TAMPEREEN TEKNILLINEN YLIOPISTO
TAMPERE UNIVERSITY OF TECHNOLOGY

TERO SOININEN

DATA-DRIVEN APPROACH TO SATELLITE SELECTION IN GNSS

Master of Science thesis

Examiners: D.Sc. Paula Syrjärinne,
D.Sc. Simo Ali-Löytty
Examiners and topic approved by the
Dean of the Faculty of
Computing and Electrical Engineering
on 9th August 2017

ABSTRACT

TERO SOININEN: Data-driven approach to satellite selection in GNSS
Tampere University of Technology
Master of Science thesis, 46 pages, 0 Appendix pages
October 2017
Master's Degree Programme in Information Technology
Major: Engineering Mathematics
Examiners: D.Sc. Paula Syrjärinne, D.Sc. Simo Ali-Löytty
Keywords: data science, GNSS, machine learning

The main goal of this work was to develop an algorithm for multi-constellation GNSS receivers that would select satellites out of the tracked ones to be used in the location solution. As the receiver has very limited computational resources, the complexity of the algorithm needed to be kept low.

The work began by exploratory analysis of GNSS data. This analysis gave insight into the differences of the various satellite navigation systems as well as into the nature of the pseudorange residuals. These observations helped in shaping the algorithm that we proposed for the problem of satellite selection. The algorithm itself was developed using data science techniques to filter out bad pseudorange measurements and borrowed some earlier ideas to optimize the geometric dilution of precision of the solution set as well.

The approach we chose was shown to work very well when applied to real data measured from road tests in varying surroundings. Even with practically non-existent parameter tuning the algorithm was able to spot almost 90% of the bad pseudorange measurements, keeping the specificity, i.e., ability to hold on to the good measurements at over 90% level.

The ability to filter out bad pseudorange measurements translated to improved location accuracy as well. All in all, the results achieved in this work proved encouraging enough to begin implementing the algorithm in actual receiver software to study the performance of the data-driven approach in action.

TIIVISTELMÄ

TERO SOININEN: Datalähtöinen lähestymistapa satelliittien valintaan globaaleissa satelliittipaikannusjärjestelmissä
Tampereen teknillinen yliopisto
Diplomityö, 46 sivua, 0 liitesivua
Lokakuu 2017
Tietotekniikan koulutusohjelma
Pääaine: Teknillinen matematiikka
Tarkastajat: D.Sc. Paula Syrjärinne, D.Sc. Simo Ali-Löytty
Avainsanat: datatiede, GNSS, koneoppiminen

Työn päällimmäisenä tavoitteena oli kehittää algoritmi paikanlaskennassa käytettävien satelliittien valintaan GNSS -vastaanottimissa, jotka kykenevät useamman järjestelmän seurantaan. Koska vastaanottimien laskentateho on varsin rajallinen täytyi algoritmin vaatima laskenta-aika pitää alhaisena.

Työ alkoi GNSS -datan tutkimuksella. Tämän myötä saatiin yleiskatsaus suorityskykyeroihin eri GNSS -järjestelmien välillä ja tyypillisiin pseudoetäisyyden virheisiin. Näiden näkemysten avulla luotiin algoritmi, joka hyödyntää data-analyysin tekniikoita suurten mittavirheiden poistamiseksi käytettävissä olevien mittausten joukosta.

Kehitetyn algoritmin osoitettiin toimivan hyvin normaaleissa GNSS -vastaanottimen käyttöympäristöissä ja olevan vaaditulta laskenta-ajaltaan erittäin kilpailukykyinen. Kokonaisuudessaan työn tulokset osoittautuivat niin rohkaiseviksi, että algoritmi tullaan toteuttamaan GNSS -vastaanottimen ohjelmistoon, jotta pääsemme varmistumaan sen tarkoituksenmukaisesta toiminnasta.

PREFACE

Looking back now I can say it was indeed a good idea to apply for the Master's thesis position at u-blox! There has been hardly a single day I wouldn't have enjoyed the atmosphere in the Tampere office, thanks to the top co-workers here.

For the thesis work opportunity I would like to thank Christoph Schmid who also provided valuable insight into the problem dealt in this work. I'd also like to express my deepest gratitude to my supervisor, Paula Syrjärinne, who made the whole work happen. Her contributions for this thesis are too numerous to be listed here. In addition I have to thank Simo Ali-Löytty, the conversations we've had shaped the work to be much more than it would have been without.

Last, I would like to thank everyone who bothers to read my work. It brings paramount joy to me to know that there are at least a few of you out there! I hope it also gives something back to you as it must take a toll.

Tampere, 24.10.2017

Tero Soininen

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND SYMBOLS

BDS	BeiDou navigation satellite system, Chinese global navigation satellite system
cno	Carrier-to-noise-density ratio
CODE	Center for orbit determination in Europe
DDA	Data-driven approach for satellite selection
DOP	Dilution of precision
GAL	Galileo, European global navigation satellite system
GDOP	Geometric dilution of precision
GLO	Globalnaja navigatsionnaja sputnikovaja sistema or GLONASS, Russian global navigation satellite system
GNSS	Global navigation satellite system, a general term for satellite navigation systems with global coverage
GPS	Global positioning system, United States' global navigation satellite system
ECEF	Earth-centered, earth-fixed coordinate system
elev	Elevation
MGEX	The multi-GNSS experiment and pilot project
QO	Quasi-optimal algorithm for satellite selection
QZSS	Quasi-zenith satellite system, regional satellite system in Japan
RAIM	Receiver autonomous integrity monitoring
SBAS	Satellite-based augmentation system
WGS 84	World geodetic system 84, latest revision of the world geodetic system

c	speed of light in vacuum, 299 792 458 m/s
diag	diagonal elements of a matrix
E	expected value
$e.$	error term, for example, $e_{\hat{p}}$ denotes the pseudorange errors
J	Jacobian matrix
$\mathcal{O}$	Big o notation
p	vector of calculated distances to satellites or probability from a logistic regression model
$p.$	distance to a single satellite
$\hat{p}$	vector of pseudoranges from the satellites
$\hat{p}.$	pseudorange from a single satellite
Pr	probability

S	set of visible satellites
$S_.$	a subset of the visible satellites
$s_.$	location of a satellite in ECEF, i.e., $[x_., y_., z_.]^T$
σ^2	variance
Σ	covariance matrix
t	GPS time
t_u	receiver clock offset
tr	trace of a matrix
u	receiver location in ECEF, i.e., $[x_u, y_u, z_u]^T$ or augmented with the clock offset term $[x_u, y_u, z_u, ct_u]^T$
$\hat{u}$	a least squares optimal solution for the receiver location in ECEF and clock offset error
w	coefficient weights from logistic regression
w_0	intercept term from logistic regression
X	set of independent variables in logistic regression
x	a realization of X
$x_.$	x-coordinate in ECEF, e.g., x_u denotes receiver x-coordinate in ECEF
Y	the dependent variable in logistic regression
y	a realization of Y
$y_.$	y-coordinate in ECEF, e.g., y_u denotes receiver y-coordinate in ECEF
$z_.$	z-coordinate in ECEF, e.g., z_u denotes receiver z-coordinate in ECEF
$\lvert \cdot \rvert$	set cardinality
$\lVert \cdot \rVert_2$	l^2 -norm

1. INTRODUCTION

Satellite systems are an integral part of positioning and navigation today. They have been in use since the 1960s when the United States military deployed the Transit system. Nowadays there are multiple Global Navigation Satellite Systems (GNSS), as well as a few regional ones. Most of the systems are made available also to civilian users. The global systems are designed so that when used alone they provide enough visible satellites for a positioning solution. Thus when using receivers that are able to track satellites from multiple systems there are often more visible satellites available than it is feasible to use. This leads to the question of selecting the optimal subset of visible satellites for the navigation solution.

The topic has properly risen only in the recent years as the BeiDou-2 and Galileo systems have not been operational for too long, and are still not in full operational capability. Most of the former work concerning the satellite selection has mainly focused on minimizing the Geometric Dilution of Precision (GDOP). GDOP depicts the multiplicative effect of the satellite geometry on the position error, i.e.,

$$\text{GDOP} = \frac{\text{location error}}{\text{measurement error}}. \tag{1.1}$$

Roughly speaking, if the satellites used are in the same direction then the GDOP is usually large and inflates the position error more. Whereas if the satellites are well spaced then the GDOP, and the positional error are both smaller.

However satellite geometry works only to inflate the initial measurement errors. The actual sources of this error include signal transmission delays caused by troposphere and ionosphere, satellite location inaccuracies and multipath caused by signal reflections on its way from the satellite to the receiver. As there are more things to consider than just the geometry when selecting the satellites a more holistic approach is wanted.

So the main goal of this work is an algorithm to minimize the location error without a considerable computational complexity. To start building towards this more comprehensive approach, we need to consider certain questions first:

- How big of a role does GDOP have in the location error? How about the measurement errors?
- What are the most important factors contributing to the measurement errors and how to use them in the satellite selection?
- How does the number of satellites affect the solution accuracy?
- Are there differences in the accuracy of the GNSS constellations?

In the following chapters we are looking to answer these questions, as well as building a background for the algorithm. The work is structured so that Chapter 2 gives a brief background into the topics deemed necessary for understanding the work. Chapter 3 explores the previous studies done concerning satellite selection and proposes a new, data driven approach for this problem. In Chapter 4 the results acquired in this work are presented. And finally, Chapter 5 discusses these results and concludes this thesis with projection into the future work on the subject.

2. BACKGROUND

This chapter, starting with Section 2.1, gives a general overview of modern Global Navigation Satellite Systems. Keeping the scope of this thesis and the present state of multiple working constellations in mind the description does not delve into too fine a detail. A more thorough explanation can be found in [20] for example. Section 2.2 formulates the problem confronted in this work and Section 2.3 briefly describes the methods selected to tackle this.

2.1 Global navigation satellite systems

The satellite navigation systems can be divided into three segments: the space segment, ground (control) segment and user segment. The space segment consists of the satellites in orbit, the numbers of which are given in Table 2.1 for the different global systems at the time of writing [2, 1]. The ground segment refers to the control and monitor stations, tracking and maintaining the satellites. User segment denotes the user receivers that process received satellite signals and perform needed operations, e.g., navigation and timing. [20]

	BDS	GAL	GLO	GPS	Σ
Number of operational satellites	15	15	23	31	84
Nominal number of satellites	35	30	24	24	113

Table 2.1 *The number of currently operational satellites (as of October 3rd 2017) and the planned nominal number of satellites for the global navigation satellite systems [2, 1].*

Modern global satellite navigation systems, GPS [26], GLONASS [29], Galileo [13] and BeiDou-2 [7], are based on time of flight measurements from multiple satellite vehicles. The navigation message broadcast by the satellites contain orbit parameters and time information that allow the user to compute the satellite location at the instant of the transmission. Now, knowing that the carrier signal travels at the speed of light the receiver can calculate the distance to the satellite by comparing the broadcast transmission and arrival times. Obtaining the distances to multiple satellites allows the receiver to trilaterate ones position.

As the receiver time and the position coordinates are unknown there are four variables that need to be solved. This means that measurements to at least four satellites are required. The satellites on the other hand carry highly accurate atomic clocks on board and the navigation message contains satellite clock bias and drift parameters that are used to correct the satellite time information. The ground based control stations also monitor the satellite positions and update them so that the ephemeris information stays accurate.

GNSS position calculation

As stated earlier, positioning in GNSS is achieved using signal time of flight measurements. These measurements give an approximation of the distances to the satellites. However, as always, there is error involved. The satellite ephemerises, satellite clocks, receiver noise, ionosphere, troposphere and multipath all induce errors in these measurements. The measured distances are referred to as pseudoranges, $\hat{p}$, rather than ranges as they contain the receiver clock bias term in addition to the distance.

Now, given some location $u \in \mathbb{R}^3$ and a location $s_i \in \mathbb{R}^3$ for a satellite vehicle i in earth-centered, earth-fixed (ECEF) coordinates, the distance between these points is the Euclidean norm

$$p_i(u) = \|s_i - u\|_2. \tag{2.1}$$

So the pseudorange to satellite i can also be given in form:

$$\hat{p}_i = p_i(u) + ct_u + e_{\hat{p}_i}, \tag{2.2}$$

where c is the speed of light, t_u the receiver clock offset and $e_{\hat{p}_i}$ denotes the error, or residual, caused by the aforementioned sources. The receiver clock offset is considered a variable that is solved jointly with the receiver location and thus does not induce error. For future reference let us now augment u with the receiver clock bias term, i.e.,

$$u = \begin{bmatrix} x_u \\ y_u \\ z_u \\ ct_u \end{bmatrix}. \tag{2.3}$$

The problem of locating a GNSS receiver boils down to minimizing the difference of the measured pseudoranges and the distances calculated from the receiver position and clock offset in the least squares sense

$$\|\hat{p} - p(u)\|_2. \tag{2.4}$$

As we have four unknowns we need a system of at least four equations in order to solve them, i.e., measurements to four or more satellites.

There are different techniques for solving this system of nonlinear equations, for example, closed form solutions [22, 15] or iterative techniques like Gauss-Newton [10]. In this work we go through the linearization based iterative technique which yields the least squares solution for the location. For a more thorough explanation see [20].

Now assume there is a location and time offset $\hat{u}$ which minimizes (2.4) in the least squares sense. This location can be given as

$$\hat{u} = u + \Delta u \tag{2.5}$$

so the pseudoranges can be calculated as

$$\hat{p} = p(\hat{u}) = p(u + \Delta u). \tag{2.6}$$

Linearizing this using a first order Taylor series yields

$$\hat{p} \approx p(u) + J_{p(u)}\Delta u, \tag{2.7}$$

where $J_{p(u)}$ is the Jacobian for $p(u)$, i.e.,

$$J_{p(u)} = \begin{bmatrix} \frac{\partial p(u)}{\partial x_u} & \frac{\partial p(u)}{\partial y_u} & \frac{\partial p(u)}{\partial z_u} & \frac{\partial p(u)}{\partial ct_u}, \end{bmatrix} = \begin{bmatrix} -\frac{x_1 - x_u}{p_1(u)} & -\frac{y_1 - y_u}{p_1(u)} & -\frac{z_1 - z_u}{p_1(u)} & 1 \\ & \cdots & & \\ -\frac{x_n - x_u}{p_n(u)} & -\frac{y_n - y_u}{p_n(u)} & -\frac{z_n - z_u}{p_n(u)} & 1 \end{bmatrix}. \tag{2.8}$$

Reorganizing (2.7) one obtains

$$\hat{p} - p(u) \approx J_{p(u)}\Delta u \tag{2.9}$$

for which the least squares minimal solution is given by

$$\Delta u = (J_{p(u)}^T J_{p(u)})^{-1} J_{p(u)}^T (\hat{p} - p(u)). \tag{2.10}$$

This leads to an iterative solution, where an initial try is updated according to (2.5) and (2.10), eventually resulting in the location and clock offset which minimizes the difference of the measured pseudoranges and the calculated ones in the least squares sense.

