Investigating the Use of Visual Information in Lane Keeping Tasks

Master of Science thesis

M.J.C. Kolff

Faculty of Aerospace Engineering

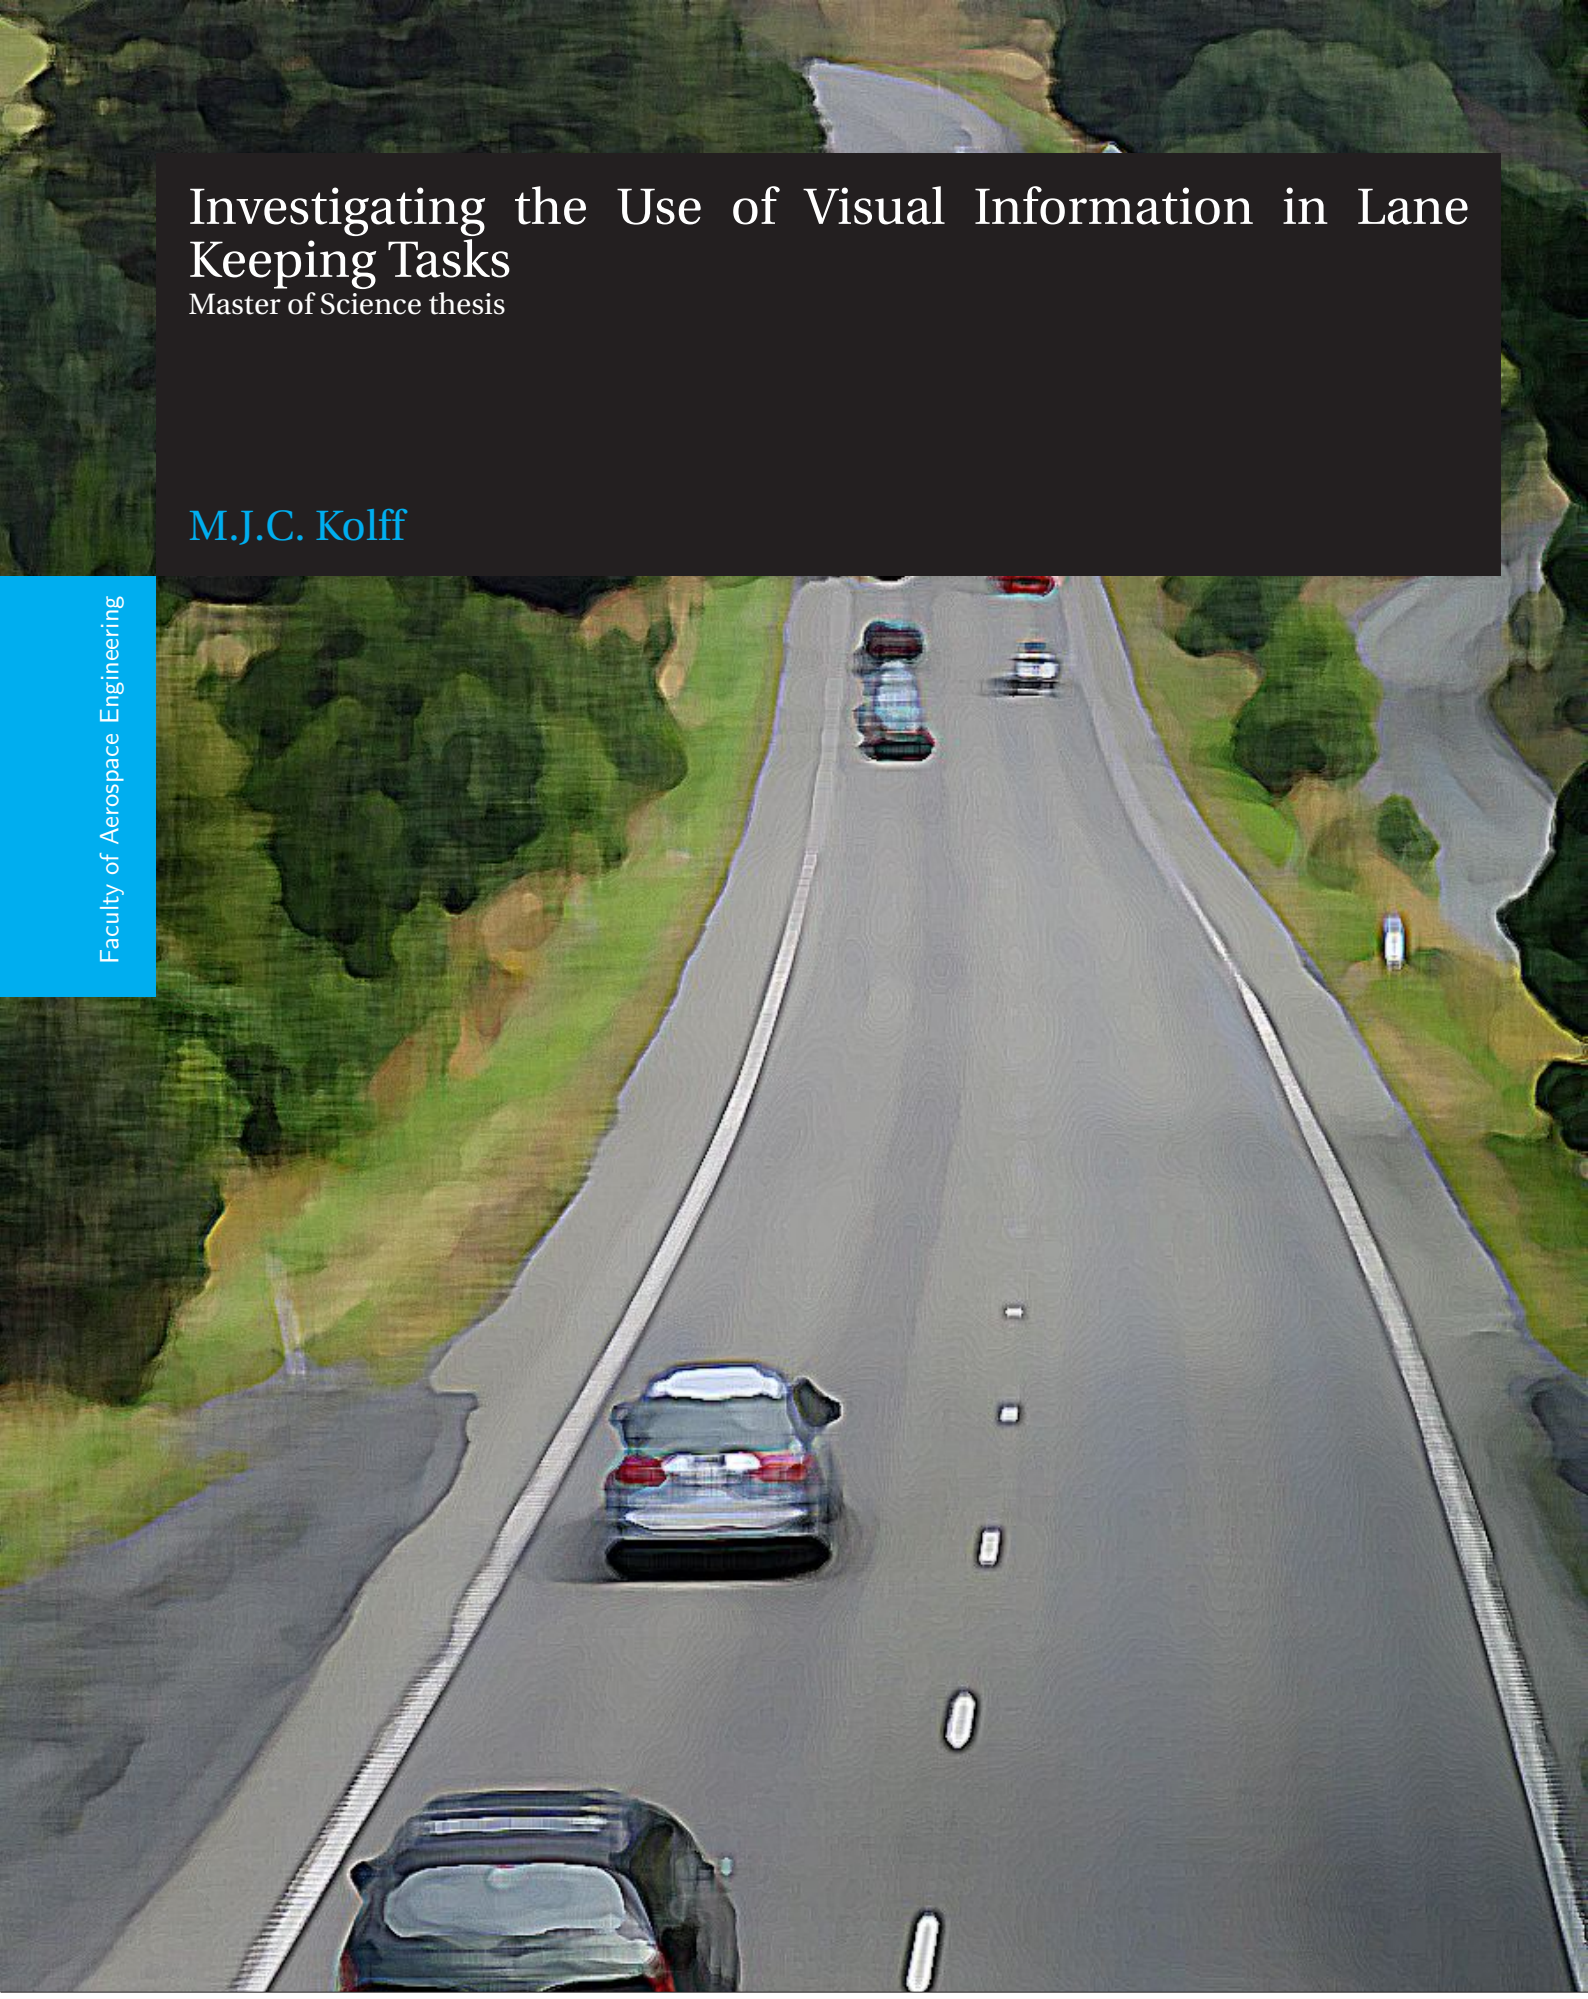

TU Delft Delft University of Technology

Challenge the future

Investigating the Use of Visual Information in Lane Keeping Tasks

A replication study of the Land & Horwood experiment

by

M.J.C. Kolff

For obtaining the degree of Master of Science
at Delft University of Technology
to be defended publicly on Monday May 27, 2019 at 14:00.

Faculty of Aerospace Engineering · Delft University of Technology

TU Delft

DELFT UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF
CONTROL AND SIMULATION

The undersigned hereby certify that they have read and recommend to the Faculty of Aerospace Engineering for acceptance a thesis entitled **"Investigating the Use of Visual Information in Lane Keeping Tasks"** by **M.J.C. Kolff**, in partial fulfilment of the requirements for the degree of **Master of Science**.

Dated: May 27, 2019

Readers:

prof.dr.ir. M. Mulder

dr.ir. D.M. Pool

dr.ir. K. van der El

dr.ir. J.C.F. de Winter

“The only true wisdom is in knowing you know nothing.”

Socrates

Preface

This thesis forms the closing piece of my time in Delft. It has been quite the learning experience, not only in terms of mathematics, simulations and other nerdy things, but also because of all the things that happened outside of the lecture rooms. I have met so many awesome and genuine people at Aerospace Engineering, Eco-runner and during my time in Munich that I value dearly.

Mainly looking back at the past nine months, a thank you is in order to those who contributed to this thesis. First of all, I thank prof.dr.ir. Max Mulder and dr.ir. Rene van Paassen for making time in their busy agendas and for giving me valuable meetings full of feedback on my plans. I feel truly honoured to have been a small part of the research performed at Control & Simulation.

I furthermore thank my weekly supervisors dr.ir. Kasper van der El and dr.ir. Daan Pool. Their support and approachability have helped me from day one. Even during the most stressful periods they helped me in staying motivated. Also a thank you to Johan Pel of the Erasmus MC for lending their Tobii glasses system. And of course thanks to everyone who participated in the experiment (Those who shall not be named...).

Finally, I thank my parents Nico and Irene, sister Anneloes and brother-in-law Zeger, for without their continuous love, motivation and support in my adventures, I would be nowhere.

Maurice Kolff,
May 27, 2019

Contents

I Paper I 1

II Paper II 9

III Literature Review 29

IV Appendices 79

A Briefing 81

B Informed Consent Form 85

C Conversion of Eye Gaze Data 87

C.1 Hardware Interface . . . 87
C.2 Data extraction. . . . 87
C.3 Conversion of Timestamps . . . 88
C.4 Reference Frame Data Conversion . . . 89
C.4.1 Detection of slit geometry . . . 89
C.4.2 Detection of baseline geometry . . . 89

D Eye gaze measurements 91

E Representative Model Fits - Subject 1 137

F Representative Model Fits - Subject 3 165

I

Paper I

Approximating Road Geometry with Multisine Signals for Driver Identification

Maurice J. C. Kolff, Kasper van der El, Daan M. Pool, [1]
Marinus M. van Paassen, and Max Mulder

Control and Simulation, Faculty of Aerospace Engineering, TU Delft,
2629 HS, Delft, The Netherlands

Abstract: The understanding of human responses to visual information in car driving tasks requires the use of system identification tools that put constraints on design of data collection experiments. Most importantly, multisine perturbation signals are required, including a multisine road geometry, to separately identify the different driver steering responses in the frequency domain. It is as of yet unclear, however, to what extent drivers steer differently along such multisine roads than they do for real roads. This paper presents a method for approximating real-world road geometries with multisine signals, and applies it to a stretch of road used in an earlier investigation into driver steering. In addition, a human-in-the-loop experiment is performed to collect driver steering data for both the realistic real-world road and its multisine approximation. Overall, the analysis of driver performance metrics and driver identification data shows that drivers adopt equivalent control behaviour when steering along both roads. Hence, the use of such multisine approximations allows for the realization of realistic roads and driver behaviour in car driving experiments, in addition to supporting the application of quantitative driver identification techniques for data analysis.

Keywords: manual control, driving, multisine signals, system identification, driver modeling

1. INTRODUCTION

Much of our understanding of driver visuomotor steering behavior in lane keeping tasks is based on models (Donges, 1978; McRuer et al., 1977; Steen et al., 2011; Van der El et al., 2019). The parameters of driver models can be estimated directly from experimental steering data, using *system identification* techniques. Such model parameters explicitly quantify control processes that are *internal* to the driver, such as their visual response delay, look-ahead time, and feedback and feedforward control gains. Unfortunately, this approach by definition requires that a model is already available, and the insight gained from the estimated parameters depends profoundly on the model's ability to capture the driver's steering dynamics.

For direct measurement of drivers' steering dynamics, without assuming a model *a priori*, black-box system identification techniques can be applied. For example, Van der El et al. (2019) obtained Frequency Response Function (FRF) estimates of drivers' separate responses to road preview, vehicle lateral position and vehicle heading, and proposed a novel, physically interpretable, driver steering model in accordance with each observed steering response.

Unfortunately, the application of such frequency-domain identification techniques requires a driving experiment where the road centerline trajectory is defined according to a *multisine* signal. It is currently unclear, however, to what extent the multisine road signals used by Van der El et al. (2019) match real-world roads and evoke realistic driver steering behavior.

[1] Corresponding author: d.m.pool@tudelft.nl

This paper investigates to what extent driver steering behavior on multisine roads is identical to steering along a winding real-world road. To do so, first, a method is introduced for approximating a real-world road trajectory as accurately as possible by a multisine signal. This method is then applied to obtain a multisine approximation of the road used in the seminal driving experiment of Land and Horwood (1995). Secondly, a human-in-the-loop simulator experiment is performed in which drivers follow both the real-world road and its multisine approximation. Measured steering behaviour on both roads is directly compared, both in the time-domain, using performance and control activity measures, as well as with driver identification results obtained with the multiloop system identification and modeling techniques from Van der El et al. (2019).

2. MULTISINE SIGNALS

2.1 Driver Identification Problem

This paper presents the work performed under the framework of a larger project in which we aim to increase understanding of driver steering behavior through a *quantitative* measuring and modeling approach (Van der El et al., 2017). For this work, we investigate the dynamics of drivers' steering behavior in continuous steering tasks, as shown in a block-diagram representation in Fig. 1.

Fig. 1 shows the vehicle dynamics, separated into the steering wheel-to-heading dynamics G_{δ}^{ψ} and heading-to-lateral-position dynamics G_{ψ}^{y}. Furthermore, the three control responses drivers use when steering through curves

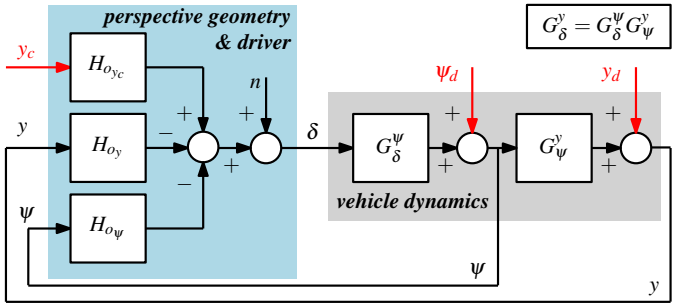

Fig. 1. Illustration of driver steering and multiloop driver dynamics, reproduced from Van der El et al. (2019).

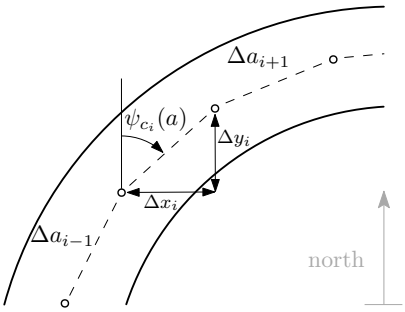

Fig. 2. Road geometry and the commanded heading angle ψ_c as function of along-track distance a.

are indicated: the feedforward (preview) response to the road ahead ($H_{o_{y_c}}$) and feedback responses to the vehicle's heading (H_{o_ψ}) and lateral position (H_{o_y}).

As explained in Van der El et al. (2019), when using instrumental variable identification techniques to estimate all *three* responses shown in Fig. 1, *three* independent multisine signals, indicated in red, are needed. These are the road geometry y_c and disturbances (i.e., wind-gusts) that perturb the vehicle's heading (ψ_d) and lateral position (y_d). This paper focuses on the first of these signals, by presenting and testing a stepwise methodology for approximating real road geometries with multisine signals.

2.2 Multisine Signal Development

The road reference signal, y_c in Fig. 1, is in practice often defined in terms of the corresponding heading angle reference of the road centerline, ψ_c (Weir and McRuer, 1970), shown in Fig. 2. Here, the road heading is a function of the along-tack distance a, similar to the approach adopted in earlier experiments (Sharp et al., 2000; Lakerveld et al., 2016; Van der El et al., 2018):

$$\psi_c(a) = \sum_{k=1}^{N_f} A[k] \sin\left(\omega[k]a + \phi[k]\right) \tag{1}$$

The signal is characterized by amplitudes $A[k]$, frequencies $\omega[k]$, and phases $\phi[k]$ of the k-th sinusoid, with N_f the number of components. We aim to find the amplitudes, frequencies, and phases such that the multsine signal approximates an original, real-world road geometry best.

In the following, we explain our proposed a four-step procedure for approximating real road geometries with a multisine signal as given by Eq. (1). Our proposed procedure is implemented for deriving a multisine signal that matches the road used in the landmark driver steering experiment of Land and Horwood (1995), which used a portion of Queens Drive (QD) in Edinburgh (detailed in personal correspondence).

Step 1: Road Coordinate Extraction To obtain the geometry of a certain stretch of road, a set of (lat/long) coordinates can be extracted from most publicly available navigation databases (e.g., GoogleMaps). First, this raw road coordinate data are converted to measures of distance. The difference in latitude ϕ_i and longitude λ_i (in deg) of two road coordinate points can be converted to relative position coordinates in meters according to:

$$\Delta x_i = \Delta\lambda_i \frac{\pi}{180} R_e \cos\Delta\phi_i \tag{2}$$

$$\Delta y_i = \Delta\phi_i \frac{\pi}{180} R_e \tag{3}$$

Here, R_e indicates the Earth's radius. These equations are valid for a spherical earth approximation, which is acceptable for "small" differences in location, i.e., $[\Delta x, \Delta y] \ll R_e$. The series of $(\Delta x_i, \Delta y_i)$ points fully define the road geometry, but can be combined through the commanded heading angle:

$$\psi_{c_i} = \arctan\left(\frac{\Delta x_i}{\Delta y_i}\right) \tag{4}$$

The corresponding along-track-distance a is defined as the cumulative sum of the length of all Δa intervals up to the i^{th} road point:

$$a = \sum_{i=1}^{N} \Delta a_i = \sum_{i=1}^{N} \sqrt{\Delta x_i^2 + \Delta y_i^2} \tag{5}$$

These Eqs. (4) and (5) together yield the signal $\psi_c(a)$, which fully defines the road centerline trajectory. The road *curvature*, another common metric for specifying road geometry, is given by $r_c = \frac{d\psi_c}{da}$.

For our application of the above methodology to the segment of Queens Drive in Edinburgh, acquired road coordinate-data from GoogleMaps consists of a total of $N = 559$ points. With an assumed constant Earth radius of $R_e = 6371$ km, the latitude and longitude coordinates were transformed to the Cartesian road geometry shown in Fig. 3 with the blue markers.

Step 2: Data Processing Exported road coordinates are generally not equally spaced along the length of the road, as is clearly visible for our raw road coordinate data in Fig. 3. An evenly-spaced data point distribution is required for applying the Fourier transform to obtain the road frequency spectrum, and approximating this spectrum with a multisine signal. Therefore, the obtained data are interpolated to constant along-tack distance ($\Delta a = 0.1389$ m, corresponding to (Van der El et al., 2019)) using shape-preserving, piecewise cubic interpolation.

Finally, our experience showed that exported raw road coordinate data contain inaccuracies and artifacts (i.e., noise), which yield unrealistically sharp changes in heading. Therefore, we smoothed the x-, y-coordinate data with a low-pass filter with a 0.007π rad/m normalized cut-off and a 0.05π rad/m stopband frequency (100 dB stopband attenuation). The interpolated and smoothed set

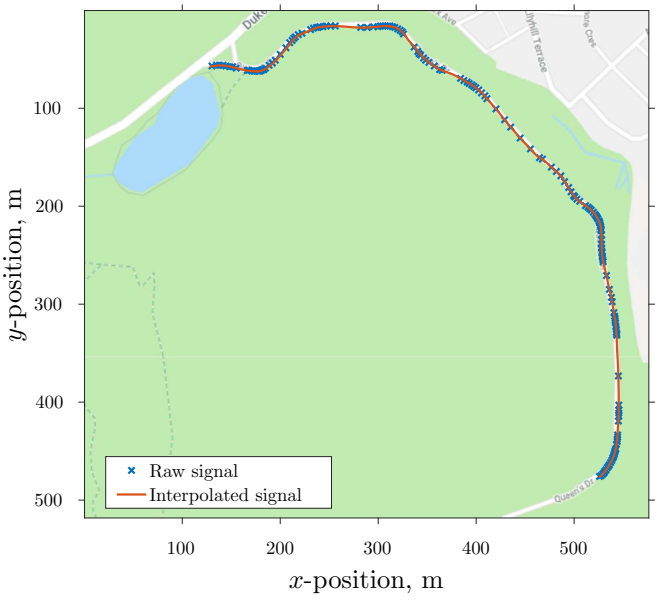

Fig. 3. Top-down map view of the Queens Drive (QD) road segment used by Land and Horwood (1995).

of road points is shown in Fig. 3 in red, and matches the underlying road trajectory well.

Step 3: Road Straightening As is clear from Fig. 2, the considered section of Queens Drive is part of a circular track, meaning that the heading angle has a linear trend (nonzero average curvature). Fig. 4 shows, in blue, this heading ψ_c calculated from the $(\Delta x_i, \Delta y_i)$ coordinates shown in Fig. 3. To follow this road, a nonzero average steering wheel deflection is required. However, a multisine signal lacks such a linear trend, which must thus be removed from the real-world road to capture the *relative* changes in heading (and the required steering) with a multisine signal. To achieve this, the overall road is "straightened" by subtracting the mean curvature of the road from the actual road curvature at each point, yielding the heading angles in Fig. 4 (red line).

Step 4: Multisine Frequency, Amplitude, and Phase Selection Finally, now, the Fourier transform of the heading can be obtained, yielding its frequency spectrum in rad/m, given in blue in Fig. 5. The spectrum reveals that

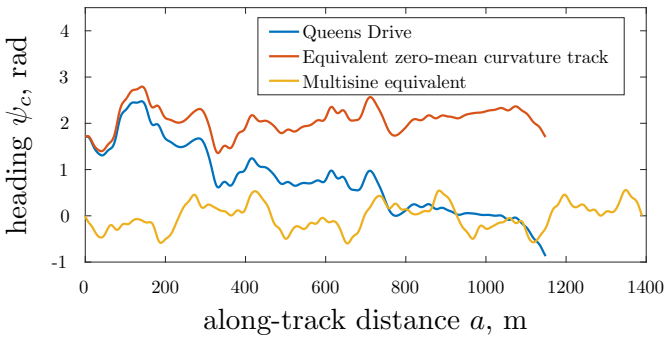

Fig. 4. Heading angle of the real-world road, the same road corrected for the curvature of the circular track and its multisine equivalent.

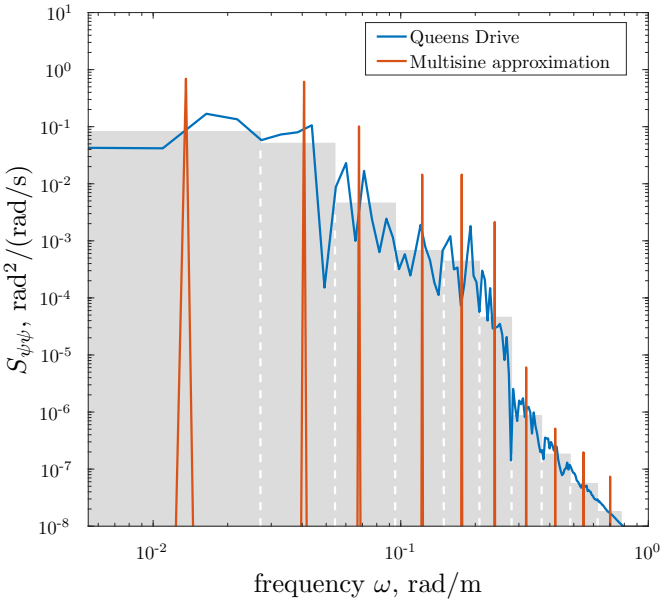

Fig. 5. Road heading signal spectra for the Queens Drive road, and its multisine signal approximation.

the Queens Drive road contains power predominantly at frequencies below 0.25 rad/m. The multisine road signal used in the driving experiments by Van der El et al. (2018, 2019) was constructed of $N_f = 10$ sine-components, of frequencies up to 0.7 rad/m, to allow measuring the driver's steering dynamics also at higher frequencies. It was decided to approximate the Queens Drive road by a multisine with these same frequencies $\omega[k]$, indicated by the red peaks in Fig. 5. Note that other choices for the frequencies are possible, see (Damveld et al., 2010) for further guidance.

Next, the multisine amplitudes $A[k]$ are determined, by computing the total power of the Queens Drive road in each frequency "bin", with boundaries in the middle between the selected multisine frequencies $\omega[k]$. The bins and total power are indicated by the gray areas in Fig. 5. Each multisine amplitude is then selected such that its power equals the total power in that frequency bin. The resulting multisine signal spectrum is shown in Fig. 5 in red. The phases $\phi[k]$ were subsequently obtained using a brute force method, in which 10.000 random phase vectors were generated with $-\pi < \phi[k] < \pi$. The phase vector for which heading error variance between the multisine and real-world road signals was smallest was selected as solution. The resulting ψ_c signal is shown in yellow in Fig. 4, and its characteristics are listed in Table 1.

3. METHODS

To verify whether the multisine (MS) approximation of the original Queens Drive road geometry evokes identical driver steering behavior, a human-in-the-loop experiment was performed in which drivers followed both road trajectories. This experiment was part of a larger study, in which also a replication of the visual occlusion experiment of Land and Horwood (1995) was performed. Here we focus on the verification of the multisine road geometry.

Table 1. Frequencies, amplitudes, and phases of the multisine approximation of the Queens Drive road in Edinburgh. The disturbance signals, reproduced from (Van der El et al., 2018, 2019), are given for completeness.

	Road center line ψ_c				Heading angle disturbance ψ_d				Lateral position disturbance y_d			
i	k	ω	A	ϕ	k	ω	A	ϕ	k	ω	A	ϕ
-	-	rad/m	deg	rad	-	rad/m	deg	rad	-	rad/m	m	rad
1	3	0.01	16.2	3.96	7	0.03	2.20	5.04	5	0.02	0.29	5.98
2	9	0.04	15.3	3.17	13	0.06	1.74	6.22	11	0.05	0.24	4.04
3	15	0.07	6.20	4.78	23	0.10	1.08	4.17	19	0.09	0.16	3.03
4	27	0.12	2.34	3.40	35	0.16	0.63	4.40	31	0.14	0.09	6.11
5	39	0.18	2.34	6.28	47	0.21	0.41	4.97	43	0.19	0.06	0.99
6	53	0.24	0.90	6.20	65	0.29	0.25	4.97	59	0.27	0.04	0.11
7	71	0.32	0.05	5.42	85	0.38	0.16	4.10	77	0.35	0.02	1.78
8	93	0.42	0.01	5.71	111	0.50	0.11	5.90	101	0.46	0.02	2.28
9	121	0.55	0.01	0.95	143	0.65	0.08	5.48	131	0.59	0.01	0.41
10	155	0.70	0.01	4.79	183	0.83	0.07	0.73	169	0.76	0.01	2.41

3.1 Experiment Setup

The experiment was performed in TU Delft's SIMONA Research Simulator (SRS), see Fig. 6. The motion system of the SRS was not used and participants performed a visual-only driving task, see Fig. 7. Visuals were presented using the collimated 60×40 deg field-of-view outside visual system (single projector), see Fig. 7; the two side projectors of the SRS were switched off to match the experiment of Land and Horwood (1995) as well as possible. Full details of the experimental setup are given in Van der El et al. (2018, 2019).

The task was performed with a constant forward vehicle velocity $U_0 = 16.9$ m/s, identical to (Land and Horwood, 1995), and the vehicle heading dynamics (G_δ^ψ in Fig. 1) were an integrator, identical to (Land and Horwood, 1995; Van der El et al., 2019). Vehicle lateral position is obtained by integrating the heading angle. The road was 3 m wide.

For the QD condition the measurement portion of each run covered 1149 m along-track-distance, see Fig. 3. This corresponds to a nominal measurement time of 68 s (when the road centerline is followed perfectly), identical to (Land and Horwood, 1995). For the MS condition, the measurement portion of the road was extended to 1389 m (82.2 s nominal measurement time), identical to (Van der El et al., 2018, 2019), to guarantee that each sine component fits exactly an integer number of times in the total measurement, for Fourier analysis. The actual trajectory driven by participants in each run was extended by a run-in (277 m) and run-out (138 m) portion, of which the data were not analyzed.

3.2 Experiment Participants and Procedure

Nine male and three female volunteers ($\mu = 25.4$ yr, $\sigma = 3.3$ yr) participated in the experiment. All participants provided informed consent prior to their participation and were in the possession of a valid driver's license ($\mu = 7.2$ yr, $\sigma = 3.0$ yr), with varying travelled distances per year ($\mu = 4420$ km, $\sigma = 7447$ km). Participants were instructed to drive as they would normally do.

They first completed five runs of the QD condition, in which the heading and lateral disturbance signals (ψ_d and y_d) were zero, identical as in (Land and Horwood, 1995). The first two runs were used to familiarize participants with the setup, and only the remaining, final three runs were used for analysis. Subsequently, four runs were performed for the multisine (MS) condition, which *did* include the additional disturbances to facilitate system identification. Here, only a single run was used for practice, and the final three runs were analyzed.

3.3 Dependent Measures

Performance and Control Activity Driver steering on the QD and MS roads is compared by analysis of the measured steering wheel rotations δ (see Fig. 1) and the car's lateral position deviation from the centerline $y_e = y_c - y$. First, measured time traces of δ, as well as the corresponding power spectra, are compared. In addition, the standard deviations of the control output (σ_δ) and

Fig. 6. The SIMONA Research Simulator (SRS).

Fig. 7. Picture of a participant performing the experiment, illustrating the experiment setup.

the lateral position deviations (σ_{y_e}) are used as metrics to compare control activity and road-following performance for both roads. The total variance has four components, e.g., for the control output:

$$\sigma_\delta^2 = \sigma_\delta^2(y_c) + \sigma_\delta^2(y_d) + \sigma_\delta^2(\psi_d) + \sigma_\delta^2(n), \qquad (6)$$

which are the variance components at the frequencies of the target centerline (y_c), lateral position disturbance (y_d), heading disturbance (ψ_d), and remnant (n). The separate contributions to the total variance of each component can be computed in the frequency domain, see (Jex et al., 1978) for details. For fair comparison with the disturbance-free QD condition, disturbance-free estimates of σ_δ and σ_{y_e} in the MS condition are obtained by subtracting the two disturbance components (y_d and ψ_d) from the total control output and lateral deviation variances using Eq. (6).

Driver Identification Driver steering in the MS condition is further analyzed using multiloop system identification techniques. The applied methods are identical to those in (Van der El et al., 2019). The three driver responses to road preview ($H_{o_{y_c}}$), vehicle heading (H_{o_ψ}), and vehicle lateral position (H_{o_y}) are estimated, with the driver steering output (in the frequency domain) defined as the linear combination of the three responses (see also Fig. 1):

$$\delta = H_{o_{y_c}} y_c - H_{o_y} y - H_{o_\psi} \psi \qquad (7)$$

First, Frequency-Response Functions (FRFs) are estimated for the three driver responses, and, second, the multiloop driver perception and steering control model from (Van der El et al., 2019) is fit to the FRF data. Only the first two measurement runs are used for estimating the driver FRFs and models, while the third measurement run is used for validation and for computing the Variance Accounted For (VAF) of the fitted models.

4. RESULTS

4.1 Performance and Control Activity

Fig. 8(a) shows that the measured control outputs for the Queens Drive (QD) road (blue data) and the multisine (MS) task (red data) are clearly different. This is confirmed by the corresponding power spectra in Fig. 8(b). For the multisine road the power spectrum shows distinct peaks at the road and disturbance input frequencies, as expected. In fact, the shape of the power spectrum at the multisine frequencies of the road centerline (black data in Fig. 8(b)) *is identical* to that of the measured control output spectrum in the QD task (blue data). The power at each multisine frequency in the MS task is higher than the power in the QD task, where the power is more spread out across frequencies. For *road following*, it can thus be concluded that the participant shows identical steering behavior on the real road (QD) as compared to the multisine approximation of this road (MS).

This is confirmed in Fig. 9. Both the *total* lateral position deviations and the control outputs (i.e., including the contributions of ψ_d and y_d) are higher on the multisine road, as compared to the Queens Drive road. However, when only considering the centerline and remnant frequencies, so excluding the disturbance components in Eq. (6), the control output and lateral deviation magnitudes are *equivalent* for both roads, for all participants.

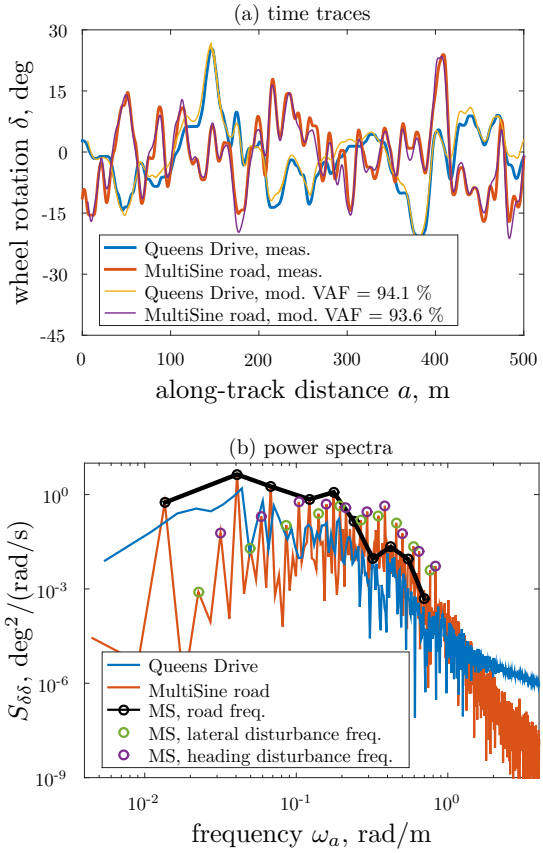

Fig. 8. Measured control output time traces and power spectra (Participant 1, Run 3).

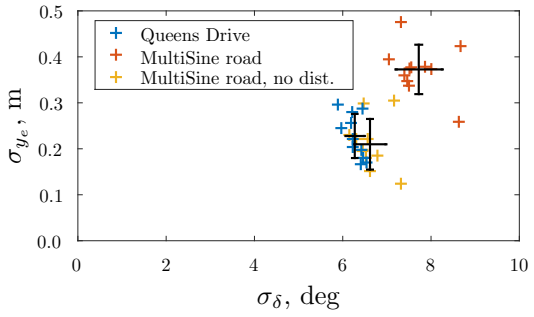

Fig. 9. Measured standard deviations of steering wheel rotations (σ_δ) and lateral deviations from the road centerline (σ_{y_e}), individual participant data and sample average with standard deviations. The multisine road data without disturbances were computed in the frequency domain (Jex et al., 1978).

4.2 Driver Identification

Fig. 10 shows the estimated multiloop driver response dynamics obtained from the MS condition data. Note again that these estimated driver responses can only be obtained because of the three uncorrelated multisine perturbations applied in this task. The estimated FRFs, given by the markers in Fig. 10, show the control dynamics of the driver in response to the road preview, vehicle lateral position, and vehicle heading, respectively. The fitted driver model of Van der El et al. (2019) evidently captures the driver dynamics very well, except the feedforward

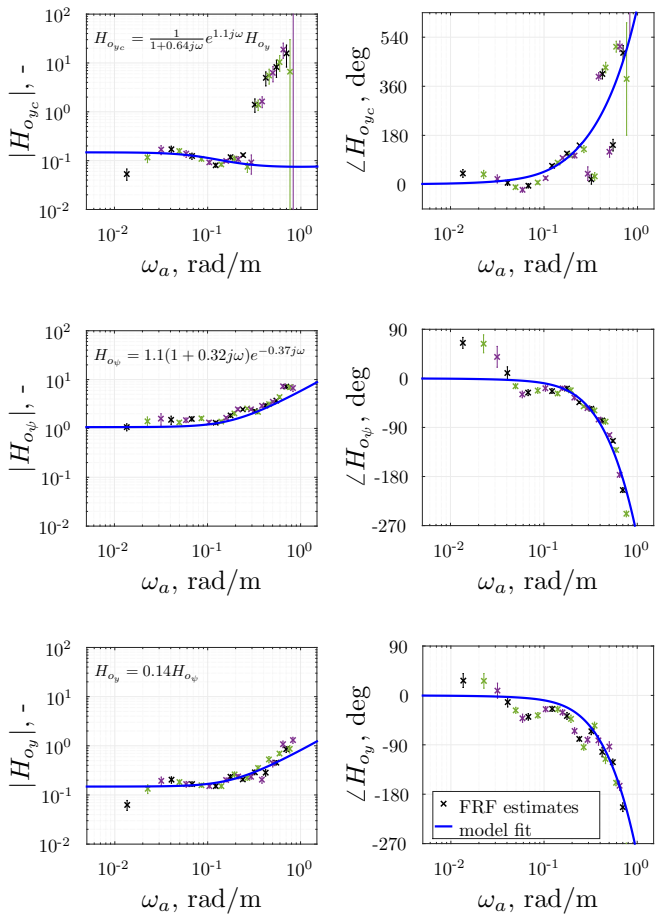

Fig. 10. Example Bode plots of estimated driver dynamics in the MS condition, based on the first two measurement runs of Participant 1.

$H_{o_{y_c}}$ at the highest frequencies, where the control output is low. The FRF estimates thereby validate that the driver model indeed describes the actual driver's multiloop control dynamics.

Using the estimated model to simulate the control output in the third measurement run (validation data) using Eq. (7) yields the "model" time traces in shown in Fig. 8(a) (yellow trace for QD, purple trace for MS). For the multisine road condition, the modeled control outputs are nearly identical to the measured control outputs (red data), with a VAF of 94.1%. Strikingly, using the same estimated model to simulate the control output in the third run of the Queens Drive condition provides an equally good match (93.6% VAF) to the measured control output. Comparable, but slightly noisier results are obtained for the other participants, with average VAFs of 86.5% and 89% in the QD and MS conditions, respectively. These results further indicate that participants adopted equivalent steering behavior in the MS condition, as compared to driving on the corresponding real-world road.

5. CONCLUSION

This paper presented a method to analyze the spectral components of real roads and construct equivalent multisine signals with matched power over selected ranges of frequencies. This method was applied to approximate a portion of Queens Drive in Edinburgh, used in the experiments of Land and Horwood (1995), with a multisine. From collected human-in-the-loop experiment data it was found that driver steering behaviour, in terms of driving performance and driver control dynamics, is equivalent between the real Queens Drive and its multisine approximation. This confirms that a multisine approximation of a real road evokes similar driver behaviour, which enables the application of frequency-domain identification techniques to analyze driver steering behaviour in future curve driving experiments.

REFERENCES

Damveld, H.J., Beerens, G.C., Van Paassen, M.M., and Mulder, M. (2010). Design of forcing functions for the identification of human control behavior. *Journal of Guidance, Control, and Dynamics*, 33(4), 1064–1081.

Donges, E. (1978). A Two-Level Model of Driver Steering Behavior. *Human Factors*, 20(6), 691–707.

Jex, H.R., Magdaleno, R.E., and Junker, A.M. (1978). Roll Tracking Effects of G-vector Tilt and Various Types of Motion Washout. In *Proc. 14th Ann. Conf. on Manual Control*, 463–502.

Lakerveld, P.R., Damveld, H.J., Pool, D.M., Van der El, K., Van Paassen, M.M., and Mulder, M. (2016). The Effects of Yaw and Sway Motion Cues in Curve Driving Simulation. In *Proc. of the 13th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Human-Machine Systems, Kyoto, Japan.*

Land, M.F. and Horwood, J. (1995). Which Parts of the Road Guide Steering? *Nature*, 377, 339 – 340.

McRuer, D.T., Allen, R.W., Weir, D.H., and Klein, R.H. (1977). New Results in Driver Steering Control Models. *Human Factors*, 19(4), 381–397.

Sharp, R.S., Casanova, D., and Symonds, P. (2000). A Mathematical Model for Driver Steering Control, with Design, Tuning and Performance Results. *Vehicle System Dynamics*, 33(5), 289–326.

Steen, J., Damveld, H.J., Happee, R., Van Paassen, M.M., and Mulder, M. (2011). A Review of Visual Driver Models for System Identification Purposes. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2011*, 2093 – 2100.

Van der El, K., Pool, D.M., Van Paassen, M.M., and Mulder, M. (2019). A Unifying Theory of Driver Perception and Steering Control on Straight and Winding Roads. *IEEE Trans. on Human-Machine Systems.* Submitted.

Van der El, K., Pool, D.M., and Mulder, M. (2017). Measuring and Modeling Driver Steering Behavior: From Compensatory Tracking to Curve Driving. *Transportation Research Part F.* Online preprint.

Van der El, K., Pool, D.M., Van Paassen, M.M., and Mulder, M. (2018). Identification and Modeling of Driver Multiloop Feedback and Preview Steering Control. In *Proc. of the IEEE International Conference on Systems, Man, and Cybernetics, Myazaki, Japan*, 1227–1232.

Weir, D.H. and McRuer, D.T. (1970). Dynamics of Driver Vehicle Steering Control. *Automatica*, 6(1), 87–98.

II

Paper II

Effects of Visual Occlusion in Lane Keeping Tasks on Driver Model Identification and Gaze Behavior

M.J.C. Kolff
Delft University of Technology
m.j.c.kolff@student.tudelft.nl
Supervisors: K. van der El, D.M. Pool, M. Mulder

***Abstract*—In car driving, manual control to keep a vehicle within its lane is dominated by the use of visual information. Linear models describing behaviour in such tasks can therefore be directly based on the human perception of the visual scene, although it is currently unclear how this perception guides control behaviour. In literature, occlusion experiments have investigated this connection by restricting the field of view of the visual scene. Such experiments were long considered an indication for two-point models, but only looked at the deviation from the road current path as a measure of performance. However, drivers are not necessarily inclined to aim for a high performance by means of perfect center line tracking, and similar performance measures can also be the result of different control strategies. In this paper an occlusion experiment is described in which a recently developed parametric model is fitted, including the estimation of the driver Frequency Response Functions (FRFs), describing the individual responses to road preview, lateral position and heading angle information. Complementary, the eye gaze is measured and compared to the interpretable model parameters. For the first time, insight is given in the behavioral changes under various occlusion conditions with respect to baseline behavior directly in relation where drivers look. It is shown that drivers adapt their modelled aim points and eye gaze to the available road geometry if only a single slit is present. For double-slit conditions, drivers place both the gaze and aim points between the occlusion slits, effectively interpolating the available visual information while still responding to a single metric, although preprocessing behavior is reduced compared to full visual driving.**

***Index Terms*—manual control, lane keeping, visual occlusion, system identification, eye gaze.**

I. Introduction

Car driving is a predominantly visual task [1], [2]. Understanding the way in which humans extract information from the visual scene [3]–[5], by responding to its cue dynamics, is required to form a comprehensive description of driver behavior [6], [7]. For that reason, it has been a long-standing question as to which areas are adequate of providing drivers enough visual information for their guidance, which visual regions are actually used in car control, and how this relates to the attention span of drivers in continuous [8]–[10] or intermittent [11], [12] tasks. Understanding this level of driver behavior can for example provide car driving automation that responds in a human-like manner [13], [14].

In literature, several research studies have aimed at relating these questions of visual perception and control theory by performing *occlusion* studies, in which changes in driver performance are measured as a function of a decrease in field of view. This can be done by increasing the transparency of far- and near visuals [15] or by completely blocking large parts of the visuals by an overlay, meaning that drivers can only perceive the road through distinct slits, such as done by Land & Horwood [16]. In this study, by varying the vertical position of the slits, it was found that to achieve similar performance as compared to a baseline (full visuals) measure, at least two ”distinctly” separated areas are required. The need of these visual areas were considered empirical proof for two-point driver control models, such as [17], in which driver behavior is explained by responses to two separate points: a near-view lateral position feedback response and a far-view feedforward response to the road curvature.

Similar occlusion experiments ([18], [19]) and simulations [20] not only failed to reproduce these results, but also measured a dependency on the experiment setup (e.g. the presence of optical flow, which affects guidance [21], and the display type). Furthermore, the presence of texture affects the control activity of drivers [22], whereas Land & Horwood used simplistic line visuals. Therefore it is currently unclear as to why driver performance changes for most occlusion conditions. More importantly, all of these occlusion experiments have only compared relative performance measures between occluded and full-visual driving tasks. However, similar performances do not automatically result from similar control behavior; drivers can apply different driving strategies to obtain the same performance. Although lumped driver responses have been measured before [23]–[25], these provide little insight in how humans process the different sources of information (such as lateral position and heading angle perceived from the scene).

As a recent development, Van der El et al. [26] introduced a novel control-theoretic driver model, based on a three-level structure of a road target signal feedforward, lateral position feedback and heading angle feedback response. In this structure, drivers respond to a single perceived error, which combines information on these three metrics. This is in contrast to the two-point models such as [17], as well as the most important result of [16]: that at least two distinctly separated areas are required for successful control. The development of the control model by [26] also provided fundamental insight in the behavioral dependency on visual cues (such as the dependency on lateral position or heading angle information). This would allow the verification that specific occlusion conditions possibly not only yield a similar performance compared to full visuals, but also require a

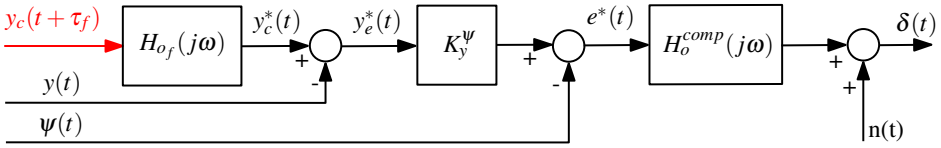

Fig. 1: Control-theoretic lane keeping model, indicating the target preview, lateral position feedback and heading angle feedback to describe human control to a single perceived error $e^*(t)$, adapted from [26].

similar driving strategy, indicating that only specific regions are required for describing lane keeping control.

For the disentanglement of such individual driver responses and determining their Frequency Response Functions (FRF), multisine signals are required [27]. In [28], a method was presented for the conversion of realistic road geometries to multisine signals for the road centerline geometry and shown that this conversion does not affect the prediction methods of the model of [26], such that drivers instigate similar driving behavior. This allows for the use of system identification and parameter estimation methods *within* a single occlusion experiment, gaining more insight in how the geometry of the road is actually used by humans and how this relates to the full visual control, and directly relating the eye gaze to interpretable model parameters. A similar study [29] showed that for preview tasks, drivers also aim their gaze close to a presented near-view point rather than the far-view point, but this has not yet been extended to full lane keeping tasks. Driver gaze behavior has been measured extensively in car driving tasks [9], [11], [12], [30], but not in direct relation to control-theoretic metrics. Given that the model of [26] provides interpretable parameters directly related to the road geometry, measuring the gaze simultaneously could therefore further increase the link between visualperception and control theory.

The goal of this research study is therefore to form a better understanding of lane keeping tasks using three main steps. First, by measuring the effects of visual occlusion, the experiment originally performed by [16] is replicated and both studies are directly compared. Second, individual driver FRF responses of these occlusion conditions are compared to their full-visual counterparts. This is extended by fitting the parametric model of [26] to see how drivers adapt their driving strategy and how this relates to the geometry of the occlusion. Finally the driver viewing gaze and the interpretable aim points of the parametric model are directly compared.

This paper is structured as follows. An introduction of driver modelling is given in Section II, of which the required system identification methods are explained in Section III. Section IV describes the methods used for the experiment. The results are shown in Section V, of which the discussion is given in Section VI. The paper ends with the conclusions.

II. Driver Model Geometry

A. Use of Visual Metrics

This paper is part of a larger research project in which the aim is to increase the understanding of driver steering behavior by describing the use of visual metrics. Previous work has investigated a step-wise build up [31] of the traditional pursuit- and preview control tasks to a representative model of driver steering behavior in continuous lane keeping tasks.

During lane keeping, the driver control task is to keep the vehicle on the road by minimizing the lateral position deviation $y_e(t)$ from the center-line at $y_c(t)$ (i.e., $y_e(t) = y_c(t) - y(t)$). This is done by rotating the steering wheel by $\delta(t)$, such that the controlled element (the vehicle dynamics) output $y(t)$ remains on the road.

The preview information of the road, perceived through the windscreen, is used to anticipate for changes of the geometry (i.e., bends). The current vehicle position with respect to current and future commanded road trajectories guides the control of the vehicle. However, this position, defined by the lateral position deviation $y_e(t)$, cannot be observed directly, as it is positioned below the vehicle, such that a form of preview is required. The attitude of the vehicle is directly visible through the windscreen and defined by the vehicle heading angle deviation, $\psi_e = \psi_c - \psi$, where ψ_c is the commanded heading angle of the road geometry, ψ is the heading angle of the vehicle and ψ_e is the resulting heading angle deviation.

Although expressions of optimal driver control exist, such as [32], drivers are in principle free to choose any strategy that suits them. However, this freedom is limited to the visual information available, meaning that hiding segments of the visual scene strongly affects driver behavior [26] or performance [16], [18], [19].

Assuming control is described by a look-ahead parameter T_{la} (Point A in Fig. 2a), a perceived angle $\eta(t, T_{la})$ describes the current vehicle position with respect to the road center-line:

$$\eta(t, T_{la}) = \arcsin\left(\frac{y_c(t+T_{la}) - y(t)}{U_0 T_{la}}\right) - \psi(t) \approx \frac{y_c(t+T_{la}) - y(t)}{U_0 T_{la}} - \psi(t) \quad (1)$$

which is valid for small angles. This optical stabilization angle describes the geometry of the road within a *single* metric, combining information on the lateral position and heading angle, with U_0 the velocity of the vehicle.

B. Relation between Visual Metrics and Control Theory

From a control-theoretic point of view, a similar deduction can be made. Previous research ([33], [34]) has shown that drivers apply a form of preprocessing on the road curvature that affects the above mentioned stabilization angle. This behavior is modelled by $H_{o_f}(j\omega)$, a low-pass filter that attenuates the response to high frequencies, while passing the low frequencies, resulting in corner-cutting behavior:

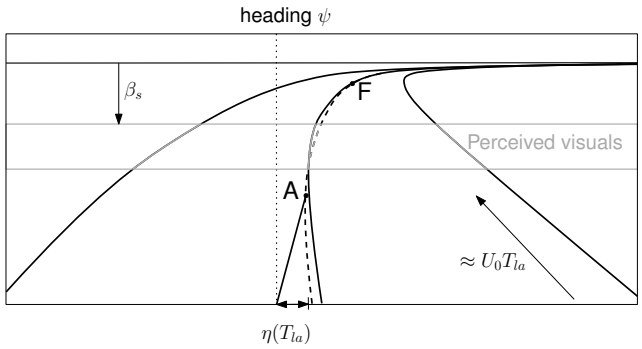

(a) Geometry of the interpretable viewing parameters not directly visible to the driver (black) and the view through the occlusion slit (grey), adapted to the full-visual geometry of [26].

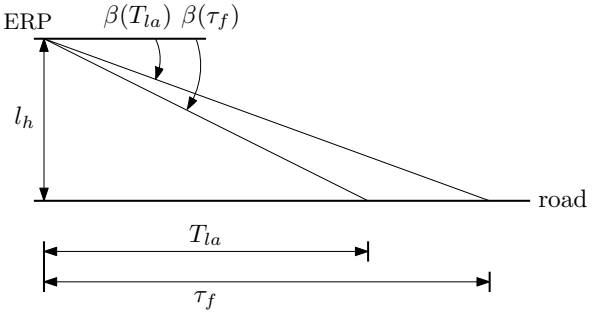

(b) Geometry of the interpretable viewing parameters linked to the occlusion angles, with exaggerated angles.

