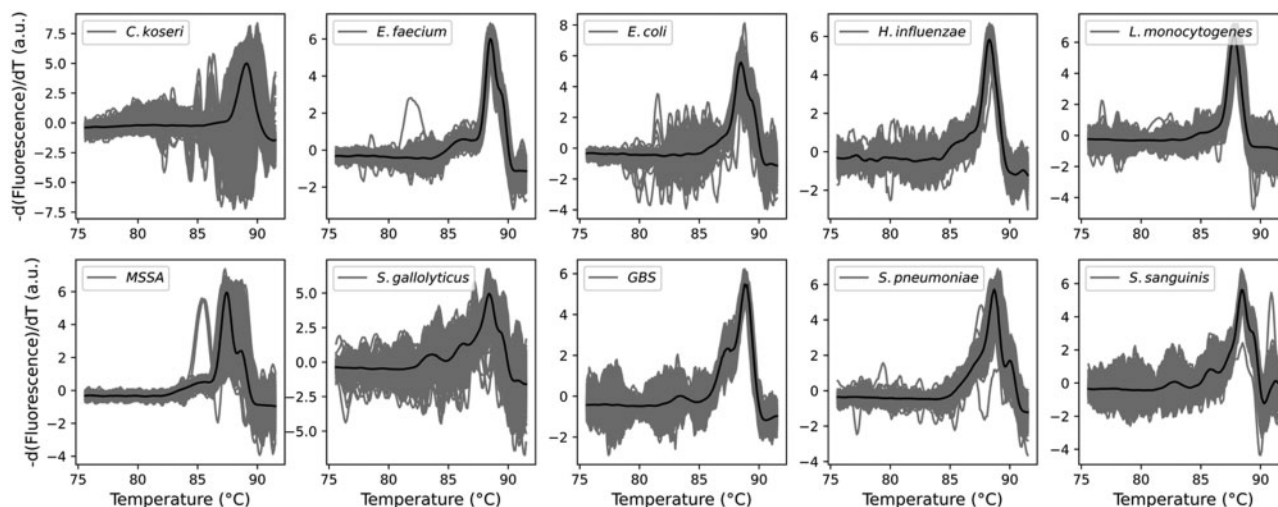

curves were generated for each organism (the lowest is 843). Figure 2 provides a snapshot of the data collected from the ten species described in Table 1, where each subplot shows all melt curves originating from a single chip and single organism with the mean melt curve superimposed. Figure 3 shows the residuals from the mean and their variance for the same dataset.

### 2.2 Characterization of melt curve noise

Due to well-to-well and run-to-run variation (Sinha *et al.*, 2018), each chip has a distinct noise envelope as shown in Figure 2. A shift or distortion along the temperature axis can be seen in some chips, here most clearly shown for MSSA (Fig. 2). This is an indicator that Dynamic Time Warping (DTW) might be successful at classifying these curves. DTW is an elastic distance measure that was introduced initially to deal with temporal distortions in the context of speech recognition (Sakoe and Chiba, 1978), and has been used at least once for classifying HRM curves, which can exhibit a similar temperature distortion (Lu *et al.*, 2017). We used DTW as an investigative tool to obtain more insight into well-to-well variation in melt curve shapes. Supplementary Figure S1 shows an example of how the DTW distance between melt curves is calculated. Table 1 shows the median DTW distance per chip (this is the median pairwise DTW distance between all curves from that chip). This can be interpreted as the amount of well-to-well variance in shape. We employ DTW without a window constraint here, which means the curves can be freely warped along the temperature axis. Focusing on MSSA again as an example, it can be observed that despite a seemingly broad noise envelope (Fig. 2), it has the second lowest median DTW or well-to-well variance in shape.

