

Bachelor Thesis

Sensor Noise of the Humanoid NAO Robot and its Effects on Camera Projection

Yannik Meinken

Matriculation number: 4592225

August 16, 2023

- 1. Gutachter:** Dr. Tim Laue
- 2. Gutachter:** Dr. Alexander Fabisch

Faculty 4: Production Engineering
Study program: Systems Engineering



Yannik Meinken

Sensor Noise of the Humanoid NAO Robot and its Effects on Camera Projection

Sensorrauschen des humanoiden NAO Roboters und seine Einflüsse auf die Kameraprojektion

Bachelor Thesis, Faculty 4: Production Engineering

University Bremen, August 16, 2023

Abstract

All sensors that measure a real-world quantity are subject to some form of noise. To improve the quality of the observed state different measurements are combined, either to improve precision or because a single sensor can not observe the full state. To respect the uncertainty of a measurement a noise model is needed. Especially in complex systems such as robots, the measured data often can not be combined directly because it is in different reference systems. In such a case, the measurement and the associated uncertainty have to be transformed.

This thesis studies the noise characteristics of the proprioceptive of the *NAO⁶* robot via a long-time experiment. The data is analyzed for different typical noise characteristics resulting in a lower-bound approximation of the sensor noise.

To deduce the influence of the sensor noise on the world model of the robot a framework for uncertainty propagation is presented. It allows propagating uncertainty through the kinematic chain and the camera projection. Thereby it respects the unique properties of the transformations and resulting distributions.

That results in a more accurate description of the input data of subsequent filters, allowing for better performance.

Eigenständigkeitserklärung

Ich versichere, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe. Alle Teile meiner Arbeit, die wortwörtlich oder dem Sinn nach anderen Werken entnommen sind, wurden unter Angabe der Quelle kenntlich gemacht. Gleiches gilt auch für Zeichnungen, Skizzen, bildliche Darstellungen sowie für Quellen aus dem Internet. Die Arbeit wurde in gleicher oder ähnlicher Form noch nicht als Prüfungsleistung eingereicht. Die elektronische Fassung der Arbeit stimmt mit der gedruckten Version überein. Mir ist bewusst, dass wahrheitswidrige Angaben als Täuschung behandelt werden.

Erklärung zur Veröffentlichung von Bachelorarbeiten

Die Abschlussarbeit wird zwei Jahre nach Studienabschluss dem Archiv der Universität Bremen zur dauerhaften Archivierung angeboten. Archiviert werden:

1. Masterarbeiten mit lokalem oder regionalem Bezug sowie pro Studienfach und Studienjahr 10% aller Abschlussarbeiten
2. Bachelorarbeiten des jeweils ersten und letzten Bachelorabschlusses pro Studienfach u. Jahr.

- Ich bin damit einverstanden, dass meine Abschlussarbeit im Universitätsarchiv für wissenschaftliche Zwecke von Dritten eingesehen werden darf.
- Ich bin damit einverstanden, dass meine Abschlussarbeit nach 30 Jahren (gem. §7 Abs. 2 BremArchivG) im Universitätsarchiv für wissenschaftliche Zwecke von Dritten eingesehen werden darf.
- Ich bin nicht damit einverstanden, dass meine Abschlussarbeit im Universitätsarchiv für wissenschaftliche Zwecke von Dritten eingesehen werden darf.

Einverständniserklärung über die Bereitstellung und Nutzung der Bachelorarbeit in elektronischer Form zur Überprüfung durch Plagiatssoftware

Eingereichte Arbeiten können mit der Software *Plagscan* auf einen hauseigenen Server auf Übereinstimmung mit externen Quellen und der institutionseigenen Datenbank untersucht werden. Zum Zweck des Abgleichs mit zukünftig zu überprüfenden Studien- und Prüfungsarbeiten kann die Arbeit dauerhaft in der institutionseigenen Datenbank der Universität Bremen gespeichert werden.

- Ich bin damit einverstanden, dass die von mir vorgelegte und verfasste Arbeit zum Zweck der Überprüfung auf Plagiate auf den *Plagscan*-Server der Universität Bremen hochgeladen wird.
- Ich bin ebenfalls damit einverstanden, dass die von mir vorgelegte und verfasste Arbeit zum o.g. Zweck auf dem *Plagscan*-Server der Universität Bremen hochgeladen u. dauerhaft auf dem *Plagscan*-Server gespeichert wird.
- Ich bin nicht damit einverstanden, dass die von mir vorgelegte u. verfasste Arbeit zum o.g. Zweck auf dem *Plagscan*-Server der Universität Bremen hochgeladen u. dauerhaft gespeichert wird.

Mit meiner Unterschrift versichere ich, dass ich die oben stehenden Erklärungen gelesen und verstanden habe. Mit meiner Unterschrift bestätige ich die Richtigkeit der oben gemachten Angaben.

Bremen, August 16, 2023



.....
(Yannik Meinken)

Contents

1	Introduction	1
1.1	Context	1
1.2	Motivation	2
1.3	Goals	3
1.4	Related Work	3
2	Sensor Noise of the Nao	5
2.1	Sensor Description	5
2.2	Analysis of Possible Error Sources	6
2.3	Long-Time Experiments	8
2.3.1	Setup	8
2.3.2	Analysis	8
2.4	Motion-Tracking	14
3	Uncertainty Propagation	18
3.1	Kinematic Chain	18
3.1.1	$SE(3)$	18
3.1.2	Uncertainty in $SE(3)$	23
3.1.3	Differences to $SO(3) + \mathbb{R}^3$	25
3.1.4	Propagation of Uncertainty in $SE(3)$	25
3.1.5	Own Implementation	27
3.2	Camera to Field	28
3.2.1	Projection Properties	28
3.2.2	Unscented Transform	29
3.2.3	Own Implementation	29
3.2.4	Differences and Advantages to Previous Approach	30
4	Future Work	33
4.1	Sensor Noise	33
4.2	Uncertainty Propagation	34
4.2.1	Kinematic Chain	34
4.2.2	Camera to Field	35
4.2.3	Others	35
4.3	New Opportunities	36

1 Introduction

Due to limited perception capabilities and sensor noise in the context of robotics, one should never fully rely on the measured state but keep in mind its uncertainty. There exist a variety of filters that try to compensate for these inaccuracies by combining measurements from different sensors or times. If the quality of the measurements differs, either due to different sensors or external factors that influence the measurement quality, these filters need a measure for the uncertainty of the different measurement points.

This thesis aims to provide methods to provide those measures in the context of a humanoid, soccer-playing robot.

1.1 Context

This thesis is connected to the student project *B-Human*¹, the robot soccer team of the *University of Bremen* (Uni Bremen) and the *German Research Center for Artificial Intelligence* (DFKI). It participates in the *Soccer Standard Platform League* (SPL) of the *RoboCup*² since 2009 (Thomas Röfer et al. 2022). The RoboCup is an international initiative to promote robotics research by providing a standard problem that captures the challenges of the real world while remaining feasible for a wide range of research groups (Kitano et al. 1997). Since 1997 yearly competitions (except 2020) in several leagues with different focus areas, are held together with a conference (Arne Hasselbring 2020). This allows to compare different approaches and share findings.

During an SPL match seven robots of each team play autonomously on a nine by six meter field (RoboCup Technical Committee 2023). They have to

¹<https://b-human.de>

²<https://www.robocup.org>

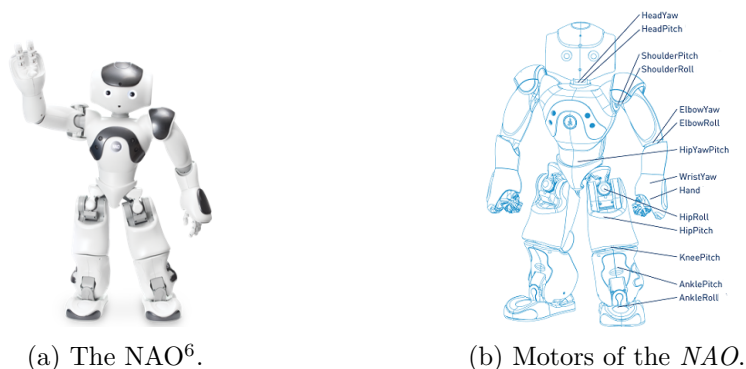


Figure 1.1: The NAO robot. Images from SoftBank Robotics (2023a,b).

perceive and model the environment and their position in it and make decisions based on this model.

