

This page intentionally left blank.

© Copyright by Peter J. Gadowski 2016

All Rights Reserved

MEASURING GLACIER SURFACE VELOCITIES WITH
LIDAR: A COMPARISON OF THREE-DIMENSIONAL
CHANGE DETECTION METHODS

A Thesis

Presented to

the Faculty of the Department of Civil and Environmental Engineering
University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Master of Science

in Geosensing Systems Engineering and Sciences

by

Peter J. Gadomski

December 2016

MEASURING GLACIER SURFACE VELOCITIES WITH LiDAR: A COMPARISON OF
THREE-DIMENSIONAL CHANGE DETECTION METHODS

Peter J. Gadowski

Approved:

Chair of the Committee
Dr. Craig L. Glennie
Associate Professor
Geosensing Systems Engineering and Sciences

Committee Members:

Dr. Preston Hartzell
Research Assistant Professor
Geosensing Systems Engineering and Sciences

David Finnegan
Research Physical Scientist
Cold Regions Research and
Engineering Laboratory
Hanover, NH

Dr. Suresh K. Khator
Associate Dean
Cullen College of Engineering

Graduate Program Director
Dr. Craig L. Glennie
Associate Professor
Geosensing Systems Engineering and Sciences

ACKNOWLEDGEMENTS

This thesis is dedicated to the memory of Dr. Gordon Hamilton, who died October 22, 2016 while working in Antarctica. He was an inspiration personally and professionally.

MEASURING GLACIER SURFACE VELOCITIES WITH
LIDAR: A COMPARISON OF THREE-DIMENSIONAL
CHANGE DETECTION METHODS

An Abstract

of a

A Thesis

Presented to

the Faculty of the Department of Civil and Environmental Engineering

University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Master of Science

in Geosensing Systems Engineering and Sciences

by

Peter J. Gadomski

December 2016

Abstract

Using airborne and terrestrial LiDAR data from glaciers in Greenland and Antarctica, we compare three change detection methods for accuracy and performance. We focus in particular on one method, Coherent Point Drift (CPD). We find that CPD outperforms Iterative Closest Point (ICP) and Particle Imaging Velocimetry (PIV) when used on a terrestrial LiDAR dataset at the Helheim Glacier in southeast Greenland. At one representative location, CPD calculated an average glacier velocity of 20.11 m d^{-1} with Root-Mean Squared Error of 2.5 m d^{-1} when compared to a GNSS-derived measurement of 20.44 m d^{-1} . All three change detection methods fail to fully capture the motion of the Canada Glacier in Antarctica, but do detect change in the fast-moving and crevassed portion of the glacier. We conclude that these change detection methods, and CPD in particular, are useful tools for measuring glacier velocity when the data have sufficient identifiable features in both epochs.

Table of Contents

Acknowledgements	v
Abstract	vii
Table of Contents	viii
List of Figures	x
List of Tables	xiii
1 Introduction	1
1.1 Measuring glacier surface velocities	1
1.2 Three-dimensional change detection	4
1.3 Objectives and outline	5
2 LiDAR	7
2.1 Types of LiDAR scanners	8
2.1.1 Pulse-based LiDAR scanners	8
2.1.2 Phase-based LiDAR scanners	13
2.2 Scanner's Own Coordinate System	14
2.3 Modes of LiDAR scanning	16
2.3.1 Terrestrial LiDAR scanning	16
2.3.2 Airborne and mobile LiDAR scanning	17
2.4 Error in pulse-based LiDAR scanning	19
3 Pointcloud change detection	21
3.1 Particle Imaging Velocimetry (PIV)	21

3.2	Iterative Closest Point (ICP)	23
4	Coherent Point Drift	26
4.1	Overview	26
4.2	Explanation	26
4.3	Rigid transformation	28
4.4	Fast Gauss Transform	30
4.5	Software implementation	30
5	Test data and experiments	32
5.1	Helheim Glacier, Greenland	32
5.1.1	Glacier surface velocities and artificial scene compression	34
5.1.2	Absolute accuracy of change detection methods	39
5.1.3	Relative accuracy of change detection methods	41
5.2	Canada Glacier, Antarctica	42
6	Results	45
6.1	Synthetic displacements	45
6.2	Time delta between scan pairs	48
6.3	Tuning CPD	48
6.4	Tuning ICP	55
6.5	Tuning PIV	58
6.6	Comparing change detection methods at the GNSS receivers at Helheim Glacier	62
6.7	Comparing change detection methods at the three additional Helheim Glacier points	66
6.8	Full-glacier change detection at Helheim Glacier	69
6.9	Full-glacier change detection at Canada Glacier	73
7	Conclusions and future work	78
	References	80

List of Figures

2.1	Outgoing (reference) pulse from a Riegl LMS-Q1560 airborne scanner.	9
2.2	Incoming (return) energy from one channel of a Riegl LMS-Q1560 airborne scanner.	10
2.3	MTA zones as a function of pulse rate.	12
2.4	The Scanner's Own Coordinate System (SOCS), a right-handed coordinate system. Left image is a side view of the scanner. Right image is a back view of the scanner looking in the direction of laser emission. θ is the direction of the laser beam as it rotates from $+z$ to $-z$. ϕ is the rotation of laser head around the vertical axis. . . .	15
3.1	Schematic of an Iterative Closest Point iteration, matching step on left, transformation step on right.	24
5.1	Helheim Glacier, Greenland.	33
5.2	Location of scanner and GNSS receivers during summer 2014 collects at the Helheim Glacier. Also labelled are three additional points used for relative velocity comparisons.	35
5.3	Timing of TLS scans and GNSS receiver location records for the summer 2014 collect at the Helheim Glacier.	36
5.4	Time offsets within a single scan of the Helheim Glacier. Scanner head rotation was counterclockwise. Time values are GPS week time.	38
5.5	Canada Glacier in the McMurdo Dry Valleys region of Antarctica.	43
6.1	Synthetic displacements (x-axis, in meters) and the residuals of CPD-calculated displacements (y-axis, in meters) at hg02 and h03 from scan data collected 2014-07-11 00:33 UTC at the Helheim Glacier, in the presence of various levels of noise.	47

6.2	CPD-calculated glacier velocities at hg02, divided vertically by velocity component (x-velocity on top, y-velocity in the middle, z-velocity on the bottom) and horizontally by nominal time gap between scans in minutes.	49
6.3	CPD-calculated glacier velocities at hg03, divided vertically by velocity component (x-velocity on top, y-velocity in the middle, z-velocity on the bottom) and horizontally by nominal time gap between scans in minutes.	50
6.4	CPD-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg02 for crop widths (horizontal divides) of 100 m, 250 m, and 500 m. Because of data gaps due to intervals of generator failure during scan collection, x-axis is scan pair ID, not time of collection.	51
6.5	CPD-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg03 for crop widths (horizontal divides) of 100 m, 250 m, and 500 m.	52
6.6	CPD-derived velocities at hg02 for three different outlier weights. The outlier weight is the single free parameter, between zero and one, in CPD. There is not discernible different in results in this range of outlier weights.	54
6.7	ICP-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg02 for a variety of crop widths (horizontal divides).	56
6.8	ICP-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg03 for a variety of crop widths (horizontal divides).	57
6.9	PIV-derived x and y velocities (vertical divide) at GNSS site hg02 for a variety of pixel sizes (horizontal divides).	58
6.10	PIV-derived x and y velocities (vertical divide) at GNSS site hg03 for a variety of pixel sizes (horizontal divides).	59
6.11	PIV-derived x and y velocities (vertical divide) at GNSS site hg02 for both peak-finding algorithms (horizontal divide).	61
6.12	Three-dimensional change detection at hg02, 240 min between scan pairs.	63
6.13	Three-dimensional change detection at hg03, 240 min between scan pairs.	64
6.14	Glacier surface velocities at the Helheim Glacier, as computed by the three change detection methods. Results are separated by location, r1, r2, and r3, in the horizontal dimension, and by velocity component x, y, and z, in the vertical dimension.	68

6.15	Full-glacier change detection of Helheim Glacier. 2014-07-13 18:17:00 to 2014-07-13 22:24:12. Arrows indicate direction and magnitude of velocity in the x-y plane, background indicates total velocity magnitude.	70
6.16	Typical point density, in points/m, for a scan of the Helheim Glacier. Scan data from 2016-07-10 20:28:08 UTC.	71
6.17	Frequency distributions of velocity vectors for change detection runs on scans from 2014-07-13 18:17:00 to 2014-07-13 22:24:12. Only x and y velocities are presented.	72
6.18	Surface roughness at the Canada Glacier. Surface roughness has the units of m and is defined as the largest inter-cell difference between a central pixel and its surrounding cells. Data from the 2001 NASA collect.	75
6.19	Surface roughness at the Helheim Glacier. Surface roughness is in units of m and is defined as the largest inter-cell difference between a central pixel and its surrounding cells. Data from the TLS scan at 2016-07-10 20:28:08 UTC.	76
6.20	Full-glacier change detection at the Canada Glacier. Velocities are presented in m/year due to the slow-moving nature of the Canada Glacier. Results are filtered to only present velocities where the average roughness is greater than 0.75 m, where roughness is defined as the largest inter-cell difference between a central pixel and its surrounding cells.	77

List of Tables

5.1	Terrestrial LiDAR scanner parameters used at the Helheim Glacier in Greenland in July 2014.	34
5.2	Three additional points for relative velocity comparisons at the Helheim Glacier. . .	41
5.3	Parameters for the Airborne Laser Terrain Mapper (ALTM), used for 2001 Canada Glacier collect by NASA.	42
5.4	Parameters for Optech Titan LiDAR scanner, used for 2014 Canada Glacier collect by NCALM.	44
6.1	Root mean squared error for CPD-calculated translations, as compared to applied translations, at various noise levels.	46
6.2	Statistics for CPD-derived velocities compared to GNSS-derived velocities at hg02 separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1}	48
6.3	Statistics for CPD-derived velocities compared to GNSS-derived velocities at hg03 separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1}	53
6.4	Statistics for ICP-derived velocities compared to GNSS-derived velocities at hg02, separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1}	55
6.5	Statistics for ICP-derived velocities compared to GNSS-derived velocities at hg03, separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1}	55
6.6	Statistics for PIV-derived velocities when compared to GNSS-derived velocities at hg02, separated by pixel size. All values in m d^{-1}	60

6.7	Statistics for PIV-derived velocities when compared to GNSS-derived velocities at hg03, separated by pixel size. All values in m d^{-1}	60
6.8	Statistics for PIV-derived velocities at hg02, compared to GNSS-derived velocities, separated by peak-finding algorithm. All values in m d^{-1}	61
6.9	Statistics for PIV-derived velocities at hg03, compared to GNSS-derived velocities, separated by peak-finding algorithm. All values in m d^{-1}	61
6.10	Mean and RMSE between change detection methods and GNSS measurements at hg02. Values in the RMSE columns for GNSS velocities are the standard deviation, calculated from the accuracy of each GNSS point propagated through to the final velocity. All values in m d^{-1}	65
6.11	Mean and RMSE between change detection methods and GNSS measurements at hg03. Values in the RMSE columns for GNSS velocities are the standard deviation, calculated from the accuracy of each GNSS point propagated through to the final velocity. All values in m d^{-1}	65
6.12	Crop radii used for ICP and CPD at the three additional relative comparison points on the Helheim Glacier.	67

1 Introduction

As the Earth's climate warms, melting ice sheets will be a primary driver behind sea level rise, a problem which has potentially devastating effects for human populations and shoreline ecosystems (IPCC 2007; Pfeffer et al., 2008). An understanding of the mass-balance and melt dynamics of these ice sheets is therefore needed to predict the magnitude and extent of sea level rise (Rignot et al., 2011).

It has been shown that the dynamics of outlet glaciers from the Greenland Ice Sheet (GIS) play a significant role in the overall mass-balance of the GIS (Parizek and Alley 2004). Glacier surface velocities are crucial for understanding the mechanics of glacier advance, retreat, and mass loss. Many efforts have been made to characterize the changeable surface velocities of these outlet glaciers from the macro (Rignot and Kanagaratnam 2006) to meso (Stearns and Hamilton 2007) scales. These surface velocities are then used to inform models (Whillans and Bindschadler 1988; M. L. Andersen et al., 2010), explain specific components of glacier motion such as glacial earthquakes and their response to calving (Nettles et al., 2008), and estimate mass loss (Howat et al., 2007). A variety of techniques are used to measure glacier surface velocity, each with its own typical coverage and resolution in the spatial and temporal domains.

1.1 Measuring glacier surface velocities

Direct measurement of glacier surface velocities is possible by placing GNSS receivers onto the ice (Ahlstrøm et al., 2013). These receivers are left on the ice for a period of time to record their position while the ice flows downstream, providing data with very high accuracy and temporal resolution. On-ice GNSS surveys are limited by their spatial scope and their logistical complexity; placing personnel and equipment on the ice is never a risk-free endeavor. Nevertheless, on-ice GNSS receivers are a valuable tool for rapid, accurate glacier velocity measurements, and they are crucial for investigating the effects of large scale calving and glacial earthquakes on a short time scale (Nettles et al., 2008). GNSS receivers can also provide insight to the magnitude and extent of movement due to flotation at the tongue of a glacier (Juan et al., 2010).

Because direct measurements with GNSS are difficult, remote sensing techniques are often used; they generally provide better spatial coverage with significantly less individual expense and risk. The three most common methods are feature or speckle tracking via imagery or radar, differential interferometric synthetic aperture radar (dInSAR), and photogrammetric techniques. These methods are often used on data collected by satellites, but uses on data from airborne and ground-based sensors are also possible. Airborne or terrestrial surveys allow for faster repeat times and better spatial resolution, but they can be expensive and limited in spatial scope.

Feature tracking, whether through passive or active sensing, uses distinct features in an electromagnetically-derived image to approximate the two-dimensional velocity field of a glacier (Lucchitta and Ferguson 1986; Scambos et al., 1992; Joughin 2002). Though passive imagery is restricted to use during daylight hours with clear skies, active methods such as radar speckle-tracking are suitable at night and through cloud cover, making these active methods a good fit for the extreme climates and latitudes at which many glaciers exist. Satellite-mounted sensors also provide excellent spatial coverage and ever-improving spatial resolution, making them ideal for meso- and macro-scale study of glacier velocities. Where these methods fall short is both capturing truly three-dimensional data and characterizing the short-term (hours to days) variations in glacier surface velocity. Most feature tracking algorithms for two-dimensional imagery rely on pixel-matching or other techniques that operate in the image plane, which can only capture motion in that image plane (Heid and Kääb 2012). The image cross-correlation necessary for good optical image matching usually forces the temporal periodicity of the data to scale with the image resolution, with coarser resolution images only able to resolve glacier displacements and velocities over longer time intervals (Heid 2011).

A second method, dInSAR, uses the phase change between two synthetic aperture radar images to calculate the spatial change in a scene (Berardino et al., 2002; Rignot et al., 2011). This complex but highly precise method can detect millimeter-scale changes in a scene, making it well-suited for wide-area, gradual changes. dInSAR falls short when a scene changes quickly or in a disorganized way, which causes a loss of coherence between the pair of SAR images (Rignot 1996; Michel and Rignot 1999). Current dInSAR change detection exercises are limited by the repeat interval of the collecting satellites, though this may improve with new satellites such as the NASA-ISRO SAR Mission (NISAR) (Hensley et al., 2014).

Finally, photogrammetry uses two or more images, taken from different positions and angles, to

extract three-dimensional data from two-dimensional imagery (Yamaguchi et al., 1993; Stearns and Hamilton 2007). Used well, photogrammetry can provide high-resolution three-dimensional data that can then be differenced to extract change vectors. In practice, retrieving data with acceptable spatial resolution and accuracy is limited by daylight, the absence of cloud cover, and sufficient ground control points for image registration. Despite its limitations, photogrammetry is a promising method that can provide a relatively affordable way to capture truly three-dimensional data, in particular from unmanned aerial sensors (UAS). UAS platforms armed with only cameras can produce high-density, reasonable accuracy Structure from Motion (SfM) point clouds, which then can be compared between epochs to extract displacements and velocities (Turner et al., 2012; Ryan et al., 2015; Ahn and Box 2010). In order to extract these velocities, we must utilize the same three-dimensional change detection techniques that we use for LiDAR data (discussed next), making the increased prevalence of SfM point clouds an additional motivation for our work.

Another method, upon which we will focus, is called Light Detection And Ranging, or LiDAR. LiDAR uses active visible or near-infrared laser energy to rapidly measure distances between a sensor and a target (Measures 1985; Glennie et al., 2013). Many LiDAR units use a time-of-flight principle to measure distance; the scanner emits a pulse of laser energy that travels through space, hits a target, and is reflected back to the sensor. This two-way travel time is then used to calculate the distance from the sensor to the reflective target. LiDAR units can be mounted on satellites, airborne platforms, or vehicles, and they can be used in a static tripod-based setup. Modern systems capture many hundreds of thousands of points a second, and modern LiDAR scanners have an effective range of several kilometers over snow (Deems et al., 2013).

Glacier surface velocities can be calculated from two LiDAR scans taken at two distinct epochs. Terrestrial LiDAR scans in particular can provide moderately good temporal resolution and excellent spatial resolution, though their spatial coverage is limited. LiDAR scanners are also significantly more expensive than, say, a photogrammetry setup, but their ease of use and high accuracy still makes them useful tools for ground-based change detection. The limiting factor for the use of high-resolution three-dimensional spatial data, such as those collected through Structure from Motion photography and by LiDAR sensors, is usually not collecting data themselves; the problem lies in the change-detection algorithms used to convert spatial data into estimated velocity vectors.

1.2 Three-dimensional change detection

One-dimensional change, such as the calculation of snow depth from Digital Elevation Models (DEMs) (Deems et al., 2013), is relatively straightforward to determine; most of the complexity lies in the computation of the DEM, not in the DEM differencing. As mentioned above, two-dimensional change detection methods such as feature/speckle tracking are well-established. Truly three-dimensional change is problematic.

One approach for calculating three-dimensional change borrows from the two-dimensional realm; extract features from the data and compute the change for those features (Kusari 2015; DeLong et al., 2015). This works well in organized environments, usually anthropogenic scenes such as cities or the interiors of buildings. Three-dimensional feature extraction change detection algorithms generally use simple shapes, such as planes and cylinders, to simplify a complex three-dimensional point cloud to a relatively small collection of parametric models, which are computationally easy to compare for spatial displacement and transformation. Feature-tracking breaks down in disorganized and non-regular environments, such as the surface of a glacier. It is possible to introduce artificial features into the environment, such as retroreflective targets for use by LiDAR scanners, but this is fraught with additional complexity; single points on the surface of a glacier are both dangerous to deploy and limited in their spatial coverage.

Another change detection method borrows from the field of hydrodynamics, specifically from the laboratory method of optical flow diagnostics: Particle Image Velocimetry (PIV) (Westerweel 1997). In the traditional usage, PIV uses two or more particle images to deduce the underlying flow field of a scene. In hydrodynamics, the particles are small tracers that are added to a fluid. For LiDAR, PIV can be used on a rasterization of the LiDAR scene where certain features (e.g. tall objects, high-intensity LiDAR returns) are augmented to serve as proxy tracer particles. By tilting the rasterization plane, PIV can be used to deduce velocity vectors in all three dimensions — however, a single run of PIV produces only two-dimensional velocity data. This method has been used to characterize the non-rigid deformation that is found in landslides (Aryal et al., 2012). As PIV is one of the only non-rigid change detection methods for LiDAR data, we will use it and compare its results to those of other methods.

The field of truly three-dimensional change detection that works with points alone is limited. A

widely-used and accepted method is the iterative closest point (ICP) algorithm (Besl and McKay 1992; Zhang et al., 2015). ICP iteratively transforms a target point cloud by using least-squares to minimize the distance between a point from the target point cloud and a point from another point cloud. ICP works well but often breaks down in the presence of noise or when applied to low-density data; many methods exist to protect ICP from noise, local minima, and other pitfalls (Rusinkiewicz and Levoy 2001; Chetverikov et al., 2002; Maier-Hein et al., 2011).

A newer three-dimensional change detection method is Coherent Point Drift (CPD) (Myronenko and Song 2010). CPD is also an iterative algorithm, but instead of using one-to-one correspondences between points it treats the target point cloud as a Gaussian Mixture Model (GMM) and uses the correspondence between a target point and all points from the other point cloud. This theoretically makes CPD less sensitive to noise, and potentially more effective for disorganized, sparse datasets, such as LiDAR scans of glaciers. A detailed study of CPD for use on real-world topography does not exist to this author’s knowledge; we endeavor to partially fill this gap with this work.

As the availability of high-density three-dimensional point clouds increases, more methods to quantitatively and accurately extract change information from those point clouds will be needed. In particular, extracting change information from disordered natural environments, such as the surface of a glacier, is more difficult than detecting movement in urban environments with well-organized, man-made features. An assessment of the absolute accuracy and performance of CPD for measuring glacier surface velocities, as well as a comparison between CPD and the more-established ICP and PIV, will provide valuable information about the utility and effectiveness of each method for use in calculating glacier surface velocities, and hopefully will provide guidance for future researchers looking to utilize these information-dense point clouds to detect three-dimensional change.

1.3 Objectives and outline

High-volume three-dimensional point clouds, such as those produced by LiDAR sensors and Structure from Motion photography, can be valuable change detection tools; but the algorithms to extract three-dimensional change from these data are not yet mature. In particular, there is a need to measure three-dimensional glacier surface velocities from point clouds and no consensus on the best way to do so exists yet. We present Coherent Point Drift (CPD), an algorithm that is not new

but that, to the best of our knowledge, has not been used before for geographic point cloud change detection. We compare the accuracy and performance of CPD with other established methods of two-dimensional and three-dimensional measurement of glacier surface velocities, including on-ice GNSS, Iterative Closest Point (ICP), and Particle Imaging Velocimetry (PIV). We assess these algorithms as used on our data, with an eye towards adding CPD to this array of commonly-used tools for three-dimensional change detection.

We will introduce LiDAR in greater depth, starting with descriptions of its use and operational concerns during collection. After LiDAR, we will discuss in-depth the iterative closest point (ICP), particle imaging velocimetry (PIV), and coherent point drift (CPD) algorithms, outlining the theory behind their operation and describing their most common variants. We will then discuss ways to apply ICP and CPD to a non-rigidly deforming body such as a glacier, since for our purposes these algorithms can only calculate rotational and translational change; CPD can calculate non-rigid change but we will not use those capabilities in this work. Finally, we will bring LiDAR and change detection together at the Helheim Glacier in eastern Greenland and in the McMurdo Dry Valleys in Antarctica. We will use ICP, PIV, and CPD to calculate glacier velocities and use on-ice GNSS data (at Helheim) to assess the accuracy of these methods.

2 LiDAR

Remote sensing, traditionally defined, is the collection and analysis of electromagnetic radiation with the purpose of measuring extrinsic and intrinsic properties of a reflective or emissive object (Lillesand et al., 2014). Remote sensing is divided into two categories, active and passive. Active remote sensing systems emit energy and measure the reflected response of that energy; passive systems collect energy from another source, usually reflected energy from the sun or energy emitted from the sensed target itself. LiDAR is active remote sensing using light amplification by stimulated emission of radiation, or laser energy. The primary use of LiDAR systems is to rapidly measure distances from a sensor to external reflective objects, thereby creating a large number of three-dimensional points called *point clouds*.

The first laser was built in 1960 by Theodore H. Maiman; soon after, lasers were used to measure distances, most strikingly by NASA as a part of the Lunar Laser Ranging (LLR) program (Dickey et al., 1994). Lasers were put to immediate use by the military for ranging, and within a decade they made their way into civilian surveying as a part of new Electronic Distance Measuring (EDM) instruments. When EDMs were equipped to also measure angles they became total stations, which are ubiquitous today and can be considered the simplest form of laser-based three-dimensional measurement instrument (Shan and Toth 2008). While a bit of a simplification, a modern LiDAR scanner can be considered an advanced total station that can collect many thousands of points a second.

In the following subsections, we will discuss the different methods by which LiDAR instruments measure distances, followed by the different operational modes of LiDAR scanning with a particular focus on terrestrial and airborne LiDAR scanning. We will then discuss error sources in LiDAR point clouds and demonstrate rigorous error propagation of these component uncertainties to final data products.

2.1 Types of LiDAR scanners

LiDAR scanners use laser energy, usually in the visible or near-infrared wavelengths, to measure distances. There are several mechanisms by which laser energy can be turned into a distance measurement, of which we recognise three primary categories: pulse-based (i.e., linear-mode), phase-based, and gieger-mode (i.e., single-photon).

2.1.1 Pulse-based LiDAR scanners

Pulse based LiDAR scanners use the two-way travel time of a pulse of laser energy to measure the distance between scanner and target, via the following relationship

$$d = c\frac{t}{2}, \tag{2.1}$$

where d is the distance between scanner and target, c is the speed of light through the transmission medium, and t is the time between pulse emission and return detection.

When a pulse-based system emits a laser pulse, part of that energy is split off and directed to an A/D converter, where the time of pulse emission is recorded. Figure 2.1 shows a typical outgoing pulse from a airborne LiDAR system. The remainder of the laser energy is emitted into the environment, where it may interact with one or more reflective surfaces and be fractionally backscattered to the LiDAR sensor. This backscattered energy is collected by the scanner via a A/D converter. Figure 2.2 shows the response of a two-channel system to backscattered laser energy. The time between pulse emission and pulse return is then used to calculate a distance via Equation 2.1.

The fundamental ranging accuracy of a pulse-based system is a function of the width (duration) of the outgoing pulse and the sampling resolution of the system’s photodetector (Baltasvias 1999). While an ideal outgoing pulse would approximate a delta function, physical limitations of the lasing device limit the outgoing energy to a roughly Gaussian profile (Hartzell et al., 2013).

As shown in Figure 2.2, the backscattered energy is often more complex than the initial outgoing roughly-Gaussian pulse. These perturbations are usually due to multiple backscattering surfaces within the finite laser beamwidth, though other factors such as suspended particles or thermal

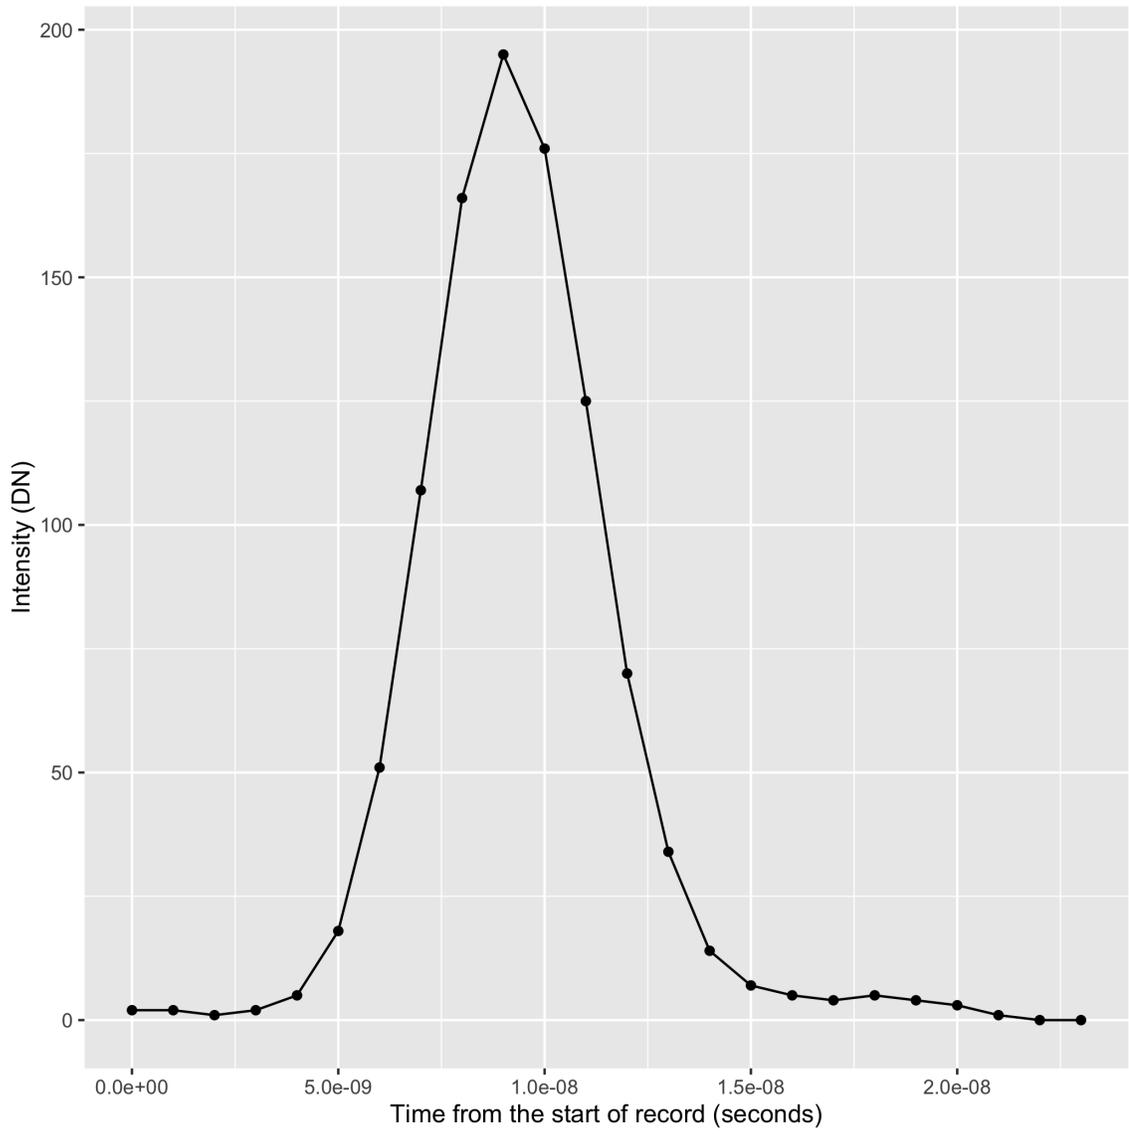


Figure 2.1: Outgoing (reference) pulse from a Riegl LMS-Q1560 airborne scanner.