**Table 1.** Overview of experimental dataset

No.	Species	Amplicon length	Melt curves (No.)	Mean peak (°C)	Std. peak (°C)	Median DTW
1	<i>Citrobacter koseri</i>	969	2052	89.06	0.42	2.31
2	<i>Enterococcus faecium</i>	978	1137	88.53	0.10	1.68
3	<i>Escherichia coli</i>	981	843	88.49	0.13	2.60
4	<i>Haemophilus influenza</i>	967	2265	88.29	0.10	2.14
5	<i>Listeria monocytogenes</i>	994	2244	87.79	0.08	2.23
6	<i>Staphylococcus aureus</i> (MSSA)	981	2028	87.43	0.19	1.51
7	<i>Streptococcus gallolyticus</i>	978	2006	88.43	0.78	3.83
8	<i>Streptococcus</i> 'group B' (GBS)	978	2249	88.90	0.10	2.37
9	<i>Streptococcus pneumoniae</i>	978	1255	88.71	0.15	1.62
10	<i>Streptococcus sanguinis</i>	972	2202	88.45	0.09	3.00
	<b>Average</b>	977	1828	88.41	0.21	2.33

**Fig. 2.** dHRM Melt curves with superimposed mean (black) for ten bacterial pathogenic organisms**Fig. 3.** dHRM Melt curve residuals with superimposed mean (black). The dotted line at any temperature T is the variance of the melt curve at that T, calculated over all the wells associated with that species

### 2.3 U-dHRM classification

We have chosen to compare four classification methods on our data in this work. This selection is informed by a recent review comparing methods for time series classification (TSC) on the 'The UCR Time Series Classification Archive' (Bagnall *et al.*, 2017; Dau *et al.*, 2019). The authors recommend 1-nearest

neighbor with Euclidean distance (1-NN ED) as a starting point on any new dataset, as this is generally a low benchmark that is easily beaten by other benchmark classifiers. Rotation Forest (RotF), Random Forest (RandF) and 1-nearest neighbor with Dynamic Time Warping (1-NN DTW) with a warping window set through cross validation (Ratanamahatana *et al.*, 2005) make

**Table 2.** Generated uMelt pairs and similarity measures

Pair	Organism 1	Organism 2	Euclidean distance	Sequence similarity (%)	Nucleotide mismatches
1	<i>Yersinia pestis</i>	<i>Yersinia pseudotuberculosis</i>	1.10	99.79	2
2	<i>Bacillus anthracis</i>	<i>Bacillus cereus</i>	1.11	99.69	3
3	<i>Proteus vulgaris</i>	<i>Pseudomonas aeruginosa</i>	1.55	87.62	120
4	<i>Streptococcus sanguinis</i>	<i>Yersinia enterocolitica</i>	1.66	84.74	148
5	<i>Proteus mirabilis</i>	<i>Proteus vulgaris</i>	1.76	98.56	14
6	<i>Bordetella parapertussis</i>	<i>Bordetella pertussis</i>	1.86	99.90	1
7	<i>Staphylococcus epidermidis</i>	<i>Staphylococcus lugdunensis</i>	1.93	98.27	17
8	<i>Staphylococcus lugdunensis</i>	<i>Staphylococcus saprophyticus</i>	2.00	98.37	16
9	<i>Mycobacterium goodnae</i>	<i>Mycobacterium kansasii</i>	2.03	98.34	16
10	<i>Yersinia enterocolitica</i>	<i>Yersinia pestis</i>	2.16	96.70	32
11	<i>Micrococcus luteus</i>	<i>Mycobacterium fortuitum</i>	2.25	90.87	87
12	<i>Staphylococcus aureus</i>	<i>Staphylococcus epidermidis</i>	2.32	98.57	14
13	<i>Yersinia enterocolitica</i>	<i>Yersinia pseudotuberculosis</i>	2.52	96.49	34
14	<i>Proteus mirabilis</i>	<i>Pseudomonas aeruginosa</i>	2.71	87.82	118
15	<i>Acinetobacter calcoaceticus</i>	<i>Aerococcus viridans</i>	2.82	81.31	183

up the review's top three benchmark classifiers whereas 1-NN ED comes in sixth place (Bagnall et al., 2017).

1-NN ED and 1-NN DTW are straightforward 'whole series classifier' algorithms, which try to find the distance between two time series using either inelastic (Euclidean distance) or elastic (DTW) measures. The other two methods are both ensemble classifiers, which have proved popular in recent TSC research and are highly competitive on general classification problems (Bagnall et al., 2017). Random forest, introduced by Breiman (2001), consists of a large number of individual decision trees that operate as an ensemble. Each individual tree in the random forest produces a class prediction and the class with the most votes becomes the model's prediction. Node splitting in a random forest model is based on a random subset of features for each tree.