Fig. 2

$$H_{o_f}(j\omega) = K_f \frac{e^{T_{l,f} j\omega}}{1 + T_{l,f} j\omega} e^{\tau_c j\omega} \tag{2}$$

where K_f is the target response gain. The time constant for the low-pass filter $T_{l,f}$ is measured from a feedforward response at τ_c, such that if $T_{l,f} = 0$ s, drivers apply feedforward control exactly on the commanded road geometry. Therefore, to be able to describe corner-cutting behavior, visual information at $\tau_f = \tau_c + T_{l,f}$ is required (Point F in Fig. 2a. This feedforward control to the commanded road geometry is, as shown in Fig. 3, complemented by the feedback of the lateral position and vehicle heading angle, such that drivers aim at minimizing a single perceived error:

$$e^*(j\omega) = K_y^\psi \left(K_f \frac{e^{T_{l,f} j\omega}}{1 + T_{l,f} j\omega} e^{\tau_c j\omega} Y_c(j\omega) - Y(j\omega) \right) - \psi(j\omega) \tag{3}$$

where $e^*(t)$ represents the perceived error, such that if $T_{l,f} = 0$ s, the equation simplifies to the response of the true error $e(t)$. K_y^ψ is the weighting between the lateral position and the heading angle information. The negation of a single perceived error also has the implication that driver behavior can be described in compensatory fashion, which is well-described by the crossover model [35]. As [26] indicated, the control description in Eq. (3) is similar in structure to the geometrical description of the optical cue $\eta(t, T_{la})$ in Eq. (1). Therefore, by equating the relations:

$$K_y^\psi = \frac{1}{U_0 T_{la}} \tag{4}$$

and:

$$e^{T_{la} j\omega} = K_f \frac{e^{(T_{l,f} + \tau_c) j\omega}}{1 + T_{l,f} j\omega} \tag{5}$$

the first equation shows that the look-ahead time T_{la} is linearly related to the weighting of the lateral position with respect to the heading angle for a constant vehicle velocity. This shows that a smaller look-ahead time, thus looking closer to the vehicle itself, uses more lateral position information. Geometrically this is true, as due to the perspective viewing in car driving, the lateral position is visually amplified close to the vehicle, but gets more compressed (and is thus harder to perceive) further away from the vehicle.

Eq. (5) shows that the point τ_c approximates the look-ahead point T_{la}, depending on the value of K_f and $T_{l,f}$. For the lower frequencies (below $1/T_{l,f}$) this equation can thus still be valid for nonzero values of $T_{l,f}$. As the parameters T_{la} and τ_c correspond to different dynamics (feedback and feedforward, respectively), the parameters can *both* be determined, in which similar values for both parameters indicate the response to the same point on the road (i.e., the use of a single cue) [26].

C. Effect of Occlusion on Visual Scene

By hiding large parts of the visuals, occlusion changes the visual scene and the presented road geometry by definition. It forces drivers to obtain the visual cues through fixed slits on the road, as shown by the grey segment in Fig. 2a, located at β_s below the horizon. For example, showing visual information only close to the vehicle (high β_s) forces drivers to directly respond to the lateral position information (high K_y^ψ). Considering Eq. (4) and Eq. (5), this also affects the interpretable viewing parameters T_{la} and τ_f.

As Fig. 2a shows, both parameters correspond to points on the road. Thereby, they can be directly linked to the geometry of the occlusion. The distances that these parameters measure from the vehicle are $T_{la} U_0$ and $\tau_f U_0$. Fig. 2b indicates the angle below the horizon of these parameters, measured from the Eye Reference Point (ERP), which is located at l_h above the road. Therefore, the relation between the angle below the horizon and the look-ahead parameter T_{la} is:

$$\beta(T_{la}) = \arctan\left(\frac{l_h}{T_{la} U_0}\right) \tag{6}$$

and similarly for $\beta(\tau_f)$. By combining the geometries of the perspective view and the occlusion, system identification and parameter estimation can be used to measure changes in behavior for changing occlusion geometries.

III. System Identification Tools

To measure three independent driver responses required for the insight in cue dependencies, the system identification of

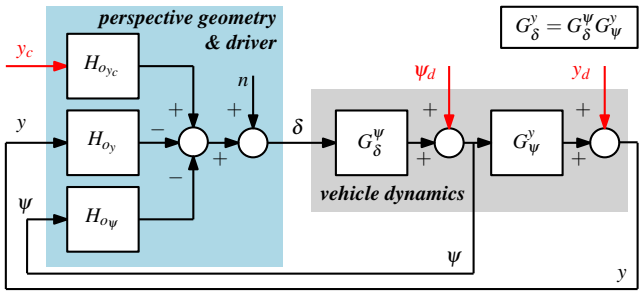

Fig. 3: Illustration of driver steering and multiloop driver dynamics, reproduced from [26].

driver data is required. An illustration of the parallel model structure as presented by [26], is shown in Fig. 3. This structure is in essence the same as in Fig. 1, but allows for the identification of the three control responses that drivers use when steering on winding roads: the feedforward (preview) response ($H_{o_{y_c}}$) to the commanded road geometry y_c, the feedback response (H_{o_ψ}) to the vehicle's heading angle ψ and the feedback response (H_{o_y}) to the lateral position y. The total driver steering output $\delta(j\omega)$ is then given as:

$$\delta(j\omega) = H_{o_{y_c}}(j\omega)Y_c(j\omega) - H_{o_y}(j\omega)Y(j\omega) - H_{o_\psi}(j\omega)\psi(j\omega) + N(j\omega) \quad (7)$$

where $Y_c(j\omega)$, $Y(j\omega)$ and $\psi(j\omega)$ are the Fourier transforms of the commanded road geometry, lateral position and vehicle heading angle signals, respectively. The remnant $N(j\omega)$ describes the non-linear part of the driver control dynamics. When comparing these dynamics to Fig. 1:

$$H_{o_\psi}(j\omega) = H_o^{comp}(j\omega) \quad (8)$$

$$H_{o_y}(j\omega) = K_y^\psi H_o^{comp}(j\omega) \quad (9)$$

$$H_{o_{y_c}}(j\omega) = H_{o_f}(j\omega)K_y^\psi H_o^{comp}(j\omega) \quad (10)$$

where $H_o^{comp}(j\omega) = K_{e^*}(1 + T_{L,e^*}j\omega)e^{-\tau_{e^*}j\omega}$, shown in Fig. 1. These dynamics are characterized by the equalization gain K_{e^*}, lead-time constant T_{L,e^*} and the human processing delay τ_{e^*}.

Fig. 3 also shows the vehicle dynamics, separated into the steering wheel-to-heading dynamics G_δ^ψ and heading-to-lateral position dynamics G_ψ^y. Double integrator are used throughout this paper, i.e., $G_\delta^\psi = \frac{K_\delta^\psi}{j\omega}$ and $G_\psi^y = \frac{U_0}{j\omega}$.

As explained in [36], when using identification techniques to estimate all three responses shown in Fig. 3, also *three*, uncorrelated multisine signals, indicated in red, are needed [27]. This is done by using not only the road geometry y_c as an independent signal, but also by adding two disturbances (acting as wind-gusts) that perturb the lateral position (by y_d) and vehicle heading (by ψ_d).

A. Identification Method

Estimating three independent responses of the control model, $H_{o_\psi}(j\omega)$, $H_{o_\psi}(j\omega)$ and $H_{o_\psi}(j\omega)$ of Eq. (7), also three equations are required [26], [37]. The remnant $N(j\omega)$ is typically relatively small at the excitation frequencies, compared to the complete linear response. Although the output $\delta(j\omega)$ is directly measurable at the excitation frequencies of the forcing function and the two disturbances, the information on the three estimated dynamics is only valid at those specific frequencies [37]. This means that, if the system of three equations is to be solved, interpolation of the signals $\delta(j\omega)$, as well as the three signals $Y_c(j\omega)$, $Y(j\omega)$ and $\psi(j\omega)$ is required to be able to estimate the dynamics at all those frequencies:

$$\begin{bmatrix} \delta(j\omega_{y_c}) \\ \tilde{\delta}(j\omega_{y_c}) \\ \check{\delta}(j\omega_{y_c}) \end{bmatrix} = \begin{bmatrix} Y_c(j\omega_{y_c}) & -\psi(j\omega_{y_c}) & -Y(j\omega_{y_c}) \\ \tilde{Y}_c(j\omega_{y_c}) & -\tilde{\psi}(j\omega_{y_c}) & -\tilde{Y}(j\omega_{y_c}) \\ \check{Y}_c(j\omega_{y_c}) & -\check{\psi}(j\omega_{y_c}) & -\check{Y}(j\omega_{y_c}) \end{bmatrix} \begin{bmatrix} H_{o_{y_c}}(j\omega_{y_c}) \\ H_{o_\psi}(j\omega_{y_c}) \\ H_{o_y}(j\omega_{y_c}) \end{bmatrix} \quad (11)$$

Here, the signals $\tilde{\delta}$, $\tilde{Y_c}$, $\tilde{Y}$ and $\tilde{\psi}$ indicate the interpolated signals from ω_{ψ_d} to ω_{y_c}. Similarly, $\hat{\delta}$, $\hat{Y_c}$, $\hat{Y}$ and $\hat{\psi}$ are the interpolated signals from ω_{y_d} to ω_{y_c}. Solving this system results in the estimation of the dynamics $H_{o_\psi}(j\omega_{y_c})$, $H_{o_\psi}(j\omega_{y_c})$ and $H_{o_\psi}(j\omega_{y_c})$, i.e. the dynamics valid at the target forcing function frequencies. Therefore, the same steps described above must also be applied to the lateral position disturbance frequencies (ω_{y_d}) and the heading angle disturbance (ω_{ψ_d}) to be able to describe the dynamics at all excitation frequencies.

B. Parameter Estimation

The parametric model of [26] is fitted using a cost function over all frequencies using a Nelder-Mead algorithm. The used cost function of the optimization is:

$$J(\hat{\Theta}) = \sum_{i=1}^{30} |\delta(j\omega) - \hat{\delta}(j\omega|\Theta)|^2 \quad (12)$$

This cost function is calculated at the frequencies of the input signal and the disturbances, resulting in 30 frequencies. $\delta(j\omega)$ and $\hat{\delta}(j\omega)$ are the measured and modelled Fourier transforms of the steering wheel output, respectively. The latter is optimized for the parameter vector $\Theta = [K_{e^*}\ T_{L,e^*}\ \tau_{e^*}\ K_y^\psi\ K_f\ \tau_f\ T_{l,f}]^T$, as defined in Section II. For the optimization, 30 sets of randomly generated initial conditions for all parameters are evaluated. The initial condition set that leads to the lowest cost function value J is used as the best fit. Furthermore, the Variance Accounted For (VAF) is calculated over the whole frequency range:

$$\text{VAF} = 1 - \frac{\sum_{k=1}^{N_f} |U(j\omega_b k) - \hat{U}(j\omega_b k)|^2}{\sum_{k=1}^{L} |U(j\omega_b k)|^2} \quad (13)$$

where N_f is the number of samples of the measured time traces and ω_b is the fundamental frequency. The VAF value represents the quality-of-the-fit; a value of 100% indicates a model that exactly describes the measured steering output.

During the optimization process, parameters were not allowed to become smaller than 0. If these situations would occur, a penalty value (10^6) was added to the cost function J, such that these model configurations would not be selected as the optimal solution.

TABLE I: Frequencies, amplitudes, and phases of the multisine approximation of the Queens Drive road in Edinburgh. The disturbance signals, reproduced from [26], [38], are given for completeness.

	Road center line ψ_c				Heading angle disturbance ψ_d				Lateral position disturbance y_d			
i	k	ω	A	ϕ	k	ω	A	ϕ	k	ω	A	ϕ
-	-	rad/m	m	rad	-	rad/m	m	rad	-	rad/m	m	rad
1	3	0.01	16.2	3.96	7	0.03	2.20	5.04	5	0.02	0.29	5.98
2	9	0.04	15.3	3.17	13	0.06	1.74	6.22	11	0.05	0.24	4.04
3	15	0.07	6.20	4.78	23	0.10	1.08	4.17	19	0.09	0.16	3.03
4	27	0.12	2.34	3.40	35	0.16	0.63	4.40	31	0.14	0.09	6.11
5	39	0.18	2.34	6.28	47	0.21	0.41	4.97	43	0.19	0.06	0.99
6	53	0.24	0.90	6.20	65	0.29	0.25	4.97	59	0.27	0.04	0.11
7	71	0.32	0.05	5.42	85	0.38	0.16	4.10	77	0.35	0.02	1.78
8	93	0.42	0.01	5.71	111	0.50	0.11	5.90	101	0.46	0.02	2.28
9	121	0.55	0.01	0.95	143	0.65	0.08	5.48	131	0.59	0.01	0.41
10	155	0.70	0.01	4.79	183	0.83	0.07	0.73	169	0.76	0.01	2.41

IV. Methods

A. Experiment Setup

1) Simulator: The experiment was performed in the SIMONA Research Simulator (SRS) at TU Delft. The motion system of the SRS was not used, such that participants performed a visual-only driving task, which was approved by the TU Delft Ethics Committee. Visuals were presented using the collimated 60×40 deg field-of-view visual system (single projector), see Fig. 4. The two side projectors of the SRS were switched off to match the experiment of [16] as well as possible.

Drivers were positioned in the left seat of the simulator, equipped with a steering wheel. The vehicle heading dynamics (G_δ^ψ in Fig. 3) was an integrator with constant $K_\delta^\psi = 1.33$ (deg/s)/deg, identical to [36]. The vehicle lateral position is obtained by integrating the heading angle (G_ψ^y), resulting in the double integrator dynamics between the steering wheel output and the lateral position on the road, as described in Section III. The road was 3 m wide and only showed the outside boundary lines in white. The task was performed with a constant forward vehicle velocity $U_0 = 16.9$ m/s, identical to [16], to remain as close as possible to their study.

2) Eye tracker: To measure the eye gaze, drivers wore a head-mounted Tobii Pro Glasses 2 eye tracker (component [a] in Fig. 4), which restricted the selection of drivers as they could not wear ordinary glasses. The eye tracker camera recorded video images of the visual scene, whereas its infrared sensor measured the eye gaze relative to the tracker. The measured gaze therefore depends on the position of the glasses on their heads. For this reason, calibration was required every time the glasses were put on. This process required drivers to look at a fixed point (component [b] in Fig. 4), of which the shape was automatically recognized by the eye tracker software and used for the correction of any offsets.

Due to the possible influence of infra-red radiation on the eye tracker, causing noise in the gaze measurements, the lights were turned off such that the only source of illumination was from the visual system. All heat-emitting devices nearby the glasses, such as the simulator internal camera surveillance system, were turned off for the same reason.

As the tracker can also be subject to measurement drift during the runs, as well as the possibility of moving the glasses by the participant, it was necessary to periodically check and possibly recalibrate the hardware. After every condition, drivers were required to look at a fixed LED light (component [c] in Fig. 4), which was visible in the dark. If this fixed point fell outside the uncertainty (0.5 deg) of the eye tracker, recalibration was required, where the lights were temporarily turned on.

3) Data Acquisition: The measurement portion of the road was 1389 m, yielding 82.0 s nominal measurement time (if the center of the road is followed perfectly), for the analysis of the control data. The actual road driven each run was extended by a run-in (277 m) and run-out (138 m) to avoid transient responses in the measurements, of which the data were not analyzed. The data of the eye tracker during the run- in and out were also not used.

B. Gaze processing algorithm

The head-mounted eye tracker measures the gaze direction with respect to the video frame. As the gaze data stream is stored at a different rate (60Hz) than the video frames (25Hz), it was first required to match these two together using a nearest-neighbour approach. Furthermore, for quantitative comparisons of the measured gaze to the characteristics of

Fig. 4: Picture of a participant performing the baseline condition, illustrating the experiment setup. Component [a] is the eye tracker, [b] is the eye tracker calibration point and [c] is the reference LED light.

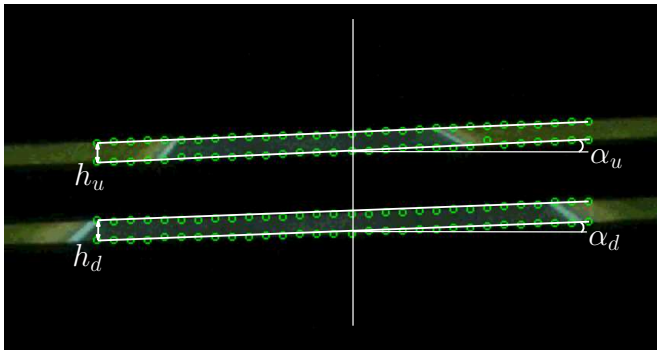

Fig. 5: Part of the video screen as recorded during the experiment, showing the road through the occlusion slits. The overlay indicates the edge detection points (green) and the geometry of the gaze frame correction algorithm (white).

the road, it is necessary to convert the eye gaze from its head-fixed to the world-fixed reference frame. As this requires the recognition of segments on the video screen of which the location is known precisely, such as the vertical position of the slits, it was not possible to correct the horizontal component as well, as there were no horizontally fixed reference points on the video screen. Therefore, only the vertical position below the horizon is analyzed is this paper. As a form of post-processing in each video frame, this was done using the following five steps:

1) To be able to detect the slits, the transition points from occluded to non-occluded parts and vice versa are detected by searching for segments in which the RGB levels change from black [0,0,0] to non-black (at least [0,40,0]) for the detection of the top of a slit, or vice versa for the bottom. To ensure that the occlusion was always visible, this step is done at 30 evenly spaced points ranging from $(\frac{5}{16})$ to $(\frac{11}{16})$ of the video frame, shown in Fig. 5. Linear regression is used to account for acquisition errors as much as possible, resulting in two lines for each slit (the top and bottom, see Fig. 5). In case of the baseline conditions, as no slits were present, the top of the video screen and the horizon were used for the edge detection, functioning as a large 'slit'.
2) For the slit, the rotation angle α_u and, if the second slit was present, α_d, are determined as the average of the angles of the slits two lines (white lines in Fig. 5) with respect to the video frame.
3) To determine the height of a slit in pixels, the average difference between the two *local* slit lines (that is, the top and bottom of a slit) is measured. As this difference is measured as the vertical difference in the video frame, it must be corrected for the rotation of the slit itself, such that the actual difference in pixels becomes:

$$\hat{h}_u = h_u \cos \alpha_u \tag{14}$$

and similarly for the bottom slit.
4) The y-coordinates of the measured eye gaze, which are also measured with respect to the video frame, are corrected for the rotation angle of the video frame, resulting in the y-coordinates of the gaze in the world-fixed frame:

$$y_w = -x_h \sin \alpha + y_h \cos \alpha \tag{15}$$

Where x_b and y_b are the measured gaze coordinates in the head-fixed reference frame. In the presence of two slits, the average of the angles α_u and α_d is used for α.
5) Finally, as the position of the slits is known for each condition, the relative position between the slits and the eye gaze can be used to extrapolate the vertical gaze position in degrees. As the estimated gaze depends on the extrapolation of the pixels-per-slit, the acquisition error amplifies if the gaze is located far away from the slits. As drivers typically look close to or between the visible slits, this is not considered a large problem. To account for this as much as possible, in the case of the presence of two slits, the steps mentioned above are performed separately for each slits and the two resulting gaze measurements are averaged.

The resulting vertical gaze position in the world-fixed frame can now be directly compared to the model parameters. The raw gaze data from the eye tracker are unfiltered and also includes the blinks of the eye, in which the view angle drastically decreases, shown by the blue line in Fig. 6. After the gaze conversion the data are therefore filtered with a low-pass filter with a passband frequency of 0.05π rad/s, stopband frequency of 0.8π rad/s and a stopband attenuation of 100 dB. This filters the large high-frequency blinks, but also some of the small saccadic eye movements, as shown by the red line in Fig. 6.

C. Forcing Function

Because a discrete-frequency forcing function is required for determining the Frequency Response Function (FRF) [26], the first part of the experiment included a comparison between a realistic road geometry of an existing route (Queens Drive in Edinburgh, United Kingdom, as used by [16]) and its multisine equivalent, structured as:

$$\psi(a) = \sum_{k=1}^{N_f} A[k] \sin\left(\omega[k]a + \phi[k]\right) \tag{16}$$

where N_f is the number of sinusoids, 10 in this case, and a is the along-track distance. The multisine frequencies $\omega[k]$ are chosen the same as in [26], whereas the amplitudes $A[k]$ were set to result in equal power along the complete spectrum of Queens Drive. Phases $\phi[k]$ were selected such that the heading angle variance of the resulting multisine signal compared to the Queens Drive road was the smallest. In [28], the first part of this experiment was analyzed by comparing modelled driver behavior for both the Queens Drive and its multisine equivalent (the baseline in this paper), showing that both geometries indeed induce similar driver behaviour.

The occlusion conditions use the same multisine signal as the baseline. The resulting amplitudes, frequencies and phases of the individual sinusoids are shown in Table I, which also includes the sinusoids of the heading angle and lateral position disturbances, which are the same as in [26].

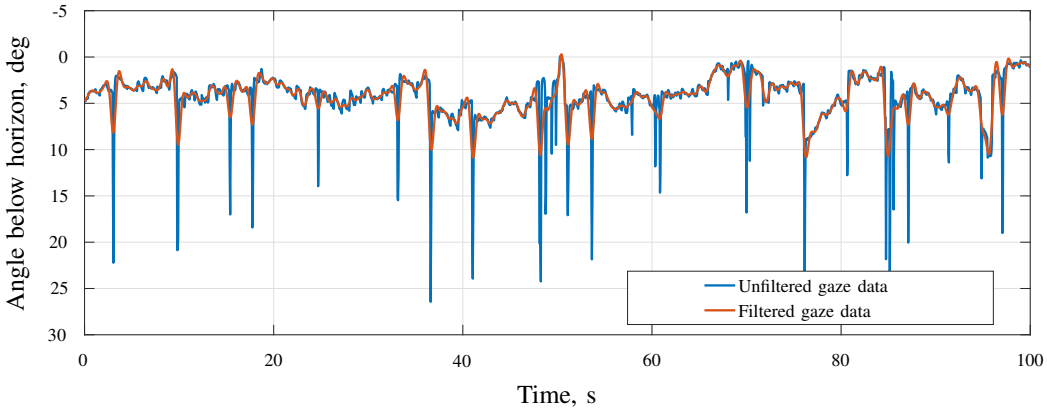

Fig. 6: Representative baseline gaze measurement showing the unfiltered and filtered data, subject 1, run 4.

D. Experiment Procedure

Nine male and three female volunteers ($\mu = 25.4$ yr, $\sigma = 3.3$ yr) participated in the experiment. All participants provided informed consent prior to their participation. As driving performance is affected by the amount of experience [39], they were all in the possession of a valid driver's license ($\mu = 7.2$ yr, $\sigma = 3.0$ yr), with varying travelled distances per year ($\mu = 4420$ km, $\sigma = 7447$ km).

Participants were instructed to drive as they would normally do and first completed five runs of the Queens Drive road of [16], which was used for the analysis in [28]. Subsequently, drivers performed four baseline (BL) runs. The first run was used to familiarize participants with the condition, and only the remaining, final three runs were used for analysis.

After both of these full visual conditions, the occlusion was added, while the road forcing function remained the same. The position of the slits was the same as in [16], such that three types of occlusion conditions exist. In single slit conditions (Denoted S), the position of a single slit was varied between 1° and 10° between the conditions. For top fixed conditions (Denoted T), a slit between 1° and 2° was always present, while the position of a second slit was varied between 1° and 10° between the conditions. Finally, for the bottom fixed conditions (Denoted B) it was the other way around; A slit at the lowest position (9°-10°) was always present, with a second slit varying in position between the conditions.

To account for training effects between the conditions, the order in which the occlusion conditions were presented varied per driver and was based on an unfinished latin-square design, shown in Table II. Participants performed four runs of each condition, of which all respective first runs were discarded, such that for all conditions three runs were used as measurements. Finally, as the baseline condition is used for relative comparison to all other conditions and is thus of high importance, an additional set of four baseline runs was performed *after* the occlusion conditions. This is to mitigate for additional training effects and fatigue over all conditions.

Due to the large amount of conditions used in the experiment (25 occlusion, two baseline and one Queens Drive condition for [28]), the experiment sessions were split up into two sessions of three hours each, but never on the same day. After every run, roughly 30 seconds were taken to relax the driver eye sight, ensure no misery effects were present and inform drivers about the performance of their last run (given as $\sigma^2_{y_e}$) to keep them motivated throughout the experiment. The completion of four conditions, in most cases corresponding to 16 runs, resulted in a 15 minute break outside of the simulator.

E. Dependent Measures

Steering behavior is compared between the occlusion conditions by various metrics, always including a comparison to the baseline (full visuals).

1) Performance and Control Activity: The performance is measured based on the vehicle lateral position y_e, which is defined as the vehicle deviation from the current centerline ($y_e = y_c - y$). For each condition, the standard deviation of the lateral position error (σ_{y_e}) is determined for each run and averaged over the three measurement runs. Then, to act as a direct comparison with [16], the performance of each condition is determined with respect to the baseline:

$$\text{Performance} = \left(\frac{\sigma_{y_e,m}}{\sigma_{y_e,bl}}\right)^{-1} \tag{17}$$

The inverse in Eq. (17) is used such that the relative performance is determined in a more intuitive manner; for a value below 1 the occlusion performance is *worse* than that of the baseline. Spreads over the subjects are used to compare between the conditions

Similar to the performance, the relative control activity is measured as the standard deviation of the steering wheel output (σ_δ) compared to the baseline:

$$\text{Control activity} = \frac{\sigma_{\delta_m}}{\sigma_{\delta_{bl}}} \tag{18}$$

In this case, the inverse is not used, as there is no advantage in doing so.

2) Driver Identification: Driver control behavior of occlusion conditions is analyzed by extracting the Frequency Response Functions (FRFs) of the measured time traces by following the approach outlined in Section III-A. The dynamics $H_{o_{y_c}}$ (target preview response), H_{o_y} (lateral position response) and H_{o_ψ} (heading angle response) are compared between the occlusion conditions and to the baseline.

TABLE II: Unfinished randomized latin-square for the condition sequence per subject. BL = Baseline, S = single slit, T = Top slit fixed, bottom varying and B = Bottom slit fixed, top varying. Numbers indicate the position of the top of the respective slit in degrees.

Driver	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19	C20	C21	C22	C23	C24	C25	C26
1	BL	S5	T5	S7	S8	S3	B2	S6	T3	T7	S1	T6	B4	B3	B5	S2	B6	T2	S4	B8	B7	T8	T9	S9	T4	BL
2	BL	T5	B5	T7	T8	T3	S2	T6	B3	B7	S9	B6	S4	S3	S5	T2	S6	B2	T4	S8	S7	B8	S1	T9	B4	BL
3	BL	T2	B2	T4	T5	S8	B7	T3	T8	B4	S6	B3	S1	B8	S2	S7	S3	T7	S9	S5	S4	B5	B6	T6	T9	BL
4	BL	B2	S2	B4	B5	T8	S7	B3	B8	S4	T6	S3	S9	S8	T2	T7	T3	B7	T9	T5	T4	S5	S6	B6	S1	BL
5	BL	T6	B6	T8	T9	T4	S3	T7	B4	B8	T2	B7	S5	S4	S6	T3	S7	B3	T5	S9	S8	S1	S2	B2	B5	BL
6	BL	B8	S8	S2	S3	B6	T5	S1	S6	T2	B4	S9	T7	T6	T8	B5	T9	S5	B7	B3	B2	T3	T4	S4	S7	BL
7	BL	B5	S5	B7	B8	B3	T2	B6	S3	S7	T9	S6	T4	T3	T5	B2	T6	S2	B4	T8	T7	S8	S9	S1	S4	BL
8	BL	T7	B7	T9	B2	T5	S4	T8	B5	S1	T3	B8	S6	S5	S7	T4	S8	B4	T6	T2	S9	S2	S3	B3	B6	BL
9	BL	T4	B4	T6	T7	T2	S1	T5	B2	B6	S8	B5	S3	S2	S4	S9	S5	T9	T3	S7	S6	B7	B8	T8	B3	BL
10	BL	S4	T4	S6	S7	S2	T9	S5	T2	T6	B8	T5	B3	B2	B4	S1	B5	S9	S3	B7	B6	T7	T8	S8	T3	BL
11	BL	S2	T2	S4	S5	B8	T7	S3	S8	T4	B6	T3	T9	T8	B2	B7	B3	S7	S1	B5	B4	T5	T6	S6	S9	BL
12	BL	B6	S6	B8	S1	B4	T3	B7	S4	S8	B2	S7	T5	T4	T6	B3	T7	S3	B5	T9	T8	S9	T2	S2	S5	BL

3) Model Fitting and Parameter Estimation: The parametric driver model from [26], as explained in Section III-B, is fitted to the steering wheel deflection $\delta(j\omega)$ data, of which the parameters of the occlusion conditions are compared to the baseline. The first two measurement runs are used for the identification and modelling, while the third run is always used to compute the Variance Accounted For (VAF) of the fitted models.

4) Gaze Behavior: Finally, the timetraces of the viewing gaze are analyzed. Due to limitations in the experiment setup only the vertical viewing gaze is considered, as mentioned in Section IV-B. Individual runs with a bad data quality (such as the eye tracker failing to detect the eye) are discarded. Over all subjects, all three measurements runs per subject are grouped together, such that every condition type has one distribution describing behavior over all subjects. Histograms containing the spread of the data are produced. In principle drivers can apply different gaze positions during a single run, resulting in inaccurate data representation by their means. Therefore the medians are used as a method of comparing between the various conditions. These values are grouped together per condition and calculated as distributions over all subjects.

F. Hypotheses

During the experiment, the following hypotheses are tested:

- H.I For single slit conditions, drivers adapt their behavior to the available information, such that the interpretable viewing parameters T_{la} as well as the mean focus of the viewing gaze, lie within the slits. τ_f will be close to 0 as no information beyond T_{la} is available.
- H.II For double-slit conditions, drivers will also adapt their behavior to the available information, but are more free to place the parameters and gaze between the slits. As visual information can be extracted from two distinct slits, the difference between T_{la} and τ_f will be more pronounced.
- H.III The changes in driver responses are still well captured by the parametric model, as it is based directly on the available geometry, resulting in comparable VAF values for all conditions.
- H.IV The results of Land & Horwood will be verified, although the shifts in the overall trends will be present, as the richer textures induce more optic flow.

V. Results

A. Relative Performance and Control Activity

To serve as a direct comparison to the original study by [16], Fig. 7 shows the measured relative performance box plots for all occlusion conditions with respect to the baseline performance, indicating the spread across the drivers. The different types of occlusion configurations (S, B, and T) are grouped together by a color (green, blue and red, respectively). Second-order polynomials are fitted for each configuration type to show general trends. The figure also includes the normalized standard deviation indicating the spread of the baseline condition runs (grey band). Finally, to compare the results of this study to that of Land & Horwood, their measured relative performances are also indicated in the same figure, indicated by points and second-order fitted polynomials (same S, B, and T color configuration, increased transparancy).

The measured relative control activity is shown in Fig. 8, which uses the same color configuration and also displays the spread of the baseline runs. As it is unclear how Land & Horwood defined their ”instability index”, these results are not directly compared.

1) Single Slit: The single slit conditions are indicated by the green box plots in Fig. 7 (performance) and Fig. 8 (control activity). In the condition with the lowest performance, S12, performance of the baseline is on average 1/0.38 = 2.6 times higher than for this occlusion condition. By placing the slit lower on the screen, performance gradually increases, up to the middle range of slit positions (S56 and S67), which approximate the bottom of the baseline performance spread. Putting the slit any higher (S78 - S910) gradually decreases performance again down to an average of 0.7 times the baseline performance, meaning that performance is not as low as for the slits located at the top of the screen (such as S12 and S23). The results show a similar but slightly better performance ratio than of Land & Horwood. The increase over all single slit conditions compared to their study is consistently around 0.1.

Similarly, the control activity is determined and shown in Fig. 8. Control activity is comparable to baseline values between S12 and S45, but sees a strong increase for lower placed slits. For slits S89 and S910, it thus shows that these conditions can be controlled at a slight decrease in

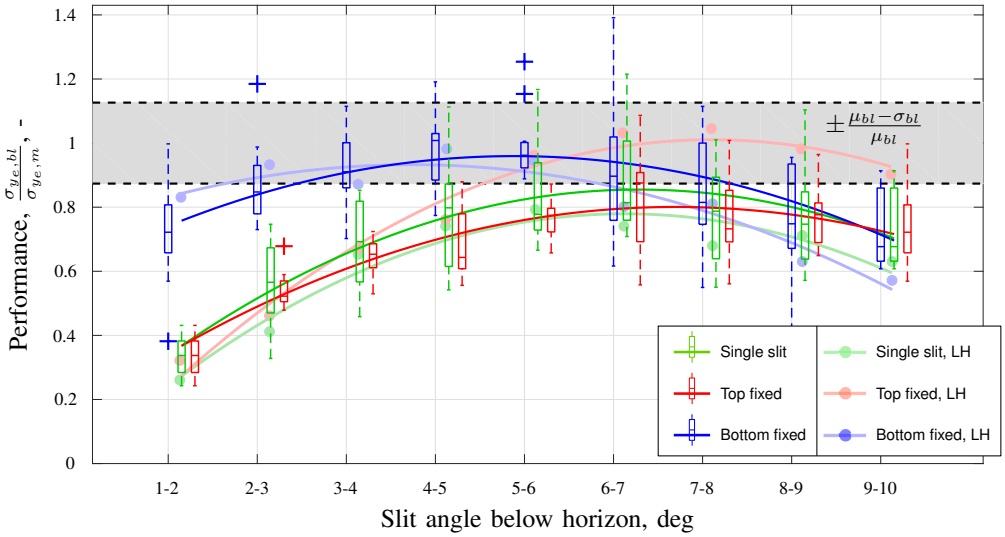

Fig. 7: Performance as function of the slit configuration relative to the baseline.

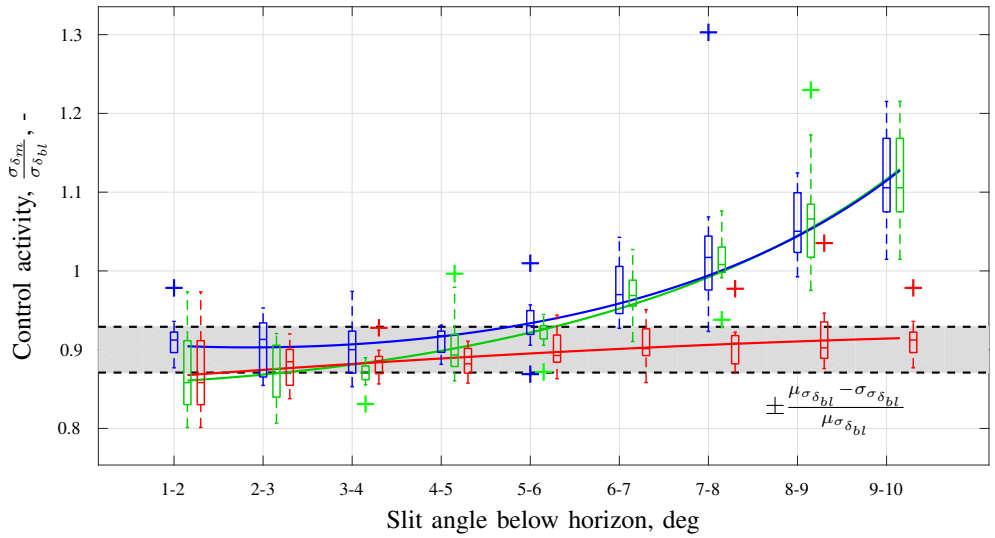

Fig. 8: Control activity as function of the slit configuration relative to the baseline.

performance (compared to baseline), but at a strong increase in *control activity*, of up to 1.5 the baseline value. An explanation is that for low-placed slits, lateral position information is more directly visible, such that drivers can more actively try to achieve a higher performance (directly defined by the lateral position deviation) compared to cases where the lateral position information is barely visible. However, the lacking feedforward information at these areas requires a high control activity to retain stable control. In that sense, drivers are forced to either have a baseline-comparable performance but high control activity, or a baseline-comparable control activity at the cost of a strongly decreased performance.

2) Double Slit: Compared to single slit conditions, adding a second slit to the visual scene has an effect on the measured performance. Surprisingly, adding a second slit at the top of the screen (the top-fixed conditions, red box plots) even slightly decreases the relative performance compared to the single slits, although the spread also consistently decreases. At the lowest slit positions (T12, T23), performance is again around 0.38 and for the highest slit positions (T89, T910) around 0.7. The best configuration is T67 (performance of 0.8), which does not reach the baseline performance. Top-fixed conditions do consistently result in a relative control activity of 1. The difference with Land & Horwood is strong, mainly for low-placed slits (T56 - T910), where they measured a better performance of up to 0.2. As [19], [22] indicate, the presence of more realistic texture affects driver behaviour. As this experiment used a textured road and grass around it, versus the simpler two-color setup used by Land & Horwood, the road itself might be more difficult to detect at larger distances. Drivers can neglect this information of the top slit, effectively only using visual information from a single slit. This argument is strengthened by the fact that in this experiment, the top-fixed conditions show an almost identical performance as the single

slit conditions.

The bottom-fixed conditions (blue box plots) generally shows the best relative performance for conditions B23 - B67, reaching values close to 1 and well within the baseline spread. Again, this indicates a strong dependency on a low-placed slit for a high performance, due to the necessity of lateral position information for reaching high performances. Control activity is close to the baseline for conditions B12 - B45, but suffers greatly for conditions below B56, and sees an increase almost identical to the single slit conditions (up to 2.5). This might indicate that for those conditions a second slit adds little to the feedforward control required for baseline-comparable stabilization. Compared to the top-fixed conditions, the difference to Land & Horwood is smaller and has higher performance over the complete range of angles.

B. Non-parametric Estimation

The non-parametric estimates of the dynamics are shown in Fig. 10, divided in the single slit (a), top-fixed (b) and the bottom-fixed (c) conditions, averaged over all subjects. The baseline is included in all plots by the grey data, to serve as a comparison. The occlusion scenarios all have three representative conditions to show the trends in the measured behavior. The colors represent the slit positions, for example dark green represents a high-placed single slit (S12) and light green a low-placed single slit (S89) in Fig. 10a.

1) Single Slit: Fig. 10a shows the responses to the single slit tasks. In condition S12, the phase response of the feedforward dynamics increases for high slit positions (S12) at high frequencies, indicating that drivers indeed look further ahead compared to conditions S56 and S89. This effect is even stronger (i.e., drivers look further ahead) than for the baseline. Nevertheless, the magnitude of that same response has the inverse trend, such that drivers react *stronger* to the little available feedforward information, if the slit is located low on the screen. Similar effects were measured in [40]. Furthermore, this condition shows a neglection of a lateral position response, indicating that for the highest slit position possible, drivers almost neglect their lateral position feedback loop.

Drivers adapt their lateral position magnitude response depending on the occlusion condition. This is as expected, as drivers will respond to the lateral position information more strongly if this type of visual information is also more present within the scene; as explained in Section II; the lateral position information becomes more difficult to perceive at larger distances. Only the heading response sees values comparable the baseline, indicating that drivers strongly rely on heading angle information.

Especially present in the feedforward dynamics $H_{o_{y_c}}$, the magnitudes have a flat trend at low frequencies, indicating that drivers steer equally strong to such frequencies. At high frequencies ($\omega > 0.30$ rad/m), there is a sudden strong increase in the magnitude by a factor of 10^2, similar as in [26]. It is contradictory to the modelled driver response explained in Section II, which predicts a *decrease* in magnitude at high frequencies due to the low-pass filter behavior by preprocessing the commanded road geometry. This behavior is present

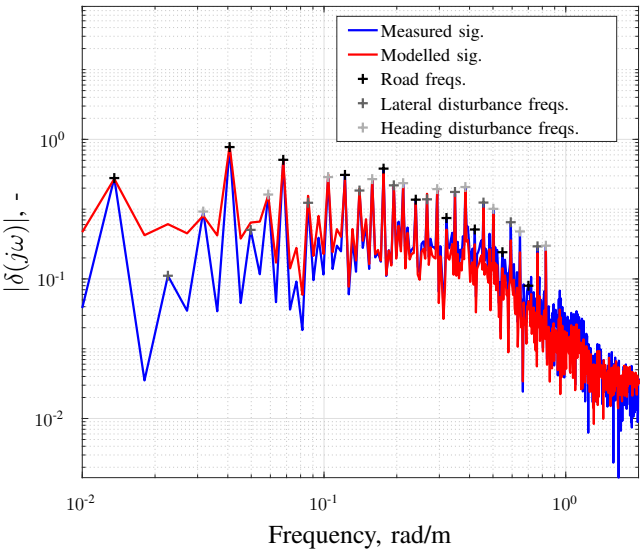

Fig. 9: Measured and modelled magnitude response of the steering wheel deflection $\delta(j\omega)$, Subject 1.

in *all* condition types, indicating that it is not an artefact of the presence of occlusion, but might be caused by drivers not directly responding to these frequencies, such that the remnant becomes dominant. This effect is indeed observed in the measurements of $\delta(j\omega)$, of which an example is shown in Fig. 9.

2) Double Slit: By adding an additional fixed slit to this configuration in principle allows drivers to extract the visual metrics from both slits. The target responses $H_{o_{y_c}}$ in Fig. 10b indicate that behavior changes little between the top-fixed conditions. Again it should be noted that for frequencies $\omega > 0.30$ rad/m the shown data do not correspond to a coherent driver response. The phase response of $H_{o_{y_c}}$ in Fig. 10b shows that for lower positions of the additional slit, drivers aim less far ahead (similar as in [29]). By comparing the feedforward phase responses of Fig. 10a and Fig. 10b, the phase of T34 is *lower* than that of S23, indicating that drivers look less far ahead if a second slit is available, and that the the additional top-fixed slit is not actively used for the feedforward response (i.e., drivers place their look-ahead parameter closer). For both condition types, there are no strong differences in the lateral position and heading angle responses H_{o_y} and H_{o_ψ}.

C. Parameter Estimation

The parametric model of [26] is fitted to the frequency signal of the steering wheel output δ. Therefore, the individual responses, described in the previous subsection, can also be modelled. Fig. 11 shows a representative example of the model fitted to measured baseline FRF data for subject 7, including estimated parameters. The complete parameter estimation sets for all conditions, displayed by boxplots of the subject spread, are shown in Fig. 12. The parameter K_f, the low-pass filter gain of between the actual and processed center-line, was included in the parametric fits, but was almost always close to one and showed little difference between the different

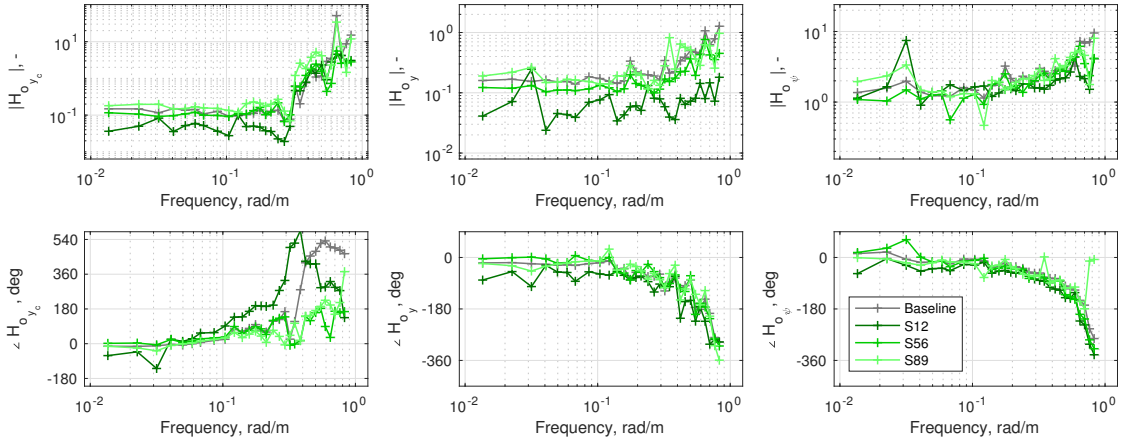

(a) Single slit

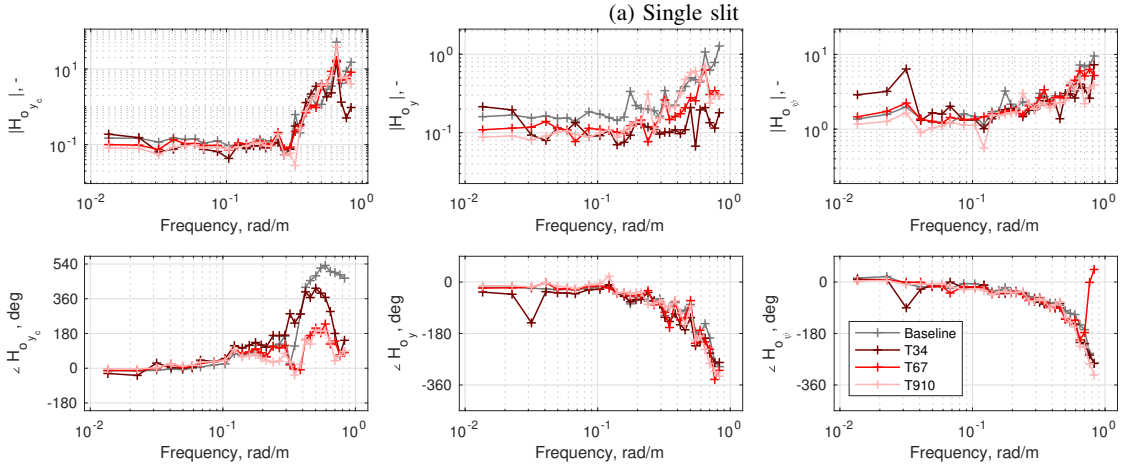

(b) Double slit, top fixed

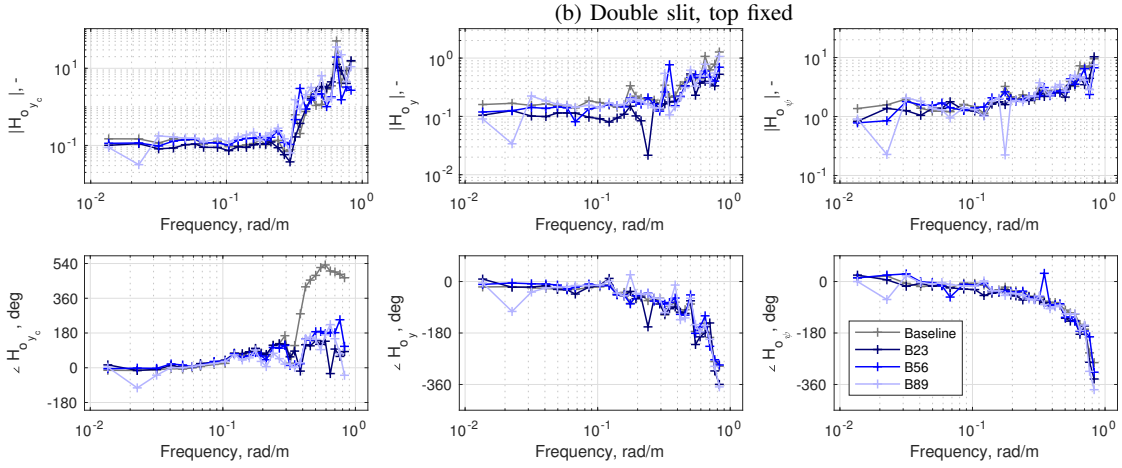

(c) Double slit, bottom fixed

Fig. 10: FRFs of selected conditions, averaged over all subjects.

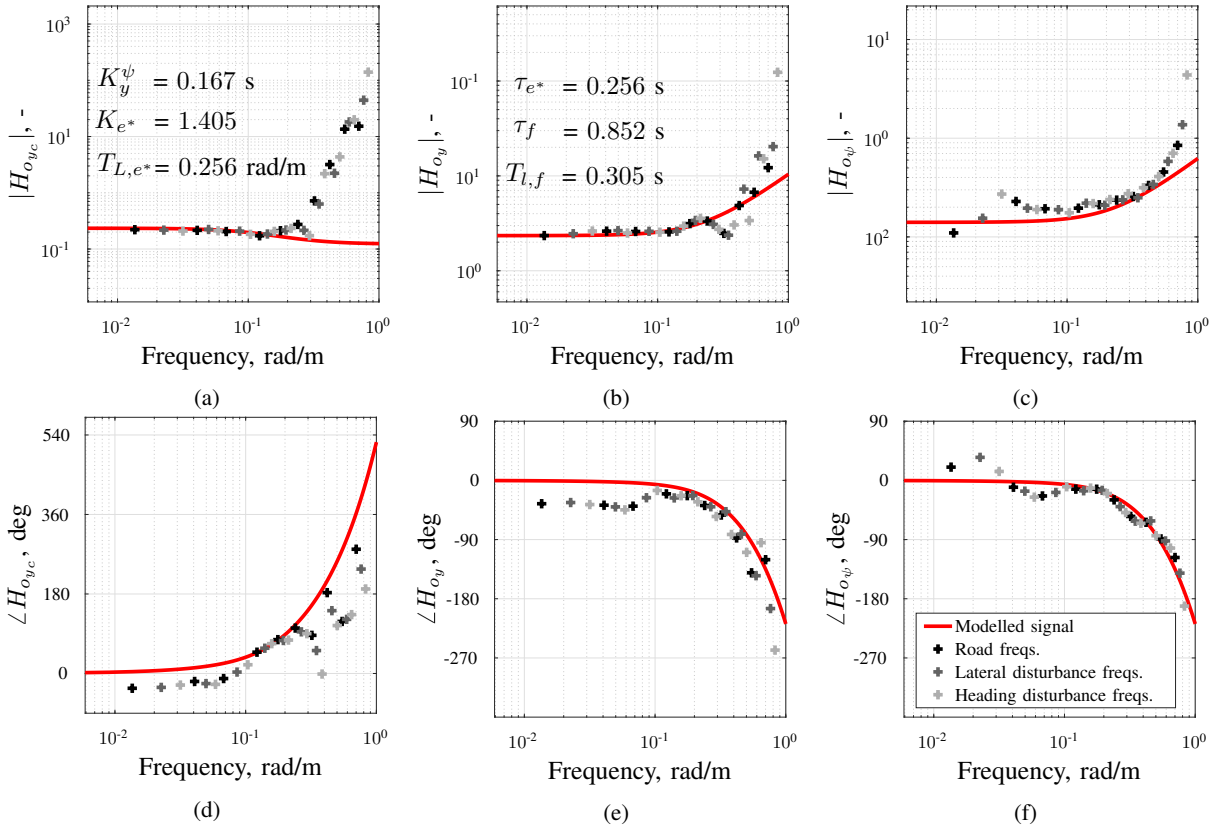

Fig. 11: Representative FRF estimates and model fits of the baseline condition, subject 7.

conditions, (i.e., $K_f = 0.99 \pm 0.02$). Therefore it was not included in the figures.

The calculated VAFs describe how well the model parameters are able to describe the measured responses and are shown in Fig. 12a for all conditions. The VAFs are generally well above 90%, although for the S12 condition (single slit closest to the horizon) it is slightly lower (ranging from 80% to 95%), indicating a larger remnant contribution. The different trends within occlusion types are separated by the left (single slit), middle (top fixed) and right (bottom fixed) plots. The baseline is also included to serve as a direction comparison.

1) Single Slit: Of the three occlusion types, the single slit conditions show the largest change in behavior. First of all, the green box plots in Fig. 12(b) shows that K_{e^*}, representing the gain to the perceived error of the aim point, has a positive increasing trend for lower slit positions. Directly affecting the cross-over frequencies of the open loop dynamics, this indicates that drivers react stronger to the amount of information that is present, as they respond to a point closer-by, and must steer stronger to still remain on the road. Surprisingly, T_{L,e^*} remains constant over all conditions. This indicates that drivers are not more actively generating lead to increase their stability.

This is confirmed by the human processing delay τ_{e^*} in Fig. 12(d), which shows no increase for low slit positions, whereas increases in lead generation typically induce higher time delays [35]. Surprisingly, the time delay only increases from 0.4 s to 0.5 s for for high placed slits. Although without lead generation the delay is expected to always be roughly the same, as it stems from the human inherent physiological constraints, it is possible that drivers deliberately delay their behavior, i.e., drivers *wait* until the perceived trajectory of the road is closer by and internally estimate their current road position to increase their performance.

The lateral position-to-heading angle gain K_y^ψ is shown in Fig. 12(e) and indicates that drivers rely more on lateral position information for low slit positions. This is in correspondence with Eq. (3), as the further down from the horizon the visual information comes, the more directly the lateral position is observable compared to the vehicle heading angle information. Only the lower positioned slits entail a similar gain compared to the baseline behavior.

The far-view point τ_f in Fig. 12(f) shows the opposite trend compared to K_y^ψ. By determining the correlations of the look-ahead parameter T_{la} between the feedback ($T_{la} = \frac{1}{U_0 K_y^\psi}$) and feedforward ($T_{la} = \tau_f - T_{l,f}$) responses in Fig. 13 (green points), the correlation is almost unity. This indicates that drivers consistently respond to a single cue for all single slit conditions, as expected for a single slit (see Section II and [26]). Finally, the time constant $T_{l,f}$ for the preprocessed corner cutting behavior is heavily affected under occlusion and never reaches the baseline, shown in Fig. 12(g). There is no single condition under which drivers approximate baseline behavior.

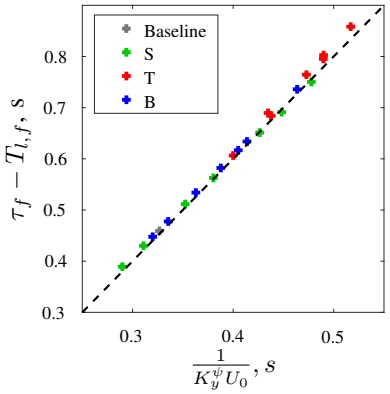

Fig. 13: Correlation between feedback and feedforward aim-point estimates for all conditions, averaged over all subjects.

2) Double Slits: The center and right plot of Fig. 12 represent the top fixed and bottom fixed conditions, respectively. Slightly increasing values of K_{e^*} in the top-fixed conditions of Fig. 12(b) again indicate a more aggressive response to the optical cue, and remains constant for bottom-fixed conditions. The time constant τ_{e^*} in Fig. 12(d) for top fixed conditions again shows the possible deliberate 'delaying' the human response, and is not visible in the bottom fixed conditions, as it is not necessary when the bottom slit is also present.

The lateral position-to-heading angle gain suprisingly that even for bottom-fixed conditions, the baseline values are not reached. Only the two conditions that together resemble (almost) a large, single slit (conditions B78 and B89) approximate the baseline weighting. This shows that drivers respond to the lateral position information more strongly in full visual conditions, similarly as was measured in their FRF estimates (Fig. 10c).

Both types of double slit conditions show a stronger presence of the corner cutting time constant $T_{l,f}$, especially in the conditions which have a clear distinction between the top and bottom slit (for example T78 and B45). This indicates that drivers only apply this corner cutting behavior if enough additional information allows them to do so. Considering that for double slit conditions the visual information sources are clearly separated from each other, it is surprising that, as Fig. 13 shows by the red points (top fixed) and blue points (bottom fixed), drivers still respond to a single cue (unity correlation). This indicates that drivers internally process the information perceived from both slits and form an approximation of the view-points at which they are responding to, i.e., interpolate the information.