2.2 Position accuracy and problem formulation

Since the position solution is calculated from the pseudorange measurements, the errors in the measurements propagate to the solution as well. Now, say we have some other source for the true location of the receiver, u. Then we can obviously calculate the pseudorange residuals after solving for the time offset

$$e_{\hat{p}} = \hat{p} - p(u), \tag{2.11}$$

but we can also calculate how much error the residuals would cause in the location solution. One way would be to go through the whole iteration process described above, however that can be rather tedious. Another would be to begin with the known location and create the Jacobian for that. Then the error in the solution, $e_{\hat{u}}$, can be acquired by multiplying the pseudorange residuals with the solution matrix

$$e_{\hat{u}} = (J_{p(u)}^T J_{p(u)})^{-1} J_{p(u)}^T e_{\hat{p}}. \tag{2.12}$$

Now for the sake of simplicity let us consider the pseudorange errors to be iid random variables from a Gaussion distribution with zero mean. Then following (2.12) the location errors are also Gaussion with zero mean, since the Jacobian, and hence the whole $(J_{p(u)}^T J_{p(u)})^{-1} J_{p(u)}^T$ is considered fixed. So calculating the covariance of $e_{\hat{u}}$ gives

$$\begin{aligned} \Sigma(e_{\hat{u}}) = \mathrm{E}(e_{\hat{u}} e_{\hat{u}}^T) &= \mathrm{E}((J_{p(u)}^T J_{p(u)})^{-1} J_{p(u)}^T e_{\hat{p}} e_{\hat{p}}^T J_{p(u)} (J_{p(u)}^T J_{p(u)})^{-1}) \\ &= (J_{p(u)}^T J_{p(u)})^{-1} J_{p(u)}^T \, \Sigma(e_{\hat{u}}) \, J_{p(u)} (J_{p(u)}^T J_{p(u)})^{-1}. \end{aligned} \tag{2.13}$$

Note in the previous that $(J_{p(u)}^T J_{p(u)})^{-1}$ is symmetric. Now due to the iid assumption

$\Sigma(e_{\hat{p}}) = I\sigma_{\hat{p}}^2$, where $\sigma_{\hat{p}}^2$ is the variance of the residuals. Thus

$$\Sigma(e_{\hat{u}}) = (J_{p(u)}^T J_{p(u)})^{-1}\sigma_{\hat{p}}^2. \tag{2.14}$$

In the covariance matrix the diagonal holds the variance values for the location errors

$$\mathrm{diag}(\Sigma(e_{\hat{u}})) = \begin{bmatrix} \sigma_{x_u}^2 & \sigma_{y_u}^2 & \sigma_{z_u}^2 & \sigma_{ct_u}^2 \end{bmatrix}. \tag{2.15}$$

Now remembering back for the definition of GDOP (1.1) we get

$$\mathrm{GDOP} = \frac{\sqrt{\sigma_{x_u}^2 + \sigma_{y_u}^2 + \sigma_{z_u}^2 + \sigma_{ct_u}^2}}{\sigma_{\hat{p}}} \tag{2.16}$$

and since $\mathrm{tr}(\Sigma(e_{\hat{u}})) = \mathrm{tr}((J_{p(u)}^T J_{p(u)})^{-1}\sigma_{\hat{p}}^2$ it follows that

$$\mathrm{GDOP} = \sqrt{\mathrm{tr}((J_{p(u)}^T J_{p(u)})^{-1})}. \tag{2.17}$$

The bold assumption of the pseudorange residuals being iid from a zero mean Gaussian distribution was made to simplify the derivation of the GDOP. This assumption does not hold in reality, which can be seen in the results of this work as well. However the GDOP proves to be of value nonetheless.

Problem formulation

Hence there are two factors playing into the location accuracy: the satellite location dependent geometric dilution of precision and the pseudorange measurement inaccuracies. The latter includes such as ionosphere, troposphere and multipath induced errors. Some of these can be diminished by applying relevant corrections in the receiver, but some are harder to mitigate. For example there are models like Klobuchar and NeQuick for taking into account the ionosphere induced errors [21, 11, 17]. And for the troposphere errors Hopfield and Saastamoinen models have been developed [18, 30]. In addition there are services like the satellite based augmentation systems (SBAS), which provide correction data for these kind of errors. The multipath delays however are often more problematic and can have greater

effects on the calculated location. These can be affected by the selection of the satellite measurements used in the solution.

The main problem of this work is thus an optimal subset selection one. We are looking for a subset of satellites from the set of all visible satellites that minimizes the error in the receiver location, $e_{\hat{u}}$, with possibly an added limit to the cardinality of this subset.

Let S now define the set of all visible satellites. Similarly let $S_i \subseteq S$ denote the subset selected for the location solution. Now we can formulate the problem as

$$\begin{aligned} \underset{S_i \subseteq S}{\text{minimize}} \quad & \|e_{\hat{u}}(S_i)\|_2 \\ \text{subject to} \quad & |S_i| \leq k \end{aligned} \tag{2.18}$$

where k is the maximum number of satellites allowed in the solution.

There are certain requirements for the selection algorithm. First, and foremost of these requirements is that the selection needs to be executable in the receiver hardware in real time. This rules out too complex solutions as well as brute force methods. Second, the selection criteria should preferably be a white box model, i.e., the selection criteria can be understood when inspecting the model.

As the interest now is in minimizing the location error, both the GDOP and pseudorange errors need to be considered. For the latter a data driven approach was selected in the form of logistic regression. How this was applied in the selection is described in Section 3.2, but a general overview of logistic regression is presented next.

2.3 Logistic regression

Logistic regression is a linear method for classification. Here we consider a binary output case, but logistic regression in general is by no means limited to just two outcome categories. The extension to multiclass cases can be achieved for example by a set of independent binary regressions. As the multinomial logistic regression is not in the scope of this work an interested reader can for example refer to [16] for a more thorough explanation.

Logistic regression works by creating probabilities for the dependent variable classes from the linear combinations of the prediction, or independent, variables. So given

random variables $Y \in \{0, 1\}$ and $X \in \mathbb{R}^n$, we are interested in modeling the posterior probability of

$$\Pr(Y = 1|X = x). \tag{2.19}$$

Here Y is called the dependent variable and X the prediction variable, for which x is some realization. In logistic regression the modeling is done by applying the logistic function,

$$\mu(t) = \frac{e^t}{e^t + 1} = \frac{1}{1 + e^{-t}}, \tag{2.20}$$

which has the useful property of $\mu(\mathbb{R}) = (0, 1)$, allowing the output to be interpreted as a probability. Using the logistic function the probability can be expressed by inserting $t = w_0 + w^T x$, where $w_0 \in \mathbb{R}$ and $w \in \mathbb{R}^n$, so

$$\Pr(Y = 1|X = x) = \mu(w_0 + w^T x) = \frac{e^{w_0 + w^T x}}{e^{w_0 + w^T x} + 1}. \tag{2.21}$$

And similarly

$$\Pr(Y = 0|X = x) = 1 - \Pr(Y = 1|X = x) = \frac{1}{e^{w_0 + w^T x} + 1}. \tag{2.22}$$

If we now look at the log-odds of the probability of "success"

$$\begin{aligned} \log \frac{\Pr(Y = 1|X = x)}{1 - \Pr(Y = 1|X = x)} &= \log \frac{\Pr(Y = 1|X = x)}{\Pr(Y = 0|X = x)} \\ &= w_0 + w^T x, \end{aligned} \tag{2.23}$$

we see that the decision boundary is a linear hyperplane defined by the affine function $w_0 + w^T x$.

Figure 2.1 shows an example of a logistic regression decision boundary in a two-dimensional case. The data for this figure is simulated and happens to be completely linearly separable. Often in real scenarios there is some overlap between the classes and a complete separation is not possible, at least with a linear boundary. A non-linear decision boundary is also possible with logistic regression by applying a so

called kernel trick in which the features are replaced with a kernel function. This is however not relevant to the current work, an interested reader can find more information about the kernel methods for example from [16, 25].

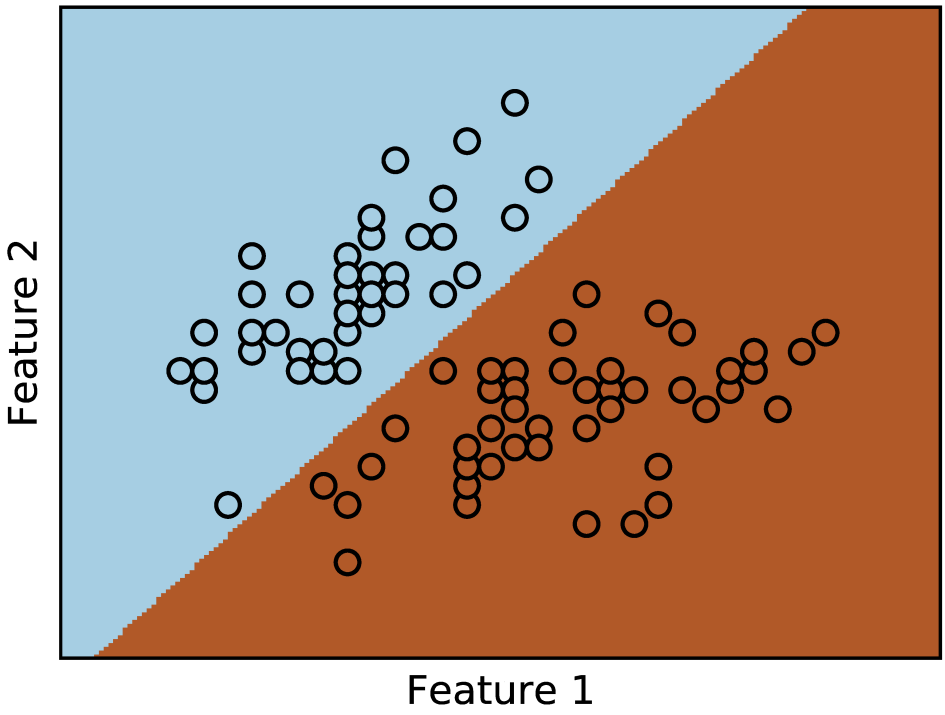

Figure 2.1 *An example of the linear decision boundary from a logistic regression model. The color of the background denotes into which class the samples falling to that area are to be classified and the color inside the circle the true class of the sample. In this case the samples are completely linearly separable.*

The logistic regression models are usually fitted in the maximum likelihood sense. If we now consider the realizations x_i to include a constant term 1 to accommodate the intercept term w_0 in w, then in the binary case

$$\begin{aligned}\mathcal{L}(w) &= \Pr(Y|X;w) \\ &= \prod_i \Pr(y_i|x_i;w) \\ &= \prod_i \left(\frac{e^{w^T x_i}}{e^{w^T x_i}+1}\right)^{y_i} \left(1-\frac{e^{w^T x_i}}{e^{w^T x_i}+1}\right)^{1-y_i}.\end{aligned} \tag{2.24}$$

However, the log-likelihood is often much easier to optimize, as is the case here as well. Now the log-likelihood is of the following form

$$\begin{aligned}\ell(w) &= \log \Pr(Y|X;w) \\ &= \sum_i y_i(w^T x_i) - \log(e^{w^T x_i} + 1).\end{aligned} \tag{2.25}$$

When the partial derivatives of this are set to zero the set of equations is non-linear and thus the solution is not entirely straightforward. There are multiple iterative algorithms for the solution, e.g. gradient descent and Newton-Raphson methods. More about the solutions can be found in [16, 19].

In the next chapter we take a look at the previous work done in the satellite selection domain and also introduce the new approach that exploits logistic regression in the selection process.

3. SATELLITE SELECTION

The issue of selecting satellites for positioning has emerged in the recent years. Before GPS had competing satellite navigation systems one could simply use all of the satellite vehicles in view for the solution. Nowadays with multi-constellation receivers the number of satellites visible is already so high that the position calculations may prove rather ponderous. Also, as shown in [35] using all of the visible satellites can yield poorer results than using a well chosen subset of those satellites. The previous points provoke the question of which satellite vehicles to use for the solution.

3.1 Previous work

Previous work concerning satellite selection in GNSS has been mainly about minimizing the effect of the satellite geometry. This has been approached in numerous ways. One can naturally go through all the possible subsets and select the optimal one but this produces huge computational load. Greedy algorithms performing backward elimination [24] or forward selection based on DOP [28] are obvious reliefs, but they still require multiple matrix inversions. To mitigate this alternative metrics have been proposed as well.

In [27] a quasi-optimal algorithm for the selection problem is proposed. The paper introduces a cost function for the satellites

$$J_j = \sum_{i=1}^{N} \cos(2\theta_{i,j}), \tag{3.1}$$

where $\theta_{i,j}$ is the angle between the line of sight vectors to satellites i and j. This stems from the idea that satellites with collinear line of sight vectors are redundant.Using this cost function the satellites are selected in a backward elimination manner, always removing the vehicle with the highest cost. The aforementioned algorithm does not often provide optimal satellite geometries but the computational load required makes it suitable for real time applications. Indeed, since the cosine term can be expressed as

$$\cos(2\theta_{i,j}) = 2\cos^2(\theta_{i,j}) - 1 = 2(v_i \cdot v_j)^2 - 1, \tag{3.2}$$

where $v_i \in \mathbb{R}^3$ and $v_j \in \mathbb{R}^3$ are the line of sight unit vectors to satellites i and j, the computation requires very low resources. The aforementioned paper has also inspired multiple other metrics using cosine for the ranking of the satellites [33, 31]. They incorporate additional ideas like requiring the highest elevation satellites to be included in the selection.

The idea of requiring the highest elevation satellites was first introduced in [23] as their investigations concluded that the optimal GDOP in four satellite case is obtained by having a satellite at zenith and the rest equally spaced at horizon. This was later expanded to n satellite cases in [34], where they select p satellites at the zenith and the rest $n-p$ satellites in the horizon. With p depending on the number of satellites to be selected.

Other strategies have also been developed in search for the best geometries. In [4] an algorithm is proposed in which those satellites are selected that span a convex hull around the visible satellite vehicles. Such algorithms have also been proposed that pick those subsets which most resemble some predefined geometries [6, 34, 14].

However the recent paper [35] exposed the shortcomings of algorithms that rely on minimizing DOP alone. Even though the GDOP would be optimal for a given size of satellite subsets, if there are satellites with very poor pseudorange measurements included, the location error can be quite significant. This is also depicted in the results of this work.

Indeed, when minimizing the location error, both the GDOP and measurement errors need to be considered. Thus we propose a more holistic approach for the satellite selection issue. This approach is presented in the next section.

3.2 Data-driven approach

The approach that was constructed for this work scores the satellites in two parts, the other tries to optimize the GDOP and the other aims at eliminating too large measurement errors. As calculating the GDOP can only happen after a subset of the visible satellites has been selected it is not practical to be computing it in the selection process. That is why we model the effect the satellites have on GDOP by something we call the redundancy score from now on. The redundancy scores depend on which satellites have already been picked in to the selection subset and

are only calculated to those satellites that have not been picked yet. The other part, the large measurement errors, we try to predict with the logistic regression model.

Since the measurement errors are not in anyway dependent on the satellites already picked for selection, we initialize the algorithm by calculating the probabilities that a pseudorange measurement is over a certain threshold. The logistic regression method for this was described in Section 2.3 on a general level. In this case, the probability, p, a satellite $j \in S$ has a bad pseudorange measurement is

$$p_j = \frac{1}{e^{w^T x_j} + 1}, \tag{3.3}$$

where x_j is information about the satellite and its signal. In this case x_j includes carrier-to-noise-density ratio, elevation and satellite system information.