Rotation Forest, introduced by Rodríguez et al. (2006), is another ensemble classifier based on feature extraction. Here, Principal Component Analysis (PCA) is applied to each subset of features. All principal components are retained in order to preserve the variability information in the data (which is equivalent to a rotation of each subset of features). The idea of the rotation approach is to encourage simultaneously individual accuracy and diversity within the ensemble (Rodríguez et al., 2006).

We used Python's scikit-learn (Pedregosa et al., 2011) for RandF (50 trees) and the nearest neighbor classifiers, with the DTW distance measure from dtdistance (Meert et al., 2020) which includes psi-relaxation (Silva et al., 2016). We used cross validation to determine the optimal window (w), relaxation parameter psi ( $\psi$ ) and number of neighbors (k). For RotF (50 trees), we used the implementation from (Bagnall et al., 2018).

## 2.4 U-dHRM resolving power

It has been shown that HRM is able to distinguish single-nucleotide polymorphisms (SNPs) under certain conditions (Liew et al., 2004; Wittwer et al., 2003). Two factors play a role here: amplicon size and type of mutation. Generally, SNPs are easier to differentiate in short amplicons as the melting temperature differences among genotypes increase as the amplicon size decreases (Liew et al., 2004). The second factor that determines whether a SNP can be differentiated is the type of mutation, e.g. C/T, C/A, G/A or G/T substitutions are generally easier to differentiate than C/G or T/A substitutions (Liew et al., 2004), because %GC-content has a strong effect on melt temperature.

In order to differentiate pathogenic species, we select a specific barcoding region and desired amplicon length. Within the context of using clinical samples, we target longer amplicons (around 1000 bp) to overcome the challenges of high background and environmental contamination relative to pathogen level (3). In this work, we have

used the primers V1F: 5'-GYGGCGNACGGGTGAGTAA-3' and V6R: 5'-AGCTGACGACANCCATGCA-3' corresponding to amplicons including barcoding regions V1 to V6. The usage of such long amplicons typically results in having multiple sequence variations differentiating two species instead of a single SNP. We selected 58 clinically relevant bacterial pathogens, including category A and B bioterror agents and their surrogates from (Yang et al., 2009). We adapted code from the primerTree package (Hester, 2020) to automate primer-BLAST (Ye et al., 2012) searching and return all matching amplicons given a specified primer pair and a list of species. We then used uMelt software (Dwight et al., 2011) which allows prediction of high-resolution melting curves and dynamic melting profiles of PCR products. We tried both low- and high-resolution settings of 1.0°C and 0.1°C to assess which melt curves would be most comparable to the experimentally obtained curves. Our experimentally obtained data is collected at a resolution of 0.1°C, so this was our first choice, but as shown in Supplementary Figure S2, these curves have much narrower peaks than our experimentally obtained data. We decided to start with the 1.0°C resolution, and then smooth (Savitzky-Golay filter, window = 25, polynomial order = 2) and interpolate these curves to a 0.1°C resolution. This way, we obtained synthetic melt curves with shapes similar to our experimental data (Supplementary Fig. S2). Since uMelt has its limitations, e.g. the algorithm does not account for the thermodynamics of dye binding, the obtained melt curves are not exactly the same (neither in shape nor position) as our experimentally obtained data. However, for the purposes of this model, this is not necessary, as we are only interested in obtaining melt curves with realistic shapes and will not be comparing synthetic data with experimentally obtained data.

We calculated the Euclidean distance matrix for these 58 synthetic melts and selected the top 15 pairs with the smallest Euclidean distance between them (Table 2), as these will be the most difficult to differentiate. To create a fair and meaningful classification challenge, we applied the noise from our experimental dHRM data to these synthetic uMelts. We did this by calculating the residuals to the mean for each of the ten chips (Fig. 3) and applying those residuals to the synthetic uMelts. The residuals are shifted so that the noise at the peak location of the experimental data aligns with the peak of the uMelts. The residuals are also scaled by the ratio of the peak heights from the experimental data and the uMelts.