D. Relation between gaze and aim point parameters

The measured eye gaze data were not always of sufficient quality to be used for analysis. For two participants, the processed data could not be used, as it showed highly noisy data. This was most likely caused by the tracker failing to detect the eye, which was confirmed by the fact that these subjects required recalibration after almost every condition. Furthermore, in the data of a third participant, part of the experiment setup was partially in the way of the video screen, caused by the participant not sitting exactly in the middle of the seat. Therefore, the edge detection algorithm failed to produce data for this participant, meaning that these data were also discarded. This resulted in nine participant data sets.

The calculated medians for each condition per subject are shown in Table III. The baseline eye gaze distributions are shown in Fig. 14a. It shows that there is a high variability between the drivers as to which segment of the road they look at for baseline control. This freedom is not available under occlusion scenario's. Fig. 14b shows another distribution, now for the condition S45, where it is clearly visible that there is a peak around the visible slit (dashed lines). A same situation is visible in B78 (Fig. 14d), for two close slits. If the difference between the slits becomes more pronounced, such as is the case in T67 (Fig. 14c), drivers place their gaze between the slits.

As a comparison to the occlusion geometry, Fig. 15 contains the previously determined spreads of T_{la} (feedforward) by means of green boxplots and the the far-view point τ_f, of which the difference is by definition the preprocessing constant $T_{l,f}$. The spread of the eye gaze medians $\tilde{\beta}$ is shown by the red boxes. The model parameters are directly related to the slits through Eq. (6) (pictured as horizontal black lines).

1) Single Slit: The left segment in Fig. 15 shows the single slit conditions gaze medians and modelled aim points. Drivers clearly adapt their gaze to the available slit, as this is where the information must be perceived from. Single slit conditions also result in an adaptation of T_{la} to the slit geometry, with almost no difference to the feedforward point τ_f.

2) Double Slit: A clear difference is present with the double slit conditions, in the middle and right of Fig. 15. If the slits are distinctly separated from each other, looking at two slits at the same time is impossible. Drivers are also no longer bound to the geometry perceived through the slit, but prefer to compensate for their perceived error *behind* the occlusion, indicating that a form of interpolation of information from both slits is used, as explained in Section V-C. This effect becomes stronger if slits are located further away from each other. Locating the eye gaze near the top slit is preferred, most likely because the information from the bottom slit can also be

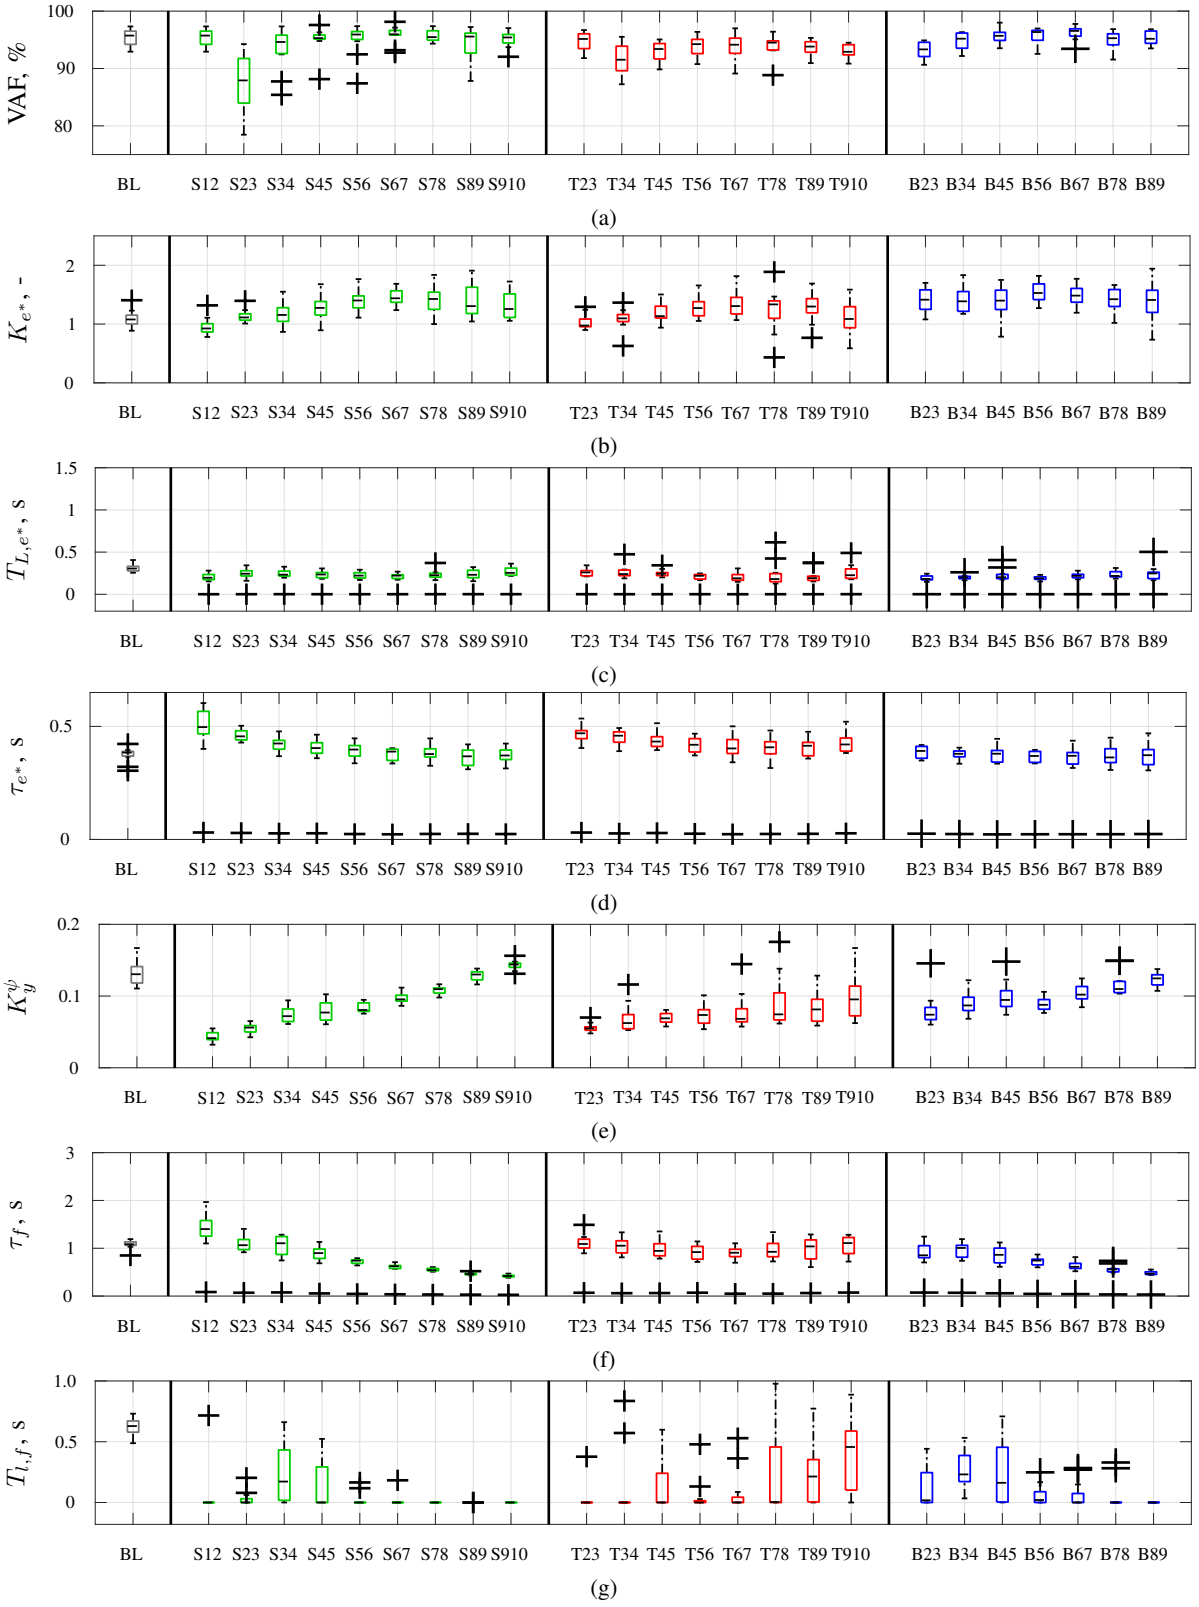

Fig. 12: Estimated parameters of the model of [26] for all subjects.

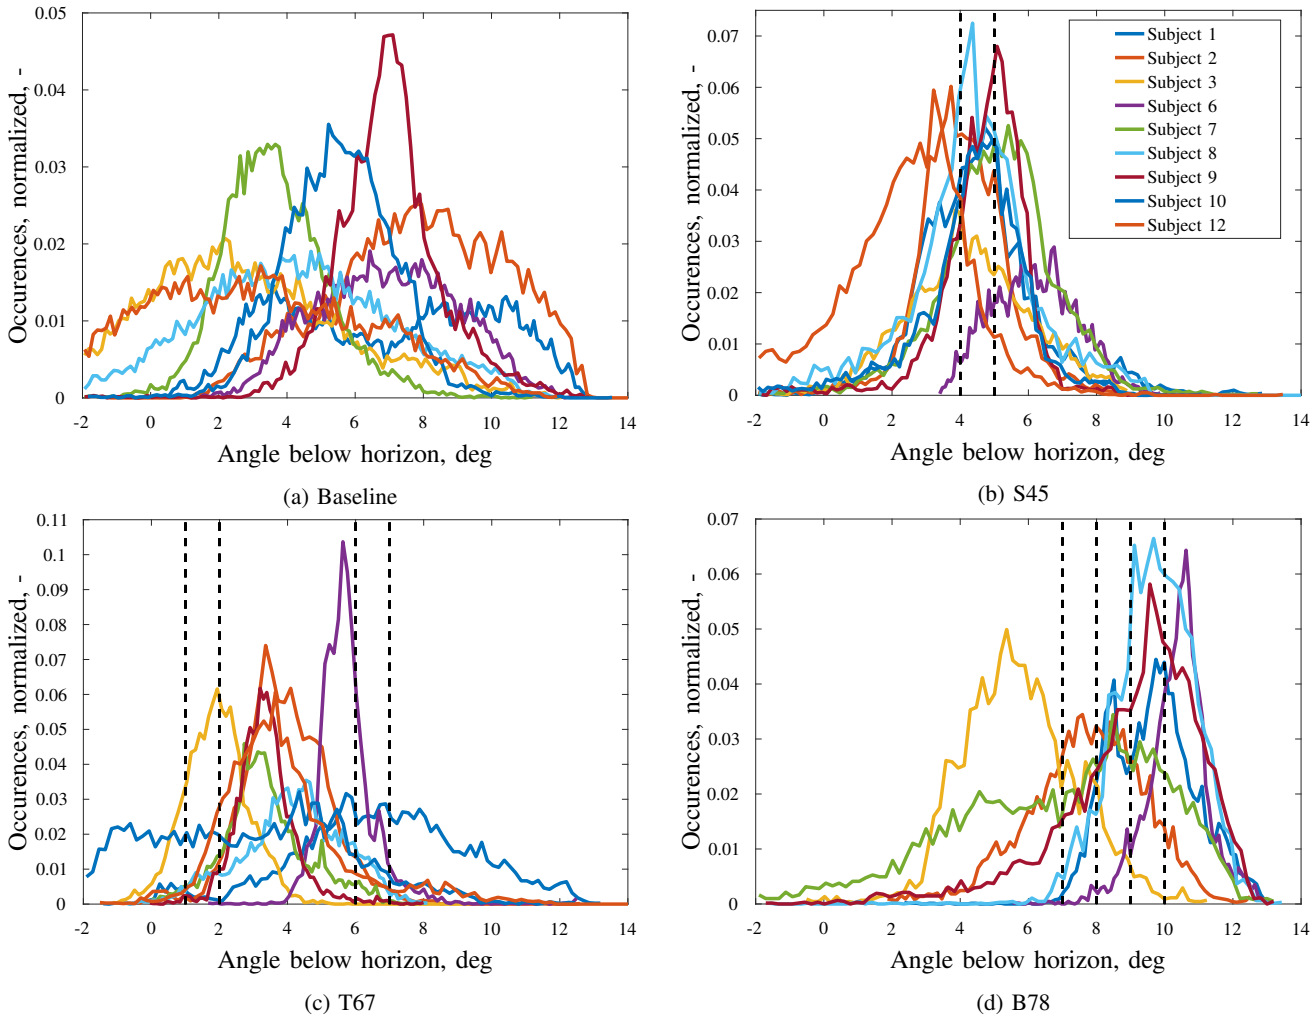

(a) Baseline

(b) S45

(c) T67

(d) B78

Fig. 14: Gaze distributions of the baseline and selected occlusion conditions for all subjects, including slit positions (dashed lines).

perceived through the peripheral vision, whereas the top slit requires more attention to be perceived due to low optic flow. If the slits move far away from each other a similar pattern as for T_{la} and τ_f is observed; drivers also aim their gaze *between* the slits, indicating that peripheral vision is used for both slits. Although gaze measurements are always subject to hardware inaccuracies, it is still observable that there is a correlation between the median gaze and the far-view point, rather than with the look-ahead constant T_{la}, especially present in the top-fixed measurements.

VI. Discussion

The goal of this paper was to replicate the original study performed by Land & Horwood [16] and increase the understanding in the use of visual perception in lane keeping tasks. By the estimation of the three individual FRF dynamics, parameter estimation based on the model of [26] and the subsequent eye gaze measurements, the perception of the visuals can be linked to the use of visual metrics guiding the control.

Drivers were in all occlusion cases able to successfully drive their vehicle along the winding road, resulting in high VAF values of the linear model. Presenting single slit conditions show that drivers adapt their aim-point and the viewing gaze to the occlusion geometry and that it is possible (optimally for condition S56) to drive a vehicle by a single aim point, perceived from a single visual cue point, as predicted by [26]. This results in only a small decrease in performance and increase in control activity. As no further information is possible above the single slit, corner cutting behaviour by placing the far-view point τ_f further than T_{la} is limited, but not necessarily zero. In principle drivers can place T_{la} after the occlusion slit with τ_f within the slit, still effectively corner cutting. These results show that drivers are never inclined to do so, and although stable driver control can be described by a single visual area, drivers seem to extract visual information from wider areas of the scene.

This is partially confirmed by the double-slit conditions, which allow drivers to use a larger range of visual regions for high performance measures. Drivers consistently place their aim-point and viewing gaze *between* the two occlusion

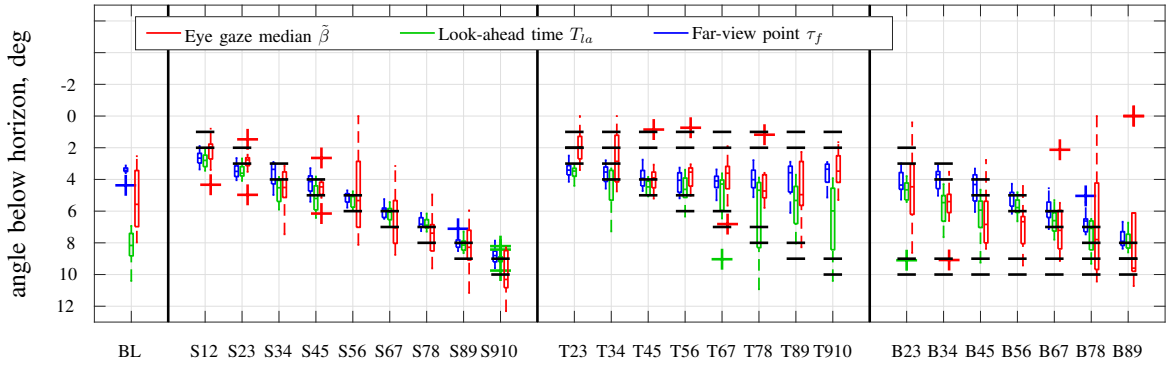

Fig. 15: Box plots of the estimated aim point parameter and median viewing gaze, relative to the slit position (outlined in black).

TABLE III: Measured medians of the eye gaze in degrees below the horizon.

	S1	S2	S3	S6	S7	S8	S9	S10	S12
BL	7,01	6,96	5,57	3,17	7,97	2,53	6,97	3,53	4,22
S12	4,33	1,5	2,17	0,80	1,86	1,90	2,48	1,96	3,40
S23	3,52	2,75	2,77	3.00	2,44	2,88	4,97	2,69	1,47
S34	7,46	3,60	5,84	4,91	3,28	3,75	4,76	4,51	3,13
S45	4,47	4,88	4,56	2,64	4,04	4,22	6,15	5,16	4,37
S56	6,08	3,81	6,98	6,30	5,01	5,33	7,14	8,13	-
S67	7,25	8,41	6,72	3,15	5,66	4,34	8,76	7,05	7,91
S78	4,95	8,50	7,68	6,44	7,01	8,57	7,39	6,99	9,61
S89	9,11	7,29	8,33	6,96	7,72	5,95	1,12	9,07	8,28
S910	1,14	1,03	8,52	1,07	8,28	8,39	1,23	1,03	9,27
T23	3,41	1,60	2,55	2,57	0,33	3,10	1,94	1,81	-
T34	4,76	1,62	2,88	1,40	2,86	1,74	4,51	4,03	-
T45	5,21	3,06	4,48	3,90	3,69	0,85	4,39	4,68	3,89
T56	3,35	3,39	4,08	3,04	5,03	0,74	4,8	3,54	4,30
T67	6,82	3,29	3,01	3,61	3,88	1,90	5,65	3,23	4,24
T78	4,91	3,74	5,80	4,03	4,99	1,18	5,66	3,58	4,73
T89	5,35	4,26	8,28	2,26	6,37	2,35	5,40	3,04	4,96
T910	3,84	1,65	3,14	2,49	5,30	3,48	4,97	2,51	3,93
B23	4,59	4,46	8,63	0,77	8,10	0,40	5,57	4.00	3,80
B34	6,62	5,25	5,92	3,51	9,08	5,40	5,90	4,13	5,21
B45	7,88	6,84	8,38	2,77	7,63	3,43	8,37	6,09	6,03
B56	7,79	6,61	7,76	6,36	6,67	4,41	9,43	8,85	6,28
B67	8,3	5,47	8,49	2,13	6,11	7,21	9,15	6,54	8,34
B78	9,66	9,74	9,49	3,46	7,80	5,63	1,04	7,48	-
B89	9,76	8,15	9,60	5,97	9,92	9,22	1,07	9,68	-

slits if the difference between the slits becomes emphasized, extracting information from both slits while still responding to a single visual cue (i.e., a direct correlation between feedforward and feedback look-ahead times), thus effectively interpolating the information from both slits. This shows that visual perception of the road and the use of its visual metrics are not necessarily the same, and provides an explanation for the contradictory results found by [26], compared to the two-level modes of driver behavior ([17]). Furthermore, values of τ_f increase compared to single slit conditions, but drivers are still not necessarily inclined to place this point in the top slit.

The final hypothesis relates back to the original work of Land & Horwood. Similar trends were measured for single slit and bottom-fixed condition, although top-fixed conditions performed significantly less. Both [22] and [19], the latter specifically for occlusion, reported a dependency on the presence of texture of the visuals, which is significantly different in the case of this experiment compared to Land & Horwood, which used simple white-on-black visuals and might make the information from the highest located slits easier to perceive.

Finally, although the experiment was performed at a constant velocity, it is in principle possible to extend the parametric model and therefore the analysis of occlusion at other (even varying) velocities. Further research is required to understand how the parametric model adapts to this change in perception for both full visuals and occlusion scenario's.

VII. Conclusions

In this paper, a replication study based on the occlusion study of Land & Horwood was described. Based on a human-in-the-loop lane keeping experiment with twelve participants, Frequency Response Functions (FRFs) estimates of the target feedforward, lateral position feedback and vehicle heading angle feedback were estimated, including the fitting of a parametric model on the driver control data. For the first time insight in the behavioral changes under occlusion was given. For both single- and double slit conditions, it was shown that drivers adapt their modelled look-ahead, far-view points as well as the eye gaze within the occlusion slits. In the presence of two distinctly separated slits, drivers aim both their eye gaze and the look-ahead points *between* the slits, effectively interpolating the information from both presents

slits while still responding to a single optical cue, although preprocessing behavior is strongly affected. These results show a high adaptability to the visual scene and although the results of Land & Horwood were partially verified, this is not an indication of two-level driver control.

REFERENCES

[1] M. Sivak, "The information that drivers use: is it indeed 90% visual?" *Perception*, vol. 25, no. 9, pp. 1081–1089, 1996.

[2] J. J. Gibson and L. E. Crooks, "A Theoretical Field-Analysis of Automobile-Driving," *The American Journal of Psychology*, vol. 51, no. 3, pp. 453–471, Jul. 1938.

[3] J. P. Wann and M. F. Land, "Steering with or without the flow: is the retrieval of heading necessary?" *Trends in Cognitive Sciences*, vol. 4, no. 8, pp. 319–324, Aug. 2000.

[4] R. M. Wilkie and J. P. Wann, "Judgments of Path, Not Heading, Guide Locomotion," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 32, no. 1, pp. 88–96, Feb. 2006.

[5] A. V. van den Berg, "Judgements of Heading," *Vision Research*, vol. 36, no. 15, pp. 2337–2350, 1996.

[6] A. J. Grunwald and S. J. Merhav, "Vehicular Control by Visual Field Cues," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 6, no. 12, pp. 835–845, 1976.

[7] M. F. Land and J. Horwood, "How Speed Affects The Way Visual Information Is Used In Steering," *Vision in Vehicles*, vol. 6, pp. 43–50, 1998.

[8] J. P. Wann, D. Swapp, and S. K. Rushton, "Heading perception and the allocation of attention," *Vision Research*, vol. 40, no. 18, pp. 2533–2543, Aug. 2000.

[9] R. M. Wilkie and J. P. Wann, "Eye-movements aid the control of locomotion," *Journal of Vision*, vol. 3, pp. 677–684, Nov. 2003.

[10] R. R. Mourant and T. H. Rockwell, "Mapping Eye-Movement Patterns to the Visual Scene in Driving: An Exploratory Study," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 12, no. 1, pp. 81–87, Feb. 1970.

[11] K. L. Macuga, A. C. Beall, R. S. Smith, and J. M. Loomis, "Visual control of steering in curve driving," *Journal of Vision*, vol. 19, no. 5, May.

[12] R. J. Jagacinski, E. Rizzi, B. J. Bloom, O. A. Turkkan, T. N. Morrison, H. Su, and J. Wang, "Drivers Attentional Instability on a Winding Roadway," *IEEE Transactions on Human-Machine Systems*, submitted.

[13] S. Kolekar, J. C. F. de Winter, and D. A. Abbink, "A Human-like Steering Model Sensitive to Uncertainty in the Environment," in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, Banff, Canada*, Oct. 2017.

[14] P. Bosetti, M. Da Lio, and A. Saroldi, "On Curve Negotiation: From Driver Support to Automation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 4, pp. 1–12, Aug. 2015.

[15] I. Frissen and F. Mars, "The effect of visual degradation on anticipatory and compensatory steering control," *The Quarterly Journal of Experimental Psychology*, pp. 1–8, Aug. 2013.

[16] M. F. Land and J. Horwood, "Which parts of the road guide steering?" *Nature*, vol. 377, pp. 339 – 340, Sep. 1995.

[17] E. Donges, "A Two-Level Model of Driver Steering Behavior," *Human Factors*, vol. 20, no. 6, pp. 691–707, Dec. 1978.

[18] S. Cloete and G. Wallis, "Visuomotor control of steering: the artefact of the matter," *Experimental Brain Research*, pp. 475–489, Jan. 2011.

[19] A. Chatziastros, G. M. Wallis, and H. H. Bülthoff, "The Effect of Field of View and Surface Texture on Driver Steering Performance," *Vision in Vehicles*, vol. VII, no. 71, 1999.

[20] D. D. Salvucci and R. Gray, "A two-point visual control model of steering," *Perception*, vol. 33, no. 10, pp. 1233–1248, Dec. 2004.

[21] C. Mole, G. Kountouriotis, J. Billington, and R. Wilkie, "Optic flow speed modulates guidance level control: New insights into two-level steering," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 42, no. 11, pp. 1818–1838, Nov. 2016.

[22] P. van Leeuwen, C. G. I. Subils, A. R. Jimenez, R. Happee, and J. C. F. de Winter, "Effects of visual fidelity on curve negotiation, gaze behaviour and simulator discomfort," *Ergonomics*, vol. 58, no. 8, pp. 1–18, feb 2015.

[23] D. T. McRuer, D. H. Weir, H. R. Jex, R. E. Magdaleno, and R. W. Allen, "Measurement of Driver-Vehicle Multiloop Response Properties with a Single Disturbance Input," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 5, no. 5, pp. 490–497, Sep. 1975.

[24] D. T. McRuer, R. W. Allen, D. H. Weir, and R. H. Klein, "New Results in Driver Steering Control Models," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 19, no. 4, pp. 381–397, Aug. 1977.

[25] D. H. Weir and D. T. McRuer, "Dynamics of Driver Vehicle Steering Control," *Automatica*, vol. 6, no. 1, pp. 87–98, Jan. 1970.

[26] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder, "Multiloop Identification and Modeling of Driver Visual Information Processing and Steering Control," 2019.

[27] H. J. Damveld, G. C. Beerens, M. M. van Paassen, and M. Mulder, "Design of Forcing Functions for the Identification of Human Control Behavior," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 4, pp. 1064–1081, Jul. 2010.

[28] M. Kolff, K. van der El, D. Pool, M. van Paassen, and M. Mulder, "Approximating Road Geometry with Multisine Signals for Driver Identification," 2019.

[29] E. Rezunenko, K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder, "Relating Human Gaze and Manual Control Behavior in Preview Tracking Tasks with Spatial Occlusion," in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, Myazaki, Japan*, Oct. 2018, pp. 3440–3445.

[30] M. F. Land and D. N. Lee, "Where we look when we steer," *Nature*, vol. 369, pp. 742 – 744, Jun. 1994.

[31] K. van der El, D. M. Pool, and M. Mulder, "Measuring and Modeling Driver Steering Behavior: From Compensatory Tracking to Curve Driving," *Transportation Research Part F*, 2017, online preprint.

[32] C. C. MacAdam, "An Optimal Preview Control for Linear Systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 102, pp. 188–190, Sep. 1980. [Online]. Available: http://hdl.handle.net/2027.42/65011

[33] K. van der El, D. M. Pool, H. J. Damveld, M. M. van Paassen, and M. Mulder, "An Empirical Human Controller Model for Preview Tracking Tasks," *IEEE Transactions on Cybernetics*, vol. 46, no. 11, pp. 2609–2621, Nov. 2016.

[34] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder, "Effects of Preview on Human Control Behavior in Tracking Tasks With Various Controlled Elements," *IEEE Transactions on Cybernetics*, vol. 48, no. 4, pp. 1242–1252, Apr. 2018.

[35] D. T. McRuer and H. R. Jex, "A Review of Quasi-Linear Pilot Models," *IEEE Transactions on Human Factors in Electronics*, vol. HFE-8, no. 3, pp. 231–249, Sep. 1967.

[36] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder, "A Unifying Theory of Driver Perception and Steering Control on Straight and Winding Roads," *IEEE Trans. on Human-Machine Systems*, 2019, submitted.

[37] R. Pintelon and J. Schoukens, *System Identification - A Frequency Domain Approach*. John Wiley and Sons, 2012.

[38] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder, "Effects of Linear Perspective on Human Use of Preview in Manual Control," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 496–508, Oct. 2018.

[39] P. van Leeuwen, R. Happee, and J. C. F. de Winter, "Changes of Driving Performance and Gaze Behavior of Novice drivers During a 30-min Simulator-based Training," *Procedia Manufacturing*, vol. 3, pp. 3325–3332, jan 2015.

[40] K. van der El, S. Padmos, D. M. Pool, M. M. van Paassen, and M. Mulder, "Effects of Preview Time in Manual Tracking Tasks," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 486–495, Oct. 2018.

III

Literature Review

This part includes the literature study report as performed for the separate AE4020 Literature Study course.

1 Introduction

Manual control forms the cornerstone of many practical applications that people use throughout their daily life. An important example is car driving, where drivers manually keep their cars within the instructed lane and away from other cars. With the relatively low level of training required to enter the domain of car transportation, it is paramount to understand how drivers steer their cars within their lane. Furthermore, the recent development of autonomously driving vehicles or lane-keeping assistance systems has given arise to the question how humans actually perform this type of control, as well as what cues determine their actions.
In the simplest variant of these situations, the information presented to the driver is available in a constant, steady flow (same rate of visual information due to fixed velocity, no discontinuous visuals such as road signs, buildings and surrounding traffic). The driver is only responding to the dynamics of the road and its visuals. Although driving on a road induces motion, the information presented by the road is mostly visual (Sivak [22]). These tasks in which drivers aim at keeping the vehicle between the outer lines is a **lane keeping task**, thus a subcomponent of any-day car driving. Modelling the human behaviour in such situations can directly be used for the design of automated systems, for example by designing these systems in such a way that they behave similarly as human drivers.

1.1. Control Models of Lane Keeping

Over the years many studies have attempted to describe this human behaviour with control models, with varying degrees of success. Although many model forms exist (see Steen et al. [23]), the most accepted form is the *two-point* model, of which multiple variations have been investigated. The first attempt by Donges [4] proposed a relatively simple variant by stating that human control can be described by two stages (hence their name): An open loop anticipatory response, where drivers purely observe the oncoming curvature, and a closed loop compensatory response, where drivers aim to minimize their lateral position deviation through feedback. The latter is the obvious task of the driver; make sure that the vehicle does not leave the road, whereas the former provides the driver with anticipatory information to smoothen its response and compensate for its delay.

(a) (b)

Figure 1.1: Fully displayed and occluded road visuals.

1.2. Visual Perception

For many years, the other side of the coin has been the study of visual perception, in which the most important goal is to show humans actually *perceive* visual information from their environment. Eye tracking and determining the gaze direction might seem as the most straightforward approach. The time distribution gives information on where drives focus their view point, but it gives no insight in how areas outside that view point (parafoveal and peripheral vision) might contribute to control. Practically, this method is also often subject to tracking problems (such as not being able to detect the eye direction) or calibration issues.

A different approach, visual occlusion, has been the most used and discussed method in the field of visual perception. By occluding large parts of the road and leaving slits open for drivers to look through, the most important visual regions of the road can be identified (see Figure 1.1). Land & Horwood [8], as well as others (Cloete & Wallis [3], Chatziastros et al. [2]), performed such experiments with varying setups. Interestingly, these occlusion studies all disagree on their findings and their underlying causes, but they might provide a form of validation of the one- or two point control models, as it forces drivers to obtain visual information from distinct areas. Land & Horwood described that the use of two slits (Figure 1.1b) not too close to each other is the minimum for an identical control compared to full visuals (Figure 1.1a), indicating that indeed at least two regions of visual information are required. Contrary to analysis of eye tracking, visual occlusion *does* allow for the investigation of the effect of non-focal view. The addition of a second slit might give a driver additional information through parafoveal or peripheral vision, increasing their performance. Until now, these occlusion experiments have investigated the effect of occlusion only by a performance measure, such as a difference in the standard deviation of the lateral position, compared to the 'full view'. This gives a limited understanding of how the individual responses to the obtained metrics actually work, as the same performance can be achieved by various control strategies.

1.3. Combining visual perception with control models

Although control-theoretic models provide useful human behaviour models, they often do not describe how visual information is perceived. A more recent research, described

in Van der El [25], proposed a model that seems to simplify this connection to the visual-perceptual approach. The use of system identification tools allowed for the disentanglement of *individual* driver responses to the signals of the lateral position, heading angle and target. Drivers minimize a single error on an 'aim point' in compensatory fashion, much as humans would do in simple compensatory tasks, already described in the sixties by McRuer & Jex [13]. It is of yet unclear how this model relates to the findings of the occlusion experiments, which are suitable in finding the the spatial regions guiding the control. Furthermore, combining a visual occlusion experiment with system identification and parameter estimation allows the disentanglement of individual, multiloop, driver responses and the prediction of control parameters as function of the presence of occlusion its location, such that these can be compared to a full-visual scenario.

1.4. Report structure

This report first gives a brief overview of the manual control for classical compensatory, pursuit & preview and precognitive control in Chapter 2, which serves as a basis for the explanation for understanding the control-theoretic models in Chapter 3. Chapter 4 describes the methods of visual perception, explaining how the understanding of the use of visual information by humans can be used to also better understand the models that describe their behaviour. The current lack of understanding, the research questions and the experiment setup are given in Chapter 5. Finally, Chapter 6 summarizes this report.

2

An Introduction to Manual Preview Control

As the use of visual information in car driving can essentially be seen as a extended preview task, this chapter aims at introducing classical control tasks in terms of their underlying characteristics and boundaries. Although manual control has various applications and variations, the hierarchy in which the perception and control occurs for general applications can be structured by the *Successive Organization of Perception (SOP)*. This framework, as presented by Krendel & McRuer [7] aims at capturing the stages in which humans can perform their control tasks. The three corresponding types of control (compensatory, pursuit and precognitive control) are discussed briefly in each setion.

2.1. Compensatory control

In the most basic form of the SOP the human responds to and controls a single element, the displayed error. The error is defined as the difference between the target signal and the output of the controlled system:

$$e(t) = f_t(t) - x(t) \tag{2.1}$$

Where $f_t(t)$ is the target input signal to be followed and $x(t)$ is the output signal of the controlled element. In compensatory tracking, the aim is to follow the input signal, meaning that a controller constantly aims at minimizing the error $e(t)$. Only this error is fed back to the human controller; it will thus not see the input signal that drives the dynamics, nor the output of the dynamics it controls.
The pilot dynamics as response to this error is denoted as $H_p(j\omega)$ and drives the steering signal $u(t)$ using a manipulator to the controlled element $H_c(j\omega)$. Figure 2.1 gives a schematic overview of this configuration. Any non-linearities and time-varying effects that occur in the system are captured by the remnant $n(t)$.
In manual control tasks the controller is often human. The complete process from perception of the visuals on the display to the output of a control signal $u(t)$ can be captured as part of the human controller $H_p(j\omega)$, thus it also includes the dynamics of the manipulator (such as a control stick or steering wheel).
Human controllers will never be able to achieve perfect error rejection for random input signals, because they are limited by their physiological constraints and the precision in

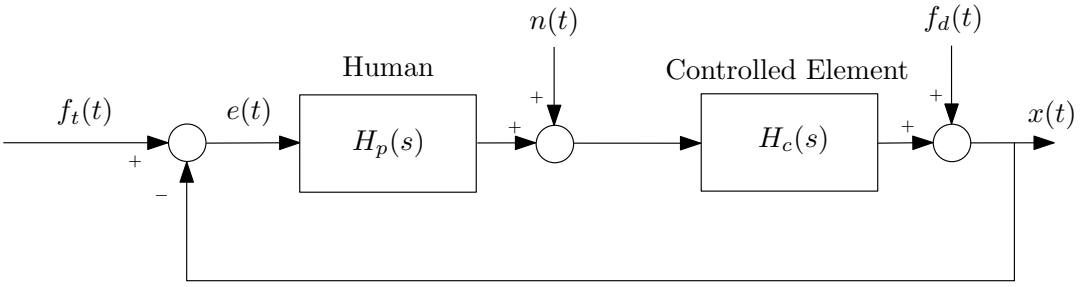

Figure 2.1: Control overview of a compensatory tracking set-up.

execution of the human and the manipulator. The delay caused by the perception of the visuals, processing of the information in the brain and execution by sending commands to the muscles is captured in a single time constant τ_e, resulting in an exponential term in the frequency domain:

$$H_d(j\omega) = e^{-j\omega\tau_e} \tag{2.2}$$

Hence a time delay induces a phase difference between the in- and output signal. Furthermore, the dynamics of the muscles themselves also induce a lag in execution. These dynamics are commonly described by a second order mass-spring-damper system:

$$H_{nms}(j\omega) = \frac{\omega_{nms}^2}{(j\omega)^2 + 2\zeta_{nms}\omega_{nms}(j\omega) + w_{nms}^2} \tag{2.3}$$

Where ω_{nms} is the natural frequency and ζ_{nms} is the damping coefficient. The dynamic relations between in the input- and output signals as closed- and open-loop representations are, in the frequency domain:

$$H_{ol,ft}(j\omega) = \frac{X(j\omega)}{E(j\omega)} = H_p(j\omega)H_c(j\omega) \tag{2.4}$$

$$H_{cl,ft}(j\omega) = \frac{X(j\omega)}{F_t(j\omega)} = \frac{H_p(j\omega)H_c(j\omega)}{1 + H_p(j\omega)H_c(j\omega)} \tag{2.5}$$

Perfect tracking is achieved when the output signal $x(t)$ matches the input signal $f_t(t)$. Signals are equal 1) if their magnitude is equal, hence $|H_{cl,ft}(j\omega)| = 1$, and 2) if there is no phase difference between the signals, $\angle H_{cl,ft}(j\omega) = 0$.
Furthermore, a signal $f_d(t)$ can be modeled such that it represents a disturbance acting on the controlled element. The addition of this disturbance becomes of importance when identifying separate underlying human controller dynamics.
McRuer and Jex [13] found that a human *adapts* its dynamics $H_p(j\omega)$ to the dynamics of the controlled element $H_c(j\omega)$ such that the dynamics of the complete open loop system approximates integrator dynamics around the crossover region, with an effective time delay:

$$H_{ol}(j\omega) = H_p(j\omega)H_c(j\omega) = \frac{\omega_c}{j\omega}e^{-j\omega\tau_e} \tag{2.6}$$

Where ω_c is the crossover frequency, which is defined as the frequency where the magnitude of the open loop dynamics becomes one, ie. $|H_{ol}(j\omega)| = 1$. Equation 2.6 is only valid

for compensatory tracking and is therefore often limited to few practical applications, although it does give the fundamental insight that the high human control adaptability can actually be measured and predicted, as long as the dynamics of the controlled element are known.
For typical controlled elements $H_c(j\omega)$ such as K_c, K_c/s or K_c/s^2 a model can be constructed that describes the adapted behaviour of the human controller, as given in equation 2.7.

$$H_p(j\omega) = K_p\left(\frac{1+T_L j\omega}{1+T_I j\omega}\right)e^{-j\omega\tau_e} \tag{2.7}$$

Containing the controller gain K_p, the delay term τ_e and the two characterizing lead (T_L) and lag (T_I) terms.

Typically, humans will aim at controlling the error that is presented to them (compensatory control) as a 'first base' (Mulder et al [18]); it is the most fundamental situation and control task if skill for the specific task is limited. If preview of the input signal is available, a controller can also use this information.

2.2. Pursuit and Preview Control

As is the case with car driving, a human controller will in fact have visual cues give more information than just the error and indicate the target signal to be followed, known as *pursuit control*. In this case, the human does not only respond to the error $e(t)$, but also to the input signal $f_t(t)$ and thus also implicitly to the output state $x(t)$, as shown in 2.1. An overview of such a control scheme is given in figure 2.2. The human controller can be seen as a 'summation' of three individual responses to the input signals.A feedforward response with dynamics $H_{p_t}(j\omega)$ describes the direct response to the target signal. A feedback response $H_{p_x}(j\omega)$ on the current system output and finally a direct compensatory response $H_{p_e}(j\omega)$ to the error. Each response is still subject to the physiological restraints of the human body.
It is of great fundamental insight to understand *how* these individual response functions are used in manual control, although due to the overdetermination of the system this is impossible, resulting in a choice which blocks are wanted for identification.

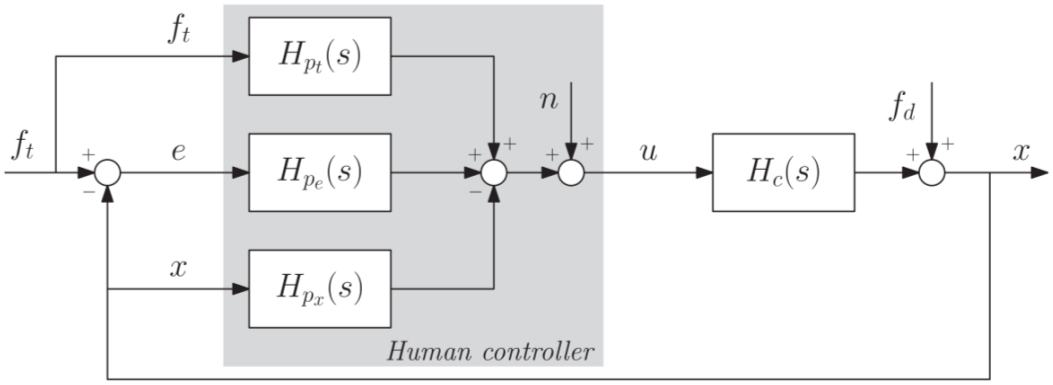

Figure 2.2: Control overview of pursuit tasks, adapted from Mulder et al. [18].

A special case of pursuit tasks is with the presence of preview. In this case, not only the forcing function at time t is shown to the human controller, but a stretch of future forcing function points up to $t+\tau_p$ is present. This way, a controller can choose what information is used to guide their behaviour, for example to allow for the compensation of their own delay by producing lead, or use no preview information at all (in which the human response equals that of its pursuit counterpart).
Shown in Figure 2.3 is the preview model by Van der El et al. [26]. Humans low-pass filter target information of a far view point into a processed error signal $e^*(t)$, which only exist internally and can smoothen the signal. The near view point response acts as a high pass filter. An overview on the identification of human dynamics in preview tasks for various controlled elements can be found in and Van der El et al. [28].

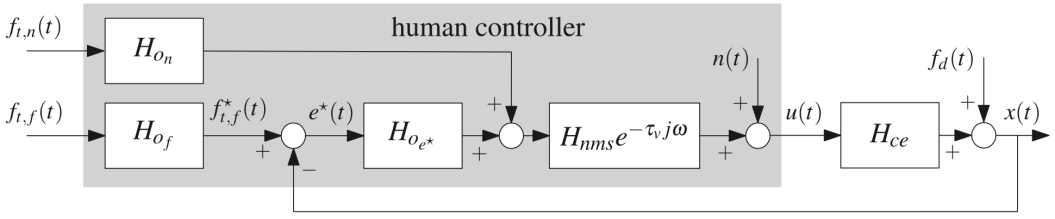

Figure 2.3: Control overview of preview tasks, adapted from Van der El et al. [28].

2.3. Precognitive Control

The final step within the SOP is the precognitive control (Figure 2.4), corresponding to the highest level of skill. In this method of controlling, drivers are directly responding to a situation within the forcing function from a learned response; it does not require feedback of the current state and/or error for successful completion and is completely open loop. Because precognitive control describes an internal prediction of the output, it requires an approximate knowledge of the controlled element, hence a deep level of experience is required.

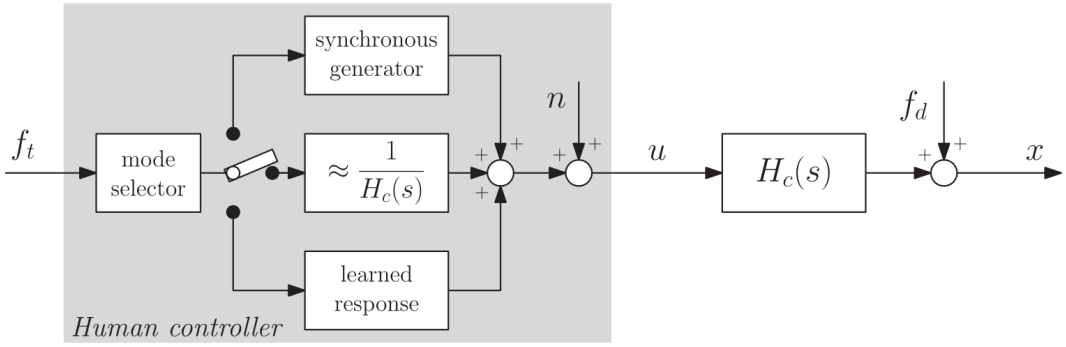

Figure 2.4: Various precognitive modes for human control, adapted from Mulder et al. [18].

Lappi [11] furthermore distinguishes between *predictive* and *ballistic* forms of precognitive behaviour. In the former, the aforementioned predictability will lead to an inverse motor code to match the forcing function based on an internal model; the driver has

knowledge on the input (such as in the case of a repetitive system) and thus directly compensates by the inverse of the controlled element from skill. Although the execution of control is generally open loop, feedback must be used, such as to see if the signal indeed acts as predicted. Ballistic precognitive is completely open loop and is a learned response executed as a motor program; an example would be a direct reflexive response to a sudden deviation in the forcing function, such as a hole in the road in car driving situations.
Predictable inputs (such as repetitive signals) can lead to precognitive behaviour, hence for the design of forcing functions it is necessary to create quasi-random signals that avoid recognition of patterns, which could lead to an entanglement of precognitive and preview control, making it difficult to quantify individual responses to the forcing function.

3

Control-Theoretic Analysis of Lane Keeping

Many studies have aimed at describing the human behaviour of lane keeping in a fundamental way, by constructing models that represent the human dynamics from input (the road geometry) to output (the position on the road). This chapter describes the step to representative models of car driving and discusses various of such constructed models from literature. The goal of these descriptions is often to describe human behaviour accurately, but ideally also in the most intuitive manner as possible, such as with parameter descriptions that match physical interpretations of lane keeping.

3.1. The Road as Control Task

This section describes how car driving can be described as a control task. First, the definition of the human control task for conventional car lane keeping is given, after which it is described how the geometry of the road determines the task definition.

3.1.1. The Control Task

In the simplest and most common practical form, drivers are given the assignment of following a road, alternating between straight and curved segments. Its pattern is clearly distinct from the current environment (such as on highways), indicated by white lines. By assuming the driver uses the shape of the road as its main visual cue, the road can be seen as an input signal that requires constant control action. Here a fundamental difference arises with the preview tracking explained in Chapter 2, as it is not necessarily clear what type of control strategy drivers use, as its control task is not clear: Drivers are free to position the vehicle anywhere on the road, as long as it roughly remains on the road and within the lines.

As Donges [4] noted, the centreline can be used as the forcing function if a driver actively tries to keep the vehicle in the middle of the road and the centreline is defined as the middle of the boundaries, but often centrelines are not visible on roads. Therefore car driving is a boundary avoidance task, known as *lane keeping*.

This freedom given to drivers allows for a high flexibility in driving strategies. Drivers can choose to aim for perfect tracking of the road, or be more selective in the components of the road that they respond to. For example, neglecting the higher frequency parts of the road can result in corner cutting behaviour. This strategy might result in a worse

performance compared to someone who aims for perfect tracking, but the lower control activity results in a more laid-back driving style.

3.1.2. Road Geometry

With the knowledge that drivers directly use the road as the target signal ahead of them, lane keeping is an extended form of preview control, in which perspective viewing must be taken into account. Figure 3.1(b) shows the same task as in Figure 3.1(a) at a perspective view, which is the view corresponding to the visual scene present in lane keeping. The perspective view changes the way how the underlying frequencies are displayed. Small differences are more difficult to perceive at larger distances, although these parts become enhanced once they get closer to the driver. This difference between perspective and top-down viewing thus affects driver behaviour, although changes in terms of model validity were found to be small (Van Der El et al. [27]), meaning that the method humans use in handling the lane keeping remains similar.

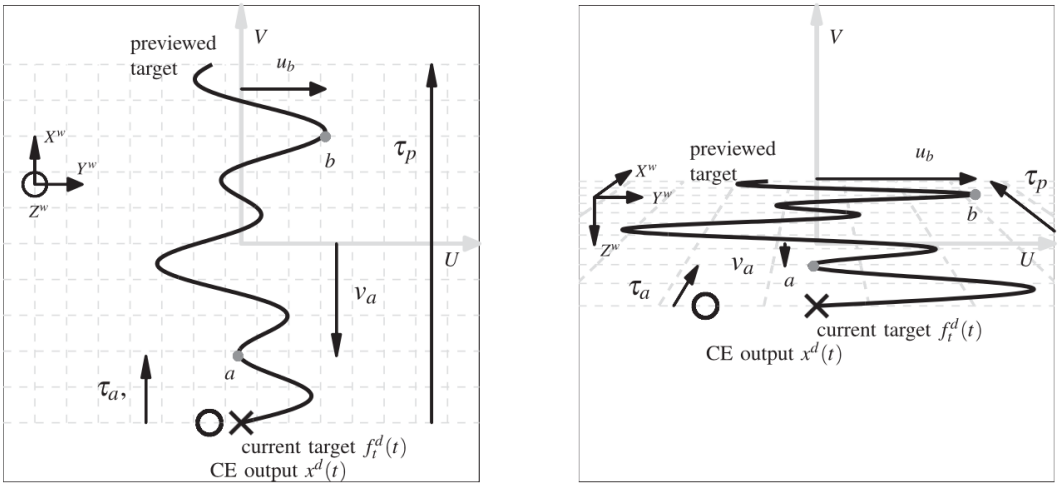

Figure 3.1: Topdown (a) and perspective (b) views of the road tracking task, adapted from Van der El et al. [27].

3.2. Metrics of Stabilization

A conceptualization of lane keeping behaviour requires the description of parameters used by drivers to stay within the lines. Donges [4] described this as *stabilization information*, meaning that drivers use these metrics to stabilize themselves to the road. As the road shape can change constantly, these metrics are functions of distance. Assuming the width of the road remains constant, the state of the vehicle can be described by the three main metrics:

1. The lateral position deviation (denoted y_e), which is the difference between the road centreline and the vehicle position, perpendicular to the local tangent of the centreline. This metric is measured from the center of gravity, meaning that is typically not visible from within the vehicle itself. It is therefore not clear to

which prediction of the lateral position driver respond. The case where the lateral position deviation is nonzero is shown in Figure 3.3(b).

2. The heading deviation (denoted ψ_e), which is the difference between the commanded heading ψ_c of the road and the actual heading ψ. An example in which this case is nonzero is shown in Figure 3.3(c). Again it should be noted that there is no clear definition to what heading humans respond, as they might use any point along a winding road as their commanded heading. Without the lateral position or heading angle errors, the case shown in Figure 3.3(a) is present.

3. The curvature deviation (denoted r_e), which is the difference between the road curvature r_c and the vehicle curvature r. Consider the situation in Figure 3.2. If a driver would directly steer towards an arbitrary point at heading ψ_c, it would still need to drive along a winded, and not straight, part of the road, to match the commanded curvature, meaning that the curvature must be included as a separate metric. However, it is possible that this curvature matching can be fully explained by taking the curvature as a range of commanded headings.

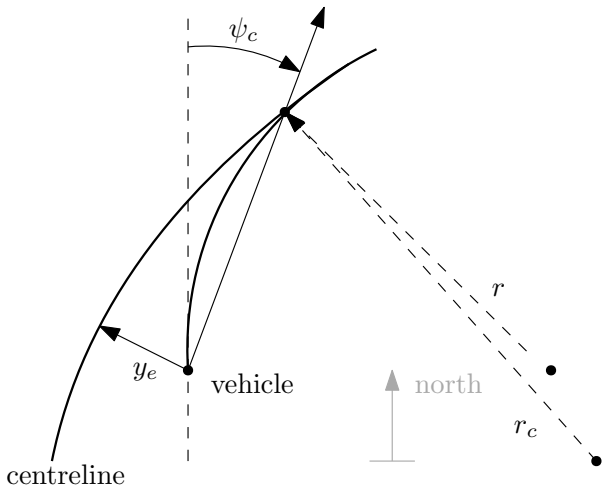

Figure 3.2: Overview of visual metrics determining the state of the vehicle with respect to the road.

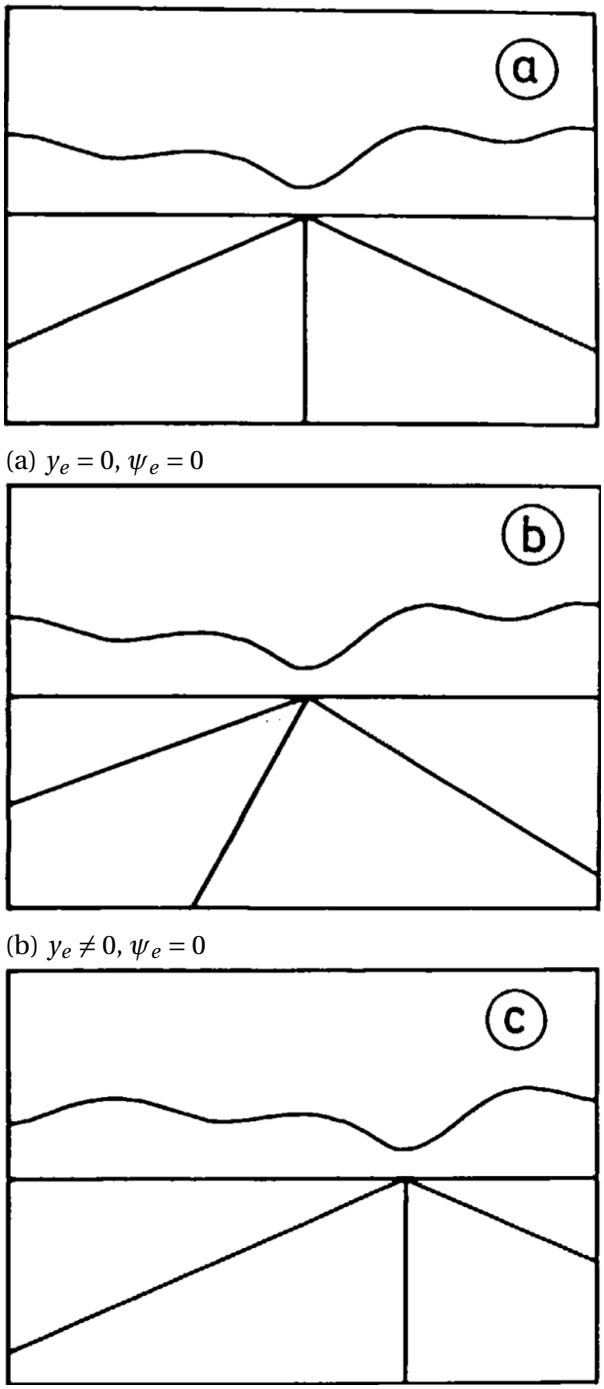

Figure 3.3: Examples of zero errors (a), non-zero lateral position error (b) and non-zero heading error (c) lane keeping, adapted from Donges [4].

3.3. General Model Structures

An important question that arises is: How do drivers perceive these stabilization metrics, and how do they use these for lane keeping? The lateral position describes how far the vehicle is located from the center of the road, which is the main metric to be minimized for drivers, although it cannot be observed directly. In car control, the lateral position

is measured from the center of gravity which is located behind the driver, meaning that drivers have to interpret the lateral position from the visuals in front of the vehicle. As this information can theoretically be perceived from any part of the road, its location might vary per driver. The heading angle has a similar problem. It indicates the relative attitude of the vehicle with respect to the shape of the road, which is directly observable through the front window. However, it still needs to be extracted from a point on the road.