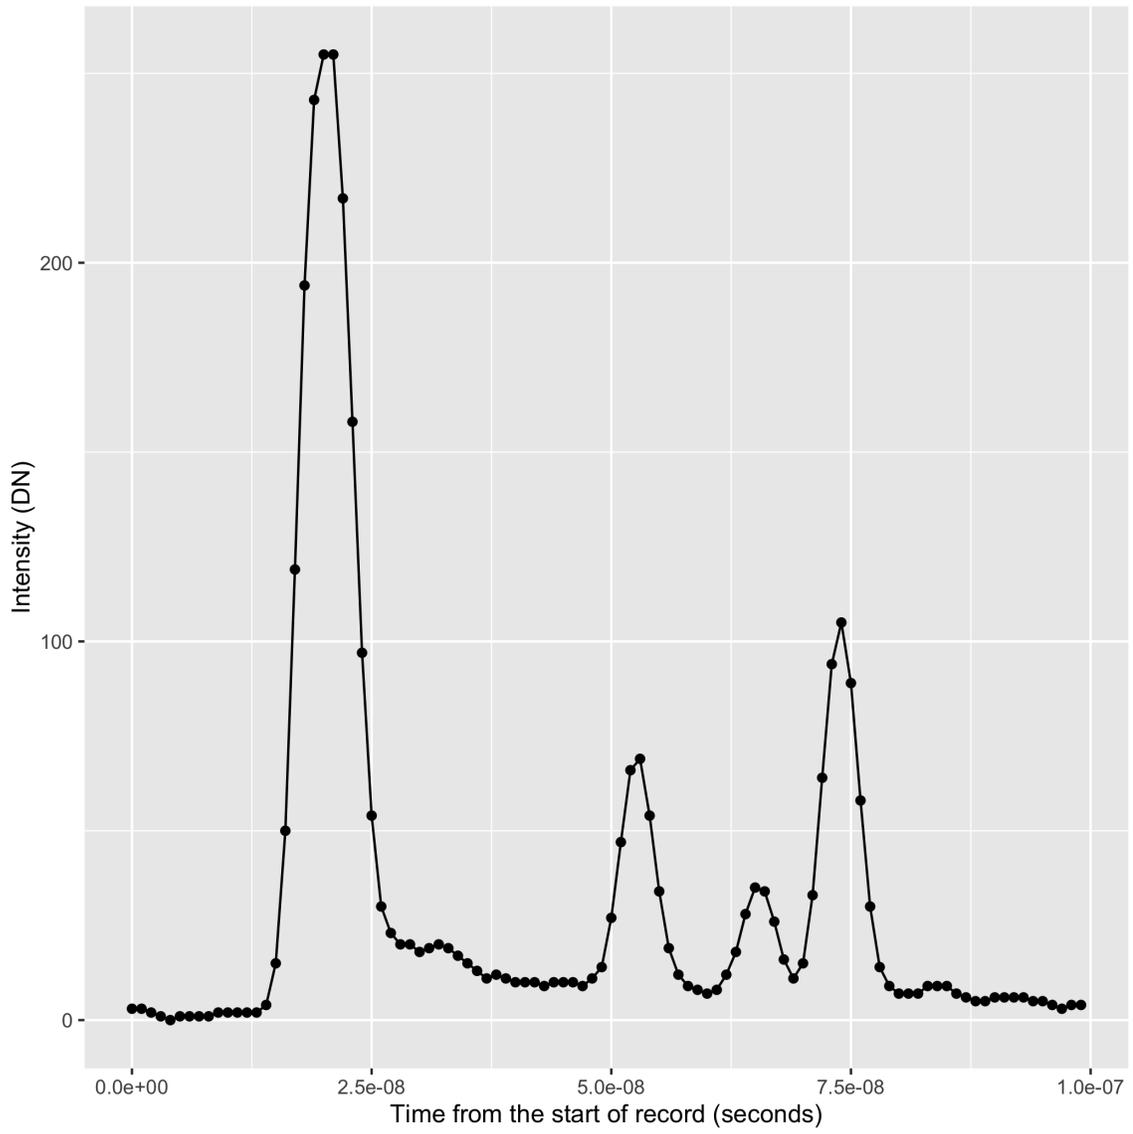


Figure 2.2: Incoming (return) energy from one channel of a Riegl LMS-Q1560 airborne scanner.

refraction can also play a part. While these “full waveform” timeseries records are used in some fields, such as bathymetry (Pan et al., 2015) and land cover classification (Wang and Glennie 2015), most data consumers want three-dimensional point clouds, not full waveforms. The process of converting full waveform data into discrete points is called “discretization.” Discretization can be done on the scanner in real-time via hardware, usually by constant fraction discrimination (Glennie et al., 2013), or in post-processing via software (Castorena and Creusere 2015; Pan et al., 2015). Regardless of their source, point cloud data from pulse-based LiDAR scanners are collections of three-dimensional points, usually with additional information such as the relative magnitude of backscattered energy and the time of data collection. It is possible, with careful calibration, to extract the absolute reflectivity of a target from the relative intensity of the backscattered energy (Wagner 2010); however, this absolute reflectance value is of limited use due to the single-wavelength nature of the laser energy (Hartzell et al., 2014).

Pulse-based scanners emit a large number of pulses per second, up to half a million in modern systems. High pulse-rates present a trade-off between data collection speeds and maximum unambiguous range; at high pulse-rates, a second pulse may be emitted before the backscattered energy from the first pulse has had time to return to the scanner. This is commonly referred to as the Multiple Time Around (MTA) problem, and it can arise during mobile, airborne, and terrestrial LiDAR scans. In order to correct for MTA problems, each point must be placed into an MTA “zone” and re-processed accordingly. Each MTA “zone” is named by the number of pulses in the air at the time of that point’s detection. The first MTA zone, starting at the scanner and expanding outwards, is zone 1; the next zone, which represents points collected with two pulses in the air simultaneously, is zone 2, and so on. Figure 2.3 presents the size of each MTA zone for a range of common pulse rates.

Pulse-based scanners have long maximum ranges and relatively small power requirements when compared to phase-based scanners. They are the most common type of LiDAR scanner for outdoor and mobile collects. All data used in this work was collected by pulse-based LiDAR scanners, and when we speak of LiDAR scanners in later sections we will be speaking about pulse-based scanners unless otherwise specified.

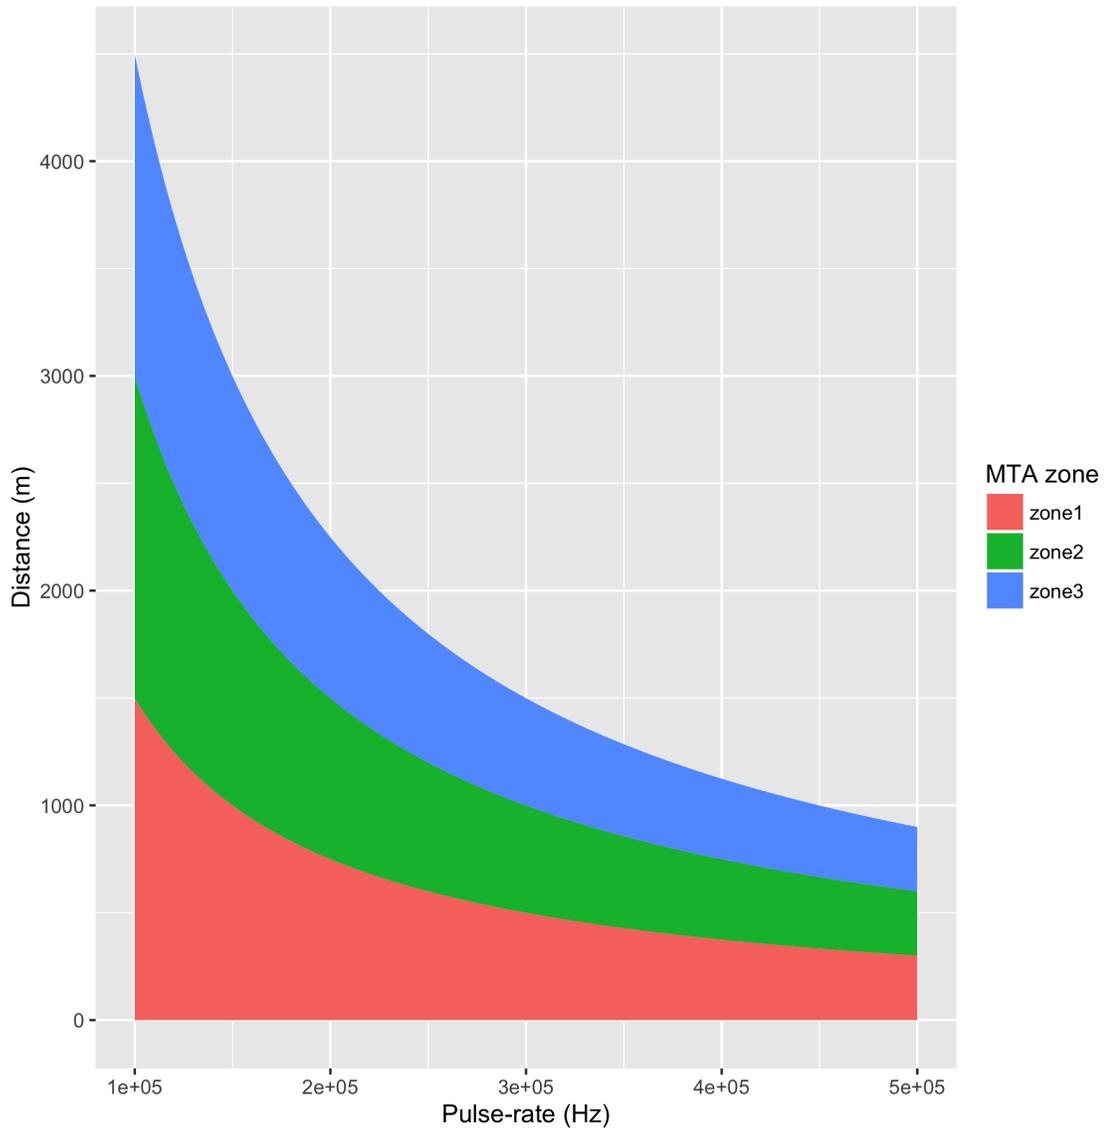


Figure 2.3: MTA zones as a function of pulse rate.

2.1.1.1 Single-photon LiDAR scanners

Pulse-based scanners require many photons of light to record a single measurement. Single-photon LiDAR systems are a subset of pulse-based scanners that make distance measurements based on the detection of a single photon of energy, using the same time-of-flight principle outlined in Equation 2.1.

One type of single-photon systems use over-voltaged avalanche photodiodes that catastrophically trigger upon the receipt of a single photon, operating in so-called “Geiger mode”; these diodes are usually arranged in array, allowing the collection of many simultaneous range measurements over a wide field of view. As might be expected, these systems are highly sensitive and very susceptible to false positives, making their resultant data inherently noisy (Aull et al., 2002; Marino and Davis 2005). Research is ongoing into feature-recognition using probabilistic methods that can extract useful information from these confused data (Cho et al., 2006). Another type of single-photon system uses a micropulse laser and photon-counting detector that does not operate in Geiger mode, but still is sensitive to single photons; the ICESat-2 project uses this type of detector (Abdalati et al., 2015).

The advantages of a useful array-based single-photon system are several, as are the disadvantages. First, the ability to detect the return of a single photon means that the required return energy is much lower, allowing either much lower-power laser emission or measurement at much longer ranges (Degnan et al., 2008). Second, the array-based detectors can collect data over a large spatial area, decreasing the time required to collect data (Herzfeld et al., 2013). Single-photon systems are limited by high cost, high processing complexity, and the inability to collect information about the radiometric properties of the target object (Marino and Davis 2005). We will not deal with any single-photon data in this study.

2.1.2 Phase-based LiDAR scanners

Phase-based, or continuous-wave, LiDAR scanners use interferometry to measure a distance from the scanner to a target. Interferometry is the use of waves to measure distance as a function of those waves’ lengths. For a phase-based LiDAR system, the waves are generated by a modulated laser signal which reflects off of a target and backscatters to the scanner’s detector. The phase of

the return signal is compared to the phase of the outgoing signal, and the distance to the target is calculated using half of the wavelength distance between the two phase measurements modulo an integer number of complete cycles.

In order to avoid a range ambiguity, the wavelength of the measured wave must be at least twice the distance from the scanner to the reflective surface. However, the precision of a continuous-wave LiDAR system is generally considered to be about 1° , which represents an unacceptable level of uncertainty for longer wavelengths (Shan and Toth 2008). In order to achieve usable maximum distances while maintaining acceptable range uncertainties, many continuous-wave LiDAR systems utilize multiple measurement wavelengths, quickly switching between them and using the longer wavelengths to help resolve the integer ambiguity in the shorter wavelengths.

Phase-based LiDAR systems generally use much more energy than pulse-based systems, and therefore are not usually used at long ranges or in outdoor settings. They are most often used in high-precision applications (Shan and Toth 2008). We have not used any phase-based LiDAR scanners in this work.

2.2 Scanner’s Own Coordinate System

Before discussing the various modes of collecting LiDAR data, it will be useful to define some terms relating to the coordinate reference systems of data collected by LiDAR scanners.

The Scanner’s Own Coordinate System (SOCS) is the coordinate reference system for all raw measurements made by the scanner. This coordinate system is usually centered somewhere inside the scanner; the most common configuration is for the x axis to point directly “out” of the scanner through the glass in the direction of laser emission, the z axis pointed vertically up out of the scanner, and the y axis orthogonal to the first two to create a right-handed coordinate system (Figure 2.4).

To produce a single three-dimensional point in SOCS, LiDAR scanners usually record three measurements in spherical coordinates:

- r : the measured range from the scanner to the target point.
- ϕ : the angle around the z axis, ranging from 0 rad to 2π rad
- θ : the angle from vertical, ranging from 0 rad to π rad

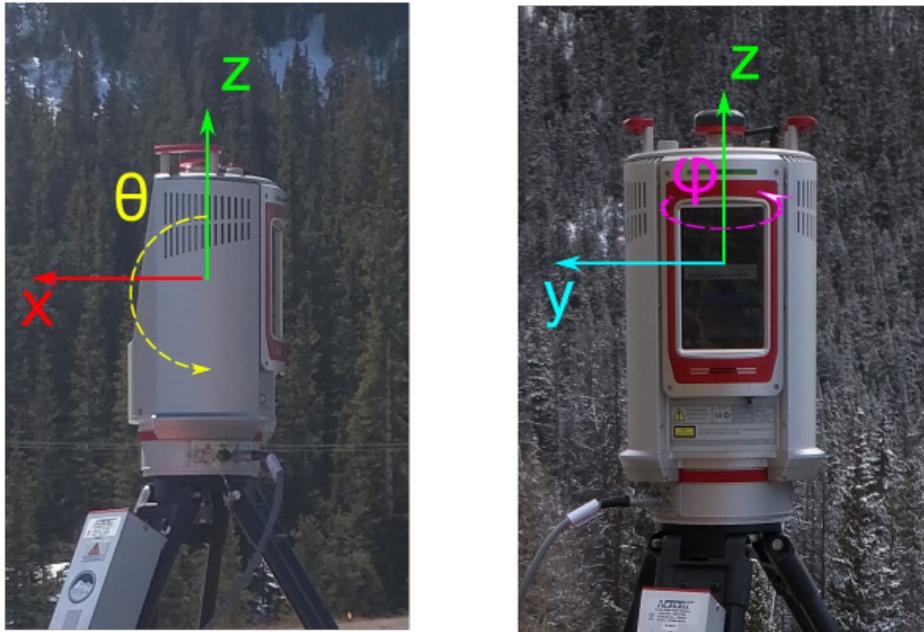


Figure 2.4: The Scanner's Own Coordinate System (SOCS), a right-handed coordinate system. Left image is a side view of the scanner. Right image is a back view of the scanner looking in the direction of laser emission. θ is the direction of the laser beam as it rotates from $+z$ to $-z$. ϕ is the rotation of laser head around the vertical axis.

Angle and range measurements (r, ϕ, θ) can be converted to SOCS Cartesian coordinates (x, y, z) by the following relationship

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} r \sin \theta \cos \phi \\ r \sin \theta \sin \phi \\ r \cos \theta \end{bmatrix}. \quad (2.2)$$

While data in SOCS is sufficient for scene-specific studies, most projects require data in some global coordinate system. The means by which data in SOCS is converted to a global coordinate system is specific to the mode of data collection, but generally requires knowing or calculating the position and orientation of the laser scanner at all times. We will discuss this conversion processes, often called “georeferencing”, in the next section.

2.3 Modes of LiDAR scanning

To this point, we have only discussed the means by which LiDAR sensors measure distances, which is only one piece of collecting full three-dimensional point data. In order to fully capture a scene, the scanner’s laser energy needs to be moved and/or rotated through space; the best means to do so vary depending on application and operational requirements. In this section will review several modes of LiDAR scanning, spending time on the operational and processing requirements for each and the methods by which LiDAR data collected in SOCS is georeferenced into a global coordinate system.

2.3.1 Terrestrial LiDAR scanning

Terrestrial LiDAR scanning, or TLS, is LiDAR scanning performed from a fixed position, usually a surveyor’s tripod or a fixed mount. Most TLS systems have a static base that is attached to the tripod, on top of which sits a rotating head that (usually) provides 360° of rotation around the vertical (z) axis. The rotation of the head is controlled by motors and is tracked with an angular encoder, recording the direction of the laser beam in the xy plane (ϕ). In order to pan the laser above and below the xy plane, TLS scanners have some sort of rotating or oscillating mirror that

directs the laser beam. In many models, this mirror is positioned inside the scanner, where it rotates and directs the laser beam out of the scanner through a pane of glass. This mirror's rotation is also measured with an angular encoder, providing θ . These two angles and the range, r , comprise the basic components required to measure a point in three-dimensional space in the Scanner's Own Coordinate System (SOCS) (Shan and Toth 2008). If the origin of laser emission moves with the mirror rotation, an additional mirror dimension s may also be required.

To georeference a TLS point cloud, the position and the orientation of the LiDAR scanner must be known. One common method is to use external tie points, or locations in the scene that have a known global xyz position. These tie points are marked in the scan using retroreflective targets, usually flat discs or volumetric cylinders, with known size and shape; the retroreflective properties of the reflectors make them easy to identify in point cloud data due to their high relative reflectance compared to the rest of the scene. These reflectors can be positioned on or over a known benchmark, or their location can be recorded using a Global Navigation Satellite System (GNSS) receiver. An array of reflectors with known positions found in a point cloud can be used to solve for the position and orientation of the scanner via least squares (Riegl LMS GmbH 2012).

Reflector arrays are an effective and robust method for georeferencing terrestrial LiDAR data, but can be cumbersome to use in rough regions or in areas where there are few benchmarks. Additionally, uncertainty in reflector position or poor reflector geometry can lead to suboptimal registration. In the case where two or more TLS scan positions need to be georegistered together, overlapping features in the point cloud can be aligned to improve registration quality (Ullrich et al., 2003; Gruen and Akca 2005).

2.3.2 Airborne and mobile LiDAR scanning

Airborne and mobile LiDAR scanning (ALS and MLS, respectively) use a LiDAR scanner attached to a moving platform. MLS generally refers to scanning from a ground-based vehicle, such as a car, while ALS includes both fixed-wing and rotor aircraft. LiDAR scanning from an unmanned airborne vehicle (a.k.a. drone) is usually referred to as ULS. In this paper, we will discuss ALS systems, but the same principles can be applied to MLS and ULS.

ALS scanners forgo the rotating head of TLS scanners and (generally) only have a rotating

or oscillating mirror to direct the laser beam. Because an ALS scanner is always moving, points in the Scanner's Own Coordinate System are unusable; the points must be georeferenced before they can be used. To properly capture the motion of the platform and to locate the point cloud in a global reference frame, ALS platforms include a Global Navigation Satellite System (GNSS) receiver integrated with an Inertial Motion Unit (IMU). GNSS/IMU units provide regular samples of a platform's position, attitude, velocity, and acceleration. Data from the GNSS/IMU are combined with LiDAR range and angle data via the georeferencing equation (Glennie 2007) to produce a point cloud

$$\mathbf{p}_G^l = \mathbf{p}_{GPS}^l + \mathbf{R}_b^l (\mathbf{R}_s^b \mathbf{r}^s - \mathbf{l}^b), \quad (2.3)$$

where:

- \mathbf{p}_G^l are the coordinates of the target point in the global reference frame.
- \mathbf{p}_{GPS}^l are the coordinates of the GNSS sensor in the global reference frame.
- \mathbf{R}_b^l is the rotation matrix from the navigation frame to the global reference frame.
- \mathbf{R}_s^b is the rotation matrix from the scanner's frame to the navigation frame (boresight matrix).
- \mathbf{r}^s is the coordinates of the laser point in the scanner's frame.
- \mathbf{l}^b is the lever-arm offset between the scanner's origin and the navigation's origin.

The boresight matrix \mathbf{R}_s^B and the lever arm \mathbf{l}^b warrant additional explanation. Because the GNSS/IMU and the LiDAR system are separate physical entities, they cannot be mounted in the same place on the aircraft. The positional offset between the two systems, the lever arm offset, is usually measured from engineering drawings of the mounting system or measured by hand (Glennie 2007). The rotational offset, the boresight matrix, is harder to measure. At typical ALS measurement distances, a small rotational offset between the GNSS/IMU and the LiDAR scanner can produce a large difference in the position of the measured point. The boresight matrix is calculated from the data themselves, usually by manually or automatically aligning matching features from multiple data collections (Skaloud and Lichti 2006).

Airborne LiDAR systems can collect large amounts of data over large areas, making them well suited for regional-scale data collection. The point spacing of ALS data is governed primarily by the

height of the scanning platform over the ground, the field of view of the scanning instrument, and the laser pulse rate. In mountainous or otherwise rugged regions, it can be difficult to maintain a consistent above-ground height in a fixed-wing aircraft; helicopters can be especially useful in these environments.

2.4 Error in pulse-based LiDAR scanning

A detailed analysis of the error sources in LiDAR scanning is necessary to assess the usability of LiDAR data as a scientific tool and to develop reasonable expectations for the precision of LiDAR-derived data products. Using Equation 2.3 as our guide, we can examine uncertainty in each component of the georeferencing equation: uncertainties in range, r ; uncertainties in the direction vector of the laser beam; uncertainties in the lever arm, \mathbf{l}^b , and boresight matrix, \mathbf{R}_s^b ; and uncertainties in the position, \mathbf{p}_{GPS}^l , and attitude, \mathbf{R}_b^l , of the scanning platform. This analysis is derived from Glennie (2007).

For pulse-based LiDAR systems, the error in the range measurement r is primarily a function of the resolution of the clock used to measure the time of flight of the laser pulse. Most modern LiDAR scanners sample at a rate of about 1 GHz, which can a range resolution as small as 1 cm (Baltsavias 1999; Mallet and Bretar 2009).

Angular uncertainty is a function of two primary components — the error due to the angular encoders used to convert the mirror (and, in the case of TLS, the scanner head) position to an angle, and uncertainty due to the finite laser beamwidth. The resolution of the angular encoders are usually provided with the manufacturer’s documentation, and tend to range between 0.01° to 0.001° . Beam divergence arises because a laser beam has a finite diameter at scanner exit and then spreads out as it travels through space. The reflective surface observed by the LiDAR scanner can exist anywhere within that finite beamwidth. D. D. D. Lichti and Gordon (2004) assume a uniform power distribution within that beam footprint, and though Glennie (2007) observes that the power distribution is actually roughly Gaussian instead of uniform, both sources agree that the uncertainty due to beam divergence can be conservatively estimated at one quarter of the beam diameter, in angular units.

The lever arm, \mathbf{l}^b , is the xyz offset between the measurement centers of the GNSS/IMU and the LiDAR unit. This value is usually measured in the field or derived from drawings of the mounting apparatus, and the derivative uncertainty is therefore a function of the error in the measurement mechanism. Glennie (2007) provides a rule-of-thumb value of 2 cm for this error.

The rotational offset between the GNSS/IMU and the LiDAR scanner, \mathbf{R}_s^b , is generally impossible to measure with necessary precision via direct measurement, and because pointing errors are magnified at distance, it is crucial to accurately capture these offsets. Usually the boresight adjustment, as it is known, is done by flying over the same features from different directions, and using the offsets between those features to compute a best-fit boresight matrix to bring the planes into alignment (Skaloud and Lichti 2006). Roofs of buildings are particularly useful for these adjustments, which can be done manually or using software automation. Regardless of the alignment mechanism, Glennie (2007) provides rule-of-thumb values of 0.001° for roll and pitch and 0.004° for yaw. Finally, error in the absolute positioning of the GNSS/IMU can usually be obtained as a byproduct of the differential GNSS solution. Typical values are about 2 cm plus 1 parts/million.

These errors can be combined via the General Law Of Propagation of Variance (GLOPOV) to produce a per-point error covariance matrix (Ghilani 2010; Hartzell et al., 2015). Note that this uncertainty does not take into account the effects of the interaction between the laser energy and the reflective surface. Reflections from oblique or poorly reflective surfaces generally contain more uncertainty than a strong orthogonal reflection. Most LiDAR collects do not include per-point uncertainty estimates as a part of their final data products; it is traditional to simply provide some collect-wide accuracy values estimated from ground control points or from the manufacturer's specifications.

3 Pointcloud change detection

Collecting and processing point cloud data, whether from a LiDAR scanner or another source, are just the first steps of full data exploitation. Change detection is a powerful tool that can be applied to multiple epochs of data collected of the same scene. Whether by manual or automated means, point clouds can be compared for translations, rotations, or radiometric changes.

The type and scope of changes than can be detected from point clouds vary with application. In the case of calculating snow depths or other vertical displacements, a simple Digital Elevation Model (DEM) subtraction may be sufficient to capture the change of interest (Deems et al., 2013). For two-dimensional changes, feature tracking algorithms on rasterized point cloud data can provide good estimates of two-dimensional flow (Schwalbe et al., 2008). But subtraction and feature tracking with point cloud-derived DEMs fall short. These methods do not provide estimates of fully three-dimensional motion, which is necessary to characterize many systems, including glacier surface velocities. They also require rasterization of the data to a DEM, which necessarily simplifies the data and potentially discards useful information.

In this section, we explore two commonly-used point cloud change detection methods. The first, Particle Imaging Velocimetry (PIV), is a two-dimensional rasterized change detection method. It is useful for capturing non-rigid change in a LiDAR scene and will provide a useful benchmark to which we can compare our other methods. The second method is Iterative Closest Point (ICP), a widely-used method for calculating rigid three-dimensional change from point clouds. A third change detection method, Coherent Point Drift (CPD), will be discussed in Chapter 4.

3.1 Particle Imaging Velocimetry (PIV)

Particle Imaging Velocimetry (PIV) is a feature-tracking algorithm derived from the field of experimental fluid dynamics. The original PIV method used flash photography to capture the motion of reflective tracer particles in a fluid (Westerweel 1997). We will briefly describe the mechanics of the method, then describe its application to point cloud data.

PIV takes two images, I_1 and I_2 , collected at two distinct epochs, and compares them for

change. Two “windows” are used, a correlation window W_C and an interrogation window W_I , such that $W_I > W_C$. The cross-correlation between the two images within W_C is computed for a range of shifts S_i and S_j

$$r_N(i_s, j_s) = \frac{\sum_{i=0}^{W_C} \sum_{j=0}^{W_C} (I_1(i, j) - \mu_{I_1})(I_2(i + i_s, j + j_s) - \mu_{I_2})}{\sqrt{\sigma_{I_1}} \sqrt{\sigma_{I_2}}}, \quad (3.1)$$

where:

- r_N is the value of the correlation matrix for the given shifts i_s and j_s .
- S_i and S_j are set so that the correlation window visits every possible spot within the interrogation window, and $-S_i \leq i_s \leq S_i$ and $-S_j \leq j_s \leq S_j$.
- μ is the average pixel value for the given image.
- σ is the standard deviation of the pixel values for the given image.

The value of r_N with the largest cross-correlation value is interpreted to be the pixel-based velocity of the pixel at the center of the original correlation window. This process is repeated for each usable pixel in the image, and thereby per-pixel two-dimensional velocities are calculated (Westerweel 1997).

Some optimizations are possible. If the sizes of the two images, I_1 and I_2 , are equal, the cross-correlation can be performed in the frequency domain via the Fast Fourier Transform (FFT) (Westerweel 1997). In order to obtain sub-pixel resolution, a Gaussian fit is used to find the correlation peak (Raffel et al., 2013); when using PIV in this work, we test two different methods for determining the Gaussian fit and use the method that proves more accurate for our data.

When using PIV for point cloud data, it is necessary to rasterize the point cloud into a two-dimensional image (Ferreira 2014). Usually this image is a DEM, where the value of each pixel is a point cloud-derived elevation (Aryal et al., 2012). It is also possible to choose planes other than the horizontal (xy) plane as the plane of rasterization; we only use a horizontal plane.

Aryal et al., (2012) demonstrated the use of PIV for landslide deformation detection at the Cleveland Corral landslide, as well as some empirical estimates of the method’s error. They found 8.6% relative error when comparing PIV results to velocity values obtained via manual feature tracking. We will perform similar comparisons between PIV and on-ice GNSS receivers for one of

our Helheim Glacier data sets.

Aryal et al., (2012) also discuss the usability of PIV for measuring large displacements, where the hyper-accurate differential InSAR method tends to decorrelate. This makes PIV attractive for use in swiftly moving regions or where the time between data collects is large. However, PIV requires rasterization of point cloud data and only works in a two-dimensional plane, limiting its use for accurately capturing three-dimensional change. Rankin and Sridhar (1997) discuss a three-dimensional PIV variant, but to our knowledge this method has not been used for LiDAR data, and we will not explore its use in this work due to the lack of a readily-available software implementation.

3.2 Iterative Closest Point (ICP)

Iterative Closest Point (ICP) is an iterative algorithm for aligning two point clouds. In its simplest form, it calculates a rigid (rotation and translation) transformation that minimizes the point-to-point distance between two point clouds. It was first presented by Chen and Medioni (1991) and Besl and McKay (1992).

Each iteration of ICP consists of two steps, the matching step and the transformation step. In the matching step, each point of the target point cloud \mathbf{P} is matched to the closest point in the source point cloud \mathbf{Q} . If p_i is a point in \mathbf{P} , q_i is the matching point in \mathbf{Q} , and d is the Euclidean distance function, then the matching point satisfies the following relationship

$$d(p_i, q_i) = \min_{j \in N_Q} d(p_i, q_j). \quad (3.2)$$

The matching step has been extended to match points with other geometric entities, such as planes or lines; we will only use the point-to-point version in this work.