Currently, the NAO⁶ is used in the SPL as it comes closest to meeting the requirements. The humanoid robot is 0.57 m high and weighs about 4.5 kg. It consists of 25 joints which are shown in figure 1.1b. They are actuated by DC motors and have integrated rotary encoders that measure the actually achieved angle. Other proprioceptive sensors include an *Inertial measurement unit* (IMU) for measuring acceleration and angular velocity and force sensors on the foot soles (Gouaillier et al. 2009). Exteroceptive sensing is provided by two cameras, an ultrasonic range sensor and microphones (SoftBank Robotics 2023a).

1.2 Motivation

The orientation of the robot cameras has a great impact on the projection of percepts in the image to the ground plane (Jüngel, Mellmann, and Spranger 2008). These percepts are input for several filters that estimate, e.g., the position of the robot on the field or the ball. These filters need, as previously stated, a measure for the uncertainty of the percepts. Currently, this is done via constant factors that are only adapted for the distance that the measurement has to the robot combined with a binary decision if the robot is walking or standing. Several aspects are ignored, like the inaccuracy of the camera pose

during walking, influences from external disturbances, or the position of the projected point inside the image. Additionally, the used values are guessed and not based on the actual performance of the sensors and filters.

It impacts the performance of the subsequent filters that take these dynamics into account and limits the physical meaning of the modeled uncertainty. This increases the difficulty to set new parameters that interact with the uncertainty of the models as they depend on the previously guessed values and can therefore only be guessed as well.

1.3 Goals

This thesis aims to improve the situation explained above. For that, a lower bound of the random noise from the proprioceptive sensors should be determined experimentally. The uncertainty should then be propagated through the kinematic chain. Special focus should thereby lie on accurately modeling the type of distribution. Lastly, the effect of an uncertain camera pose on the projection from the image to the field should be modeled.

To examine the noise of the IMU, a long time series will be recorded and analyzed. The joint sensors should be compared to a ground truth provided by a motion tracking system. The *Probabilistic product-of-exponentials* (PPOE) (Meyer, Strobl, and Triebel 2022) should be used to propagate uncertainty through the kinematic chain and handle the *banana-shape* distribution. For the uncertainty propagation through the camera projection, the *Unscented transform* (UT) will be used.

1.4 Related Work

Jünger, Mellmann, and Spranger (2008) addressed the error introduced on the projection from a camera image to a ground plane for the *Sony Aibo* robot. They also determined the kinematic chain as one possible error source, but instead of modeling the uncertainty they tried to reduce it by exploiting known properties of objects perceived in the image.

Meyer, Strobl, and Triebel (2022) introduced the PPOE to model and propagate the uncertainty in a robot kinematic chain, making use of the representation of uncertainty for the $SE(3)$ Lie group presented by Barfoot and Furgale (2014). The propagation of uncertainty through the kinematic chain presented in this thesis uses the same approach. In addition to that, the obtained uncertainty of the camera pose is further propagated through the camera projection to obtain a measure of uncertainty of a perceived point on the ground.

Laue et al. (2009) lists some particularities of the IMU in the NAO and how its data is used to calculate the torso pose of the robot. This thesis further investigates the properties of the IMU of the NAO⁶ to derive a noise model that improves the performance of the filters presented by Laue et al. (2009) further.

2 Sensor Noise of the Nao

In this chapter, the error of the rotary encoders at the joints and the IMU are examined. First, the different sensor types are introduced, and the already known properties are put into context. Then, possible sources of errors that can influence the measurements are identified. Afterward, the execution and results of a long-time experiment to evaluate which noise types occur and which have the most impact are presented. The last section deals with the possible application of a motion-tracking system and its limitations.

2.1 Sensor Description

An IMU consists of accelerometers that measure the linear acceleration in one direction, and gyroscopes, that sense the rotational velocity about one axis. They are often used to track the position and, more importantly, the orientation of the sensor. There exists a variety of different measurement principles (Woodman 2007). Although there is no information from the manufacturer, it is most likely that the one used in the NAO uses *Micromachined electro-mechanical systems* (MEMS) technology as it is compact, rugged and cheap (Woodman 2007). The IMU in the NAO⁶ consists of one accelerometer and gyroscope for each of the three axes. They are located on the *inertial board* that communicates with the *chest board*. Additionally, the orientation is already computed from the raw sensor data (SoftBank Robotics 2023a). The computed orientation does not include the vertical Z-axis, this might be a relict from older versions of the NAO that did not have a gyroscope for that axis (SoftBank Robotics Europe 2023). One specialty of the NAO is that the data from the accelerometer and the gyroscope is not updated every frame but alternates between these two. This effectively reduces the frequency by a factor of two

(Jesse Richter-Klug 2018). The manufacturer gives no further information about the expectable uncertainty.

On each active joint of the NAO, a rotary encoder is attached that measures the angle through the Hall effect. The measured angle is represented as a 12 bit value (SoftBank Robotics 2023a) resulting in a maximum precision of 1.53×10^3 rad.

$$\frac{2\pi\text{rad}}{2^{12}} = 1.53 \times 10^3 \text{ rad} \quad (2.1)$$

The noise resulting from this is called discretization noise.

2.2 Analysis of Possible Error Sources

Woodman (2007) identified different error sources for IMUs, with a focus on the cheap MEMS as probably used in the NAO (see section 2.1). They found five error sources that are listed below.

- Bias
- White Noise
- Temperature Effects
- Calibration
- Bias Instability

This thesis focuses on white noise and bias instability, as the temperature effects can not be observed with the experimental possibilities. The errors introduced by the bias and calibration are not considered as they can be corrected by calibration and therefore are zero for a lower-bound approximation.

White noise is a random noise with an expected value of zero that follows a normal distribution. For MEMS IMUs the source is mostly high-frequency noise caused by thermo-mechanical effects.

Bias instability, on the other hand, describes a slow change of the bias over time. It can be interpreted as the standard deviation of the bias after a given time (Woodman 2007).

Additionally to these errors, one possible error source might be a time

delay between a change of the observed variable and the time the received measurement changes. If the delay for all sensor types was the same, and therefore the data comes in the correct order, a subsequent filter would have the same delay. Otherwise, the correct order has to be recreated, accepting the maximal delay or incorporating the delay by increasing the uncertainty of the delayed measurements. The slower update rate of the raw IMU data (see section 2.1) is one example of a delayed measurement.

Another type of noise that every digital sensor has is quantization errors due to their discrete representation of the measured value.

The sensors for the joint angles could potentially have the same error sources mentioned above. But there might be other error sources that are specific to joint angles. Depending on the mounting point of the sensors, there might be play due to backlash in the gears that is not measured. This would cause an error between the measured joint angle and the actual one. This type of error is hard to model as the amount of backlash depends on the wear of the joint, which in turn changes over time, between robots and even depends on the part of the gear that is currently in mesh. Furthermore, the error is correlated, which violates the assumptions of most popular filtering methods. As a movement of a joint by hand without moving the motor can be observed in the sensor data it can be assumed that at least most of the joint play is measured (Reichenberg 2020).

Besides the active joints that are measured, there might be unobserved deformations or passive joints that introduce further uncertainty. In the case of the NAO, this might be the connection between legs and torso that in some worn-out robots is partially broken, resulting in play. Another possible error source that affects the kinematic chain are manufacturing inaccuracies such as wrong link dimensions or rotated sensors.

As the goal is to find a lower-bound approximation of the uncertainty this thesis will only focus on random errors and neglect biases and other errors that could be compensated by calibration.