Over the years various models have been developed, of which most assume that drivers obtain these metrics of stabilization from one or multiple fixed points on the road. The position of the lateral position and heading angle viewpoints do not necessarily have to coincide. A comprehensive overview was given by Steen et al. [23]. Using control theory optimal control models can be derived that describe human control accuractely (such as Sheridan [21] or MacAdam [12]), although these models often lack physical interpretations of the inputs and thus provide little information on what guides control. The most accurate predictors of human control that still use some form of physical interpretations of the driver are the *two-point models*, in which the control is described by a combinational use of two points. Donges [4] similarly described driver behaviour in two stages: An open loop anticipatory control of the curvature of the road, as well as a closed loop compensatory control loop of the lateral position, in which drivers minimize their lateral position deviation by forming a closed loop. The model structure proposed by Donges is shown in Figure 3.4.

Although two-point models as proposed by Donges, or similarly as proposed by Sentouh et al. [20] often provide relatively accurate description of human control, it is unclear what the physical interpretation should be of these two points. Humans cannot focus on more than one point at the same time, but might spread their attention or simply perceive the information through their peripheral/parafoveal or even focal view. To partially solve this problem, Salvucci & Gray [19] thought of a combined 'visual angle' to perceive the information required for the two-point model. However, all of these models are only confirmed in a performance sense, meaning that until now there have been no actual measurements of the individual response dynamics, such as to the lateral position or heading angle. This requires new interpretations of how humans actually perceive and process the visual information in lane keeping tasks.

3.4. Van der El model

A different explanation on the handling of road geometry was recently described by Van der El [25], in which driver behaviour is described similarly as in the previously mentioned combined visual angle: by describing the negation of a single perceived viewing error.

3.4.1. Model Geometry

Consider the situation in Figure 3.6. Van der El stated that the driving data corresponded to the response on a single aim point A, located at some $(t + T_{la})$ ahead. This T_{la} is the *look-ahead time*, a directly interpretable metric on how far drivers look when driving down a winding road. The reduction of this perceived error is achieved by a *compensatory* fashion, referring to the traditional findings of McRuer & Jex [13]. This

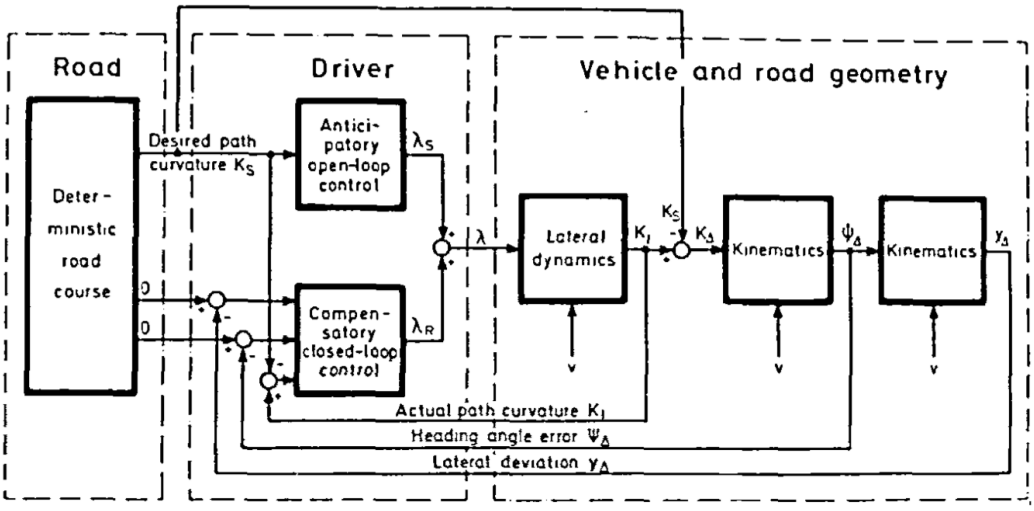

Figure 3.4: Two-level model structure, adapted from Donges [4].

perceived error, located at angle η from the vehicle heading, is in a sense a *combination* of the metrics described in Section 3.2; it provides drivers with information on the lateral position, heading and curvature deviations. From the geometry in Figure 3.5, this angle $\eta(t, T_{la})$ equals:

$$\eta(t, T_{la}) = \arcsin\left(\frac{y_c(t+T_{la}) - y(t)}{T_{la}U_0}\right) - \psi(t) \approx \frac{y_c(t+T_{la}) - y(t)}{T_{la}U_0} - \psi(t) \tag{3.1}$$

For small angles. Here, T_{la} is the look-ahead time and U_0 is the velocity. If T_{la} becomes large (i.e. approaches the vanishing point), this angle approaches the heading angle error:

$$\eta(t, T_{la}) = -\psi(t) \tag{3.2}$$

This also confirms the notion that $\eta(t, T_{la})$ contains information on the lateral position as well as the heading angle, whose weighting depends on the location of the aim point: $1/(T_{la}U_0)$.

3.4.2. Control Model

The metric $\eta(t, T_{la})$ can be used by drivers for control, to which Van der El hypothesized a model description as response to the heading and lateral position information, shown in Figure 3.7. The precision model used consists of a gain-lead structure:

$$H_o^{comp}(j\omega) = K_{e^*}(1 + T_{L,e^*}\, j\omega)e^{-\tau_v j\omega}\frac{\omega_{nms}^2}{(j\omega)^2 + 2\zeta_{nms}\omega_{nms}j\omega + \omega_{nms}^2} \tag{3.3}$$

With τ_v the visual response delay, ω_{nms} and ζ_{nms} are the parameters of the neuromuscular system. Note that the parameter K_{e^*} and the lead term $(1 + T_{L,e^*})$ already correspond to the following of a smoothed, low-pass filtered signal, slightly deviating from the original road input signal and leading to corner cutting behaviour:

$$H_{o_f} = K_f \frac{1}{1 + T_{l,f}\, j\omega} e^{\tau_f j\omega} \tag{3.4}$$

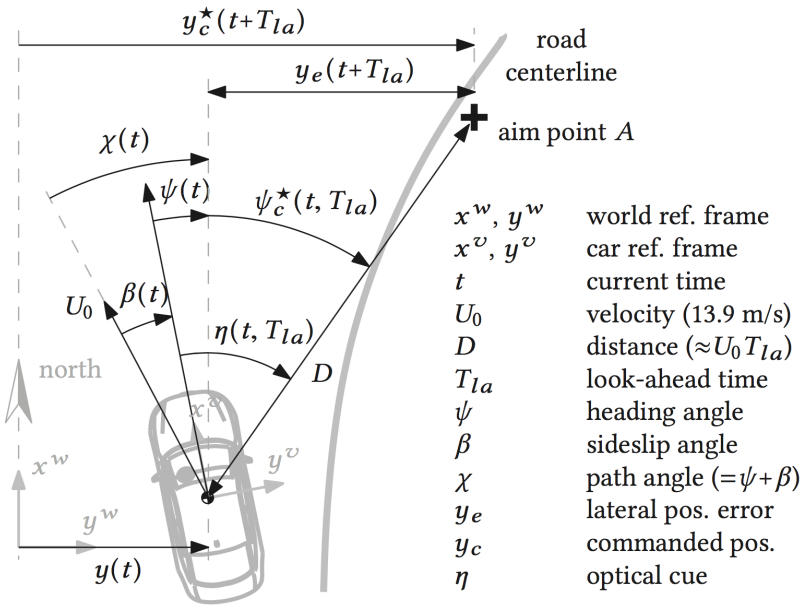

Figure 3.5: Top-down overview of the vehicle state geometry, adapted from Van der El et al. [25].

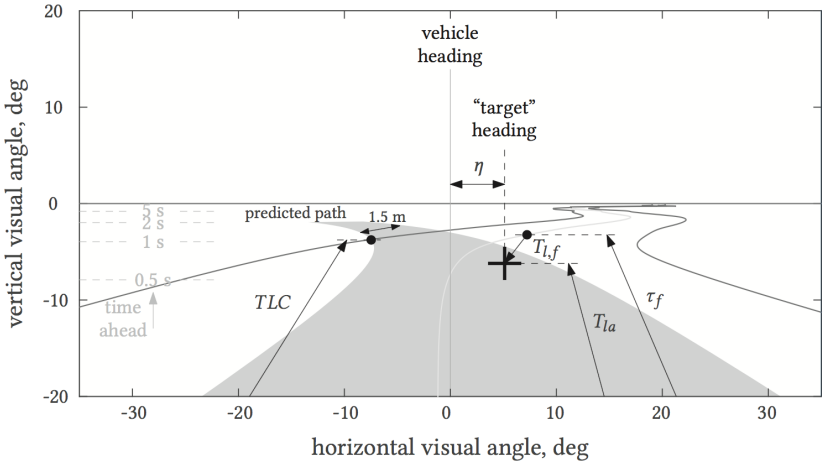

Figure 3.6: Isometric overview of the vehicle state geometry, adapted from Van der El [25].

This prefiltered signal is thus a low pass filter with time constant $T_{l,f}$. As shown in Figure 3.7, the complete dynamics between the input and the control output equal:

$$e^*(j\omega) = \left(K_f \frac{1}{1 + T_{l,f}\, j\omega} e^{\tau_f j\omega} Y_c(j\omega) - Y(j\omega) \right) K_y^\psi - \psi(j\omega) \tag{3.5}$$

Which is subsequently the input for the precision model of Equation 3.3. The constant τ_f indicates the farthest point on the trajectory used for control. As can be observed, this has a similar structure as in Equation 3.1, in which the constants of the heading angle, lateral position and forcing function responses can be written geometrically:

$$K_y^\psi = \frac{1}{T_{la} U_0} \tag{3.6}$$

And:

$$K_f \frac{e^{\tau_f j\omega}}{1 + T_{l,f} j\omega} = e^{j\omega T_{la}} \tag{3.7}$$

This again confirms the notion that the factor T_{la} determines the weighting between the response to the heading angle and the lateral position. From Figure 3.7, the responses to the individual signal can also be determined:

$$H_{o_\psi}(j\omega) = H_o^{comp}(j\omega) \tag{3.8}$$

$$H_{o_y}(j\omega) = K_y^\psi H_o^{comp}(j\omega) \tag{3.9}$$

$$H_{o_{y_c}}(j\omega) = H_{o_f} K_y^\psi H_o^{comp}(j\omega) \tag{3.10}$$

Where $H_{o_\psi}(j\omega)$, $H_{o_y}(j\omega)$ and $H_{o_{y_c}}(j\omega)$ represent the control of the heading angle, lateral position and target feedforward, respectively.

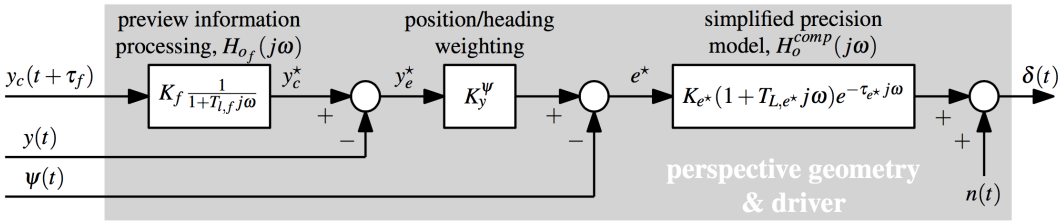

Figure 3.7: Control block diagram, adapted from Van der El [25].

This result by Van der El provides a fundamental quasi-linear control model of a multi-loop phenomenon that forms a greater connection with the visual-perceptual approach (Chapter 4), as the modelled parameters are directly physically interpretable and for the first time gave fundamental insight in the disentangled dynamics of each response.

4

Visual-Perceptual Analysis of Lane Keeping

The study of human control in car driving tasks is incomplete without a description of how information guides the control and what information is selected by drivers. Car driving is a predominantly visual control task (Sivak [22]), meaning that the study of visual perception is important in understanding how stable lane keeping control is achieved. The main focus of this chapter is thus to describe the underlying principles of visual perception and methods to analyze to disentangle this flow of information.
The first two sections discuss the concepts of focal levels and optic flow. The third section gives an overview of gaze attraction models. The last section discusses visual occlusion, which is the partial hindrance of visual information to what visual regions are important for successful driving.

4.1. Types of Focal Vision

The visual information, by an almost discrete distribution of rods and cones (either rods or cones are present at an angle, shown in Figure 4.1) in the eye, is sent to the brain, where the signals of both two eyes are fused into one single image. This causes a view that has decent sharpness and color levels, although humans distinguish between colours best when looking right in front, and minimal differences in intensity are still detected the best at the sides of the perceived image. This fusion of information allows humans to perceive large amounts of visual information at the same time from a wide Field of View (FoV). The 5° area directly around the focal point provides the *fovea* (Wandell [30]), corresponding to the yellow spot that provides the sharpest image. The area directly next to the fovea is the parafovea, which still provides a decent sharpness of an object. For example, with the parafoveal vision it will not be possible to read a word from a book, but one might still be able to determine whether it is a long word, how many words there are and what colours exist. The perifovea compromises the largest area of the complete visual image, as shown in figure 4.2. In this area it is still possible to perceive objects or their movements in space, but at a largely decreased sharpness.

4.2. Characteristics of Optic Flow

Car driving is above all a dynamic task, as the human must constantly respond to what is perceived from its surroundings, such as the road. This change of visual information

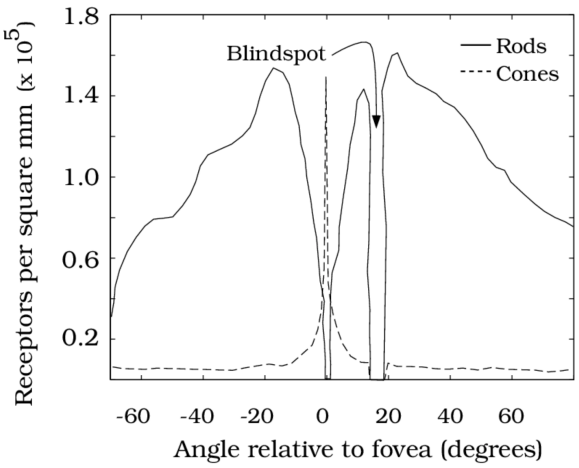

Figure 4.1: Schematic overview of the cone and rode distribution on the retina, adapted from Wandell [30].

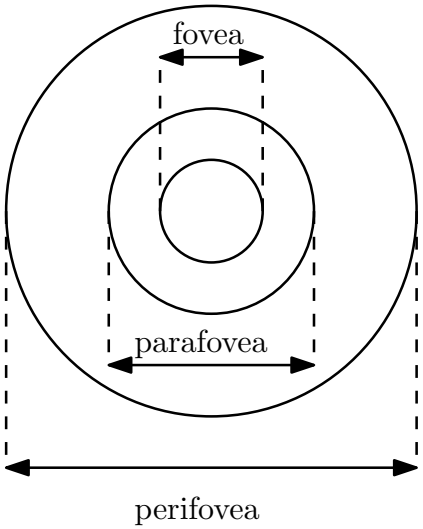

Figure 4.2: Schematic overview of the focal types from Wandell[30].

is determined by the optic flow, which is the perception of a *moving* object, or in the case of car driving, the movement of surroundings. Optic flow can be categorized by two main metrics (Gibson [6]): The amount mainly stems from the physical properties of the objects. A stone road is more grainy and thus provides more flow compared to a perfectly flat road. Secondly, the relative direction or pattern of the optic flow expands from the focus of expansion (FoE) and provides drivers with the perception of movement, affecting the drivers' sense of direction.
This sense of direction can be used to extract heading information (See Warren [34] or Van den Berg [24]), although it is not agreed upon how accurate this method is for drivers. Wann et al. [31] as well as Wann & Land [33] developed the idea of optic flow being the actual guidance for heading control, as shown in Figure 4.3. Wilkie & Wann [36] stated that humans adapt their heading in such a way that the resulting optic flow received by the drivers approaches in a straight fashion, indicating that the driver locomotes directly towards its intended target. Humans can judge the heading as well as path information with an accuracy of ~3° for straight roads, but on curved section this

judgement dropped to ~5° for the path and to ~13° for the identification of heading. This study also noted that it is difficult to quantify how large the changes in optic flow are when imperfections with respect to the intended model arise, such as when drivers are distracted and move their head. These components of optic flow cannot just be extracted as separate translations or rotations. This free gaze behaviour affects driver control behaviour to an extent that it cannot be ignored as potential explanation on differences found between driving experiments (Mole et al. [16]).

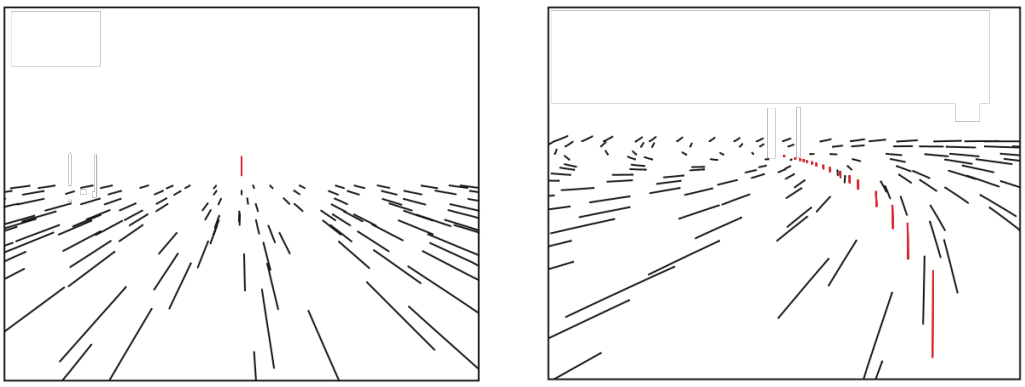

(a) Straight road. (b) Curved road.

Figure 4.3: Optic flow patterns in straight and curved road segments, adapted from Wann & Land [31].

4.3. Gaze Attraction Models

The most direct method to obtain driver data on the use of visual information is eye tracking, by measuring the gaze of the driver. Using this method one can see where drivers look when steering. Although these methods often yield a single gaze point, drivers can also perceive information around that point through their peripheral or parafoveal vision, first noted by Mourant & Rockwell [17], which is something not captured by measuring a single gaze point. Although these vision types do not yield vision with the same sharpness as focal vision, shapes and colours can still be seen and can provide a useful source of information.
The time distribution of the eye gaze yields a direct spatial distribution such as shown in Figure 4.4. Similar to the theories of heading extraction from optic flow, Land & Lee [10] as well as Wilkie & Swapp [32] noted that when drivers approached bends, their focus would initially be placed on the tangent point of the inner curvature line, shown by the location of the gaze distribution in the same figure. The geometry links the curvature of the road segments to the viewing distance from according to:

$$C = \frac{1}{R} = \frac{1}{d\cos\theta} - \frac{1}{d} \approx \frac{\theta^2}{2d} \tag{4.1}$$

Where C and R are the curvature and the radius of the corner respectively, d is the lateral distance to the tangent point and θ is the relative angle from the driver to the tangent point (as shown in Figure 4.5). Although this behaviour occurs for the participants in those studies, drivers often only applied this method for a short period of time; only

when approaching a bend. Therefore it seems that such models only provide a part of understanding human control. Salvucci & Gray [19] implemented these ideas within the framework of a two-point control model and even managed to account for lane changing.

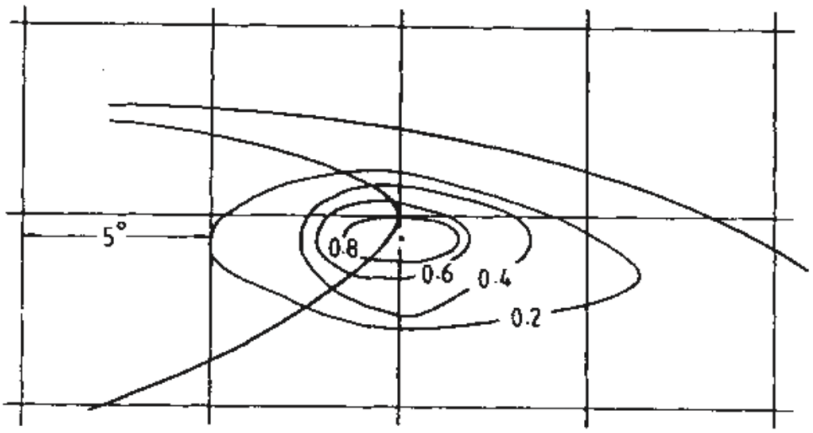

Figure 4.4: Statistical occurrences of gaze direction, adapted from Land & Lee [10].

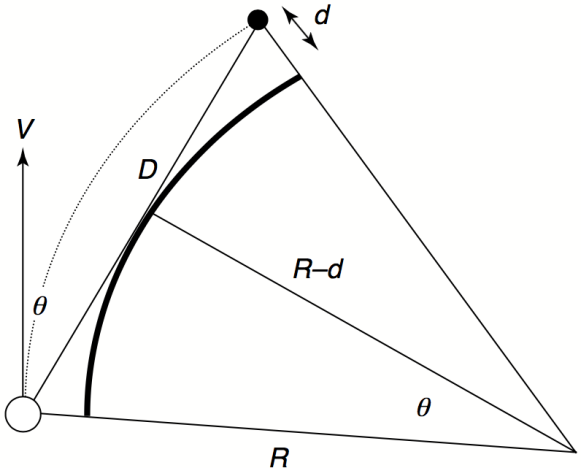

Figure 4.5: Geometry of the tangent-point attractor model, adapted from Wann & Land [31].

4.4. Visual Occlusion

Considering that the previous sections have formed a short introduction on the use of visual perception, this section describes how the selection (or *occlusion*) of the visuals presented to the drivers can be used to determine how humans use this perception in lane keeping. The method of visual occlusion relies on the separation of visual information to identify the importance of spatial regions of the road on driver behaviour. Only specific segments of the road are presented to the driver, whereas all other visuals are hidden. It is notable that often these experiments have formed the empirical evidence of one- or two point control theoretic models, such as from Donges [4].
The most used method for occlusion is longitudinal segmental visual occlusion, in which the longitudinal position of the slits is varied for each condition. By varying the position

of this overlay, and thus the displayed longitudinal segment of the road, for example the effect on the lateral position deviation over time can be measured.

4.4.1. Study by Land & Horwood

Land & Horwood [8] were one of the first to perform a simple occlusion experiment for lane keeping tasks, partially motivated by the findings of Donges [4]. In this experiment a lane keeping task was simulated with simple visuals, where white lines indicated the boundaries of the road, all else was black. There was no center line present and drivers drove down a fixed-width (3m) winding road[1], while keeping the vehicle between the lines. No other scenery was displayed. Four types of measurement conditions were investigated, with the available visual information as independent variable.

Measurement Setup

First the full visuals were displayed to the driver, with no occlusion at all. This is the closest approximation to the normal car driving situation and serves as the baseline of the driver behaviour. For the other three conditions the occlusion was presented. Small *apertures* of 1° in height, but at full width, formed the only way of looking at the road, all other points on the visuals were occluded by a black overlay. As if one is watching through a mailbox, the underlying visuals remain exactly the same, meaning that the same experiment is performed simply with a decreased amount of visual information, as shown in Figure 4.6(a-b). By varying the vertical position of these slits below the horizon, the researchers aimed at describing how important each of these vertically spaced areas are for control.
There were nine possible positions for the segments (Figure 4.7), spanning from the horizon to ten degrees down, where the field of view always remained constant. Geometrically speaking, the true horizon is the horizontal line through the convergence point of the road side lines, known as the vanishing point. In reality, a perfect sight towards the vanishing point is unlikely and that might have been the reason for Land & Horwood to locate the highest opened slit at 1° below this true horizon, such as shown in Figure 4.7. Although this means that the effect of an opened slit between 0° and 1° was not investigated, this area *was* included in the baseline visuals.

Independent Measures

After the baseline (full visuals) was complete, the second type of condition showed only one of these slits, with its middle point at varying position between 1.5° and 9.5°. As the third step the farthest aperture at 1.5° was visible but fixed over runs, with a second aperture at varying angle per condition, known as 'top fixed' and vice versa for the fourth step, where the bottom aperture was always at a fixed position, known as 'bottom fixed'. These last two types of conditions thus have two apertures present, providing drivers with a mixed selection of longitudinal visual information. Sometimes the slits were located relatively high on the screen, such as in a top-fixed situation with a second aperture at 3.5°, or vice versa. There are also conditions where their distribution is more evenly

[1] The track used by [8] is based on a model of Queens Drive in Edinburgh, United Kingdom, which is a narrow road around Arthur's Seat in Holyrood Park.

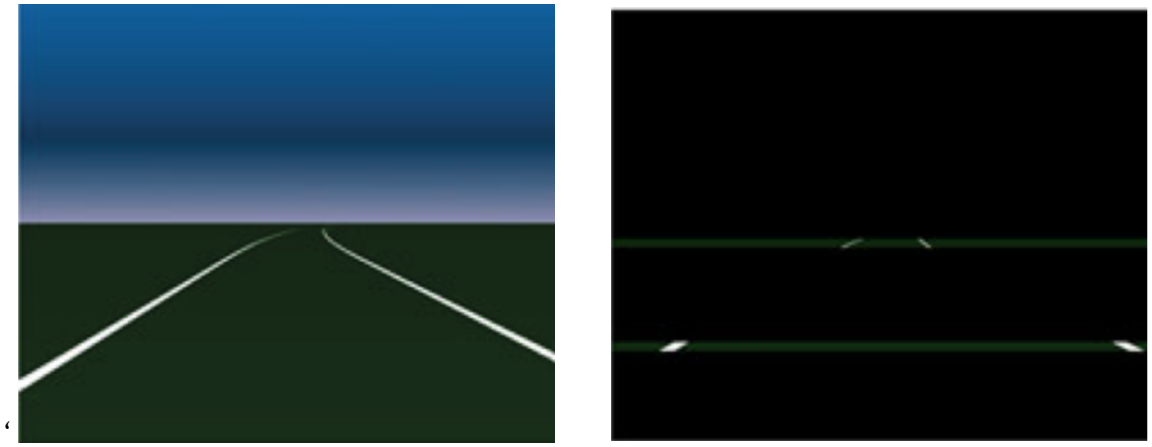

(a) (b)

Figure 4.6: Road segmentation example with non-occluded (a) and occluded (b) visuals, reproduced from Cloete & Wallis [3].

spread; one bar at the top and one at the bottom.

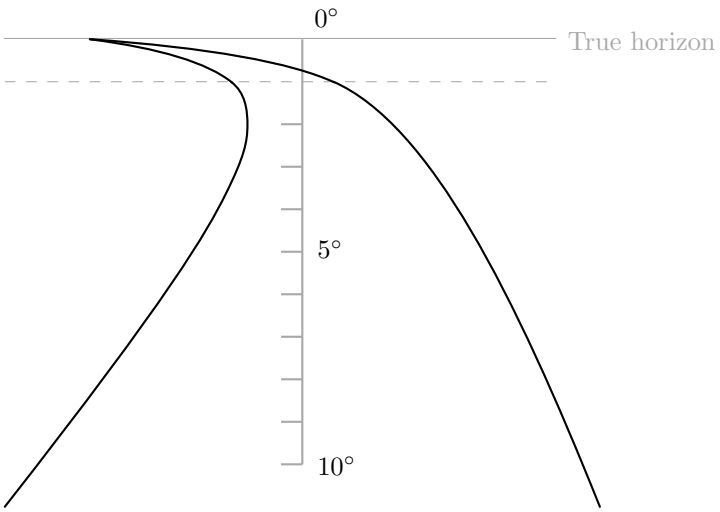

Figure 4.7: Aperture road segment division, as function of longitudinal angle, reproduced from to Land & Horwood [8].

Control Variables

The car always drove at a fixed velocity of 16.9ms^{-1} (61kmh^{-1}), whereas a second study (Land & Horwood [9]) investigated the same conditions at 12.5ms^{-1} (45kmh^{-1}), and 19.7ms^{-1} (71kmh^{-1}). The visuals were presented on a 60cm wide screen at a distance of 80cm and the refresh rate of the display was 7Hz. Drivers aimed their sight from a position of $l_h = 1.1\text{m}$ above the road. A double integrator setup was used as vehicle model, although its integration constants are unknown.

Dependent Measures

Land & Horwood measured and compared the standard deviation of the cars lateral position error between the various conditions for three different drivers. The accuracy ratio $\frac{\sigma_{bl}}{\sigma}$ was calculated as the standard deviation of the baseline divided by the standard deviation of the selected condition; it is the reciprocal of standard deviation with respect to the baseline, hence a value of one indicates an equal performance compared to the baseline conditions, but anything lower indicates a worse performance. It is possible to achieve a better performance than the baseline ($\frac{\sigma_{bl}}{\sigma} > 1$), but one might expect that reducing the visual information almost always leads to a worse performance. The results of the original experiment can be seen in Figure 4.8(a-c).

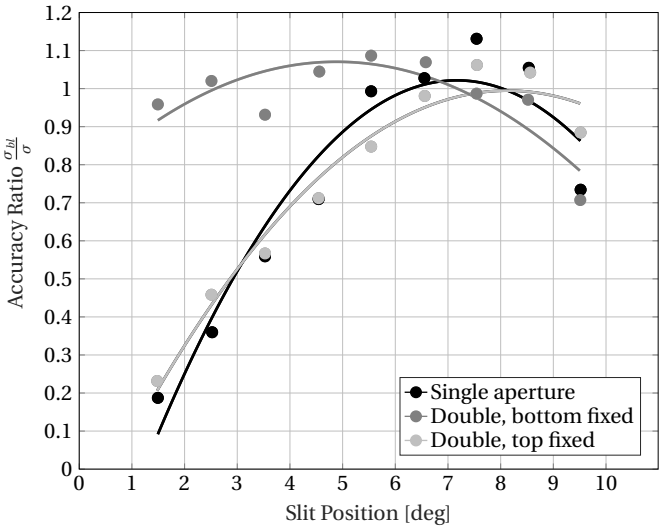

(a) v = 12.5ms^{-1}

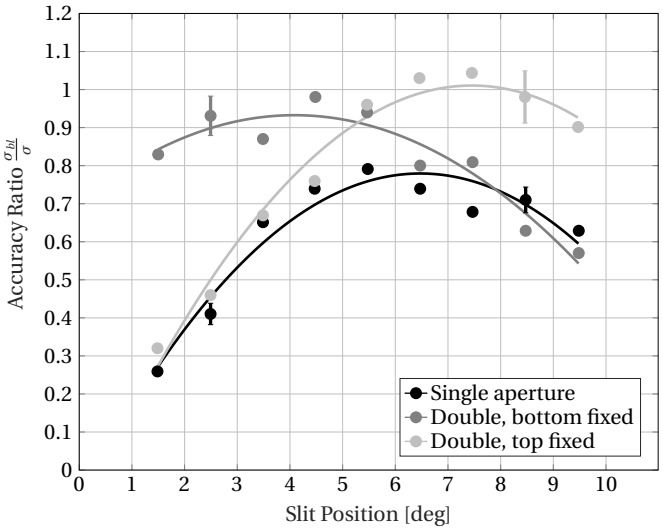

(b) v = 16.9ms^{-1}

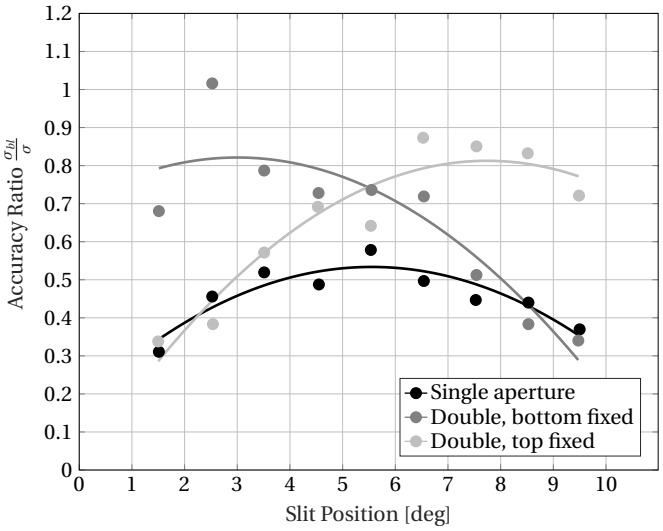

(c) v = 19.7ms^{-1}

Figure 4.8: Accuracy ratios as function of longitudinal aperture angle for three different velocities, adapted from Land & Horwood [9].

4.4.2. Discussion of Results

For a fixed velocity of v = 16.9ms^{-1}, shown in Figure 4.8(b), presenting only one aperture decreases the driving performance with respect to the baseline at all angles, although it is at best in the middle of the road ($5° - 6°$). However, if the nearest aperture is shown and with that a second aperture as function of position (bottom fixed), the accuracy ratio increases, with an optimum somewhere in the middle aperture ($3° - 5°$). A similar observation is made for the reverse case; showing the farthest aperture and varying the other apertures one at a time (open circles) yields a decent performance and is in two cases even better than viewing the full road. Here the optimum is somewhere close to the bottom of the display ($6° - 9°$). Hence, these results show that a two-segmental occlusion scenario can describe the same level of performance as the baseline condition, for some angles.

At a higher velocity (v = 19.7ms^{-1}), presented in Figure 4.8(c), the same effects can be observed. The problem of having only one aperture is enhanced and degrades the performance ratio to a maximum of 0.6, whereas two apertures also do not provide drivers with a similar performance compared to the baseline. However, it should be noted that the position of the maxima still occur at roughly the same angles.
For the lowest velocity setting of v = 12.5ms^{-1}, shown in Figure 4.8(a), drivers seem to be able to obtain enough information from a single point, being almost identical to the top fixed variant, although both have distinct optima. The bottom fixed variant has, with the exception of the lowest setting, an almost horizontal distribution in the performance ratio. Land & Horwood explain this by noting that at lower speeds, it is possible to extract feedback from the road edges from the nearest slit, thus decreasing the need of a second slit. This strategy becomes increasingly unreliable when the velocity increases, because this leads to unstable driving.
Although not directly visible in these figures, it was also noted that the position of these slits determines the strategy that humans apply to keep within the lanes. Only showing the nearest slit was characterized by a jittery control output, whereas a high placed slit caused drivers to apply smooth steering. It is not necessarily true that leads to a difference in performance; different control strategies can work equally well.

The conclusion made is that for the middle velocity (v = 16.9ms^{-1}) drivers are able to successfully steer cars and approximate the baseline with limited information available, such that these two conditions become almost indistinguishable from a performance perspective. This is as long as this information comes from two separated visual points, with a separation optimum of roughly $5°$.

4.4.3. Comments regarding validity

The findings of Land & Horwood are important in the sense that for many years they have assumed to be the basic verification of two-point models. Quite crucially, however, some important remarks must be made about the *execution* of the experiment and its subsequent claims.

1. **Low display refresh rate** - The refresh rate of the display was set at 7Hz, which results in a jittery simulation of an otherwise continuous process. On every

non-simulated frame the eye perceives a lack of new information, leading to a difference between the observed frame of the simulation and the expected visual information based on experiences of car driving. On every 'lacking' frame, where drivers cannot perceive visual feedback, this effectively results in a delay. The way in which humans change their behaviour depends on the magnitude of this delay (Miall et al. [14]), meaning that the low refresh rate can have a profound influence on the behaviour of the driver. Cloete & Wallis [3] describe the way in which types of control can occur by three stages: smooth, step-and-hold (intermittent corrections) and 'bang-bang', where the latter indicates a control strategy where drivers discretely vary their input between two (near) opposite values. Land & Horwood [8] also describe the the occurrence of this last type of control, which might be partially explained by this low refresh rate, although as mentioned, their explanation is purely based on the interpretation of the control task.

2. **Lack of optic flow** - The presence of optic flow (such as from the road itself or surrounding scenery) might have an effect on the performance of the driver, mainly in the areas where the flow is the heaviest. The aperture close to the vehicle, having the largest longitudinal angle, show more optic flow than the parts farther away. The 'aggressiveness' of these apertures will be higher, and driver might be more motivated to also attain a more aggressive control strategy or to increase the control activity. Drivers can also be more selective in the areas they look at to decrease the 'heavy' optic flow meeting their eyes. Finally, as explained in Section , flow can be important in the extraction of the heading angle. As a comparison is made to real-world driving, where optic flow is also present, it would be better to implement an optic flow component in the experiment as well.

Furthermore, there are also remarks concerning the *analysis* of the obtained data.

1. **Small data set** - Due to the small amount of subjects (N = 3), it is questionable if the data holds any statistical significance, also due to the lack of (consistent use of) error bars. Although the main trends were the same for all three drivers, it is more difficult to argue whether the collected data were outliers and thus how convincing the results are.

2. **No representative population** - The authors participated in the experiment themselves, providing them with unwanted information on control strategies leading to a confound in the experiment. By designing the experiment it is likely they were subject to training effects. Most importantly, they were aware of the hypothesis and could thus adapt their driving strategy to this information. Preferably, one would look for subjects unaware of this type of research and its measuring methods.

3. **Inconsistent data** - In two special cases of the top- and bottom fixed configurations, there are two apertures overlapping each other: The top fixed exactly overlaps second aperture at 1.5°, whereas the bottom fixed configuration with a second aperture at 9.5° results in the same problem. These are indistinguishable from their single aperture counterparts, and, if tested both, should thus lead to the same driver performance. However, as can be seen in Figures 4.8 (a-c), there

are discrepancies between these conditions. The call for error bars in these figures is invigorated by this observation.

Comments about the *interpretation* will be discussed in section 4.4.7.

4.4.4. Study by Cloete & Wallis

As mentioned, the refresh rate used (7Hz) in the original experiment might have a profound effect on the measured performance. In that case drivers could apply a different driving strategy, for example by using more aggressive steering or acquire more lead to compensate for the 'delayed' visuals. In that case, the presence of occlusion will investigate the wrong control task due to the lower refresh rate. It is possible that this effect is also measurable in the performance ratio's calculated by Land & Horwood, although this is not necessarily the case. Drivers can perform equally well while controlling a different task.

The main work by Cloete & Wallis [3] aimed at the partial reproduction of Land

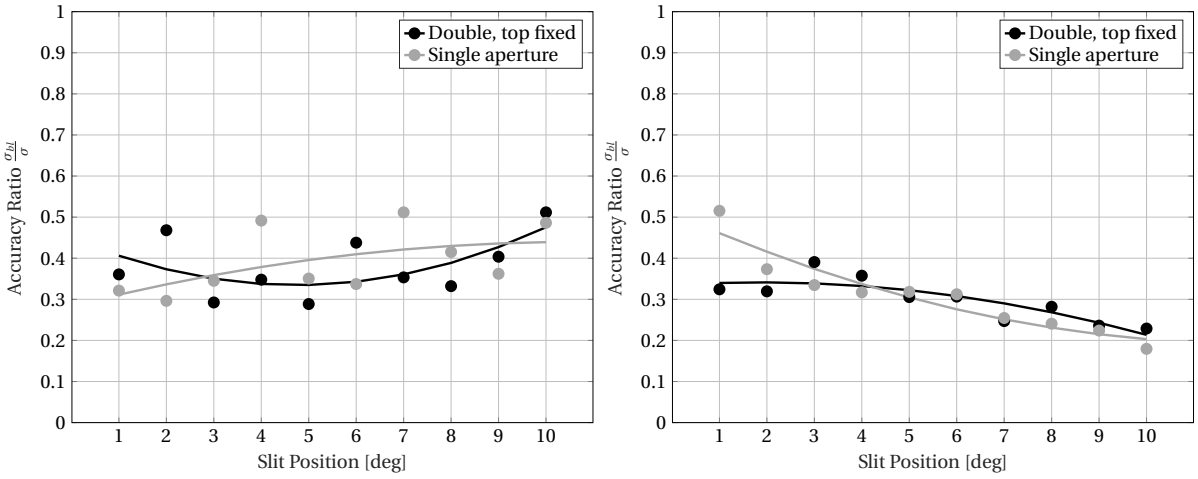

(a) $f_r = 72$Hz (b) $f_r = 7.2$Hz

Figure 4.9: Accuracy ratios as function of longitudinal aperture angle for three different velocities, from Cloete & Wallis [3].

& Horwood's experiment, by comparing between simulations with a low (7.2Hz) and high refresh rate (72Hz). As mentioned before, the 'stepped', jittery sequence inputs of the former, due to the low refresh rate, can give rise to a summation of individual responses by human controllers (Miall et al. [15]), such that they lead to an oscillatory overcompensating response, instead of a smooth and continuous response to the information nearby. This can be an explanation for the jerky response as measured by Land & Horwood. In the experiment, where the eye level was also set at 1.1m, drivers looked at a 3.4m by 2.7m projector screen. It was confirmed by Cloete & Wallis that for higher refresh rates (77Hz) the relative performance drops significantly, as they found a significantly higher standard deviation of the lateral error with respect to the baseline performance, for all aperture positions.

Again, some concerns about this experiment must be noted:

- The simplistic visuals with white-on-green road presentation (Figure 4.6) were comparable to the set-up of Land & Horwood, but also partially lack optic flow.

- Trials were only set at an 8s duration, possibly inducing relatively large transient response measurements.

4.4.5. Study by Chatziastros et al.

Finally, a more detailed study into individual effects was performed by Chatziastros et al. [2]. In this study, three different comparisons were made, for an otherwise highly similar measurement setup as Land & Horwood:

1. **Replication** - Similarly to Cloete & Wallis, a replication experiment was performed on the slit experiment by Land & Horwood. Chatziastros et al. did not measure the dependency on a second far-away point, meaning that this information appears to add nothing to the performance of the driver and that the conclusions of Land & Horwood could not be replicated.

2. **Texture type** - Both of the previously mentioned experiments were based on a simplified representation of car driving. Roads were displayed as lines on a single-colour background. This comparison directly measures the difference in performance if the road also contains textures, increasing the amount of optic flow to the driver. As shown in Figure 4.10(b), this yields a constant offset for all single slit positions, resulting in a slightly better performance if textures are present. It is not clear if ground or grass textures were also included.

3. **Display type** - Finally, a comparison was made between the type of display presented to the driver (Figure 4.10(c)). Relatively large differences in lateral deviation were found between a projector and a monitor, at the same field of view. An explanation given by the authors, is that the presence of the monitor edges provides drivers with a 'stronger' reference frame, corresponding to the optical invariant inner structure of a vehicle. Interestingly, a comparison of the texture types was also made for both display types, and the positive effect on the lateral deviation, caused by the additional textures, could not be found. An explanation for this is the decrease in resolution in the case of the projector screen, as it lowers the perceived optic flow by the drivers.

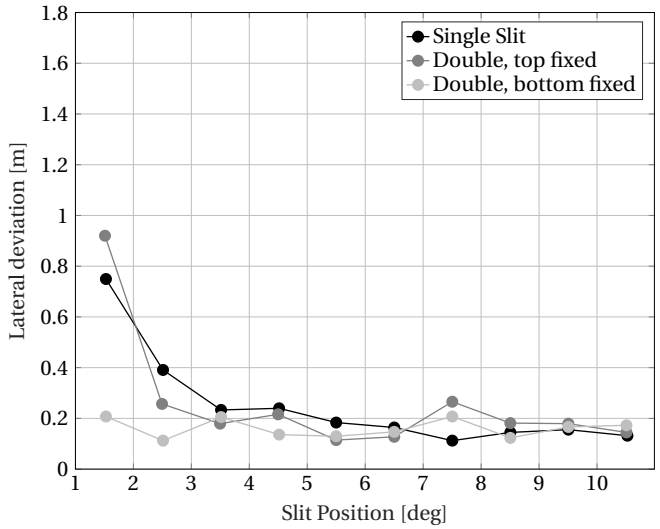

(a) Lateral deviation for different slit configurations

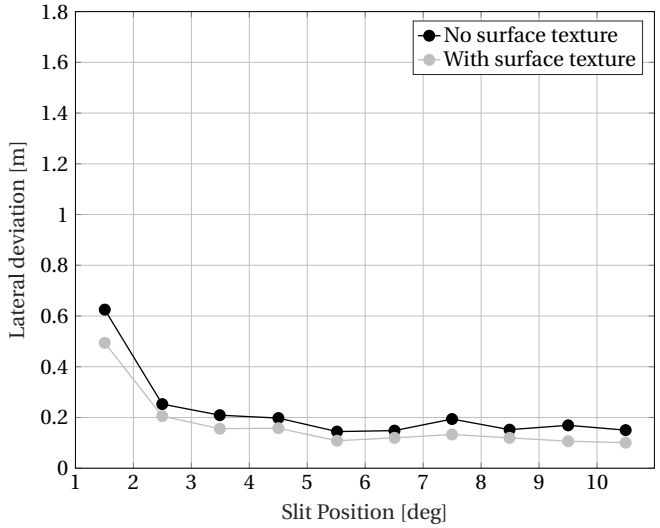

(b) Lateral deviation for different surface textures

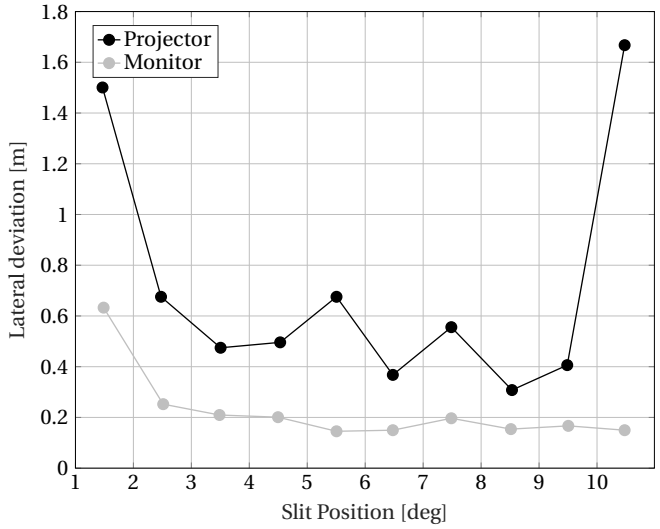

(c) Lateral deviation for display types

Figure 4.10: Lateral deviation measurements between various setup types, adapted from Chatziastros et al. [2].

4.4.6. Varying Transparancy as Occlusion

Segmental visual occlusion is not the only method to measure the differences in human steering capacity. Other types of occlusion can be performed, such as by Frissen & Mars [5], who performed an experiment in which not the position of the available visuals were varied, but its transparancy. The authors reported that up to 60% transparency the visuals still provide adequate information for drivers to successfully anticipate their path. Beyond that point, faster actions indicate a higher dependency on direct lateral position control, which seems intuitive. Segmental visual occlusion, the practice described in the studies above, has been a much more popular method due to its simplicity in splitting up the spatial regions of the road, and over the years has been assumed to be a verification method for two-point control models. However, there are some rather important remarks that must be made about the interpretation of this method.

4.4.7. Discussion of Visual Occlusion

Visual occlusion relies on the idea of forcing drivers to use specific visual information for control. In 'real' lane keeping drivers can use the whole road visuals at their discretion, meaning that the location of the points of attention used for information extraction may not be constant. Furthermore, drivers can even completely deviate from their task and quickly scan their surroundings, while peripherally perceiving information of the road or confidently assuming that deviations do not become large, given that the driver does not look away for too long. This is a pitfall of visual occlusion, as it only answers the question which parts of the road describe control models, but it does not answer questions about the visual attention span of drivers over time. This is not a fundamental problem for the validity of visual occlusion as a measurement method. Driver models verified by this method typically do not need to predict this type of behaviour. It does mean that not many conclusions can be made about time distribution of where humans look.

First of all it is clear that there is no agreement on the validity of the various occlusion experiments. The studies by Land & Horwood and comparable model simulations (such as Salvucci & Gray[19]) showed that a second slit is required for adequate performance and provides drivers with an almost identical performance compared to the baseline results, but Cloete & Wallis [3] and Chatziastros et al. [2] failed to confirm this. The former measured a worse performance at all angles, whereas the latter also showed a dependency on display type and use of textures. Parts that are located the 'highest' on the screen, corresponding to the farthest part of the road (and to the smallest angle) may be more difficult to perceive if the presented screen is small. Finally, there was also no optic flow present in these experiments (except for the condition with road textures present, by Chatziastros et al.). It would be better to directly include this in the experiment, as this provides the strongest link to real lane keeping. Only showing lines reduces the realism, and more importantly, avoids the matter of compensating for the effects of optic flow afterwards, as the flow itself might affect steering behaviour.

But more importantly, there is a fundamental difference in understanding which information segments *can* be used by drivers in occlusion experiments for an adequate

performance and the information that *actually* describes normal car operations. This is not measurable by comparing occlusion scenario's with the baseline performance. A different method of steering can lead to the same results in terms of the lateral position deviation, and thus the performance scores. In a simulation setting it is also possible that drivers put in more effort or are triggered to constantly look at the apertures, possibly affecting the way humans perform their task. This fixation of the eye gaze can have a profound effect on measured performance, as shown by Wilkie & Wann [35], compared to free gaze. Visual occlusion partially resembles the scenario of tracking gaze, as described by Wilkie & Wann. Drivers can feel inclined to constantly look at the points in a vertical sense, but can use the horizontal space at their discretion.

5

Experiment Proposal

The previous two chapters have attempted to provide an introduction to two main research fields within lane keeping:

- Control-theoretic analysis: How the behaviour of humans in lane keeping tasks can be described using control models.
- Visual-perceptual analysis: How the guidance information for such models perceived from the visual scene can be analyzed.

This chapter aims to explain the current lack of understanding in the link between these two fields and subsequently proposes the experiment. The first Section concerns the interpretation of the Van der El viewing parameters, whereas the second Section explains why an occlusion replication experiment is required. Section 5.3 states the corresponding research questions, followed by the experiment setup in Section 5.4.

5.1. Interpreting the parameters of the Van der El model

In Section 3.4 the stabilization metrics of the Van der El model were discussed. The model is based on the reduction of the aim point angle $\eta(t, T_{la})$, located at T_{la} seconds ahead on the road, the use of the second viewing parameter $T_{l,f}$ provides drivers with the option to smoothen their response (ie. the corner cutting behaviour). Although providing unparalleled insight in the individual responses of the driver (to the heading angle, lateral position and target signals), the model does not explain how these viewing points correspond to the perception by drivers and, one step further, their actual focal viewing points.
The main focus of this research project is therefore to investigate how visual information on these control metrics is perceived, and if this in any way corresponds to previously determined results from literature. The key point here is that such visual perception studies have so far mainly looked at the change in performance measures and were never able to determine the *individual* responses to the stabilization metrics at the same time. This is an important difference, as identical performances do not necessarily entail identical driving strategies.

5.2. Occlusion Replication

Occlusion, the method explained in the previous chapter, can help answer this question by determining from which regions drivers obtain their visual information. By varying the amount of slits that can be used by drivers between one and two, as well as their position, the use of the predicted parameters T_{la} and $T_{l,f}$ can be manipulated, such that drivers are forced to perceive the road information from the non-occluded parts and might place their viewing parameters within the slits. This last statement is not necessarily true, as even in occlusion experiments drivers are still free to place these parameters wherever suits their driving style, even outside the slits, while still perceiving the visual information from through the slits (for example, drivers might extra/interpolate the available information and *predict* the road geometry behind the occlusion. The presence of a single slit can also have a profound effect on the presence of the prefiltering point $T_{l,f}$, as this is forced to be located in the same slit as the aim point at T_{la}, possibly resulting in no filtering at all. This might lead to a higher control activity, as drivers lose the possibility to smoothen the target signal.

As the study by Land & Horwood [8] has been the most cited and accepted occlusion experiment, a replication experiment is preferably close to their setup, while still improving upon its shortcomings stated in the previous Chapter. More importantly, the system identification of the underlying driver responses of the Van der El model require the use of multisines, to be able to determine the dynamics at distinctly measured frequencies. The road by Land & Horwood was based on a real road in Edinburgh, meaning that it is necessary to approximate this signal as a multisine to be able to remain close to the original Land & Horwood setup and still be able to perform system identification and parameter estimation.
For this task, changes to the visual scene of the simulated environment are made, shown in Figure 5.1, meaning that it must be ensured that these affect driver behaviour in an understood manner. The presence of occlusion is likely to change the identified driver responses by definition. As of now there is no proof that a multisine approximation of a real road induces similar driver behaviour. Therefore, as part of the experiment, it must be confirmed that this assumption on the forcing function is valid.
Finally, as explained in the previous chapter, humans can obtain visual information in various ways, such as through focal as well as peripheral or parafoveal vision at the same time. By measuring the eye gaze of drivers during both the full visual and occlusion scenario's, it can be determined whether information is actually perceived through the focal vision and how this corresponds to the model view points.

5.3. Research Questions

Considering the statements of the previous two Sections, the main question to be answered is:

- How does visual information on the stabilization metrics guide the control in lane keeping tasks?

The aim is to answer this broad question by a set of subquestions:

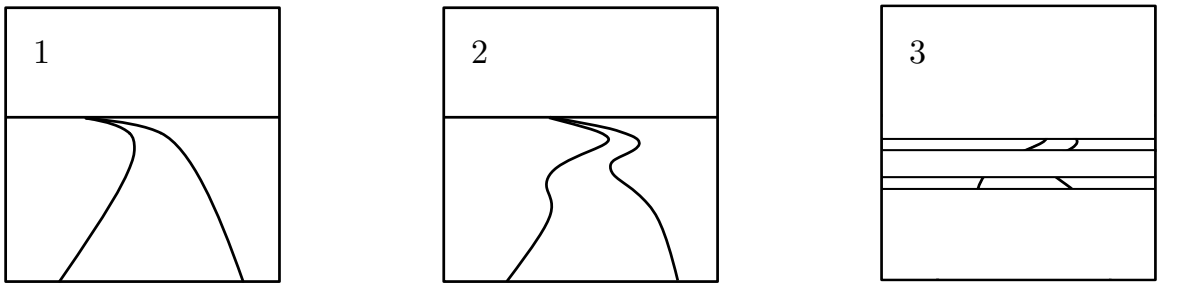

Figure 5.1: The different stages of the visual scene within the proposed experiment.