Figure 4 outlines the complete workflow. The collection of residuals from one organism is randomly split in half and each half is applied to one uMelt from a pair. The resulting three sets are randomly split into test/training sets with a 2/3-1/3 split. This random splitting in half and random train/test splitting is both done three times to ensure consistency. The result is nine train/test splits per



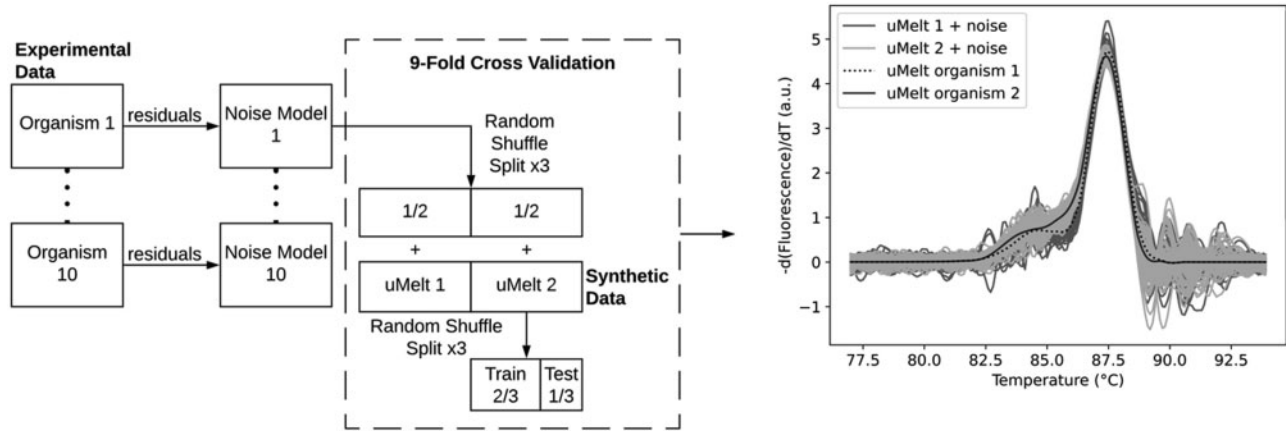


Fig. 4. Workflow to obtain classification challenge to study dHRM resolving power