ICP is generally used to calculate the optimal rigid transformation, that is a transformation with only rotations and translations. This can be done using quaternions, in the case of two- and three-dimensional ICP, or using Singular Value Decomposition (SVD) (Besl and McKay 1992). Given a rotation matrix \mathbf{R} and a translation \mathbf{t} , successive iterations of ICP minimize the objective function, which is the root-mean-square distance between corresponding points

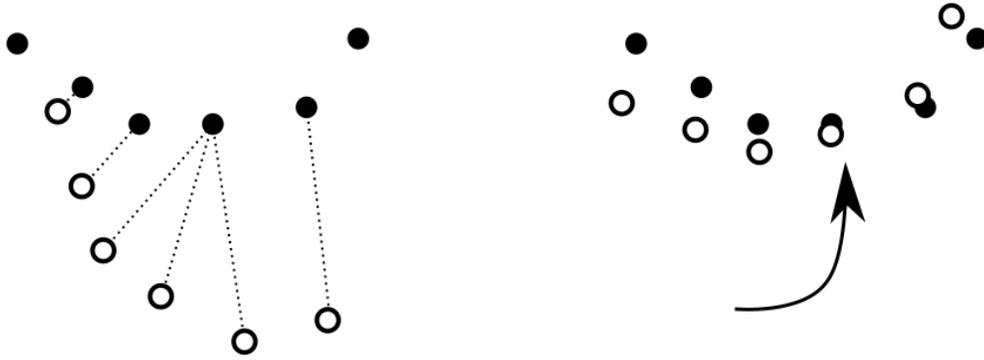


Figure 3.1: Schematic of an Iterative Closest Point iteration, matching step on left, transformation step on right.

$$\sqrt{\sum_{i=1}^{N_p} \|\mathbf{R}p_i + \mathbf{t} - q_i\|^2}. \quad (3.3)$$

This equation is minimized using least-squares. If the change in the objective is less than a predefined threshold, then the algorithm terminates; otherwise, it returns to the matching step and continues the process (Besl and McKay 1992).

ICP has been used with LiDAR data to calculate landslide deformation (Teza et al., 2007) and earthquake displacements and rotations (Nissen et al., 2012). In order to calculate non-rigid deformation over a large area, it is common to divide the datasets into pieces and calculate ICP transformations for each piece; this approach is often called “piecewise” or “segmented” ICP. Zhang et al., (2015) recommend the intelligent use of moving windows when segmenting the point cloud in order to reduce the effects of spatially-varying point cloud parameters, such as heterogeneous data density or shadowing.

ICP only guarantees that it will converge to a local minimum, not a global one (Besl and McKay 1992). A wide variety of variations exist to reduce the likelihood of being trapped in a local minima, or to improve ICP in other ways; the reader is referred to Rusinkiewicz and Levoy (2001) for an overview of some alternatives. One area of particular interest is the use of per-point uncertainty

values to effectively weight the ICP, thereby preferentially weighting the algorithm towards “better” points. Zhang et al., (2015) discuss the use of anisotropic-ICP (A-ICP) with rigorously-calculated per-point uncertainty values to improve the quality of “poor” datasets. This approach is particularly useful when comparing two datasets with divergent collection parameters or accuracy — older, legacy datasets tend to have more uncertainty than newer datasets. In this work, almost all of our TLS data was collected by the same instrument under almost identical collection conditions, and so we will not consider the per-point uncertainty when performing ICP. Per-point uncertainty analysis could be of use for our Canada Glacier dataset, but its application is beyond the scope of this research.

The accuracy of an ICP-derived transformation can be difficult to assess. It is possible to use synthetic displacements to gain a subjective sense of ICP quality, as been done by Nissen et al., (2012) and Zhang et al., (2015). A measure of absolute accuracy is only possible with external control, which we will utilize in the form of in-situ GNSS sensors for our Helheim Glacier data.

ICP is a promising method that uses true three-dimensional point cloud data to calculate three-dimensional change. As it is popular for use for point cloud change detection (Prokop and Panholzer 2009; Ferreira 2014), we will utilize it and compare it to our other methods.

4 Coherent Point Drift

Coherent Point Drift is an application of Gaussian Mixture Models (GMM) and Expectation Maximization (EM) to point cloud change detection. Rather than minimizing point-to-point distances, as in ICP, it maximises the posterior probability of each source point, given a set of Gaussians centered at the target points. It was initially developed by Myronenko and Song (2010) for use in the medical imaging field, and to date has not been used extensively for geographic-scale point cloud data. What follows is a summary of the Coherent Point Drift algorithm for three-dimensional point registration, followed by a brief discussion of software optimizations constructed for this work.

4.1 Overview

CPD is an Expectation Maximization (EM) algorithm that calculates a transformation between two datasets that best aligns them. For each iteration, the algorithm takes the following actions.

1. Calculates correspondence (error) between two datasets, \mathbf{X} and \mathbf{Y}
2. Finds a transformation that will minimize the error between \mathbf{X} and \mathbf{Y}
3. Applies that optimal transformation to \mathbf{Y}
4. If some error bound is met, exit the algorithm, otherwise return to step 1

The CPD algorithm can be tailored for a specific use by picking the appropriate transformation that will be iteratively applied to \mathbf{Y} , e.g. an affine transformation or a non-rigid transformation based in Motion Coherence Theory (Yuille and Grzywacz 1988).

4.2 Explanation

The target dataset \mathbf{Y} is treated as a Gaussian Mixture Model (GMM), and the source dataset \mathbf{X} is treated as points that were generated by the GMM. The probability density function for \mathbf{X} is

$$p(\mathbf{x}) = w \frac{1}{N} + (1 - w) \sum_{m=1}^M \frac{1}{M} \frac{1}{(2\pi\sigma^2)^{\frac{D}{2}}} \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}_m\|^2}{2\sigma^2}\right), \quad (4.1)$$

where

- w is the weight given to an additional uniform distribution to account for noise and outliers, $0 \leq w \leq 1$
- N is the number of points in \mathbf{X} , the source dataset
- M is the number of points in \mathbf{Y} , the target dataset
- D is the dimensionality of \mathbf{X} and \mathbf{Y}
- σ^2 is the bandwidth of the GMM

The GMM centroid locations, i.e., the points of \mathbf{Y} , are then parameterized by θ . The number and type of the parameters in θ will depend on the type of transformation that will be iteratively applied to \mathbf{Y} ; CPD has been implemented for several types of transformations, including affine, rigid, and nonrigid transformations. With the points in \mathbf{Y} parameterized by θ , we can then define an objective function Q , which is an upper bound of the negative log-likelihood function for our system. The objective function is defined as

$$Q(\theta, \sigma^2) = \frac{1}{2\sigma^2} \sum_{n=1}^N \sum_{m=1}^M P^{old}(m|\mathbf{x}_n) \|\mathbf{x}_n - \mathcal{T}(\mathbf{y}_m, \theta)\|^2 + \frac{N_P D}{2} \log \sigma^2, \quad (4.2)$$

where

$$N_P = \sum_{n=1}^N \sum_{m=1}^M P^{old}(m|\mathbf{x}_n) \leq N, \quad (4.3)$$

$$P^{old}(m|\mathbf{x}_n) = \frac{\exp\left(-\frac{1}{2} \left\| \frac{\mathbf{x}_n - \mathcal{T}(\mathbf{y}_m, \theta^{old})}{\sigma^{old}} \right\|^2\right)}{\sum_{k=1}^M \exp\left(-\frac{1}{2} \left\| \frac{\mathbf{x}_n - \mathcal{T}(\mathbf{y}_k, \theta^{old})}{\sigma^{old}} \right\|^2\right) + c}, \quad (4.4)$$

and

$$c = (2\pi\sigma^2)^{\frac{D}{2}} \frac{w}{1-w} \frac{M}{N}. \quad (4.5)$$

The first term of the objective function (Equation 4.2), the double sum, is the misalignment term, providing a numerical estimate of the correspondence between the two data sets given the

transformation Q and bandwidth σ^2 . The second term is a regularization term to help the algorithm from getting stuck in local minima.

Successive iterations of our algorithm minimize Q , which necessarily minimizes the negative log-likelihood for our system; the algorithm exits when some pre-defined error threshold has been met. Specific implementations of Q depend on the type of transformation applied, \mathcal{T} . The original implementation of CPD defines rigid, affine, nonrigid, and nonrigid low-rank transformations. The nonrigid implementations compute a flexible transformation based in Motion Coherence Theory (Yuille and Grzywacz 1988). While glaciers deform nonrigidly, the large number of points in a typical LiDAR glacier point cloud exceeds the capabilities of our implemented CPD algorithm. Because the naive implementation of Equation 4.2 is $\mathcal{O}(MN)$, computational complexity grows exponentially with data size; approximations such as the Fast Gauss Transform (see Section 4.4) can reduce the computational complexity, but large data sizes are as of yet not practical for use with CPD. An analysis of the performance of CPD for large data sets and increasing the upper limit of data size are areas for future work. Therefore we must apply CPD in a segmented or piecewise manner, dividing the data into an irregular grid and calculating the change vector for each piece. We make the assumption that within each piece of the piecewise algorithm, the glacier moves rigidly, and so therefore we only use the rigid CPD algorithm in this work; we make the same approximation for Iterative Closest Point.

Note that CPD is defined for any number of dimensions, D ; we only use the case $D = 3$. It would be possible to use other dimensions of the point clouds, such as laser return intensity, as additional dimensions in the CPD algorithm. Care would need to be taken to weight these additional dimensions correctly, since they will most likely exist in a different unit space than the three spatial dimensions. This would be interesting to investigate in future work.

4.3 Rigid transformation

To perform a rigid transformation, \mathcal{T} is defined as

$$\mathcal{T} = s\mathbf{R}\mathbf{y}_m + \mathbf{t}, \tag{4.6}$$

where:

- \mathbf{R} is a $D \times D$ rotation matrix.
- \mathbf{t} is a $D \times 1$ translation vector.
- s is a scaling parameter.

This presents us with an objective function Q of the following form

$$Q = \frac{1}{2\sigma^2} \sum_{m,n=1}^{M,N} P^{old}(m|\mathbf{x}_n) \|\mathbf{x}_n - s\mathbf{R}\mathbf{y}_m - \mathbf{t}\|^2 + \frac{N_P D}{2} \log \sigma^2. \quad (4.7)$$

The objective function Q must be minimized with respect to \mathbf{R} , s , \mathbf{t} , and σ^2 , and contains one free parameter w , which ranges between zero and one and represents our assumption about the amount of noise present in the data sets. The optimal R is given by

$$\mathbf{R} = \mathbf{UCV}^T \quad \text{and} \quad (4.8)$$

$$\mathbf{USSV}^T = \text{svd}(\hat{\mathbf{X}}^T \mathbf{P}^T \hat{\mathbf{Y}}), \quad (4.9)$$

where \mathbf{P} has elements $p_{mn} = P^{old}(m|\mathbf{x}_n)$, and

$$\mathbf{C} = d(1, \dots, 1, \det(\mathbf{UV}^T)), \quad (4.10)$$

where d is a diagonal matrix, and

$$\hat{\mathbf{X}} = \mathbf{X} - \mathbf{1}\mu_x^T, \quad \hat{\mathbf{Y}} = \mathbf{Y} - \mathbf{1}\mu_y^T, \quad (4.11)$$

where $\mathbf{1}$ is an appropriately-sized identity matrix, and

$$\mu_x = \mathbf{E}(\mathbf{X}) = \frac{1}{N} \mathbf{X}^T \mathbf{P}^T \mathbf{1}, \quad \mu_y = \mathbf{E}(\mathbf{Y}) = \frac{1}{N} \mathbf{Y}^T \mathbf{P} \mathbf{1}. \quad (4.12)$$

The remaining parameters are then found by setting the corresponding partial derivative of Q to zero. For a complete derivation, the reader is referred to Myronenko and Song (2010).

4.4 Fast Gauss Transform

All forms of CPD contain a computational bottleneck at the repeated calculation of the discrete Gauss transform between the source and target datasets (Equation 4.1)

$$G(y_j) = \sum_{i=1}^N q_i \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{y}_i\|^2\right). \quad (4.13)$$

The computational complexity of this operation can be reduced from $\mathcal{O}(MN)$ to $\mathcal{O}(M + N)$ with use of the Fast Gauss Transform algorithm, which produces an approximation of the discrete Gauss transform within a certain error tolerance, ϵ (Greengard and Strain 1991). The Fast Gauss Transform has itself been improved by the Improved Fast Gauss Transform (IFGT) (Raykar et al., 1991), which uses k-center clustering and a different expansion method to speed up computations.

The performance of IFGT is sensitive to parameter selection; it also performs poorly for very small bandwidths, when IFGT requires significantly more polynomial expansion terms to honor the error bound. Since CPD iteratively decreases the bandwidth σ^2 with successive iterations, it is preferred to maintain good discrete Gauss transform performance through a wide range of bandwidths. At a certain bandwidth cutoff point (later in a CPD run), we switch from using the IFGT to using the direct Gauss transform that only considers points that are close to the source point, under the assumption that for small bandwidths, far away points will have a negligible impact on the overall result of the Gauss transform (Morariu et al., 2009).

4.5 Software implementation

The Coherent Point Drift algorithm was implemented as a GPL2 (“GNU General Public License, version 2” 1991) C++ library¹, based on a Matlab implementation² by Andriy Myronenko. The

¹<https://github.com/gadomski/cpd>

²<https://sites.google.com/site/myronenko/research/cpd>

library provides a modern C++ Application Programming Interface (API), with documentation, for running CPD from external software. CPD was integrated to the BSD (“The BSD 3-Clause License”) library PDAL (“PDAL - Point Data Abstraction Library”) as a native plugin, which allows usage of CPD on data contained in many point cloud storage formats, including the commonly used `las` format (ASPRS 2013).

The truncated direct Gauss transform and Improved Fast Gauss Transform were implemented as an LGPL (“GNU Lesser General Public License” 2007) C++ library³, based on the `figtree`⁴ library by Vlad Morariu. Our library leverages OpenMP (OpenMP 2013) for parallelization.

³<https://github.com/gadomski/fgt>

⁴<https://github.com/vmorariu/figtree>

5 Test data and experiments

In order to assess the change detection methods with a variety of source data, we will use LiDAR data collected from two sites. The first site is the Helheim Glacier in Greenland. It was surveyed in the summer of 2014 with a terrestrial scanner. The second site is the Canada Glacier in Antarctica and was surveyed in 2001 and in 2014 with airborne scanners. These two sites provide a variety of data densities, collection parameters, and shadowing effects to sufficiently exercise our change detection methods.

5.1 Helheim Glacier, Greenland

The Helheim Glacier is a tidewater glacier in southeast Greenland (66.38 N, 38.8 W) draining the Greenland Ice Sheet (GIS) (Figure 5.1). Because of its size, accessibility, velocity of more than 20 m d^{-1} (Bevan et al., 2012), and high rate of iceberg discharge (Andresen et al., 2012), the Helheim Glacier has been the subject of many studies, including on-ice GNSS velocity measurements (Nettles et al., 2008; Juan et al., 2010) and remote sensing observations of mass-balance and velocity (Bevan et al., 2012; M. L. Andersen et al., 2010; Moon et al., 2014).

Terrestrial LiDAR data was collected in the summer of 2014 from July 10th through July 14th. Instrument parameters are presented in Table 5.1. The scanner was mounted on a tripod on the south side of the glacier (Figure 5.2). Scans were nominally collected at thirty-minute intervals, but due to generator issues there are several gaps in data collection. During a portion of the time the scanner was active, two GNSS receivers were placed on the glacier surface, designated **hg02** and **hg03**. These receivers collected position information at fifteen-second intervals (with some gaps), and were recovered at the end of data collection. The temporal coverage of scans and GNSS position data are presented in Figure 5.3.

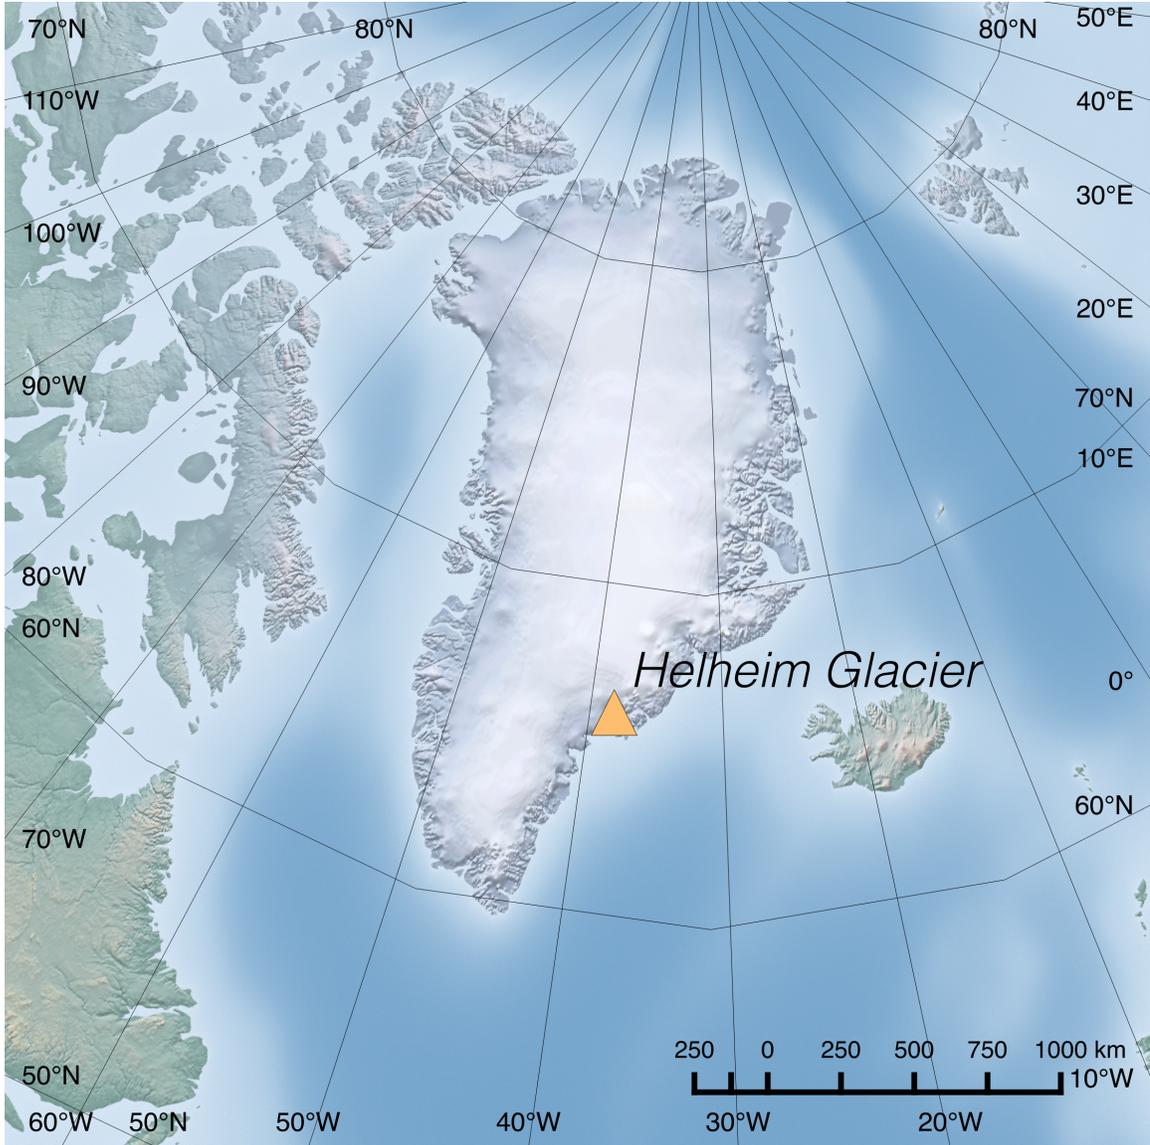


Figure 5.1: Helheim Glacier, Greenland.

Table 5.1: Terrestrial LiDAR scanner parameters used at the Helheim Glacier in Greenland in July 2014.

Manufacturer	RIEGL Laser Measurement Systems
Model	VZ-6000
Laser wavelength	1064 nm
Pulse rate	50 kHz
Max range (approximate)	6 km
Points per scan (approximate)	50 000 000
Time per scan (approximate)	30 min
ϕ range	135° to 234°
Angular velocity of scanner head	3.3 ° min ⁻¹
Beam divergence, γ	0.12 mrad ¹
Total number of scans	92

Each LiDAR scan was processed for Multiple Time Around (MTA) returns and manually cleaned of noisy data. In order to georeference the TLS scans, the 2014 scan position was registered against similar scans taken in 2013 at the same site using Riegl’s Multi-Station Adjustment (MSA) tool (Ullrich et al., 2003). The 2013 scans themselves were georeferenced into WGS84 UTM 24N using cylindrical reflectors in the scan scene with positions surveyed via GNSS.

The availability of coincident in-scene GNSS data makes this Helheim dataset particularly valuable. We will use the two GNSS records to assess the absolute accuracy of each of our three change detection methods (PIV, ICP, and CPD). We will also do intra-method comparisons, assessing the effects of a variety of parameters on change detection performance, including data density and magnitude of change (i.e., the amount of time between two scans). For all measurements, we will use measured surface ice velocity as our primary metric.

5.1.1 Glacier surface velocities and artificial scene compression

Velocity v has a simple relationship to distance d and time t

¹0.12 mrad corresponds to 12 mm increase in laser beam width at 100 m range.

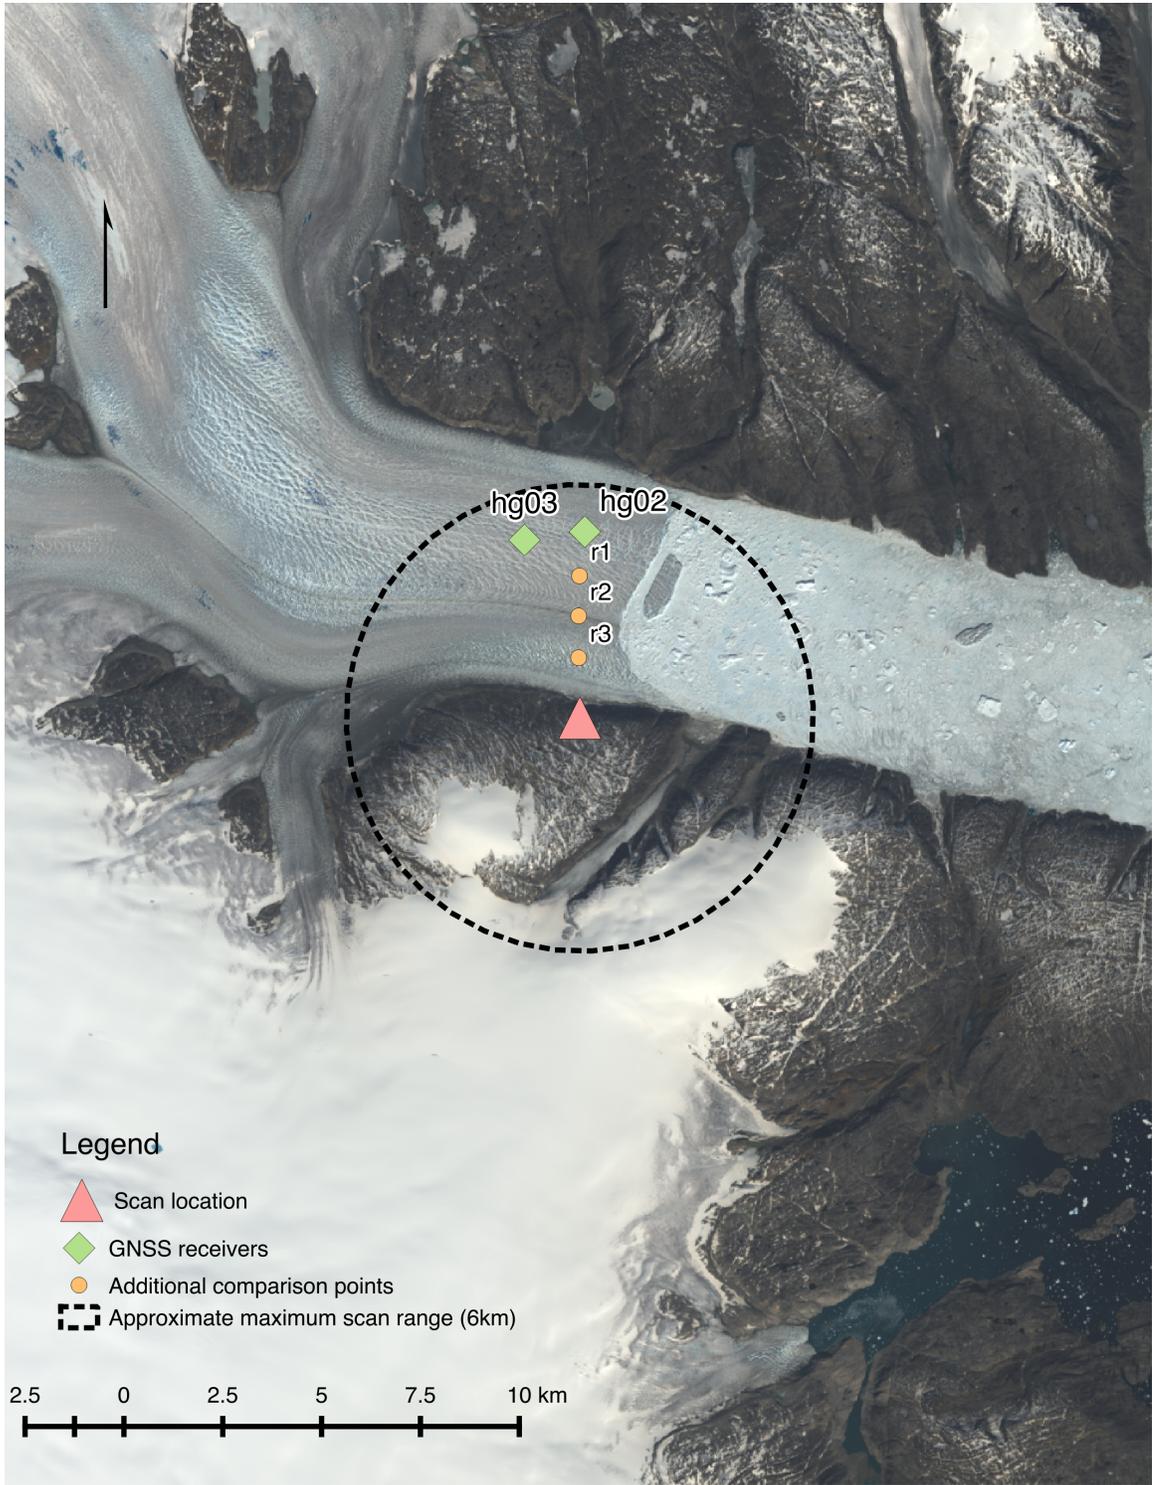


Figure 5.2: Location of scanner and GNSS receivers during summer 2014 collects at the Helheim Glacier. Also labelled are three additional points used for relative velocity comparisons.

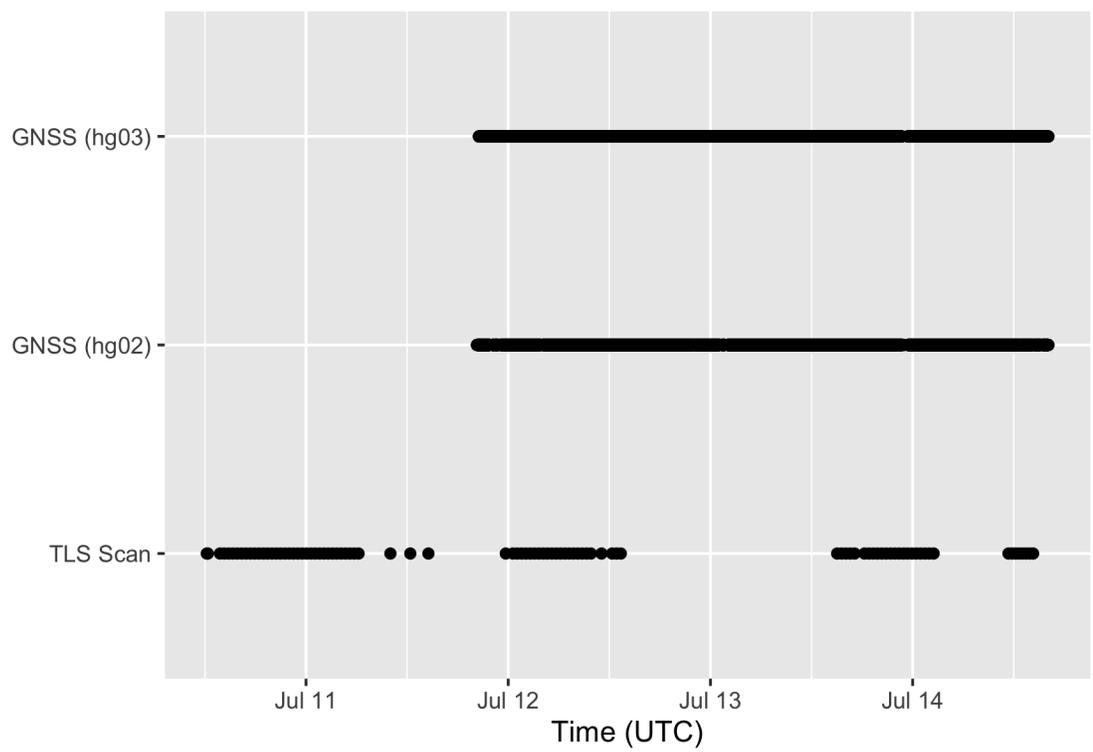


Figure 5.3: Timing of TLS scans and GNSS receiver location records for the summer 2014 collect at the Helheim Glacier.

$$v = \frac{d}{t}. \quad (5.1)$$

Change detection methods retrieve a distance d displacement between two point clouds. In order to find a v , t must be known. Because a terrestrial LiDAR scan takes a finite amount of time, about 30 min in the case of our Helheim Glacier collect, each point of the glacier is measured at some positive time offset from the start of that scan (see Figure 5.4 for a visualization of these time offsets with real data). Because the glacier is constantly moving, this leads to either an extensional or compressional effect in the overall scan data, depending on the direction of scanner head rotation. For our 2014 collect, all scans were collected with a counterclockwise scanner head rotation, leading to an overall compressional effect. If we assume a average glacier velocity of 1 m h^{-1} , the glacier will appear roughly 0.5 m short when measured by a 30 min scan.