1. What is the effect of using a multisine approximation of real roads on driver control behaviour? (Measurable by the VAF of the measured wheel rotation δ and its simulated response of the dynamics based on the multisine signal)

2. What is the effect of occlusion on driver behaviour, and is there an occlusion condition that provides the same response compared to the full visuals? (Measurable by comparing the responses H_{o_y}, $H_{o_{yc}}$ and H_{o_ψ} as well as T_{la} and $T_{l,f}$ from the Van der El model between the full visual and occlusion conditions)

3. Does the presence of a single occlusion slit limit the use of prefiltering view point? (Measurable by determining T_{la} and $T_{l,f}$ as well as the control activity for single slit conditions compared to conditions where $T_{l,f}$ is nonzero)

4. To what extent do drivers focus their eye gaze on the view point metrics of the Van der El model? (Measurable by determining the spread of the eye gaze time-distribution compared to the parameter estimations of T_{la} and $T_{l,f}$ and the position of the apertures)

5.4. Experiment Setup

The main driving tasks resembles the setup of Land & Horwood with identical velocity (16.9m/s) and aperture angles, also with the idea of improving on its shortcomings and checking their findings. The experiment includes optic flow due to the road, grass and sky textures. It also contains a fairer and larger selection of participants.
Drivers will repeatedly drive double integrator dynamics, the same as in Van der El et al. [29]. They are first asked to drive five consecutive runs on the original Land & Horwood road, named *Queens Drive*. After that, four runs of the multisine baseline road are taken. In these runs the wind disturbances are also present, which are the same as in Van der El et al. [29]. It is assumed that these nine runs provide drivers with enough 'feeling' of the simulator to perform the occlusion conditions, which follow afterwards. Of all 24 occlusion conditions (shown in Table 5.2) four runs are taken. The order in which these conditions are presented is based on a randomized latin-square design for thirteen participants, shown in Table 5.3.
The last four runs are again the multisine baseline condition. This is to ensure that the effects of exhaustion as well as training can be taken into account, considering the baseline condition is important as a reference to which all other conditions are compared.

Table 5.1: Multisine components of the forcing function and the disturbances in heading angle and lateral position.

	Road center line ψ_c				Heading angle disturbance ψ_d				Lateral position disturbance y_d			
i	k	ω	A	ϕ	k	ω	A	ϕ	k	ω	A	ϕ
-	-	rad/m	deg	rad	-	rad/m	deg	rad	-	rad/m	deg	rad
1	3	0.01	16.2	3.96	7	0.03	2.20	5.04	5	0.02	0.29	5.98
2	9	0.04	15.3	3.17	13	0.06	1.74	6.22	11	0.05	0.24	4.04
3	15	0.07	6.20	4.78	23	0.10	1.08	4.17	19	0.09	0.16	3.03
4	27	0.12	2.34	3.40	35	0.16	0.63	4.40	31	0.14	0.09	6.11
5	39	0.18	2.34	6.28	47	0.21	0.41	4.97	43	0.19	0.06	0.99
6	53	0.24	0.90	6.20	65	0.29	0.25	4.97	59	0.27	0.04	0.11
7	71	0.32	0.05	5.42	85	0.38	0.16	4.10	77	0.35	0.02	1.78
8	93	0.42	0.01	5.71	111	0.50	0.11	5.90	101	0.46	0.02	2.28
9	121	0.55	0.01	0.95	143	0.65	0.08	5.48	131	0.59	0.01	0.41
10	155	0.70	0.01	4.79	183	0.83	0.07	0.73	169	0.76	0.01	2.41

Table 5.2: Overview of the occlusion conditions.

Variable angle [deg]	Single Slit	Double slit, top fixed	Double slit, bottom fixed
1-2	S12	-	-
2-3	S23	T23	B23
3-4	S34	T34	B34
4-5	S45	T45	B45
5-6	S56	T56	B56
6-7	S67	T67	B67
7-8	S78	T78	B78
8-9	S89	T89	B89
9-10	S910	T910	-

Furthermore, every first measurement run of a new condition (first two in the case of Queens Drive) are not taken into account to ensure that only the data is selected in which drivers already have some idea about the task they are about to control.
The experiment takes place in the SIMONA Research Simulator. During the whole experiment, drivers are equipped with a Tobii Pro Glasses 2 eye tracker, which stores both the video recording of the experiment as well as the measured eye gaze of the driver, relative to the head or video. The conversion of the eye gaze from the head-fixed frame of reference to the world frame of reference is performed separately.
Directly related to the original experiment, the standard deviation of the lateral position deviation is measured, together with the control activity. Furthermore, the parameters of the control model are estimated based on the control output. The eye gaze data comes from a different interface and thus must be linked in the time domain to the control data. This data provides the longitudinal position of the eye gaze point in degrees.

Table 5.3: Randomized latin-square for the condition sequence.

Driver	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19	C20	C21	C22	C23	C24
1	S56	T56	S78	S89	S34	B23	S67	T34	T78	S12	T67	B45	B34	B56	S23	B67	T23	S45	B89	B78	T89	T910	S910	T45
2	T56	B56	T78	T89	T34	S23	T67	B34	B78	S910	B67	S45	S34	S56	T23	S67	B23	T45	S89	S78	B89	S12	T910	B45
3	T23	B23	T45	T56	S89	B78	T34	T89	B45	S67	B34	S12	B89	S23	S78	S34	T78	S910	S56	S45	B56	B67	T67	T910
4	B78	S78	S12	S23	B56	T45	B89	S56	S910	B34	S89	T67	T56	T78	B45	T89	S45	B67	B23	T910	T23	T34	S34	S67
5	B23	S23	B45	B56	T89	S78	B34	B89	S45	T67	S34	S910	S89	T23	T78	T34	B78	T910	T56	T45	S56	S67	B67	S12
6	T67	B67	T89	T910	T45	S34	T78	B45	B89	T23	B78	S56	S45	S67	T34	S78	B34	T56	S910	S89	S12	S23	B23	B56
7	B89	S89	S23	S34	B67	T56	S12	S67	T23	B45	S910	T78	T67	T89	B56	T910	S56	B78	B34	B23	T34	T45	S45	S78
8	B56	S56	B78	B89	B34	T23	B67	S34	S78	T910	S67	T45	T34	T56	B23	T67	S23	B45	T89	T78	S89	S910	S12	S45
9	T78	B78	T910	B23	T56	S45	T89	B56	S12	T34	B89	S67	S56	S78	T45	S89	B45	T67	T23	S910	S23	S34	B34	B67
10	T45	B45	T67	T78	T23	S12	T56	B23	B67	S89	B56	S34	S23	S45	S910	S56	T910	T34	S78	S67	B78	B89	T89	B34
11	S45	T45	S67	S78	S23	T910	S56	T23	T67	B89	T56	B34	B23	B45	S12	B56	S910	S34	B78	B67	T78	T89	S89	T34
12	S23	T23	S45	S56	B89	T78	S34	S89	T45	B67	T34	T910	T89	B23	B78	B34	S78	S12	B56	B45	T56	T67	S67	S910
13	B67	S67	B89	S12	B45	T34	B78	S45	S89	B23	S78	T56	T45	T67	B34	T78	S34	B56	T910	T89	S910	T23	S23	S56

6

Summary

This chapter summarizes the literature survey and experiment proposal of this report. Car driving is a manual control task accessible to almost everyone, in which anticipatory information must be used. Understanding how drivers control cars within lane keeping, where they are tasked with keeping their vehicle between the boundary lines of the road, requires an extended understanding of human control with respect to preview control.

For many years, studies have aimed at providing such an understanding by constructing models that describe and predict control behaviour, of which an overview is given by Steen et al. [23]. The most generally accepted form were two-point models, first noted by Donges [4], where driver control is described by both an open loop anticipatory and a closed loop lateral position deviation response. However, these studies often only focussed on driver responses by comparing performance measures, meaning that the underlying responses of drivers to the visual information of the road were not investigated and thus provide little proof on the proposed structure.
The recent development of a multiloop driver control model by Van der El [25] for the first time allowed to accurately determine the *individual* responses of the lateral position, heading angle and target signal using system identification. The subsequent parameter estimation describes the use of a single modelled aim point, from which the stabilization metrics are determined. Furthermore, drivers can adopt corner cutting behaviour by prefiltering the road signal, such that the human target response can partially neglect higher frequencies, resulting in a more laid-back driving style (lower control activity).
What this model does not describe is how these modelled points correspond to the actual visual perception of the driver. Therefore, the main question to be answered in this experiment is:

- How does visual information on the stabilization metrics guide the control in lane keeping tasks?

An empirical study originally performed by Land & Horwood [8] had a similar goal. Land & Horwood performed an occlusion experiment, in which large parts of the visuals were occluded by an overlay, such that the importance of spatial regions on driver control could be determined. Depending on the run, drivers were shown either one or two slits to look through of which the position was varied. It was determined that to

reach similar performance as in the baseline condition (full visuals), drivers require at least two opened slits separated by roughly 5°. This dependency on a second slit became larger for increasing velocities.
Land & Horwood, as well as various replication experiments, only measured the change in driver behaviour by a performance metric (the change in the standard deviation of the lateral position error compared to the full visuals). However, it is not necessarily true that an identical performance between occlusion scenarios and the full visuals also entails identical driver behaviour. Drivers can achieve the same performance for different driving strategies.

The proposed experiment is therefore a combination of the Van der El and Land & Horwood experiments. As the use of system identification requires the use of multisine signals, it must first be confirmed that the road geometry of real roads, such as used by Land & Horwood, can be approximated by a multisine equivalent, and that this change in road geometry has no effect on driver behaviour.
The presence of occlusion manipulates the information that can be used by drivers, such that it is likely that the modelled aim points of the Van der El model are affected. Furthermore, it is possible that only showing a single slit affects the possibility of humans adopting a corner cutting strategy, as the leading information used for this smoothing on the aim point is not available anymore or becomes small. This is likely to have an effect on the control activity of the human, as it is forced to respond to the road without any pre-smoothing.
Finally, it is not a necessity that drivers aim their focal point, the gaze, directly at either one of the aim points. Considering that the visual perception of information might be through peripheral of parafoveal vision, the experiment will also include the measurement of the driver eye gaze in the time domain, of which the longitudinal statistical distribution can be compared to the model aim points and, if present, the position of the slits.

Bibliography

[1] Bosetti, P., Da Lio, M. and Saroldi, A. [2015], 'On Curve Negotiation: From Driver Support to Automation', *IEEE Transactions on Intelligent Transportation Systems* **16**(4), 1–12.

[2] Chatziastros, A., Wallis, G. M. and Bülthoff, H. H. [1999], 'The Effect of Field of View and Surface Texture on Driver Steering Performance', *Vision in Vehicles* **VII**(71).

[3] Cloete, S. and Wallis, G. [2011], 'Visuomotor control of steering: the artefact of the matter', *Experimental Brain Research* pp. 475–489.

[4] Donges, E. [1978], 'A Two-Level Model of Driver Steering Behavior', *Human Factors* **20**(6), 691–707.

[5] Frissen, I. and Mars, F. [2013], 'The effect of visual degradation on anticipatory and compensatory steering control', *The Quarterly Journal of Experimental Psychology* pp. 1–8.

[6] Gibson, J. J., Olum, P. and Rosenblatt, F. [1955], 'Parallax and Perspective during Aircraft Landings', *The American Journal of Psychology* **69**(3), 372–385.

[7] Krendel, E. S. and McRuer, D. T. [1960], 'A Servomechanics Approach to Skill Development', *Journal of the Franklin Institute* **269**(1), 24–42.

[8] Land, M. F. and Horwood, J. [1995], 'Which parts of the road guide steering?', *Nature* **377**, 339 – 340.

[9] Land, M. F. and Horwood, J. [1998], 'How Speed Affects The Way Visual Information Is Used In Steering', *Vision in Vehicles* **6**, 43–50.

[10] Land, M. F. and Lee, D. N. [1994], 'Where we look when we steer', *Nature* **369**, 742 – 744.

[11] Lappi, O. and Mole, C. D. [2018], 'Visuomotor Control, Eye movements, and Steering: A Unified Approach for Incorperating Feedback, Feedforward, and Internal Models'.

[12] MacAdam, C. C. [1980], 'An Optimal Preview Control for Linear Systems', *Journal of Dynamic Systems, Measurement, and Control* **102**, 188–190.
URL: *http://hdl.handle.net/2027.42/65011*

[13] McRuer, D. T. and Jex, H. R. [1967], 'A Review of Quasi-Linear Pilot Models', *IEEE Transactions on Human Factors in Electronics* **HFE-8**(3), 231–249.

[14] Miall, R. C., Weir, D. J. and Stein, J. F. [1985], 'Visuomotor tracking with delayed visual feedback', *Neuroscience* **16**(3), 511–520.

[15] Miall, R. C., Weir, D. J. and Stein, J. F. [1993], 'Intermittency in Human Manual Tracking Tasks', *Journal of Motor Behavior* **25**(1), 53–63.

[16] Mole, C., Kountouriotis, G., Billington, J. and Wilkie, R. [2016], 'Optic flow speed modulates guidance level control: New insights into two-level steering', *Journal of Experimental Psychology: Human Perception and Performance* **42**(11), 1818–1838.

[17] Mourant, R. R. and Rockwell, T. H. [1970], 'Mapping Eye-Movement Patterns to the Visual Scene in Driving: An Exploratory Study', *Human Factors: The Journal of the Human Factors and Ergonomics Society* **12**(1), 81–87.

[18] Mulder, M., Pool, D. M., Abbink, D. A., Boer, E. R., Zaal, P. M. T., Drop, F. M., Van der El, K. and van Paassen, M. M. [2017], 'Manual Control Cybernetics: State-of-the-Art and Current Trends', *IEEE Transactions on Human-Machine Systems* .

[19] Salvucci, D. D. and Gray, R. [2004], 'A two-point visual control model of steering', *Perception* **33**(10), 1233–1248.

[20] Sentouh, C., Chevrel, P., Mars, F. and Claveau, F. [2009], A Sensorimotor Driver Model for Steering Control, *in* 'Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics', pp. 2462 – 2467.

[21] Sheridan, T. B. [1966], 'Three Models of Preview Control', *IEEE Transactions on Human Factors in Electronics* **7**(2), 91–102.

[22] Sivak, M. [1996], 'The information that drivers use: is it indeed 90% visual?', *Perception* **25**(9), 1081–1089.

[23] Steen, J., Damveld, H. J., Happee, R., van Paassen, M. M. and Mulder, M. [2011], A Review of Visual Driver Models for System Identification Purposes, *in* 'IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2011', pp. 2093 – 2100.

[24] van den Berg, A. V. [1996], 'Judgements of Heading', *Vision Research* **36**(15), 2337–2350.

[25] van der El, K. [2018], How Humans Use Preview Information in Manual Control, PhD thesis, Delft University of Technology, Faculty of Aerospace Engineering.

[26] van der El, K., Pool, D., Damveld, H., Van Paassen, M. and Mulder, M. [2016], 'An empirical human controller model for preview tracking tasks', *IEEE Transactions on Cybernetics* **46**(11), 2609–2621.

[27] van der El, K., Pool, D. M., van Paassen, M. M. and Mulder, M. [2018*a*], 'Effects of Linear Perspective on Human Use of Preview in Manual Control', *IEEE Transactions on Human-Machine Systems* **48**(5), 496–508.

[28] van der El, K., Pool, D. M., van Paassen, M. M. and Mulder, M. [2018*b*], 'Effects of Preview on Human Control Behavior in Tracking Tasks With Various Controlled Elements', *IEEE Transactions on Cybernetics* **48**(4), 1242–1252.

[29] van der El, K., Pool, D. M., van Paassen, M. M. and Mulder, M. [2019], Multiloop Identification and Modeling of Driver Visual Information Processing and Steering Control.

[30] Wandell, B. A. [1995], *Foundations of Vision*, Sinauer Associates.

[31] Wann, J. P. and Land, M. F. [2000], 'Steering with or without the flow: is the retrieval of heading necessary?', *Trends in Cognitive Sciences* **4**(8), 319–324.

[32] Wann, J. P. and Swapp, D. [2000], 'Why You Should Look Where You Are Going', *Nature Neuroscience* **3**(7), 647–648.

[33] Wann, J. P., Swapp, D. and Rushton, S. K. [2000], 'Heading perception and the allocation of attention', *Vision Research* **40**(18), 2533–2543.

[34] Warren, W. H., Morris, M. W. and Kalish, M. [1988], 'Perception of Translational Heading From Optical Flow', *Journal of Experimental Psychology* **14**(4), 646–660.

[35] Wilkie, R. M. and Wann, J. P. [2003], 'Eye-movements aid the control of locomotion', *Journal of Vision* **3**, 677–684.

[36] Wilkie, R. M. and Wann, J. P. [2006], 'Judgments of Path, Not Heading, Guide Locomotion', *Journal of Experimental Psychology: Human Perception and Performance* **32**(1), 88–96.

IV

Appendices

A

Briefing

This chapter includes the briefing as presented to the participants before the experiment.

SIMONA experiment briefing

Investigating the Use of Visual Information in Car Driving

The goal of this experiment is to better understand what visual information is important for lane-keeping in regular car driving. Car drivers directly respond to the road ahead of them, although we do not yet fully understand which parts of the road provide important information for drivers to stay within their lane. By occluding parts of the road, we will investigate which visual information is used for simple lane keeping tasks.

Control Task

During the experiment you will drive along a winding road, in the presence of wind disturbances that perturb the trajectory of your car. Your goal is to keep the car on the road and between the lines that mark the road boundaries, see Figure 1a.

In most conditions, large parts of the visuals will be occluded by black boxes, resulting in vertically spaced horizontal *slits* to look through. You can see parts of the road through these slits (see Figure 1b), but the occlusion blocks all other visual information that would normally be available for you when driving a car. By varying the vertical position of these slits, we measure how well you are still able to control the car. In some cases two slits are available on screen, in others only one slit is visible. Furthermore, we will also collect reference data when no occlusion is present, so when the full visual is available for you to use, as in Figure 1a. This serves as a baseline condition with which all other conditions are compared. In all these conditions your eye movements are measured using an eye tracker.

You are only able to steer the position of the car, meaning that the velocity of the vehicle remains fixed and constant throughout the experiment.

(a) (b)

Figure 1: The road visuals fully displayed (a) and partially occluded (b).

Apparatus
During the experiment you will sit in the left seat of the SIMONA research simulator (see Figure 2. You will use a steering wheel in front of you to steer the car. The velocity of the car will remain fixed during the experiment, so there are no pedals for accelerating or braking.
Furthermore, the motion of the simulator is turned off during the experiment, so you can only respond to what is happening on the screen in front of you.
The road is presented on the simulator's out-of-the-window viewing system, either fully visible as in Figure 1a, or partially occluded, such as in Figure 1b.
You will be asked to wear an eye tracker in the shape of regular glasses, which will remain on throughout the experiment. In some cases it might be necessary to calibrate this hardware between runs, requiring you to look at a fixed aimpoint located on the right of the simulator cabin.

Figure 2: The SIMONA Research Simulator

Experimental Procedure
The experiment will start with a training phase, to get you used to the task and the simulator. After your performance stabilizes, the experiment will start. A single driving run lasts about 2.5 to 3 minutes. At the end of each run, you will be informed of your score. As the experiment may induce motion sickness for some people, the experimenter will ask you for any signs of motion sickness after each run, on a scale from 0 to 10, see Table 1. The experiment will be stopped if your subjective evaluation of your sickness reaches 5 or higher.
In total, there are 26 experimental conditions (including the baseline). First, you will perform eight runs with full visuals (no occlusion), after which the occlusion conditions are tested in a randomized order. For each experimental condition with occlusion present there will be four runs. At the end of the experiment you will perform another four runs with full visuals. Short (about 10-15 minutes) breaks will be taken regularly after completing several conditions. In addition, the experiment is split over two days. The full experiment will take approximately 6 hours, resulting in 3 hours for both days.

Your Rights
Participation in the experiment is voluntary. This means that you can terminate your cooperation at any time. By participating in the experiment you agree that the collected data may be published. Your data will remain confidential and anonymous, so only the experimenter can link the results to a particular participant. To make sure that you understand and comply with the conditions of the experiment, you will be asked to sign an *informed consent* form.

Table 1: Misery Scale (MISC) scores.

symptom		**score**
no problems		0
slight discomfort but no specific symptoms		1
dizziness, warm, headache, stomach awareness, sweating, etc.	vague	2
	some	3
	medium	4
	severe	5
nausea	some	6
	medium	7
	severe	8
	retching	9
vomiting		10

B

Informed Consent Form

This chapter includes the Informed Consent Form as presented to and signed by the participants before the experiment.

Consent Form for *Investigating the use of visual information in car driving*

Please tick the appropriate boxes	Yes	No
Taking part in the study		
I have read and understood the experiment briefing dated 24/09/2018, or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	☐	☐
I consent voluntarily to be a participant in this study conducted by **Maurice Kolff,** under the supervision of **dr.ir. Daan Pool**, and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.	☐	☐
I understand that the study takes place in the SIMONA Research Simulator and involves the measurement of my steering behaviour in a car driving task, as well as the recording of my eye movements with a head-worn eye tracker, which will both be stored for analysis.	☐	☐
Risks associated with participating in the study		
I understand that taking part in the study involves a small risk of physical discomfort and nausea due to simulator motion sickness. I understand that this is measured using the Misery Scale and the experiment ends when my subjective score reaches 5 or higher.	☐	☐
I confirm that the researcher has provided me with a detailed safety and operational instructions on the use of the simulator.	☐	☐
Use of the information in the study		
I understand that information I provide will be used for the scientific publication of the data.	☐	☐
I understand that personal information collected about me that can identify me, such as my name, will not be shared beyond the study team.	☐	☐
Future use and reuse of the information by others		
I give permission for the steering behaviour data and eye view video recordings that I provide to be archived at Delft University of Technology, so it can be used for future research and learning.	☐	☐
Eligibility to drive		
I confirm that I am in the possession of a European driving licence, category B.	☐	☐

Signatures

_____	_____	_____
Name of participant	Signature	Date

I have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

_____	_____	_____
Researcher name	Signature	Date

TU Delft

C

Conversion of Eye Gaze Data

The measurement of the vertical eye gaze position with respect to the road is measured by a Tobii pro glasses 2 eye tracker. This appendix gives additional information on the gaze processing algorithm as explained in Paper II.

C.1. Hardware Interface

All drivers wore the eye glasses during the experiment, for the full duration of the sessions. This head-worn tracker records a full HD (1920x1080) video stream and was connected to a computer laptop in the SIMONA control room through a LAN-cable, which stored the video and gaze data. These video files do not contain the gaze points itself (i.e., it was the raw video data), meaning that it is necessary to import to gaze data and add it to the video files frame by frame.

However, there are two main problems that occur when using combining the frames and gaze data together:

- The measurement of the data is stored within a separate interface from the driver, at a different rate (60Hz) compared to the video frames (25Hz)
- The gaze data is measured in a head-fixed reference frame, compared to the body-fixed reference frame of the simulator required for quantitatively determining where drivers look with respect to the visual scene.

Therefore, an algorithm was written that can link the video and data together, and subsequently determine the gaze position relative to the slits or the full visuals.

C.2. Data extraction

The Tobii eye tracker stores its data within a .json data structure. An example of a single timestamp (in nanoseconds) within a stored timetrace is shown in Figure C.1. The system does not only measure the [x,y]-coordinates of the gaze [1], but also:

- **pc** - Pupil center, measures the position of the pupil with respect to the eye in mm, for both eyes separately.

[1] Tobi Pro Glasses 2 API Developer's Guide v.1.12.2. TobiiAB, Stockholm, 2016

```
{"ts":1345832874,"s":0,"pc":[31.62,-27.68,-28.18],"eye":"left"}
{"ts":1345832874,"s":0,"pd":4.98,"eye":"left"}
{"ts":1345832874,"s":0,"gd":[0.0173,0.1677,0.9857],"eye":"left"}
{"ts":1345832874,"s":0,"pc":[-32.56,-23.66,-26.08],"eye":"right"}
{"ts":1345832874,"s":0,"pd":4.65,"eye":"right"}
{"ts":1345832874,"s":0,"gd":[0.0227,0.1334,0.9908],"eye":"right"}
{"ts":1345832874,"s":0,"gp":[0.4925,0.4418],"l":92022}
{"ts":1345832874,"s":0,"gp3":[7.26,32.55,354.99]}
```


Figure C.1: Example of a single data timetrace of the Tobii eye tracker.

- **pd** - Pupil dilation, which measures the size of the pupil in mm, for both eyes separately.
- **gp** - Gaze position, measures the [x,y]-coordinates of the gaze based on the relative orientation of the eyes with respect to the glasses, as fractions of the video frame size, for both eyes averaged.
- **gp3** - Gaze position 3D, measures the [x,y,z]-coordinates similarly as in gp, but also estimates the z-coordinate based on the pupil dilation, in mm.
- **gd** - Gaze direction, measures the [x,y,z] as a vector from the eyes, for both eyes separately.

As the position must be related to the geometry of the slits for quantitative comparison, the 'gp' structure provides the most direct approach.

C.3. Conversion of Timestamps

The complete data file is imported and only the 'gp' data was used for analysis. To link the video and data together, the algorithm loops through all video frames and approximates the gaze data at this point. The videos of runs are stored together, such that it would often occur that 8 or even more runs were recorded within a single video (and a corresponding data file). If calibration was required, the recording was stopped and a new file was created. This means that for a single participant, typically 8-10 different video and corresponding data files exist.

The starting point of each run in seconds (t_s) are written down manually and imported through an Excel-sheet, together with a label that indicates in what files the run is recorded. At the first frame, the algorithm starts at $i = t_s \cdot \dot{d}$, where $\dot{d}$ is the data rate (measured as the length of the data file divided by the total increase of the timestamp within the file). Within that single frame, the detection of the slit geometry and the reference frame conversion take place. Once the algorithm is finished, it jumps to the next frame and the data counter i is increased by $\frac{\dot{d}}{\dot{v}}$, i.e., how much data measurements were taken within a single frame, where $\dot{v}$ is the frame rate of the video. This value is rounded, as only integer values can be used to evaluate the gaze data, and thus results in a nearest-neighbour approach.

C.4. Reference Frame Data Conversion

Once the data is available in a frame, the algorithm detects the geometry of the visuals, either for the occlusion slits or the full-visuals.

C.4.1. Detection of slit geometry

To be able to relate this data to the control model metrics, it is necessary to post-process the coördinates by detecting parts of the video frame of which the positions are already known. The occlusion slits provide a perfect method of doing so, as the angles at which the slits are located below the horizon are known beforehand for each condition.

The core of the algorithm works by detecting the slits in every single video frame, using edge detection to find transitions from occluded to non-occluded parts and vice versa. Firstly, the frame RGB-levels are determined. The algorithm then loops through vertical segments at thirty evenly spaced points, ranging from $\left(\frac{5}{16}\right)$ to $\left(\frac{11}{16}\right)$ as fractions of the frame width. Data outside of this range is not analyzed, as this typically corresponds to more faded-out slits (due to only one projector being active) and thus poor detection quality.

The points where the RGB-values switch from (0,0,0) to (0,40,0) or higher is used as the transition point, being able to detect both the grass (green) and the road (light grey). This results in 60 transition points, 30 at the top of a slit and 30 at the bottom. Linear regression is used to approximate the overall transition position and orientation, i.e. determine the position of the slit on the screen (and measure its pixels-per-degree ratio), for both the top and bottom of a slit. The average of these metrics is taken to approximate the slit position as accurately as possible. The head can freely move, meaning that the coordinates of the gaze can be rotated with respect of the slit geometry. Therefore, the gaze coordinates must be corrected by an Euler rotation by the rotation angle of the slit α:

$$y_w = -x_b \sin\alpha + y_b \cos\alpha \tag{C.1}$$

Now that the position of the gaze and the slits are both known within the same reference frame, it is possible to relate these together. This is done by measuring the difference in vertical position and multiplying by the pixel-per-degree ratio, as determined earlier. In the presence of two slits, these steps are performed separately for each slit and the resulting gaze is averaged.

C.4.2. Detection of baseline geometry

The algorithm to extract the gaze in baseline (full visual) conditions is the same as for the slits, with some exceptions. Without the presence of slits, other characteristics of the scene must be used. In this case, the transition from the top of the video screen and the horizon are used, as shown in Figure C.2. For the edge of the video screen the RGB-threshold is set to (0,60,0), and for the horizon to (0,0,100). The horizon position is typically difficult to approximate, as the color levels between the sky and grass/road are closer to each other. Furthermore, the horizon includes trees and other scenery that affect the estimation, making it even more important to use enough regression points. Specifically, a high tree section within the scene highly affected the regression by locally placing the transition points higher than parts that contained almost no trees. Therefore, the algorithm is also able to discard these high tree sections by looking at the gradient *between* the regression points and subtracting a 10-pixel 'penalty' at these positions.

Figure C.2: Example of a recorded baseline frame, showing the transition points (green) regressions of the scene edge and the horizon (white lines), as well as the gaze position (red).

