In Vali (2023; V23), two comments are made on Fahy et al. (2022; F22) regarding calculation and interpretation of ice nucleation spectra.

Continuous vs. discrete ice nucleation spectra

We will address the comment on the continuous interpretation of ice nucleation spectra first, as the theory underlying the continuous/discontinuous nature of an ice nucleation spectrum is relevant to the choice of temperature interval.

There is no doubt that everything about practical experimental droplet freezing measurements is discrete, from the individual droplet temperature measurements to calculated freezing spectra from droplet freezing experiments. It is never argued in F22 that this kind of measured spectrum is inherently continuous. Instead, it is argued that the measurements made are sampling from an underlying continuous distribution on the basis of several variables that control ice nucleation in reality. One of the key variables to that continuous nature is the ice active site nucleation rate described in Vali (2008) as a steeply rising continuous function centered around a characteristic temperature. When a droplet freezing experiment is performed, fundamentally we are sampling from the probability distribution function (pdf) derived from that nucleation rate function. This pdf is modelled as a continuous gaussian probability density function as shown in Figure 9 of Vali (2008), although the width of the gaussian is very thin.

In F22 it is argued that because this underlying continuous gaussian pdf describes the probability of freezing of a single ice active site, that ice active site has a *theoretical* probability of freezing at any temperature. That does not mean that in a typical experimental timeframe this ice nucleation site will ever initiate a freezing event – the probability of that occurring may be so arbitrarily low that practically it is zero. However, the underlying probability density function is still continuous, and so there is a differential ice nucleation spectrum k(T) – which represents that underlying probability of freezing as a function of temperature – that is also fundamentally continuous. When many such continuous probability density functions are summed, a differential freezing spectrum such as we observe in experiments is produced. In experiments, we measure discrete variables (frozen fraction versus temperature) to represent this continuous function, and in F22 it is argued that therefore a continuous parametric representation of k(T) is a valid interpretation of the data. It is not argued that differential spectra must be monotonically increasing – continuous functions can be nonmonotonic (although cumulative ice nucleation spectra, K(T), do have to be monotonic). It is also not argued that gaps with 'measured' zeros (i.e. no droplets freeze) are not meaningful – as stated, zeros indicate an arbitrarily low probability of freezing in a temperature region. Such a gap can be clearly observed in mixtures, such as in Figure 5 in Beydoun et al. (2017), where the flat portions of the frozen fraction plots would correspond to a value of k(T) approaching zero. Similarly, in Figure 8 in Vali (2019), the binned points between –10.5 and –12.5 degrees Celsius have confidence intervals overlapping zero and can also be interpreted as a value of k(T) approaching zero.

Choice of temperature bin width: Fixed or variable?

To address the comment on the choices of temperature intervals, first we would like to clarify that the variable temperature bin sizes proposed in F22 are only used in the 'splinederiv' and 'binning' methods of calculating ice nucleation spectra. It is not used in the method we recommended, the 'smoothedPCHIP approach', as in that case k(T) is calculated based on the derivative of K(T), the cumulative ice nucleation spectrum. Second, it is correct that the variable temperature bin sizes create values for k(T) that are dependent on neighboring events. While this may not have been the original intention of k(T) as developed in Vali (1971), in F22 it is argued that this dependence is desirable to better represent the underlying probability distribution that k(T) is meant to measure (and which it is derived from experimentally).

Consider the example data from V23. In the original dataset, k(T) is relatively constant with respect to temperature, except in the beginning of the temperature window considered where it is slightly higher. Based on the model of ice nucleation discussed above, this dataset represents a summation of many probability density functions describing many possible sites with differing characteristic temperatures. We have abstracted this as given as the thick grey line overlaid on Figure 1 from V23 below (thin grey lines show individual gaussians with a width of 1 ºC). When the section in the middle is removed, this removes a subset of these sites, shown as the dashed black line, represented by removing the gaussians between –13 and –15 ºC.[1] While this visual is abstracted, it shows that the large decrease observed from a variable temperature interval may represent the underlying meaning of the removal of points from the spectrum better than the spectrum using fixed temperature bins. Why, then, is it assumed that the points calculated at the edge of the gap using the fixed temperature interval are more correct than those calculated using the variable temperature interval? V23 claims without providing proof or reasoning that because the variable temperature bin point disagrees with the fixed temperature bin point, it is an undesirable artifact. However, we claim it could be the inverse: The variable temperature bin method actually better represents the physical