Next the three best scoring satellites are added in to the selection subset S_i. So the first three satellites are selected purely based on their probability to have a good measurement. In the following selections also the redundancy score is taken into account. In the redundancy score part we borrow the idea of using line of sight vector cosines from the quasi-optimal algorithm [27].

The quasi-optimal algorithm has the nice idea of modeling the effect the satellite has on the GDOP by the cosines of the line of sight vectors. As the cosine can be easily computed via the dot product it is suitable for the requirement of low computational complexity present in this work as well. A major drawback, also in the GDOP sense, in the algorithm is that it leaves the highest elevation satellites out too often. Quick simulations showed that the GDOP-wise optimal subsets included the highest elevation satellite more than 90% of the time, when selecting more than 5 satellites. This investigation was inspired by the Zhang paper [34] about fast satellite selection, in which they describe GDOP-wise optimal satellite subsets and conclude that it is beneficial to always include high elevation satellites in the selection.

This problem is eliminated in our approach by the initial selection of the three satellites that have the least probability for large measurement errors, since those satellites often also sit at high elevations. Unlike $\cos(2\theta_{i,j})$ in the quasi-optimal algorithm a simple $\cos(\theta_{i,j})$ is used here. Figure 3.1 visualizes the differences of these redundancy cost functions. As can be seen in the aforementioned figure, the $\cos(2\theta)$ cost function penalizes satellites if the angle between them gets over 90°. This is not desirable in our scenarios and so the plain dot product is good enough in this case.

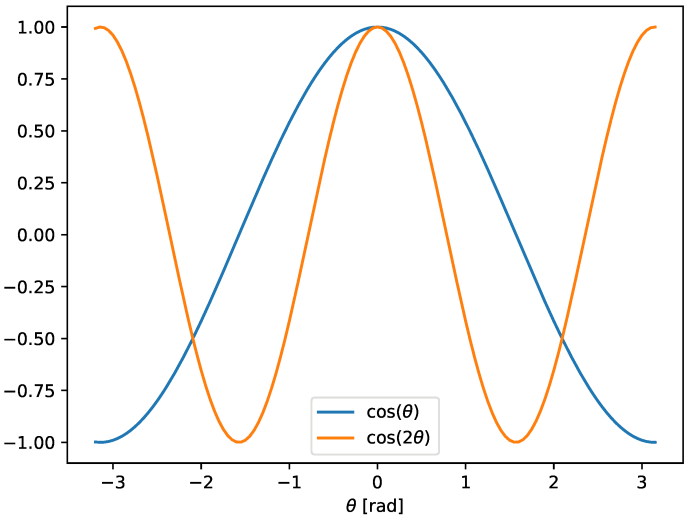

Figure 3.1 *Cost function for the satellite redundancy used in the quasi-optimal algorithm* ($\cos(2\theta)$) *in comparison with the cost function* ($\cos(\theta)$) *used in our approach.*

So now the satellites $j \in S \setminus S_i$ get a redundancy score, r_j, which we define with regards to the satellites already in the selection, $i \in S_i$. That is, the redundancy is now

$$\begin{aligned} r_j &= \frac{1}{|S_i|} \sum_{i=1}^{|S_i|} \cos(\theta_{i,j}) + 1 \\ &= \frac{1}{|S_i|} \sum_{i=1}^{|S_i|} v_i \cdot v_j + 1, \end{aligned} \tag{3.4}$$

where $v_\cdot$ are the line of sight unit vectors to the satellites.

As the redundancy score was defined as the mean, it now suitably lies in the interval $[0, 2]$. The end score, s_j, is now simply the sum

$$s_j = p_j + r_j/2. \tag{3.5}$$

Now the satellite with the lowest end score s_j is added to the selection set.

Then for a new satellite to be selected the redundancy scores for the remaining satellites need to be recalculated as they are dependent on S_i. This is repeated until the lowest end score gets too high or the selection subset reaches a predetermined

size. A simplified summarization of the selection algorithm could be given as

1. Calculate p_j, Eq. (3.3)
2. Add those three satellites with the lowest p_j to S_i
3. Calculate r_j, Eq. (3.4), and then s_j, Eq. (3.5), for the remaining satellites
4. Add the satellite with the lowest s_j to S_i
5. Repeat steps 3. and 4. until $\min(s_j) > s_{th}$ or $|S_i| = k$

where s_{th} is a score threshold after which it is considered that adding a satellite into the selection yields no profit and k is a predetermined maximum size for the selection set.

4. RESULTS

The work for this thesis began by analysing data gathered from fixed antennas. This was done to gain a reference point of the GNSS signals in near optimal situations. Reflecting the more realistic scenarios to these near optimal situations helps in achieving an insight into the factors contributing to the location error. After the fixed position analyses our attention was directed to more realistic scenarios. These consisted of three road test recordings. In the road test recordings the location truth was acquired with an Applanix POS LV product and post-processed with POSPac software to reach decimeter level accuracies [3]. The further analyses for the rooftop and road test data are done in python, for which a small framework was developed to streamline things. This framework was also used to benchmark the data-driven algorithm (DDA) that we presented earlier.

Now the rest of this chapter is constructed so that the results from the rooftop analyses are displayed in Section 4.1 and the road test results in Section 4.2. Also a quick look in to the computational requirements of the satellite selection algorithms compared in this work is given in Section 4.3.

4.1 Rooftop data analysis

For the rooftop data analysis work we acquired data from two distinct locations, 24 hours from each. Another location was in Reston, US and the other in Singapore. The locations of the antennas used in the measurements are known down to a level of few centimeters. This allows accurate distance calculations to the visible satellites and thus an accurate determination of the measured pseudorange residuals. The receivers tracked satellites from five most common satellite navigation systems: BeiDou-2, Galileo, GLONASS, GPS and QZSS as well as a few SBAS satellites, although they are not used in the navigation solution.

Table 4.1 shows the numbers of measurements from the different satellite constellations in the two locations. Measurements with elevation below five degrees were eliminated. Singapore has a higher total number mainly due to more BeiDou-2 satellites visible there. Systemwise, GPS is encountered most often in this setup.

	BDS	GAL	GLO	GPS	QZSS	SBAS	Σ
Reston	19415	374412	634333	833221	0	258770	2120151
Singapore	503916	346508	605676	912433	86381	86474	2541388
Σ	523331	720920	1240009	1745654	86381	345244	4661539

Table 4.1 *The number of measurements for the different GNSSs in the two measurement locations.*

In Figure 4.1 the pseudorange residuals are shown as a box plot, grouped according to the satellite systems. The pseudoranges here are as they would be used in the solution calculations, i.e., they are corrected with respect to atmospheric, clock bias and ephemeris errors. The boxes extend from the lower quartile to the upper quartile of the data, with a line at the median. The whiskers extend from the 5th percentile to the 95th percentile. To keep the y-axis readable the fliers are excluded. Already here the BeiDou-2, Galileo and GPS systems exhibit superior behavior over the GLONASS system, which seems to suffer from much higher residuals.

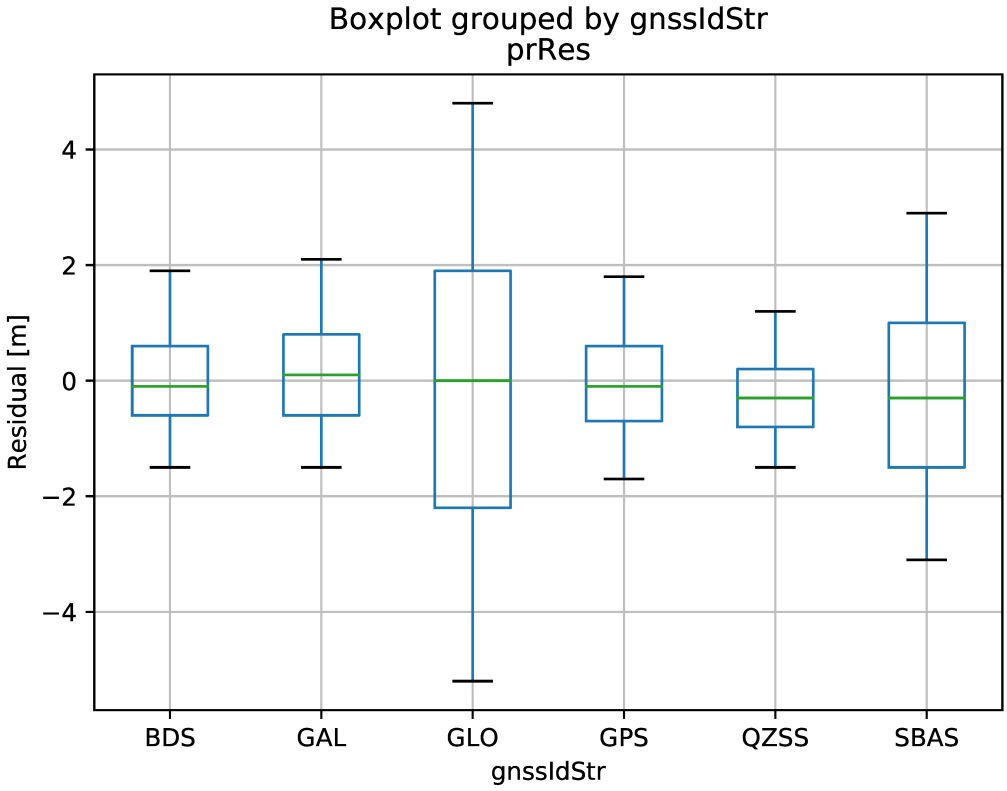

Figure 4.1 *Box plots of the pseudorange residuals for different Global Navigation Satellite Systems recorded from the roof top antennas. The boxes extend from the lower to upper quartile of the data, with a line at the median. The whiskers extend from the 5th percentile to the 95th percentile. The fliers are excluded to keep the y-axis in check.*

Another perspective is given in Figure 4.2 where the pseudorange residuals are plotted as histograms with normal distributions fitted, divided according to the Satellite Navigation Systems. The binwidth used here is 0.1m to correspond with the resolution of the residual data. Now we see that the shape of the GLONASS histogram differs from the others quite significantly, being clearly asymmetric. Also an interesting artifact is present as the bin centered at zero clearly stands out from the adjacent ones. The SBAS residuals are rather peculiarly distributed as well but that is not a concern in this work since they are not used in the positioning solution. A slight irregularity is visible in the BDS histogram too but as a whole it stays quite compact.

Going down to the satellite level in Figure 4.3 to have a look at the pseudorange residuals sheds some light to the irregularity observed in the GLONASS histogram. The measurement errors from satellites belonging to the GLONASS system are far from being zero biased. Some satellites seem to invariably err to the positive side and some to the negative side. This kind of behavior is not apparent in satellites from the other systems. Rather the other satellites, especially from Galileo and GPS systems, stay reasonably uniform with very close to zero means. There are a couple irregularities apparent in the BDS satellites however.

What is also clearly visible is that the whiskers representing the 5th and 95th percentiles are noticeably larger for the GLONASS satellites. This actually applies to the boxes as well, which present the lower and upper quartiles. All in all, besides not having a zero mean the pseudorange measurements from the GLONASS satellites seem to also suffer from more deviation resulting in higher residuals.

Figure 4.4 presents satellites only from the GLONASS system to highlight the aforementioned problems. The biases seen in this system may be due to the different channel access method employed in the GLONASS transmissions. A more thorough reflection of this is given in Chapter 5.

Worth mentioning here is that the pseudorange residuals from the other systems are not normally distributed either. The normality hypotheses were rejected on basis of the D'Agostino and Pearson's test combining skew and curtosis [9], with the p-values being in the 10^{-5} range at best.

Next, in Figure 4.5 are lag plots for the pseudorange residuals, divided again by the different satellite navigation systems. The strong clustering of the points along the diagonal points to a rather strong autocorrelation. Figure 4.6 endorses this observation, showing the autocorrelations for a single GLONASS satellite at varying lags. Achieving strong autocorrelations suggests that an autoregressive model, i.e.,

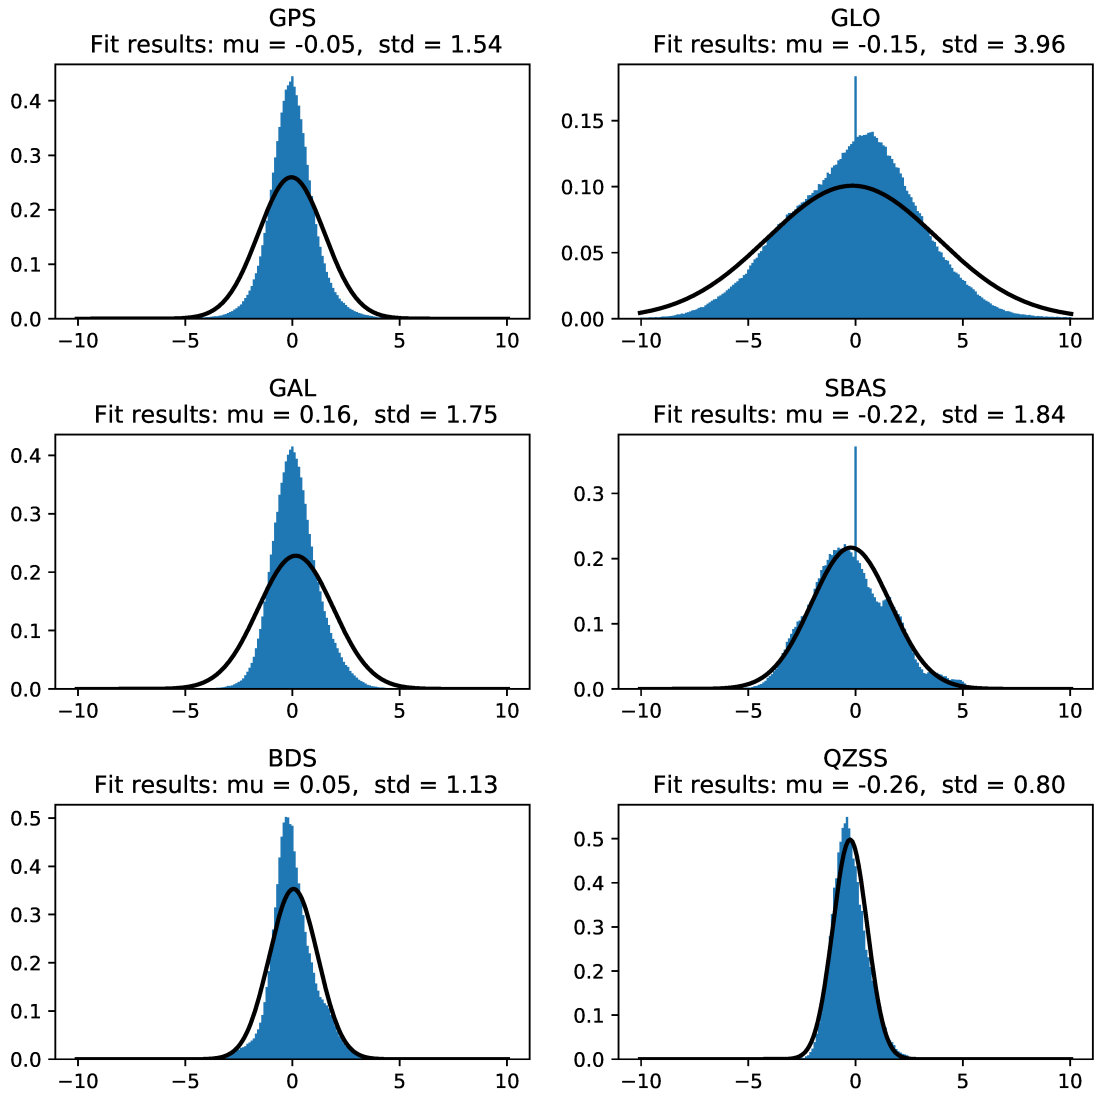

Figure 4.2 *Histograms of the pseudorange residuals for the different Global Navigation Satellite Systems recorded from the roof top antennas and normal distributions fitted to them. The binwidth used for the histograms is 0.1m, which corresponds to the resolution of the residual data.*

the past values of the pseudorange residual could be used to estimate the future residuals.