D

Eye gaze measurements

Subject 1

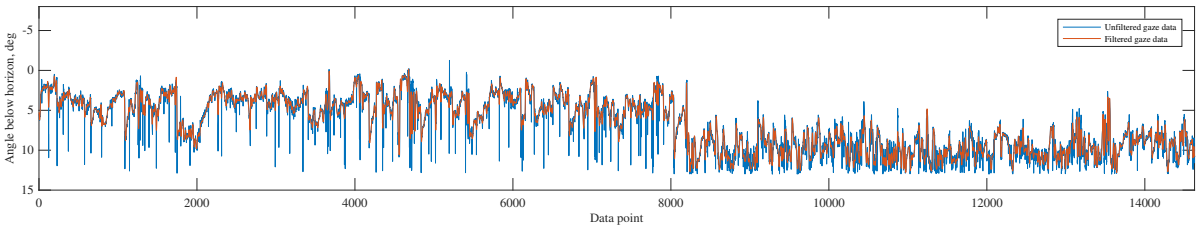

Figure D.1: Eye gaze timetrace, subject 1, Condition NO

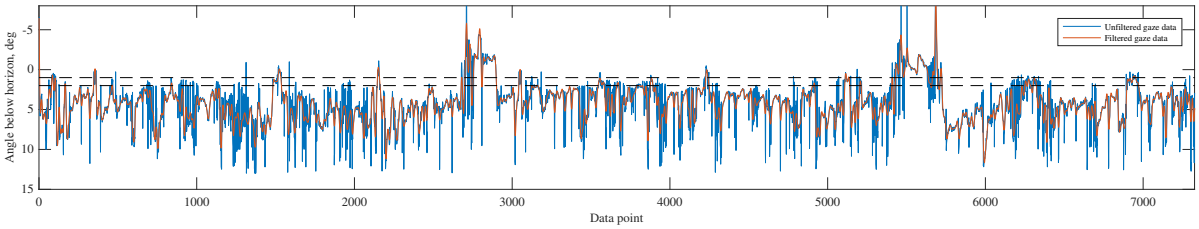

Figure D.2: Eye gaze timetrace, subject 1, Condition S12

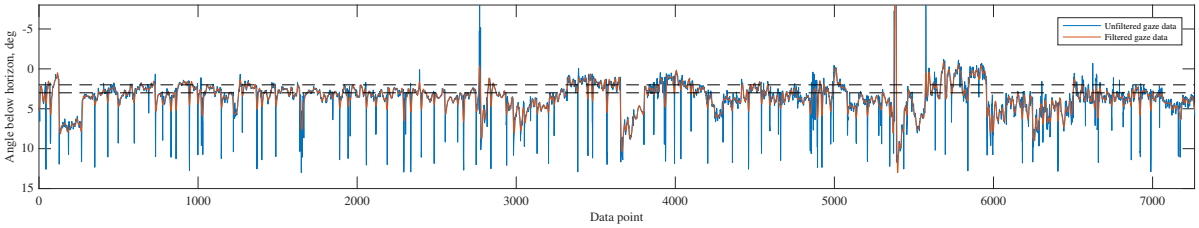

Figure D.3: Eye gaze timetrace, subject 1, Condition S23

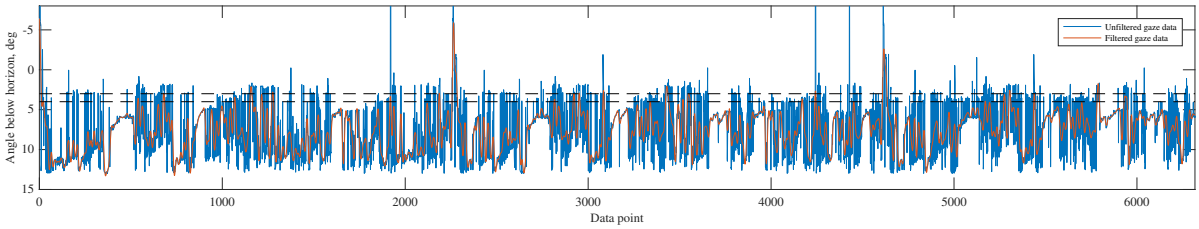

Figure D.4: Eye gaze timetrace, subject 1, Condition S34

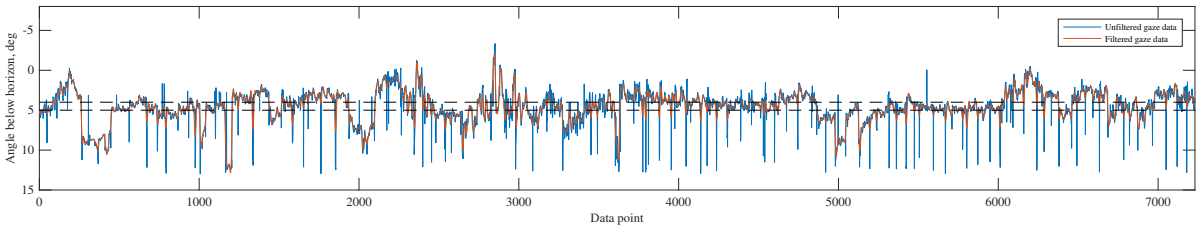

Figure D.5: Eye gaze timetrace, subject 1, Condition S45

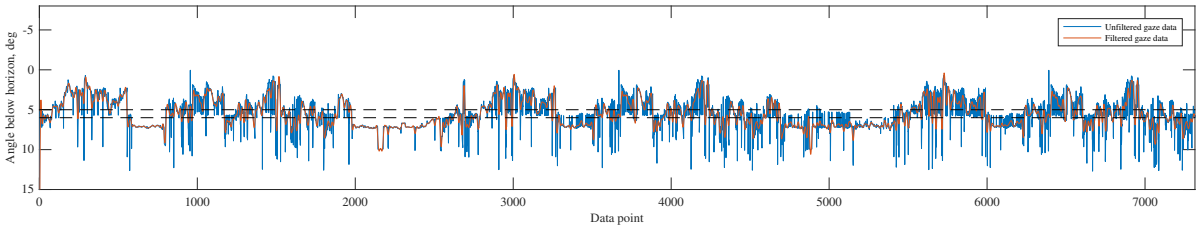

Figure D.6: Eye gaze timetrace, subject 1, Condition S56

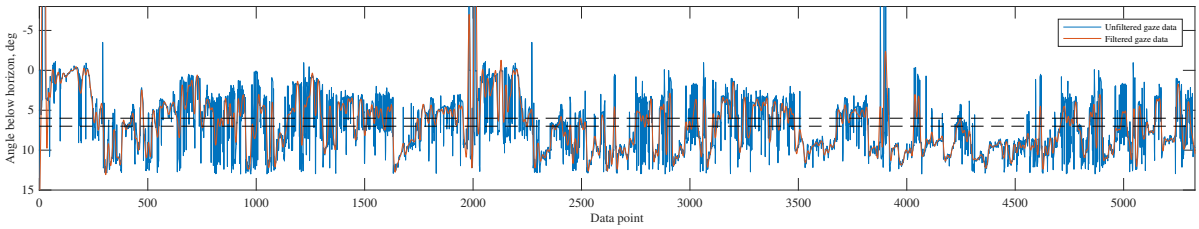

Figure D.7: Eye gaze timetrace, subject 1, Condition S67

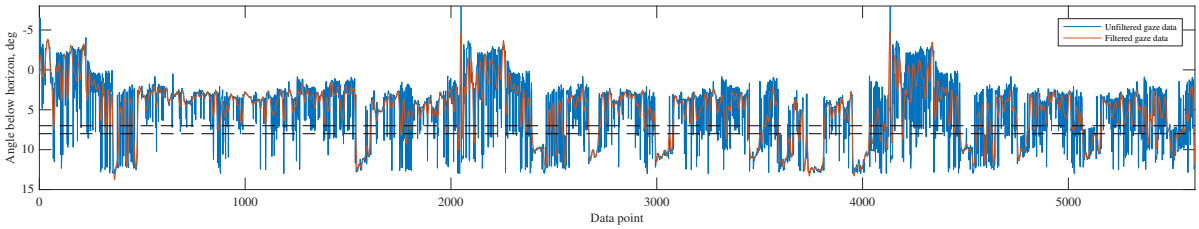

Figure D.8: Eye gaze timetrace, subject 1, Condition S78

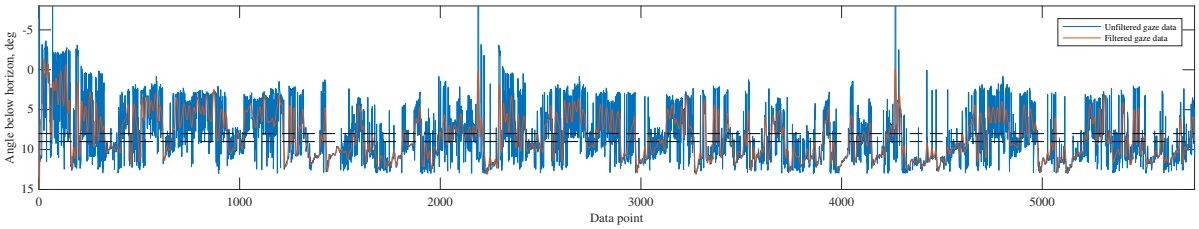

Figure D.9: Eye gaze timetrace, subject 1, Condition S89

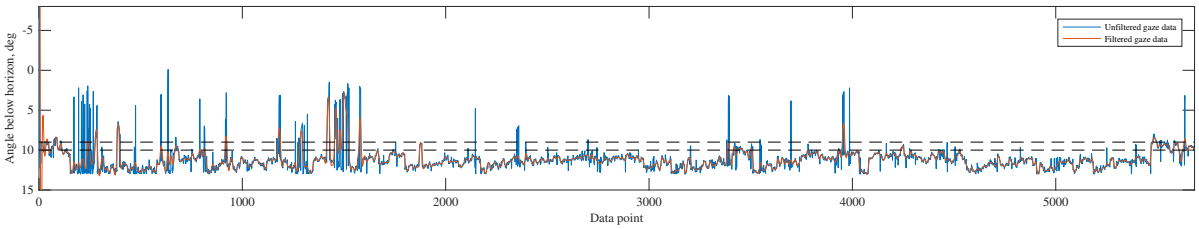

Figure D.10: Eye gaze timetrace, subject 1, Condition S910

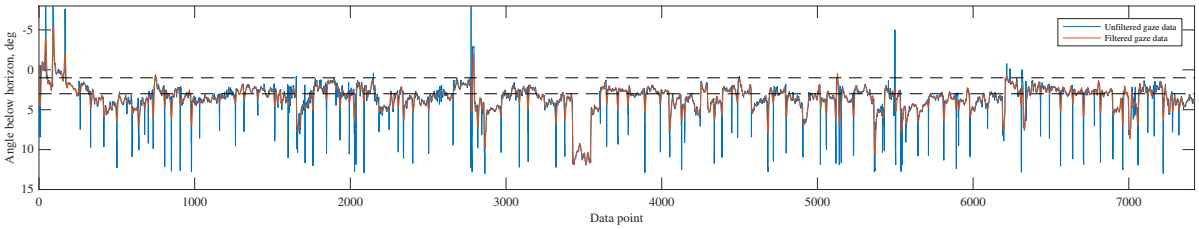

Figure D.11: Eye gaze timetrace, subject 1, Condition S13

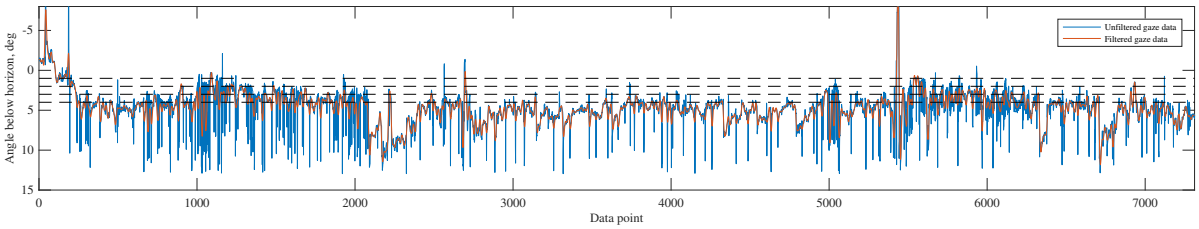

Figure D.12: Eye gaze timetrace, subject 1, Condition T34

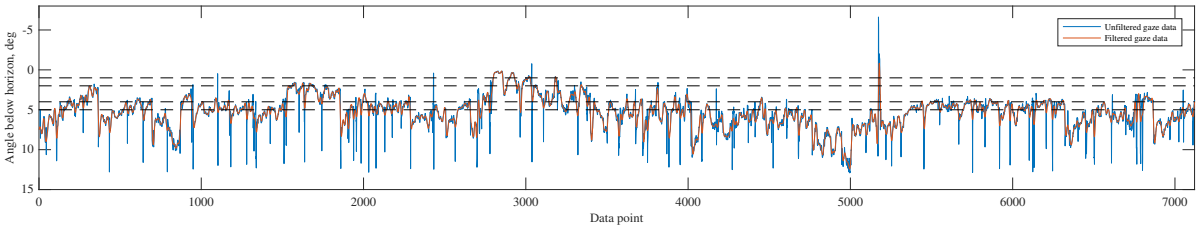

Figure D.13: Eye gaze timetrace, subject 1, Condition T45

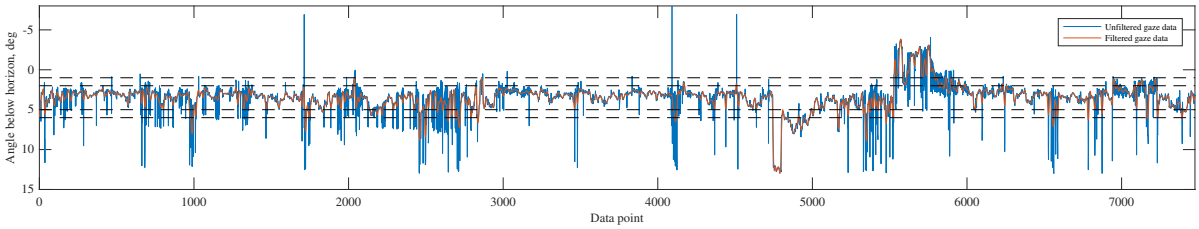

Figure D.14: Eye gaze timetrace, subject 1, Condition T56

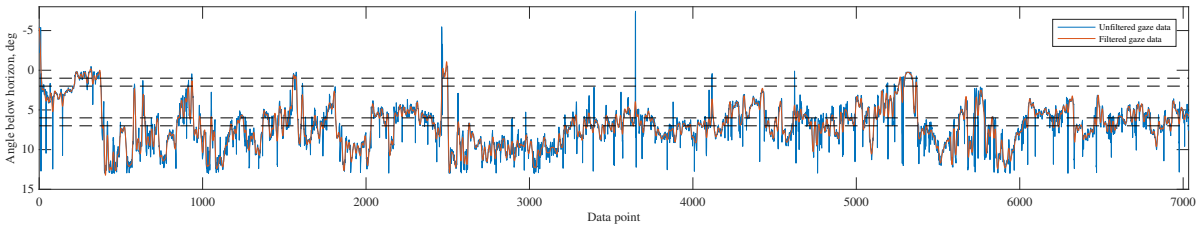

Figure D.15: Eye gaze timetrace, subject 1, Condition T67

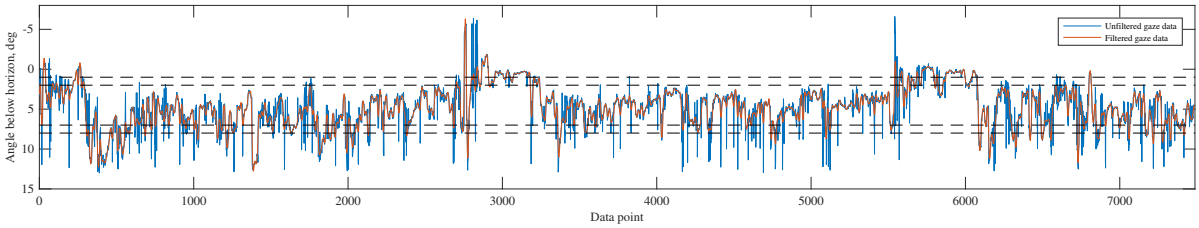

Figure D.16: Eye gaze timetrace, subject 1, Condition T78

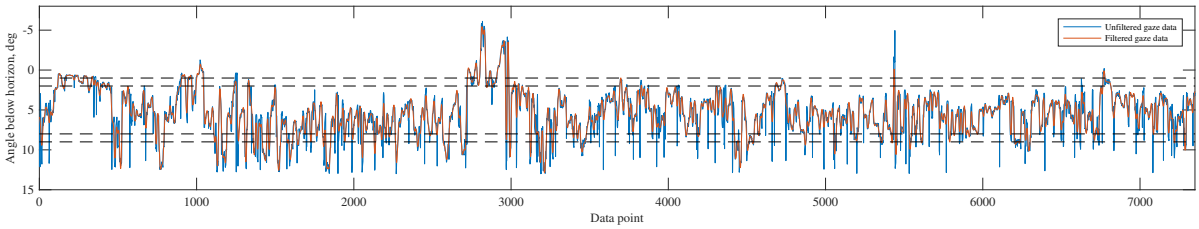

Figure D.17: Eye gaze timetrace, subject 1, Condition T89

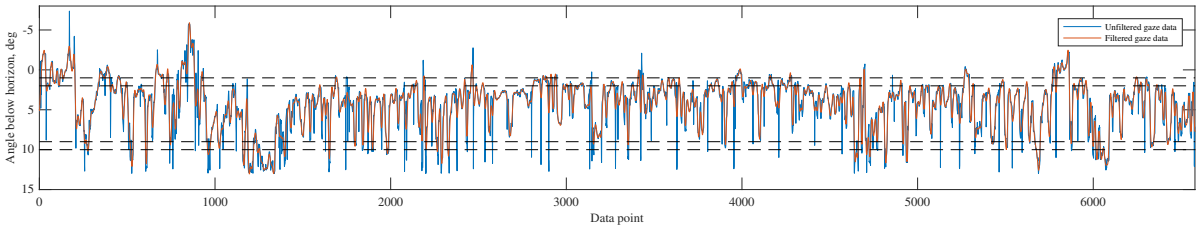

Figure D.18: Eye gaze timetrace, subject 1, Condition T910

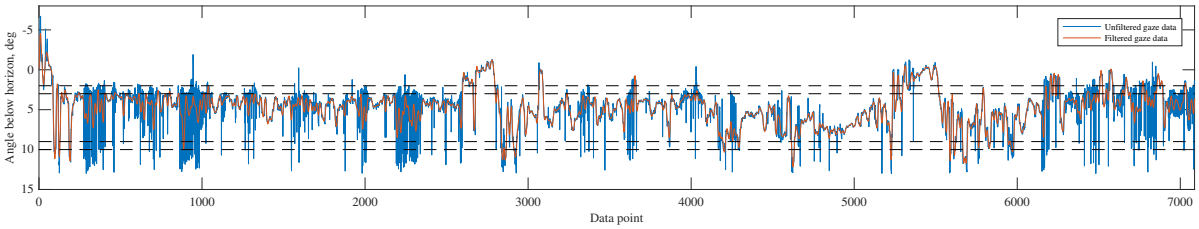

Figure D.19: Eye gaze timetrace, subject 1, Condition B23

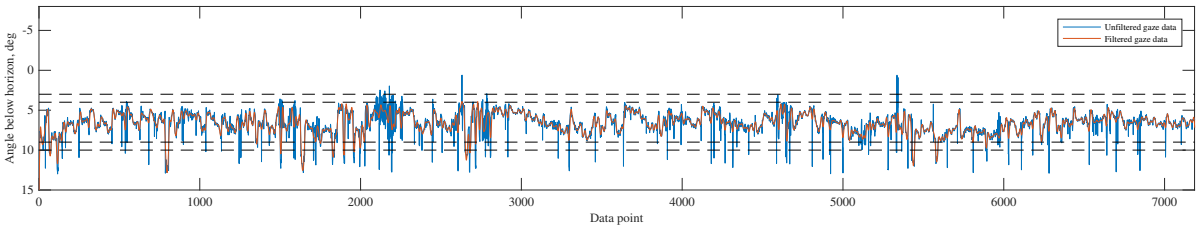

Figure D.20: Eye gaze timetrace, subject 1, Condition B34

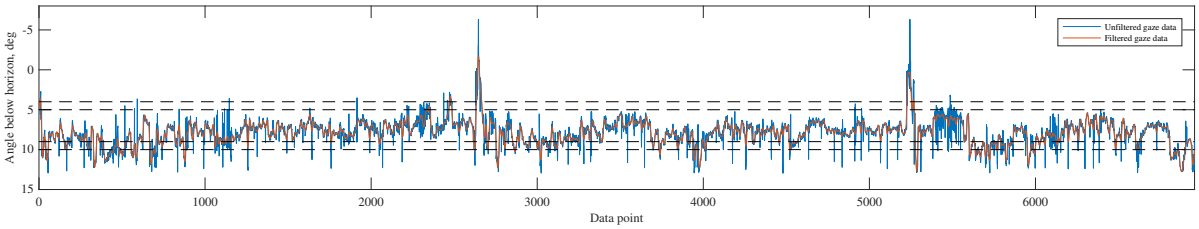

Figure D.21: Eye gaze timetrace, subject 1, Condition B45

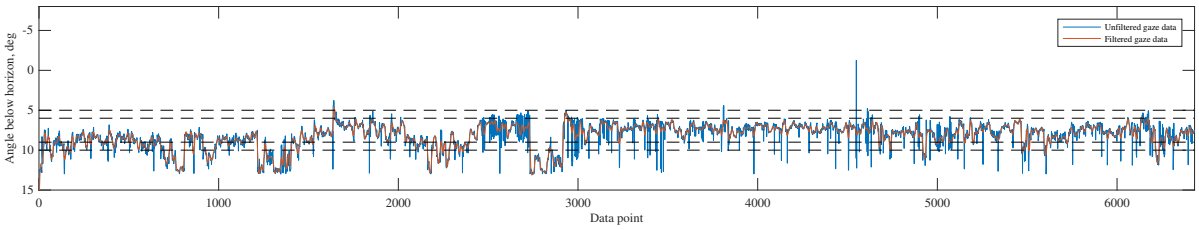

Figure D.22: Eye gaze timetrace, subject 1, Condition B56

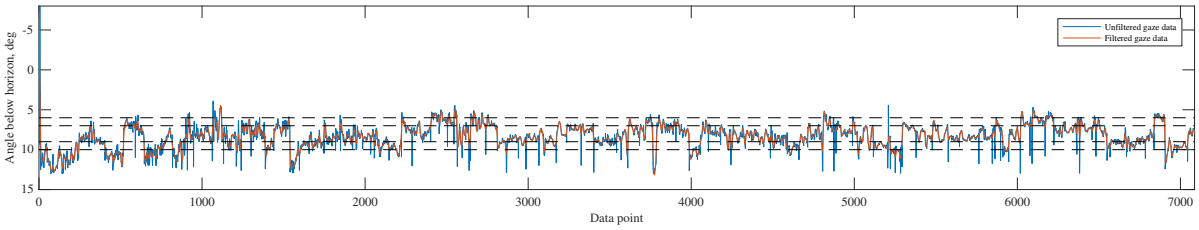

Figure D.23: Eye gaze timetrace, subject 1, Condition B67

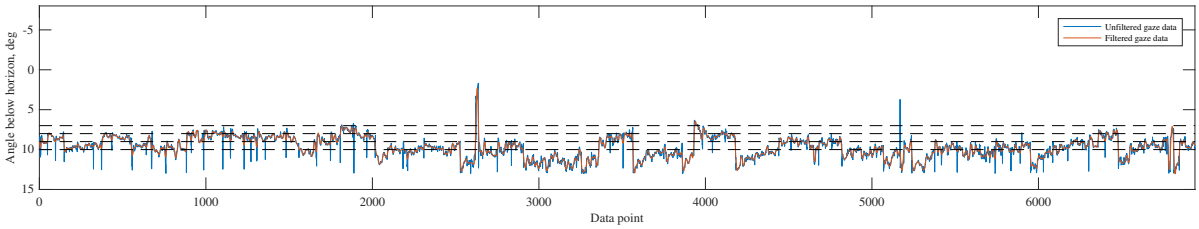

Figure D.24: Eye gaze timetrace, subject 1, Condition B78

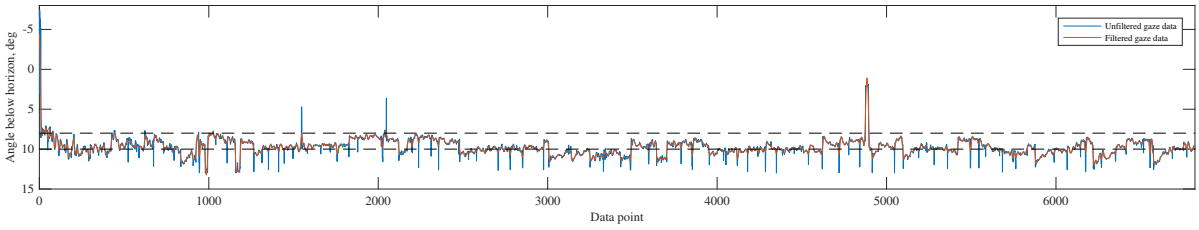

Figure D.25: Eye gaze timetrace, subject 1, Condition S810

Subject 2

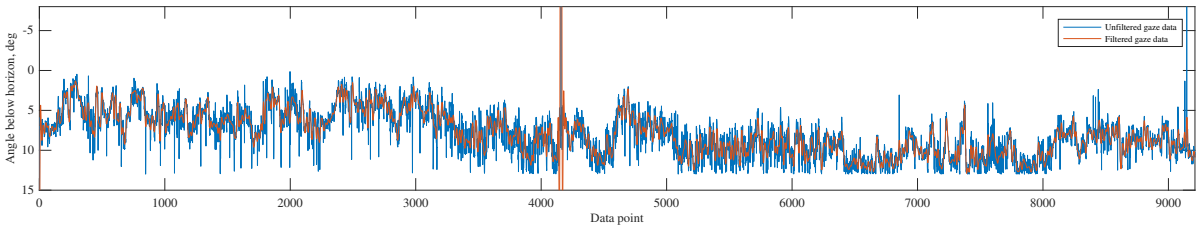

Figure D.26: Eye gaze timetrace, subject 2, Condition NO

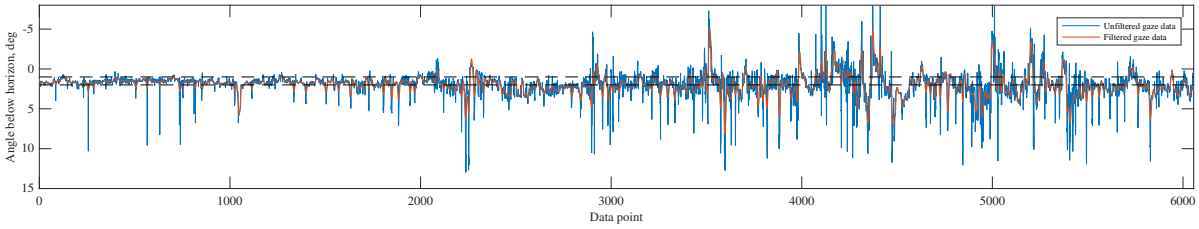

Figure D.27: Eye gaze timetrace, subject 2, Condition S12

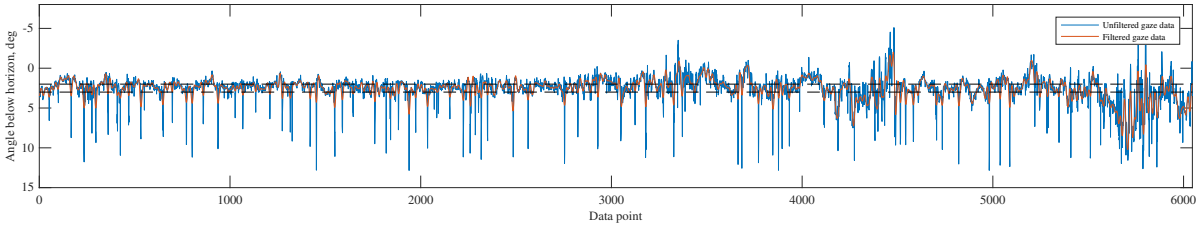

Figure D.28: Eye gaze timetrace, subject 2, Condition S23

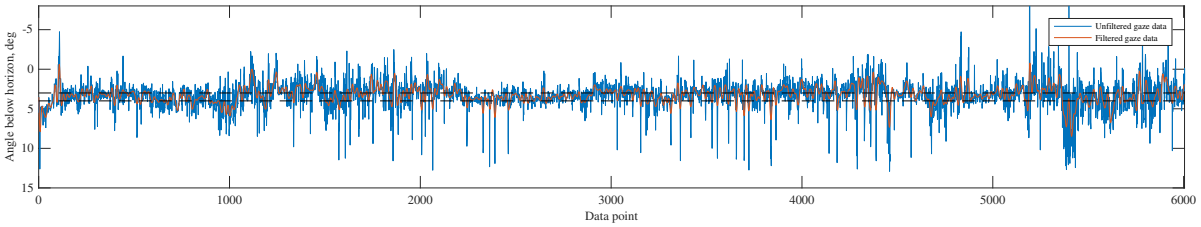

Figure D.29: Eye gaze timetrace, subject 2, Condition S34

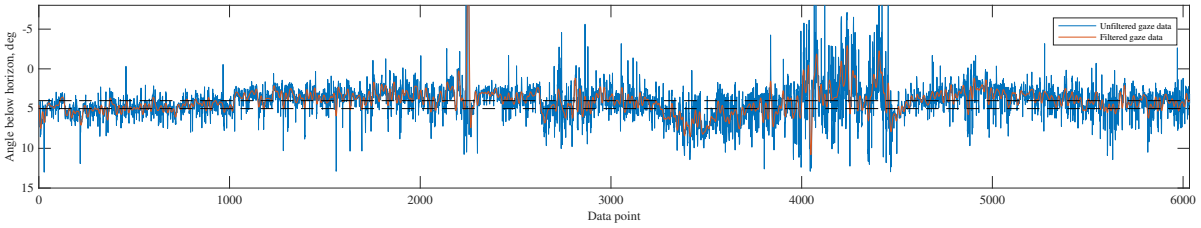

Figure D.30: Eye gaze timetrace, subject 2, Condition S45

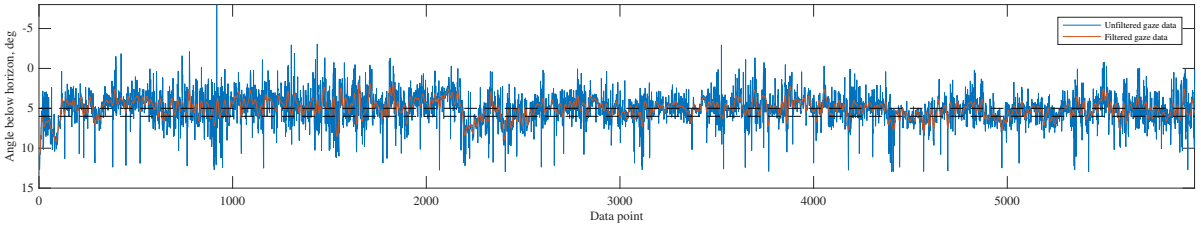

Figure D.31: Eye gaze timetrace, subject 2, Condition S56

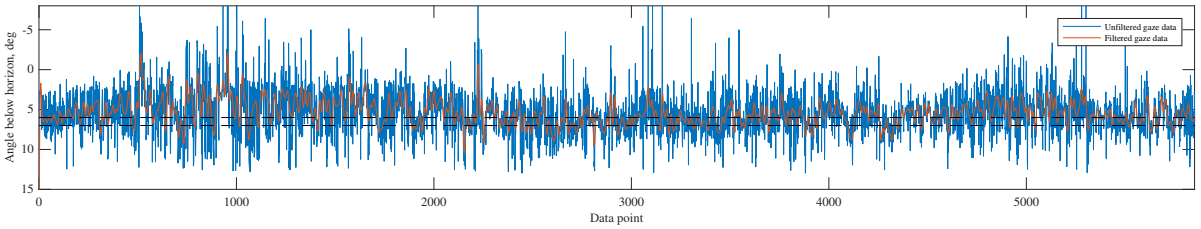

Figure D.32: Eye gaze timetrace, subject 2, Condition S67

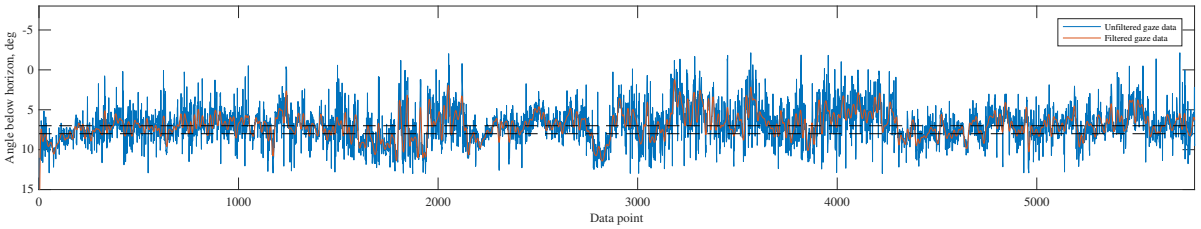

Figure D.33: Eye gaze timetrace, subject 2, Condition S78

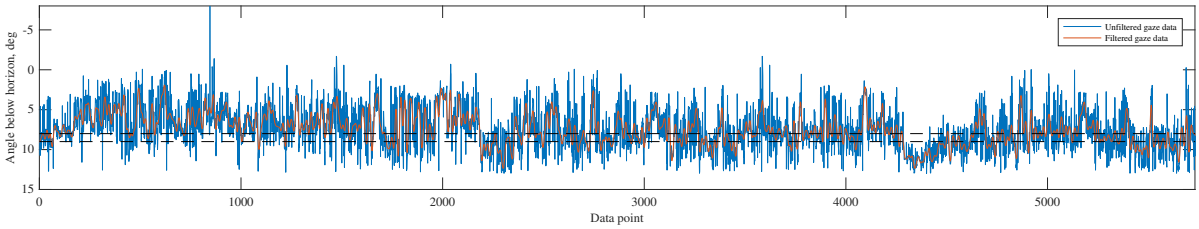

Figure D.34: Eye gaze timetrace, subject 2, Condition S89

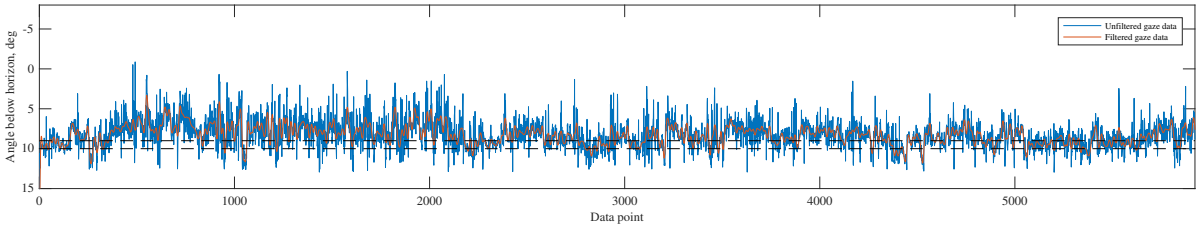

Figure D.35: Eye gaze timetrace, subject 2, Condition S910

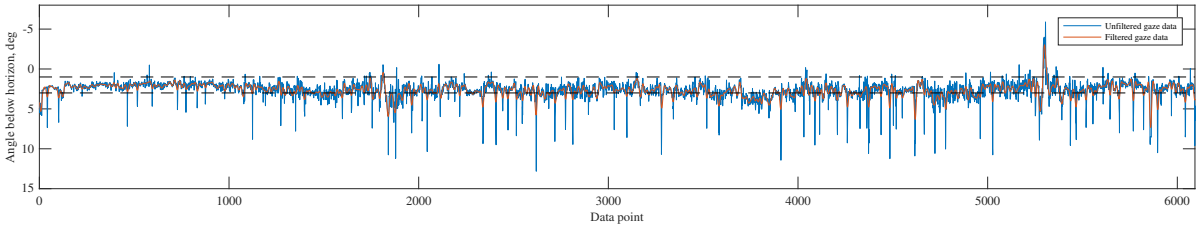

Figure D.36: Eye gaze timetrace, subject 2, Condition S13

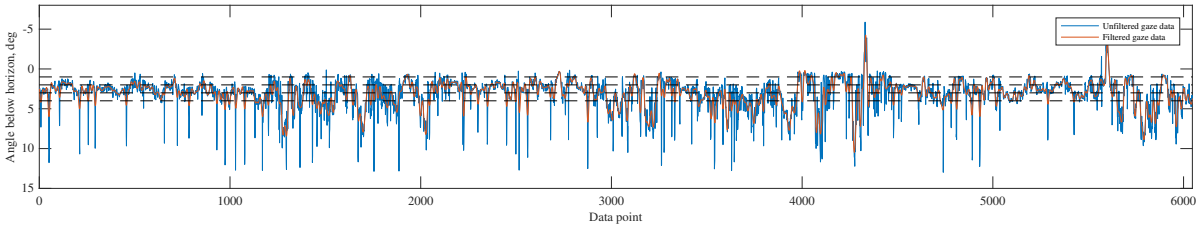

Figure D.37: Eye gaze timetrace, subject 2, Condition T34

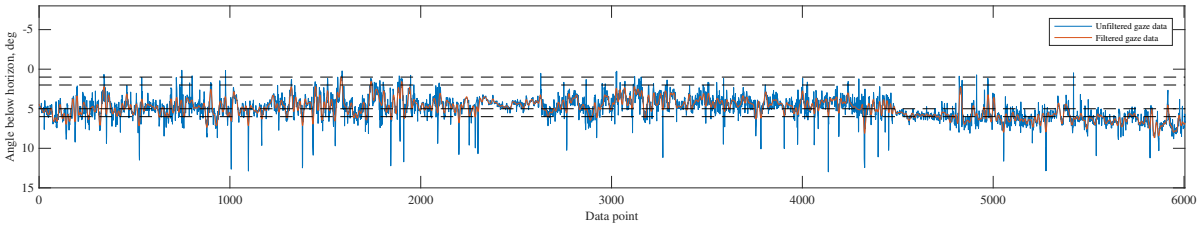

Figure D.38: Eye gaze timetrace, subject 2, Condition T56

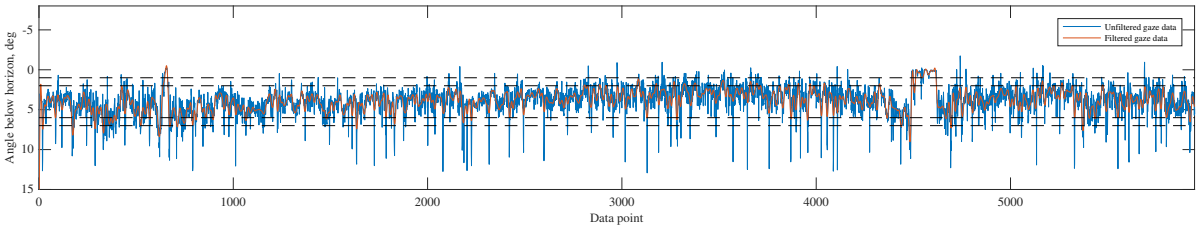

Figure D.39: Eye gaze timetrace, subject 2, Condition T67

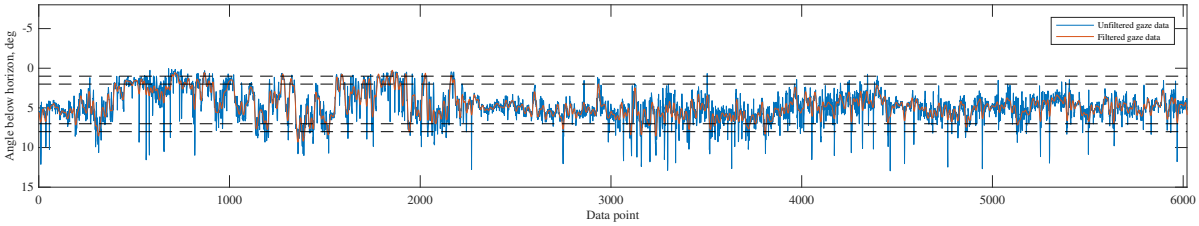

Figure D.40: Eye gaze timetrace, subject 2, Condition T78

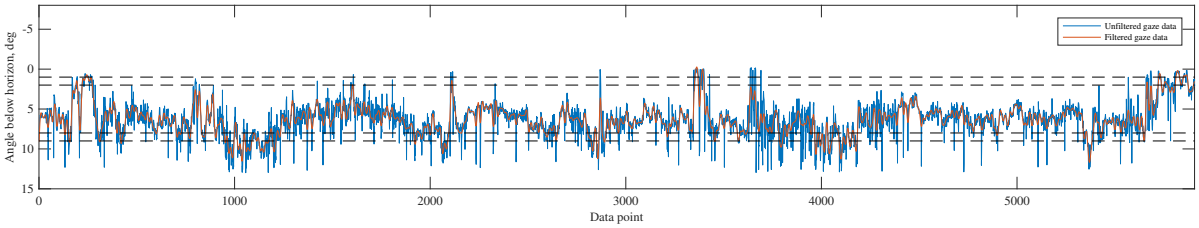

Figure D.41: Eye gaze timetrace, subject 2, Condition T89

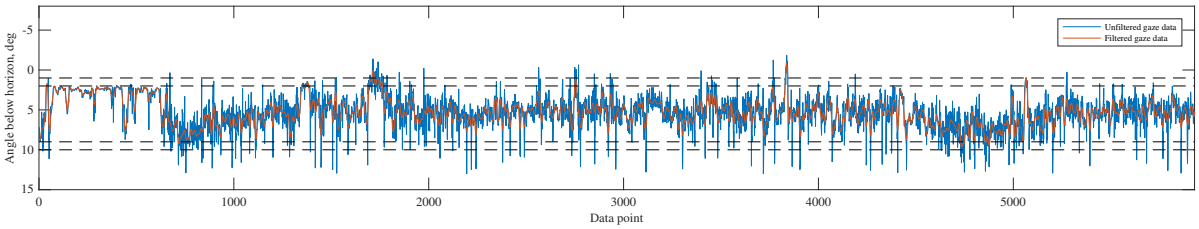

Figure D.42: Eye gaze timetrace, subject 2, Condition T910

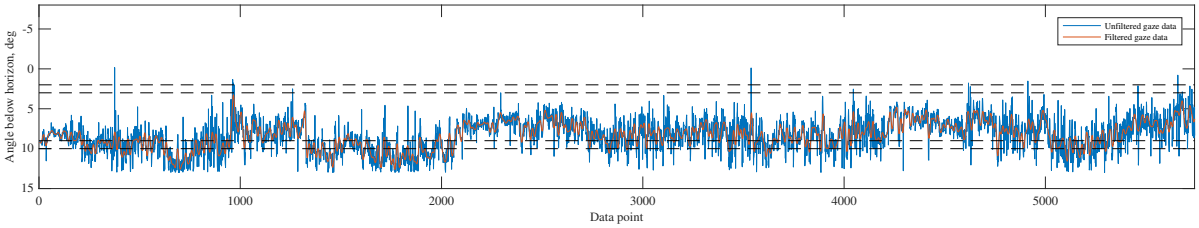

Figure D.43: Eye gaze timetrace, subject 2, Condition B23

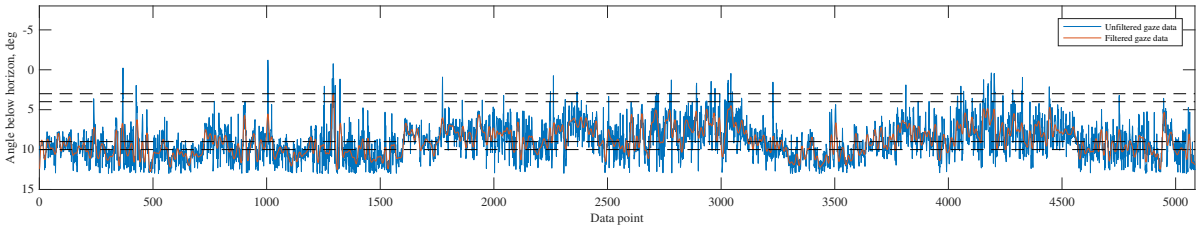

Figure D.44: Eye gaze timetrace, subject 2, Condition B34

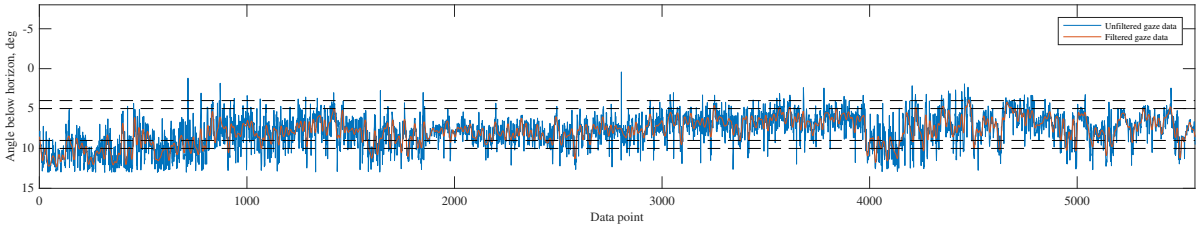

Figure D.45: Eye gaze timetrace, subject 2, Condition B45

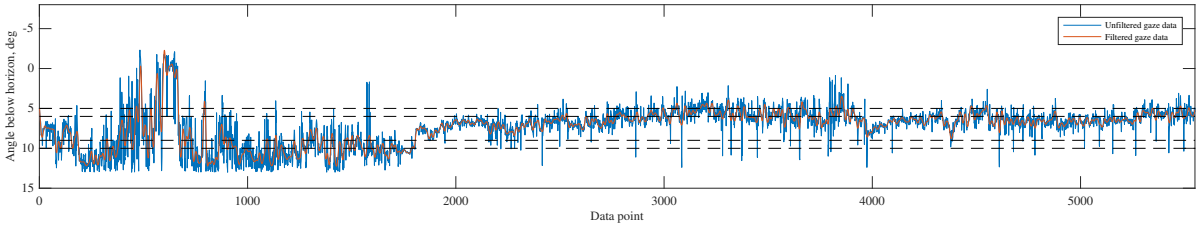

Figure D.46: Eye gaze timetrace, subject 2, Condition B56

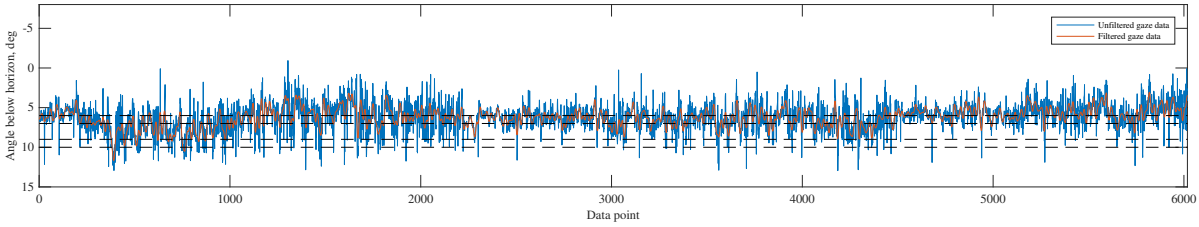

Figure D.47: Eye gaze timetrace, subject 2, Condition B67

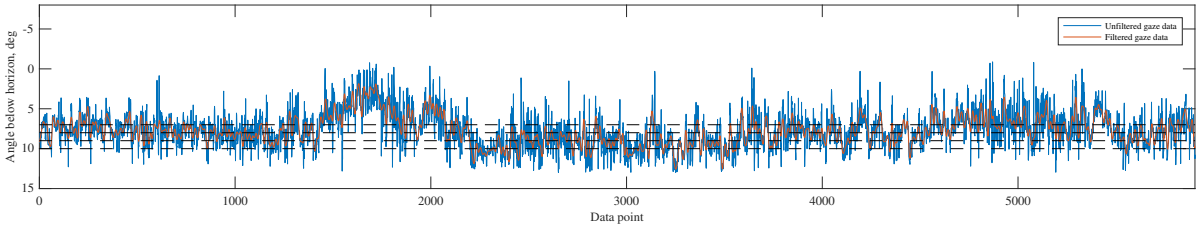

Figure D.48: Eye gaze timetrace, subject 2, Condition B78

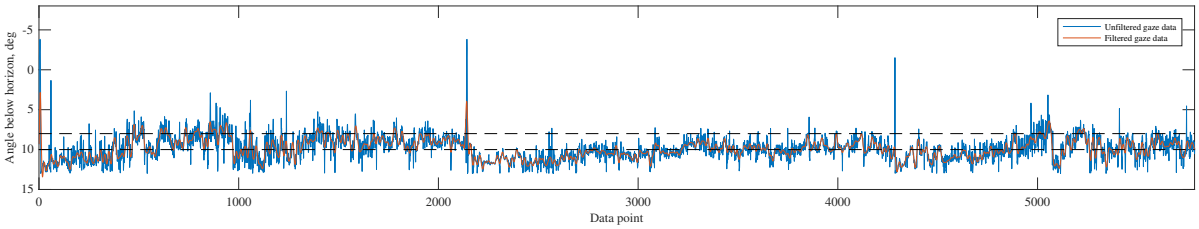

Figure D.49: Eye gaze timetrace, subject 2, Condition S810

Subject 3

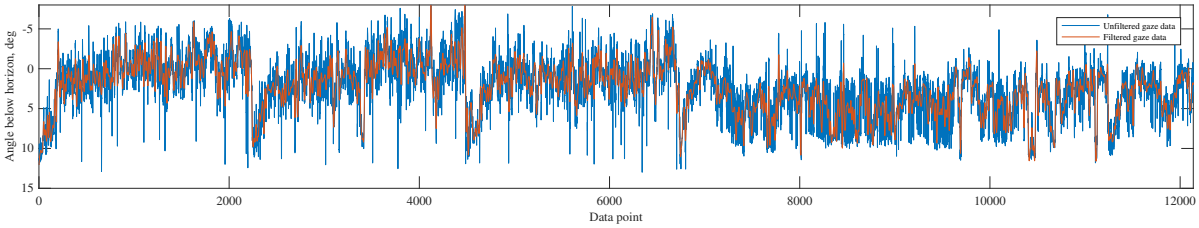

Figure D.50: Eye gaze timetrace, subject 3, Condition NO

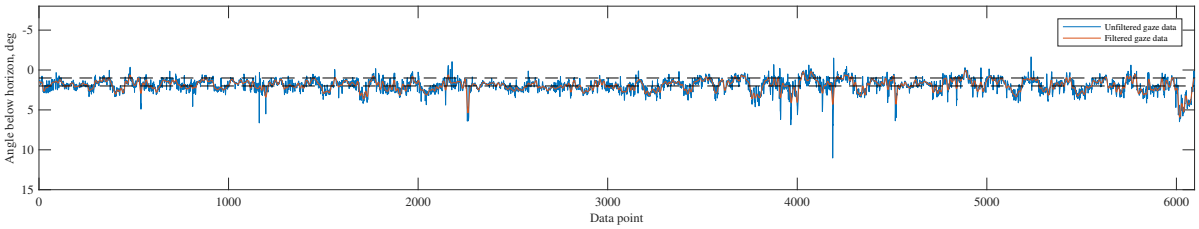

Figure D.51: Eye gaze timetrace, subject 3, Condition S12

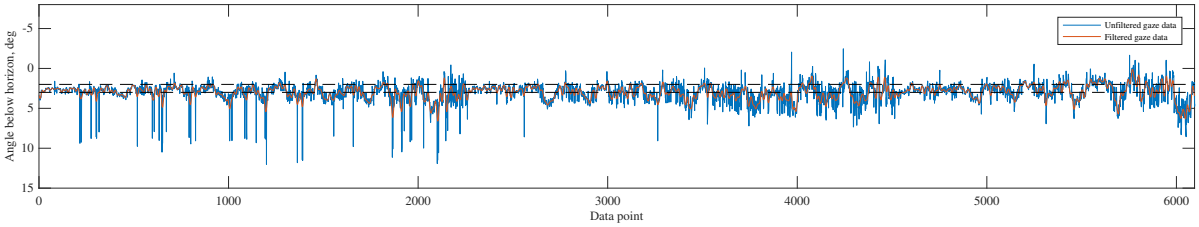

Figure D.52: Eye gaze timetrace, subject 3, Condition S23

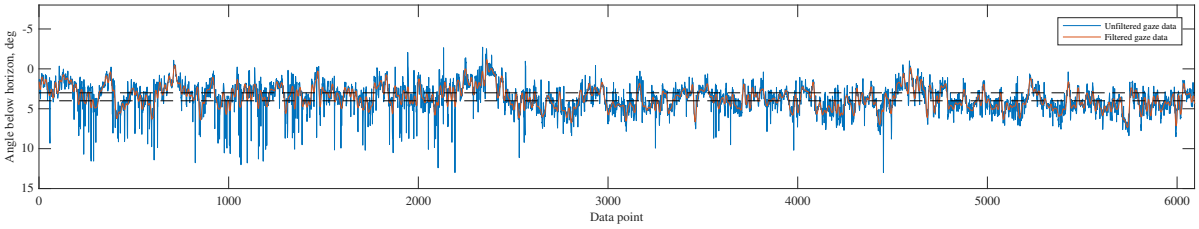

Figure D.53: Eye gaze timetrace, subject 3, Condition S34

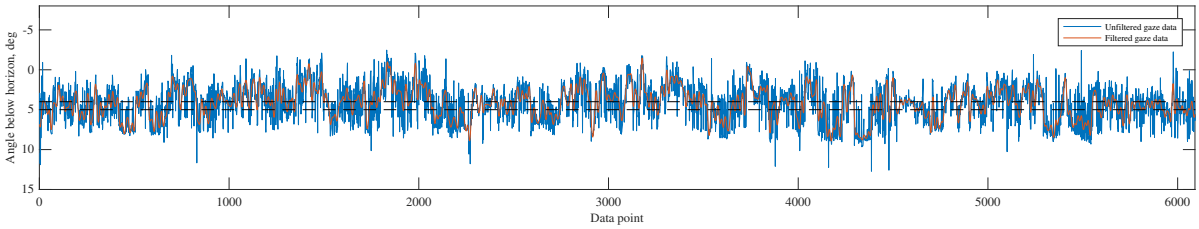

Figure D.54: Eye gaze timetrace, subject 3, Condition S45

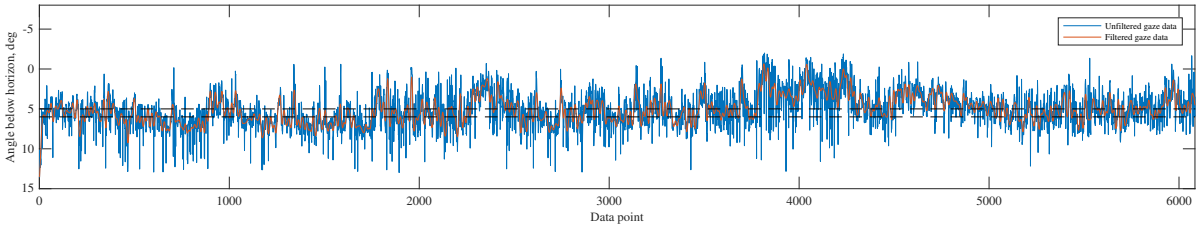

Figure D.55: Eye gaze timetrace, subject 3, Condition S56

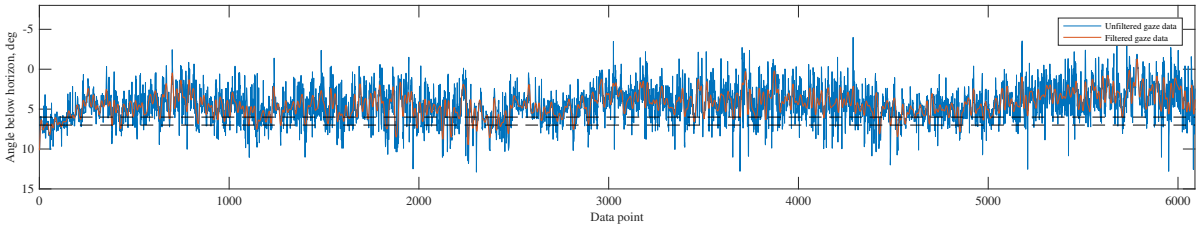

Figure D.56: Eye gaze timetrace, subject 3, Condition S67

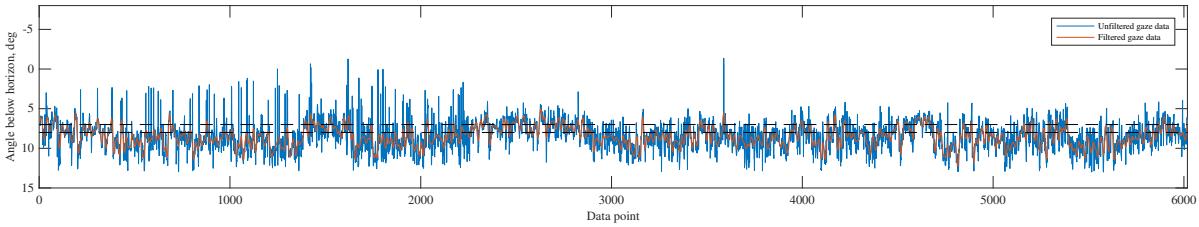

Figure D.57: Eye gaze timetrace, subject 3, Condition S78

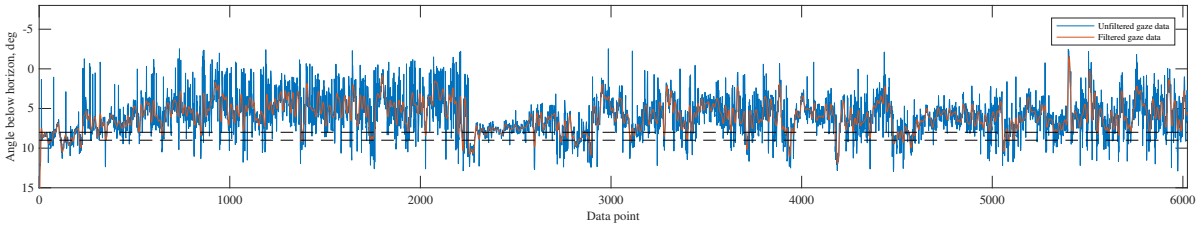

Figure D.58: Eye gaze timetrace, subject 3, Condition S89

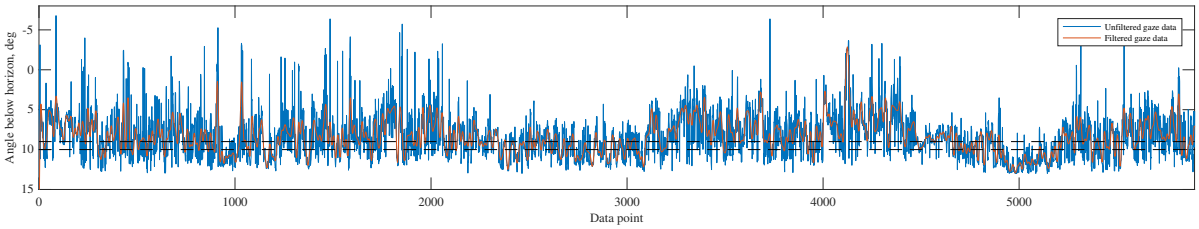

Figure D.59: Eye gaze timetrace, subject 3, Condition S910

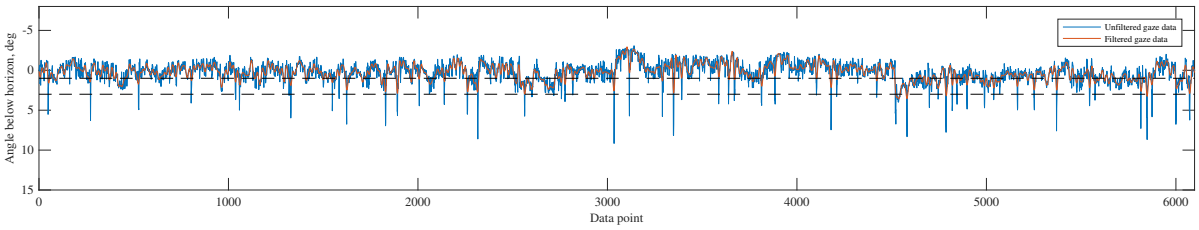

Figure D.60: Eye gaze timetrace, subject 3, Condition S13

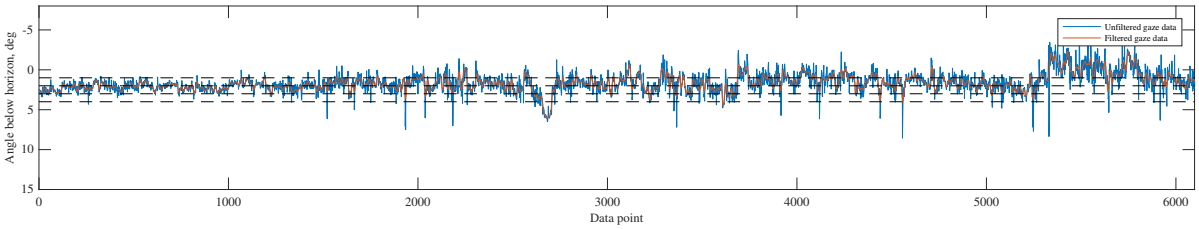

Figure D.61: Eye gaze timetrace, subject 3, Condition T34

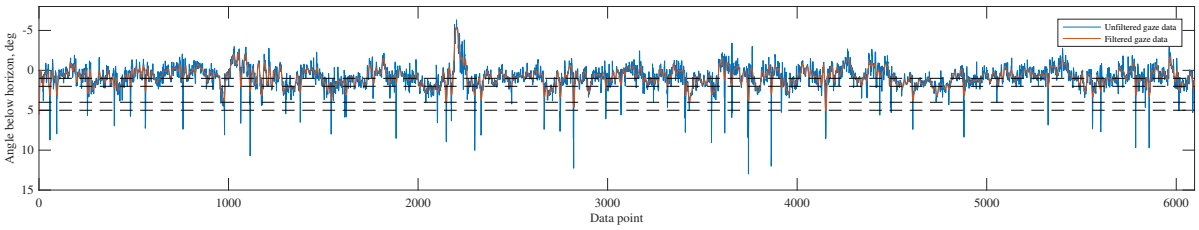

Figure D.62: Eye gaze timetrace, subject 3, Condition T45

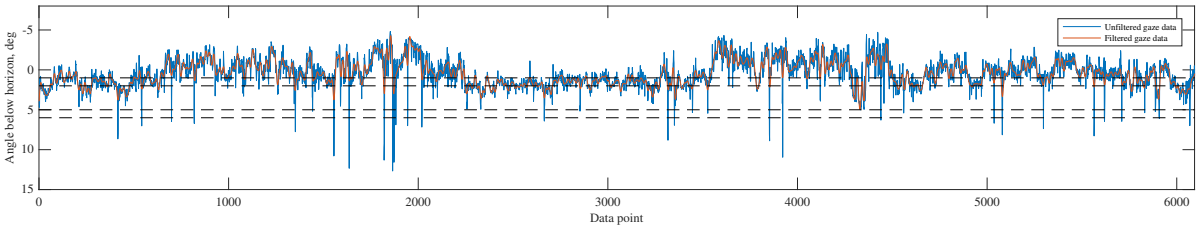

Figure D.63: Eye gaze timetrace, subject 3, Condition T56

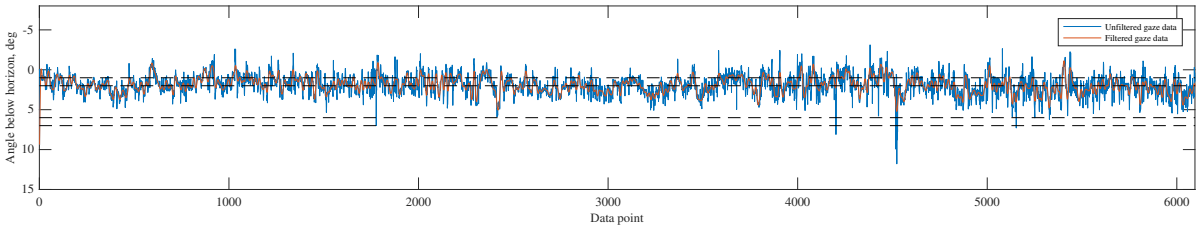