While this compressional effect matters for measuring glacier-wide distances, it can be safely ignored for smaller interrogation windows where the magnitude of the artificial compression is less than the accuracy of the measurement instrument. For instance, if we again assume 1 m h^{-1} glacier surface velocity, a region of the glacier that takes 1.2 min to scan will compress roughly 2 cm, which happens to be an optimistic estimate for the minimum error in measurable distance by our LiDAR scanner (Riegl LMS GmbH 2012). The relationship between range of target and maximum interrogation window width at that range can be approximated by the following relationship

$$w = r \tan \frac{\omega \sigma}{v}, \quad (5.2)$$

where w is the maximum interrogation window width, r is the range of the target, ω is the angular velocity of the scanner head, σ is the maximum measurable distance of the LiDAR scanner, and v is the surface velocity of the glacier in the target area. In order to ignore the effects of artificial scan compression in this work, we keep our interrogation window widths less than w using conservative values of 2 cm for σ and 1 m min^{-1} for v . At the range of our GNSS sensors, about 5 km, our maximum interrogation window width w must be less than 350 m if we assume σ of 2 cm. However, that σ value is very optimistic for a 5 km range; we could reasonably use larger interrogation windows if we had a better sense of the true accuracy of the instrument at a given range. For this simplified

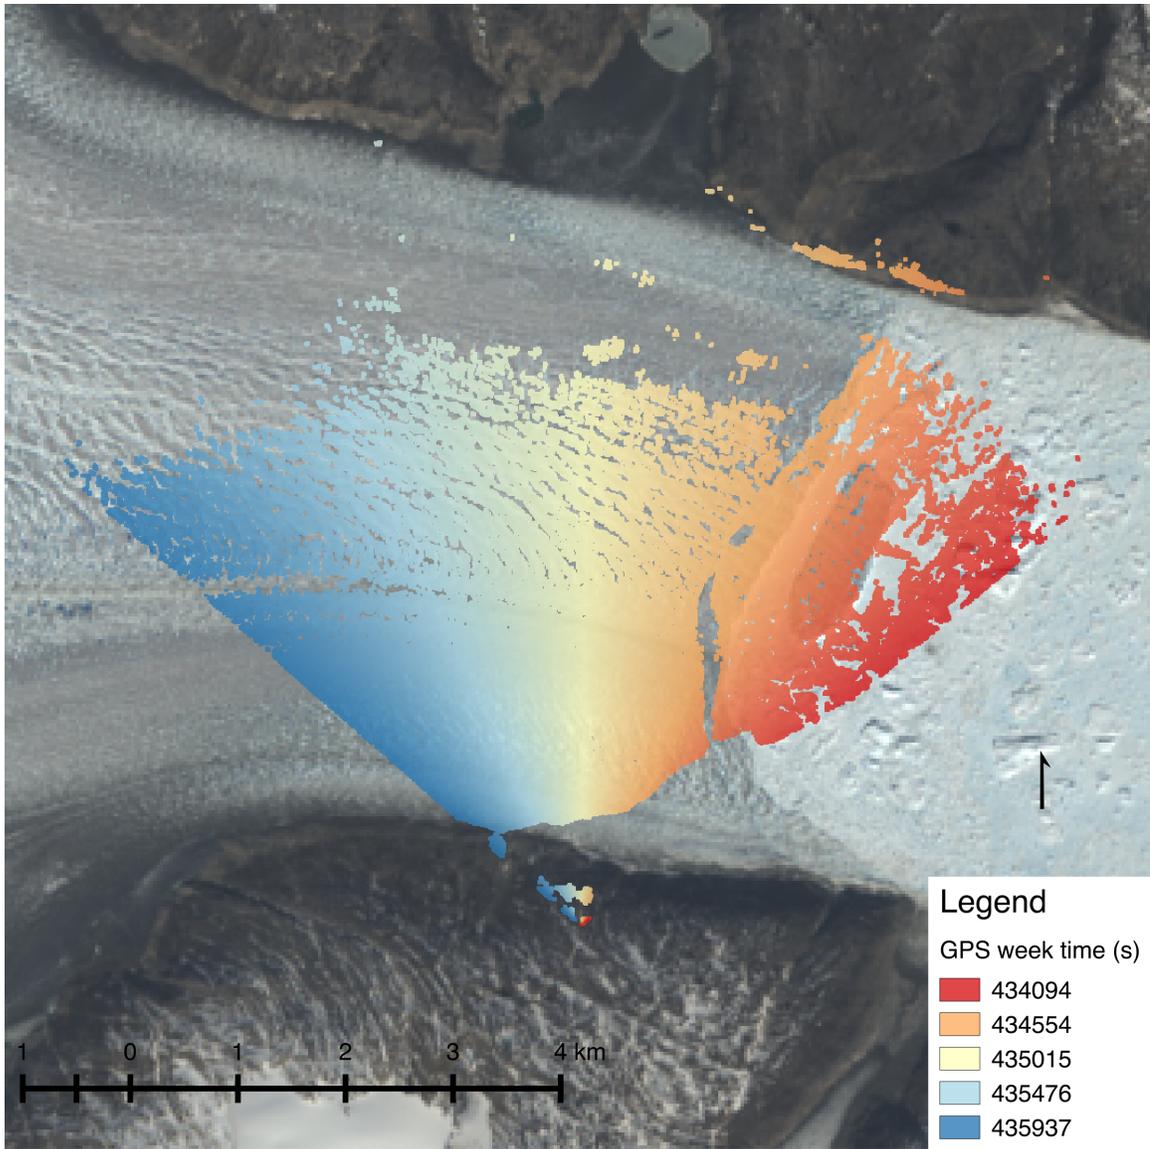


Figure 5.4: Time offsets within a single scan of the Helheim Glacier. Scanner head rotation was counterclockwise. Time values are GPS week time.

calculation, we can calculate the contribution of the laser’s finite beamwidth to angular uncertainty, and subsequently calculate the corresponding positional uncertainty. If this positional uncertainty is greater than the manufacturer’s quoted range accuracy, we can use the positional uncertainty due to beamwidth when calculating our maximum interrogation window width.

The two-dimensional power distribution of a LiDAR laser pulse can be approximated by a Gaussian (Glennie 2007). The divergence of this laser beam, γ , is usually provided as part of the manufacturer’s specifications. Per Hartzell et al., (2015), if the beam diameter is defined from the $\frac{1}{\exp^2}$ points in the Gaussian power profile (as it is for the VZ-6000 used to collect the data), the angular variance can be defined as

$$\sigma_{BW_{angle}}^2 = \left(\frac{\gamma}{4}\right)^2. \quad (5.3)$$

For our γ of 0.12mrad, this produces a $\sigma_{BW_{angle}}^2$ of 0.0009 rad. At a 5km of range this corresponds to a positional uncertainty of 5 cm due *only* to angular uncertainty. We can comfortably use this value as our minimum allowable artificial compression, increasing our maximum allowable interrogation window width to 880 m. A more rigorous analysis of per-point uncertainty at various ranges could produce more precise values for maximum allowable interrogation window width, but we contend that the approximation presented above is a conservative estimate and therefore reasonable.

5.1.2 Absolute accuracy of change detection methods

In order to assess the absolute accuracy of each change detection method, we selected the 54 scans that were taken during the collect time of the hg02 and hg03 GNSS receivers. These scans were organized into pairs based upon set time intervals between scans, such as thirty minutes or one hour — a variety of intervals were tested to examine the effect of total displacement on each change detection method. Each set of pairs was then processed with each of the three change detection methods. A slightly different procedure was used to compare the scans with PIV, which operates on a rasterized simplification of the scan data, and ICP and CPD, which operate on the point clouds themselves.

For PIV, we compare the effect of changing the size of the pixels and other parameters to

find the configuration that produces the most accurate result when compared to our on-ice GNSS data at the Helheim Glacier. A raster size of 25 m^2 square pixels using inverse-distance weighted rasterization algorithm was found to have the lowest RMSE when compared to the GNSS data; all rasterization was done horizontally. Each pair of rasters was then analyzed with PIVLab, an open-source Matlab tool for running PIV (Thielicke and Stamhuis 2014). We experimented with a variety of parameters, presented in the results section, and ended up using a window size of 16 pixels and three point Gaussian subpixel matching, as these parameters produced the least error when compared to our known velocities. From the resultant displacement (XY change) raster, the value of the pixel containing each GNSS receiver’s location at time of laser incidence was extracted, and these values are reported as the displacement for that scan pair at that GNSS receiver location.

For ICP and CPD, a region equal to or smaller than the maximum allowable interrogation window size around the GNSS receiver was cropped from each source point cloud, and these regions were run through the respective change detection algorithm. The ICP method used was provided by the Point Cloud Library (PCL) (Rusu and Cousins 2011). We tested a variety of ICP parameters and present the results, then picked the best set as our tuned parameters. The CPD method was developed by the author, based upon work by Myronenko and Song (2010). The same tuning procedure is done with CPD, where a variety of outlier weights and window sizes are tested and presented as results. The average displacement of the region of interest was reported as the detected change.

Once a two- (in the case of PIV) or three-dimensional (in the case of ICP and CPD) displacement vector has been calculated, the vector can be converted to a velocity vector using the time difference between the two scans. For each scan, the location of hg02 and hg03 was found, and the timestamp of the closest point was extracted and used as the timestamp for that region’s change vector. The difference in the timestamps between two scans was combined with the calculated distance change to produce a velocity measurement.

To compute a corresponding velocity measurement from the GNSS data, a one-minute window of data was selected around the scan’s sampling timestamp. These positional data were averaged to determine an average position of the receiver for that epoch, and these positions and times were used to calculate a GNSS velocity measurement for each scan pair.

5.1.3 Relative accuracy of change detection methods

Both hg02 and hg03 were a long distance (~ 5 km) from the scanner. At that range, point density and point accuracy is relatively low compared to regions closer to the scanner. The performance of each change detection algorithm may improve or degrade with increased point density. In addition, the glacial surface is undulating, leading to a variety of incidence angles between the laser beam and the mean glacier surface, creating a variety of shadowing regimes. While we cannot compute absolute accuracies for each change detection method in regions where there is no absolute control, we can compare the change detection methods to each other and our prior knowledge and understanding of the Helheim Glacier’s velocity field. We use two methods to compare the relative accuracy and precision of our three change detection methods: point-based comparison and full-glacier qualitative analysis.

For our point-based comparisons, we choose a set of additional points on the glacier at which to compare our change-detection-derived velocities (Table 5.2). These points are chosen for their variety in average glacier surface velocities, surface roughnesses, angle between dominant serac trend and laser beam incidence angle, and point density. The velocity timeseries are compared and analysed for spread/noise. Only three points were chosen because these points lack validation and are therefore only useful for human-based qualitative analysis; if validation or statistical analysis methods were to be employed, additional points could be chosen.

Table 5.2: Three additional points for relative velocity comparisons at the Helheim Glacier.

Name	Easting	Northing	Distance from scanner
r1	537180	7360710	3758 m
r2	537158	7359697	2747 m
r3	537158	7358640	1700 m

For full-glacier surface velocities, it is necessary to break the glacier into segments. PIV naturally provides full-scan velocity information with no additional segmentation beyond rasterization, albeit in two dimensions only. For ICP and CPD, we segment full-scan data into chunks of roughly equal *point counts*. We choose to segment by point count rather than area for two reasons. First, both

algorithms, but CPD in particular, are very sensitive to large point counts. Our TLS scan data has much higher point densities close to the scanner, meaning that equal-area segments would have widely varying point counts, causing undesirable algorithm behavior in either the far or near fields. Second, as discussed in Section 5.1.1, larger interrogation regions are susceptible to artificial scene compression due to the finite scan length and constantly moving glacier. By keeping near-field regions smaller, we reduce the total scan time of those regions, keeping them below our maximum allowable interrogation window width as defined in Equation 5.2.

The full-glacier change detection results are presented visually for qualitative inspection and gross error identification. By visual inspection, one can see where each change detection method performs well and poorly. Histograms of velocities measurements are also presented to give a sense of velocity distribution over the entire glacier. Differences in these histograms can indicate the types of errors likely for a given change detection method.

5.2 Canada Glacier, Antarctica

As a second test case, we use two airborne collects of the Canada Glacier in the McMurdo Dry Valleys region of Antarctica (Figure 5.5). The first collect was performed by NASA in 2001 using the Airborne Laser Terrain Mapper (ALTM) (Table 5.3) (Kennett and Eiken 1997), and the second was performed by the National Center for Airborne Laser Mapping (NCALM) in 2014 using the Optech Titan (Table 5.4) (Rees 2015). Both collects were airborne collects.

Table 5.3: Parameters for the Airborne Laser Terrain Mapper (ALTM), used for 2001 Canada Glacier collect by NASA.

Manufacturer	Teledyne Optech
Model	Airborne Laser Terrain Mapping System
Wavelength	1047 nm
Beam divergence	Not specified
Maximum operating altitude	1800 m AGL

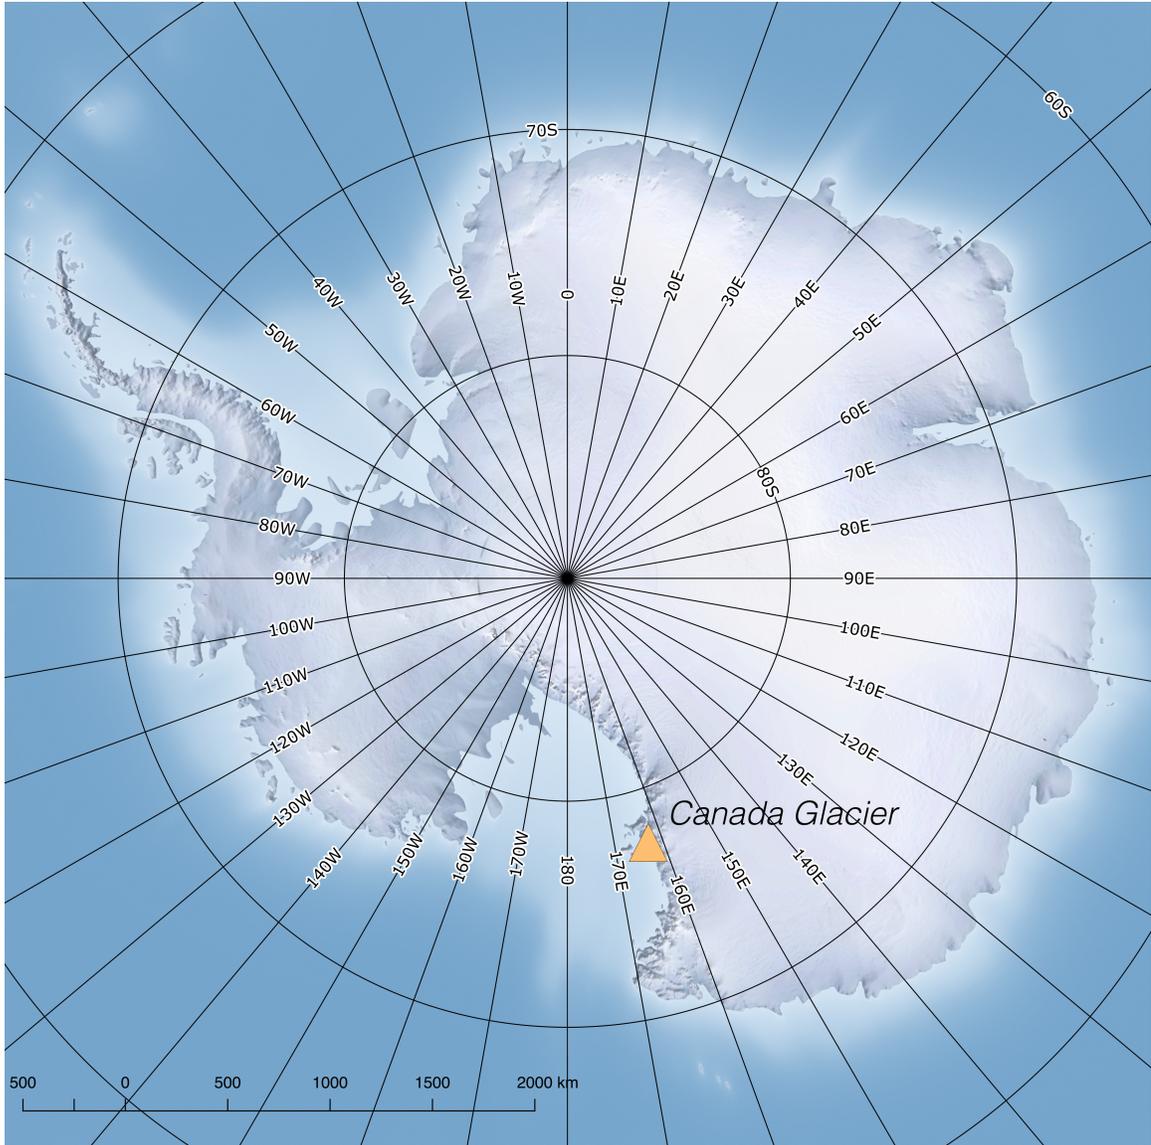


Figure 5.5: Canada Glacier in the McMurdo Dry Valleys region of Antarctica.

Table 5.4: Parameters for Optech Titan LiDAR scanner, used for 2014 Canada Glacier collect by NCALM.

Manufacturer	Teledyne Optech
Model	Titan
Wavelengths	532 nm, 1064 nm, 1550 nm
Beam divergence	0.35 mrad for 1064 and 1550, 0.7 mrad for 532
Maximum operating altitude	2000 m AGL

The data from the Canada Glacier is very different from the TLS Helheim collects. Data densities are relatively constant within each scan, but vary widely between the two scans. The 2001 NASA collect has an average data density of about 0.2 points/m², whereas the 2014 NCALM collect has an average data density of about 1.2 points/m². The data were also collected from roughly nadir-looking platforms, meaning that serac shadowing is greatly reduced.

Because we do not have any absolute control for the Canada Glacier collects, we must rely on relative comparisons between the change detection methods as well as a-priori measurements of glacier velocity. Fountain et al., (2006) observed velocities between 0.3 m/year to 9.0 m/year in the Taylor Valley glaciers generally, and velocities close to the maximum value at the Canada Glacier in specific. We expect our derived velocities to be within this range. As with the Helheim Glacier collect, we qualitatively compare full-glacier change detection maps from each method. We use the same parameters for ICP and CPD that we used on the Helheim Glacier, as we have no ground-truth data to assess the accuracy of any new parameters. For PIV, we used a 2 m rasterization for PIV.

6 Results

We first run a set of trials to determine the optimum parameters for each change detection method. We then examine change detection results for the Helheim Glacier, and analyze these results for absolute accuracy (as compared to GNSS-derived velocities) and relative accuracy (inter-method comparisons). We then do relative analysis for change detection results from the Canada Glacier, and compare these results to known velocities from the literature. We draw conclusions about the usefulness of each change detection method and the factors that influence each method’s performance.

6.1 Synthetic displacements

In order to verify that CPD can extract displacements from these data, we applied synthetic displacements and noise to a subset of the Helheim point cloud data in order to test whether CPD could retrieve those synthetic displacements. At both GNSS sites hg02 and hg03 at the Helheim Glacier, we cropped a 500 m square area, then conducted a number of randomized trials. For each trial, the point cloud was translated a random distance in each dimension, where the distance translated was randomly selected from a Gaussian distribution with a mean of zero and a standard deviation of 10 m. Each point in the translated point cloud was then shifted by an additional random amount, as determined by another Gaussian distribution with a mean of zero and a standard deviation between 0.1 m to 1.0 m — trials were conducted using discrete standard deviations inside that range. CPD was then run on the untranslated and the translated pointclouds, and the average CPD-calculated translation was compared to the applied translation to determine a residual. Trials were at both hg02 and hg03 with noise standard deviation values of 0.1 m, 0.25 m, 0.5 m, 1.0 m; 25 trials were conducted with each set of parameters. Results are shown in Figure 6.1 and Table 6.1. These results show that CPD is able to recover translations for these data even in the presence of noise levels exceeding that what which we would expect from a terrestrial LiDAR sensor at that range.

Table 6.1: Root mean squared error for CPD-calculated translations, as compared to applied translations, at various noise levels.

Noise σ (m)	RMSE (m)
0.1	0.002
0.25	0.006
0.5	0.012
1.0	0.025

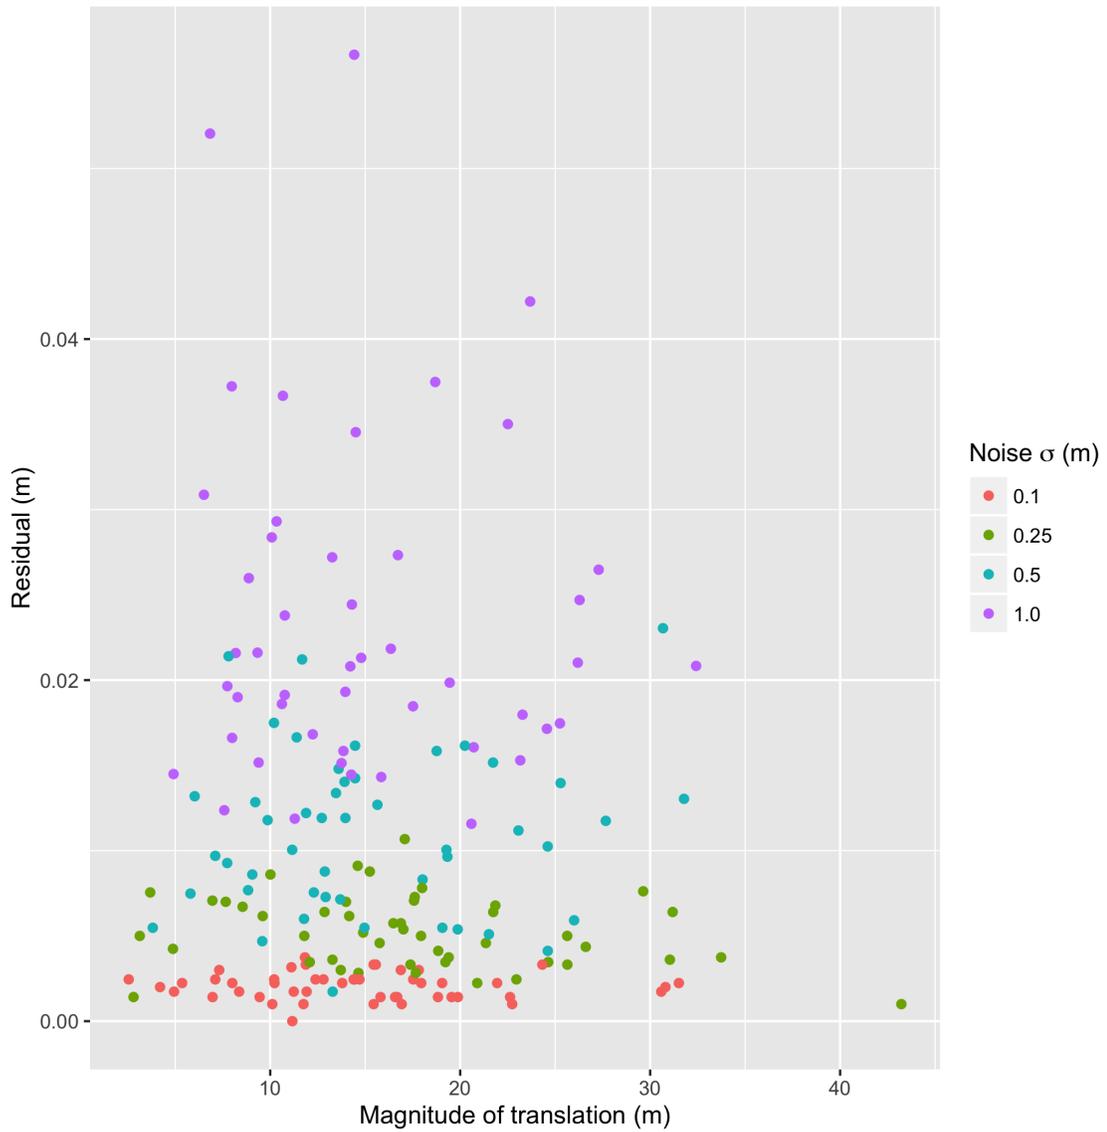


Figure 6.1: Synthetic displacements (x -axis, in meters) and the residuals of CPD-calculated displacements (y -axis, in meters) at hg02 and h03 from scan data collected 2014-07-11 00:33 UTC at the Helheim Glacier, in the presence of various levels of noise.

6.2 Time delta between scan pairs

Using data from the Helheim Glacier in the WGS 84 / UTM 24N coordinate system, we examine the effect of absolute displacement on change detection results. If the time interval between two scans is larger, the glacier will have moved a greater distance. As seen in Figure 6.2 and Figure 6.3, running CPD on two scans with a smaller time gap, e.g. 30 min, produces noisier results than the same calculations with a larger time gap, e.g. 240 min. By choosing larger intervals, we increase the signal-to-noise ratio — for a given amount of noise, a larger time period produces a larger displacement signal to detect. For subsequent results we have chosen a 240 min interval between scans; we do not choose a longer interval because we hope to detect the known twice daily tidal signal in the glacier velocity (Juan et al., 2010). Unfortunately, we are not able to detect the tidal signal via this analysis.

6.3 Tuning CPD

We find the optimal size of the cropped area around each point of interest (GNSS receivers) by testing a variety of crop widths and comparing the Root-Mean Squared Error (RMSE) for each width. The crop width with lowest RMSE is used as the “best” crop width. Results are presented in Figures 6.4 and 6.5, and Tables 6.2 and 6.3. The RMSE drops slightly with larger crop radii at both GNSS locations, and so we will use a 500 m crop radius for point tests of CPD. We did not test crop regions larger than 500 m in an effort to avoid artificial data compression due to glacier motion during scans.

Table 6.2: Statistics for CPD-derived velocities compared to GNSS-derived velocities at hg02 separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1} .

Crop radius	min	mean	max	σ	RMSE	Avg. number of points
100 m	14.96	20.02	25.31	2.32	2.65	350
250 m	15.40	20.64	26.08	2.47	2.51	1920
500 m	15.87	20.70	26.34	2.47	2.51	4348

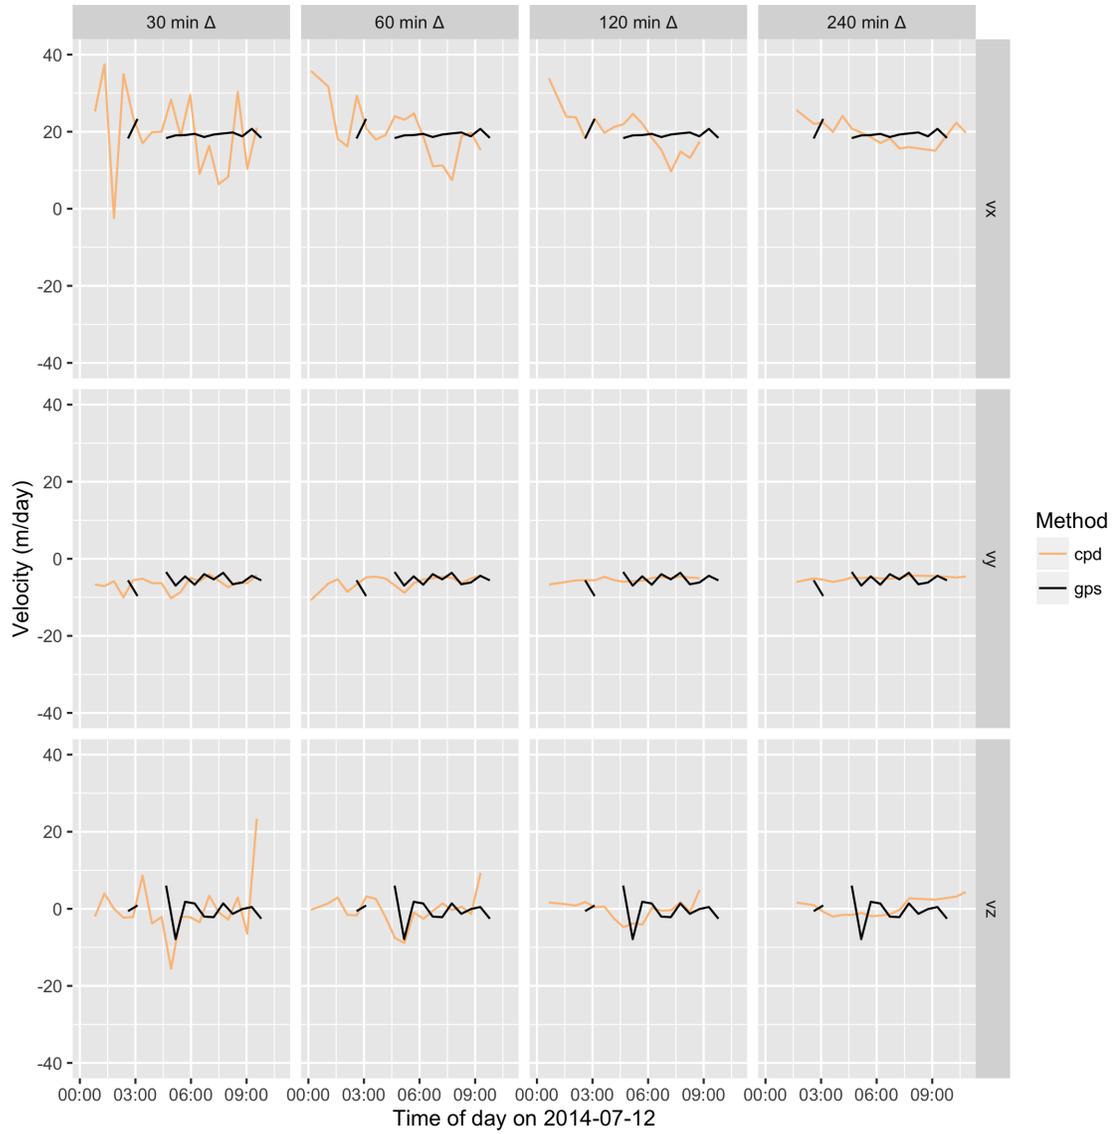


Figure 6.2: CPD-calculated glacier velocities at hg02, divided vertically by velocity component (x-velocity on top, y-velocity in the middle, z-velocity on the bottom) and horizontally by nominal time gap between scans in minutes.