In theory the pseudorange measurements from the different satellites are considered to be independent. This was assumed in the position accuracy formulations in Section 2.2 as well. That would lead to the position accuracy improving as more satellites are added in to the solution. Although the independency assumption is not entirely true in real scenarios, due to some of the error sources being common, adding more satellites to the solution does indeed increase the obtained accuracy.

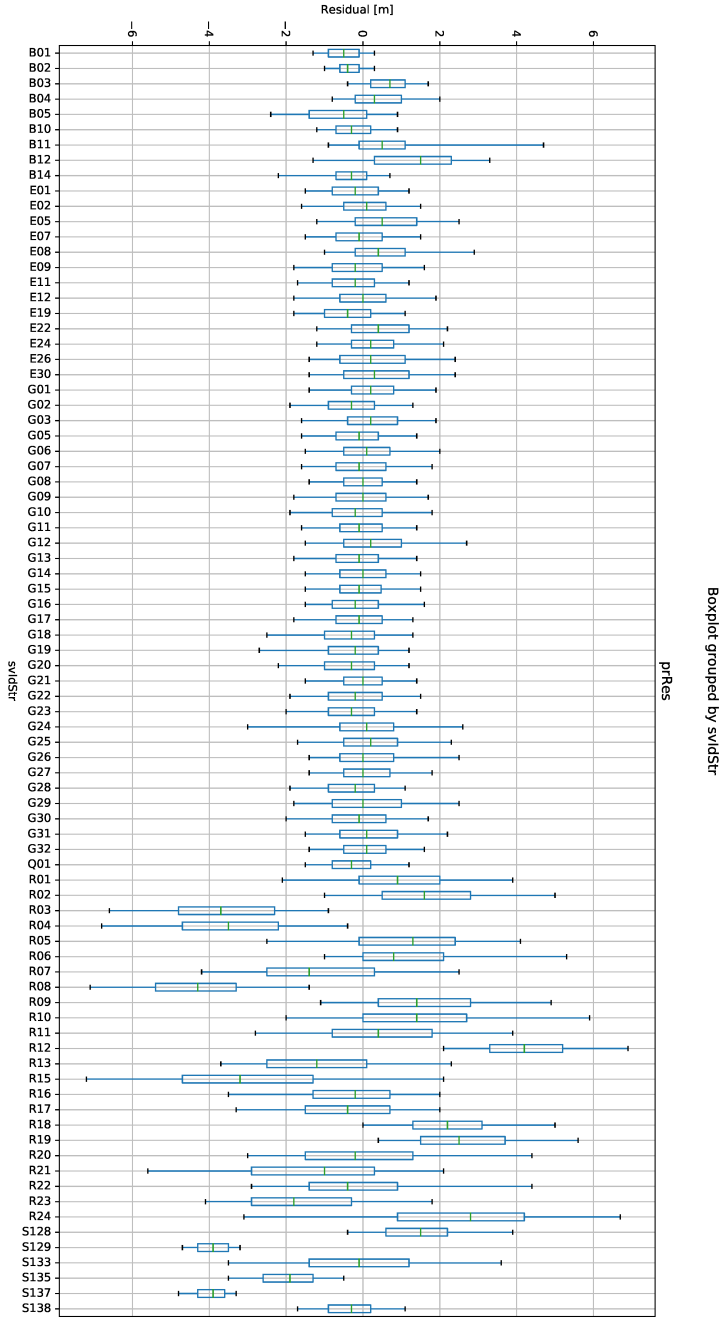

Figure 4.3 *Box plots of the pseudorange residuals for the individual satellites seen in the fixed locations. The boxes extend from the lower to upper quartile of the data, with a line at the median. The whiskers extend from the 5th percentile to the 95th percentile. The fliers are excluded to keep the y-axis in check.*

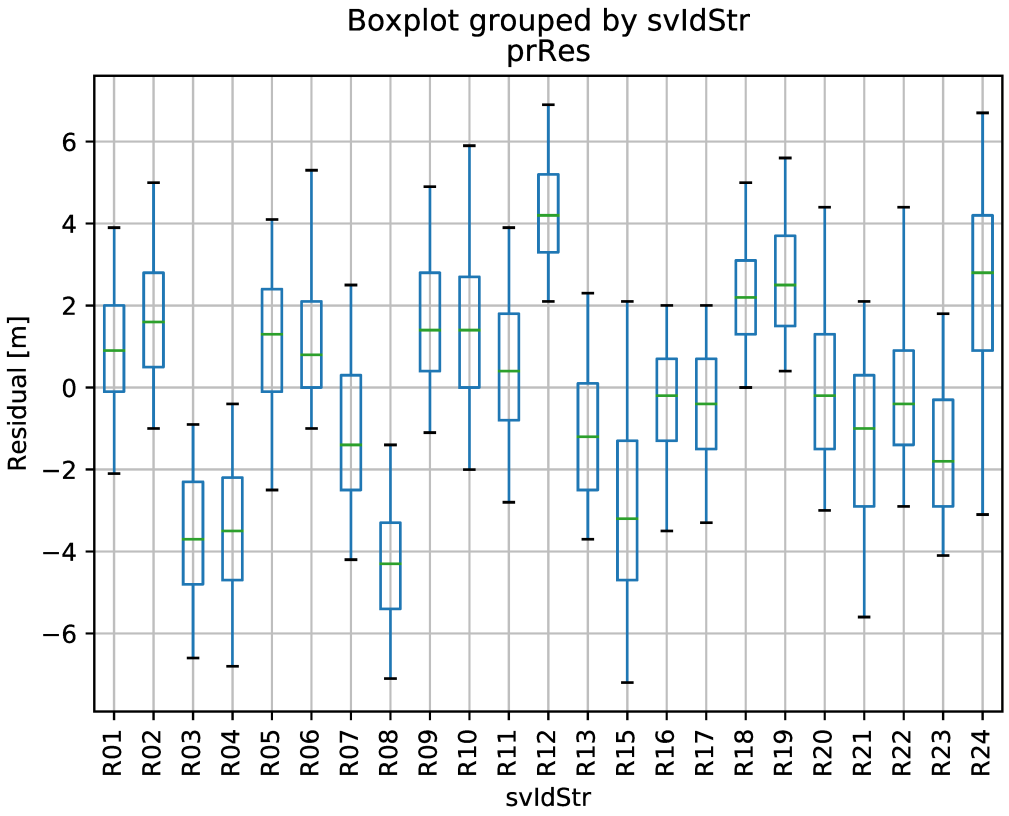

Figure 4.4 *Box plots of the pseudorange residuals for satellites from the GLONASS system recorded from the roof top antennas. The boxes extend from the lower to upper quartile of the data, with a line at the median. The whiskers extend from the 5th percentile to the 95th percentile. The fliers are excluded to keep the y-axis in check.*

In part this works by improving the achieved GDOP and was verified to hold at least in the cases that the errors are drawn from a Gaussian distribution or from the rooftop data.

Figure 4.7 shows how the location accuracy behaves as the number of satellites used in the solution is increased. The blue line shows the location errors when the pseudorange errors are sampled from the rooftop data and in the orange case the pseudorange residuals are modeled as Gaussian with zero mean and standard deviation of 2m. The pseudorange errors are translated to position errors according to Equation (2.12). The satellite geometries for this simulation are from the Multi-GNSS Experiment and Pilot Project (MGEX) data made available by Center for Orbit Determination in Europe (CODE) [12]. The MGEX project is set up to track, collate and analyze all GNSS signals. They also generate precise data products for the orbit and clock of the satellites.

The precision of the orbit data is not of utmost concern in the simulations for these

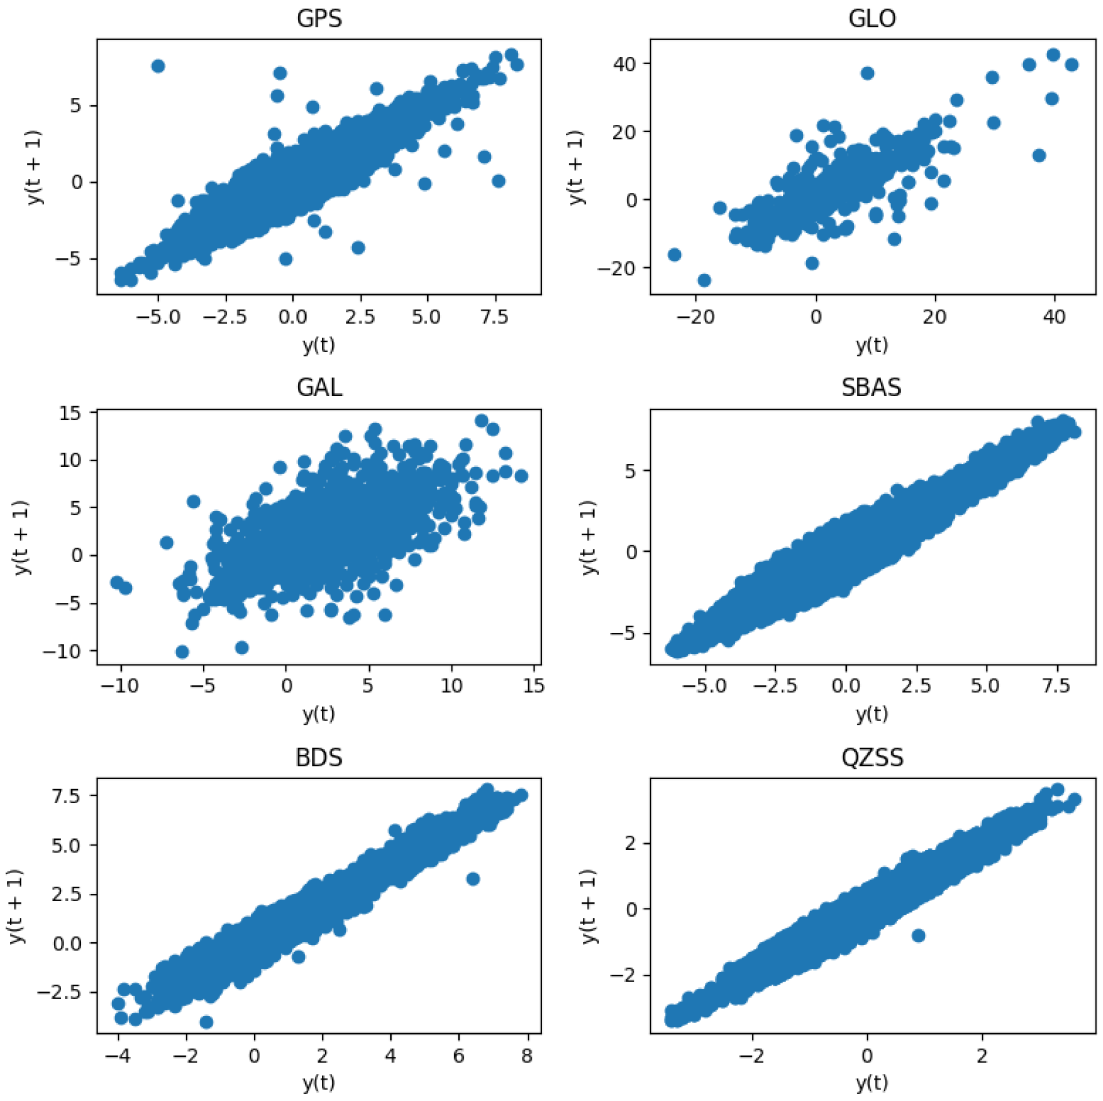

Figure 4.5 *Lag plots of the pseudorange residuals. QZSS data is from the Singapore station, the others in this case only from the Reston station.*

figures, and the MGEX data was selected mainly due to its easy incorporation into the python framework. For Figure 4.7 a time period of 24 hours in the beginning of GPS week 1941 was used. The satellite selection for the location solution in this case is done simply according to their elevation angle, picking the highest ones.

The same data from CODE is used in Figure 4.8 to visualize the number of satellites nowadays available. Here the receiver was simulated to be in 1.3°, 104.0°, 0 location on the WGS 84 ellipsoid and visibility to a satellite was assumed if the elevation from the surface of the sphere exceeded 5°.

From Figure 4.7 we see that the location error starts to pretty much saturate after

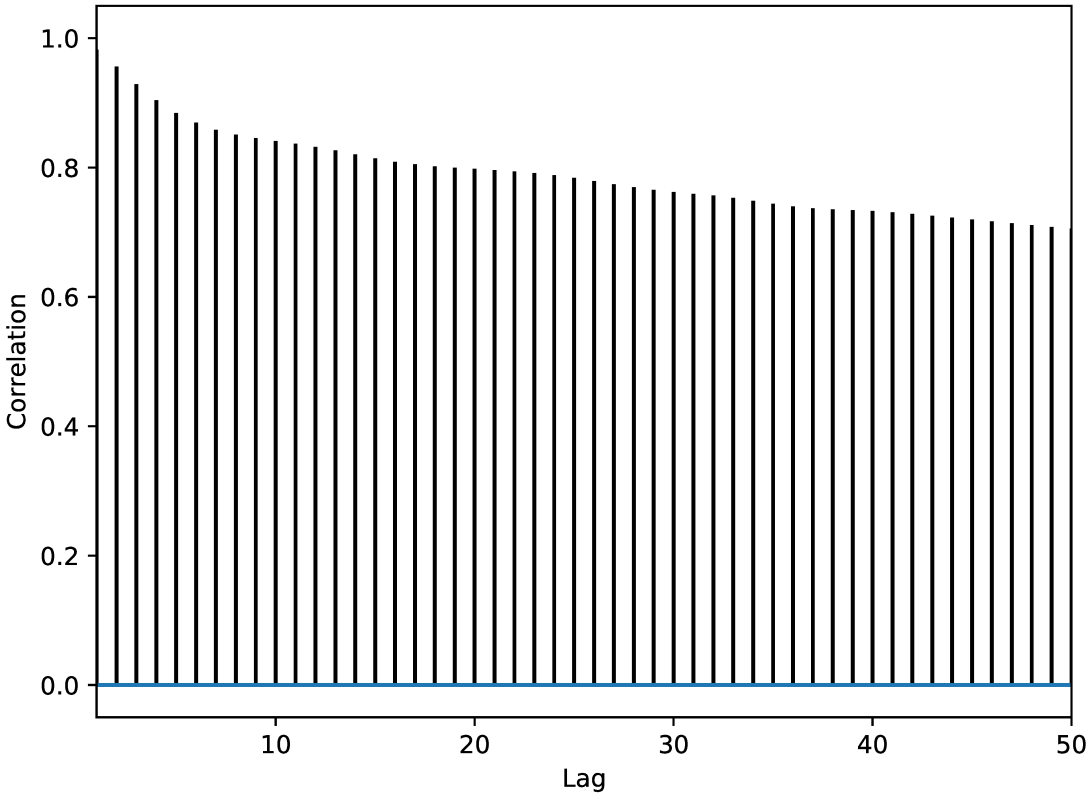

Figure 4.6 *Autocorrelation plot for the pseudorange residuals of a single satellite vehicle from GLONASS constellation.*

20 satellites in the solution, especially in the sampled error case. Combining this knowledge with the numbers we see in Figure 4.8 it is apparent that, if called for, those numbers of satellites in the solution are not a rarity nowadays. Also visible in Figure 4.8 is that GPS still remains the dominant system when it comes to the number of satellites. This is true even in those regions where the geostationary BeiDou satellites are visible, that is, mainly in Asia.