Figure D.64: Eye gaze timetrace, subject 3, Condition T67

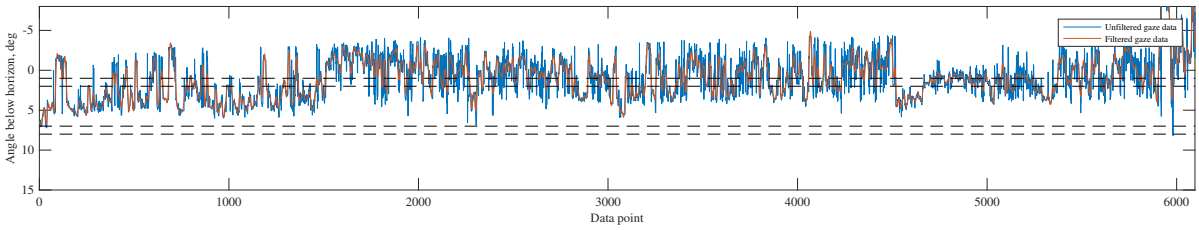

Figure D.65: Eye gaze timetrace, subject 3, Condition T78

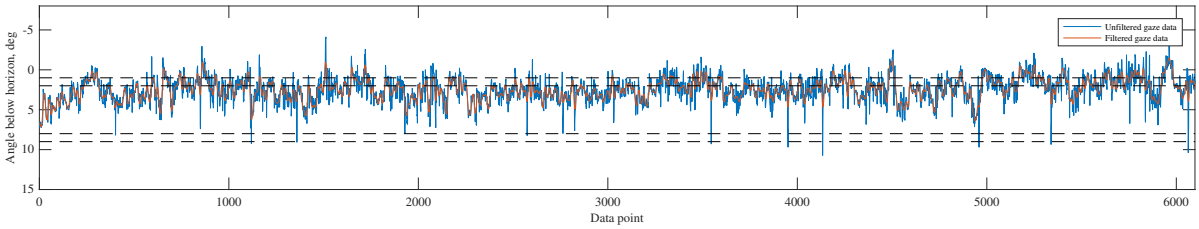

Figure D.66: Eye gaze timetrace, subject 3, Condition T89

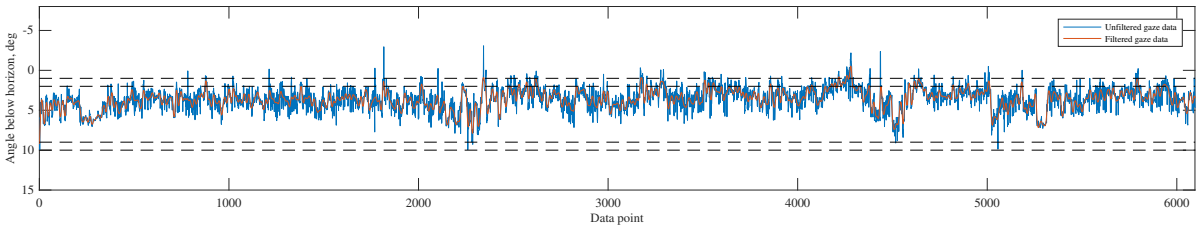

Figure D.67: Eye gaze timetrace, subject 3, Condition T910

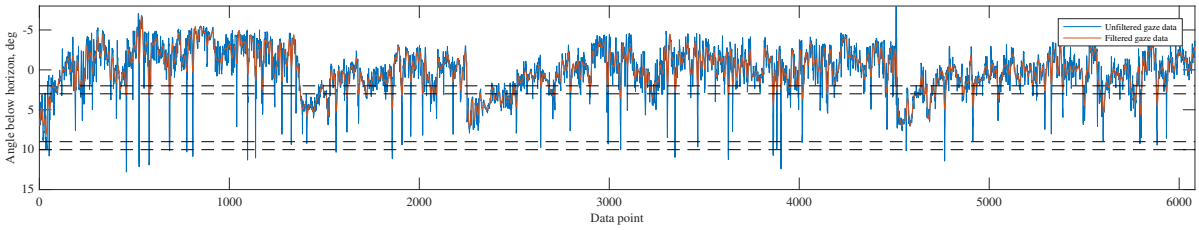

Figure D.68: Eye gaze timetrace, subject 3, Condition B23

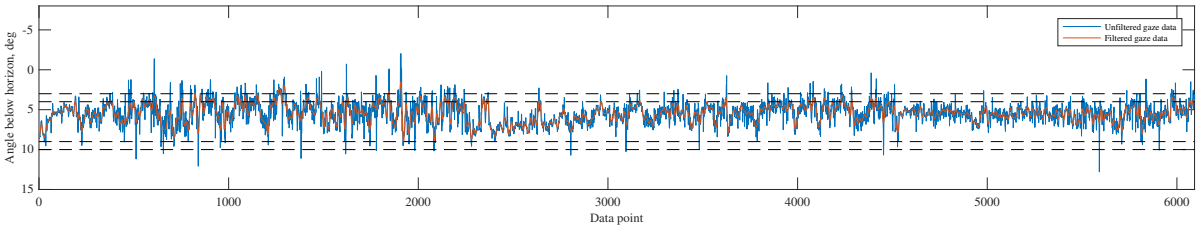

Figure D.69: Eye gaze timetrace, subject 3, Condition B34

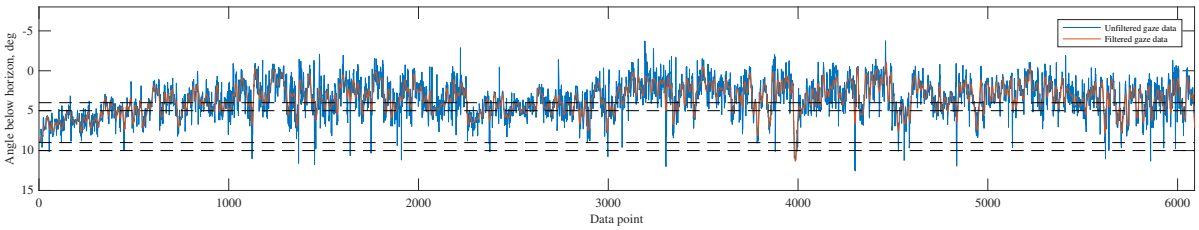

Figure D.70: Eye gaze timetrace, subject 3, Condition B45

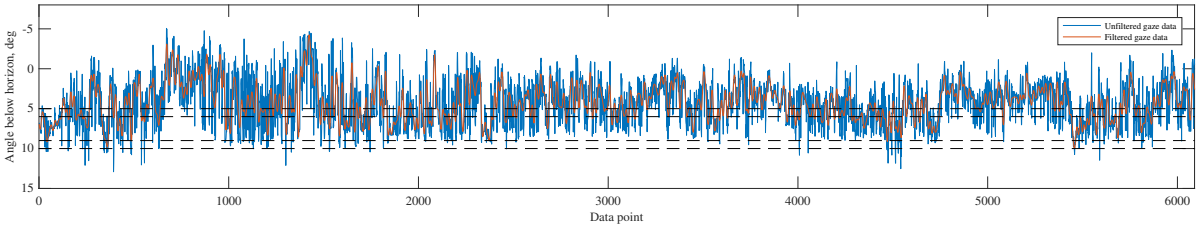

Figure D.71: Eye gaze timetrace, subject 3, Condition B56

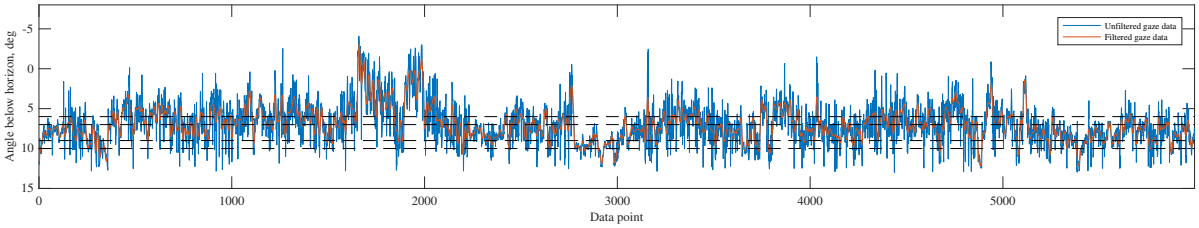

Figure D.72: Eye gaze timetrace, subject 3, Condition B67

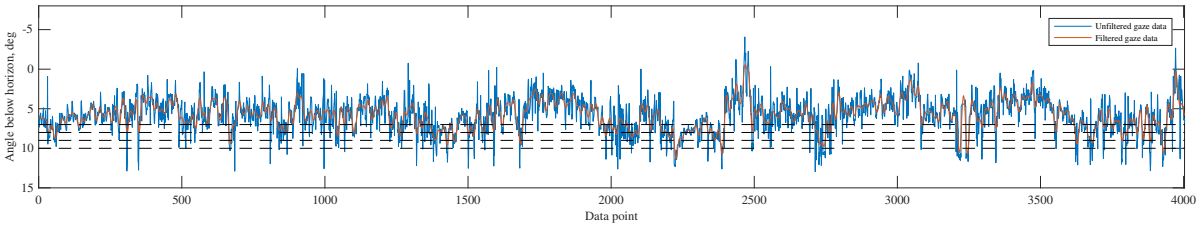

Figure D.73: Eye gaze timetrace, subject 3, Condition B78

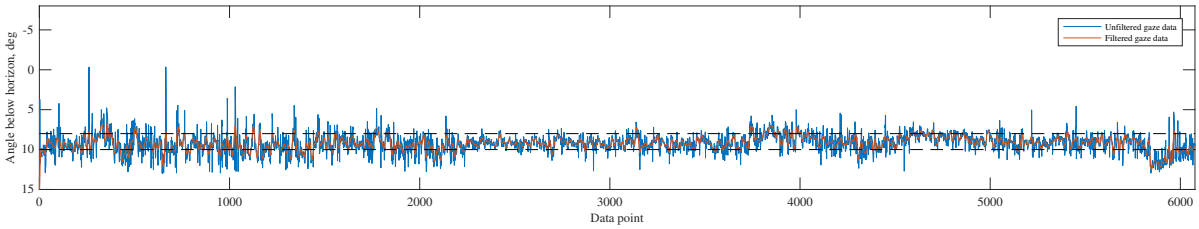

Figure D.74: Eye gaze timetrace, subject 3, Condition S810

Subject 6

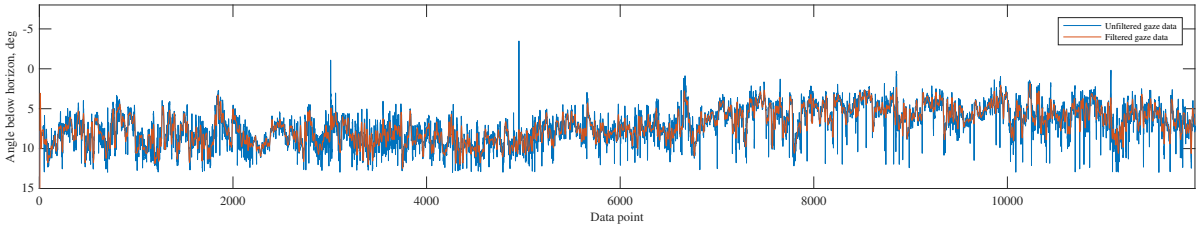

Figure D.75: Eye gaze timetrace, subject 7, Condition NO

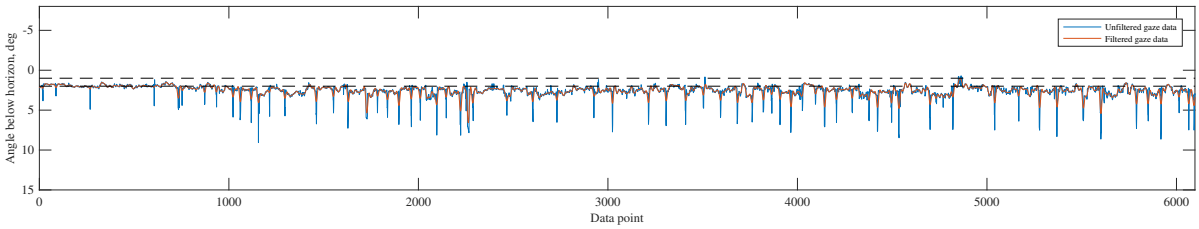

Figure D.76: Eye gaze timetrace, subject 7, Condition S12

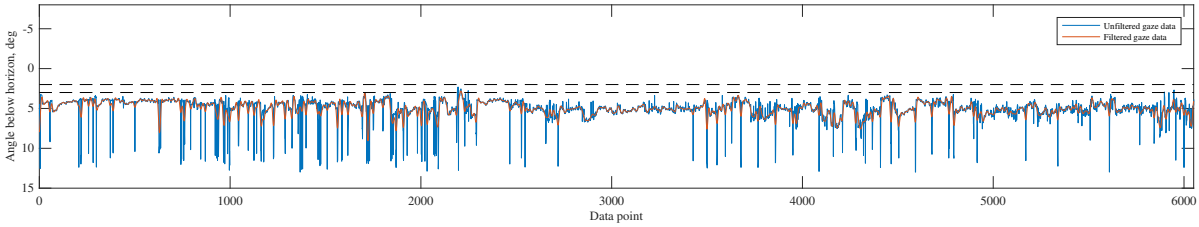

Figure D.77: Eye gaze timetrace, subject 7, Condition S23

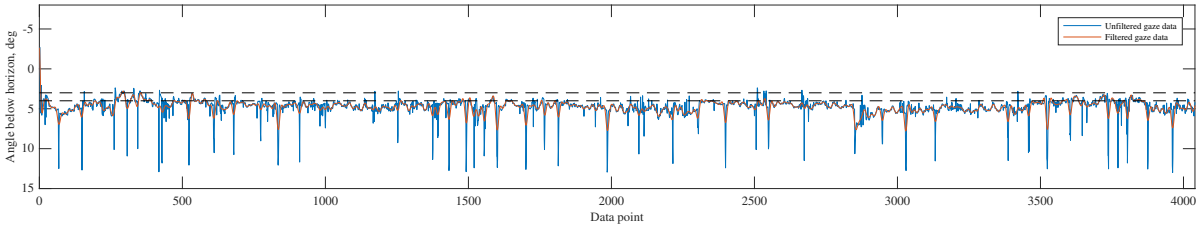

Figure D.78: Eye gaze timetrace, subject 7, Condition S34

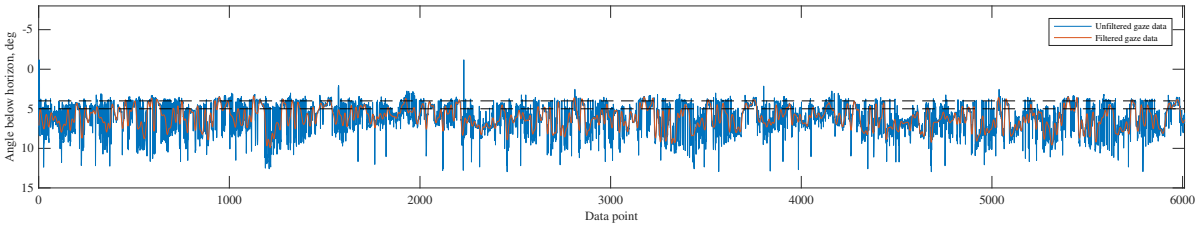

Figure D.79: Eye gaze timetrace, subject 7, Condition S45

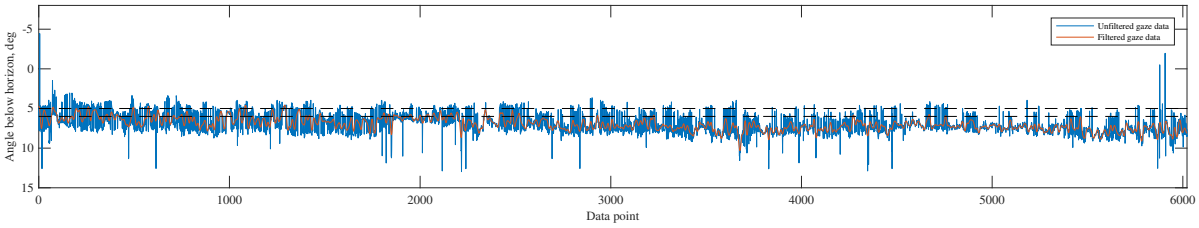

Figure D.80: Eye gaze timetrace, subject 7, Condition S56

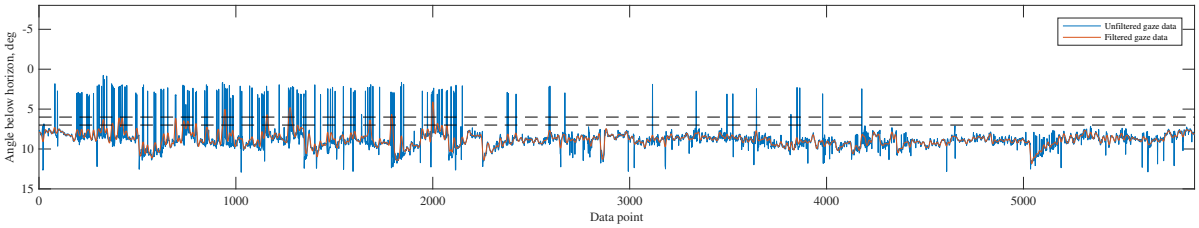

Figure D.81: Eye gaze timetrace, subject 7, Condition S67

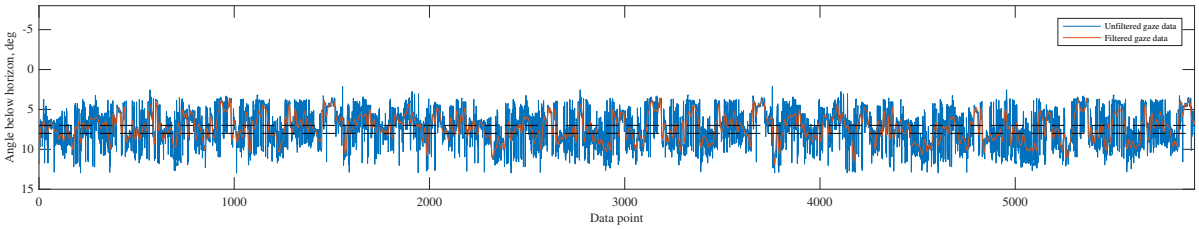

Figure D.82: Eye gaze timetrace, subject 7, Condition S78

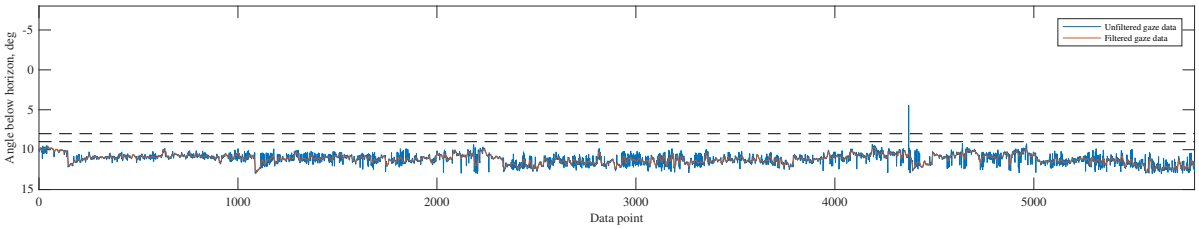

Figure D.83: Eye gaze timetrace, subject 7, Condition S89

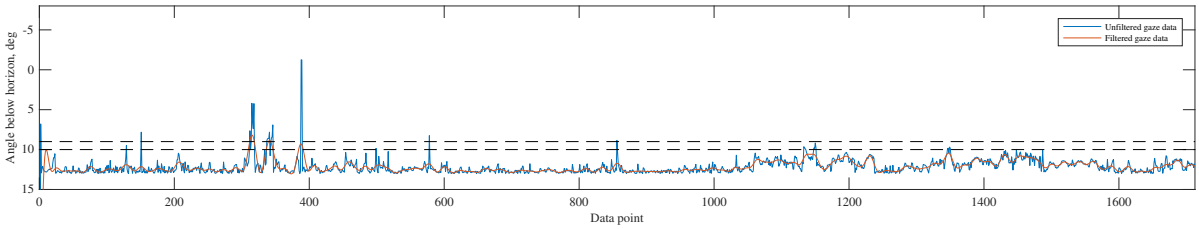

Figure D.84: Eye gaze timetrace, subject 7, Condition S910

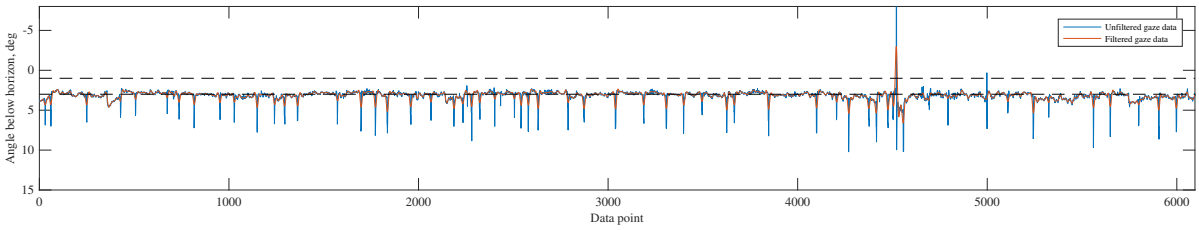

Figure D.85: Eye gaze timetrace, subject 7, Condition S13

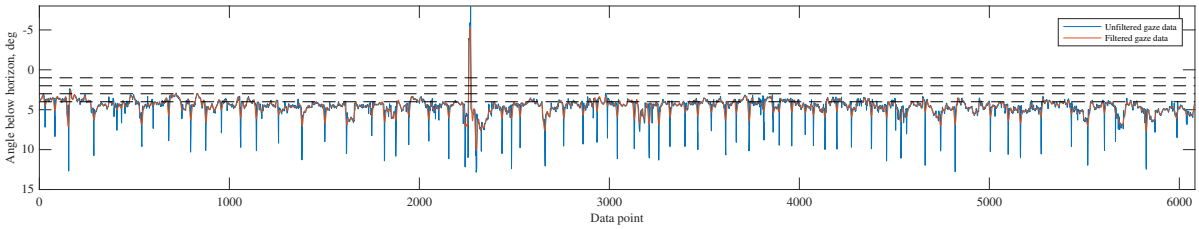

Figure D.86: Eye gaze timetrace, subject 7, Condition T34

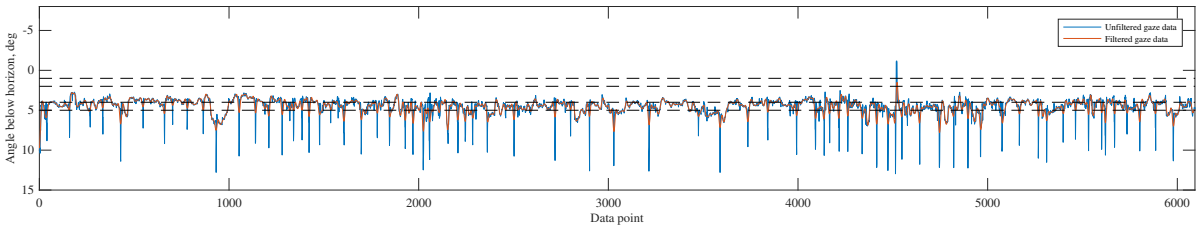

Figure D.87: Eye gaze timetrace, subject 7, Condition T45

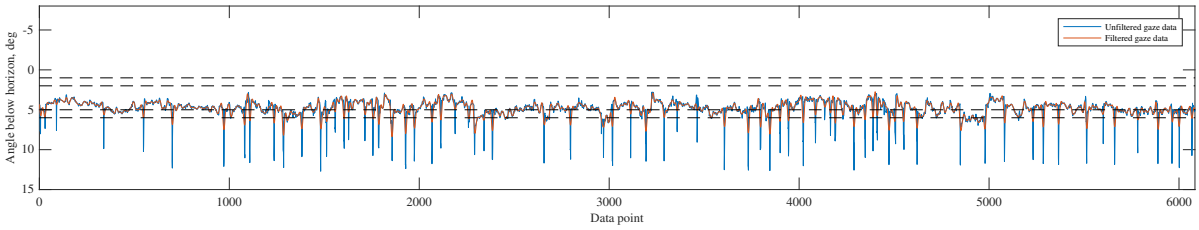

Figure D.88: Eye gaze timetrace, subject 7, Condition T56

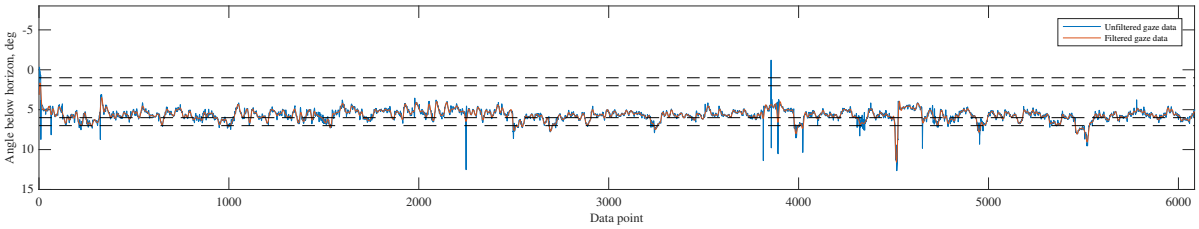

Figure D.89: Eye gaze timetrace, subject 7, Condition T67

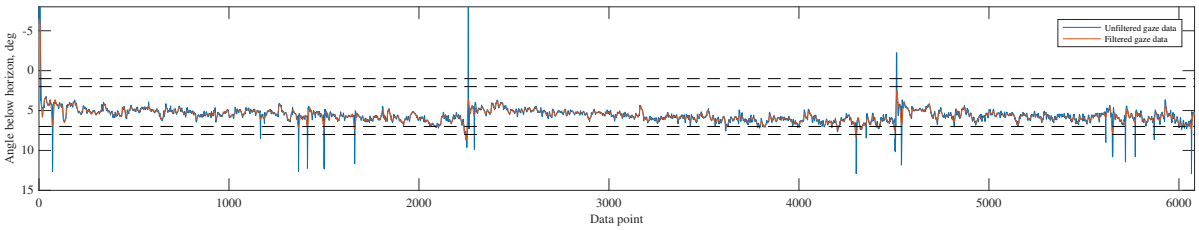

Figure D.90: Eye gaze timetrace, subject 7, Condition T78

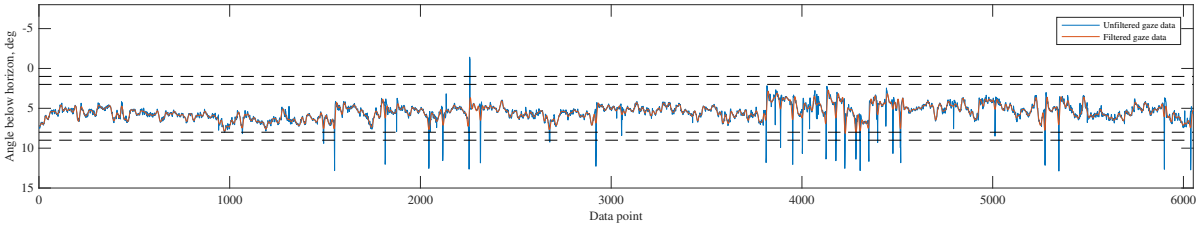

Figure D.91: Eye gaze timetrace, subject 7, Condition T89

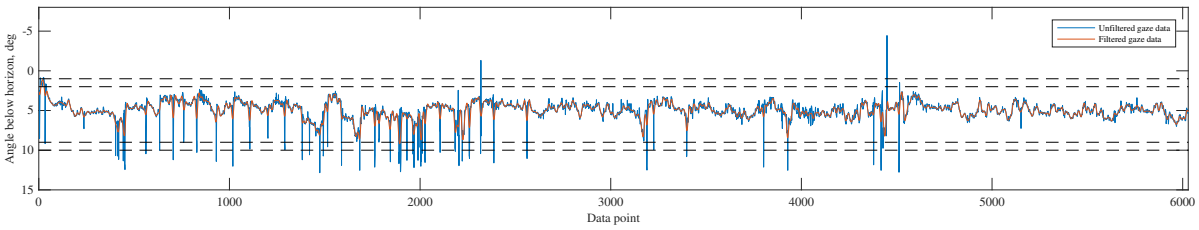

Figure D.92: Eye gaze timetrace, subject 7, Condition T910

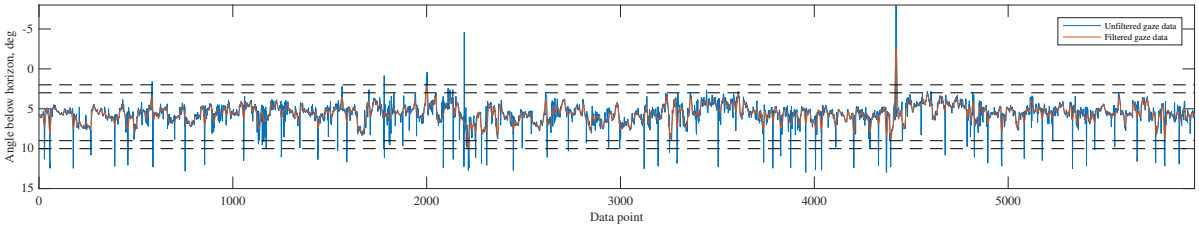

Figure D.93: Eye gaze timetrace, subject 7, Condition B23

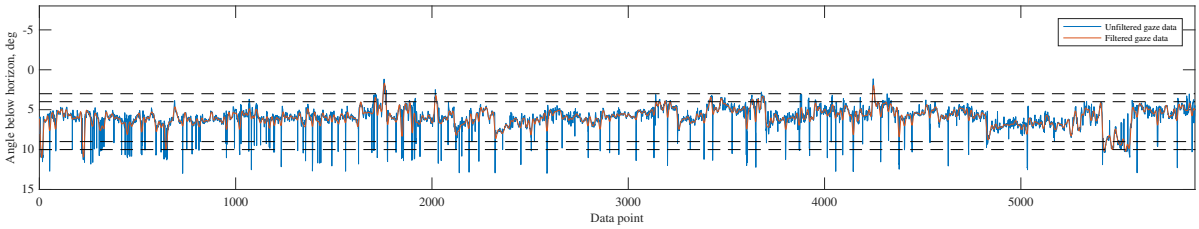

Figure D.94: Eye gaze timetrace, subject 7, Condition B34

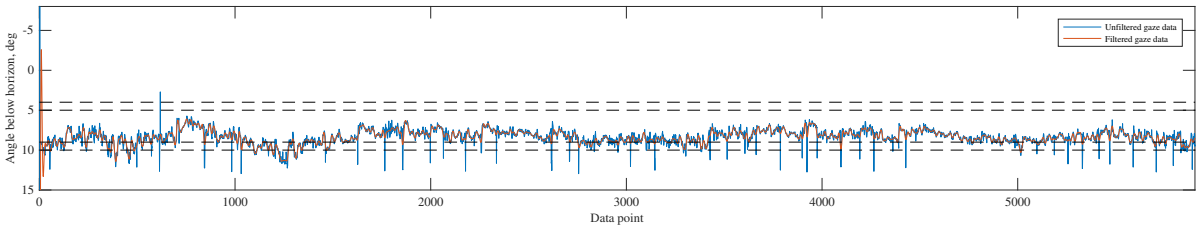

Figure D.95: Eye gaze timetrace, subject 7, Condition B45

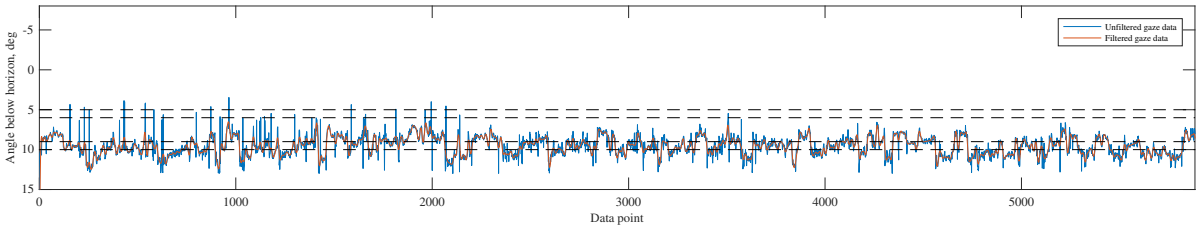

Figure D.96: Eye gaze timetrace, subject 7, Condition B56

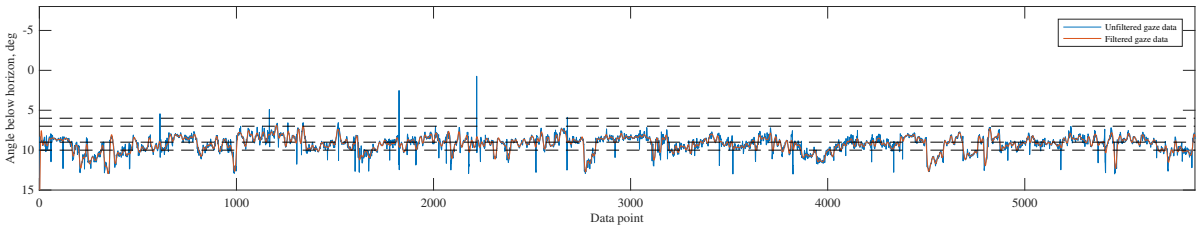

Figure D.97: Eye gaze timetrace, subject 7, Condition B67

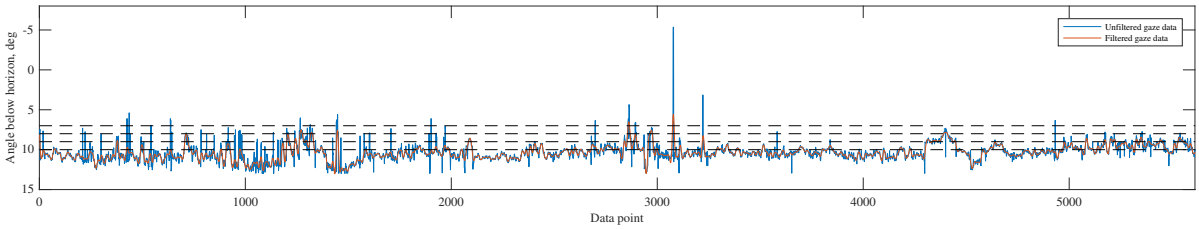

Figure D.98: Eye gaze timetrace, subject 7, Condition B78

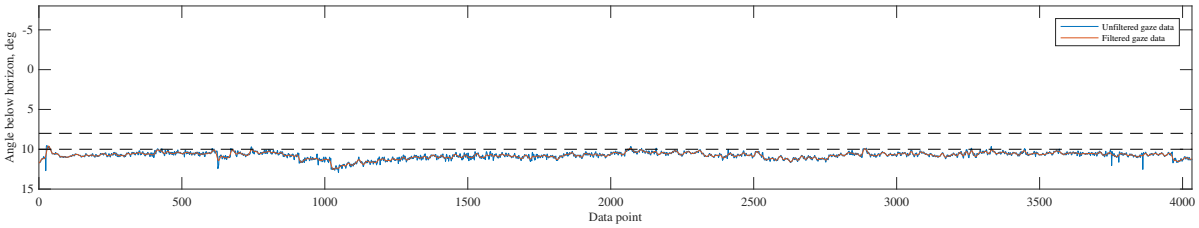

Figure D.99: Eye gaze timetrace, subject 7, Condition S810

Subject 7

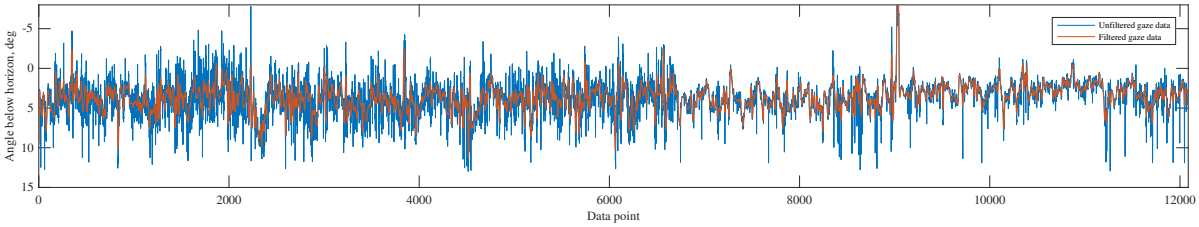

Figure D.100: Eye gaze timetrace, subject 8, Condition NO

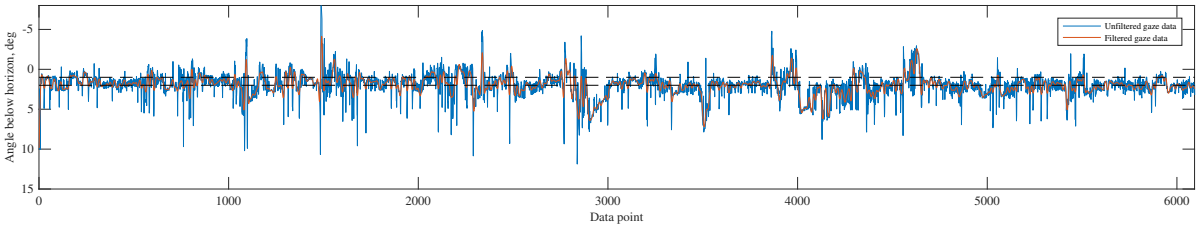

Figure D.101: Eye gaze timetrace, subject 8, Condition S12

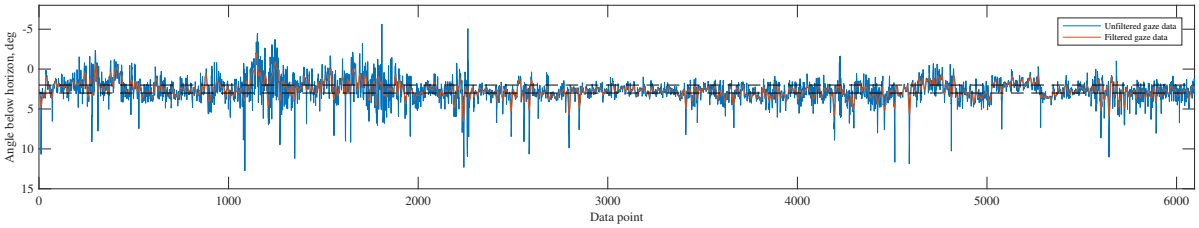

Figure D.102: Eye gaze timetrace, subject 8, Condition S23

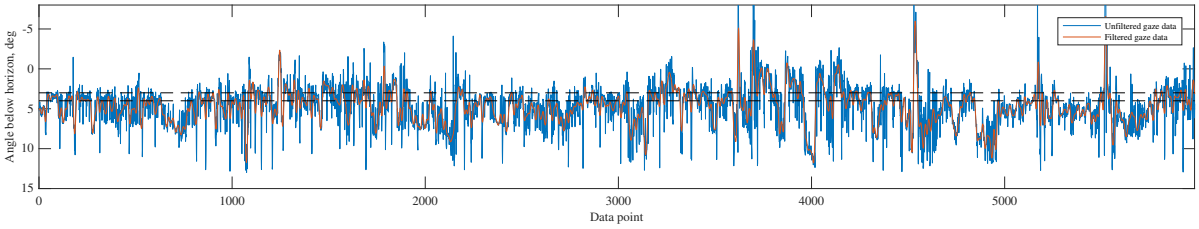

Figure D.103: Eye gaze timetrace, subject 8, Condition S34

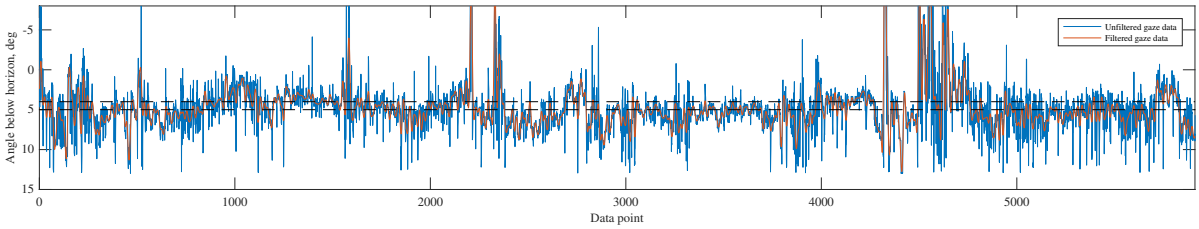

Figure D.104: Eye gaze timetrace, subject 8, Condition S45

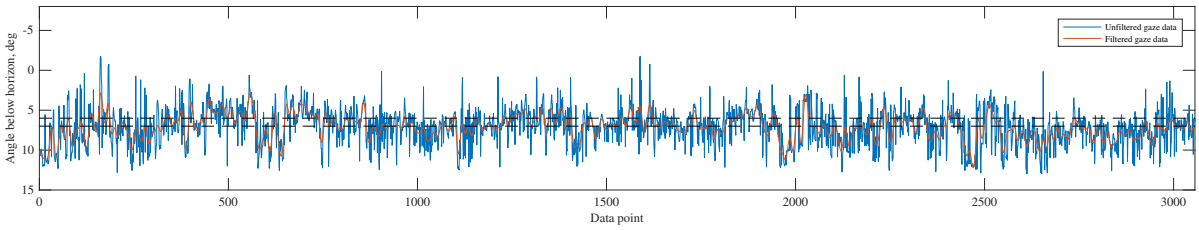

Figure D.105: Eye gaze timetrace, subject 8, Condition S67

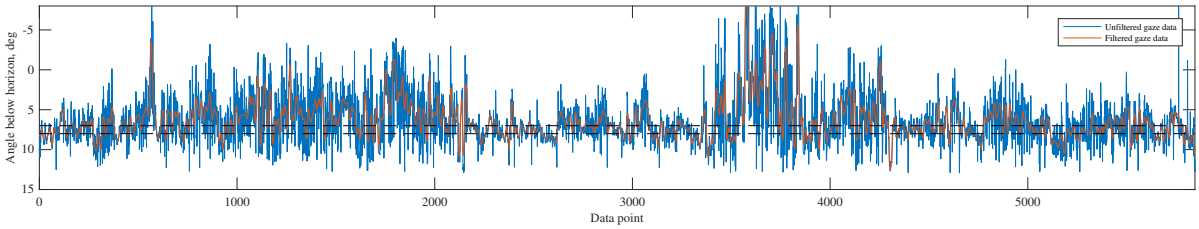

Figure D.106: Eye gaze timetrace, subject 8, Condition S78

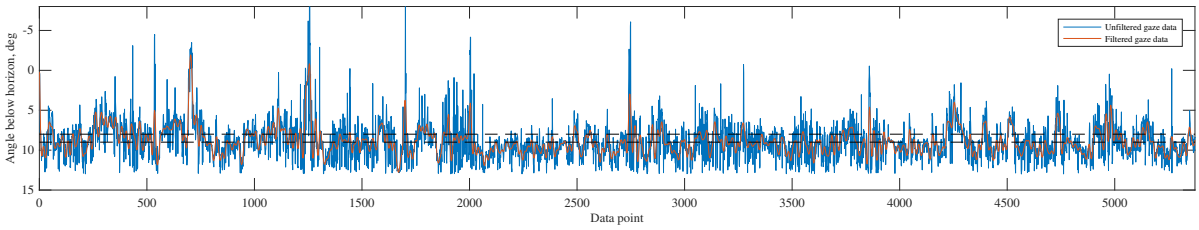

Figure D.107: Eye gaze timetrace, subject 8, Condition S89

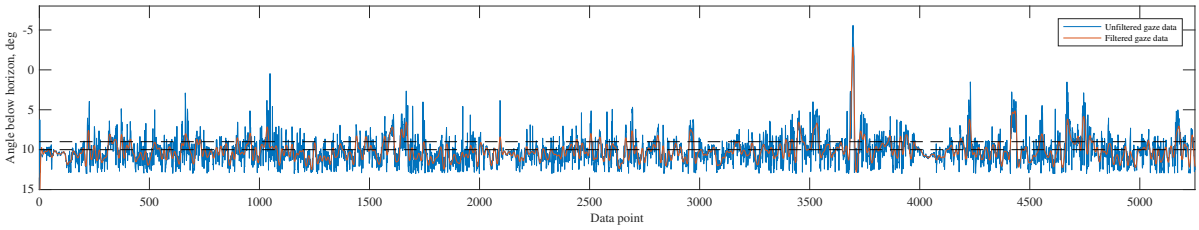

Figure D.108: Eye gaze timetrace, subject 8, Condition S910

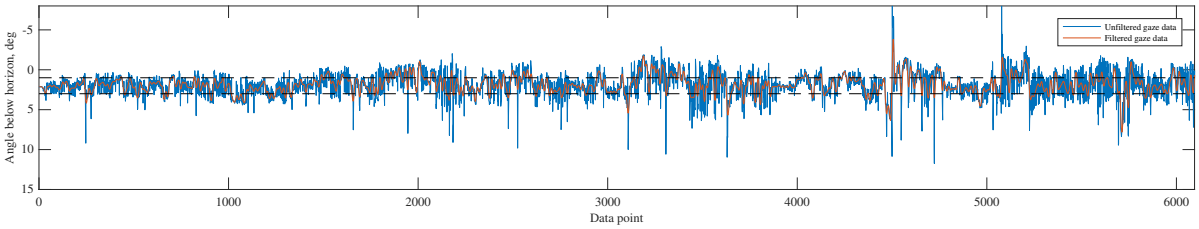

Figure D.109: Eye gaze timetrace, subject 8, Condition S13

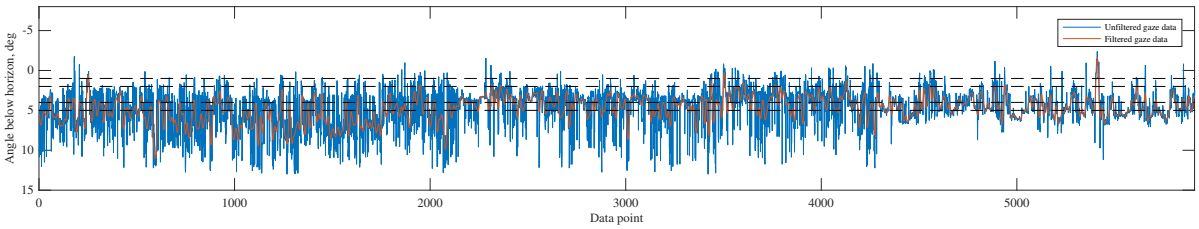

Figure D.110: Eye gaze timetrace, subject 8, Condition T45

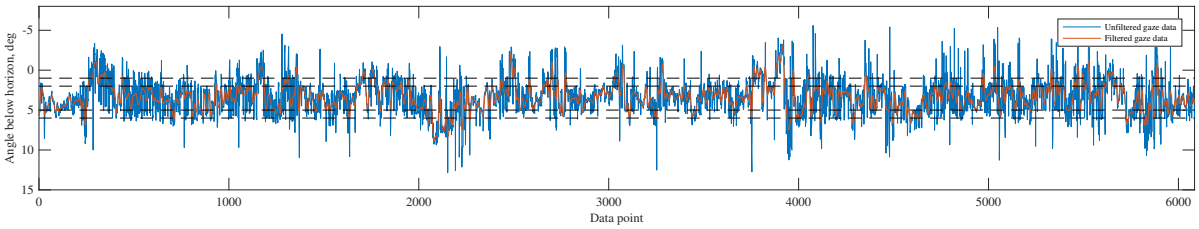

Figure D.111: Eye gaze timetrace, subject 8, Condition T56

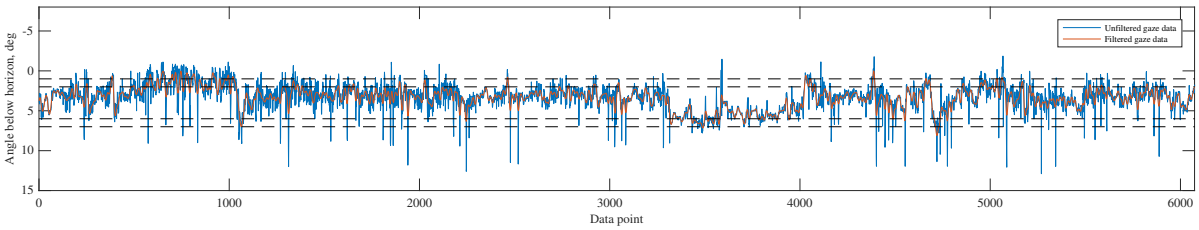

Figure D.112: Eye gaze timetrace, subject 8, Condition T67

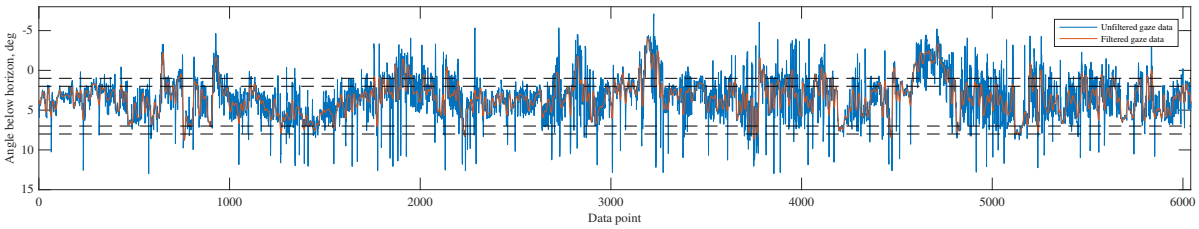

Figure D.113: Eye gaze timetrace, subject 8, Condition T78

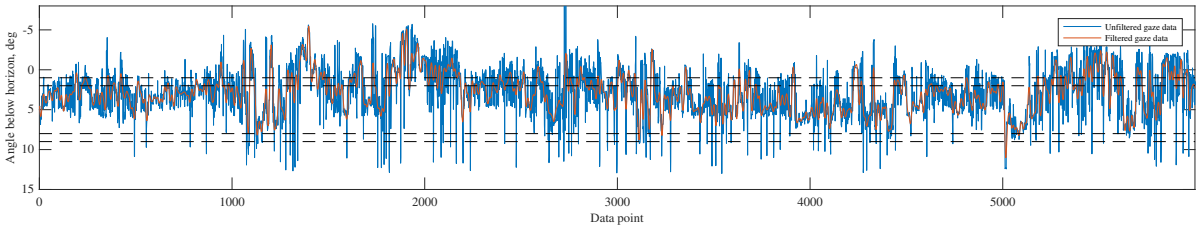

Figure D.114: Eye gaze timetrace, subject 8, Condition T89

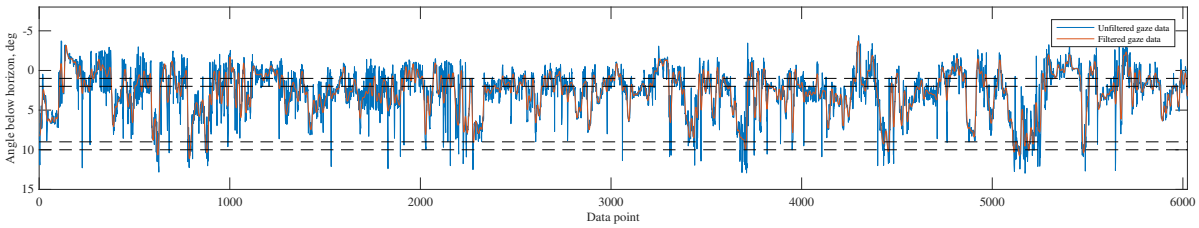

Figure D.115: Eye gaze timetrace, subject 8, Condition T910

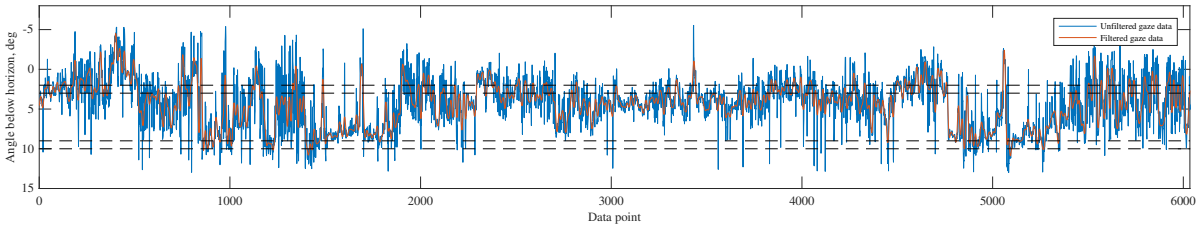

Figure D.116: Eye gaze timetrace, subject 8, Condition B23

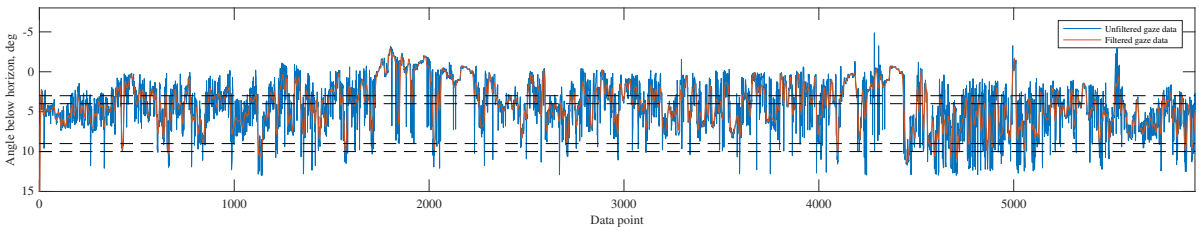

Figure D.117: Eye gaze timetrace, subject 8, Condition B34

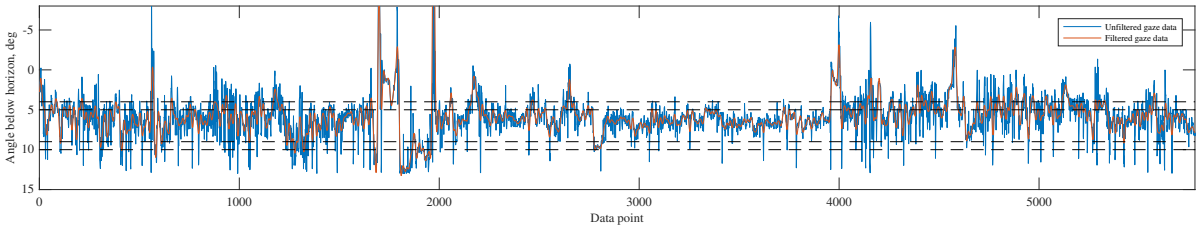

Figure D.118: Eye gaze timetrace, subject 8, Condition B45

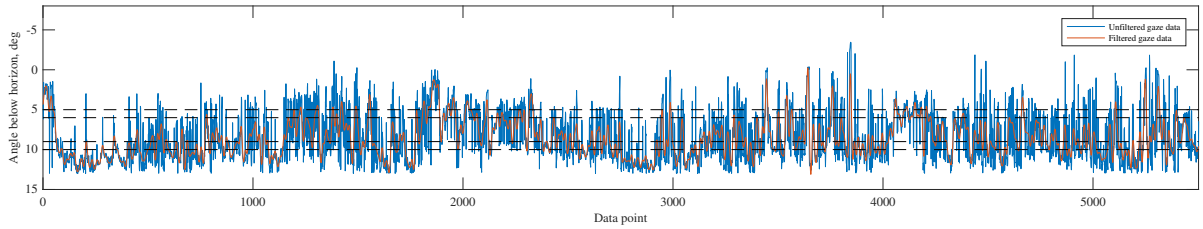

Figure D.119: Eye gaze timetrace, subject 8, Condition B56

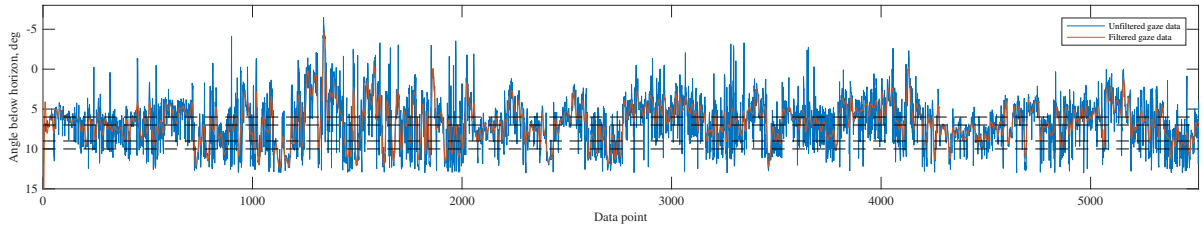

Figure D.120: Eye gaze timetrace, subject 8, Condition B67

Subject 8

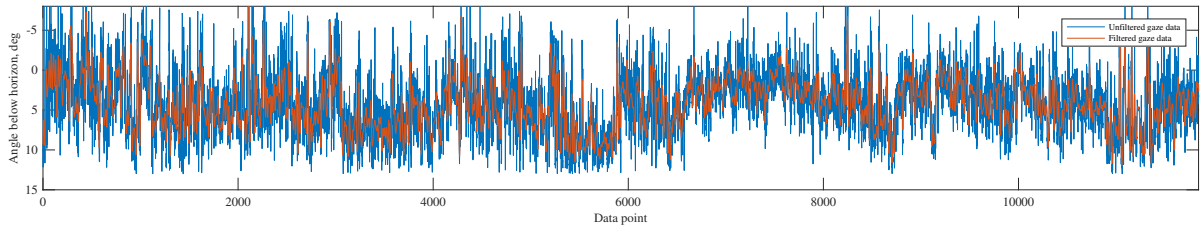

Figure D.121: Eye gaze timetrace, subject 9, Condition NO

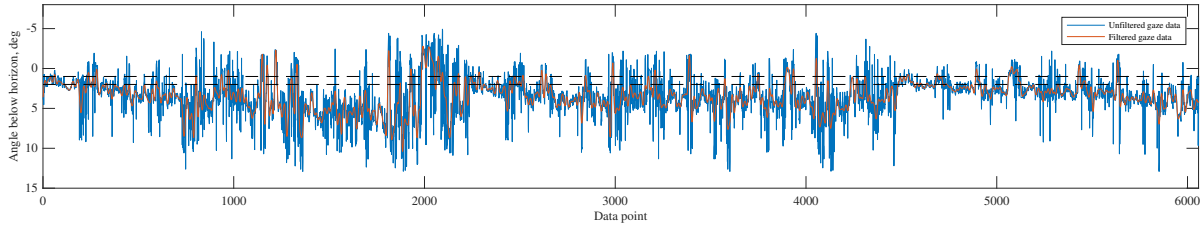

Figure D.122: Eye gaze timetrace, subject 9, Condition S12

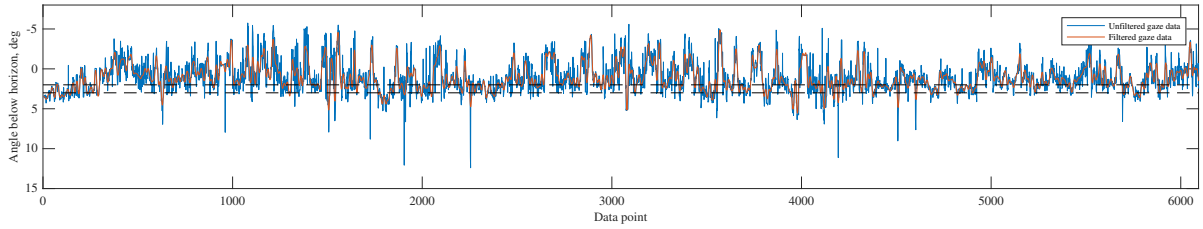

Figure D.123: Eye gaze timetrace, subject 9, Condition S23

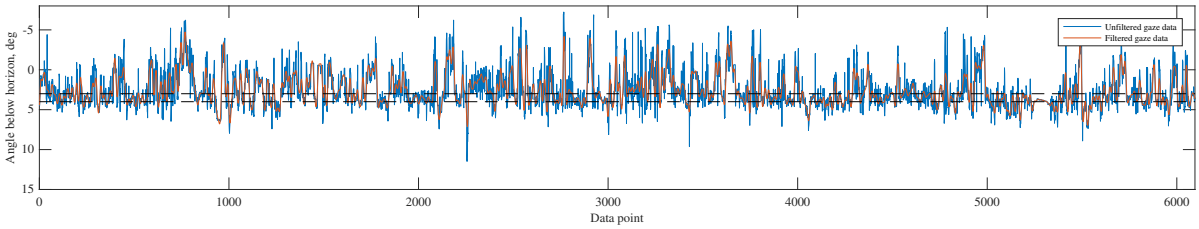

Figure D.124: Eye gaze timetrace, subject 9, Condition S34

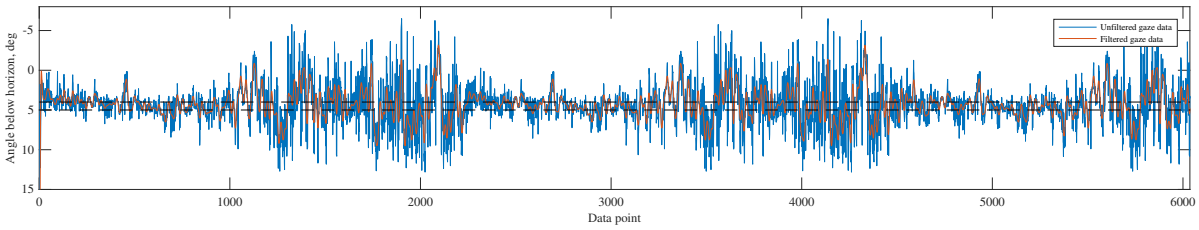

Figure D.125: Eye gaze timetrace, subject 9, Condition S45

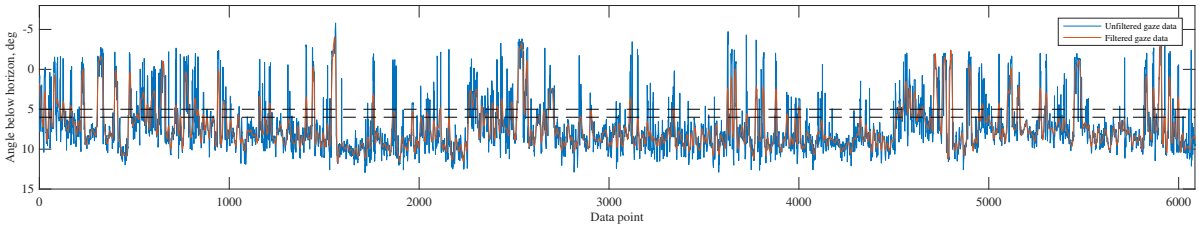

Figure D.126: Eye gaze timetrace, subject 9, Condition S56

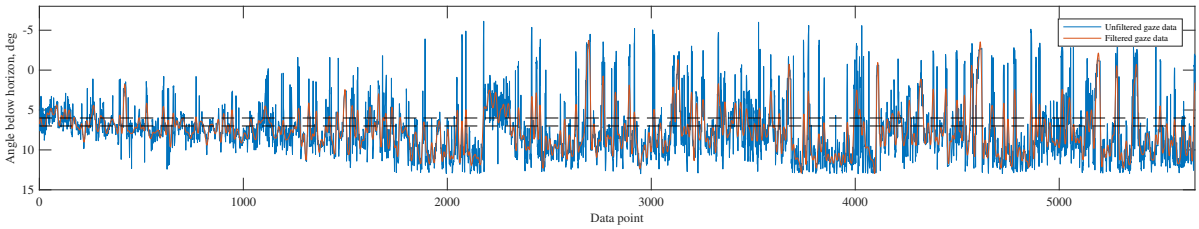

Figure D.127: Eye gaze timetrace, subject 9, Condition S67