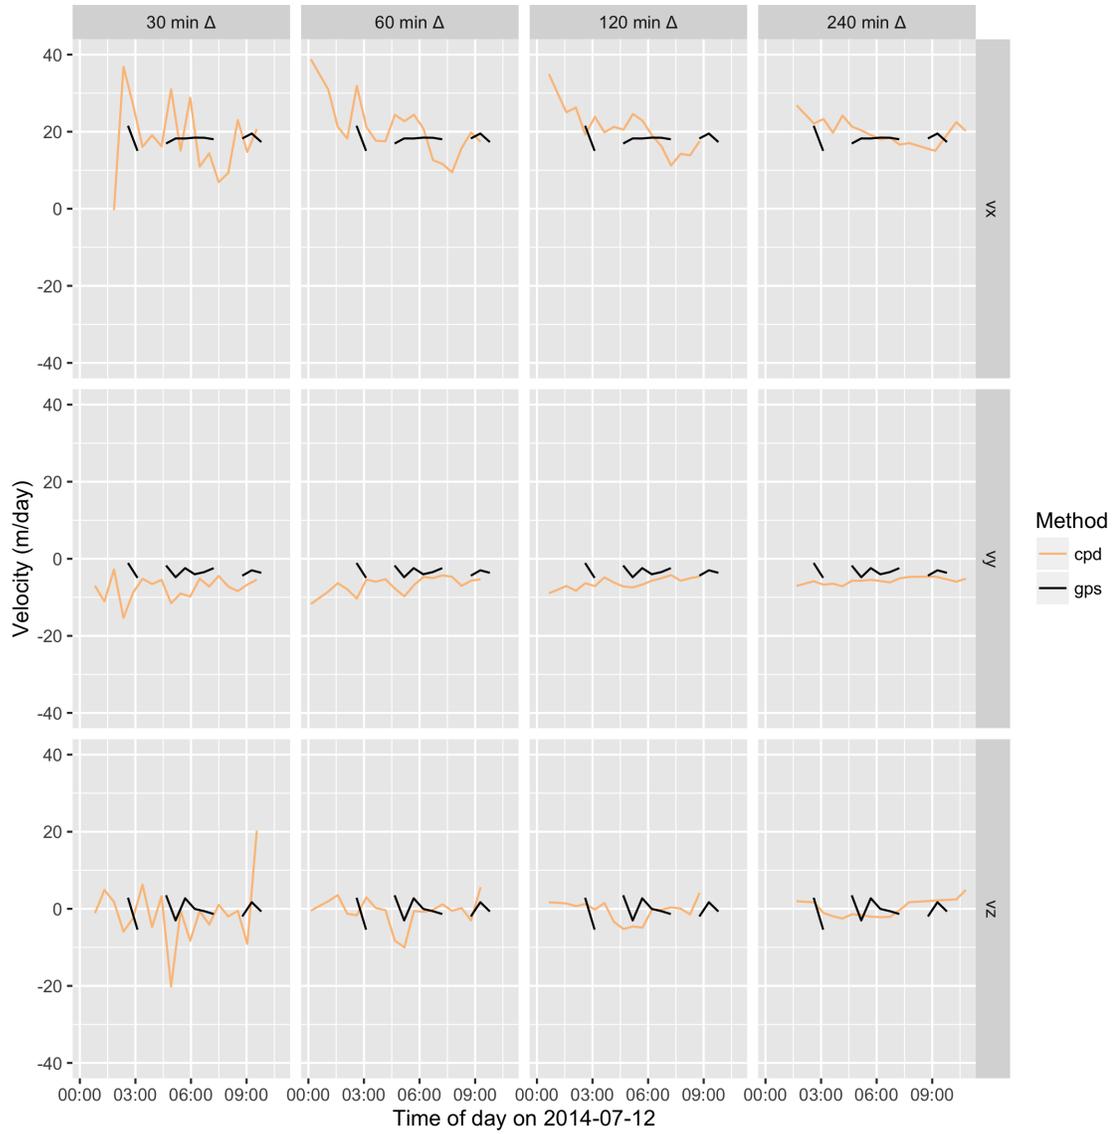


Figure 6.3: CPD-calculated glacier velocities at hg03, divided vertically by velocity component (x-velocity on top, y-velocity in the middle, z-velocity on the bottom) and horizontally by nominal time gap between scans in minutes.

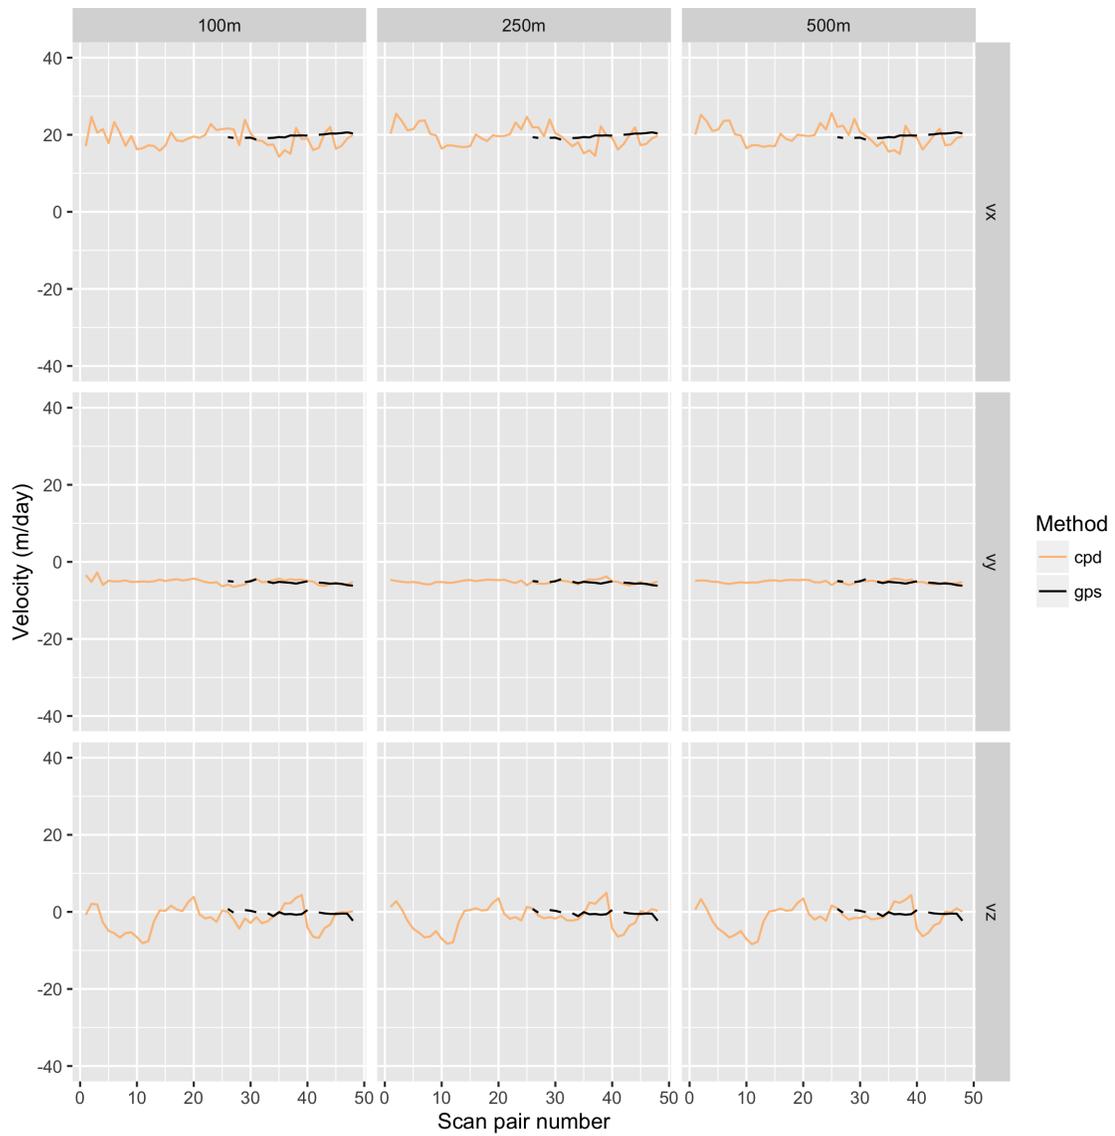


Figure 6.4: CPD-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg02 for crop widths (horizontal divides) of 100 m, 250 m, and 500 m. Because of data gaps due to intervals of generator failure during scan collection, x-axis is scan pair ID, not time of collection.

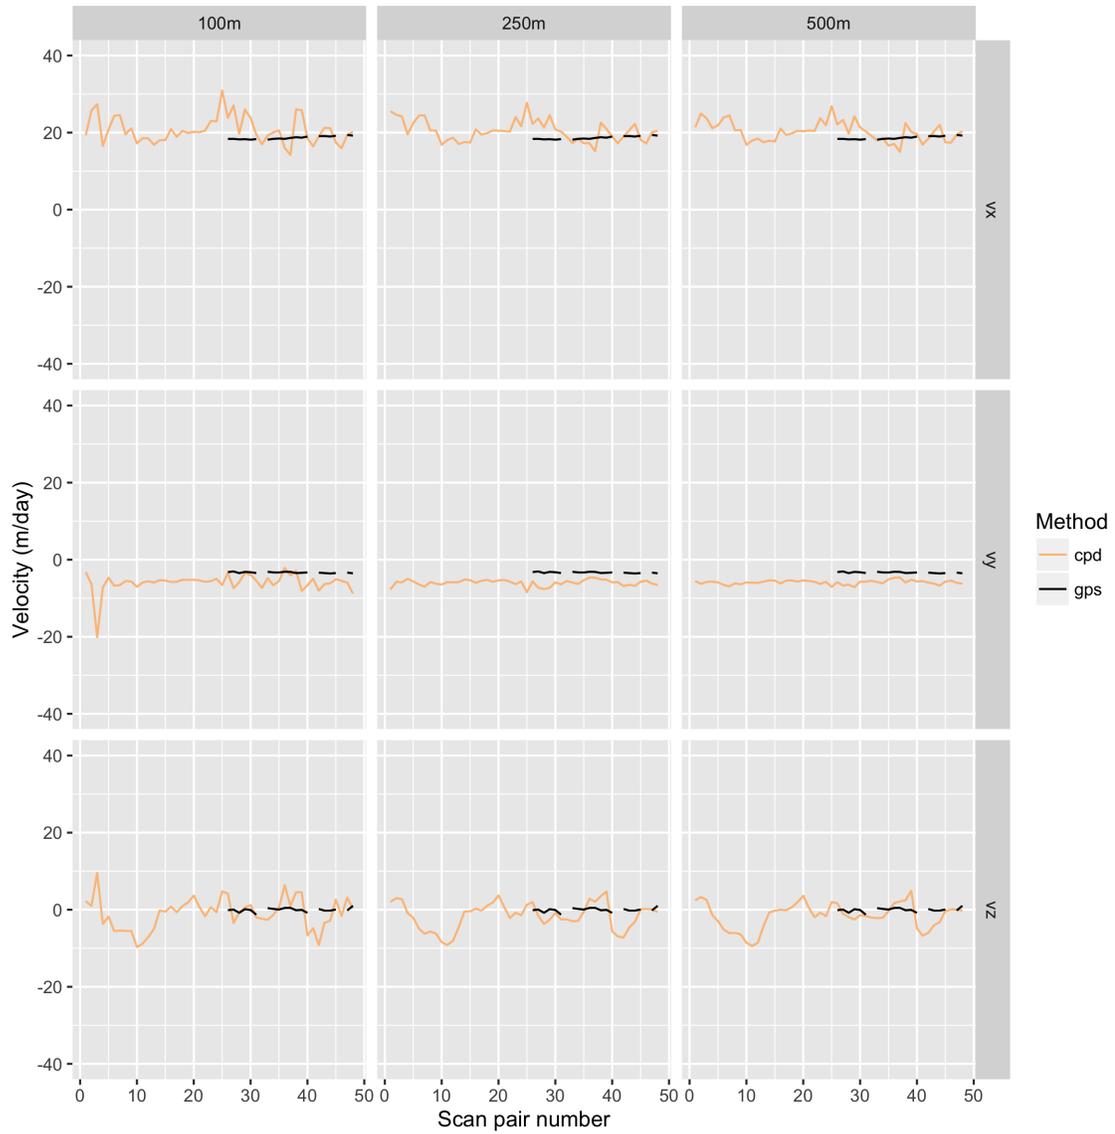


Figure 6.5: CPD-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg03 for crop widths (horizontal divides) of 100 m, 250 m, and 500 m.

Table 6.3: Statistics for CPD-derived velocities compared to GNSS-derived velocities at hg03 separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1} .

Crop radius	min	mean	max	σ	RMSE	Avg. number of points
100 m	14.80	22.00	35.24	3.77	4.59	104
250 m	16.02	21.67	28.98	2.68	3.31	714
500 m	15.85	21.44	27.82	2.47	3.04	3945

The one free parameter in rigid CPD is an outlier weight, a value between zero and one that roughly represents the amount of noise in the dataset. We tested a variety of outlier values and examine their effect on our results by computing the RMSE between CPD results and the GNSS-derived velocities. All outlier values between 0.1 and 0.8 produces nearly identical output (Figure 6.6), and so we used the author-recommended value of 0.1 (Myronenko and Song 2010).

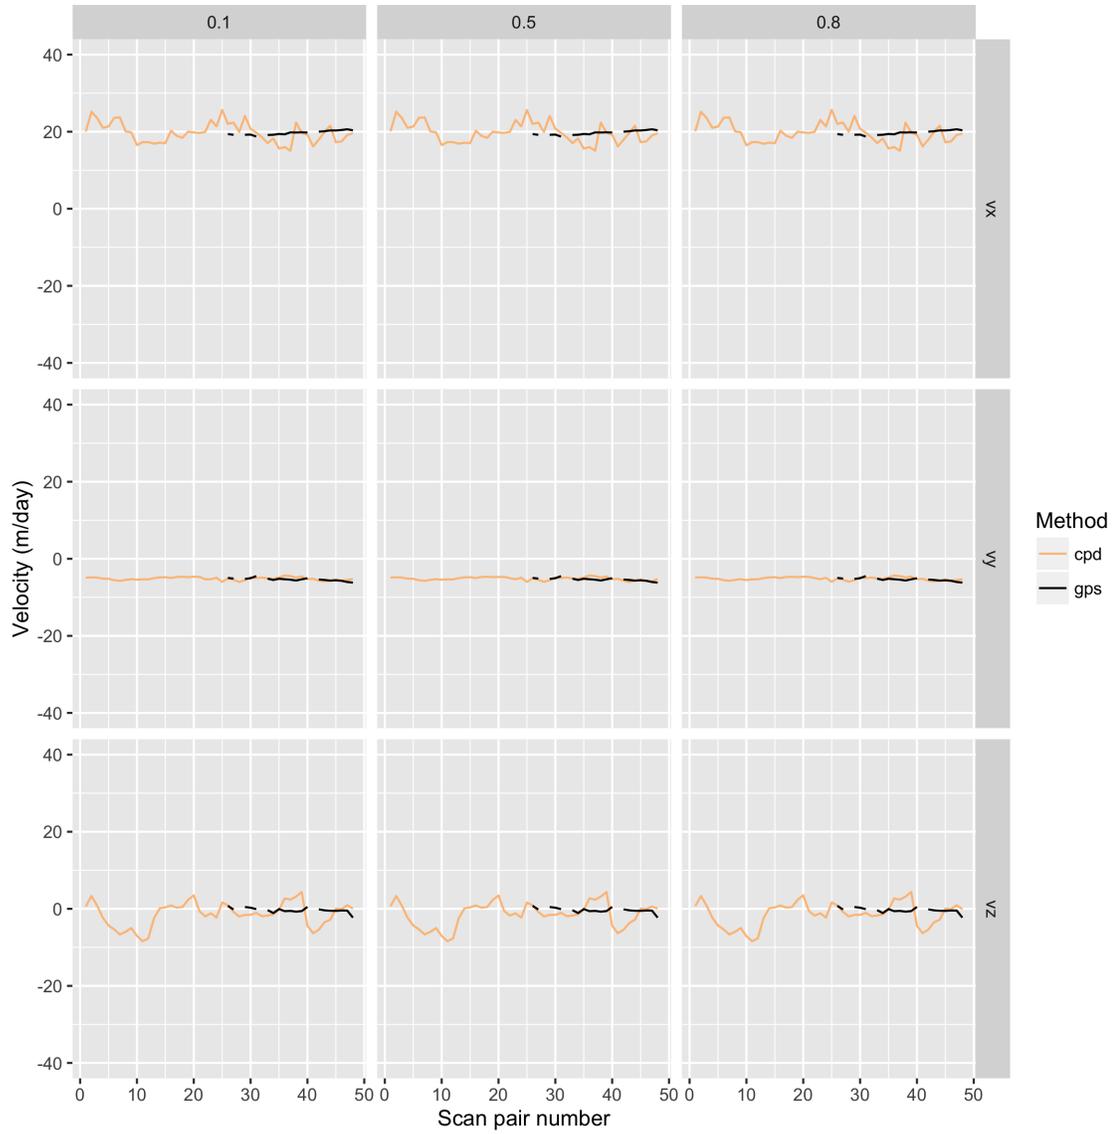


Figure 6.6: CPD-derived velocities at hg02 for three different outlier weights. The outlier weight is the single free parameter, between zero and one, in CPD. There is not discernible different in results in this range of outlier weights.

6.4 Tuning ICP

We applied the same test from CPD to ICP, namely a comparison of the RMSE between ICP and GNSS velocities for a set of crop widths. The results are presented in Figures 6.7 and 6.8, and Tables 6.4 and 6.5. ICP performs worse at **hg03** for smaller window sizes, where point densities are consistently lower and therefore fewer points are used for change detection; **hg02**, which has a higher point density, is relatively insensitive to crop radius. For this reason, we will use 500 m crop radius for all ICP runs.

Table 6.4: Statistics for ICP-derived velocities compared to GNSS-derived velocities at **hg02**, separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1} .

Crop radius	min	mean	max	σ	RMSE	Avg. number of points
100 m	4.54	21.68	62.21	9.27	3.42	350
250 m	12.41	24.49	91.24	13.22	3.15	1920
500 m	11.72	23.94	56.28	8.27	3.69	4348

Table 6.5: Statistics for ICP-derived velocities compared to GNSS-derived velocities at **hg03**, separated by the distance around the GNSS receiver to use for change detection. All values in m d^{-1} .

Crop radius	min	mean	max	σ	RMSE	Avg. number of points
100 m	8.93	24.59	111.22	16.66	12.27	104
250 m	11.89	29.61	144.46	22.73	12.61	714
500 m	11.31	26.57	74.79	13.76	11.10	3945

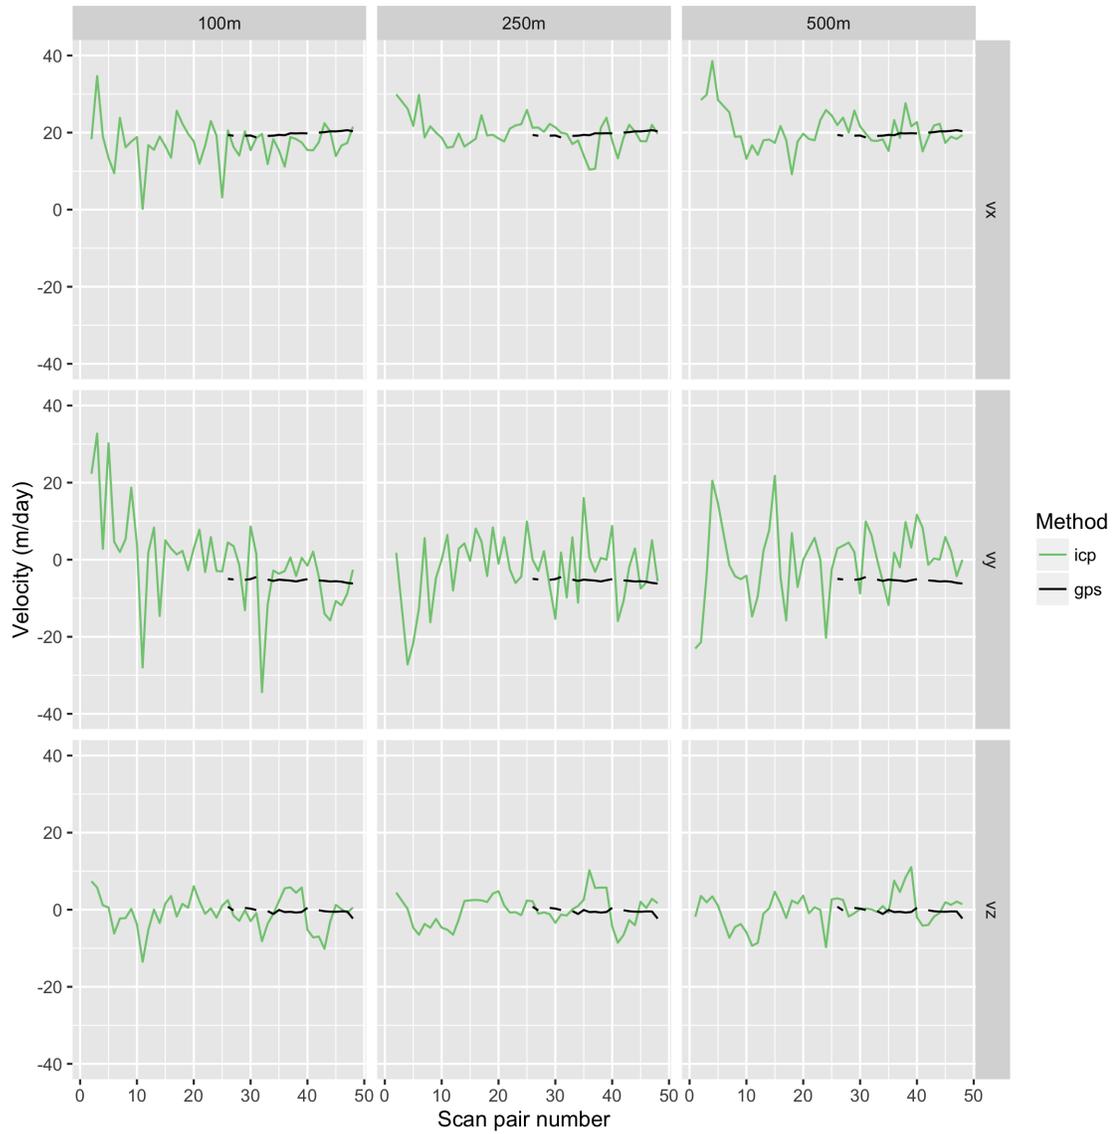


Figure 6.7: ICP-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg02 for a variety of crop widths (horizontal divides).

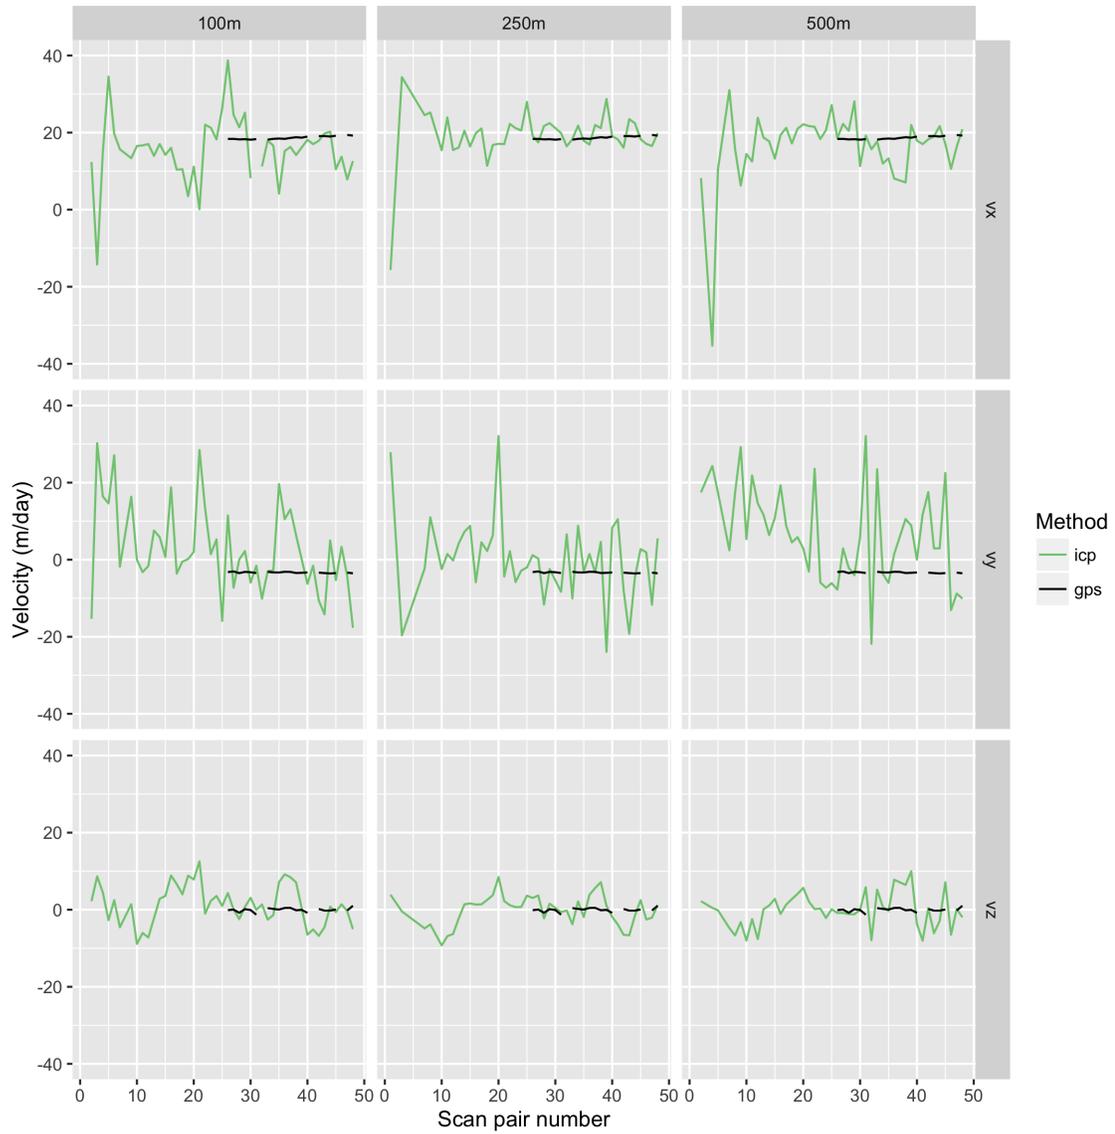


Figure 6.8: ICP-derived glacier velocities (vertical divides) compared with GNSS-derived velocities at GNSS site hg03 for a variety of crop widths (horizontal divides).

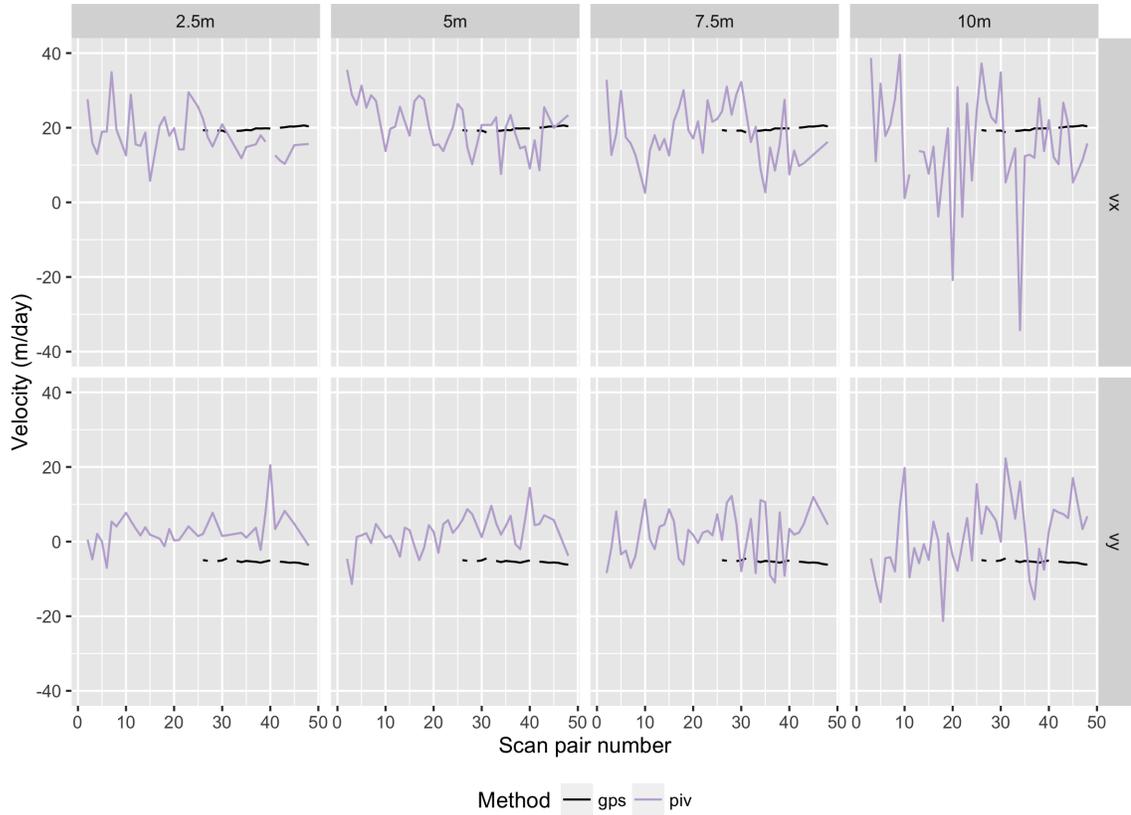


Figure 6.9: PIV-derived x and y velocities (vertical divide) at GNSS site hg02 for a variety of pixel sizes (horizontal divides).

6.5 Tuning PIV

In order to find the optimal set of parameters for PIV, we first examined the effect of varying the size of each pixel in the point cloud rasterization (Figure 6.9). RMSE values for the PIV-to-GNSS comparison are presented in Tables 6.6 and 6.7. At hg02 the 5 m rasterization performs best, while at hg03 the 10 m rasterization is best. However, the spread of RMSE at hg02 is much larger than that at hg03, so we concluded that 5 m pixel size is best for these data and used that value for our PIV rasterizations.