Besides the differing numbers of satellites in the different systems, this section also revealed differences in the measurement accuracies of the systems. Already in the near optimal situations with static antennas in open sky environments it was visible how the GLONASS system had significantly larger pseudorange residuals than the other systems. The measurements from satellites belonging to GPS, Galileo and BeiDou systems were comparatively even, although the BeiDou did indicate slight biases for some of the satellites. It was also noted that the pseudorange residuals are strongly autocorrelated and thus an autoregressive model would most likely be appropriate in predicting the errors. Finally a short look was taken at how the number of satellites in the solution affects the location error. It seemed that having

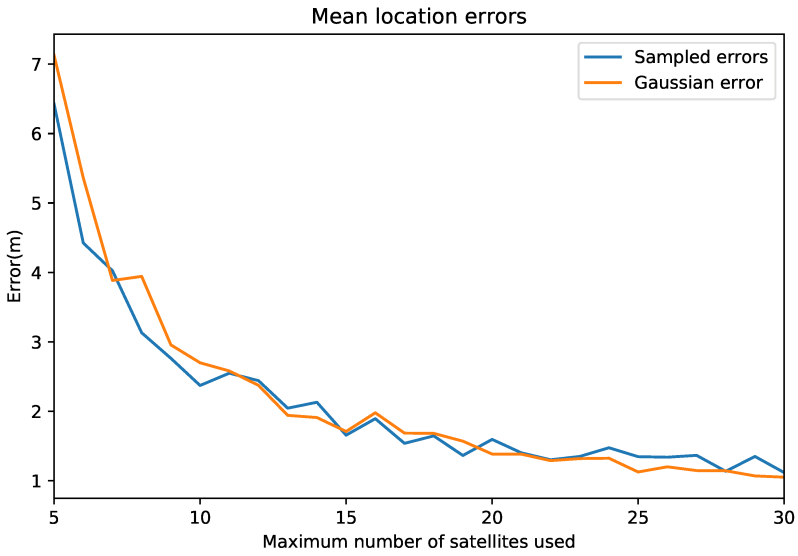

Figure 4.7 *Mean location error as a function of the number of satellites used for the solution. The numbers are averaged from 96 simulations. The blue line gives the mean location errors for the case that the pseudorange errors are sampled from the rooftop data. In the orange case the pseudorange residuals are modeled as Gaussian errors with zero mean and 2m standard deviation. The satellite selection in this case is done by picking the vehicles with the highest elevation angles.*

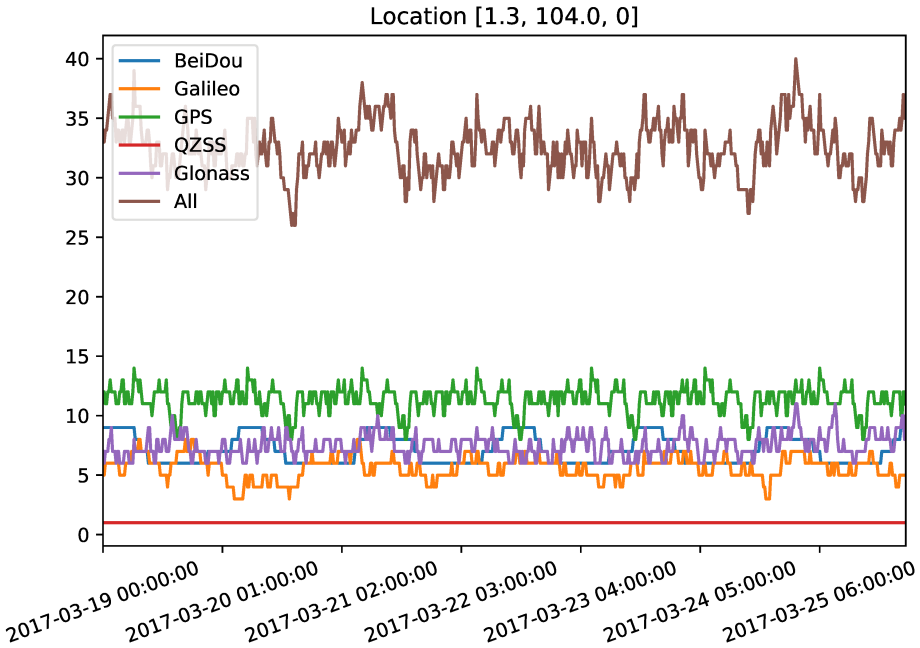

Figure 4.8 *The number of theoretically visible satellites in location 1.3, 104.0 (decimal degrees, WGS 84) during GPS week 1941 in March 2017. The satellite locations are from MGEX data made available by Center for Orbit Determination in Europe (CODE). An elevation of over five degrees was required for the satellite to be considered visible.*

over 20 satellites in the solution had very negligible improvements for the accuracy.

In the next section the environments the data was recorded in switch to a more realistic level, as we look at data that was gathered driving in and around Zürich.

4.2 Road test results

The road test data used for the results in this section comes from three recordings done in and around Zürich. They contain changing environment from urban settings to open skies on country roads. These recordings give a very realistic picture about the situations the GNSS receiver faces when used for navigation in a moving vehicle.

The data is now completely acquired from the receiver, the satellite locations and signal informations are sampled at 1Hz intervals. The location truth is from an Applanix POS LV product, post-processed with POSPac software and is accurate down to a decimeter level [3]. This truth recording is used to form the pseudorange residuals for the satellites at each epoch.

First we looked at the pseudorange residuals in these more realistic scenarios, repeating the analyses of the previous section for the road test data as well. The three recordings were treated separately in the actual algorithm benchmarks but combined for the initial pseudorange residual analysis.

Figure 4.9 shows the pseudorange residuals as a box plot grouped according to the systems. As expected, all of the systems generate higher residuals than were seen in the open sky environments. Also, the GLONASS system involves significantly higher residuals compared to the three other global navigation satellite systems, which also is in line with the results of the previous section.

In Figure 4.10 are the histograms of the pseudorange residuals in these road test datasets. Here again the three systems; GPS, Galileo and BeiDou seem comparable, whereas the GLONASS has a clear disadvantage. However the normal distribution fitted to Galileo measurements embodies a substantial standard deviation. This would suggest that there are relatively many observations in the tails of the Galileo histogram.

This suggestion is confirmed in Figure 4.11, where the pseudorange residuals are given as a box plot grouped by the individual satellites. The satellite E12 exhibits highly deviant results here, as the 95th percentile of its residuals extends all the way up to 60m. Not even the satellites from the GLONASS system reach this level of inaccuracy, although in general being much more imprecise.

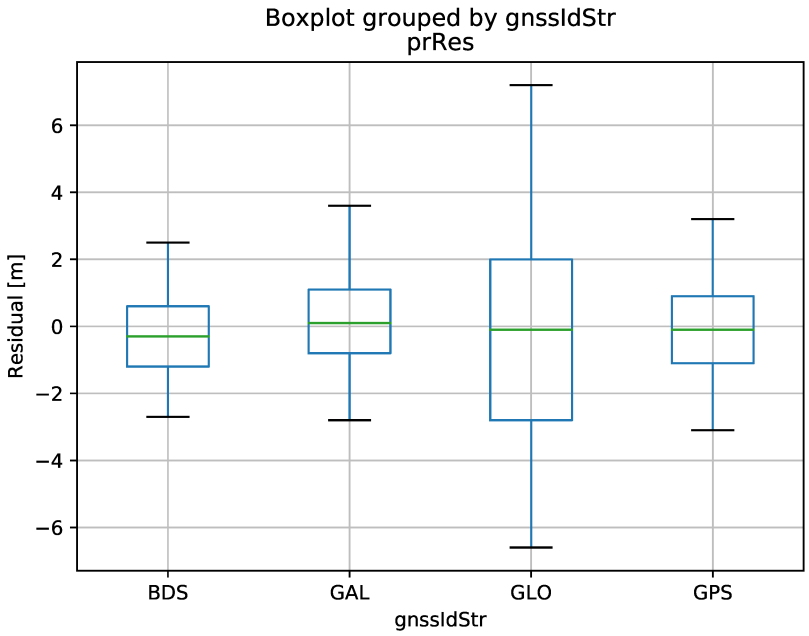

Figure 4.9 *Box plots of the pseudorange residuals for different Global Navigation Satellite Systems recorded in road test scenarios. The boxes extend from the lower to upper quartile of the data, with a line at the median. The whiskers extend from the 5th percentile to the 95th percentile. The fliers are excluded to keep the y-axis in check.*

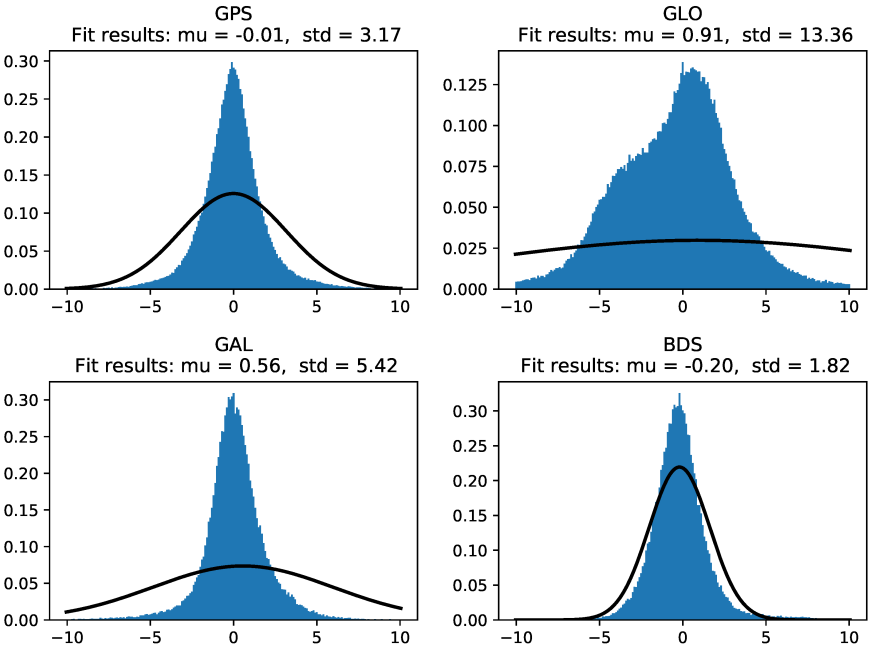

Figure 4.10 *Histograms of the pseudorange residuals for the different Global Navigation Satellite Systems recorded in road test scenarios and normal distributions fitted to them. The binwidth used for the histograms is 0.1m, which corresponds to the resolution of the residual data.*

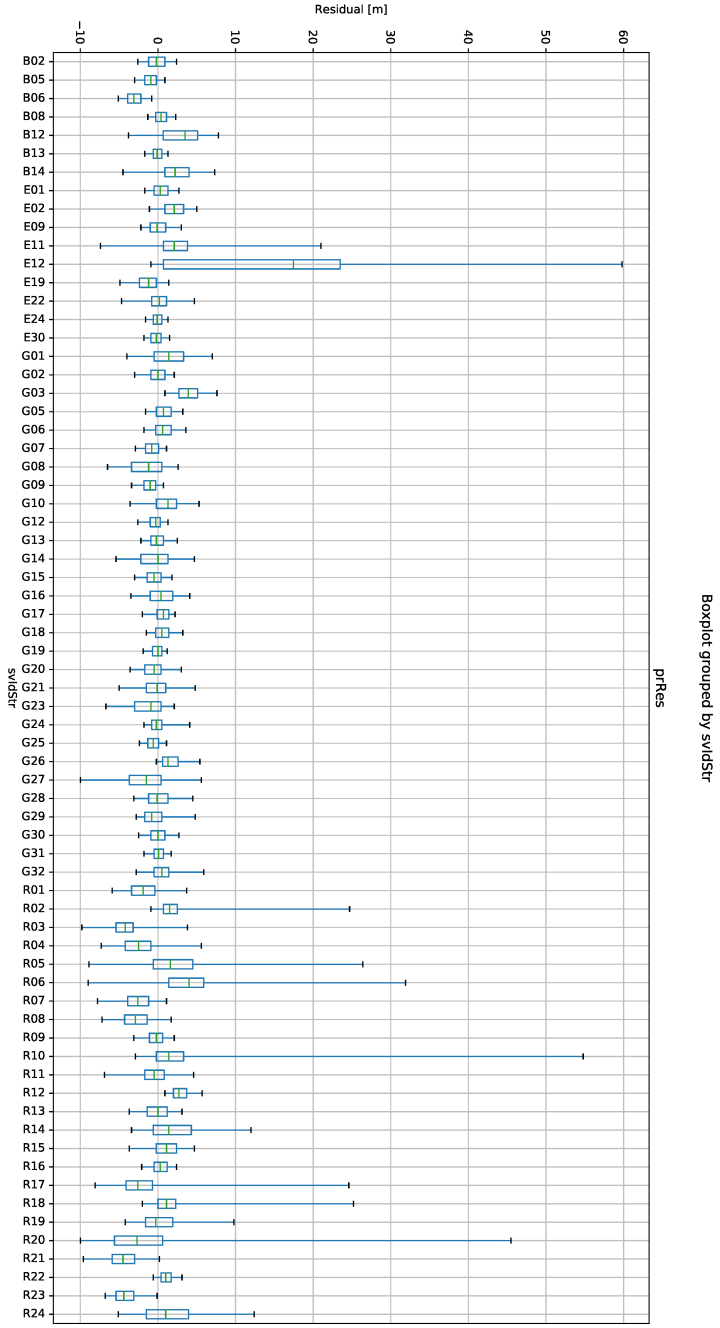

Figure 4.11 *Box plots of the pseudorange residuals for the individual satellites recorded in the road test scenarios. The boxes extend from the lower to upper quartile of the data, with a line at the median. The whiskers extend from the 5th percentile to the 95th percentile. The fliers are excluded to keep the y-axis in check.*

Replicating the residual analyses for the data from the road tests illustrated how the environmental factors affect the accuracy of the pseudorange measurements. The residuals measured when driving around Zürich were considerably poorer when compared to those from the rooftop measurements.