2.3 Long-Time Experiments

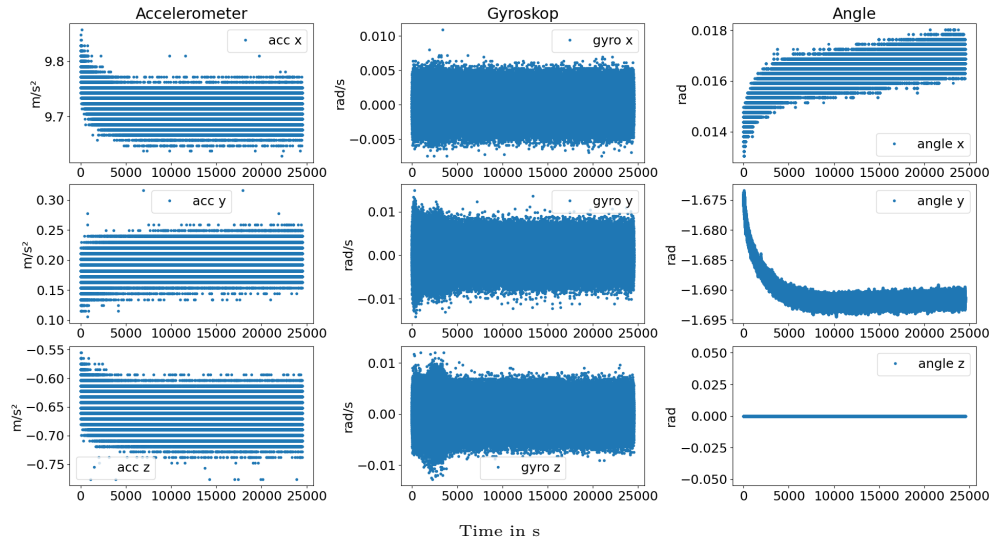
This section describes how a long-time experiment to investigate random, additive noise in a static environment was conducted. The data is first displayed using *Matplotlib* (Hunter 2007) and intuitively analyzed before the results are validated using the Allan variance.

2.3.1 Setup

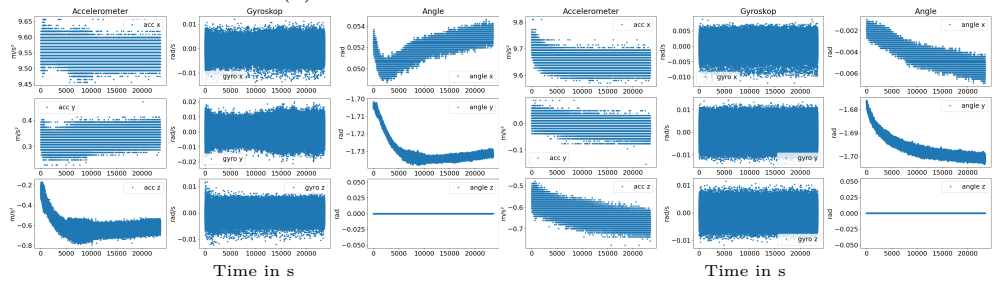
During the experiment, the sensor readings for three different robots were tracked over an extended period (more than 6 h). To capture not only the high-frequency noise types but also slow changes of the bias. Additionally, a longer sample period results in a more accurate determination of the parameters for high-frequency noise. In order to have a ground truth to compare the measured values against, the robots laid on the ground and were not moved during the experiment. Therefore, the expected measurement of the gyroscope was zero. The accelerometer still measures an acceleration due to the gravitational force, but it should stay constant.

2.3.2 Analysis

Based on the progression of the time series shown in figure 2.1, first statements over the types of noise present can be derived. The accelerometer data seems to be subject to white noise as well as to a bias drift. But as the drift of the bias seems to be mostly at the start, it might also be caused by a preceding filter with an infinite impulse response so that the previous measurements still have an impact. Another possible source might be temperature effects as the sensor heats up before reaching saturation. Additionally, it is obvious that the data is discrete as only specific values are reached, but as the gaps between them are relatively small the other noise types seem to have a greater effect. For the gyroscope, the only prominent type of noise seems to be white noise. The computed angle is not important as it is only used for initial calibration, but the strong change over time can possibly be explained by the dependency from the accelerometer data, combined with another (probably infinite impulse response) filtering stage.



(a) IMU data from FriedrichWilhelm.



(b) IMU data from Daskänguru.

(c) IMU data from OttoVon.

Figure 2.1: Raw data of the IMUs during the experiment.

To validate this first intuition about the existing noise types, Allan variance can be used (Woodman 2007). It allows to analyze the dependency between deviation and averaging time. It was first developed to analyze oscillators but can be adapted to a variety of devices such as IMUs (El-Sheimy, Hou, and Niu 2008). To compute the Allan variance for a given averaging time T , the time series is split into N bins of the size of the averaging time. The average of each bin $\bar{\Omega}_i$ is computed, and the change between consecutive bins is calculated. The Allan variance is computed as the variance of these differences.

$$\sigma_{Allan}^2(T) = \frac{1}{2(N-1)} \sum_{i=1}^{N-1} (\bar{\Omega}_i - \bar{\Omega}_{i+1})^2 \quad (2.2)$$

The square root of the Allan variance is called Allan deviation (σ_{Allan}) (El-Sheimy, Hou, and Niu 2008; Woodman 2007).

To achieve a higher confidence, especially for high averaging times, there exists an overlapping Allan deviation that is commonly used in practice (Anders E.E. Wallin et al. 2019).

When plotting the Allan deviation of a time series as a function of averaging time different noise types cause different slopes on a log-log scale. El-Sheimy, Hou, and Niu (2008) derived how different noise types typical for an IMU express on the Allan deviation plot. White noise is represented as a slope with a gradient of $-\frac{1}{2}$ and the numerical value of the deviation can be read off at $T = 1$. An instability of the bias is represented by a flat region in the plot, the quantifying value is the minimum.

Based on the observations above the Allan plot for the accelerometers should descend with a slope of $-\frac{1}{2}$ before reaching a minimum. For the gyroscope, it should be descending with a slope of $-\frac{1}{2}$ everywhere.

The behavior of the Allan plots for the gyroscope data shown in figure 2.2 fits the expectations. The slight deviations for higher averaging times could be caused by the lower confidence of the Allan deviation for high averaging time as fewer bins can be created. In the plots for the accelerometers, the slope of $-\frac{1}{2}$ representing white noise dominates for short averaging times as expected. For longer averaging times, instead of reaching a flat, it ascends at a slope of

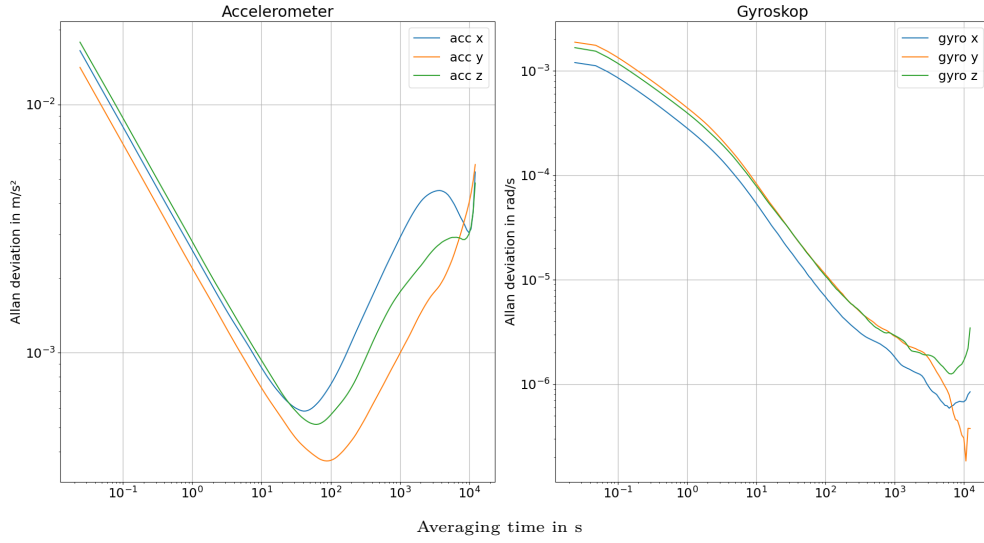


Figure 2.2: Allan plots for the IMU data of FriedrichWilhelm.

$\frac{1}{2}$, indicating not a bias instability but another type of noise. According to El-Sheimy, Hou, and Niu (2008) the corresponding noise type is rate random walk.

The values for the joint variables shown in figure 2.3 seem to jump between discrete values. The magnitude of the jumps is the same as the maximum resolution of the sensor given by the manufacturer (see section 2.1). This indicates that the main random error in a constant scenario is the quantization noise. To find additional error types the Allan variance can be used, in figure 2.4 a slope of $-\frac{1}{2}$ representing white noise can be observed.

The next step is to derive explicit values that can be used by filters. For the IMU the covariance matrices for the gyroscope and accelerometer measurements were computed. The values for FriedrichWilhelm are shown in table 2.1. The variances on the main diagonal are, as expected, at least one order of magnitude higher than the covariances. This means it is valid to assume that the errors of the different axes are uncorrelated. In order to obtain a single value, the variance for each axis and robot are plotted in figure 2.5. To detect outliers the distance of each sample to the median is computed and compared to the median of these distances (MAD) (Rousseeuw 1991). If the point differs more




Figure 2.3: Measurements of the joint angles of FriedrichWilhelm during the experiment.