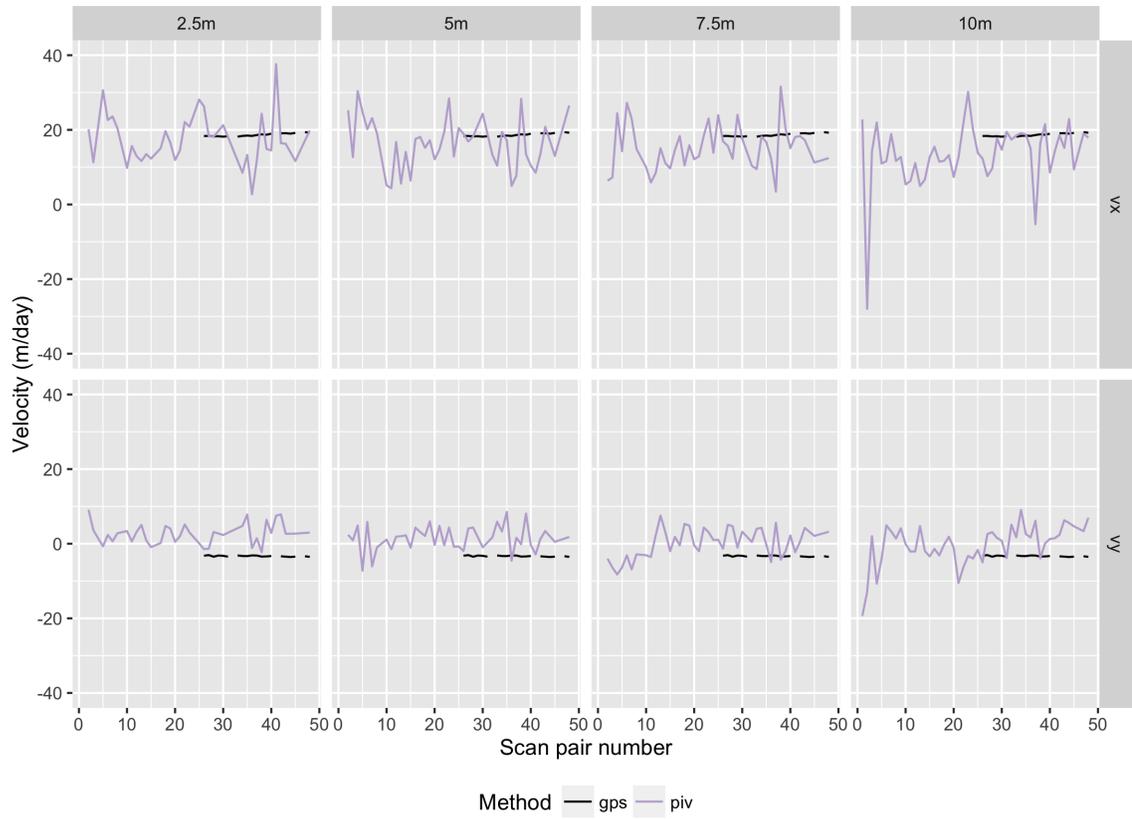


Figure 6.10: PIV-derived x and y velocities (vertical divide) at GNSS site hg03 for a variety of pixel sizes (horizontal divides).

Table 6.6: Statistics for PIV-derived velocities when compared to GNSS-derived velocities at hg02, separated by pixel size. All values in m d^{-1} .

Pixel size	min	mean	max	σ	RMSE
2.5 m	6.07	20.23	56.42	9.48	12.83
5 m	7.84	21.47	35.81	6.13	5.45
7.5 m	8.22	19.66	33.88	6.90	8.26
10 m	3.80	38.17	706.20	100.41	8.38

Table 6.7: Statistics for PIV-derived velocities when compared to GNSS-derived velocities at hg03, separated by pixel size. All values in m d^{-1} .

Pixel size	min	mean	max	σ	RMSE
2.5 m	2.97	19.03	58.77	9.41	6.20
5 m	4.59	16.77	30.75	6.77	6.34
7.5 m	6.59	15.87	31.86	5.78	6.01
10 m	5.37	15.63	30.82	6.21	5.47

Two peak-finding algorithms are available: three-point Gaussian fit or a two-dimensional Gaussian fit. The three-point Gaussian algorithm fits a Gaussian function to the distribution of pixel intensities in the neighboring pixels to achieve sub-pixel precision. The two-dimensional Gaussian is a more complex method that fits a two-dimensional Gaussian to a neighborhood, modeling elliptical intensity distributions (Thielicke and Stamhuis 2014). We tested both algorithms at both GNSS sites using a 5 m pixel size (Figure 6.11). RMSE results are presented in Tables 6.8 and 6.9. At both sites the three-point Gaussian fit performed better when compared with RMSE, and so we will use the three-point algorithm for this work.

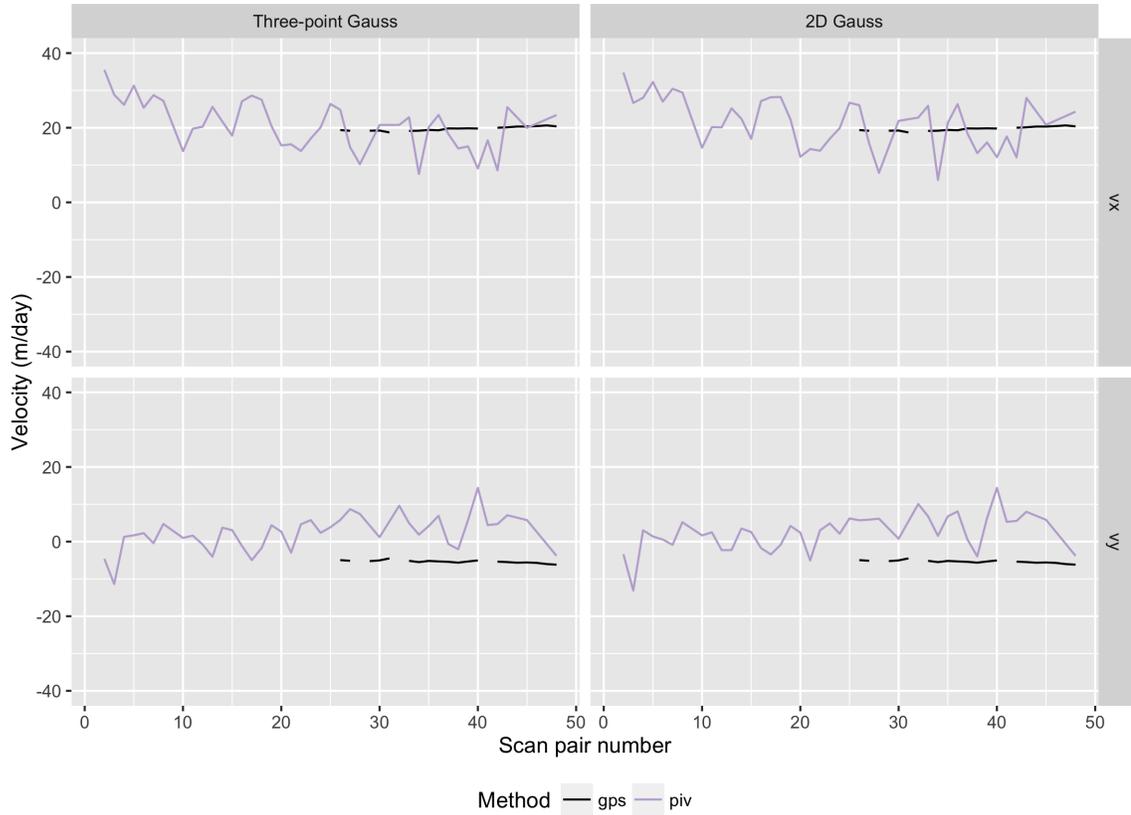


Figure 6.11: PIV-derived x and y velocities (vertical divide) at GNSS site hg02 for both peak-finding algorithms (horizontal divide).

Table 6.8: Statistics for PIV-derived velocities at hg02, compared to GNSS-derived velocities, separated by peak-finding algorithm. All values in m d^{-1} .

Peak-finding algorithm	min	mean	max	σ	RMSE
Three-point Gaussian	7.84	21.47	35.81	6.13	5.45
Two-dimensional Gaussian	6.20	22.06	34.96	6.60	6.08

Table 6.9: Statistics for PIV-derived velocities at hg03, compared to GNSS-derived velocities, separated by peak-finding algorithm. All values in m d^{-1} .

Peak-finding algorithm	min	mean	max	σ	RMSE
Three-point Gaussian	4.59	16.77	30.75	6.77	6.34
Two-dimensional Gaussian	3.56	16.03	36.91	7.86	7.17

6.6 Comparing change detection methods at the GNSS receivers at Helheim Glacier

Figure 6.12 shows three-dimensional CPD, ICP, and GNSS velocity measurements, and two-dimensional PIV measurements, in m d^{-1} , from the region surrounding **hg02**. Figure 6.13 shows the same around **hg03**. From visual inspection, ICP and CPD perform reasonably well when compared to the GNSS data, with ICP providing a noisier signal. Means and RMSE results for both **hg02** and **hg03** are presented in Table 6.10 and Table 6.11, respectively. The RMSE column for GNSS is the uncertainty in the GNSS-derived velocity, calculated by propagating the per-measurement uncertainty of the constituent GNSS points through to the final velocity using the General Law Of Propagation Of Variance (GLOPOV). Note that there are some data gaps in the GPS data that prevent us from producing a GPS velocity for the entire interval.

All three methods have mean x velocity values within 2 m d^{-1} of the GNSS-derived values, and mean overall velocity values within 3.5 m d^{-1} . These are acceptable values that indicate that each method is able to recover the overall velocity trend of the glacier. CPD has the lowest RMSE of the three methods, with 2.51 m d^{-1} at **hg02** and 3.05 m d^{-1} at **hg03**. In the x-direction, the velocity component with the largest magnitude, all three methods have RMSE values of less than 7 m d^{-1} at both locations, but in the y direction both ICP and PIV produce RMSE values larger than 10 m d^{-1} , when the y-velocities measured by GNSS are on the order of 3 m d^{-1} to 5 m d^{-1} at the two locations. While the average magnitudes may be off, CPD correctly identified the average direction of motion in all three dimensions at both locations; the same cannot be said of ICP or PIV.

From comparing change detection methods to GNSS records, CPD is more effective at extracting change vectors from these point cloud data. The poor performance of ICP and PIV in the y-direction could be due to the smaller absolute change in the y-direction, or possibly something to do with the directionality of the serac pattern. Around the GPS receivers, the seracs run roughly north-south, leading to greater data heterogeneity in the x-direction and more homogeneity in the y-direction; if ICP or PIV require more heterogeneity than CPD to work effectively, this could lead to the direction-dependent accuracy of the methods for these data. CPD performs *very* well in the y direction; the surprisingly tight match between CPD in the y direction and the GNSS records at both GNSS receivers bears additional investigation in future work.

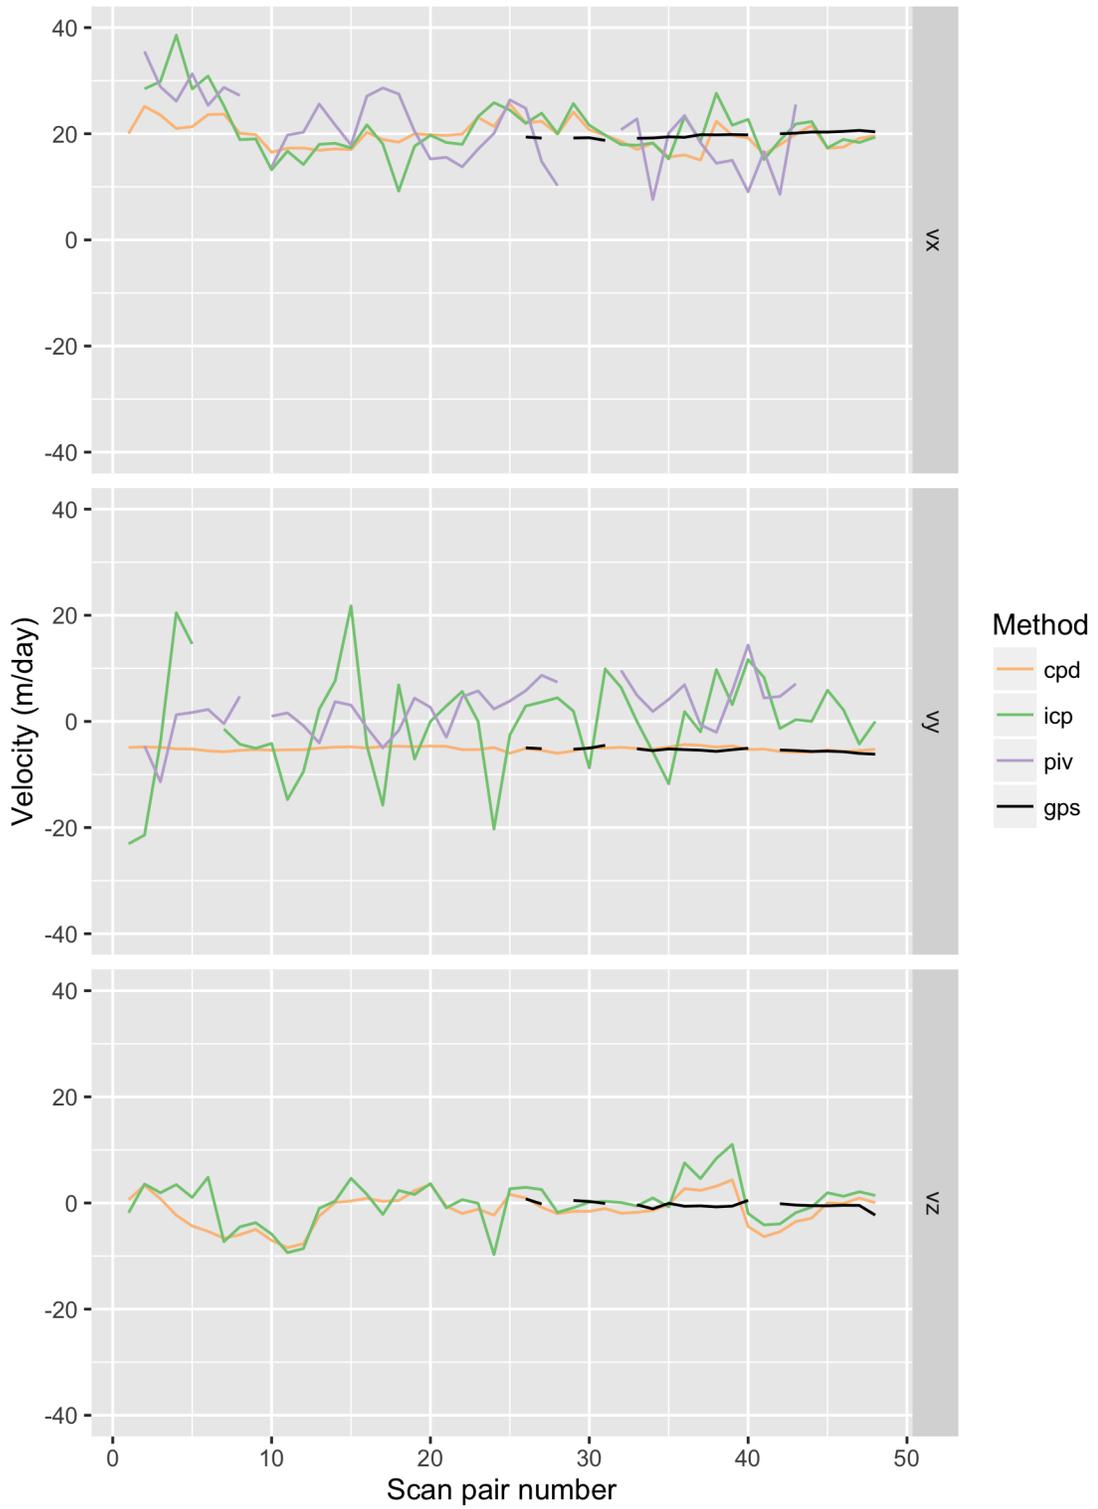


Figure 6.12: Three-dimensional change detection at hg02, 240 min between scan pairs.

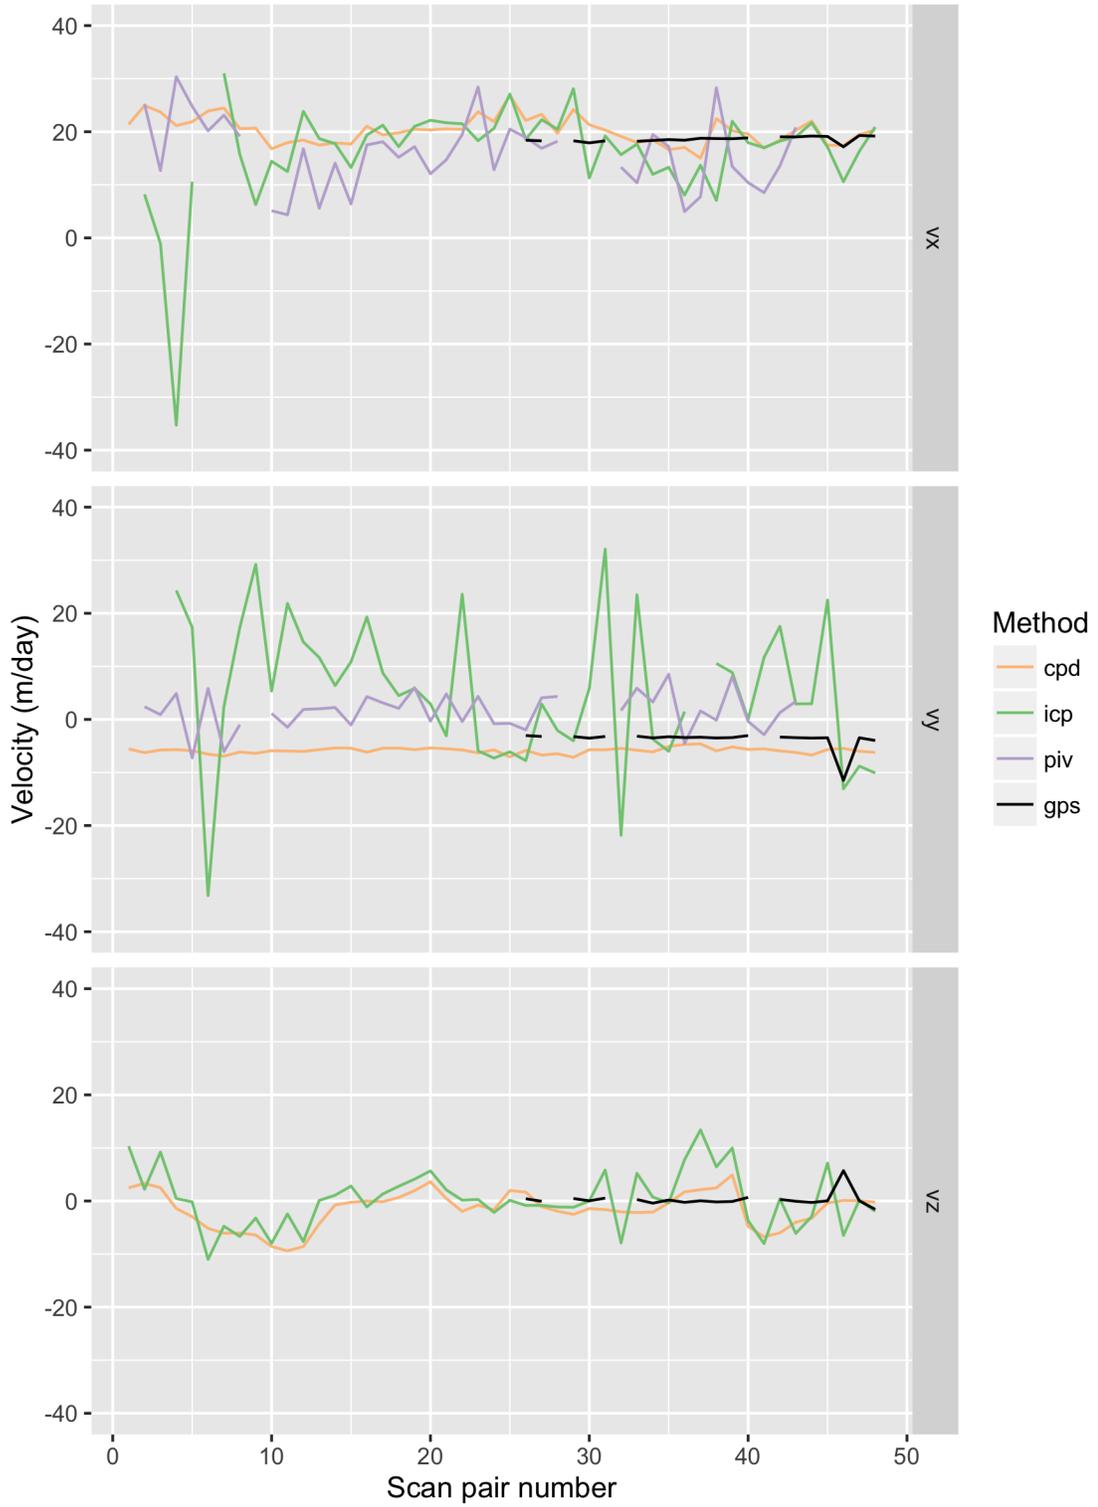


Figure 6.13: Three-dimensional change detection at hg03, 240 min between scan pairs.

Table 6.10: Mean and RMSE between change detection methods and GNSS measurements at hg02. Values in the RMSE columns for GNSS velocities are the standard deviation, calculated from the accuracy of each GNSS point propagated through to the final velocity. All values in m d^{-1} .

Method	μ_{vx}	$RMSE_{vx}$	μ_{vy}	$RMSE_{vy}$	μ_{vz}	$RMSE_{vz}$	μ_v	$RMSE_v$
GNSS	19.71	0.08	-5.37	0.08	-0.35	0.14	20.45	0.18
CPD	19.25	2.56	-5.19	0.51	-0.51	2.70	20.11	2.51
ICP	20.75	3.26	0.97	8.51	1.72	4.37	21.89	3.69
PIV	17.91	6.08	4.32	10.55	—	—	—	—

Table 6.11: Mean and RMSE between change detection methods and GNSS measurements at hg03. Values in the RMSE columns for GNSS velocities are the standard deviation, calculated from the accuracy of each GNSS point propagated through to the final velocity. All values in m d^{-1} .

Method	μ_{vx}	$RMSE_{vx}$	μ_{vy}	$RMSE_{vy}$	μ_{vz}	$RMSE_{vz}$	μ_v	$RMSE_v$
GNSS	18.66	0.07	-3.33	0.07	-0.04	0.14	18.96	0.18
CPD	19.81	2.63	-5.87	2.61	-0.95	2.66	20.85	3.05
ICP	17.24	5.18	7.14	18.78	1.90	5.25	23.24	11.10
PIV	16.51	6.70	2.17	6.49	—	—	—	—

While CPD is better than ICP and PIV at determining the velocity of the glacier when compared to in-situ GNSS data, an overall velocity RMSE of 2.51 m d^{-1} at hg02 indicates substantial velocity detection errors that may mask smaller signals, such as a 2 m d^{-1} change after a glacial earthquake (Juan et al., 2010). This indicates that CPD should be preferred to ICP and PIV for extracting glacier velocities, but that it is no replacement for highly accurate GNSS sensors placed on-ice. Further work needs to be done to determine if errors in CPD could be improved through higher data densities; we expect that higher data densities and longer intervals between scan pairs could lead to better accuracy.

6.7 Comparing change detection methods at the three additional Helheim Glacier points

We selected three additional points (Table 5.2) at which we conducted relative comparisons of the three change detection methods. While more points could have been selected for inspection, we chose three points so that the results could be easily examined in chart form; because we are not doing statistical analysis on the glacier velocities query locations, we feel as though additional locations would not provide additional information to the reader. The three points represent different data densities, with $r1$ providing lower data densities than $r3$; data density is roughly proportional to the range from the scanner. Due to the increased point density and closer range of these points, we used smaller crop radii for ICP and CPD, listed in Table 6.12. Results are shown in Figure 6.14.

At $r1$, the location about 3758 m from the scanner, all three change detection methods capture the x velocity well, with PIV providing a noisier signal than the other two. As we saw with the absolute accuracy test, both PIV and ICP produce a very noisy y velocity, while CPD produces a consistent signal that, at least for our absolute accuracy tests, has less error than the other two methods when compared to the GNSS-observed values. At the other two locations, PIV tends to break down, possibly because the higher data densities at the closer ranges leave fewer “data gaps” in the raster. Because PIV was developed to track particles as they move through a scene, alternating light and dark portions of the image (such as spaced-apart seracs) serve as particles that can be tracked by the cross-correlation portion of the algorithm. Without distinct light and dark regions in the image, the cross-correlation is unable to detect similar but translated regions of the image. Both ICP and CPD track each other well, with CPD showing a consistently larger x velocity than ICP. For z velocities, both ICP and CPD capture a very similar signal, with the two z velocity signals at the high-density location $r3$ tracking each other almost identically.

Because ICP is commonly used in other research for change detection purposes (Zhang et al., 2015), we are encouraged by the close agreement between ICP and CPD in the x and z velocities, and as we have shown in the previous section CPD’s consistent y velocity signal is likely a better measure of true glacier velocity.

Table 6.12: Crop radii used for ICP and CPD at the three additional relative comparison points on the Helheim Glacier.

Name	Distance from scanner	Crop radius
r1	3758 m	250 m
r2	2747 m	250 m
r3	1700 m	100 m

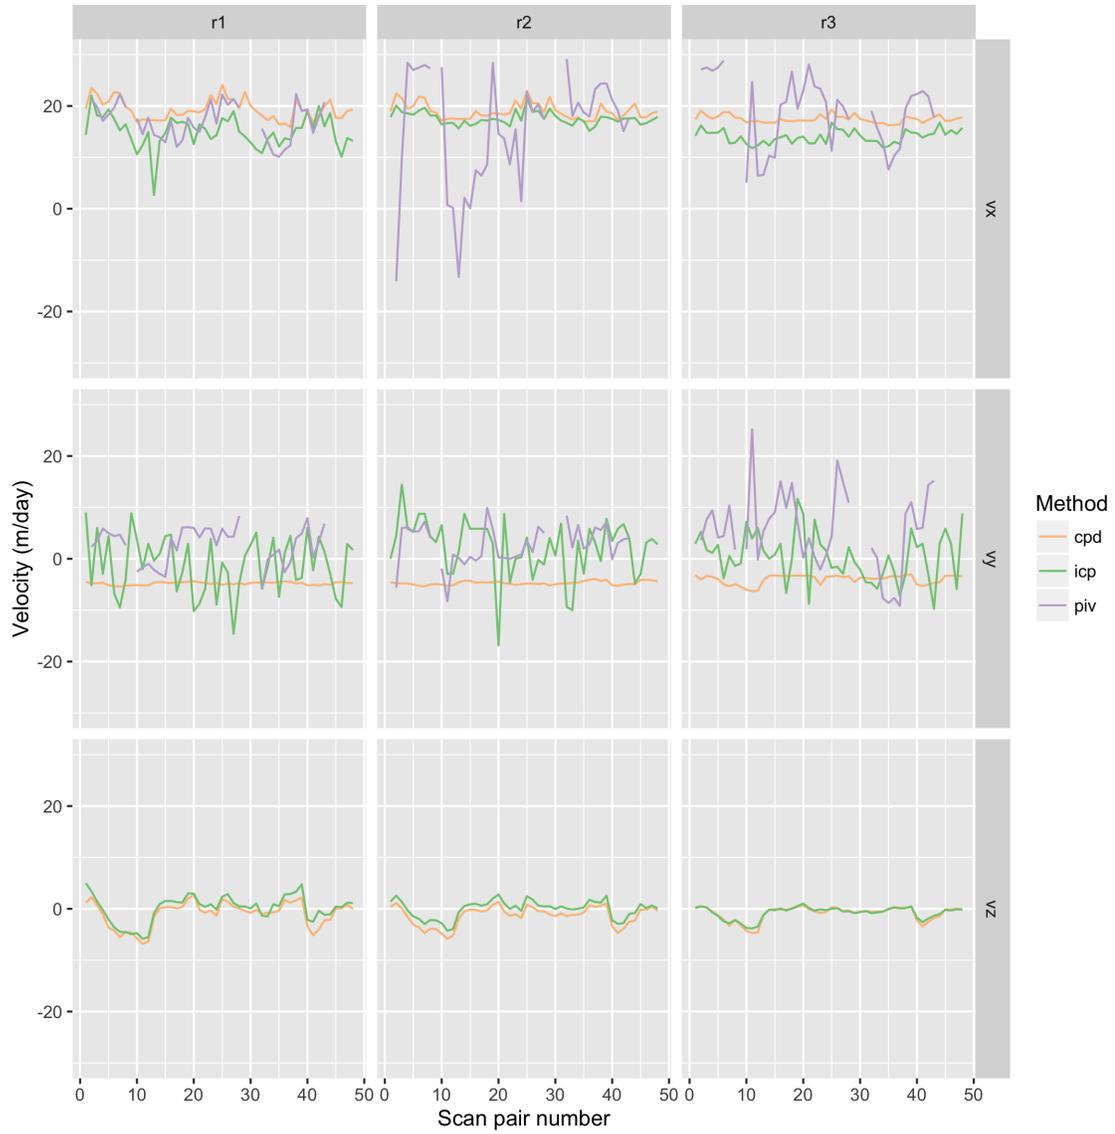


Figure 6.14: Glacier surface velocities at the Helheim Glacier, as computed by the three change detection methods. Results are separated by location, r1, r2, and r3, in the horizontal dimension, and by velocity component x, y, and z, in the vertical dimension.

6.8 Full-glacier change detection at Helheim Glacier

When assessing the relative effectiveness of each change detection method, we will look at the accuracy of the overall velocity field using a-priori knowledge of the glacier flow, the variations and noise within the overall pattern, and the density and distribution of the extracted change vectors.

The Helheim Glacier generally flows east to west, with faster flow in the middle of the glacier and slower flow near the shores (Howat et al., 2005). As we can see from Figure 6.15, all three change detection methods capture these general patterns. Both the PIV and CPD results additionally detect a speedup of the glacier close to the terminus.

