The Cryosphere Discuss., 5, C1721–C1728, 2012
www.the-cryosphere-discuss.net/5/C1721/2012/

The Cryosphere
Discussions

# *Interactive comment on* "Use of a thermal imager for snow pit temperatures" *by* C. Shea et al.

**C. Shea et al.**

cashea@ucalgary.ca

Received and published: 14 January 2012

Previous comments are *in italics*, responses are in normal text.

*The use of thermal imaging for snow studies is very interesting work. The instant visual feedback reveals something that can't be seen with any other method.*

*Below are some remarks and suggestions for clarifications and things to reconsider for the final version.*

Thank you for your examination of our error sources, this has seeded what we feel to be a productive discussion below.

*Section 7 in general:*
*What is the effect and level of instrument noise to point gradients calculated from the images? Analyze the overall accuracy of the gained results based on errors mitigating from different sources.*

Instrument noise is related to the sensitivity of the voltage measuring circuits within the thermal imager. For the FLIR B300, the voltage across the microbolometer pixels can be measured precisely enough to measure equivalent temperature differences between pixels of 0.05 C or better. This is the NETD in the specifications, which expanded means Noise Equivalent Temperature Difference. This is the measure of noise that pertains to gradients, and to relative temperatures within one image, and to relative temperatures between images where the microbolometer is at the same temperature. Pixels within a single image have even more accuracy in measuring their differences because pixels within the same image at short distances are subject to the same atmospheric temperature, humidity, and reflected temperature.

The NETD is based upon a flat field, and not the interactions between a complex surface which may selectively inset crystals of different temperatures during preparation with a shovel (i.e. snow). As pixels are not perfectly aligned with crystal edges, this uneven surface affects both temperatures and gradients within a single image. Therefore, we additionally analyze the accuracy of the results with respect to a complex snow surface in the old Lens effects section. As per that section, we find that the variation across a complex snow surface is on the order of the NETD. From your comments and the comments of the other reviewers – which are very helpful – we now see that this two-stage analysis was not clear. Hence, we have re-named the Lens effects section to be Secondary lens effects, and rephrased the introduction of that section (new line 225-226, and new lines 287–290) to address both this and some of your comments below.

In addition, we have added some clarification on the difference between NETD and overall absolute temperature accuracy from an engineering perspective. This may be found in the new version of the Atmospheric effects section, starting on new line 228.

*Section 7.2:*

*In the (sub-)millimeter pixel scale there really is no 0 degree photographic angle, as the crystal facets are in semi-random orientation. To assume a single incidence angle, perhaps a larger area should be investigated (averaged) to get a statistically significant amount of orientations so that discrete crystal facet angles are not dominating. These values could be compared with the higher resolution data to see whether it has any significance to results.*

We are using a relatively large area as suggested - around one quarter of each image. Using an even larger area would not allow the differences within the same scene to be measured at multiple places within one image, and so we chose this as a compromise between sample size and consistency.

*Use of natural snow when correcting the lens effects is not a good thing. Natural snow has spatially varying thermal structures (the very thing you are looking for!) which will not give satisfactory calibration as the target used for calibration is unknown. The calibration for lens effects should be done using a isothermal target (tens of images, median stack the temperature arrays, normalize, use that as a flat-field correction for all the data). From this target the authors would also gain the level of instrument noise (pixel-to-pixel variance in measured temperature).*

If our goal was to correct for only the lens, then yes. However, the lens is already corrected for at manufacture, by using a flat field emitter, as you describe. Earlier, you pointed out that there is no 0 degree photographic angle on snow, and this is the effect we desired to assess the magnitude of. Due to the presence of lens angles, the lens can 'see' to a different depth in the crystal interstitial spaces at different points on the lens. Our basic investigation with the lens corrections is: does this slightly different 'depth' of imaging on a complex snow surface create large or unpredictable thermal differences across the lens as to be uncorrectable? As shown in the paper (new lines 319-323), we found this not to be the case, that it was correctable, and on the order of the equipment noise.