Figure 2.4: Allan plots of the joint measurements of FriedrichWilhelm.

x	y	z
4.13e-4	-4.53e-5	6.55e-5
-4.53e-5	2.31e-4	-4.91e-5
6.55e-5	-4.91e-5	3.77e-4

(a) Accelerometer covariance
in $(\text{m/s}^2)^2$.

x	y	z
2.13e-6	-4.91e-8	1.51e-7
-4.91e-8	5.23e-6	3.37e-8
1.51e-7	3.37e-8	4.07e-6

(b) Gyrometer covariance
in $(\text{rad/s})^2$.

Table 2.1: Covariance matrices for the IMU data from FriedrichWilhelm.

than three times the MAD from the median (indicated by the red lines) it is considered an outlier. By using the median instead of the mean for outlier detection it is far more robust against outliers (Rousseeuw 1991). Except for a few outliers, the values seem to be roughly the same, so after removing the outliers, the mean is calculated to obtain a single value that can be used by filters. Because the outliers were removed the higher accuracy of the mean can be utilized despite its sensitivity to outliers (Rousseeuw 1991).


As the discretization noise was identified to be the main error source for the joint angles, and it is specified by the manufacturer the specified values can be used. However, the *Probability density function* (PDF) of discretization noise is a uniform distribution which is described by the width Δa instead of a variance σ . The variance can be computed using the definition of variance.

$$\begin{aligned}
 \sigma &= \int_{-\infty}^{\infty} (y - \bar{y})^2 p(y) dy \\
 &= \int_{-\frac{\Delta a}{2}}^{\frac{\Delta a}{2}} y^2 \frac{1}{\Delta a} dy \\
 &= \frac{\Delta a^2}{12} \\
 &= \frac{(1.53 \times 10^{-3} \text{ rad})^2}{12} = 1.95 \times 10^{-7} \text{ rad}^2
 \end{aligned} \tag{2.3}$$


2.4 Motion-Tracking

A motion-tracking system can track the 3D position of markers that are observed by at least two cameras. By using special targets that consist of multiple markers also the orientation can be observed. Such a system should be used in order to have a ground truth measurement for joint angles in motion by mounting a target before and behind the measured joint and assuming that the movement is constrained to the axis defined by the joint. Due to the lack of a common reference frame, the data from the tracking system could only be used as a relative measurement. With this setup, previously unobservable joint play could be measured.

To get a first impression of the accuracy of the system a sequence of



(a) Accelerometer



(b) Gyroscope

Figure 2.5: IMU variances of different robots and axis. The red lines indicating the range of usable data, points outside are considered outliers and displayed in red. The green line marks the calculated mean.




Figure 2.6: Sketch of angle computation using position measurements.

measurements of two markers that were fixated together was taken. Because they were fixated together, the relative position and orientation of the markers should not change. During the measurement, the markers were first recorded in the static case and then slowly moved. During the phase where the targets are slowly moved the measured relative orientation of them changed up to 2 degrees. This is multiple orders of magnitude more than the inaccuracy of the joint sensors, therefore this method is not usable to investigate the error of the sensor.

To improve the accuracy for measuring the joint angle one leg was suspended (as seen in figure 2.7) to allow to measure the angle of the hip joint via the measured displacement of the end of the lever. The angle φ as seen in figure 2.6 can be computed as follows:

$$\varphi = \arccos\left(\frac{1 - l^2}{2r^2}\right) \quad (2.4)$$

Although this only allows measuring a specific joint it might provide insight that can be generalized to the other joint as well.

To validate the robustness of this approach the impact of slightly different values for the radius or traveled distance, within an expected error range, can be compared for an exemplary input. One can see in table 2.2 that an error of the radius of 1 cm causes the angle to be wrong by more than a tenth degree and therefore more than the discretization error of the joint sensor. An error of the radius by 1 cm is plausible as it is measured by hand with



Figure 2.7: Setup of the Motion-tracking system.

r	l	φ
1645 mm	500 mm	17.48°
1635 mm	500 mm	17.59°
1645 mm	501 mm	17.52°

Table 2.2: Error analysis for the suspended leg.

a ruler. To investigate the influence of the accuracy of the motion-tracking system on the measured angle an error of 1 mm was assumed. Although the manufacturer does not provide a specification for the accuracy this value is based on experience in practice. The angular error introduced is in the same order of magnitude as the discretization error.

Overall, is the accuracy of this method still higher than the discretization, but it might be good enough to find additional error sources that introduce a higher error.

3 Uncertainty Propagation

This chapter describes how uncertainty in the system is handled and propagated to different reference frames. The first section covers the kinematic chain and its effects on proprioceptive measurements. The concept of representing 3-dimensional poses as members of the Special Euclidean Group $SE3$ is presented and the usual model of additive uncertainty is expanded. The second section deals with the propagation of uncertainty from the camera pose to a percept that is projected from a camera image to the field plane.

3.1 Kinematic Chain

The kinematics of most robots can be represented as a chain of rigid bodies that are connected via joints (revolute or prismatic). The transformation from one link to another consists of a translational part as well as an orientation, together called a *pose*. If the transformation from A to B (\mathbf{T}_1) as well as the transformation from B to C (\mathbf{T}_2) is known, the transformation \mathbf{T}_{12} from A to B can be derived by chaining the intermediate transformations, as shown by equation (3.1) and illustrated in figure 3.1. It is important to note here that the order of operations is important, it is not commutative!

$$\mathbf{T}_{12} = \mathbf{T}_1 \circ \mathbf{T}_2 \tag{3.1}$$

3.1.1 $SE(3)$

As the orientation can not be expressed as a vector space, because it is not commutative, this applies to a pose as well. One possibility to express an orientation is as a rotation matrix. For an orientation in three dimensions




Figure 3.1: Compounding of poses.

the rotation matrix is a three by three matrix. It is an overparameterized representation constrained by the fact that it has to be an orthogonal matrix with determinant 1. This fact makes it very easy to compute the inverse of a rotation matrix \mathbf{R} , as $\mathbf{R}\mathbf{R}^T = \mathbf{I}$ and therefore its inverse \mathbf{R}^{-1} is:

$$\mathbf{R}^{-1} = \mathbf{R}^T \quad (3.2)$$

The group of 3-D rotations is also known as the *Special orthogonal group* ($SO(3)$).

To extend this to a proper 3-dimensional pose \mathbf{T} which is also known as the *Special euclidian group* ($SE(3)$), the 3×3 rotation matrix can be extended in both dimensions by one to form a four-by-four matrix that combines rotation and translation. In the upper left corner, the matrix representation contains a rotation matrix $\mathbf{R} \in SO(3)$ as described above, the first three elements of the rightmost column contain the translation vector $\mathbf{t} \in \mathbb{R}^3$ and the lower row consists of three zeros and a one in the down-right corner.

$$\mathbf{T}_{\in SE(3)} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \quad (3.3)$$

This representation allows to compound poses via multiplication according to equation (3.1) where the compounding operation \circ is matrix multiplication. The inversion can no longer be achieved by simple transposition, but the effect

from the rotation on the translation has to be considered.

$$\mathbf{T}^{-1} = \begin{bmatrix} \mathbf{R}^T & -\mathbf{R}^T \mathbf{t} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \quad (3.4)$$

This is still much faster than classical matrix inversion.

As $SE(3)$ and $SO(3)$ are Lie groups (and hence manifolds) (Blanco 2014), various mathematical concepts related to them, most notably the concept of Lie algebras, can be used.

Each Lie group has an associated Lie algebra which is the tangent space around the identity element of the group. In contrast to the group itself it can be expressed as a vector space. It is possible to map elements from the algebra onto the associated group with the help of the exponential map.

As $SO(3)$ is a part of $SE(3)$ it will be considered first. The Lie algebra associated with $SO(3)$ is $\mathfrak{so}(3)$ which also can be expressed as a matrix $\mathbf{M} \in \mathbb{R}^{3 \times 3}$, but is subject to different constraints¹. But, as it is a vector space, each element can also be expressed as a vector $\phi \in \mathbb{R}^3$. The operator \wedge (as defined by equation (3.5)) maps from the vector representation to the skew-symmetric matrix whereas the operator \vee reverses that.

$$\phi^\wedge := \begin{bmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \end{bmatrix}^\wedge = \begin{bmatrix} 0 & -\phi_3 & \phi_2 \\ \phi_3 & 0 & -\phi_1 \\ -\phi_2 & \phi_1 & 0 \end{bmatrix} \quad (3.5)$$

Each element of $\mathfrak{so}(3)$ can be constructed by a linear combination of the three base elements.