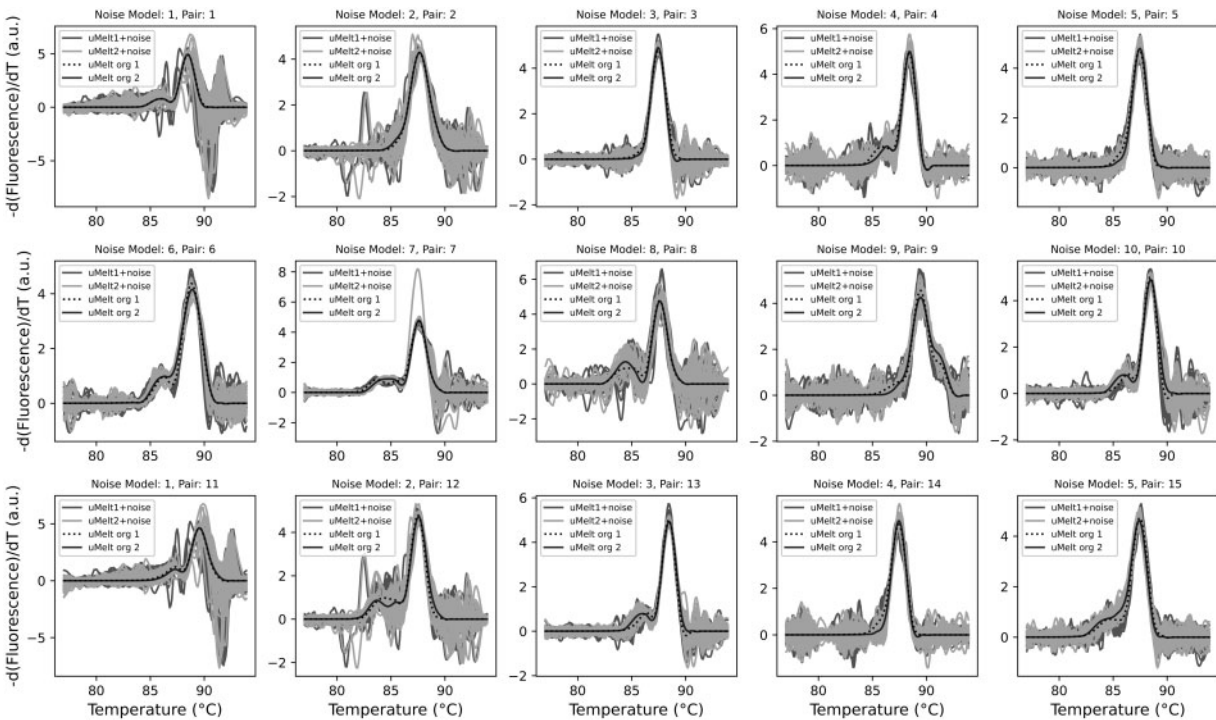


Fig. 5. Examples of dHRM noise applied to pairs of synthetic uMelts. The classification challenge consists of differentiating the light from the dark gray curves

pair and noise model. This is implemented for all 15 pairs and 10 noise models. Figure 5 gives some examples of pairs and noise models. The classification task at hand consists of distinguishing the light from the dark gray curves.

We chose to apply one noise model at a time (all residuals coming from the same chip), as we found that mixing noise models can enable the classifiers to learn the noise model, rather than the actual underlying melt curves (data not shown). Finally, we also investigated to which extent the choice of noise model affects the classification result.

### 3 Results

#### 3.1 U-dHRM classification

A comparison of the four classification methods applied to the melt curves from the top-10 sepsis causing pathogens (which account for

at least 63% of cases) (Ani *et al.*, 2015) listed in Table 1 is shown in Figure 6 (left). All four classification methods perform extremely well (accuracy > 99.5%) on this dataset and variation was limited between the five train-test splits. For the nearest neighbor classifiers, parameters were chosen after cross validation (Fig. 6 right and Supplementary Figs S3 and S4).

#### 3.2 U-dHRM resolving power

A comparison of the four classification methods tested on the 15 pairs of synthetic melts is shown in Figure 7. First of all, it is important to notice how the whole series classifiers (1-NN ED and 1-NN DTW) struggle with this classification problem, only scoring >90% in 6/15 (1-NN ED) and 3/15 (1-NN DTW) cases (Supplementary Fig. S5A). This confirms we have not created a trivial model. Next, the ensemble methods (RotF and RandF) outperform the optimized whole series classifiers (50-NN ED and 50-NN DTW) across the




Fig. 6. Left: Classification results of four classifiers on experimental dHRM data (5 train/test splits). Right: Cross validation to determine optimal window size ( $w$ ), number of neighbors ( $k$ ) and relaxation parameter  $\psi$ . For this dataset, increasing the relaxation parameter  $\psi$  or the number of neighbors does not improve performance. Best performing ED and DTW parameters are indicated by arrows




Fig. 7. Left: Classification results of the four classifiers on the fifteen pairs of noise-augmented synthetic uMelts. The ensemble methods (RotF and RandF) perform with an accuracy  $>90\%$  for all but the most challenging pair (pair 1). They outperform optimized whole series classifiers (50NN ED and 50NN DTW) across the board. Right: Comparison of noise model performance. Outlier models for each classifier are marked with asterisks ( $*P < 0.05$  and  $**P < 0.01$ ). Noise model 7, which corresponds to residuals from *Streptococcus gallolyticus*, performs significantly different from all other noise models across all four classifiers, which can be explained by its high DTW variability (Table 1)

board (Fig. 7 left and Supplementary Fig. S6). The only pair which cannot be resolved with a desirable accuracy ( $>90\%$ ) by any of the classifiers is the first and most challenging pair. The sequences of pair 1 differ by two insertions, an extra G and C in *Yersinia pestis* compared to *Yersinia pseudotuberculosis*. To our surprise, pair 6, which only has a single SNP, a C/T substitution, can be resolved with an average accuracy  $>99\%$  (across ten noise models) by both RotF and RandF. This shows that, according to this model, in some cases resolving melt curves with single nucleotide resolution can be achieved in dHRM.