*All the corrections (atmospheric, lens effects, exposure time) should be done to absolute values, not derivative quantities. Some of the errors are additive (atmosphere), some multiplicative (lens effects) and some function of external forcing (exposure: time of being exposed before imaging and the temperature gradient between snow and air). This is to say: no linear adjustment will give a correct result! The corrections presented in the manuscript seem to be negligible compared to the accuracy (+-2 %) given by the manufacturer (is that +-0.2 C at -10 C or +-5.3 K at 263 K? I hope the first. Please check and present in the paper). The data might not need all the corrections, but if a correction is applied, it should be based on theoretically sound principles.*

We agree that any correction should be based on sound principles, but not necessarily theoretical ones. For example, FLIR applies an empirical correction to translate between observed microbolometer voltage and actual flat field temperature, and this correction is obtained by empirically calibrating the microbolometers. This calibration occurs at the factory by shining a known temperature near-blackbody at a known-temperature microbolometer, and programming (or re-programming) the relation into the integrated circuits. This type of correction is not based on theoretical principles; rather, it is based on extensive engineering ones.

The 2% error comes from the variation of the points used to measure voltage (or current, if voltage is kept constant) across the resistive pixel, or other connections. The electrodes used to actually do the measuring actually change in resistance themselves with their own temperature. The common materials used for these electrodes vary approximately 2% across the thermal spectrum that the high-temperature FLIR B300 can measure. You are correct that this, among the many other errors that the thermal measurement is subject to, are not linear. Therefore, as stated on old page 2533, line 23-25, atmospheric corrections must occur on the raw voltages so that the factory scaling to temperature may then be used on the new voltages.

That said, you (and second reviewer J. Dozier) have a very good point that it is certainly misleading to describe this empirical correction as directly coming from physical laws such as the Stefan Boltzman Law. The physical laws are related and accounted for in the empirical relation, but the relation also accounts for many engineering details unrelated to emissivity laws. So, we have removed references to this physical basis with wattages, and refer to it simply as a calibration. FLIR does, however, state explicitly that the modification of the raw voltages occurs as per new Equation 2 (old Equation 3), and so we have retained that description.

It may be useful to further distinguish between the 2% absolute temperature accuracy versus the 0.05 between-pixel accuracy from an engineering perspective. We have added a summary to this extent in the paper, new lines 228-260. Concerning the 2%, it is more common for the entire microbolometer itself to be at a different, unknown temperature than a steep gradient to exist across the microbolometer itself. As long as the microbolometer itself is at a relatively homogeneous temperature with only few degrees or so variation (relatively easy to achieve when being encased in a camera body), and the measurement of interest is the difference in temperature (resistance) between pixels, the error improves to well below 2% – 0.05 C in the case of the B300. This is because the limiting question is no longer the physical change in electrode material (or the effect of its heat capacity) over a wide range of temperatures, which is related to its physical structure and the absolute (and unknown) temperature of the microbolometer itself. Instead, when the microbolometer varies by only a few degress across its width, the differences in resistance between pixels due to the thermal variation and heat capacity of the electrodes is negligible.

So instead, the accuracy of the difference between pixels depends on the sensitivity of the voltage-difference-measuring integrated circuits, which is then a question of price. As described earlier within this response, the B300 advertises that it provides 0.05 C NETD (the base noise in a reasonably thermally homogenous microbolometer) in the worst case. This is probably based on the voltage difference sensitivity of the circuit that FLIR chose to put in the camera, along with other internal engineering principles, and so it is not something that can be theoretically re-derived by the user.

Therefore, for purposes of measuring gradients, the (engineered) accuracy is 0.05 C or better, and for purposes of measuring absolute temperatures, the (engineered) accuracy is 2% over the entire range, as the internal microbolometer measurement varies by that amount. Figure 3 does a fairly good job of showing the relative independence of absolute temperatures and their fluctuation, compared with the relative stability of between-pixel measurements over time; however, it is unclear (and somewhat beyond the scope of a paper concerned primarily with NETD, or absolute temperature observations supported by NETD) how much of this is due to sensor variation and how much to the snow's equalization upon exposure.