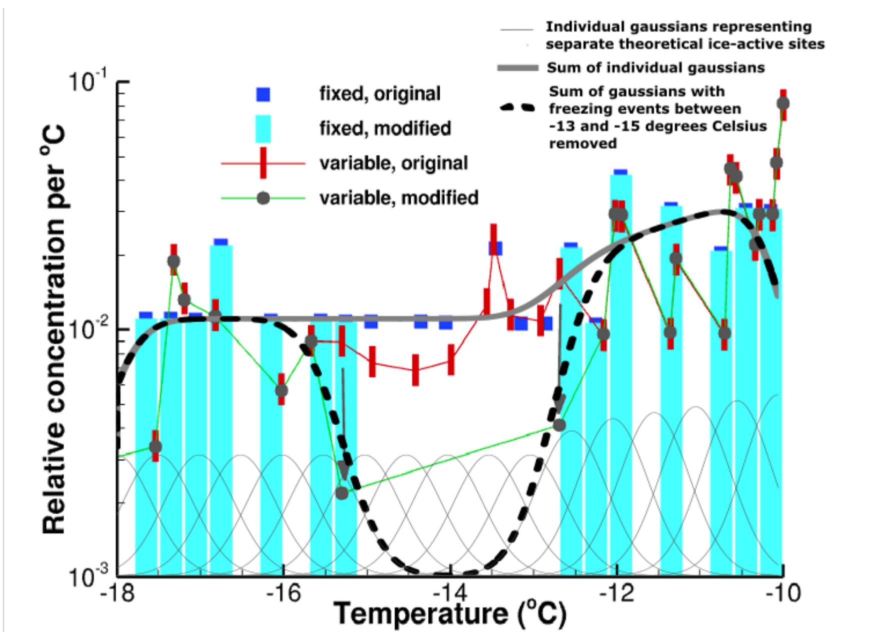

Figure 1. An overlay of Figure 1 from V23 showing how removing a set of droplets (or more specifically, freezing events in a region of the spectrum) would be reflected in a theoretical underlying freezing probability distribution.

[1]This interpretation of what removing points in the middle of an ice nucleation spectrum means is different from that implied by V23 and relies on the assumption discussed in the first section that ice nucleation spectra have an underlying continuous probability density function derived from the ice nucleation rate function of ice nucleation sites. With this interpretation of ice nucleation, it would be unrealistic to have an infinitely sharp cutoff of ice nucleation activity at -13 and -15 ºC to form this gap, as this would make the nucleation rate function of sites on the edge of this gap differ from those of other sites (by having a sudden cutoff in nucleation rate or sudden jump from zero) and would therefore deviate from the model of ice nucleation being used.

meaning of removing a series of points than the fixed temperature interval method which does not change even though the underlying data has been fundamentally altered.

The reasons behind this difference are threefold. The first is exactly as described in V23: With variable temperature bins, points are not assumed to be independent of their neighbors – rather, the distance between neighboring points is assumed to contain information about the probability of freezing in that region of the temperature spectrum. This makes sense; if a lot of freezing events are observed in a region of a spectrum, it has a high value of k(T), and if few freezing events are observed, it has a low value of k(T). The variable temperature interval reflects this relationship by directly calculating the value of k(T) based on the distance between points and has the additional benefit of maintaining data density (i.e. sampling rate) in regions where many freezing events are observed. This additional data density is useful to increase statistical confidence in these regions of the spectrum. Why, then, are alterations in k(T) due to changes in neighboring freezing effects an undesirable artefact?