We then moved on to test the applicability of logistic regression for the prediction of those poorer pseudoranges. As logistic regression is a classification model, the pseudorange residuals needed to be categorized. For a bad pseudorange we now chose a residual threshold of 10m. That is, if the receiver registered a pseudorange that deviated more than 10m from the distance calculated between the location truth and the satellite ephemeris then the pseudorange was considered bad. The threshold of 10m was chosen as a trade-off between having a reasonable number of measurements considered bad and having a sufficient difference to the accuracy of the truth data.

The prediction was done in a threefold cross-validation manner as there are three recordings. One of the recordings at a time was used as the test set whereas the two others were used for training the model. Figure 4.12 shows the coefficients the model gave for the different features, i.e., the vector w from Section 2.3. The bars represent the means from the three folds and the whiskers are one-sigma errors. Notable here is that the system identifiers are now binary variables telling if a satellite belongs to a given satellite navigation system and the carrier-to-noise-density ratio (cno) and elevation (elev) features are continuous. The coefficients are to be interpreted so that in the case of cno and elev a greater variable value increases the probability of a good pseudorange and in the system identifier case a greater negative coefficient is worse.

As the runs were done in Europe there were not that many BDS satellites available and the conclusions done from the coefficients for that system would require some caution. The others seem to be in line with the observations done in the previous section, although more data would be needed to narrow down the deviation in the estimates.

The results of the predictions are given in Table 4.2 for each of the cross-validation folds accompanied by the means and standard deviations. The means in each column exceed 90%, which is beyond our initial expectations. The training and test accuracies are very close to each other suggesting that the model does not overfit. The recall and specificity numbers further support the applicability of the model. In the recall number we see that on average the model can spot around 90% of the bad pseudoranges and specificity tells us that only about 10% of the good pseudoranges

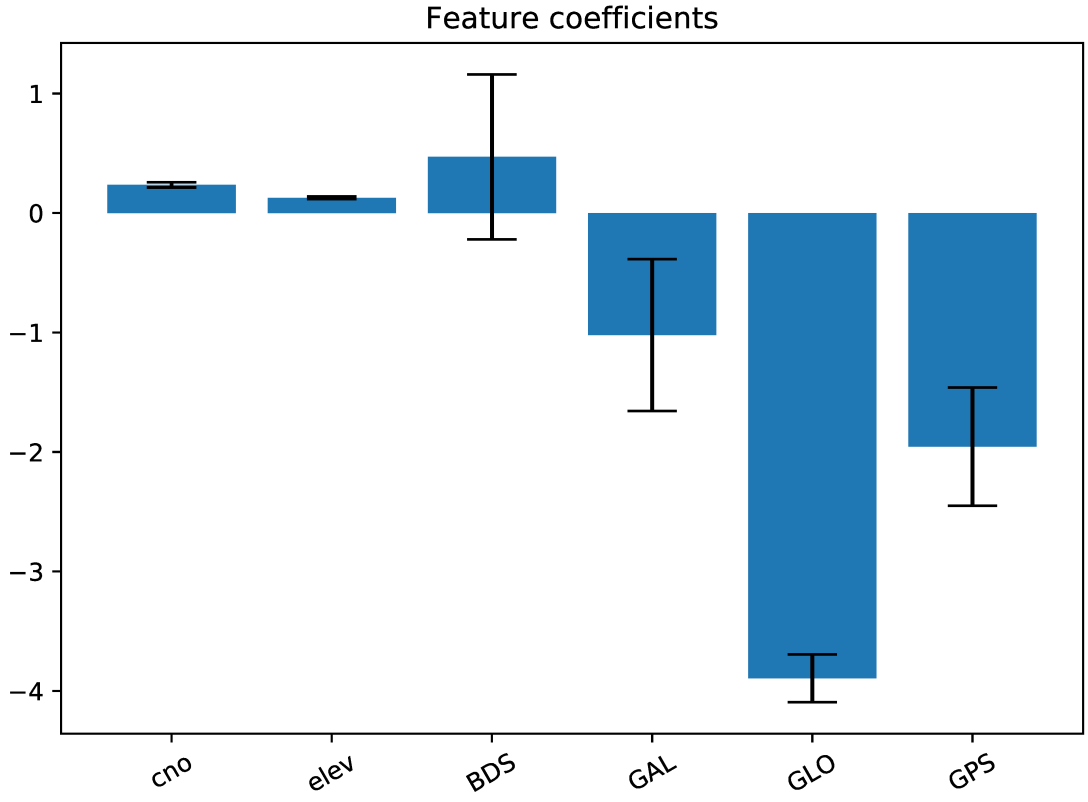

Figure 4.12 *The logistic regression coefficients used for predicting if the residual of a pseudorange measurement is going to be over 10m. The bars are the means from the cross-validation folds and the whiskers present one-sigma errors.*

are falsely rejected.

	Training accuracy	Test accuracy	Test recall	Test specificity
Fold 1	0.911	0.900	0.899	0.918
Fold 2	0.909	0.946	0.947	0.885
Fold 3	0.925	0.899	0.899	0.903
mean ± std	0.915±0.007	0.915±0.022	0.915±0.023	0.902±0.013

Table 4.2 *The results from predicting over 10m pseudorange residuals in the road test data.*

As was described in Section 2.3 the logistic regression works by coercing the linear combination of the features to a probability using the logistic function. Figure 4.13 shows histograms of these probabilities for the prediction in the test set of the third fold. The data is divided in two according to the actual pseudorange residual, under 10m (good) and over 10m (bad). Here we see that the data is quite well separable even with a linear decision surface.

Next came the part where the most important results of this work were reaped,

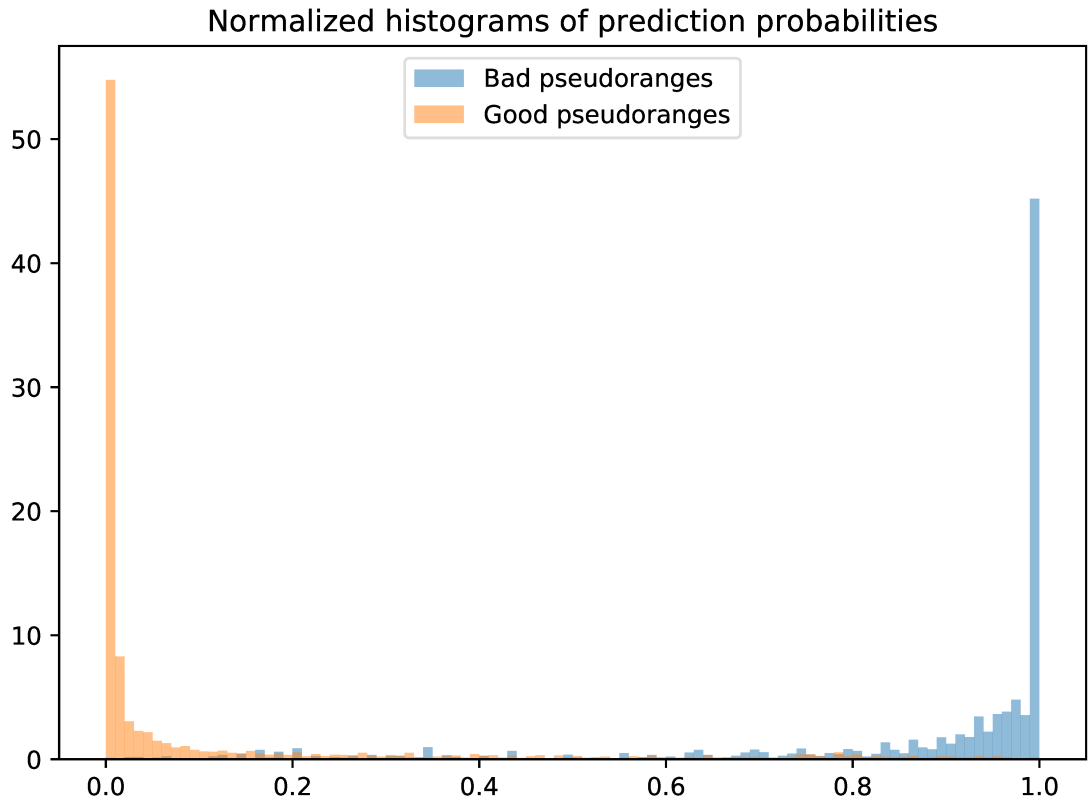

Figure 4.13 *Normalized histograms of the logistic regression probabilities divided according to if the actual pseudorange residual is under 10m (good) or over it (bad).*

applying the data driven satellite selection algorithm to the road test data. The following figures show results from the test set of the third cross-validation fold, i.e., the first two recordings were used for the training of the model and then the remaining recording is the test case. The third fold was selected as it was considered the worst scenario for the logistic regression according to the numbers seen in Table 4.2.

We also applied the quasi-optimal algorithm (QO) to the test case to obtain a reference for the performance. The quasi-optimal algorithm does not produce optimal accuracies and certain algorithms exist that exceed it, even GDOP wise. However it is well established in the field being a point of reference for virtually all the other algorithms. More importantly, it allows real time implementation even in the computationally limited receiver hardware [27].

Table 4.3 gives a summary of the results obtained from the test set with different satellite selection strategies. Overall the data-driven approach achieves very good accuracy results here. The mean results of under 2 meters in the 20 and 30 satellite

cases can be considered very good and the median shows that 50% of the fixes are even more accurate than that. The 30 satellite quasi-optimal case has identical numbers compared to the all in view up to the precision reported here.

The row with the maximum errors shows perhaps the clearest differences between the methods, with the data driven approach achieving numbers that are under one third of the all in view and QO maximum errors. An interesting aspect here is that when given the chance to use up to 30 satellites for the data-driven method it produces slightly bigger maximum errors than in the 10 and 20 satellite cases. All in all, raising the maximum number of satellites from 20 to 30 in the data-driven case produces rather minimal improvements in the location accuracy, the only notable differences seen in the minimum and the first quartiles.

Error [m]	AIV	DDA 10	DDA 20	DDA 30	QO 10	QO 20	QO 30
mean	3.38	2.01	1.37	1.37	8.51	3.72	3.38
std	3.37	0.90	0.87	0.87	7.83	3.55	3.37
min	0.08	0.11	0.07	0.05	0.10	0.14	0.08
25%	1.29	1.40	0.83	0.82	3.58	1.48	1.29
50%	2.18	1.86	1.18	1.18	5.99	2.53	2.18
75%	4.36	2.44	1.66	1.67	11.21	4.78	4.36
max	28.52	7.40	7.69	7.78	90.18	30.26	28.52

Table 4.3 *A summary of the location errors obtained from the test data using different methods for the satellite selection. The errors are measured as distance from the truth and given in meters. AIV refers to the All In View case, DDA to Data-Driven Approach and QO to the Quasi-Optimal method. The numbers behind the latter two refer to the number of satellites selected.*

Figure 4.14 shows the location errors as a function of time with different satellite selections from the quasi-optimal algorithm compared to our data driven approach selections. In addition the all in view case is given, where all of the visible satellites would be used for the location solution. Here we see that the test scenario has some epochs where the location errors grow notably in the all in view and quasi-optimal cases. The data driven approach however handles these very well.

Figures 4.15 and 4.16 provide vision into the reasons behind the poor location accuracies observed in the previous figure. In Figure 4.15 are the geometric dilution of precision numbers plotted for the different selections as a function of time. The quasi-optimal algorithm performs slightly worse than the data driven approach in the 10 and 20 satellite cases, which can be considered surprising as the QO method is designed purely for GDOP optimization. However it achieves identical numbers with the all in view selection in the 30 satellite case (the lines overlap completely), beating the data-driven method here.

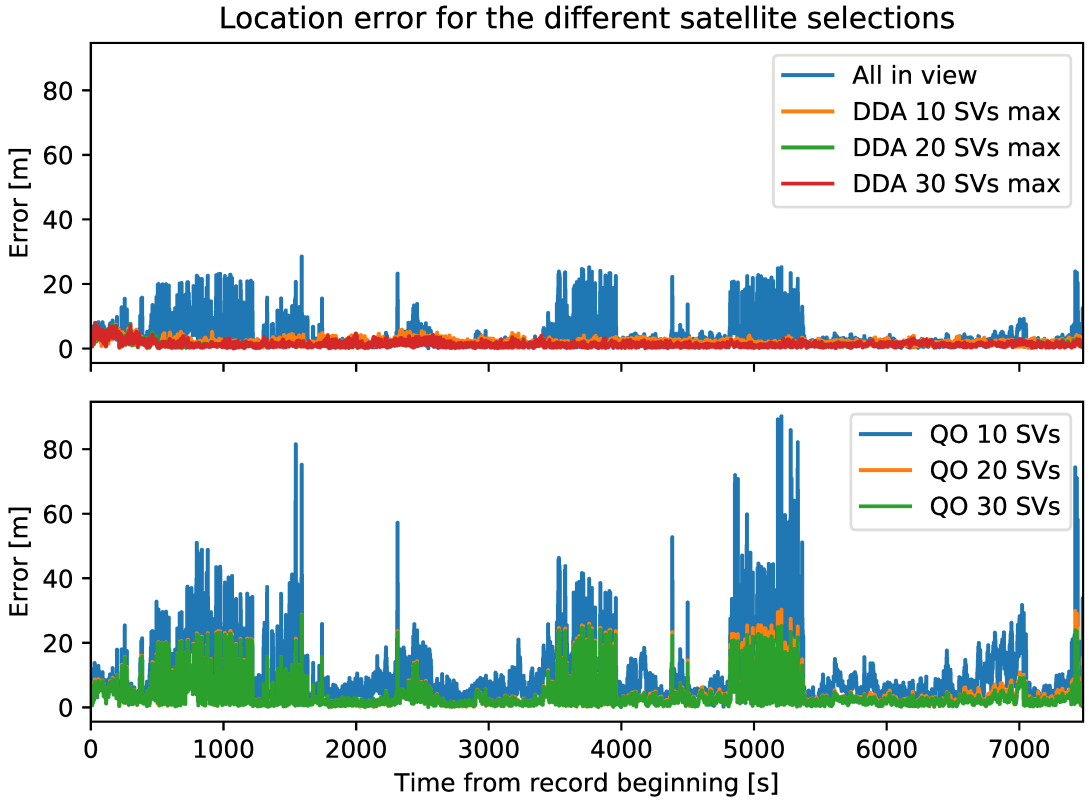

Figure 4.14 *Location errors acquired in the test data with the all in view, our data driven and the quasi-optimal satellite selection methods.*

When we look at the quadratic means of the pseudorange residuals in Figure 4.16 we see that the bad residuals overlap in time with the spikes in the location error plot. This figure also highlights how the data driven approach manages to remove those bad residuals in the solution set of measurements.

Figure 4.17 gives the location errors in a box plot form. This makes it easier to compare the average performances of the different algorithms at varying selection set sizes. The boxes extend from the lower to the upper quartile and the whiskers from the 5th percentile to the 95th percentile. Also the outliers are presented now. Apparent here is how the data driven approach performs noticeably better than the quasi-optimal algorithm and does not suffer from location errors of over 5m. An interesting observation are the few outliers in the data driven approach with a maximum of 30 satellites that are not present in the maximum of 20 satellites case.