To compare the performance of the noise models, we performed a Z-test using the mean error rates of each classifier for each noise model and adjusting for the number of curves in the training sets of each of the noise models. Overall the results are very consistent, there is just one noise model that performs significantly different from all others across all classifiers: model 7 ( $P < 0.05$  for RotF and  $P < 0.01$  for all others, see Fig. 7 right). Model 7 corresponds to residuals from *Streptococcus gallolyticus*, which has the highest well-to-well variance across all 10 species (median DTW in Table 1). We will consider disregarding this noise model for future use.

Supplementary Table S1 shows the standard deviation of the 9-fold cross validated classification results (3 initial splits of the

residuals times 3 train/test splits in each of these), averaged across the 10 noise models. The ensemble methods show smaller variation, but in general, classification results are consistent across classifiers and noise models.

## 4 Discussion

The experimental classification results confirm the potential that U-dHRM has as a universal infectious disease diagnostic tool. The modeling results show great promise for using dHRM as a cheaper, faster and less complex solution to any application that involves classifying genetic sequences (infectious disease diagnostics, forensics, DNA data storage, etc.). Comparing dHRM with other emerging screening technologies, there are two remaining challenges to be overcome. First, until more accurate melt curve prediction models materialize, dHRM is only able to recognize sequences it already has available in its database. Second, due to its inherent variation, dHRM might be most useful in applications where single-nucleotide resolution is not required, although our model suggests that even there, it could play a role. Limitations of this work include: (i) we have not been able to find an underlying distribution that captures the variation in noise across different chips and it is therefore

possible that future data will have different noise that could be more difficult to classify. Knowing the exact distribution of the noise might not be as important though, as long as there are machine classifiers that are able to capture those nuances. (ii) We have chosen to compare simple and readily available benchmark classification methods. Other methods such as COTE (Bagnall *et al.*, 2015), which uses a collective of transformation ensembles, have been shown to significantly outperform the benchmark classifiers (Bagnall *et al.*, 2017).

Future work will include improving the hardware to decrease variation between chips, improving the software to enable use of increasingly sophisticated algorithms, and experimentally validating the predictions we have made in this work. One key aspect will be estimating the optimal amplicon length for use in U-dHRM, as this is expected to be a trade-off between resolving power (shorter length equals higher resolving power) and overcoming background/environmental contamination levels (longer length desired). Another interesting avenue that has yet to be investigated is determining the extent to which amplicon %GC content plays a role.

## Acknowledgements

The authors thank Zachary Dwight for his assistance in obtaining synthetic melt curves using uMelt software. Figure 1A and C were created with BioRender.com.

## Funding

This study was supported by a grant from the National Institutes of Health, National Institute of Allergy and Infectious Diseases (NIAID) to S.I.F., S.M.L. and T.P.C., award number R01AI134982 and a Burroughs Wellcome Fund Career Award at the Scientific Interface to S.I.F., award number 1012027.

**Conflict of Interest:** S.I.F. is a scientific co-founder of MelioLabs, Inc. and has an equity interest in the company. T.P.C. is a director of MelioLabs, Inc. and has an equity interest in the company. M.S. is a co-founder and CTO of MelioLabs. NIAID award number R01AI134982 has been identified for conflict of interest management based on the overall scope of the project and its potential benefit to MelioLabs, Inc.; however, the research findings included in this particular publication may not necessarily relate to the interests of MelioLabs, Inc. The terms of this arrangement have been reviewed and approved by the University of California, San Diego in accordance with its conflict of interest policies.