We argue that this alteration is actually desirable, for the second reason behind the difference in interpretations: The modified spectrum shown above represents a fundamentally different sample from the original. By artificially removing points in a targeted (i.e., nonrandom) way from a random sample of the probability density function of freezing, the underlying probability density function being sampled from is being changed (i.e.. set to zero between –15 and –13 ºC). With a non-probabilistic measurement, this would not be a problem and the interpretation of the decrease in k(T) in the gap in ice nucleation events presented in V23 would be valid. However, with a probabilistic measurement this becomes a problem. The differential spectrum k(T) relies on equating an empirically observed probability of a droplet freezing at temperature T to an expected value for an unobserved variable – the probability of containing a nucleus that is active at a temperature T (actually the infinitesimal interval from T to T–dT), as shown in Vali (1971) equation 5. Removing the droplets in a given range for the modified spectrum is changing one side of the equation without changing the interpretation of the expected value.

In plain terms, only those droplets that froze between –13 and –15 ºC were removed, ignoring the fact that droplets that froze before –13 ºC could have contained nuclei that would freeze from –13 and –15 ºC, that the droplets that froze between –13 and –15 ºC could have been late (or early) freezers from the nuclei that are active on the edges of either of the remaining ice modes of ice nucleation activity, or that the droplets that remain just before –13 ºC and just after –15 ºC could have been early (or late) freezers from the region of nuclei that was removed. This does not influence the droplets far from the temperature region that was removed, because it was very unlikely they would freeze in the removed region. However, close to the edges of the removed region, the probability that *only* those droplets that froze in the region of interest and *none* of the droplets outside of the region were removed is unlikely. The variable temperature interval partially corrects for this, as the sudden increase in distance between points results in a drop in calculated ice nucleation probability represented by k(T). The fixed temperature interval method as shown does not correct for this.

Third, there is also a nuance hidden in the fixed temperature interval example given in V23: Consider an ice nucleation spectrum similar to the original shown in Figure 1 with a 10x larger population of droplets using the same fixed temperature intervals used above. Again, let the points between –13 and –15 ºC be removed and the spectra recalculated. Instead of 1 or 2 droplets per bin, now there are 10 or 20. Now, instead of the gap in ice nucleation being exactly on the edge of the fixed temperature bins, shift the temperature bins by half their width in either direction. Now, half of the droplets in each of the edge bins have been removed. What is the resulting change to the spectrum? *Exactly the same as that observed using* *the variable temperature bins.* The bins on either edge of the gap will be reduced by approximately 50%, simply because the frame of reference from which the bin locations have been calculated has shifted. If the reduction in k(T) using the variable temperature bins is an undesirable artifact, surely this would be too? However, we argue it is the inverse: The unchanged edge points when the fixed temperature bins are used is the undesirable artifact, and the variable temperature bin or 'shifted' fixed temperature bins are the more accurate response.

Final thoughts and caveats

None of the arguments above should be interpreted to mean that representing an ice nucleation spectrum discretely or using fixed temperature bins for calculating k(T) is invalid. Whether continuous or discrete, fixed or variable temperature intervals, each calculation method makes its own assumptions that must be carefully considered. In F22, it is argued that the continuous interpolations have increased statistical power than discrete binning methods of interpolation, not that the discrete binning methods are inherently incorrect. The reason continuous interpolations are adopted is to facilitate comparisons between ice nucleation spectra and to avoid loss of information about the shape of the ice nucleation spectrum, which coarse binning approaches might obfuscate. This benefit is likely minimal in traditional pipetted droplet-on-substrate approaches, but we posit it will be increasingly useful with microfluidic or similar approaches that analyze large ensembles of droplets. Similarly, it is argued that variable temperature intervals allow for better use of the collected data and avoids loss of information by treating each point separately, not that fixed temperature intervals are incorrect. When experimental uncertainty is considered, the two calculation methods may be statistically insignificantly different as discussed by Markus Petters in his review of V23. If the assumptions made using variable temperature intervals or continuous spectra are incorrect or misaligned with the meaning of ice nucleation spectra, this argument must be clearly presented; it was not in V23.