The CPD results are less noisy than those of ICP and PIV; both ICP and PIV are noisy, particularly in the y-direction. ICP also has several regions of obviously incorrect results, indicating more inherent instability for these data than either PIV or CPD. When inspecting the typical point densities in those locations (Figure 6.16), there are “holes” in the data where ICP performs poorly. This may indicate that CPD and PIV perform better than ICP in regions of lower point density; additional work should be done to quantitatively assess the robustness of all three algorithms to changes in point density.

PIV has a distinct advantage in terms of velocity vector density — because the data are rasterized before processing, velocity vector density is standardized over the entire dataset. Conversely, ICP and CPD require smaller segments in the near field and larger segments in the far field for this TLS data set. When looking at the frequency distribution of velocity vectors from each change detection method, Figure 6.17, it is interesting to note some apparent clipping effects in the y velocity values for ICP; there are clearly some artifacts in the ICP results. This may be due to the ICP algorithm finding local minima in the y-direction, though tests with different time intervals between scans produced inconclusive results.

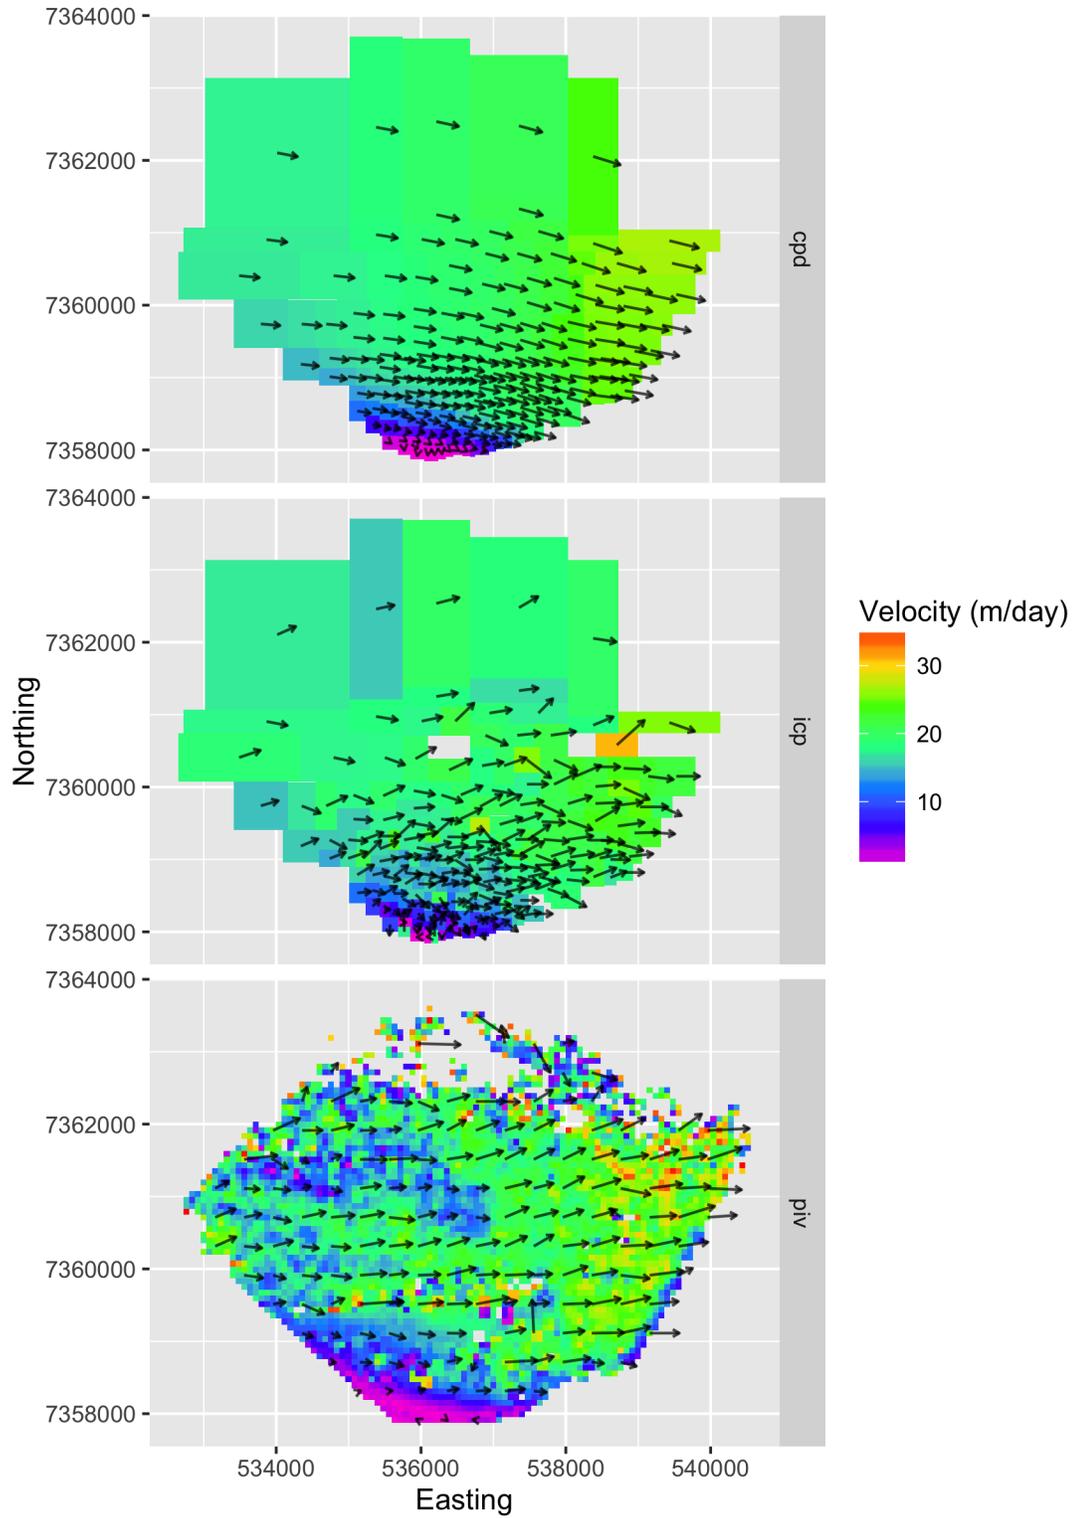


Figure 6.15: Full-glacier change detection of Helheim Glacier. 2014-07-13 18:17:00 to 2014-07-13 22:24:12. Arrows indicate direction and magnitude of velocity in the x-y plane, background indicates total velocity magnitude.

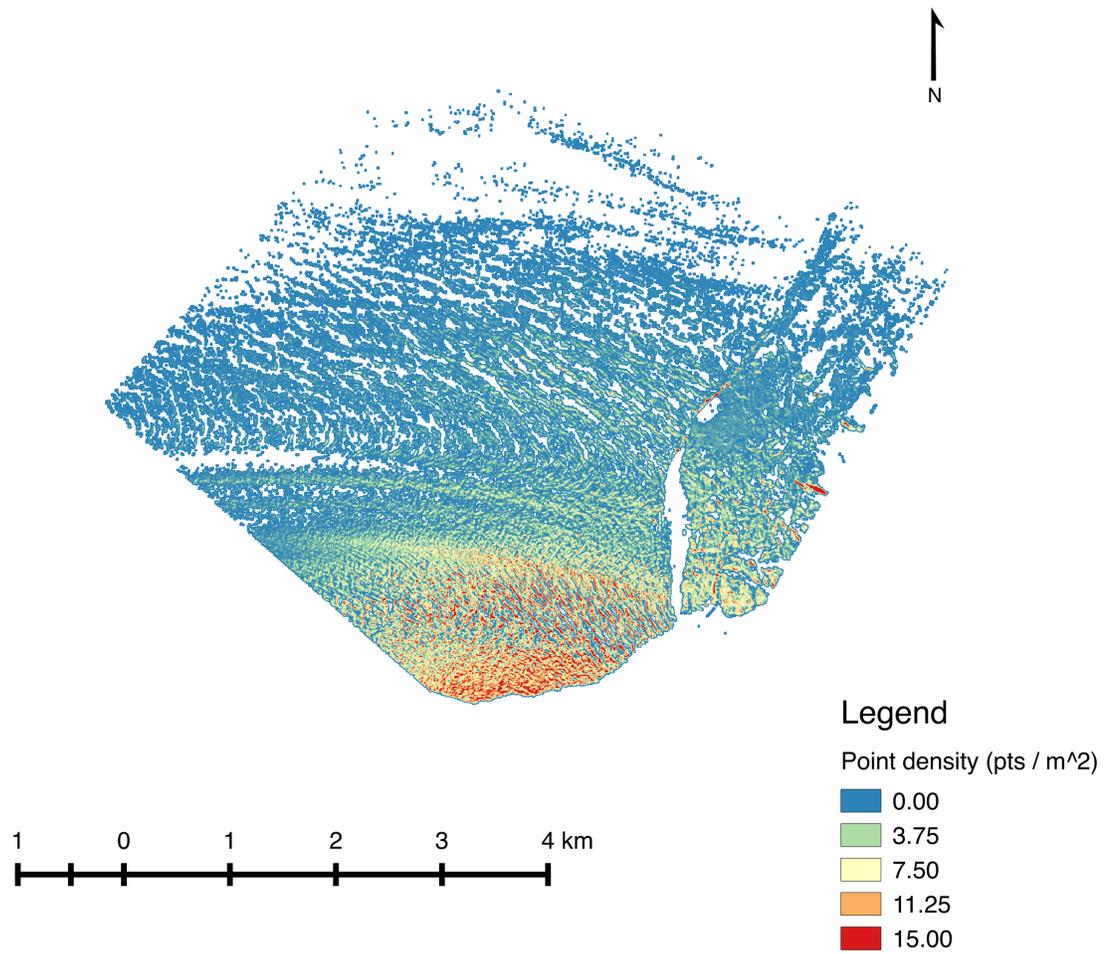


Figure 6.16: Typical point density, in points/m, for a scan of the Helheim Glacier. Scan data from 2016-07-10 20:28:08 UTC.

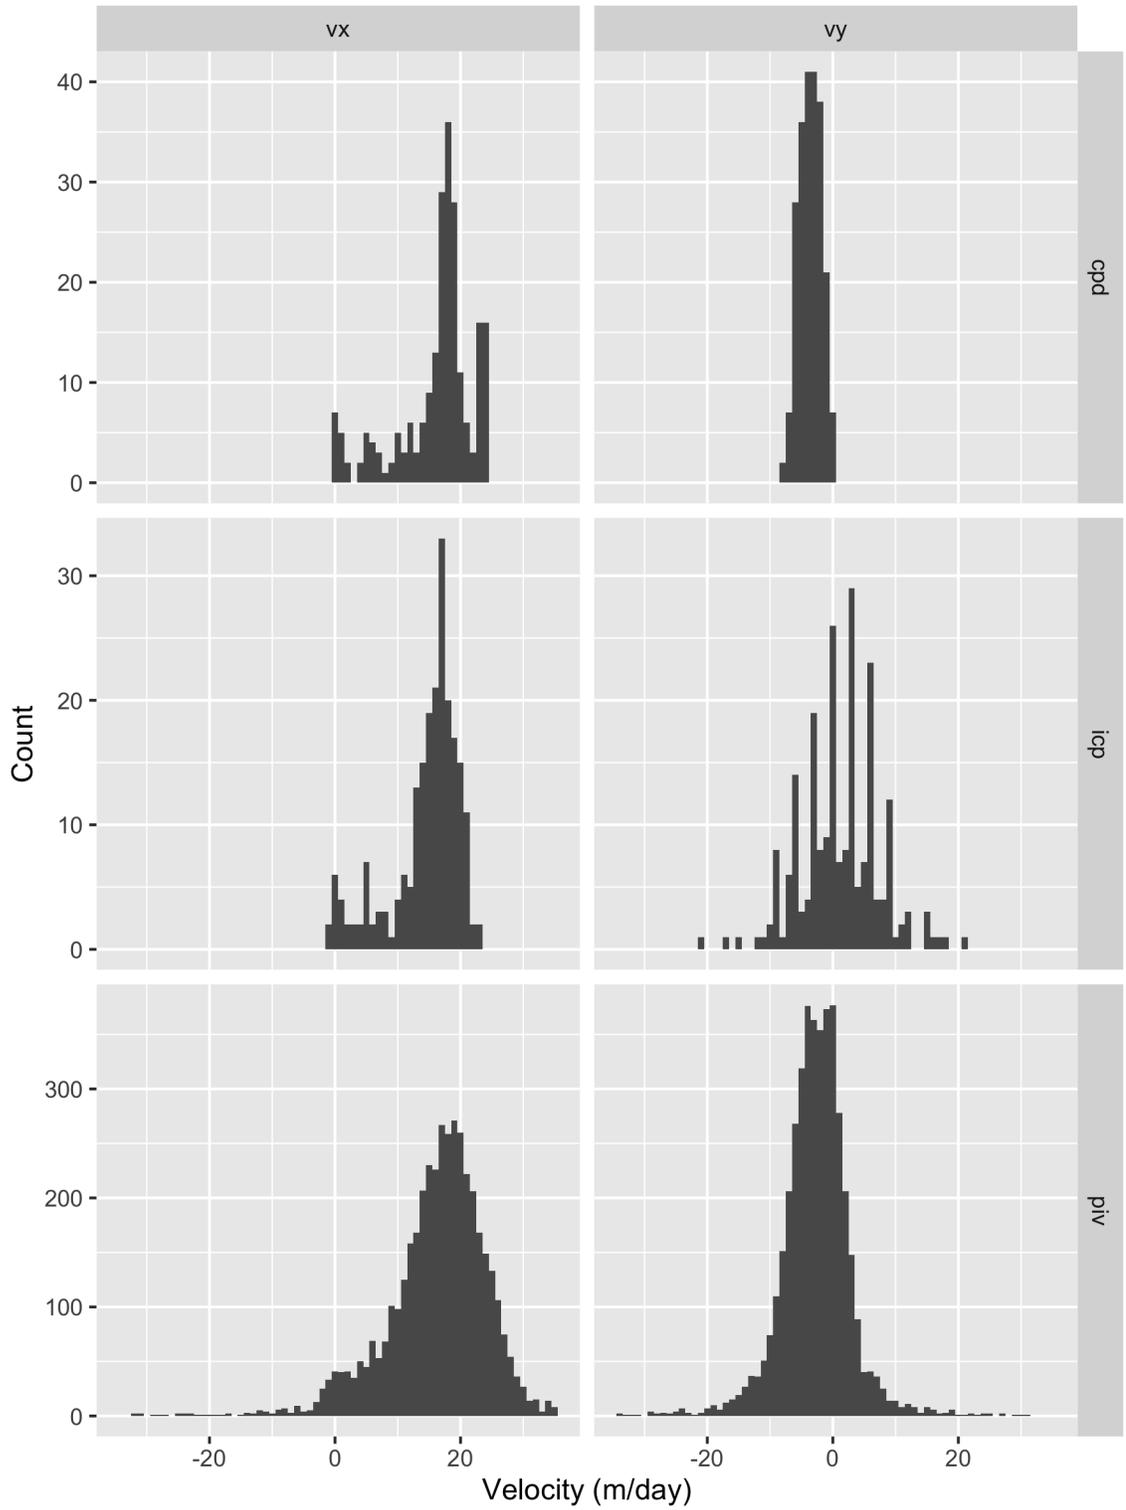


Figure 6.17: Frequency distributions of velocity vectors for change detection runs on scans from 2014-07-13 18:17:00 to 2014-07-13 22:24:12. Only x and y velocities are presented.

6.9 Full-glacier change detection at Canada Glacier

Change detection is more difficult on our data from the Canada Glacier for a few of reasons. First, the low data density of the 2001 dataset (about 0.2 points/m) greatly reduces the amount of usable features in the LiDAR dataset. Second, because both datasets were collected by airborne scanners, there are fewer shadows or other artificial features for the change detection algorithms to use. All three change detection algorithms perform poorly in regions of low surface roughness, and so we filter our results by the average surface roughness in the region of interest. Surface roughness is defined as the largest inter-cell difference of a central pixel and its surrounding cells, per M. F. J. Wilson et al., (2007); the surface roughness of the Canada Glacier is presented in Figure 6.18. Comparing the roughness of the Canada Glacier with that of the Helheim Glacier in Figure 6.19; the Helheim Glacier is much rougher. Full-glacier change detection results for the Canada Glacier are presented in Figure 6.20, filtered so that only results with an average surface roughness of 0.75 m or more are shown. Future research is necessary to determine the effects of surface roughness on each of the three change detection algorithms.

PIV was run with a raster cell size of 5 m, and ICP and CPD were run with an average window size of about 165 m to a side. The PIV and CPD results provide information about a gradually-moving glacier front (near the south part of the image). The ICP results do pick up the motion on the glacier front, but the magnitude of motion is less than that detected by CPD and PIV. Both PIV and CPD methods detect general southerly motion of the glacier near the tongue, and divergent flow in the x direction, with the western portion of the tongue flowing west and the eastern portion flowing east. The magnitude of the velocities observed by PIV and CPD are similar, between 1 m/year to 2 m/year; ICP velocities are much less.

Fountain et al., (2006) found a general southward trend in glacier velocity and faster velocities in the “neck” of the glacier, north of the tongue, with slower velocities near the terminus of the glacier. While we also observed a general southward trend and a very rough appearance of horizontal (x-direction) spreading near the terminus of the glacier, our results are clearly hampered by the lack of surface roughness in the LiDAR data over much of the glacier. Without significant roughness on center of the glacier surface at the Canada Glacier, all three change detection algorithms struggle to produce acceptable results. However, near the tongue (southmost portion) of the glacier, where

surface roughness is higher, we find that both CPD and PIV produce reasonable results when compared to Fountain et al., (2006), about 1 to 2 meters per year.

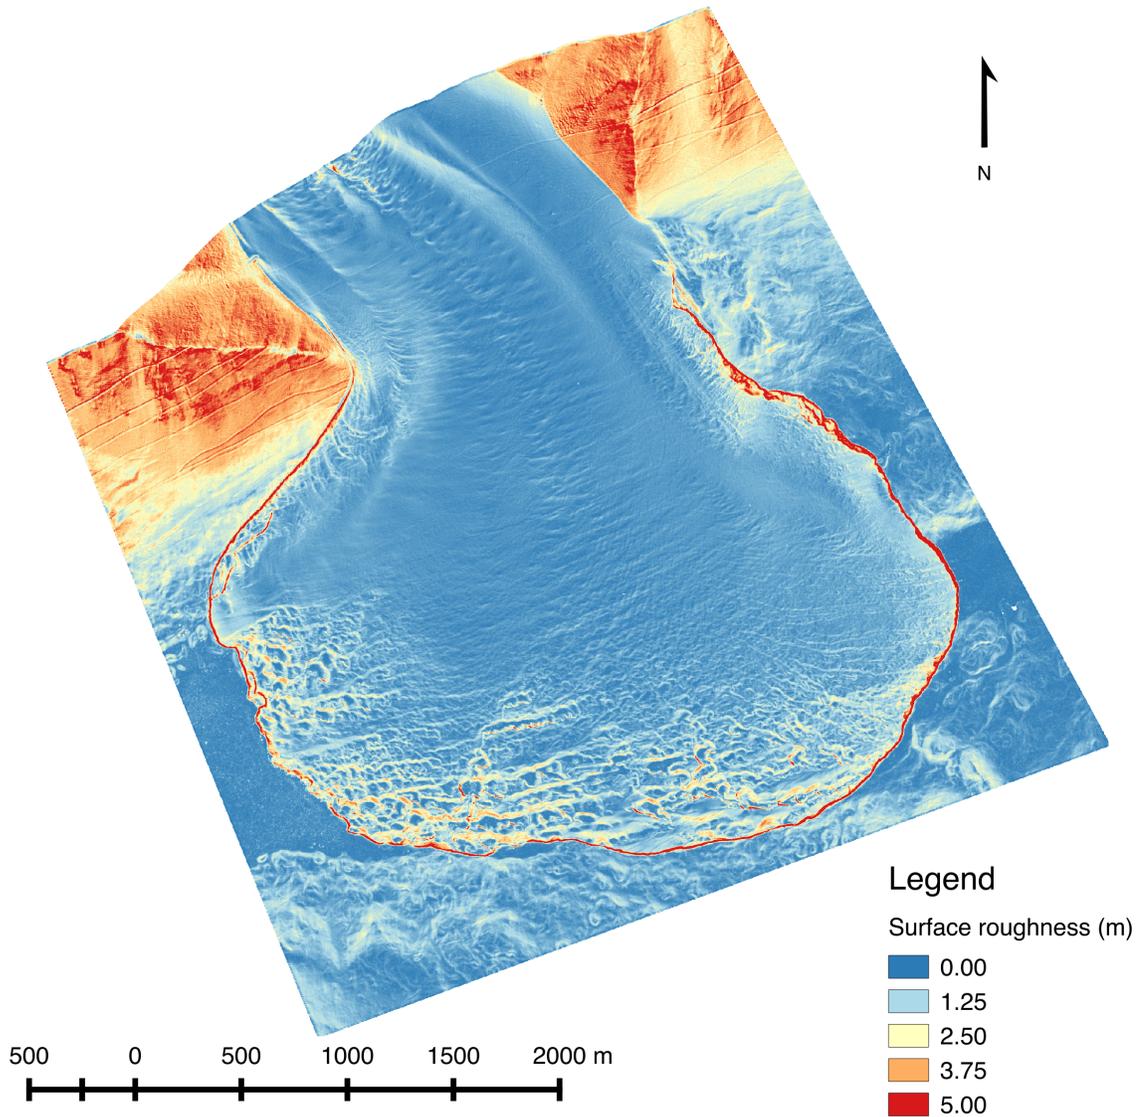


Figure 6.18: Surface roughness at the Canada Glacier. Surface roughness has the units of m and is defined as the largest inter-cell difference between a central pixel and its surrounding cells. Data from the 2001 NASA collect.

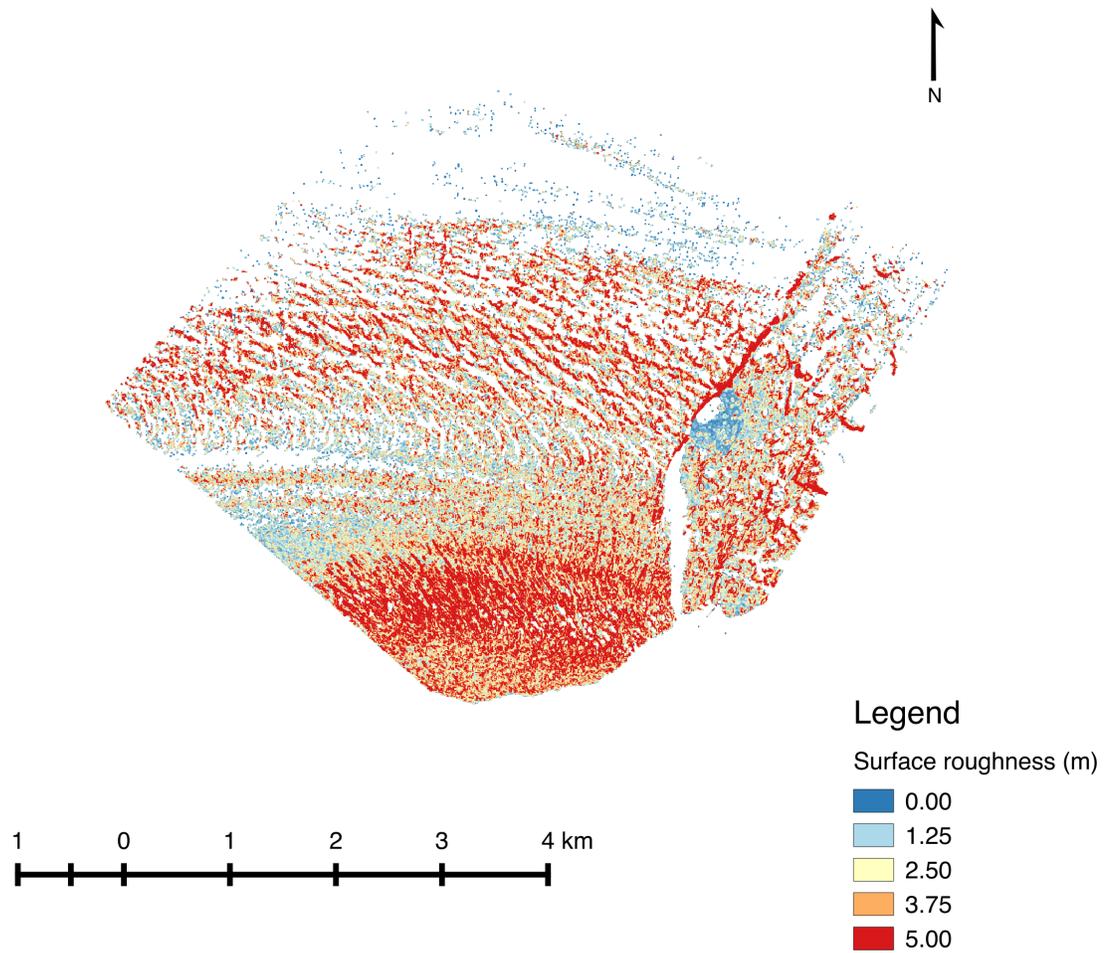


Figure 6.19: Surface roughness at the Helheim Glacier. Surface roughness is in units of m and is defined as the largest inter-cell difference between a central pixel and its surrounding cells. Data from the TLS scan at 2016-07-10 20:28:08 UTC.

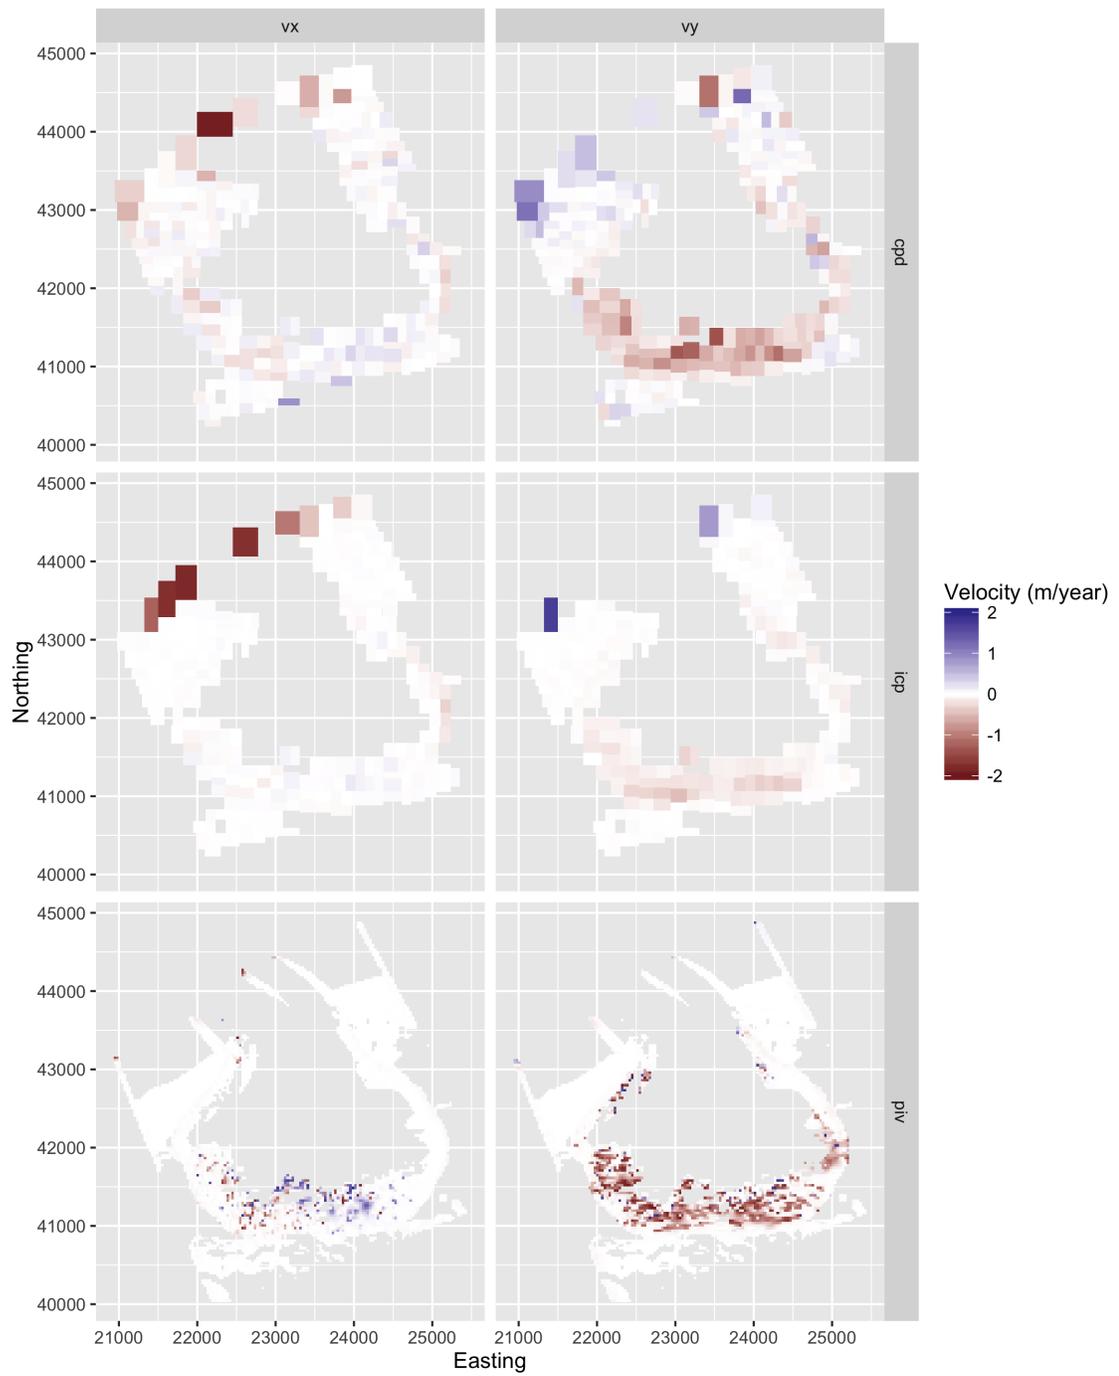


Figure 6.20: Full-glacier change detection at the Canada Glacier. Velocities are presented in m/year due to the slow-moving nature of the Canada Glacier. Results are filtered to only present velocities where the average roughness is greater than 0.75 m, where roughness is defined as the largest inter-cell difference between a central pixel and its surrounding cells.