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}^\wedge \quad \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}^\wedge \quad \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}^\wedge \quad (3.6)$$

Which corresponds to a rotation around the corresponding base vector.

¹ \mathbf{M} has to be skew-symmetric

As stated previously the exponential map transforms elements from the algebra onto the group, for $SO(3)$ it can be expressed as:

$$\begin{aligned} \exp : \mathfrak{so}(3) &\mapsto SO(3) \\ \exp(\phi^\wedge) &= \mathbf{I}_{3 \times 3} + \frac{\sin(|\phi|)}{|\phi|} \phi^\wedge + \frac{1 - \cos(|\phi|)}{|\phi|^2} \phi^\wedge^2 \end{aligned} \quad (3.7)$$

It can be noted, that $|\phi|$ describes the amount of rotation around the axis given by the direction of ϕ .

The inverse of this operation is given by the logarithmic map \ln which maps a group element $\mathbf{R} \in SO(3)$ to the algebra:

$$\begin{aligned} \ln : SO(3) &\mapsto \mathfrak{so}(3) \\ \ln(\mathbf{R}) &= \frac{\theta}{2 \sin(\theta)} (\mathbf{R} - \mathbf{R}^T) \\ \text{with } \cos(\theta) &= \frac{\text{tr}(\mathbf{R}) - 1}{2} \end{aligned} \quad (3.8)$$

When implementing these, special care has to be given to avoid division by zero.

For the $SE(3)$ case, the associated algebra $\mathfrak{se}(3)$ can also be expressed as a four-by-four matrix which is subject to a different set of constraints. But as $\mathfrak{se}(3)$ is a vector space it can also be expressed as a vector $\xi \in \mathbb{R}^6$, consisting of a translational part $\mathbf{p} \in \mathbb{R}^3$ and rotation ϕ . By overloading the operators \wedge and \vee according to equation (3.9) the same mapping as for $\mathfrak{so}(3)$ can be achieved.

$$\xi^\wedge := \begin{bmatrix} \mathbf{p} \\ \phi \end{bmatrix}^\wedge = \begin{bmatrix} \phi^\wedge & \mathbf{p} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix} \quad (3.9)$$

The exponential map is given as:

$$\begin{aligned} \exp : \mathfrak{se}(3) &\mapsto SE(3) \\ \exp(\xi^\wedge) &= \begin{bmatrix} \exp(\phi^\wedge) & \mathbf{V}\mathbf{p} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \\ \text{with } \mathbf{V}_{\in \mathbb{R}^{3 \times 3}} &= \mathbf{I}_{3 \times 3} + \frac{1 - \cos(|\phi|)}{|\phi|^2} \phi^\wedge + \frac{|\phi| - \sin(|\phi|)}{|\phi|^3} \phi^\wedge^2 \end{aligned} \quad (3.10)$$

Note how the exponential mapping of $SO(3)$ (equation (3.7)) is used for the three times three upper left part. It is also worth stressing that the rotational part affects the translation via \mathbf{V} .

The logarithmic map as the inverse of \exp can be computed by:

$$\begin{aligned} \ln SE(3) &\mapsto \mathfrak{se}(3) \\ \ln(\mathbf{T}) &= \begin{bmatrix} \ln(R) & \mathbf{V}^{-1}\mathbf{t} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix} \end{aligned} \quad (3.11)$$

Most formulas in this section were adapted based on the work of Blanco (2014) where a more detailed explanation can be found.

Apart from the tangential space around the identity of the group, it is also possible to form a tangential space around any other element of the group because it is a smooth manifold. In these, the classical tools for vector spaces can be used. To transform a vector from the tangent space around a group element \mathcal{X} to the group algebra (tangent space around identity) the adjoint action can be used. As this is a linear action it can be expressed as matrix multiplication, where the matrix depends on the point around which the tangent space is utilized. For $SE(3)$ the adjoint matrix for a pose \mathbf{T} as defined in equation (3.3) can be computed as (Solà, Deray, and Atchuthan 2021):

$$\mathbf{Ad}_{\mathbf{T}} = \begin{bmatrix} \mathbf{R} & \mathbf{t}^\wedge \mathbf{R} \\ \mathbf{0}_{3 \times 3} & \mathbf{R} \end{bmatrix} \in \mathbb{R}^{6 \times 6} \quad (3.12)$$

The operator \wedge that transforms a 3-dimensional vector in a skew-symmetric matrix is here used on the translational part rather than the orientation as

before. It is worth pointing out that the position of the $\mathbf{0}_{3 \times 3}$ submatrix depends on the definition of ξ . For our choice in equation (3.9) with the translational component in the upper half, the $\mathbf{0}_{3 \times 3}$ submatrix has to be in the lower left corner. If \mathbf{p} and ϕ would be flipped it would be flipped with the upper right submatrix.

3.1.2 Uncertainty in $SE(3)$

The standard approach to define an uncertain variable x is to *add* a small perturbation δ to the noise-free value \bar{x} .

$$x = \bar{x} + \delta \tag{3.13}$$

Where δ follows a zero-mean probability distribution (Gaussian, Uniform, ...). This approach is not directly suitable for variables in $SE(3)$ as it is not a vector space and addition is not defined on it (Barfoot and Furgale 2014). Barfoot and Furgale (2014) propose to express the uncertainty for $SE(3)$ in the tangent space associated with the noise-free point, which meets the requirement of a vector space, and map it into the group via the exponential map. This results in a disturbance around the identity, which is then applied to the noise-free value via multiplication. Barfoot and Furgale (2014) propose to apply it via left-side multiplication which results in a disturbance with respect to the global frame. Contrary Meyer, Strobl, and Triebel (2022) suggest applying it via right multiplication to represent local perturbations in a robot's kinematic chain.

$$\mathcal{X} = \underbrace{\exp(\xi)\bar{\mathcal{X}}}_{\text{global perturbation}} \qquad \mathcal{X} = \bar{\mathcal{X}} \underbrace{\exp(\xi)}_{\text{local perturbation}} \tag{3.14}$$

In general, the better suited representation depends on the problem at hand and is not always obvious (Brossard, Bonnabel, and Condomines 2017). The rest of this thesis uses the right multiplicative representation as suggested by Meyer, Strobl, and Triebel (2022) for kinematic chains.

To simplify notation the \boxplus operator introduced by Hertzberg et al. (2013)




Figure 3.2: Visualization of a Gaussian on a Lie group. Image from Bourmaud et al. (2013).

can be defined for $SE(3)$ as:

$$\begin{aligned} \boxplus : SE(3) \times \mathbb{R}^6 &\rightarrow SE(3) \\ \mathcal{X} \boxplus \xi &= \mathcal{X} \exp(\xi^\wedge) \end{aligned} \quad (3.15)$$

The perturbation ξ can be modeled by a zero-mean Gaussian distribution which enables the use of many common techniques. This is only possible because ξ lies in a vector space.

$$p(\xi) = \mathcal{N}(0, \Sigma) \quad (3.16)$$

where $\Sigma \in \mathbb{R}^{6 \times 6}$ is the covariance matrix.

The different steps on how to define a random distribution on a Lie group is well illustrated by Bourmaud et al. (2013) on a simple example with only one degree of freedom. Figure 3.2 shows how a Gaussian on the Lie algebra is mapped to the group via the exponential transform and afterwards moved.

3.1.3 Differences to $SO(3) + \mathbb{R}^3$

Hertzberg et al. (2013) made an example where they represented a pose by compounding the manifolds \mathbb{R}^3 for translation and $SO(3)$ for orientation. The \boxplus operator was defined component-wise. This means that an isolated disturbance of the orientation does not affect the position and vice versa. Correlations between them are only modeled in the corresponding elements of the covariance matrix (upper left or lower right 3×3 submatrix). When only looking at the translation the confidence ellipsoids in the tangent space keep their general geometry when transformed.