Also histograms of the location errors were drawn up, these are presented in Figure 4.18. The errors are binned into 9 bins with a width of one meter. The 1-2m error range has the most counts for almost all of the selections, only the quasi-optimal

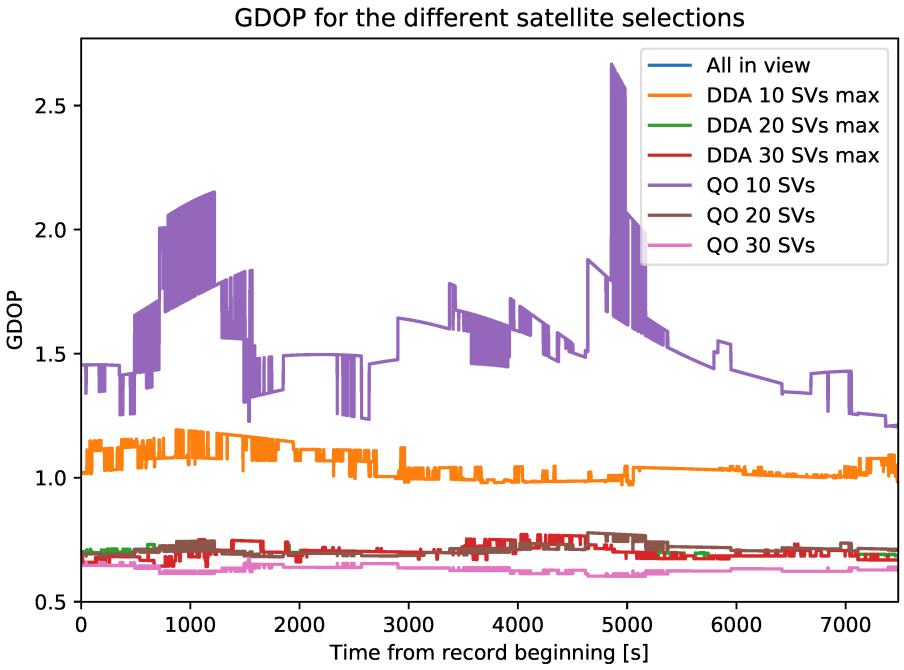

Figure 4.15 *GDOPs acquired in the test data with the all in view, our data driven and the quasi-optimal satellite selection methods.*

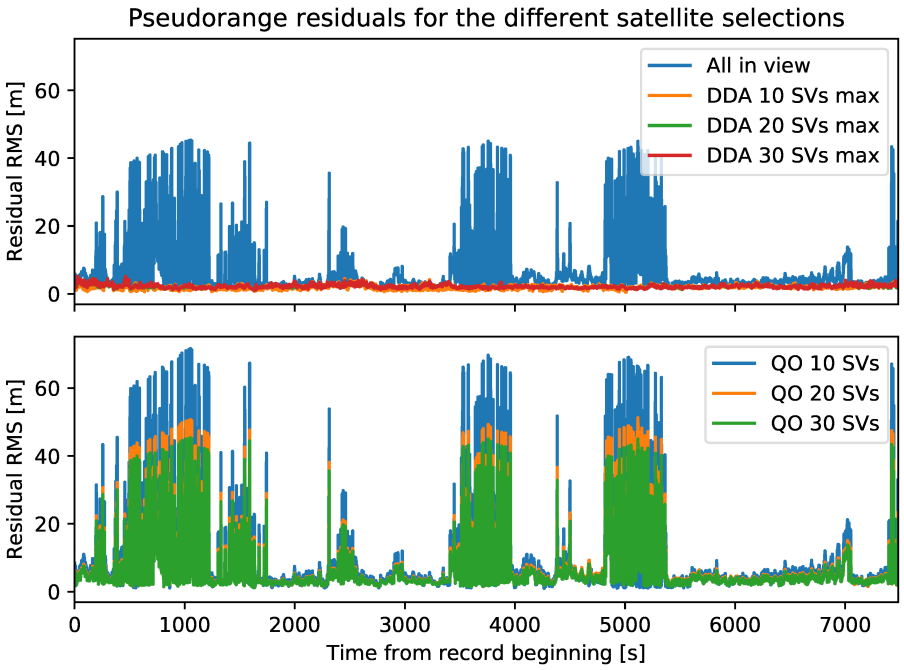

Figure 4.16 *Pseudorange residual root mean squares at each epoch acquired in the test data with the all in view, our data driven and the quasi-optimal satellite selection methods.*

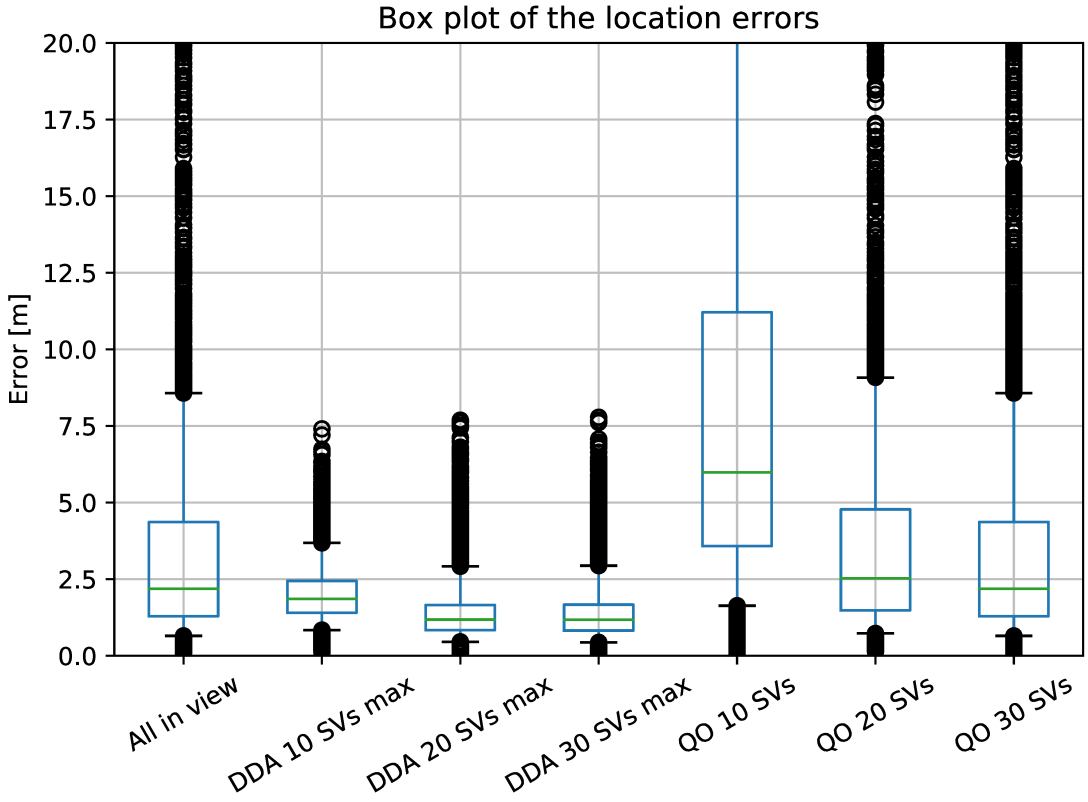

Figure 4.17 *Box plot of the location errors acquired in the test data with the all in view, our data driven and the quasi-optimal satellite selection methods.*

with 10 satellites differs in this sense. This representation also shows how the data driven approach with a maximum of 30 satellites selected has slightly higher counts in the first two bins than when allowing a maximum of 20 satellites. But these differences are very minor and as a whole the 20 and 30 satellite cases for the data-driven method can be considered equally accurate. And again can be seen how neither of them induce errors of over 8m.

These results indicated how the logistic regression works surprisingly well in predicting the bad pseudoranges. The good and bad residuals seem to be rather well linearly separable with those prediction variables that were selected for the model. This primes the data-driven method to succeed in weeding out the bad pseudorange measurements and thus provides for a more accurate positioning solution. In addition the redundancy score based optimization for the geometric dilution of precision appeared to work adequately well.

The next section looks at the computational complexity of the data-driven method and again compares it the quasi-optimal algorithm. After all, one of the requirements

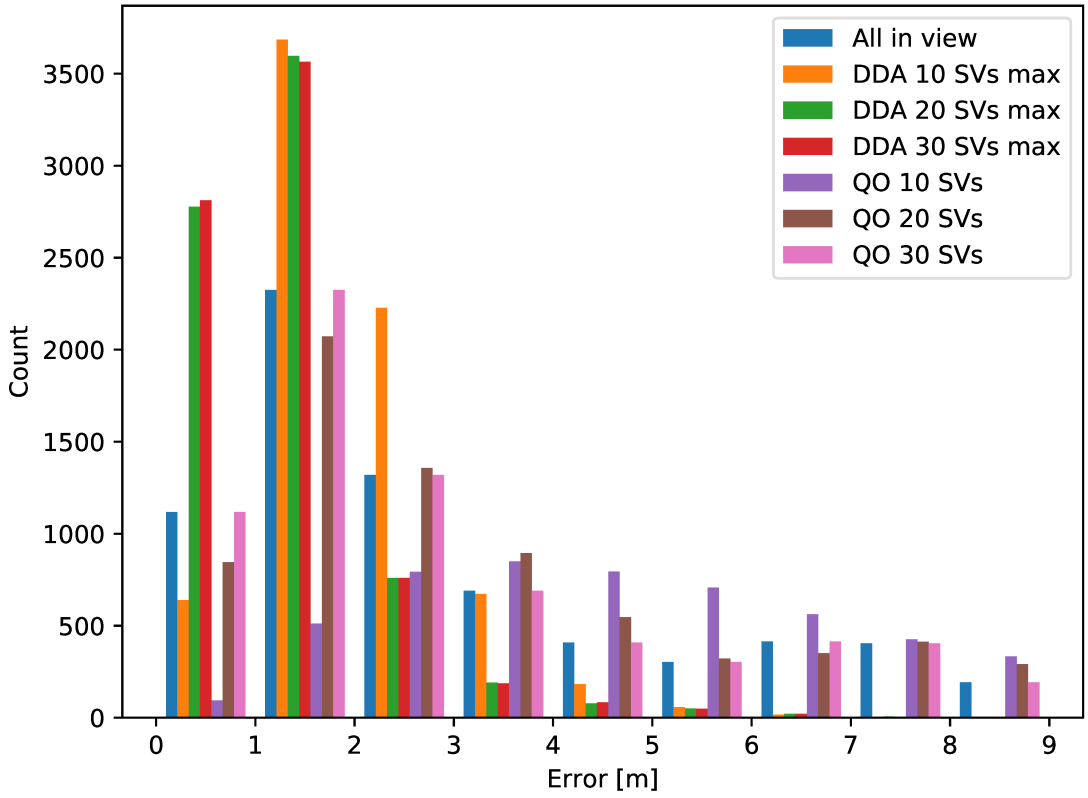

Figure 4.18 *Histograms of the location errors acquired in the test data with the all in view, our data driven and the quasi-optimal satellite selection methods.*

for the satellite selection was that it needs to work in real time in the receiver hardware.

4.3 Computational complexity

Let us consider the quasi-optimal algorithm first. In the trivial case that the number of visible satellites happens to be the same as (or less than) the number of satellites we wish to select then there is no computation needed, we just simply select all of the satellites in view. Then consider that we need to drop one satellite, i.e., the number of visible satellites, $|S|$, is one greater than the number of satellites we wish to select for the solution. Now, since the quasi-optimal algorithm works via backward elimination, we need to calculate the cosines for all of the angles between the visible satellites, see Equation 3.2. As one angle requires two satellite line of sight vectors there are $\binom{|S|}{2}$ angles. Thus when we consider the number of satellites to be selected constant, as it is, the complexity of the quasi-optimal algorithm w.r.t. the number of visible satellites is

$$\mathcal{O}\left(\binom{|S|}{2}\right) = \mathcal{O}\left(\frac{|S|(|S|-1)}{2!}\right) = \mathcal{O}\left(|S|^2\right). \tag{4.1}$$

Considering the computational complexity of the data-driven method is not quite as straightforward. In the initial stage, we need to compute the signal scores for all of the satellites, hence that part is $\mathcal{O}(|S|)$. As was stated in Section 3.2 the first three satellites are added to the selection according to these signal scores. Then, for more satellites to be added we start calculating the redundancy scores. So for the fourth satellite the number of angles we need to consider is $3(|S|-3)$. In general, for $|S_i|$ satellites to be selected out of $|S|$ we need to consider

$$\sum_{i=3}^{|S_i|-1} i(|S|-i) \tag{4.2}$$

angles. For this the complexity can be given as $\mathcal{O}(|S_i||S|)$, and since $|S_i|$ is a constant it can be simplified to $\mathcal{O}(|S|)$. Thus the complexity for the data-driven method is

$$\mathcal{O}(|S|+|S|) = \mathcal{O}(|S|). \tag{4.3}$$

Worth remembering is that the signal scores need to be calculated only once in an epoch.

To verify, we calculated the numbers of floating point multiplications needed when selecting 20 satellites out of a varying number of visible satellites. Figure 4.19 illustrates these numbers which support the complexities attained above. When there is only one satellite to dismiss the backward elimination of the quasi-optimal algorithm does this with far fewer calculations than the data-driven approach, which works in a forward selection manner. However, the number of multiplications needed in the quasi-optimal algorithm grows rapidly as the number of visible satellites increases, whereas in the data-driven approach the growth is rather linear. If we were to select only 10 satellites the difference would grow even more substantial in favor of the data-driven approach.

The numbers of visible satellite in this figure are completely realistic for multi-constellation receivers as could be seen in Figure 4.8. Thus we can conclude that the data-driven approach requires far less computational time than the quasi-optimal method in most common scenarios.

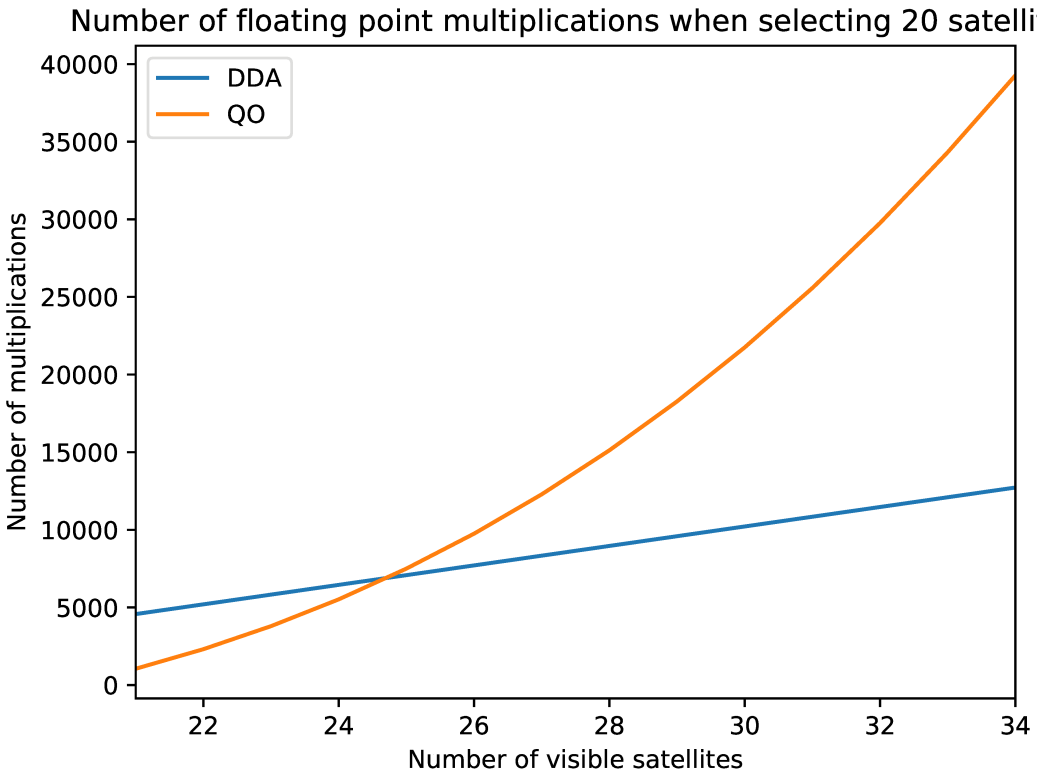

Figure 4.19 *Number of floating point operations needed as a function of the number of visible satellites when 20 satellites are to be selected.*

Next, Chapter 5 concludes this work after a brief discussion about the results presented here. We also take a look in to the possible future improvements for the data-driven approach we suggested for the satellite selection problem.