7 Conclusions and future work

When comparing the three change detection methods, Coherent Point Drift (CPD), Iterative Closest Point (ICP), and Particle Imaging Velocimetry (PIV), on terrestrial LiDAR data of the Helheim Glacier, CPD outperformed the other two methods both in terms of absolute accuracy when compared to GNSS data, and when compared qualitatively on a full-glacier basis; the x and y velocity vectors retrieved by CPD agreed best with the known behavior of the Helheim Glacier. PIV does provide a more uniform and denser distribution of velocity vectors, which could be more useful for some applications; however, we use PIV only to produce two-dimensional velocity vectors. There is active research into using PIV to get two-dimensional velocities, then subtracting the adjusted point clouds to extract the third dimension of motion, but nothing has been published on the topic as of this writing. In terms of overall accuracy, all three methods generally provide acceptable results, particularly in the dimension with the most change (x velocities), with CPD providing the lowest RMSE and closest mean to spatially coincident data collected with GNSS receivers. Therefore, we conclude that CPD is the best of the three methods for extracting change from TLS glacier datasets, if relatively sparse velocity data is acceptable. For dense velocity data, PIV provides the best option, though it works only on a rasterization of the point cloud data, not the points themselves.

For the Canada Glacier datasets, while PIV and CPD produced reasonable results in areas of high surface roughness when compared to existing literature, ICP underestimated the glacier velocity. All three methods struggled to produce a coherent picture of glacier change in regions of low surface roughness.

Improvements to all three methods are possible; in particular, refinements such as anisotropic analysis could improve ICP and PIV. CPD is more computationally intensive than the other methods, and takes much longer to run on the same data. Additional performance improvements such as increased parallelization and finer tuning of the Fast Gauss Transform algorithm, would make CPD more usable for large datasets and in quick-turnaround environments.

Additional work needs to be done to refine the accuracy assessments of each method for various environments. On-ice GNSS proved to be a useful tool for quantifying the accuracy of these methods — additional studies with coincident on-ice GNSS, particularly with GNSS in regions of varying

point cloud density and glacier surface character, would establish the effects of certain terrain features and data characteristics. More methods should be brought in as well, in particular the popular photogrammetry and Structure from Motion methods, to provide more bounds on reasonable expectations for remotely-sensed glacier velocities.

Each change detection method can be adapted to incorporate per-point uncertainties, as Zhang et al., (2015) did for Iterative Closest Point. Future work will include more analysis of the segmentation method for ICP and CPD, and whether the segmentation can affect final results.

Coherent Point Drift has a non-rigid variant, which was not used in this study due to its poor performance for large datasets. Exploration of non-rigid CPD, including performance enhancements, would provide a change detection method that is able to capture higher-resolution non-rigid deformation, e.g. as provided by PIV, in three dimensions using the original source point cloud data. Of additional interest would be the use of more dimensions in CPD analysis, such as the inclusion of point intensity information; while intensity information may not be as useful in a snow and ice covered environment, it might be more useful in a scene with greater heterogeneity of reflectivity, such as a landslide with trees and rocks.

Finally, additional datasets always provide more forums for exploration. At the Helheim Glacier in particular work continues to collect more TLS data of the glacier. As of July 2015 a remote, automated LiDAR system has been installed near the survey site used in this work, collecting TLS scans at 6 h intervals. The 6 h interval was chosen as a compromise between collecting the most data possible within the power constraints of the system. This unprecedented amount of TLS data of a glacier can be analyzed with these change detection methods to retrieve glacier velocity over a large period of time.

For terrestrial LiDAR data in particular Coherent Point Drift (CPD) proved to be the most accurate when compared to coincident on-ice GNSS. For airborne LiDAR data, all three methods failed to fully capture the glacier velocity, though the results did contain some useful velocities in regions of higher crevassing and speeds. Though imperfect, these three change detection methods (Iterative Closest Point, Particle Imaging Velocimetry, and Coherent Point Drift), provide tools for calculating the velocity of glaciers from point cloud data.

References

Abdalati, By Waleed, H Jay Zwally, Robert Bindshadler, Bea Csatho, Sinead Louise Farrell, Helen Amanda Fricker, David Harding, Ronald Kwok, Michael Lefsky, Thorsten Markus, Alexander Marshak, Thomas Neumann, Stephen Palm, Bob Schutz, Ben Smith, James Spinhirne, and Charles Webb. 2015. “The ICESat-2 Laser Altimetry Mission” 98 (5): 735–51.

Ahlstrøm, Andreas P., S. B. Andersen, Morten L. Andersen, H. Machguth, F. M. Nick, Ian Joughin, C. H. Reijmer, R. S W Van De Wal, J. P. Merryman Boncori, J. E. Box, M. Citterio, D. Van As, R. S. Fausto, and A. Hubbard. 2013. “Seasonal velocities of eight major marine-terminating outlet glaciers of the Greenland ice sheet from continuous in situ GPS instruments.” *Earth System Science Data* 5 (2): 277–87. doi:10.5194/essd-5-277-2013¹.

Ahn, Yushin, and Jason E. Box. 2010. “Instruments and Methods Glacier velocities from time-lapse photos: Technique development and first results from the Extreme Ice Survey (EIS) in Greenland.” *Journal of Glaciology* 56 (198): 723–34. doi:10.3189/002214310793146313².

Andersen, Morten L., Tine B. Larsen, Meredith Nettles, Pedro Elósegui, D. Van As, Gordon S. Hamilton, Leigh A. Stearns, James L. Davis, A. P. Ahlstrm, J. De Juan, Göran Ekström, L. Stenseng, S. Abbas Khan, René Forsberg, and D. Dahl-Jensen. 2010. “Spatial and temporal melt variability at Helheim Glacier, East Greenland, and its effect on ice dynamics.” *Journal of Geophysical Research: Earth Surface* 115 (4): 1–18. doi:10.1029/2010JF001760³.

Andresen, Camilla S., Fiammetta Straneo, Mads Hvid Ribergaard, Anders A. Bjørk, Thorbjørn J. Andersen, Antoon Kuijpers, Niels Nørgaard-Pedersen, Kurt H. Kjær, Frands Schjøth, Kaarina Weckström, and Andreas P. Ahlstrøm. 2012. “Rapid response of Helheim Glacier in Greenland to climate variability over the past century.” *Nature Geoscience* 5 (1). Nature Publishing Group: 37–41. doi:10.1038/ngeo1349⁴.

Aryal, Arjun, Benjamin A. Brooks, Mark E Reid, Gerald W Bawden, and Geno R Pawlak. 2012. “Displacement fields from point cloud data: Application of particle imaging velocimetry to landslide

¹<https://doi.org/10.5194/essd-5-277-2013>

²<https://doi.org/10.3189/002214310793146313>

³<https://doi.org/10.1029/2010JF001760>

⁴<https://doi.org/10.1038/ngeo1349>

geodesy.” *Journal of Geophysical Research: Earth Surface* 117 (F1).

ASPRS. 2013. “LAS Specification: Version 1.4 - R13.”

Aull, Brian F, Andrew H Loomis, Douglas J Young, Richard M Heinrichs, Bradley J Felton, Peter J Daniels, and Deborah J Landers. 2002. “Geiger-Mode Avalanche Photodiodes for Three-Dimensional Imaging.” *Lincoln Laboratory Journal* 13 (2): 335–50.

Baltsavias, Emmanuel P. 1999. “Airborne laser scanning: basic relations and formulas.” *ISPRS Journal of Photogrammetry and Remote Sensing* 54 (2-3): 199–214.

Berardino, Paolo, Giafranco Fornaro, Riccardo Lanari, and Eugenio Sansosti. 2002. “A new algorithm for monitoring localized deformation phenomena based on small baseline differential SAR interferograms.” *IEEE International Geoscience and Remote Sensing Symposium* 2 (11): 2375–83. doi:10.1109/IGARSS.2002.1025900⁵.

Besl, Paul J, and Neil D McKay. 1992. “A Method for Registration of 3D-Shapes.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14 (2): 239–56. doi:10.1109/34.121791⁶.

Bevan, S. L., A. J. Luckman, and T. Murray. 2012. “Glacier dynamics over the last quarter of a century at Helheim, Kangerdlugssuaq and 14 other major Greenland outlet glaciers.” *Cryosphere* 6 (5): 923–37. doi:10.5194/tc-6-923-2012⁷.

Castorena, Juan, and Charles D Creusere. 2015. “Sampling of Time-Resolved Full-Waveform LIDAR Signals at Sub-Nyquist Rates” 53 (7): 3791–3802.

Chen, Y., and G. Medioni. 1991. “Object modeling by registration of multiple range images.” doi:10.1109/ROBOT.1991.132043⁸.

Chetverikov, D., D. Svirko, D. Stepanov, and P. Krsek. 2002. “The Trimmed Iterative Closest Point algorithm.” *Object Recognition Supported by User Interaction for Service Robots* 3 (c): 0–3. doi:10.1109/ICPR.2002.1047997⁹.

Cho, Peter, Hyrum Anderson, Robert Hatch, and Prem Ramaswami. 2006. “Real-time 3-D ladar imaging.” *Proceedings - HPCMP Users Group Conference, UGC 2006* 16 (1): 321–26.

⁵<https://doi.org/10.1109/IGARSS.2002.1025900>

⁶<https://doi.org/10.1109/34.121791>

⁷<https://doi.org/10.5194/tc-6-923-2012>

⁸<https://doi.org/10.1109/ROBOT.1991.132043>

⁹<https://doi.org/10.1109/ICPR.2002.1047997>

doi:10.1109/HPCMP-UGC.2006.63¹⁰.

Deems, Jeffrey S., Thomas H. Painter, and David C. Finnegan. 2013. "Lidar measurement of snow depth: A review." *Journal of Glaciology* 59 (215): 467–79.

Degnan, John, Roman Machan, Ed Leventhal, David Lawrence, Gabriel Jodor, and Christopher Field. 2008. "Inflight performance of a second-generation photoncounting 3D imaging lidar." *SPIE Defense and Security Symposium*, 695007–7–9. doi:10.1117/12.784759¹¹.

DeLong, Stephen B, James J Lienkaemper, Alexandra J Pickering, and Nikita N Avdievitch. 2015. "Rates and patterns of surface deformation from laser scanning following the South Napa earthquake, California." *Geosphere* 11 (6): 2015–30. doi:10.1130/GES01189.1¹².

Dickey, Jean O., P L Bender, J E Faller, X. X. Newhall, R L Ricklefs, J G Ries, P J Shelus, C Veillet, A L Whipple, J R Wiant, James G. Williams, and C F Yoder. 1994. "Lunar laser ranging: a continuing legacy of the apollo program." *Science (New York, N.Y.)* 265 (5171): 482–90. doi:10.1126/science.265.5171.482¹³.

Ferreira, Benjamin R. 2014. "High-resolution lidar mapping and analysis to quantify surface movement of Swift Creek landslide , Whatcom County , WA."

Fountain, Andrew G, Thomas H Nylen, K MacClune, and Gayle L Dana. 2006. "Glacier mass balances (1993-2001) Taylor Valley, McMurdo Dry Valleys, Antarctica." *Journal of Glaciology* 52 (178): 451. doi:LTER¹⁴.

Ghilani, Charles D. 2010. *Adjustment Computations*. doi:10.1002/9780470121498¹⁵.

Glennie, Craig L. 2007. "Rigorous 3D error analysis of kinematic scanning LIDAR systems." *Journal of Applied Geodesy* 1 (3). Walter de Gruyter GmbH.

Glennie, Craig L., William E. Carter, Ramesh L. Shrestha, and W E Dietrich. 2013. "Geodetic imaging with airborne LiDAR: the Earth's surface revealed." *Reports on Progress in Physics* 76 (8):

¹⁰<https://doi.org/10.1109/HPCMP-UGC.2006.63>

¹¹<https://doi.org/10.1117/12.784759>

¹²<https://doi.org/10.1130/GES01189.1>

¹³<https://doi.org/10.1126/science.265.5171.482>

¹⁴<https://doi.org/LTER>

¹⁵<https://doi.org/10.1002/9780470121498>

86801.

“GNU General Public License, version 2.” 1991. <https://opensource.org/licenses/BSD-3-Clause>.
<https://www.gnu.org/licenses/old-licenses/gpl-2.0.en.html>.

“GNU Lesser General Public License.” 2007. <https://www.gnu.org/licenses/lgpl-3.0.en.html>.

Greengard, Leslie, and John Strain. 1991. “The Fast Gauss Transform.” *Statistics and Computing* 12 (1): 79–94.

Gruen, Armin, and Devrim Akca. 2005. “Least squares 3D surface and curve matching.” *ISPRS Journal of Photogrammetry and Remote Sensing* 59 (3): 151–74. doi:10.1016/j.isprsjprs.2005.02.006¹⁶.

Hartzell, Preston J., Peter J. Gadowski, Craig L. Glennie, David C. Finnegan, and Jeffrey S. Deems. 2015. “Rigorous Error Propagation for Terrestrial Laser Scanning with Application to Snow Volume Uncertainty.” *Journal of Glaciology* 61 (230): 1147–58. doi:10.3189/2015JoG15J031¹⁷.

Hartzell, Preston J., Craig L. Glennie, and David C. Finnegan. 2013. “Calibration of a Terrestrial Full Waveform Laser Scanner.” *ASPRS 2013 Annual Conference Proceedings*, p. 7.

Hartzell, Preston J., Craig L. Glennie, Kivanc Biber, and Shuhab Khan. 2014. “Application of multispectral LiDAR to automated virtual outcrop geology.” *ISPRS Journal of Photogrammetry and Remote Sensing* 88. International Society for Photogrammetry; Remote Sensing, Inc. (ISPRS): 147–55. doi:10.1016/j.isprsjprs.2013.12.004¹⁸.

Heid, Torborg. 2011. “Deriving glacier surface velocities from repeat optical images.”

Heid, Torborg, and A. Kääh. 2012. “Evaluation of existing image matching methods for deriving glacier surface displacements globally from optical satellite imagery.” *Remote Sensing of Environment* 118. Elsevier Inc.: 339–55. doi:10.1016/j.rse.2011.11.024¹⁹.

Hensley, Scott, Scott Shaffer, Louise Veilleux, Paul Rosen, Sang-ho Yun, Sean Buckley, Adam Freeman, Dragana Perkovic-martin, Eric Gurrola, M Chakraborty, T Misra, R V Sagi, R Satish, R Bhan, and D Putrevu. 2014. “Science performance of the NASA-ISRO spaceborne SAR mission concept.” *Proceedings of IEEE International Geoscience and Remote Sensing Symposium (IGARSS'14)*,

¹⁶<https://doi.org/10.1016/j.isprsjprs.2005.02.006>

¹⁷<https://doi.org/10.3189/2015JoG15J031>

¹⁸<https://doi.org/10.1016/j.isprsjprs.2013.12.004>

¹⁹<https://doi.org/10.1016/j.rse.2011.11.024>

1–4.

Herzfeld, U C, B W McDonald, B F Wallin, T A Neumann, T Markus, A Brenner, and C Field. 2013. “Algorithm for Detection of Ground and Canopy Cover in Micropulse Photon-Counting Lidar Altimeter Data in Preparation for the ICESat-2 Mission.” *IEEE Transactions in Geoscience and Remote Sensing* 52 (4): 2109–25. doi:10.1109/TGRS.2013.2258350²⁰.

Howat, Ian M., Ian Joughin, and Theodore A. Scambos. 2007. “Rapid changes in ice discharge from Greenland outlet glaciers.” *Science (New York, N.Y.)* 315 (5818): 1559–61. doi:10.1126/science.1138478²¹.

Howat, Ian M., Ian Joughin, Slawek Tulaczyk, and S. Gogineni. 2005. “Rapid retreat and acceleration of Helheim Glacier, east Greenland.” *Geophysical Research Letters* 32: 1–4.

IPCC, Intergovernmental Panel On Climate Change. 2007. “Climate Change 2007 - The Physical Science Basis: Working Group I Contribution to the Fourth Assessment Report of the IPCC.” *Science*, 1009. <http://www.amazon.com/Climate-Change-2007-Contribution-Assessment/dp/0521880092>.

Joughin, Ian. 2002. “Ice-sheet velocity mapping: a combined interferometric and speckle-tracking approach.” *Annals of Glaciology* 34 (1): 195–201. doi:10.3189/172756402781817978²².

Juan, Julia De, Pedro Elósegui, Meredith Nettles, Tine B. Larsen, James L. Davis, Gordon S. Hamilton, Leigh A. Stearns, Morten L. Andersen, Göran Ekström, Andreas P. Ahlström, Lars Stenseng, S. Abbas Khan, and René Forsberg. 2010. “Sudden increase in tidal response linked to calving and acceleration at a large Greenland outlet glacier.” *Geophysical Research Letters* 37 (12). Wiley Blackwell (John Wiley & Sons).

Kennett, Michael, and Trond Eiken. 1997. “Airborne measurement of glacier surface elevation by scanning laser altimeter.” *Annals of Glaciology* 24: 293–96.

Kusari, Arpan. 2015. “Precise Registration of Laser Mapping Data by Planar Feature Extraction for Deformation Mapping.” PhD thesis, University of Houston.

Lichti, Dd Derek D, and Sj Gordon. 2004. “Error propagation in directly georeferenced terrestrial laser scanner point clouds for cultural heritage recording.” *Proc. of FIG Working Week, Athens, Greece*,

²⁰<https://doi.org/10.1109/TGRS.2013.2258350>

²¹<https://doi.org/10.1126/science.1138478>

²²<https://doi.org/10.3189/172756402781817978>

May, 1–16. <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Error+Propagation+in+Directly>

Lillesand, Thomas, Ralph W Kiefer, and Jonathan Chipman. 2014. *Remote sensing and image interpretation*. John Wiley & Sons.

Lucchitta, B K, and H M Ferguson. 1986. “Antarctica: measuring glacier velocity from satellite images.” *Science (New York, N.Y.)* 234 (4780): 1105–8. doi:10.1126/science.234.4780.1105²⁴.

Maier-Hein, L, A Franz, T dos Santos, M Schmidt, M Fangerau, H Meinzer, and J Fitzpatrick. 2011. “Convergent Iterative Closest-Point Algorithm to Accomodate Anisotropic and Inhomogenous Localization Errors” PP (99): 1. doi:10.1109/TPAMI.2011.248²⁵.

Mallet, Clément, and Frédéric Bretar. 2009. “Full-waveform topographic lidar: State-of-the-art.” *ISPRS Journal of Photogrammetry and Remote Sensing* 64 (1). International Society for Photogrammetry; Remote Sensing, Inc. ISPRS: 1–16. doi:10.1016/j.isprsjprs.2008.09.007²⁶.

Marino, Richard M., and William R. Davis. 2005. “Jigsaw: a foliage-penetrating 3D imaging laser radar system.” *Lincoln Laboratory Journal* 15 (1): 23–36.

Measures, Raymon M. 1985. “Laser Remote Sensing: Fundamentals and Applications.” *Eos, Transactions American Geophysical Union* 66 (40): 686. doi:10.1029/EO066i040p00686-05²⁷.

Michel, Remi, and Eric Rignot. 1999. “Flow of Glaciar Moreno, Argentina, from repeat-pass Shuttle Imaging Radar images: Comparison of the phase correlation method with radar interferometry.” *Journal of Glaciology* 45 (149).

Moon, Twila, Ian Joughin, Benjamin E. Smith, Michiel R. van den Broeke, Willem Jan van de Berg, Brice Noël, and Mika Usher. 2014. “Distinct patterns of seasonal Greenland glacier velocity.” *Geophysical Research Letters* 41 (20): 7209–16. doi:10.1002/2014GL061836²⁸.

Morariu, Vlad I VI, BV Balaji Vasani Srinivasan, Vikas Chandrakant Raykar, Ramani Duraiswami, and Larry S Davis. 2009. “Automatic online tuning for fast Gaussian summation.” *Advances*

²³<http://scholar.google.com/scholar?hl=en&&btnG=Search&&q=intitle:Error+Propagation+in+Directly+Georeferenced+Terrestrial+Laser+Scanner+Point+Clouds+for+Cultural+Heritage+RecordingD0>

²⁴<https://doi.org/10.1126/science.234.4780.1105>

²⁵<https://doi.org/10.1109/TPAMI.2011.248>

²⁶<https://doi.org/10.1016/j.isprsjprs.2008.09.007>

²⁷<https://doi.org/10.1029/EO066i040p00686-05>

²⁸<https://doi.org/10.1002/2014GL061836>

in *Neural . . .*, 1–8.

Myronenko, Andriy, and Xubo Song. 2010. “Point set registration: coherent point drift.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32 (12). Institute of Electrical & Electronics Engineers (IEEE): 2262–75.

Nettles, Meredith, Tine B. Larsen, Pedro Elósegui, Gordon S. Hamilton, Leigh A. Stearns, Andreas P. Ahlstrøm, James L. Davis, Morten L. Andersen, J. De Juan, S. Abbas Khan, L. Stenseng, Göran Ekström, and René Forsberg. 2008. “Step-wise changes in glacier flow speed coincide with calving and glacial earthquakes at Helheim Glacier, Greenland.” *Geophysical Research Letters* 35.

Nissen, Edwin, Aravindhan K. Krishnan, J. Ramón Arrowsmith, and Srikanth Saripalli. 2012. “Three-dimensional surface displacements and rotations from differencing pre- and post-earthquake LiDAR point clouds.” *Geophysical Research Letters* 39: 1–6.

OpenMP. 2013. “OpenMP Application Program Interface,” no. July: 320.

Pan, Zhigang, Craig L. Glennie, Preston J. Hartzell, Juan Carlos Fernandez-diaz, Carl J. Legleiter, and Brandon T. Overstreet. 2015. “Performance Assessment of High Resolution Airborne Full Waveform LiDAR for Shallow River Bathymetry,” 5133–59. doi:10.3390/rs70505133²⁹.

Parizek, Byron R., and Richard B. Alley. 2004. “Implications of increased Greenland surface melt under global-warming scenarios: Ice-sheet simulations.” *Quaternary Science Reviews* 23: 1013–27.

“PDAL - Point Data Abstraction Library.” <http://www.pdal.io/>.

Pfeffer, W Tad, J T Harper, and Shad O’Neel. 2008. “Kinematic constraints on glacier contributions to 21st-century sea-level rise.” *Science* 321 (5894). American Association for the Advancement of Science: 1340–3.

Prokop, Alexander, and H. Panholzer. 2009. “Assessing the capability of terrestrial laser scanning for monitoring slow moving landslides.” *Natural Hazards and Earth System Science* 9 (6): 1921–8.

Raffel, Markus, Christian E Willert, Jürgen Kompenhans, and Others. 2013. *Particle image*

²⁹<https://doi.org/10.3390/rs70505133>

velocimetry: a practical guide. Springer.

Rankin, G W, and K Sridhar. 1997. "A New Particle Image Velocimetry Technique for Three-dimensional Flows," no. 510.

Raykar, Vikas Chandrakant, Changjiang Yang, and Ramani Duraiswami. 1991. "Fast computation of sums of Gaussians in high dimensions."

Rees, Eric van. 2015. "The first multispectral airborne lidar sensor." *GeoInformatics* 18 (1): 10.

Riegl LMS GmbH. 2012. "3D Terrestrial Laser Scanner, RIEGL VZ-4000 / RIEGL VZ-6000." *General Description and Data Interfaces*. RIEGL Laser Measurement Systems, 1–493.

Rignot, Eric. 1996. "Tidal motion, ice velocity and melt rate of Petermann Gletscher, Greenland, measured from radar interferometry." *Journal of Glaciology* 42 (142).

Rignot, Eric, and Pannir Kanagaratnam. 2006. "Changes in the velocity structure of the Greenland Ice Sheet." *Science (New York, N. Y.)* 311: 986–90.

Rignot, Eric, Isabella Velicogna, Michiel R. van den Broeke, A. Monaghan, and J. Lenaerts. 2011. "Acceleration of the contribution of the Greenland and Antarctic ice sheets to sea level rise." *Geophysical Research Letters* 38 (5).

Rusinkiewicz, S., and M. Levoy. 2001. "Efficient variants of the ICP algorithm." *Proceedings of International Conference on 3-D Digital Imaging and Modeling, 3DIM 2001-Janua*: 145–52. doi:10.1109/IM.2001.924423³⁰.

Rusu, Radu Bogdan, and Steve Cousins. 2011. "3D is here: Point Cloud Library (PCL)." *2011 IEEE International Conference on Robotics and Automation*, May. Ieee, 1–4. doi:10.1109/ICRA.2011.5980567³¹.

Ryan, J. C., A. L. Hubbard, J. E. Box, J. Todd, P. Christoffersen, J. R. Carr, T. O. Holt, and N. Snooke. 2015. "UAV photogrammetry and structure from motion to assess calving dynamics at Store Glacier, a large outlet draining the Greenland ice sheet." *Cryosphere* 9 (1): 1–11. doi:10.5194/tc-9-1-

³⁰<https://doi.org/10.1109/IM.2001.924423>

³¹<https://doi.org/10.1109/ICRA.2011.5980567>

2015³².

Scambos, Theodore A., Melanie J. Dutkiewicz, Jeremy C. Wilson, and Robert A. Bindschadler. 1992. "Application of image cross-correlation to the measurement of glacier velocity using satellite image data." *Remote Sensing of Environment* 42 (3): 177–86. doi:10.1016/0034-4257(92)90101-O³³.

Schwalbe, Ellen, Hg G Maas, Reinhard Dietrich, and H Ewert. 2008. "Glacier velocity determination from multi temporal terrestrial long range laser scanner point clouds." *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 37: 457–62.

Shan, Jie, and Charles K. Toth. 2008. *TOPOGRAPHIC LASER RANGING Principles and Processing*.

Skaloud, Jan, and Derek D Lichti. 2006. "Rigorous approach to bore-sight self-calibration in airborne laser scanning." *ISPRS Journal of Photogrammetry and Remote Sensing* 61: 47–59. doi:10.1016/j.isprsjprs.2006.07.003³⁴.

Stearns, Leigh A., and Gordon S. Hamilton. 2007. "Rapid volume loss from two East Greenland outlet glaciers quantified using repeat stereo satellite imagery." *Geophysical Research Letters* 34.

Teza, G., A. Galgaro, N. Zaltron, and R. Genevois. 2007. "Terrestrial laser scanner to detect landslide displacement fields: a new approach." *International Journal of Remote Sensing* 28 (16): 3425–46. doi:10.1080/01431160601024234³⁵.

"The BSD 3-Clause License." <https://opensource.org/licenses/BSD-3-Clause>.

Thielicke, William, and Eize J. Stamhuis. 2014. "PIVlab - Towards User-friendly, Affordable and Accurate Digital Particle Image Velocimetry in MATLAB." *Journal of Open Research Software* 2 (1): e30. doi:10.5334/jors.bl³⁶.

Turner, Darren, Arko Lucieer, and Christopher Watson. 2012. "An automated technique for generating georectified mosaics from ultra-high resolution Unmanned Aerial Vehicle (UAV) imagery, based on Structure from Motion (SFM) point clouds." *Remote Sensing* 4 (5): 1392–1410.

³²<https://doi.org/10.5194/tc-9-1-2015>

³³[https://doi.org/10.1016/0034-4257\(92\)90101-O](https://doi.org/10.1016/0034-4257(92)90101-O)

³⁴<https://doi.org/10.1016/j.isprsjprs.2006.07.003>

³⁵<https://doi.org/10.1080/01431160601024234>

³⁶<https://doi.org/10.5334/jors.bl>

doi:10.3390/rs4051392³⁷.

Ullrich, Andreas, Roland Schwarz, and Helmut Kager. 2003. *Using hybrid multi-station adjustment for an integrated camera laser-scanner system*. na.

Wagner, Wolfgang. 2010. “Radiometric calibration of small-footprint full-waveform airborne laser scanner measurements: Basic physical concepts.” *ISPRS Journal of Photogrammetry and Remote Sensing* 65 (6). Elsevier B.V.: 505–13. doi:10.1016/j.isprsjprs.2010.06.007³⁸.

Wang, Hongzhou, and Craig L. Glennie. 2015. “Fusion of waveform LiDAR data and hyperspectral imagery for land cover classification.” *ISPRS Journal of Photogrammetry and Remote Sensing* 108. International Society for Photogrammetry; Remote Sensing, Inc. (ISPRS): 1–11. doi:10.1016/j.isprsjprs.2015.05.012³⁹.

Westerweel, J. 1997. “Fundamentals of digital particle image velocimetry.” *Measurement Science and Technology* 8 (12): 1379–92. doi:10.1088/0957-0233/8/12/002⁴⁰.

Whillans, I. M., and Robert A. Bindschadler. 1988. “Mass balance of Ice Stream B. West Antarctica.” *Annals of Glaciology* 11: 187–93.

Wilson, Margaret F. J., Brian O’Connell, Colin Brown, Janine C. Guinan, and Anthony J. Grehan. 2007. *Multiscale Terrain Analysis of Multibeam Bathymetry Data for Habitat Mapping on the Continental Slope*. Vol. 30. 1-2. doi:10.1080/01490410701295962⁴¹.

Yamaguchi, Yasushi, Anne B. Kahle, Hiroji Tsu, Toru Kawakami, and Moshe Pniel. 1993. “The Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER)” 36 (4): 1062–71. doi:10.1002/ima.1850030210⁴².

Yuille, Alan L., and Norberto M. Grzywacz. 1988. “The Motion Coherence Theory.”

Zhang, Xiao, Craig L. Glennie, and Arpan Kusari. 2015. “Change detection from differential airborne LiDAR using a weighted anisotropic iterative closest point algorithm.” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 8 (7): 3338–46.

³⁷<https://doi.org/10.3390/rs4051392>

³⁸<https://doi.org/10.1016/j.isprsjprs.2010.06.007>

³⁹<https://doi.org/10.1016/j.isprsjprs.2015.05.012>

⁴⁰<https://doi.org/10.1088/0957-0233/8/12/002>

⁴¹<https://doi.org/10.1080/01490410701295962>

⁴²<https://doi.org/10.1002/ima.1850030210>

doi:10.1109/JSTARS.2015.2398317⁴³.

⁴³<https://doi.org/10.1109/JSTARS.2015.2398317>

This page intentionally left blank.