In the case of $SE(3)$, the transformation is done via the exponential map which is a nonlinear function and the orientation directly affects the translational part (see equation (3.10)). Due to this, the ellipsoids in the tangent space are warped into a *banana-like* shape during the transformation.

Which representation fits better most likely depends on the problem at hand. If the translation heavily depends on the orientation in a way captured by the exponential map the representation utilizing it better fits the real distribution (Long et al. 2012). On the other hand, if position and orientation are independent of each other, or the dependency is unknown, the composition of \mathbb{R}^3 and $SO(3)$ likely performs better.

3.1.4 Propagation of Uncertainty in $SE(3)$

As already discussed, the compound of two poses can be computed via matrix multiplication. But for the goal of this thesis also the combination of the associated uncertainties has to be considered. When considering the example from equation (3.1) but now with \mathbf{T}_1 and \mathbf{T}_2 as uncertain transformations as modeled by equation (3.14) with the corresponding $\boldsymbol{\xi}_1 \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_1)$ and $\boldsymbol{\xi}_2 \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_2)$. It is reasonable to assume that the new expected value is the compound of the expected values (Barfoot and Furgale 2014)

$$\bar{\mathbf{T}}_{12} = \bar{\mathbf{T}}_1 \bar{\mathbf{T}}_2 \tag{3.17}$$

But to also propagate the uncertainty, the full model derived in section 3.1.2 has to be used.

$$\mathbf{T}_{12} = \mathbf{T}_1 \mathbf{T}_2 \quad (3.18)$$

$$\bar{\mathbf{T}}_{12} \boxplus \hat{\boldsymbol{\xi}}_{12} = (\bar{\mathbf{T}}_1 \boxplus \hat{\boldsymbol{\xi}}_1)(\bar{\mathbf{T}}_2 \boxplus \hat{\boldsymbol{\xi}}_2) \quad (3.19)$$

$$\bar{\mathbf{T}}_{12} \exp(\hat{\boldsymbol{\xi}}_{12}) = \bar{\mathbf{T}}_1 \exp(\hat{\boldsymbol{\xi}}_1) \bar{\mathbf{T}}_2 \exp(\hat{\boldsymbol{\xi}}_2) \quad (3.20)$$

$$= \bar{\mathbf{T}}_1 \bar{\mathbf{T}}_2 \exp((\mathbf{Ad}_{\bar{\mathbf{T}}_2^{-1}} \hat{\boldsymbol{\xi}}_1)^\wedge) \exp(\hat{\boldsymbol{\xi}}_2) \quad (3.21)$$

$$\exp(\hat{\boldsymbol{\xi}}_{12}) = \exp((\mathbf{Ad}_{\bar{\mathbf{T}}_2^{-1}} \hat{\boldsymbol{\xi}}_1)^\wedge) \exp(\hat{\boldsymbol{\xi}}_2) \quad (3.22)$$

$$\hat{\boldsymbol{\xi}}_{12} = \ln \left(\exp((\mathbf{Ad}_{\bar{\mathbf{T}}_2^{-1}} \hat{\boldsymbol{\xi}}_1)^\wedge) \exp(\hat{\boldsymbol{\xi}}_2) \right) \quad (3.23)$$

As the uncertainty for each pose is defined in the corresponding tangent space they can not be directly combined but first have to be transformed onto the same tangent space. This can be done via the adjoint action and therefore multiplication with the adjoint matrix as seen in 3.21. With the help of equation (3.17), the noise-free values can be eliminated and only the random perturbations remain. Now it can be solved for the perturbation of the new, compounded pose (Meyer, Strobl, and Triebel 2022).

The compounding of two matrix exponentials can be done using the *Baker-Campbell-Hausdorff formula* (DCH), which provides an infinite series (Barfoot and Furgale 2014). The result can be inserted into the definition of covariance and evaluated up to the desired order. Considering only the first two terms the compounded covariance can be expressed as (Meyer, Strobl, and Triebel 2022)

$$\boldsymbol{\Sigma}_{12} = \mathbf{Ad}_{\bar{\mathbf{T}}_2^{-1}} \boldsymbol{\Sigma}_1 \mathbf{Ad}_{\bar{\mathbf{T}}_2^{-1}}^T + \boldsymbol{\Sigma}_2 \quad (3.24)$$

The higher-order terms become relevant for high perturbation (Barfoot and Furgale 2014), but this thesis focuses on the first-order terms to show that the concept is applicable.

3.1.5 Own Implementation

In order to decide how to implement the proposed method in the framework of B-Human the current method for compounding poses has to be considered. Currently, a pose is represented as a class consisting of a translation vector and a rotation represented as a rotation matrix. The class provides multiple constructors for easy and readable construction of general poses, as well as often-used special cases. It also takes care of compounding poses by overloading operators. To extend the functionality to also take care of the uncertainty and how it is handled during compounding, a new class `SE3WithCov` has been written that inherits from `Pose` and adds the desired functionality. By extending the existing class it is ensured that the new class can be used everywhere where the old class was used before. To use the new functionality in the system only the type had to be changed from the old `Pose` class to the new `SE3WithCov` everywhere uncertainty should be handled. The new constructors need only be used where new uncertainty is introduced which is in one central file that describes the kinematic chain and therefore the uncertain joint angles. All static transformations are considered noise-free for this work.

The new class additionally contains a covariance matrix to represent the uncertainty. It only defines basic constructors benefitting from the construction options of `Pose` by defining one that takes a pose (that can be constructed in various ways) and the associated uncertainty. The operators for compounding are defined so that the associated uncertainty is propagated according to equation (3.24). Here it is important to note that the propagation of uncertainty not only needs to take place when compounded with another object of the same type but also when composed with a `Pose` as the transformation itself affects the covariance via the adjoint matrix. This case can be treated as compounding with a `SE3WithCov` where the covariance is zero.

3.2 Camera to Field

Based on the findings of Barfoot and Furgale (2014) that a sigma point method works best for the propagation through a nonlinear camera model, I propose to use a similar method for the projection from the camera image to the field plane. Because this transformation has some challenging properties described in section 3.2.1, I argue that the *Unscented transform* (UT) is more suitable than linear models. The idea of the UT is summarized in section 3.2.2 before section 3.2.3 covers the specifics of implementing it in the B-Human system. After that, the new approach is compared to the currently used method.

3.2.1 Projection Properties

The transformation $h(p_i, c)$ from a position in the image (p_i) to a position on the field/ground plane (p_f), given a camera pose (c), has some special properties that are difficult to linearize.

$$p_f = h(p_i, c) \tag{3.25}$$

First of all, it is nonlinear but, especially with respect to the camera tilt, also nonsymmetric. This results in a significantly higher absolute error for a positive distortion in tilt than for a tilt distortion of the same amount but in the opposite direction (Jüngel, Mellmann, and Spranger 2008). Due to this, a linearization would result in a bias of the transformed mean (Julier 2002). Another effect is that the mean of the transformed distribution is not the same as the mode as seen by the differences between the red (mode) and green (mean) points in figure 3.3.

As the considered uncertainty is from the camera pose and not (at least for now) from the position in the image, the dimensions of the covariance change during the transformation. The covariance matrix associated with the pose is 6×6 whereas the position on the field is two-dimensional and therefore the corresponding covariance matrix has the shape 2×2 .

3.2.2 Unscented Transform

The UT was first introduced for use in combination with a *Kalman filter* (KF) (Julier and Uhlmann 2004), which still is its most widespread use case, but it can also be used on its own to propagate an uncertain variable through a nonlinear function.

The main idea of the UT is that “it is easier to approximate a probability distribution than it is to approximate an arbitrary nonlinear function or transformation” (Uhlmann 1994 quoted from Julier and Uhlmann 2004). It uses so-called sigma points that are sampled to capture the first two moments (mean and covariance) of the distribution before the transformation. These points are then transformed through the nonlinear function, and the mean and covariance of the transformed distribution are estimated using the transformed points. By choosing a suitable set of sigma points, also higher-order moments can be taken into account to some degree (Julier 2002).