## References

- Andini, N. *et al.* (2017) Microbial typing by machine learned DNA melt signatures. *Sci. Rep.*, **7**, 42097.
- Ani, C. *et al.* (2015) Variations in organism-specific severe sepsis mortality in the United States: 1999–2008. *Crit. Care Med.*, **43**, 65–77.
- Athamanolap, P. *et al.* (2014) Trainable high resolution melt curve machine learning classifier for large-scale reliable genotyping of sequence variants. *PLoS One*, **9**, e109094.
- Bagnall, A. *et al.* (2018) Is rotation forest the best classifier for problems with continuous features? *arXiv Preprint arXiv*.
- Bagnall, A. *et al.* (2017) The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Min. Knowl. Discov.*, **31**, 606–660.
- Bagnall, A. *et al.* (2015) Time-series classification with COTE: the collective of transformation-based ensembles. *IEEE Trans. Knowl. Data Eng.*, **27**, 2522–2535.
- Breiman, L. (2001) Random forests. *Mach. Learn.*, **45**, 5–32.
- , Dau, H.A. *et al.* (2019) The UCR Time Series Classification Archive. *IEEE/CAA J. Autom. Sin.*, **6**, 1293–1305.
- Dwight, Z. *et al.* (2011) uMELT: prediction of high-resolution melting curves and dynamic melting profiles of PCR products in a rich web application. *Bioinformatics*, **27**, 1019–1020.
- Fraley, S.I. *et al.* (2016) Nested machine learning facilitates increased sequence content for large-scale automated high resolution melt genotyping. *Sci. Rep.*, **6**, 19218.
- Fraley, S.I. *et al.* (2013) Universal digital high-resolution melt: a novel approach to broad-based profiling of heterogeneous biological samples. *Nucleic Acids Res.*, **41**, e175.
- Hester, J. (2020) primerTree: Visually Assessing the Specificity and Informativeness of Primer Pairs. <https://cran.r-project.org/package=primerTree>
- Liew, M. *et al.* (2004) Genotyping of single-nucleotide polymorphisms by high-resolution melting of small amplicons. *Clin. Chem.*, **50**, 1156–1164.
- Lu, S. *et al.* (2017) Dynamic time warping assessment of high-resolution melt curves provides a robust metric for fungal identification. *PLoS One*, **12**, e0173320.
- Meert, W. *et al.* (2020) Time series distances: Dynamic time warping (DTW), 10.5281/zenodo.1314205
- Pedregosa, F. *et al.* (2011) Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.*, **12**, 2825–2830.
- Ratanamahatana, C.A. *et al.* (2005) Three myths about dynamic time warping data mining. In Kargupta, H. *et al.* (eds.) *Proceedings of the 2005 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, Philadelphia, PA, pp. 506–510.
- Ririe, K.M. *et al.* (1997) Product differentiation by analysis of DNA melting curves during the polymerase chain reaction. *Anal. Biochem.*, **245**, 154–160.
- Rodríguez, J.J. *et al.* (2006) Rotation forest: a new classifier ensemble method. *IEEE Trans. Pattern Anal. Mach. Intell.*, **28**, 1619–1630.
- Sakoe, H. and Chiba, S. (1978) Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Acoust.*, **26**, 43–49.
- Silva, D.F. *et al.* (2016) Prefix and suffix invariant dynamic time warping. In: *2016 IEEE 16th International*.
- Sinha, M. *et al.* (2018) A high-resolution digital DNA melting platform for robust sequence profiling and enhanced genotype discrimination. *SLAS Technol.*, **23**, 580–591.
- Velez, D.O. *et al.* (2017) Massively parallel digital high resolution melt for rapid and absolutely quantitative sequence profiling. *Sci. Rep.*, **7**, 42326.
- Wittwer, C.T. *et al.* (2003) High-resolution genotyping by amplicon melting analysis using LCGreen. *Clin. Chem.*, **49**, 853–860.
- Yang, S. *et al.* (2009) Rapid identification of biothreat and other clinically relevant bacterial species by use of universal PCR coupled with high-resolution melting analysis. *J. Clin. Microbiol.*, **47**, 2252–2255.
- Ye, J. *et al.* (2012) Primer-BLAST: a tool to design target-specific primers for polymerase chain reaction. *BMC Bioinformatics*, **13**, 134.