There is one major caveat that must be considered when variable temperature intervals are used that V23 does not discuss, but Markus Petters does in his review: When there is a lack of freezing events for a long period, it is never explicitly calculated or expressed in that region of the ice nucleation spectrum that the value of k(T) is zero. As such, when variable temperature intervals are used to calculate k(T), the resulting points should be interpreted as measurements of k(T) at a given point of T, not over the entire 'bin' represented by the temperature interval. Indeed, since more than one point is never in a temperature interval at a time, this method is not really a binning approach. Instead, it treats the value ΔT as a variable, not a bin width. A lack of measurements in a region should then be interpreted literally: there were no freezing events there, so k(T) is essentially zero. A question should then be raised about how wide of a gap truly represents a value of zero, and we propose that this should be related to two factors: (1) the uncertainty in temperature measurements, and (2) the width of the probability density function for a given type of ice nucleation site measured in an experiment. The second variable could be measured using an ice nucleating material such as Snomax®, which is known to have distinctly observable types of ice nucleants (Beydoun et al., 2017).

A minor additional caveat as discussed by Markus Petters in his review of V23 is the increased noise resulting from using variable temperature intervals. This is a side effect of minimizing the loss of information from binning: the noise is essentially information about the random variability within the experiment. As was discussed in F22, noise can be minimized with the use of a smoothing algorithm, which has its own assumptions and can result in the loss of data and in many cases may not be superior to the smoothing provided by the fixed temperature interval binning approach. Noise is not useful, but this does show how variable temperature intervals could be used to avoid information loss in experiments with less uncertainty and higher density information. If there is nonrandom variability within a small temperature range, a fixed binning approach may miss it while the variable temperature approach will show it clearly. Of course, this usefulness is limited by the uncertainty of a given experimental setup.

We do acknowledge that in the footnote on line 15 of V23, Gabor Vali is correct that there is an error in F22's description of the k(T) equation. It was originally meant to be expressed as $N_0 - N(T)$, not N(T) alone. This error is not present in the actual data or code presented in F22, only in the descriptive text. We will contact the journal to make a correction to F22.

Finally, regardless of whether continuous spectra or variable temperature intervals are used, the methods for determining confidence limits presented in F22 are unchanged. Continuous (but in the case of k(T), not necessarily monotonic) spectra are used for ease of developing statistical tests to compare ice nucleation spectra and to increase statistical power, and variable temperature bins for calculating k(T) are not actually used in the methods recommended, as in that case k(T) is calculated from the derivative of K(T).

William Fahy, University of Toronto (formerly Carnegie Mellon University)

Ryan Sullivan, Carnegie Mellon University

References

Beydoun, H., Polen, M., and Sullivan, R. C.: A new multicomponent heterogeneous ice nucleation model and its application to Snomax bacterial particles and a Snomax-illite mineral particle mixture, Atmospheric Chemistry and Physics, 17, 13545–13557, https://doi.org/10.5194/acp-17-13545-2017, 2017.

Fahy, W. D., Shalizi, C. R., and Sullivan, R. C.: A universally applicable method of calculating confidence bands for ice nucleation spectra derived from droplet freezing experiments, Atmos. Meas. Tech., 15, 6819–6836, https://doi.org/10.5194/amt-15-6819-2022, 2022.

Petters, M.: Comment on amt-2023-138, Atmos. Meas. Tech. Discuss. [review] https://doi.org/10.5194/amt-2023-138-RC2, 2023

Vali, G.: Comment on "A universally applicable method of calculating confidence bands for ice nucleation spectra derived from droplet freezing experiments" by Fahy, Shalizi and Sullivan (2022), Atmos. Meas. Tech. Discuss. [preprint], https://doi.org/10.5194/amt-2023-138, in review, 2023.

Vali, G.: Repeatability and randomness in heterogeneous freezing nucleation. Atmos. Chem. Phys., *8*(16), 5017–5031. https://doi.org/10.5194/acp-8-5017-2008, 2008.

Vali, G.: Quantitative Evaluation of Experimental Results on the Heterogeneous Freezing Nucleation of Supercooled Liquids. J. Atmos. Sci., *28*(3), 402–409. https://doi.org/10.1175/1520-0469(1971)028<0402:QEOERA>2.0.CO;2, 1971.