3.2.3 Own Implementation

The functionality is implemented as a data structure (representation) that provides a method to transform a point in the image to the field plane and provides the covariance of that point. As the covariance of all points in a single frame depends on the uncertainty of the same camera pose, the Sigmapoints (in this case different camera poses) can be precomputed for the entire frame. The Sigmapoints are chosen as suggested by Julier and Uhlmann (2004). The first point is the mean and the other $2N$ points are sampled symmetrically around it. The weights of the points are chosen to be equal for all points. As the calculation of the positions of the Sigmapoints requires a matrix square root, which may have multiple solutions, an algorithm has to be found that is fast and numerically stable. Julier and Uhlmann (2004) suggest using the Cholesky decomposition, which is efficient and, in the case of a covariance matrix (which is always positive semidefinite), can always be computed. But for some unknown reason, the calculation failed almost all the time. Therefore, the Eigen decomposition was used instead, which allows for easy calculation of the square root because it only needs to be computed for a diagonal matrix

where the element-wise square root can be used.

If a point in the image is to be transformed, the transformation is calculated for all precalculated camera poses. This results in a set of $2N + 1$ points on the field plane shown in figure 3.3. Based on these points the new mean and covariance are calculated.


3.2.4 Differences and Advantages to Previous Approach

In contrast to the previous approach, the new method uses the propagated covariance of the camera pose instead of a parameter as a measure of the uncertainty. This eliminates the need for a parameter that is hard to determine experimentally, as well as enabling to better capture the dynamic changes of the uncertainty of the camera pose, e.g., during different locomotion phases.


The input of the algorithm changes as well. The previous approach calculated the covariance for an already transformed point on the field. Therefore, it scaled the resulting covariance with the distance to the robot. The new method does this implicitly, as a small disturbance has a greater effect on high-pitch angles. Furthermore, the direct usage of the point in the image enables to incorporate the position relative to the image center point as well, e.g., if the point is far from the center, distortions in the camera roll have a greater effect (Jünger, Mellmann, and Spranger 2008). This can be seen in figure 3.3b. All in all, the new interface allows a more natural formulation of the problem.

As the transformation is not linearized, the method is able to capture the asymmetries, and the resulting mean does not suffer from a bias like the linearizing method. This implies also that the calculated transformed mean (green point in figure 3.3) is not the same as the transformation of the point in the image through the mean camera pose (red point in figure 3.3), as it was the case for the previous method. This has to be kept in mind when integrating it into the rest of the system, as there might be assumptions that no longer hold.

Other more indirect advantages include the independence of the UT from the specific used transformation. It could easily be adjusted, for example if a point in the image was no longer mapped to a point on the ground plane but instead to a position in 3D space. It can also be used with other distribution



(a) Transformation for points in the horizontal center of the image.



(b) Transformation for points not in the horizontal center of the image.

Figure 3.3: UT from camera images to field plane. The directly transformed point is shown in red, the other sigmapoints are blue. The new mean of the distribution is the green point.

models that the Gaussian (Julier and Uhlmann 2004) which might fit the actual distribution better.

One last advantage is the fact that it is already widely known in the robotics community due to its prominent use in the *Unscented Kalman filter* (UKF). This makes it easy for others to understand the concept and adapt it for their needs.

4 Future Work

This chapter gives ideas on how to extend the different aspects of the work presented. It indicates methods on how to evaluate the results of the work and suggestions that might improve them or provide further insight. At the end new potential options, that were made feasible by the results of the thesis, are presented.

4.1 Sensor Noise

The motion-tracking system turned out to be not accurate enough to serve as a ground truth for the rotation of a joint. By using the special setup shown in figure 2.7 the error introduced by the motion-tracking system could be limited. The new biggest error source turned out to be the radius as it is measured by hand. To improve the accuracy of this approach one could make use of the fact that the marker moves on a circular track around the joint. This could be used to improve the radius measurement.

Another possible source of error that could not be observed in the long-time static experiment is the delay between different measurements. To determine a possible delay between joint measurements and the corresponding IMU data, the robot could move its hip while the legs are fixated on the ground. The resulting measurements from the gyroscope should correspond to the differentiated hip joint measurements. By maximizing the correlation between these, a potential delay could be determined.

If there is a notable delay its implications on the uncertainty have to be determined. As the introduced error is given by the difference between the current measurement and the one after the delay, it is heavily dependent on the current rate of change. A single value can not fit the static case, where the error

introduced by the delay is zero, and the dynamic case, where the measurement changes at different rates. Instead, the error approximation should scale with the current rate of change as an approximation for the change during the delay.

4.2 Uncertainty Propagation

One focus of this thesis was the propagation of uncertainty through various transformations that have an associated uncertainty. This section describes how an evaluation could be performed and gives impulses which experiments could be performed in the future. The last part introduces another type of transformation that has not been covered in this thesis.

4.2.1 Kinematic Chain

The evaluation of the error propagation through the kinematic chain was planned to be performed using the motion-tracking system as ground truth and evaluating how often the real transformation is inside different confidence intervals provided by the implementation. As the motion-tracking system is not accurate enough, this form of evaluation could not be performed.

Instead of evaluating directly the approximated uncertainty of the kinematic chain the influence on the rest of the system could be evaluated. This could be done by placing a ball at a known location on the field and evaluating how accurately the robot knows the correct location of the ball in contrast to the old system. Besides the location also the modeled uncertainty could be evaluated. This setup would not only evaluate the influence on the ball model but also on the localization of the robot as it is needed to transform the ball position from robot relative coordinated to global coordinates relative to the field.

If a reliable evaluation method has been found, different methods for error propagation can be compared. For example, the performance of the chosen local perturbation to the global representation could be compared. Alternatively, to the representation in $SE(3)$, one could also try to represent the disturbance in $SO(3) + \mathbb{R}^3$ and evaluate to which extent the *banana-like* shape helps to improve

performance. To maintain a good approximation even for high perturbations, the effect of including higher order terms when computing the new covariance matrix of two compounded poses in contrast to the first order in equation (3.24) should be evaluated.

4.2.2 Camera to Field

A direct comparison of the previous approach for computing the covariance of a camera precept on the field plane and the new, UT-based approach is not possible due to the different inputs as described in section 3.2.4. Instead, the applicability of the UT could be verified by comparing it with the gold standard Monte Carlo method.

Besides giving each sigma point the same weight as chosen here different weightings, especially with respect for the mean, could improve the performance. A special case for different weights would be to only consider the rotational dimensions as they probably have the most impact. This would reduce the computational cost by around a factor of two.

An additional thing to try would be to compare the UT with a linearization. But as this would require linearizing the complex projection function with respect to the camera pose and the advantages of the UT are already studied (Julier and Uhlmann 2004) the reward is very likely not worth the effort, especially since one of the advantages of the UT is that it does not need an analytical linearization.

4.2.3 Others

Most objects perceived by the robot are modeled in the coordinate frame relative to the robot, but sometimes their position relative to the global field is required. This is especially the case when they are communicated between the robots of one team, but also when their relative position to fixed points on the field is needed, for example, to check whether the ball is still near a touchline, indicating that a kick-in was not yet performed. The transformation between these reference frames is also a source of uncertainty, as the robot usually does not know its exact position and orientation on the field. Therefore,

a method that propagates the uncertainty from one reference frame to the other and incorporates the uncertainty of the transformation itself is needed for the future.

4.3 New Opportunities

The new way of representing uncertainty and propagating it through the system allows to improve the filters for the robots' orientation by incorporating new measurements. The strong correlation between translation and orientation, which can be expressed due to the use of $SE(3)$, allows gaining richer information from low dimensional measurements (Meyer, Strobl, and Triebel 2022). Low dimensional measurements of the camera pose can be extracted from special features in the image as shown by (Jünger, Mellmann, and Spranger 2008), but instead of directly calculating the angle I propose to use it as an additional measurement that can be used by a KF, fusing it with data from the kinematic chain. This allows to not only improve the angle but also the estimation for the other dimensions due to the correlation expressed in the covariance matrix.

This could be used to improve the estimation of the camera pose, but also, if the data is fed back through the kinematic chain, to improve the estimation of the torso pose, allowing for better reactions in case of disturbances.