5. DISCUSSION AND CONCLUSIONS

The roof top scenarios, from which the data for Section 4.1 was collected, presents a very close to optimal situation for the receiver. The high vantage point provides a situation where there are no obstructions on the signal path nor foliage of any kind. This scenario was analysed as we wanted to have a reference point where especially the multipath errors would be minimized.

The first thing these results confirmed was the poorer quality of measurements from the satellites belonging to the GLONASS system, even in the clear sky conditions. As could be seen in Figures 4.1 and 4.2 the pseudorange residuals for GLONASS were of inferior quality compared to the other global navigation satellite systems. An idea about one of the reasons behind this can be drawn from the box plots for individual satellites seen in Figures 4.3 and 4.4. Here the pseudorange residual distributions for the GLONASS satellites are further away from being zero mean than in the case of other systems.

This is most likely due to the fact that GLONASS uses frequency divided multiple access (FDMA) modulation for the signals [29] whereas the other systems rely on code division multiple access (CDMA) [26, 7, 13]. The utilization of FDMA means that the signals from different satellites are transmitted on different frequencies. This has been reported to cause inter-frequency or inter-channel biases [32], which also cause bias in the pseudorange measurements [8].

Often, at least in theoretical settings, the pseudorange residuals are considered i.i.d. with a zero mean Gaussian distribution. This was also checked with a D'Agostino and Pearson's test that combines skew and curtosis [9]. The tests rejected the hypotheses that even for a single satellite the residuals would be Gaussian. This provoked a further investigation in to the nature of the errors. Figures 4.5 and 4.6 illustrate a strong autocorrelation in the residuals. This suggests a strong inclination towards an autoregressive model. Such a model would possibly yield even better prediction results for the pseudorange errors, and thus better results for our satellite selection approach. However that would bring about certain awkward questions for the implementation, e.g., how reliable are the pseudorange residuals measured at

the previous epoch.

In a theoretical setting, with i.i.d. pseudorange residuals, it is obvious that adding more satellites to the solution increases the accuracy. As the i.i.d. assumption is rather arguable in a more realistic setting we had a look at the location accuracy as a function of the number of satellites in the solution in Figure 4.7. The plots here show that after roughly 20 satellites adding more measurements to the solution induces a very marginal improvement in the location accuracy. And when the added computational complexity is considered, limiting the number of satellites in the solution is well-advised already from the computational point of view.

Figure 4.8 showed that the number of visible satellites in multi-constellation receivers can already be as high as 40 so limiting the number in the solution is indeed called for. This number will only keep growing in future as there are more satellites launched in especially Galileo and BeiDou systems to reach their full designed extent. Hence the importance of the satellite selection question will become even more prominent as the number of satellites keeps on growing.

Previous work concerning the selection has been predominantly about optimizing the geometrical dilution of precision in the solution set. However that is only a scaling factor in the ensuing location accuracy, the errors stem from the pseudorange inaccuracies. Nonetheless there are very few papers published where the satellite selection has happened from the measurement error minimization point of view. This directed most of our efforts in to the direction of the pseudorange residuals.

The initial analyses from the rooftop data that were presented in Section 4.1 provided valuable insight into the selection issue. It affirmed our hypotheses that the pseudorange residuals, i.e., the measurement errors are somewhat correlated to factors like the constellation the satellite belongs to, carrier-to-noise-density ratio and elevation. This information gave rise to the idea of trying to predict the pseudorange errors using data that is available already before the location solution. The first approach was naturally to apply linear regression to the data. This however produced rather poor results in the initial tests. Discretizing the dependent variable proved a simple solution and logistic regression gave surprisingly good results already from the first tests onwards.

The regression runs were done on a more realistic data sets collected driving in various environments in and around Zürich. This road test data presented larger errors than the measurements from the rooftops but was otherwise in line with the observations done in Section 4.1. After describing the residuals in these realistic use cases we moved on to the prediction part.

Table 4.2 presented the results from the logistic regression model when predicting over 10m pseudorange residuals. It can be seen that the model works well in the road test scenarios and that the results from the three cross-validation folds are in unison. The recall of nearly 90% is especially pleasing as it means that the model was able to spot 9 out of 10 bad pseudorange measurements. However when looking at the coefficients in Figure 4.12 we see that there is considerable variation in them, especially for the BeiDou system. To get more rigid estimates for the coefficients more data would be needed for training. Also if more data of sufficient quality would be available there would be a possibility to even calculate weights for each individual satellite instead of using the systems as we do here. All in all, the efficiency of applying logistic regression for the problem at hand proved sufficient to say the least.

The location accuracies, given in Table 4.3, obtained with the data-driven approach were in a completely different class when compared to the satellite selections of the quasi-optimal algorithm. This was mainly due to the ability of the DDA to eliminate the measurements with significant error in them. However it needs to be noted that when using 10 or 20 satellites the data-driven approach achieved lower DOP as well. Worth remembering is that all the results are without the receiver autonomous integrity monitoring (RAIM) applied. The fault detection and exclusion part of RAIM performs certain basic checks to the satellite signals and can weed out some of the bad pseudorange measurements [5]. The addition of RAIM would most likely improve the acquired results to some extent.

Another interesting aspect in the location accuracy results was that there was virtually no difference in using 20 or 30 satellites in the data-driven approach. This goes well hand in hand with the observation done in Figure 4.7. Even in few epochs the accuracy was slightly poorer when allowing the DDA to use 30 satellites in the solution. This would indicate to having a slightly too high overall score(Eq. 3.5) threshold after which no more satellites are to be added to the solution.

As one of the main reasons for wanting to limit the amount of measurements in the solution was to cut down on the cost of the computations the satellite selection algorithm needs to be computationally light as well. This was studied in Section 4.3, where it was shown that the complexity of the data-driven approach grows linearly with the number of visible satellites. For the quasi-optimal method the growth is quadratic, causing it to take up more computation time with the numbers of visible satellites there are available nowadays for multi-constellation receivers.

Conclusions and future work

The main goal of this work was to develop an algorithm for multi-constellation GNSS receivers that would select satellites out of the tracked ones to be used in the location solution. As the receiver has very limited computational resources, the complexity of the algorithm needed to be kept low. The work showed that optimizing the GDOP only lacks the important aspect of considering the pseudorange measurement errors. These have been shown to have an even greater influence on the accuracy of the position solution than optimizing the GDOP does [35].

The work began by exploratory analysis of GNSS data from near optimal situations. This analysis gave already some insight into the differences of the various satellite navigation systems as well as into the nature of the pseudorange residuals. These observations helped in shaping the algorithm that we proposed for the problem of satellite selection. The algorithm itself was developed using data science techniques to filter out bad pseudorange measurements and borrowed some earlier ideas to optimize the geometric dilution of precision of the solution set as well.

The approach we chose was shown to work very well when applied to real data measured from road tests in varying surroundings. Even with practically non-existent parameter tuning the algorithm was able to spot almost 90% of the bad pseudorange measurements, keeping the specificity, i.e., ability to hold on to the good measurements at over 90% level. As there was practically no parameter tuning done here, optimizing the weights and thresholds of the selection model would most likely improve the results even further. Another development aspect would be to have weights for each of the individual satellites instead of the global navigation satellite systems. This would however require far more data of sufficient quality.

The ability to filter out bad pseudorange measurements translated to improved location accuracy as well. The data-driven approach outdid the quasi-optimal method clearly and in addition was shown to have lower computational complexity with the present number of navigation satellites already.

All in all, the results achieved in this work proved encouraging enough to begin implementing the algorithm in actual receiver software to study the performance of the data-driven approach in action.

BIBLIOGRAPHY

[1] "European GNSS Service Centre," https://www.gsc-europa.eu/system-status/Constellation-Information, accessed: 3-October-2017.

[2] "Information and analysis center for positioning, navigation and timing," https://www.glonass-iac.ru/en/, accessed: 3-October-2017.

[3] *POSLV Specifications*, Applanix corporation, February 2017.

[4] N. Blanco-Delgado and F. D. Nunes, "Satellite selection method for multi-constellation GNSS using convex geometry," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 9, pp. 4289–4297, 2010.

[5] R. G. Brown, "A baseline gps raim scheme and a note on the equivalence of three raim methods," *Navigation*, vol. 39, no. 3, pp. 301–316, 1992.

[6] Y. Cheng, J. Dambeck, and F. Holzapfel, "Satellite Selection in Multi-GNSS Positioning," in *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, 2016, pp. 140–151.

[7] *BeiDou Navigation Satellite System Signal In Space Interface Control Document Open Service Signal*, China Satellite Navigation Office, November 2016, Version 2.1.

[8] S. Chuang, Y. Wenting, S. Weiwei, L. Yidong, Z. Rui, *et al.*, "GLONASS pseudorange inter-channel biases and their effects on combined GPS/GLONASS precise point positioning," *GPS solutions*, vol. 17, no. 4, pp. 439–451, 2013.

[9] R. D'Agostino and E. S. Pearson, "Tests for departure from normality. Empirical results for the distributions of b^2 and $\sqrt{b}$," *Biometrika*, vol. 60, no. 3, pp. 613–622, 1973.

[10] D. Dailey and B. Bell, "A method for GPS positioning," *IEEE transactions on aerospace and electronic systems*, vol. 32, no. 3, pp. 1148–1154, 1996.

[11] G. Di Giovanni and S. Radicella, "An analytical model of the electron density profile in the ionosphere," *Advances in Space Research*, vol. 10, no. 11, pp. 27–30, 1990.

[12] J. M. Dow, R. E. Neilan, and C. Rizos, "The International GNSS Service in a changing landscape of Global Navigation Satellite Systems," *Journal*

of Geodesy, vol. 83, no. 3, pp. 191–198, Mar 2009. [Online]. Available: https://doi.org/10.1007/s00190-008-0300-3

[13] *European GNSS (Galileo) Open Service Signal In Space Interface Control Document*, European Space Agency, November 2015, Issue 1.2.

[14] A. Fernández, P. D'Angelo, J. Diez, L. Marradi, and V. Gabaglio, "Navigation Algorithm Optimisation for Combined Galileo/GPS Receivers with the GRANADA Environment and Navigation Simulator," in *Proc. of the 18 th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2005)*, 2005, pp. 1939–1944.

[15] E. Grafarend and J. Shan, "Closed-form solution of the nonlinear pseudo-ranging equations GPS," *Artificial satellites, Planetary geodesy*, vol. 31, no. 3, pp. 133–147, 1996.

[16] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference and prediction*, 2nd ed. Springer series in statistics, New York, 2009.

[17] G. Hochegger, B. Nava, S. Radicella, and R. Leitinger, "A family of ionospheric models for different uses," *Physics and Chemistry of the Earth, Part C: Solar, Terrestrial & Planetary Science*, vol. 25, no. 4, pp. 307–310, 2000.

[18] H. Hopfield, "Two-quartic tropospheric refractivity profile for correcting satellite data," *Journal of Geophysical research*, vol. 74, no. 18, pp. 4487–4499, 1969.

[19] A. J. Izenman, *Modern multivariate statistical techniques*. Springer, 2008, vol. 1.

[20] E. Kaplan, *Understanding GPS - Principles and applications*, 2nd ed. Artech House, December 2005.

[21] J. A. Klobuchar, "Ionospheric time-delay algorithm for single-frequency gps users," *IEEE Transactions on aerospace and electronic systems*, no. 3, pp. 325–331, 1987.

[22] J. L. Leva, "An alternative closed-form solution to the GPS pseudo-range equations," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, no. 4, pp. 1430–1439, 1996.

[23] J. Li, A. Ndili, L. Ward, and S. Buchman, "GPS receiver satellite/antenna selection algorithm for the Stanford gravity probe B relativity mission," in *National Technical Meeting'Vision 2010: Present and Future*, 1999, pp. 541–550.

[24] M. Liu, M. Fortin, and R. Landry, "A recursive quasi-optimal fast satellite selection method for GNSS receivers," *Washington: Inst Navigation*, pp. 2061–71, 2009.

[25] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.

[26] *Global Positioning Systems Directorate System Engineering & Integration Interface Specification IS-GPS-200*, Navstar GPS Space Segment/Navigation User Interfaces, September 2013, Revision H.

[27] C.-W. Park and J. P. How, "Quasi-optimal satellite selection algorithm for real-time applications," in *14 th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 2001), Salt Lake City, UT*, 2001, pp. 3018–3028.

[28] A. Peng, G. Ou, and G. Li, "Fast satellite selection method for multi-constellation Global Navigation Satellite System under obstacle environments," *IET Radar, Sonar & Navigation*, vol. 8, no. 9, pp. 1051–1058, 2014.

[29] *GLONASS Interface Control Document*, Russian Institute of Space Device Engineering: Moscow, Russia, 2008, Edition 5.1.

[30] J. Saastamoinen, "Atmospheric Correction for Troposphere and Stratosphere in Radio Ranging of Satellites," in *The Use of Artificial Satellites for Geodesy*, ser. Washington DC American Geophysical Union Geophysical Monograph Series, S. W. Henriksen, A. Mancini, and B. H. Chovitz, Eds., vol. 15, 1972, p. 247.

[31] A.-L. Tao and S.-S. Jan, "Optimal Navigation with Multi-constellation GNSS: A Satellite Selection Algorithm," in *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, 2016, pp. 128–139.

[32] L. Wanninger, "Carrier-phase inter-frequency biases of GLONASS receivers," *Journal of Geodesy*, vol. 86, no. 2, pp. 139–148, 2012.

[33] M. Wei, J. Wang, and J. Li, "A new satellite selection algorithm for real-time application," in *Systems and Informatics (ICSAI), 2012 International Conference on*. IEEE, 2012, pp. 2567–2570.

[34] M. Zhang and J. Zhang, "A fast satellite selection algorithm: beyond four satellites," *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, no. 5, pp. 740–747, 2009.

[35] N. I. Ziedan, "Multipath and NLOS Signals Identification and Satellite Selection Algorithms for Multi-Constellation Receivers," in *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, 2016, pp. 521–533.