*Section 7.3:*
*In Chapter 1 the authors state that there is error in traditional point measurements due to temporal changes, which is absent in thermographs since the imager records instantaneously the whole field of view.*

In Shea and Jamieson (2011), we show that heat from the observer penetrates deeper than the length of thermometer shafts within 30 minutes of pit wall exposure. It is not uncommon for field practitioners to move their thermometers slowly down the pit wall as they perform other observations. However, this means these thermometer point measurements can easily be split with some prior to observer heat penetration and some afterwards, skewing the findings.

*However, there is no indication how to take temporal changes into account when thermographs from different heights are combined, or how these images are related to natural undisturbed condition prior to excavations.*

In practice, we have found thermal images spaced within a minute or so from each other at different heights to combine easily and cleanly, especially when the iterative exposure of parts of a pit wall is done (e.g. digging out only as much as can be imaged quickly). However, all of our case studies intentionally focus on obtaining data around a single layer of interest within a single image, or between images where the camera sensor has been allowed to equalize and only then in conjunction with supporting gradient data, as in the case of Figure 4 and Figure 6. This is to focus on the basics of using a thermal imager for a single area, rather than adding complexity with combining images. In practice, combining images seems rather straightforward, stable, and therefore feasible, but to treat the topic formally it deserves further attention after the basics within this paper are completed.

We do discuss how these images are related to natural undisturbed conditions starting on old page 2536, line 26, and again on old page 2538, line 5. This content now appears starting on new line 334, and again at new line 364.

*In traditional measurements the snow itself acts as a protective insulator towards external disturbance. The temperature is normally measured 10 - 20 cm within the snow. For the exposed wall, there is no such mechanism to stabilize the snow, and the imaged surface starts to react to the change immediately. Authors should offer an explanation to the jump in absolute value of orange curve in Fig. 3a. By intuition the observed pit wall temperature should move asymptotically towards the air temperature at the same height.*

*A time series (image every 5 or 10 seconds for a few minutes using a tripod to minimize operator heating and camera movement) from the exposed wall (and simultaneous air temperature at different heights!) should also give a good estimate how to correct the absolute values in respect to time of exposure, and yield a method to estimate the true natural thermal state of the undisturbed snow.*

This is not necessarily true. For example, another effect on snow temperature post-exposure is the difference between humidity within the snowpack and in the atmosphere. Even if the outside air was warmer than the snow, if it was at very low humidity, then the snow – being immersed before exposure within the near-saturation environment of pore spaces – would first succumb to cooling via sublimation. Then, after that equalized, the snow would begin warming via the less strong effect of conduction and

convection from the warm air. One can see how exposure creates a complex equalization process, and so we feel that proposing the order of causes for Figure 3 would be little more than guessing. This is why we focus on gradients primarily, and why we also show that gradients, by far and large, decrease in magnitude over time.

Our solution to the equalization issue is our suggestion in using the first, fastest image (assuming subsequent images show gradient smoothing) rather than proposing a correction for later images to correct for time. The time lapse homogenization to eventual air temperature is neither linear nor constant between different condition sets. So, correcting later images to be equivalent to fresh images seems infeasible. Instead we take fresh images – having the minimum time possible to change from their unexposed state – to be the most accurate representation of the natural conditions prior to exposure. As gradient smoothing over time is by far the norm, it seems to be a reasonable assumption that we are gaining a smoothed glimpse of an even more extreme natural buried state. This is described starting at new line 368.

*Reviewer suggestion: Either recalibrate the temperature data with a scientifically sound method or show that no calibration in post-processing is needed to reach the conclusions presented.*

Between the general engineering design principles of resistive temperature devices, as well as our investigations of secondary lens effects and time effects after exposure specific to snow, we feel we have shown the corrections used to be reasonable for the conclusions we present. That said, we appreciate your attention to ensuring the process is both well understood and valid. These details were certainly not clear in the original discussion paper, and your open discussion here will hopefully help future users of thermal imagery improve upon our processes.

---

Interactive comment on The Cryosphere Discuss., 5, 2523, 2011.