Bibliography

- Anders E.E. Wallin, Danny Price, Cantwell G. Carson, Frédéric Meynadier, Yan Xie, and Erik Benkler (Aug. 3, 2019). *AllanTools*. Version 2019.9 (cit. on p. 10).
- Arne Hasselbring (2020). “Optimierung des Verhaltens von Fußballrobotern mittels simulierter Spiele und Gegnerverhaltensmodellen”. MA thesis. Bremen: University of Bremen (cit. on p. 1).
- Barfoot, Timothy D. and Paul T. Furgale (June 2014). “Associating Uncertainty With Three-Dimensional Poses for Use in Estimation Problems”. In: *IEEE Transactions on Robotics* 30.3, pp. 679–693. DOI: [10.1109/TR0.2014.2298059](https://doi.org/10.1109/TR0.2014.2298059) (cit. on pp. 4, 23, 25, 26, 28).
- Blanco, Jose-Luis (Oct. 21, 2014). *A Tutorial on SE(3) Transformation Parameterizations and on-Manifold Optimization*. Technical report #012010. Málaga: MAPIR: Grupo de Percepción y Robótica Dpto. de Ingeniería de Sistemas y Automática, p. 56 (cit. on pp. 20, 22).
- Bourmaud, Guillaume, Rémi Mégret, Audrey Giremus, and Yannick Berthoumiou (Sept. 2013). “Discrete Extended Kalman Filter on Lie Groups”. In: *21st European Signal Processing Conference (EUSIPCO 2013)*, pp. 1–5 (cit. on p. 24).
- Brossard, Martin, Silvère Bonnabel, and Jean-Philippe Condomines (Sept. 2017). “Unscented Kalman Filtering on Lie Groups”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2485–2491. DOI: [10.1109/IROS.2017.8206066](https://doi.org/10.1109/IROS.2017.8206066) (cit. on p. 23).
- Gouaillier, David, Vincent Hugel, Pierre Blazevic, Chris Kilner, Jerome Monceaux, Pascal Lafourcade, Brice Marnier, Julien Serre, and Bruno Maisonnier (May 2009). “Mechatronic Design of NAO Humanoid”. In: *2009 IEEE International Conference on Robotics and Automation*. 2009 IEEE International Conference on Robotics and Automation, pp. 769–774. DOI: [10.1109/ROBOT.2009.5152516](https://doi.org/10.1109/ROBOT.2009.5152516) (cit. on p. 2).
- Hertzberg, Christoph, René Wagner, Udo Frese, and Lutz Schröder (Jan. 1, 2013). “Integrating Generic Sensor Fusion Algorithms with Sound State Representations through Encapsulation of Manifolds”. In: *Information Fusion* 14.1, pp. 57–77. DOI: [10.1016/j.inffus.2011.08.003](https://doi.org/10.1016/j.inffus.2011.08.003) (cit. on pp. 23, 25).

- Hunter, John D. (May 2007). “Matplotlib: A 2D Graphics Environment”. In: *Computing in Science & Engineering* 9.3, pp. 90–95. DOI: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55) (cit. on p. 8).
- Jesse Richter-Klug (Feb. 26, 2018). “Visuelle Odometrie in der RoboCup Standard Platform League”. MA thesis. Bremen: University of Bremen. 93 pp. (cit. on p. 6).
- Julier, S.J. (May 2002). “The Scaled Unscented Transformation”. In: *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*. Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301). Vol. 6, 4555–4559 vol.6. DOI: [10.1109/ACC.2002.1025369](https://doi.org/10.1109/ACC.2002.1025369) (cit. on pp. 28, 29).
- Julier, S.J. and J.K. Uhlmann (Mar. 2004). “Unscented Filtering and Nonlinear Estimation”. In: *Proceedings of the IEEE* 92.3, pp. 401–422. DOI: [10.1109/JPROC.2003.823141](https://doi.org/10.1109/JPROC.2003.823141) (cit. on pp. 29, 32, 35).
- Jünger, Matthias, Heinrich Mellmann, and Michael Spranger (2008). “Improving Vision-Based Distance Measurements Using Reference Objects”. In: *RoboCup 2007: Robot Soccer World Cup XI*. Ed. by Ubbo Visser, Fernando Ribeiro, Takeshi Ohashi, and Frank Dellaert. Vol. 5001. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 89–100. DOI: [10.1007/978-3-540-68847-1_8](https://doi.org/10.1007/978-3-540-68847-1_8) (cit. on pp. 2, 3, 28, 30, 36).
- Kitano, Hiroaki, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, and Eiichi Osawa (Feb. 8, 1997). “RoboCup: The Robot World Cup Initiative”. In: *Proceedings of the First International Conference on Autonomous Agents*. AGENTS ’97. New York, NY, USA: Association for Computing Machinery, pp. 340–347. DOI: [10.1145/267658.267738](https://doi.org/10.1145/267658.267738) (cit. on p. 1).
- Laue, Tim, Thijs Jeffry de Haas, Armin Burchardt, Colin Graf, Thomas Rofer, Alexander Hartl, and Andrik Rieskamp (2009). “Efficient and Reliable Sensor Models for Humanoid Soccer Robot Self-Localization”. In: *Zhou, C., Pagello, E., Menegatti, E., Behnke, S., Röfer, T. (Eds.) Proceedings of the Fourth Workshop on Humanoid Soccer Robots in Conjunction with the 2009 IEEE-RAS International Conference on Humanoid Robots*. Paris, France, pp. 22–29 (cit. on p. 4).
- Long, Andrew, Kevin Wolfe, Michael Mashner, and Gregory Chirikjian (July 2012). “The Banana Distribution Is Gaussian: A Localization Study with Exponential Coordinates”. In: *Proceedings of Robotics: Science and Systems*. Sydney, Australia. DOI: [10.15607/RSS.2012.VIII.034](https://doi.org/10.15607/RSS.2012.VIII.034) (cit. on p. 25).
- Meyer, Lukas, Klaus H. Strobl, and Rudolph Triebel (Oct. 2022). “The Probabilistic Robot Kinematics Model and Its Application to Sensor Fusion”. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and*

- Systems (IROS)*, pp. 3263–3270. DOI: [10.1109/IROS47612.2022.9981399](https://doi.org/10.1109/IROS47612.2022.9981399) (cit. on pp. [3](#), [4](#), [23](#), [26](#), [36](#)).
- Reichenberg, Philip (Feb. 12, 2020). “Entwicklung eines dynamischen Aufstehens unter Vermeidung inkorrektter Bewegungszustände”. Bremen: University of Bremen. 52 pp. (cit. on p. [7](#)).
- RoboCup Technical Committee (June 1, 2023). *RoboCup Standard Platform League (NAO) Rule Book* (cit. on p. [1](#)).
- Rousseeuw, Peter J. (1991). “Tutorial to Robust Statistics”. In: *Journal of Chemometrics* 5.1, pp. 1–20. DOI: [10.1002/cem.1180050103](https://doi.org/10.1002/cem.1180050103) (cit. on pp. [11](#), [14](#)).
- El-Sheimy, Naser, Haiying Hou, and Xiaoji Niu (Jan. 2008). “Analysis and Modeling of Inertial Sensors Using Allan Variance”. In: *IEEE Transactions on Instrumentation and Measurement* 57.1, pp. 140–149. DOI: [10.1109/TIM.2007.908635](https://doi.org/10.1109/TIM.2007.908635) (cit. on pp. [10](#), [11](#)).
- SoftBank Robotics (2023a). *NAO - Documentation — Aldebaran 2.8.7.4 Documentation*. URL: http://doc.aldebaran.com/2-8/home_ao.html (visited on 08/12/2023) (cit. on pp. [2](#), [5](#), [6](#)).
- SoftBank Robotics (2023b). *Overview — NAO Software 1.14.5 Documentation*. URL: <http://doc.aldebaran.com/1-14/index.html> (visited on 08/12/2023) (cit. on p. [2](#)).
- SoftBank Robotics Europe (2023). *Inertial Unit — Aldebaran 2.1.4.13 Documentation*. URL: http://doc.aldebaran.com/2-1/family/robots/inertial_robot.html#robot-inertial (visited on 07/27/2023) (cit. on p. [5](#)).
- Solà, Joan, Jeremie Deray, and Dinesh Atchuthan (Dec. 8, 2021). *A Micro Lie Theory for State Estimation in Robotics*. URL: <http://arxiv.org/abs/1812.01537> (visited on 05/05/2023). preprint (cit. on p. [22](#)).
- Thomas Röfer, Tim Laue, Finn Marvin Ewers, Enrico Göhrs, Michelle Gusev, Arne Hasselbring, Jo Lienhoop, Ayleen Lührsens, Yannik Meincken, Philip Reichenberg, Laurens Schiefelbein, Florian Scholz, Sina Schreiber, and Simon Werner (Nov. 1, 2022). *B-Human Team Report and Code Release 2022* (cit. on p. [1](#)).
- Woodman, Oliver J (Aug. 2007). *An Introduction to Inertial Navigation*. Technical Report 696, p. 37 (cit. on pp. [5](#), [6](#), [10](#)).