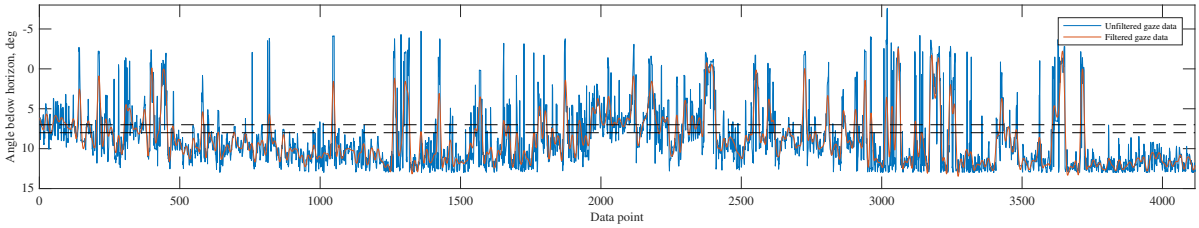

Figure D.128: Eye gaze timetrace, subject 9, Condition S78

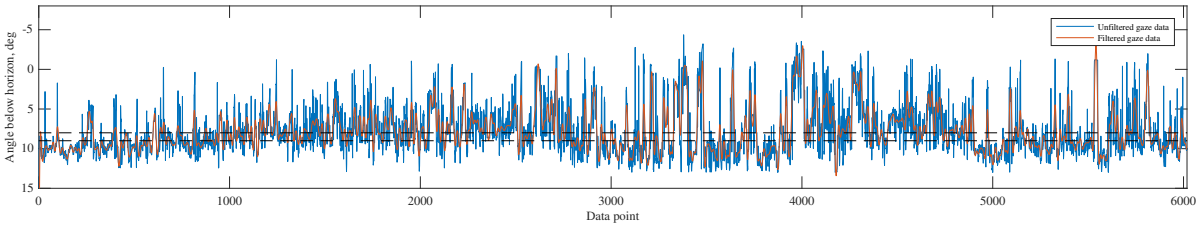

Figure D.129: Eye gaze timetrace, subject 9, Condition S89

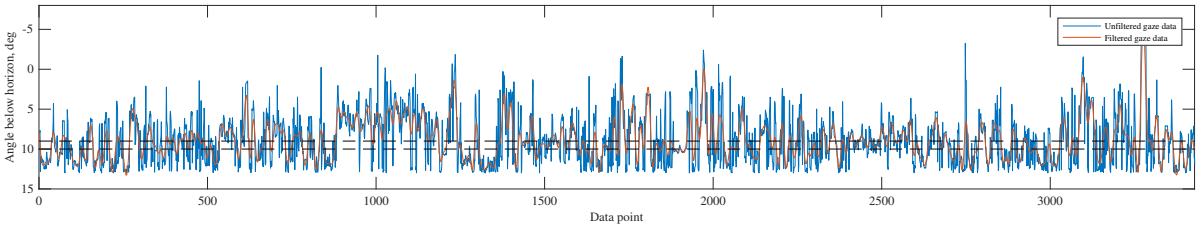

Figure D.130: Eye gaze timetrace, subject 9, Condition S910

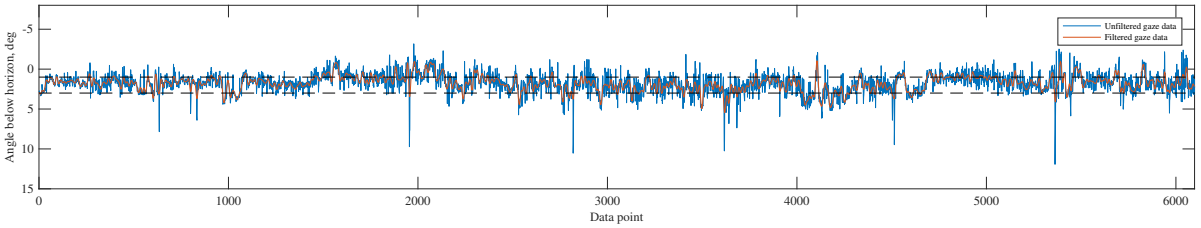

Figure D.131: Eye gaze timetrace, subject 9, Condition S13

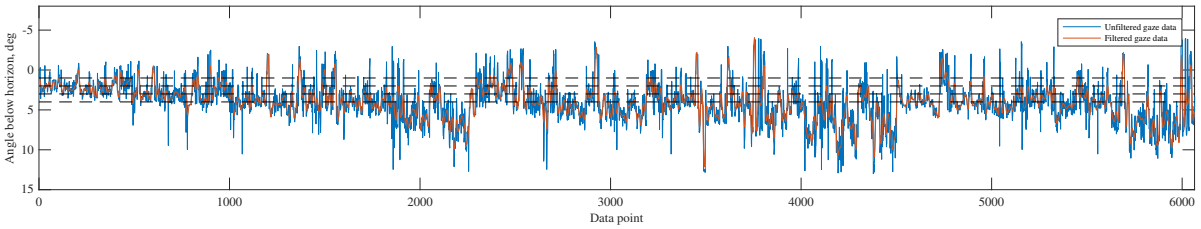

Figure D.132: Eye gaze timetrace, subject 9, Condition T34

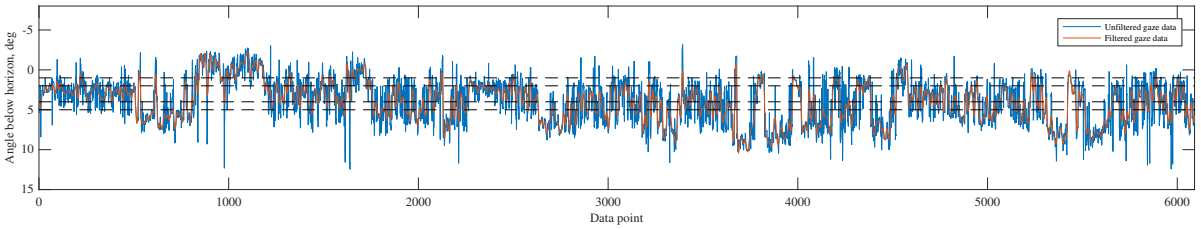

Figure D.133: Eye gaze timetrace, subject 9, Condition T45

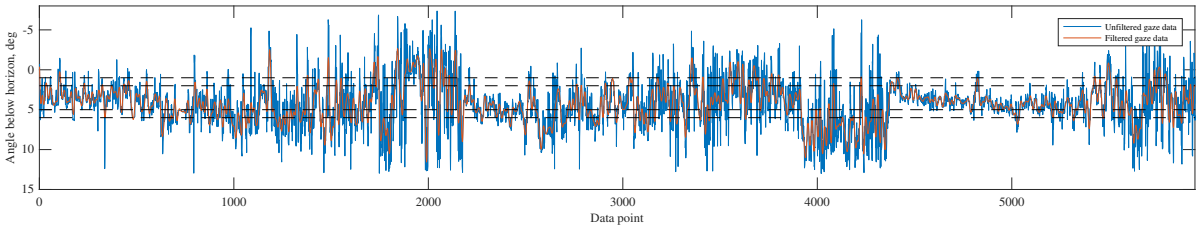

Figure D.134: Eye gaze timetrace, subject 9, Condition T56

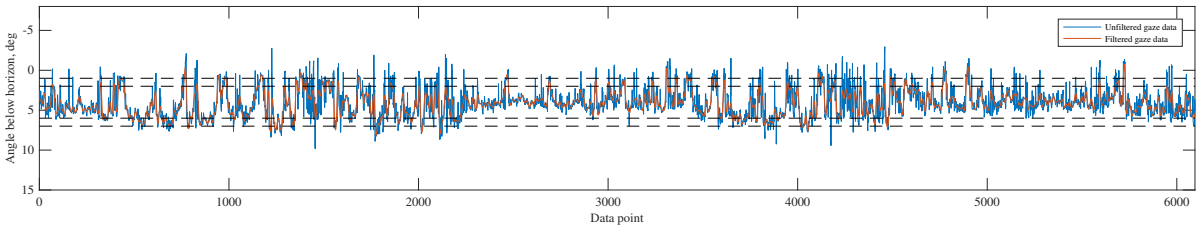

Figure D.135: Eye gaze timetrace, subject 9, Condition T67

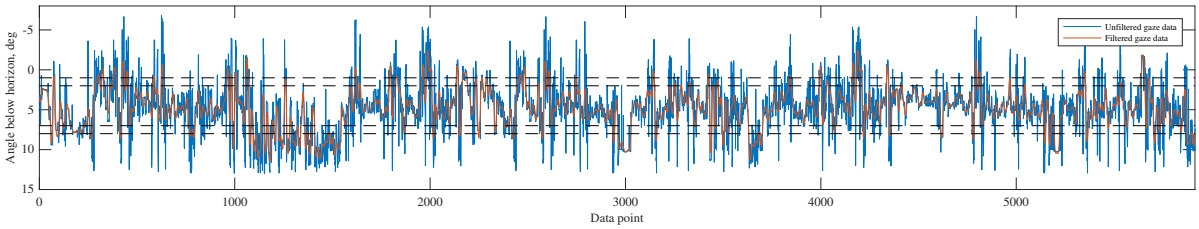

Figure D.136: Eye gaze timetrace, subject 9, Condition T78

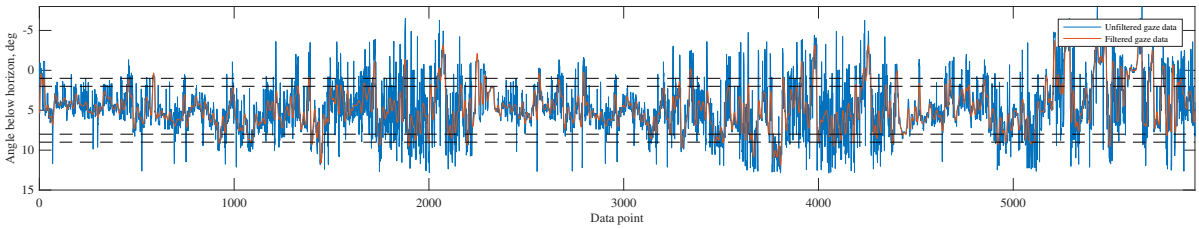

Figure D.137: Eye gaze timetrace, subject 9, Condition T89

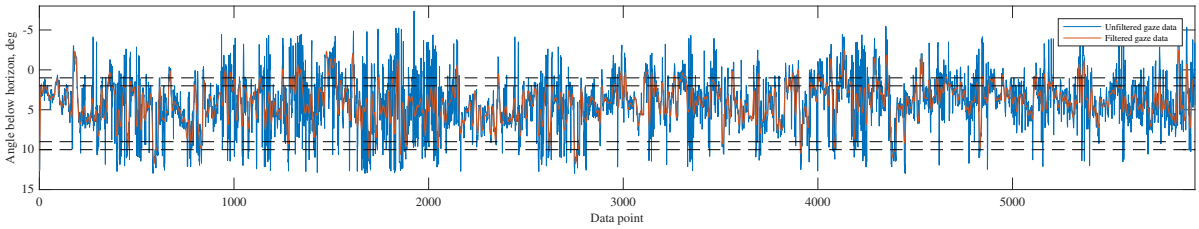

Figure D.138: Eye gaze timetrace, subject 9, Condition T910

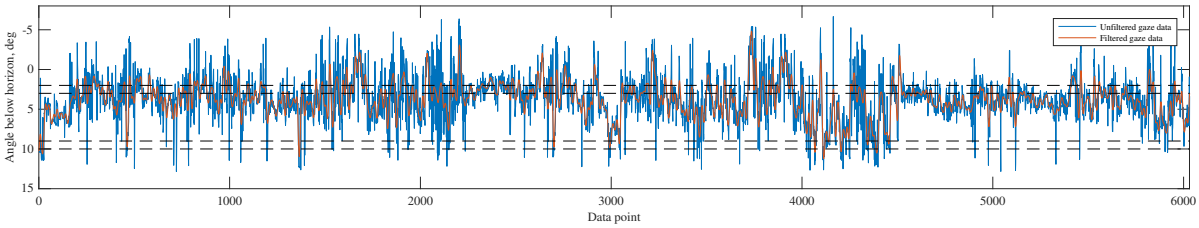

Figure D.139: Eye gaze timetrace, subject 9, Condition B23

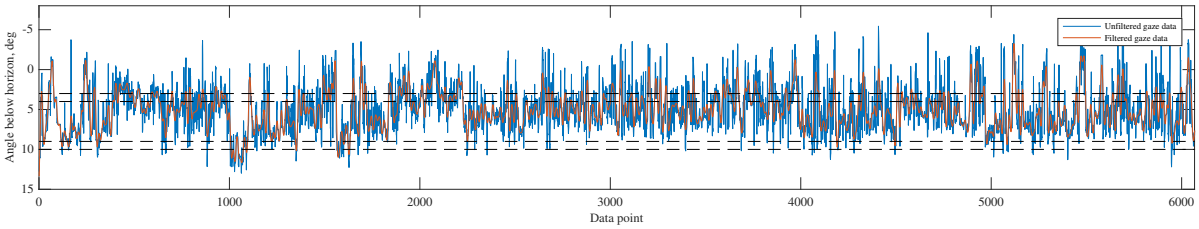

Figure D.140: Eye gaze timetrace, subject 9, Condition B34

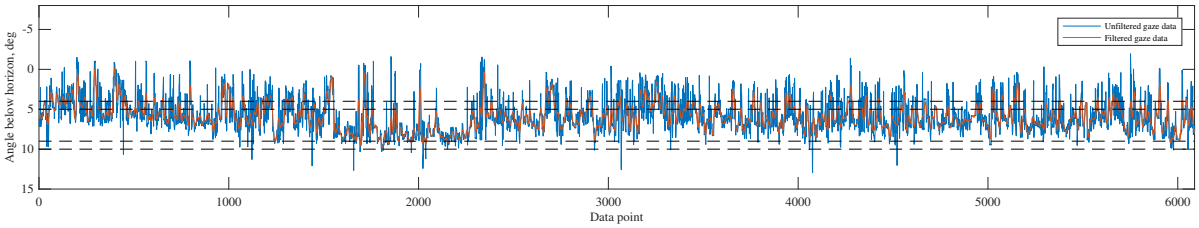

Figure D.141: Eye gaze timetrace, subject 9, Condition B45

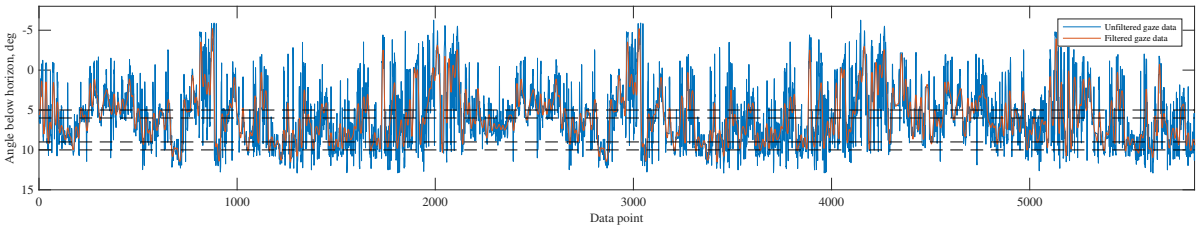

Figure D.142: Eye gaze timetrace, subject 9, Condition B56

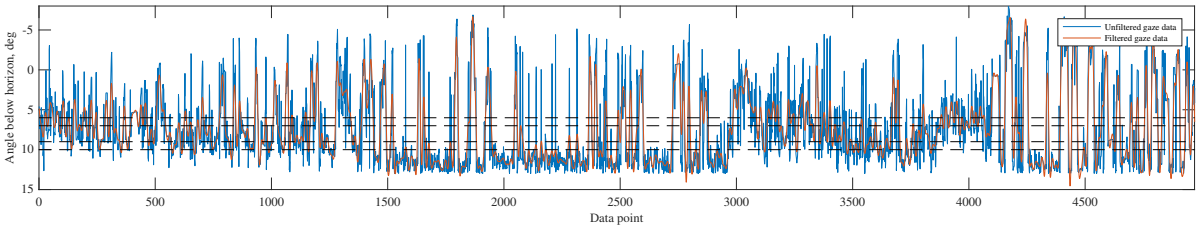

Figure D.143: Eye gaze timetrace, subject 9, Condition B67

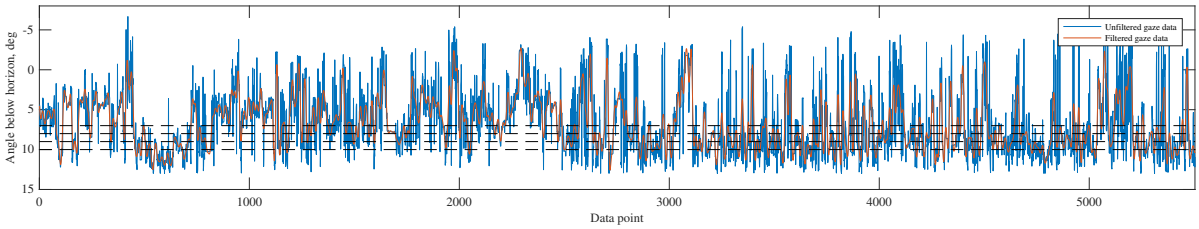

Figure D.144: Eye gaze timetrace, subject 9, Condition B78

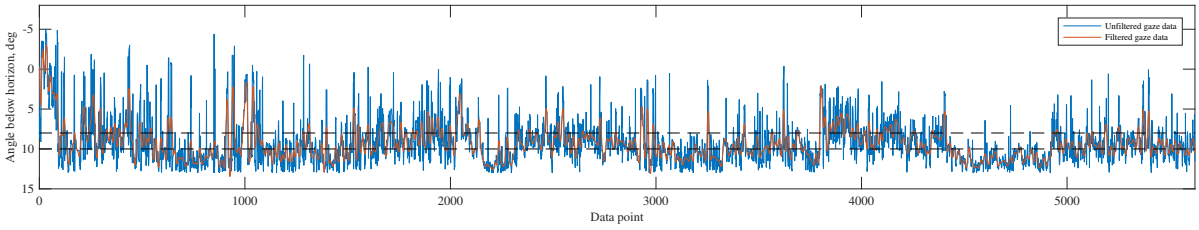

Figure D.145: Eye gaze timetrace, subject 9, Condition S810

Subject 9

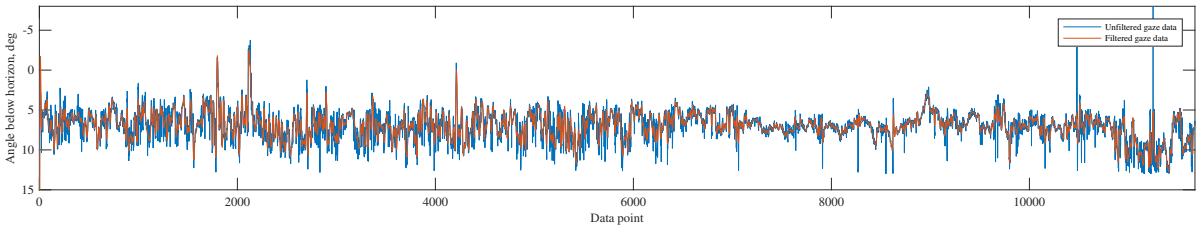

Figure D.146: Eye gaze timetrace, subject 10, Condition NO

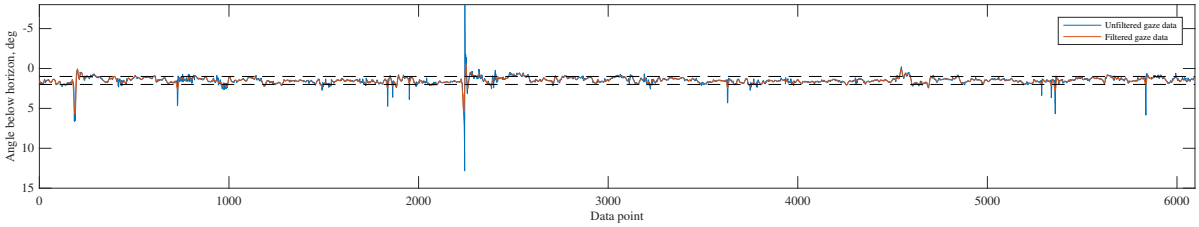

Figure D.147: Eye gaze timetrace, subject 10, Condition S12

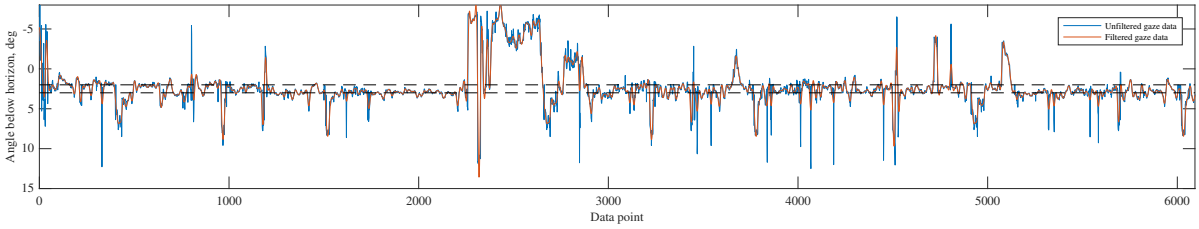

Figure D.148: Eye gaze timetrace, subject 10, Condition S23

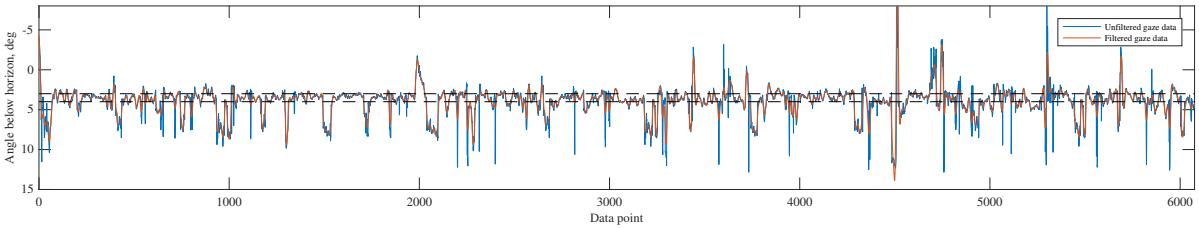

Figure D.149: Eye gaze timetrace, subject 10, Condition S34

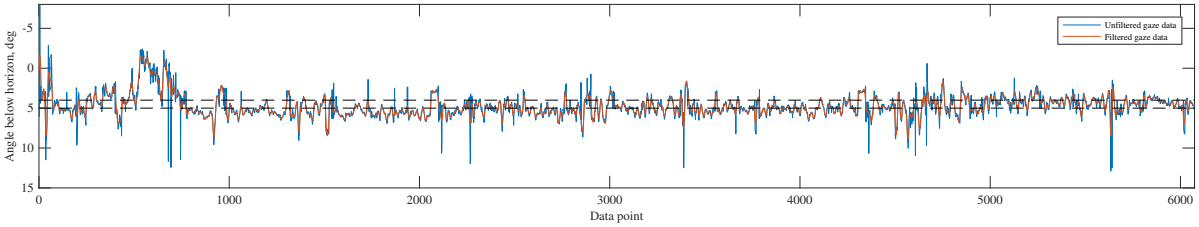

Figure D.150: Eye gaze timetrace, subject 10, Condition S45

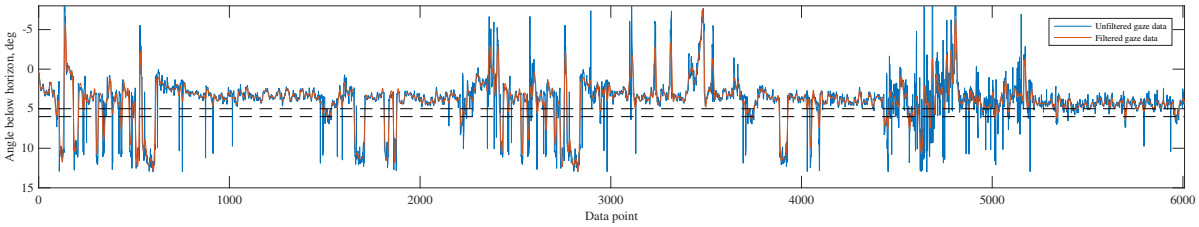

Figure D.151: Eye gaze timetrace, subject 10, Condition S56

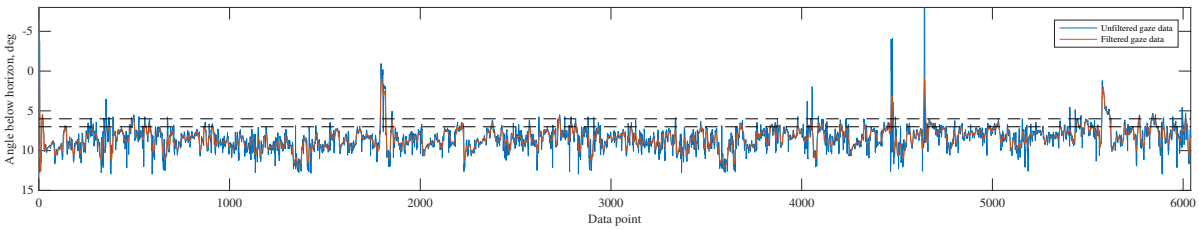

Figure D.152: Eye gaze timetrace, subject 10, Condition S67

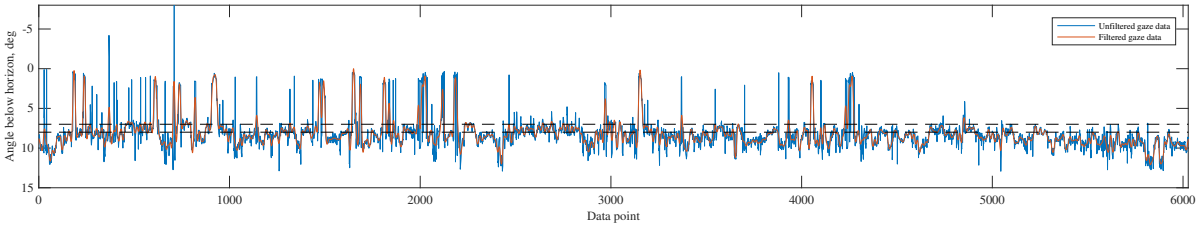

Figure D.153: Eye gaze timetrace, subject 10, Condition S78

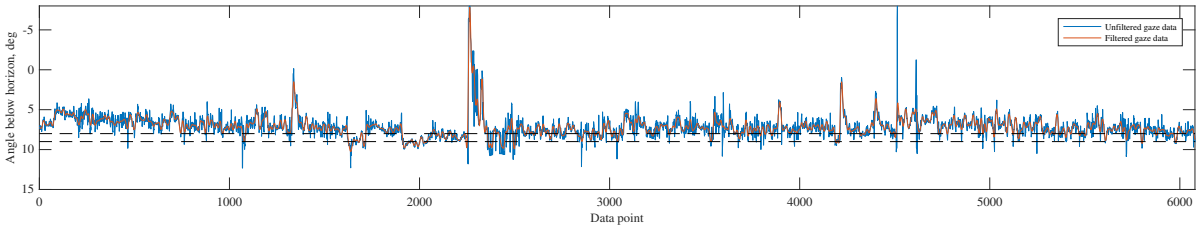

Figure D.154: Eye gaze timetrace, subject 10, Condition S89

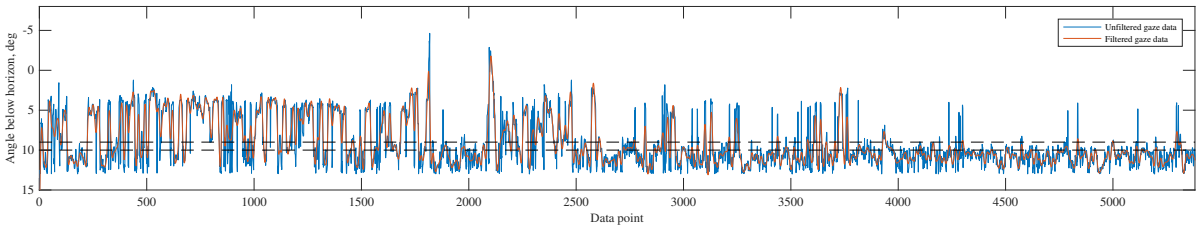

Figure D.155: Eye gaze timetrace, subject 10, Condition S910

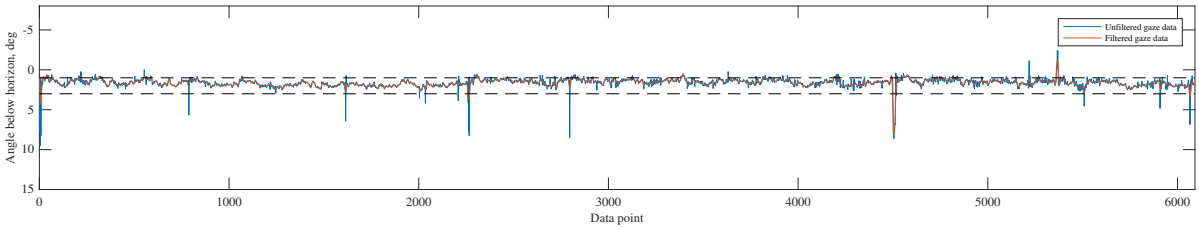

Figure D.156: Eye gaze timetrace, subject 10, Condition S13

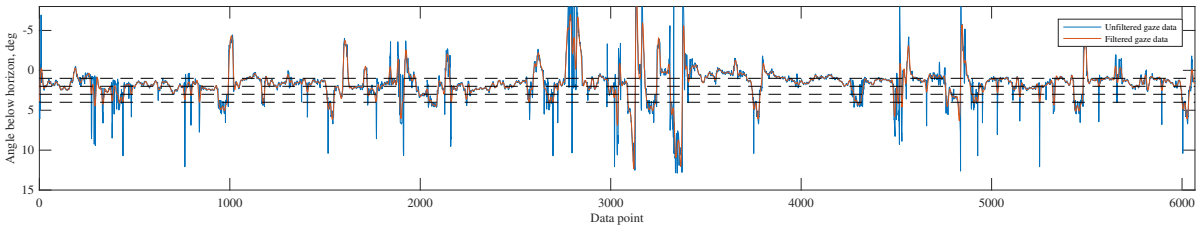

Figure D.157: Eye gaze timetrace, subject 10, Condition T34

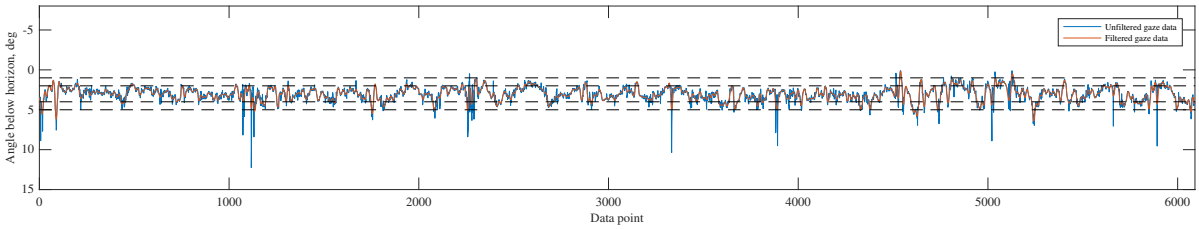

Figure D.158: Eye gaze timetrace, subject 10, Condition T45

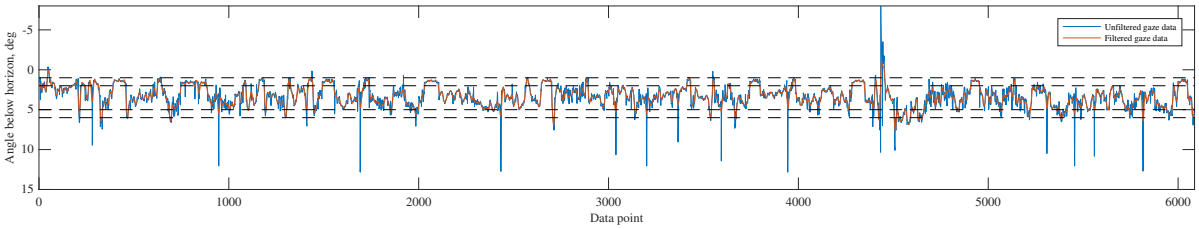

Figure D.159: Eye gaze timetrace, subject 10, Condition T56

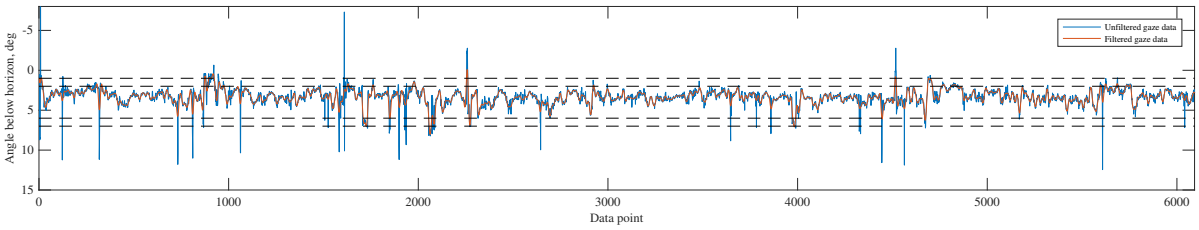

Figure D.160: Eye gaze timetrace, subject 10, Condition T67

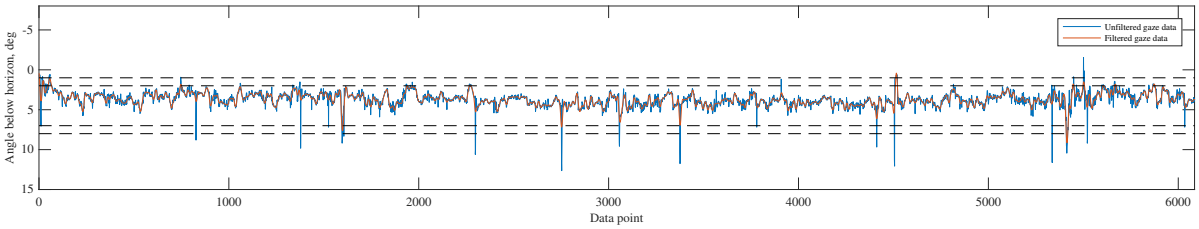

Figure D.161: Eye gaze timetrace, subject 10, Condition T78

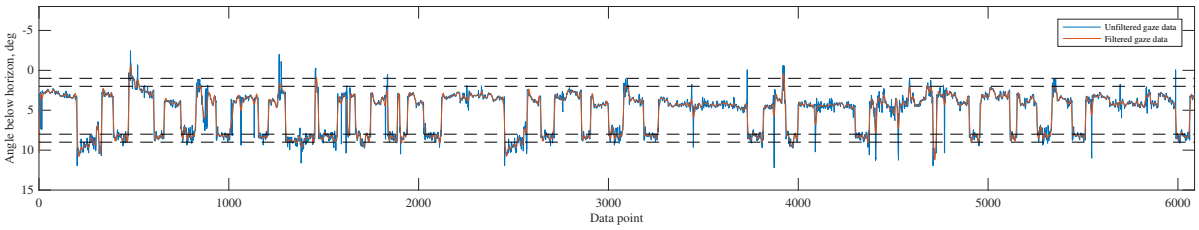

Figure D.162: Eye gaze timetrace, subject 10, Condition T89

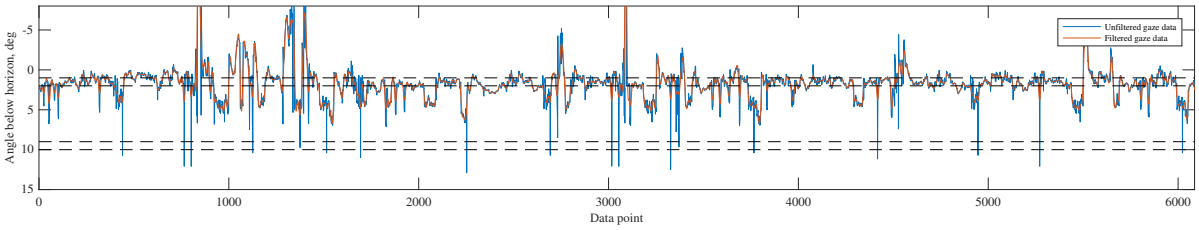

Figure D.163: Eye gaze timetrace, subject 10, Condition T910

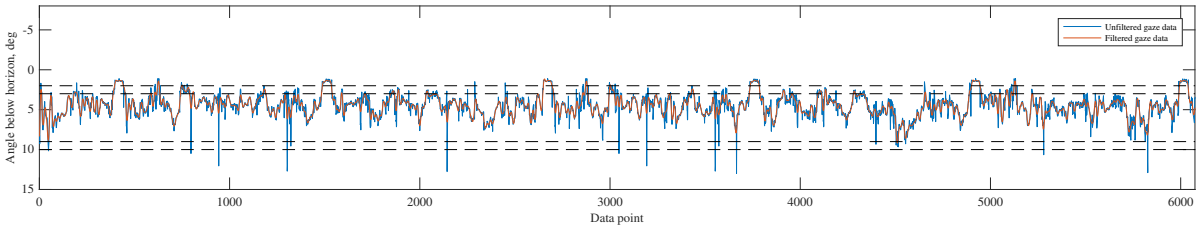

Figure D.164: Eye gaze timetrace, subject 10, Condition B23

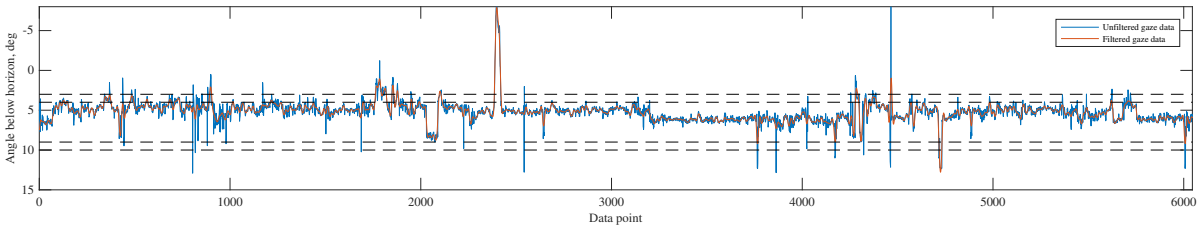

Figure D.165: Eye gaze timetrace, subject 10, Condition B34

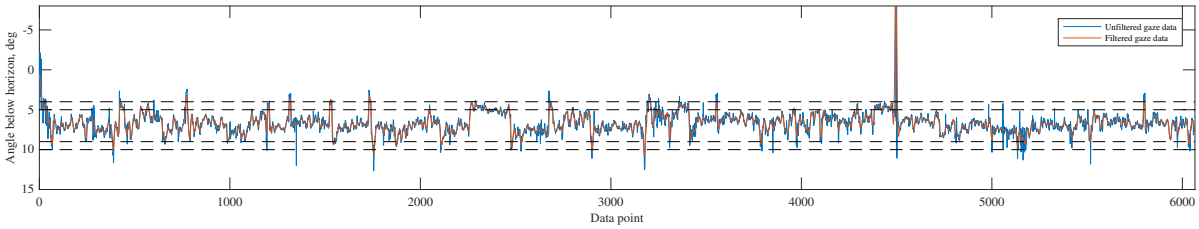

Figure D.166: Eye gaze timetrace, subject 10, Condition B45

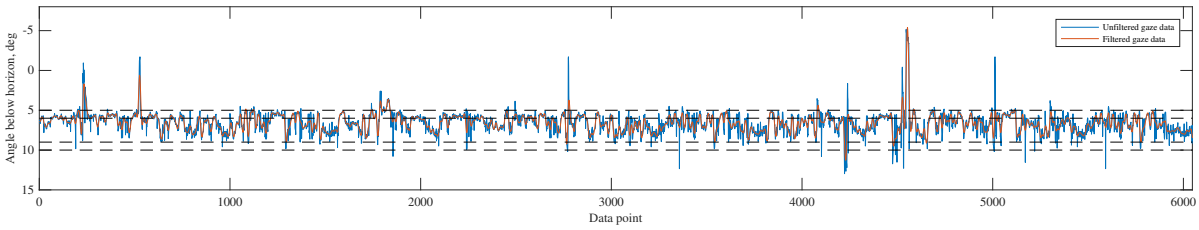

Figure D.167: Eye gaze timetrace, subject 10, Condition B56

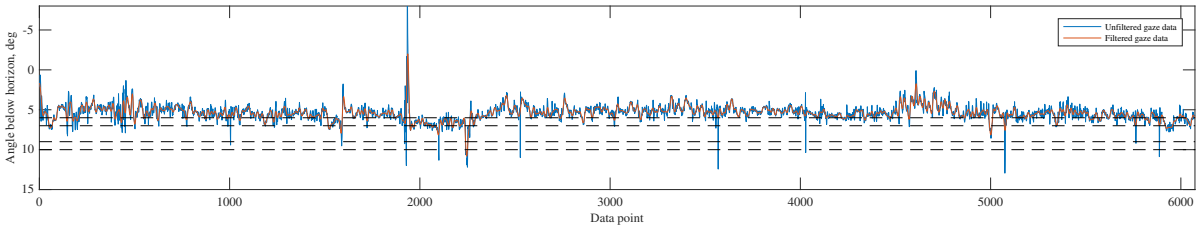

Figure D.168: Eye gaze timetrace, subject 10, Condition B67

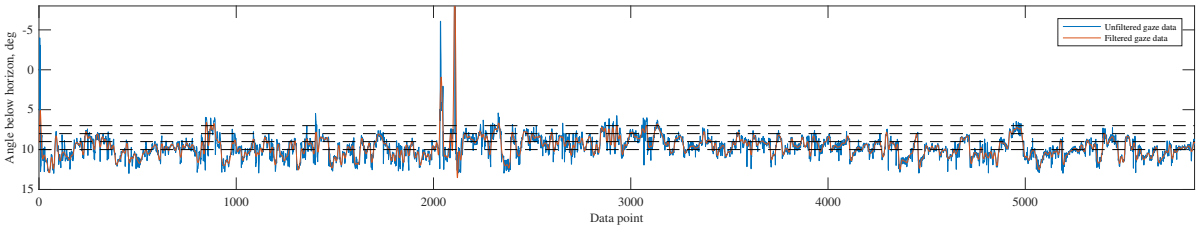

Figure D.169: Eye gaze timetrace, subject 10, Condition B78

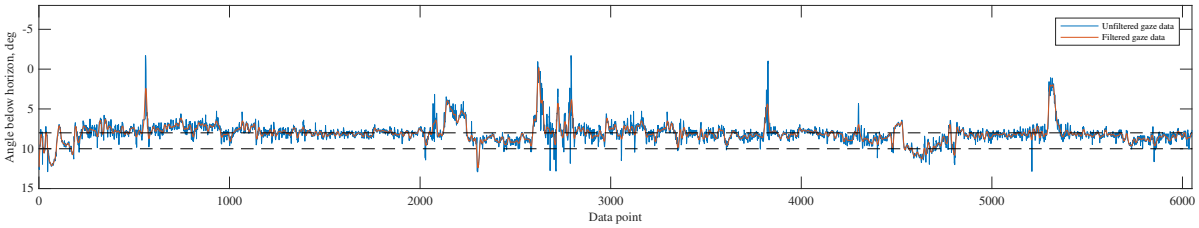

Figure D.170: Eye gaze timetrace, subject 10, Condition S810

Subject 10

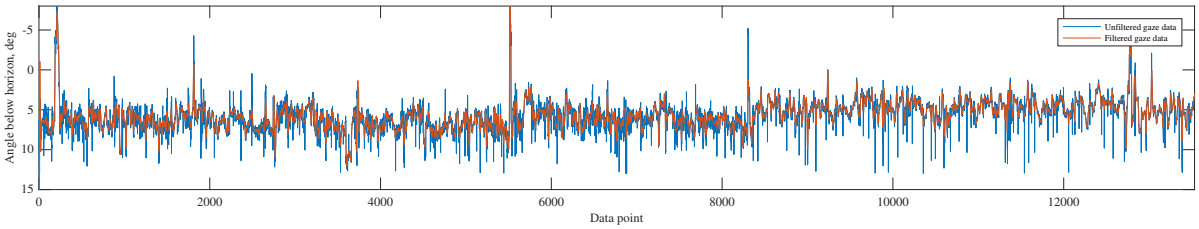

Figure D.171: Eye gaze timetrace, subject 11, Condition NO

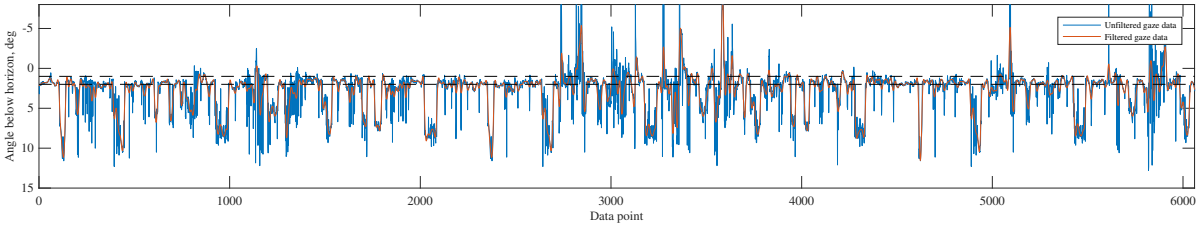

Figure D.172: Eye gaze timetrace, subject 11, Condition S12

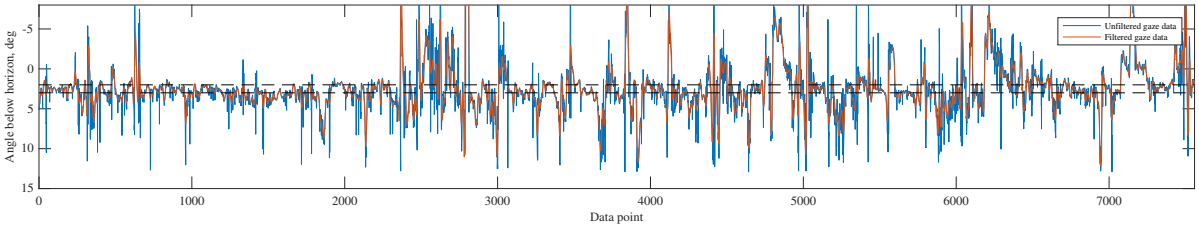

Figure D.173: Eye gaze timetrace, subject 11, Condition S23

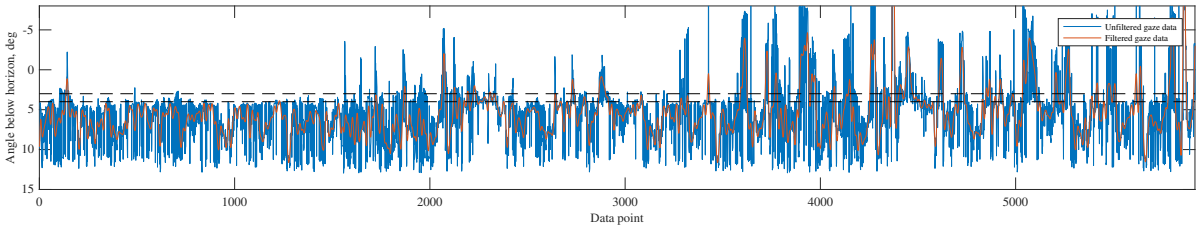

Figure D.174: Eye gaze timetrace, subject 11, Condition S34

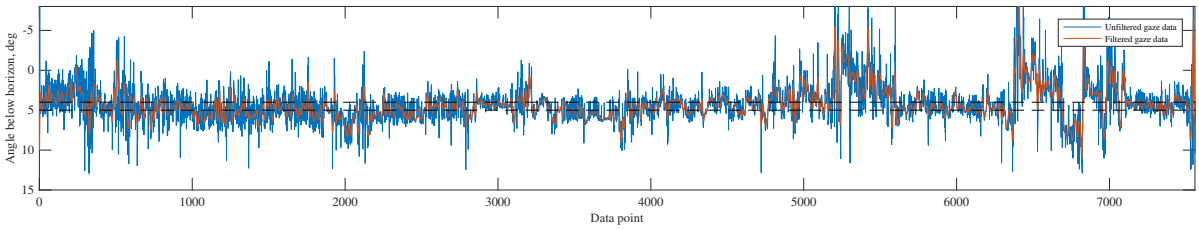

Figure D.175: Eye gaze timetrace, subject 11, Condition S45

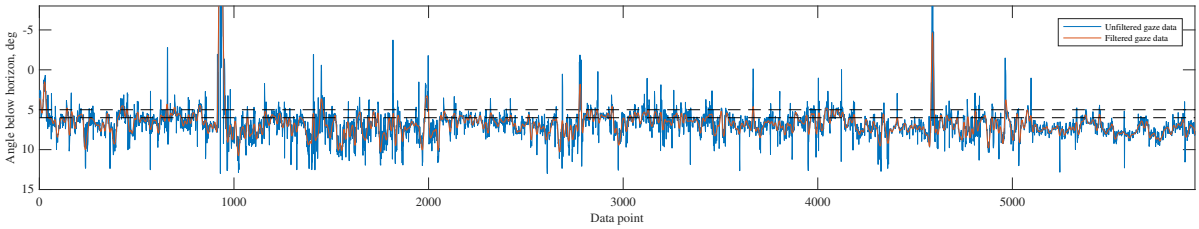

Figure D.176: Eye gaze timetrace, subject 11, Condition S56

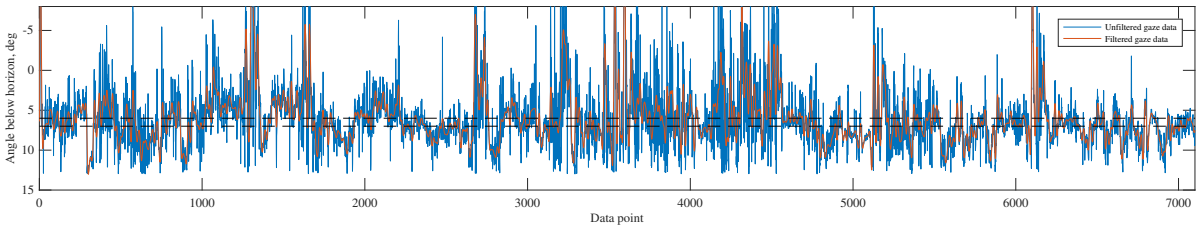

Figure D.177: Eye gaze timetrace, subject 11, Condition S67

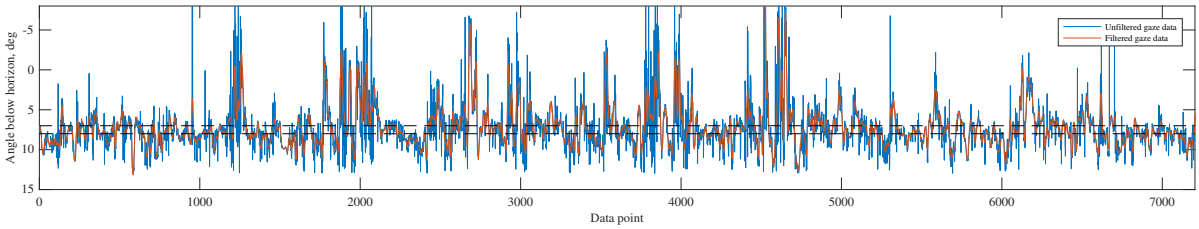

Figure D.178: Eye gaze timetrace, subject 11, Condition S78

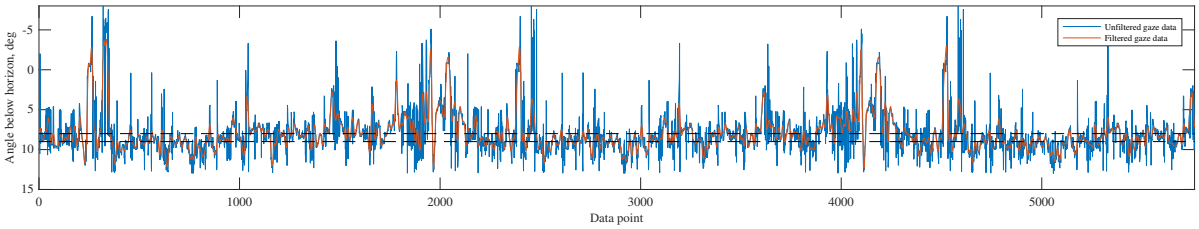

Figure D.179: Eye gaze timetrace, subject 11, Condition S89

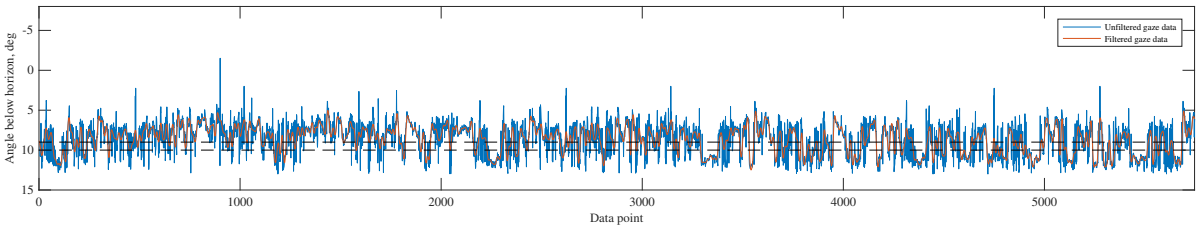

Figure D.180: Eye gaze timetrace, subject 11, Condition S910

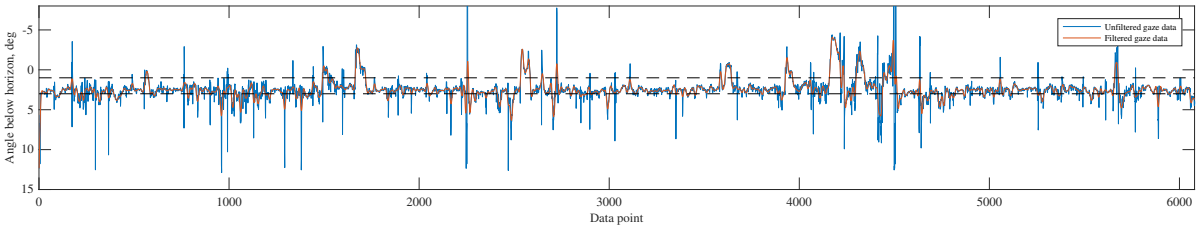

Figure D.181: Eye gaze timetrace, subject 11, Condition S13

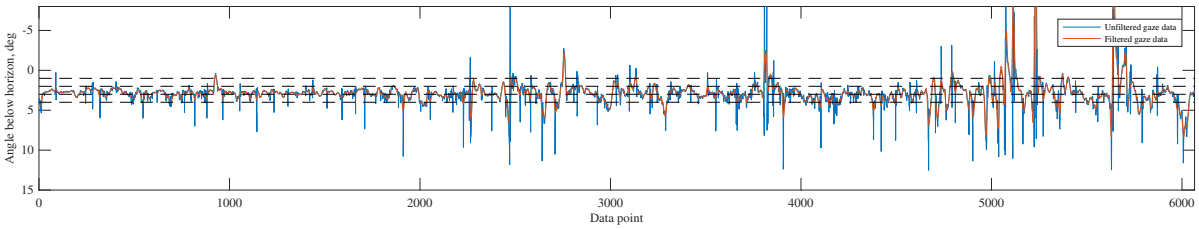

Figure D.182: Eye gaze timetrace, subject 11, Condition T34

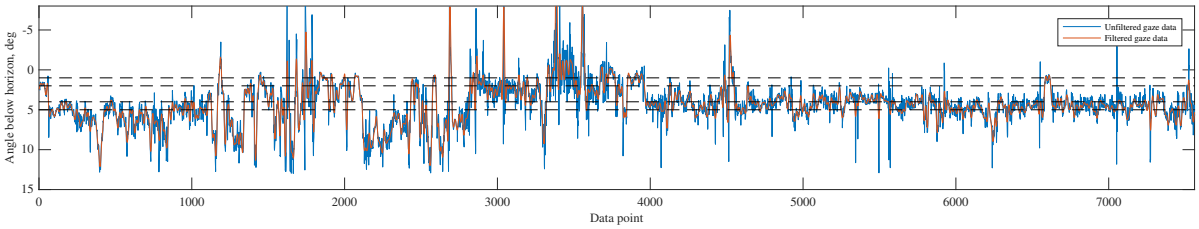

Figure D.183: Eye gaze timetrace, subject 11, Condition T45

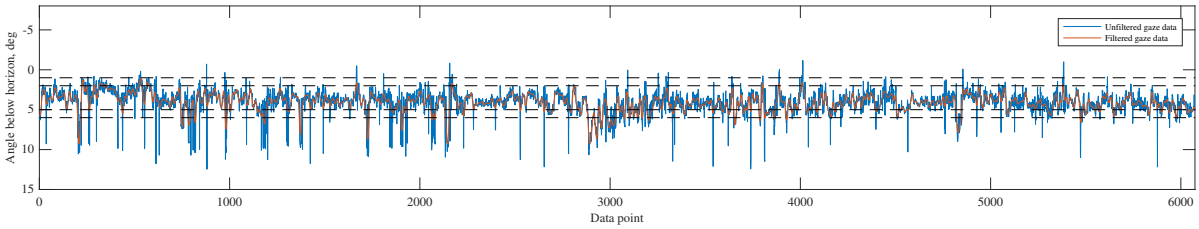

Figure D.184: Eye gaze timetrace, subject 11, Condition T56

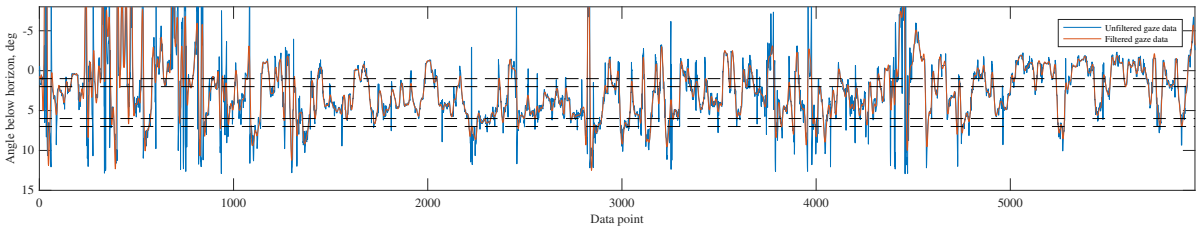

Figure D.185: Eye gaze timetrace, subject 11, Condition T67

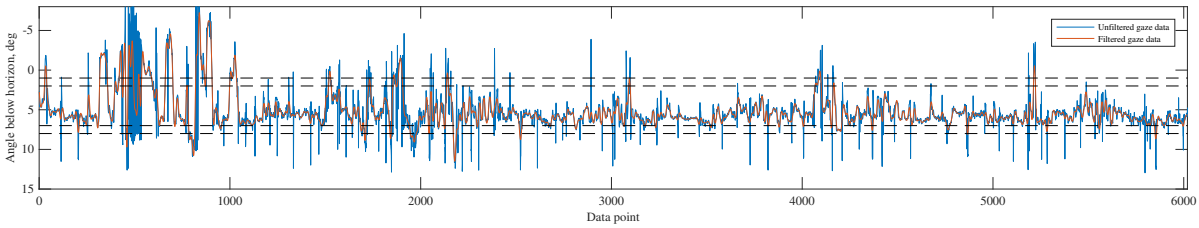

Figure D.186: Eye gaze timetrace, subject 11, Condition T78

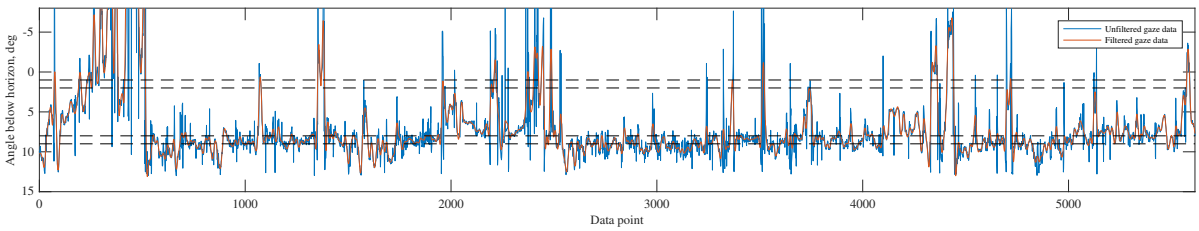

Figure D.187: Eye gaze timetrace, subject 11, Condition T89

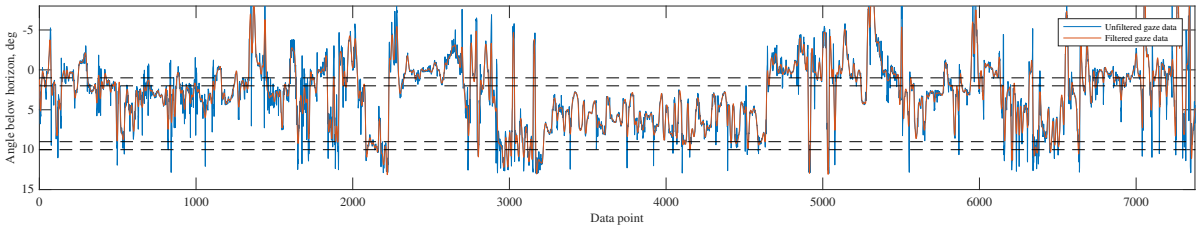

Figure D.188: Eye gaze timetrace, subject 11, Condition T910

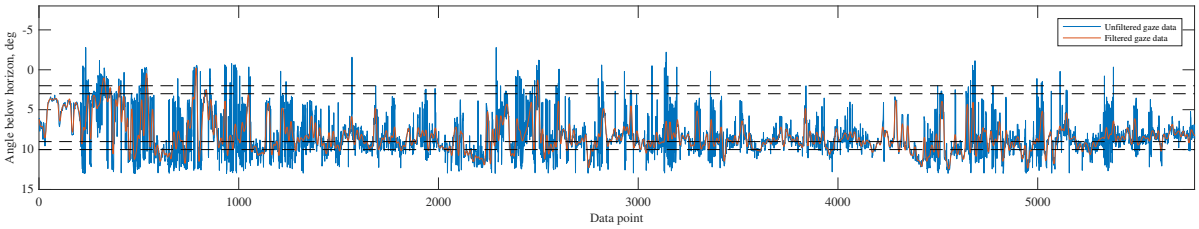

Figure D.189: Eye gaze timetrace, subject 11, Condition B23

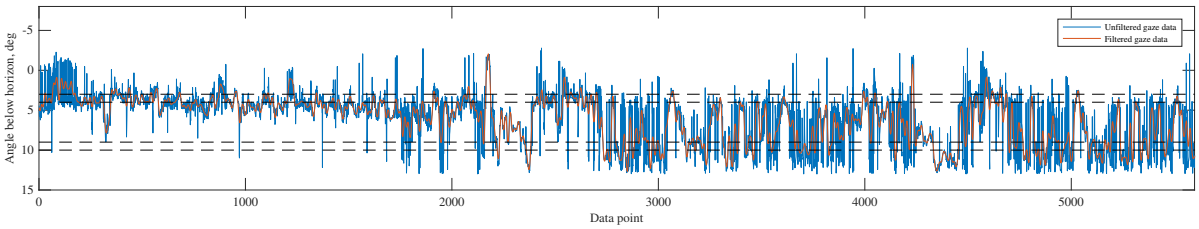

Figure D.190: Eye gaze timetrace, subject 11, Condition B34

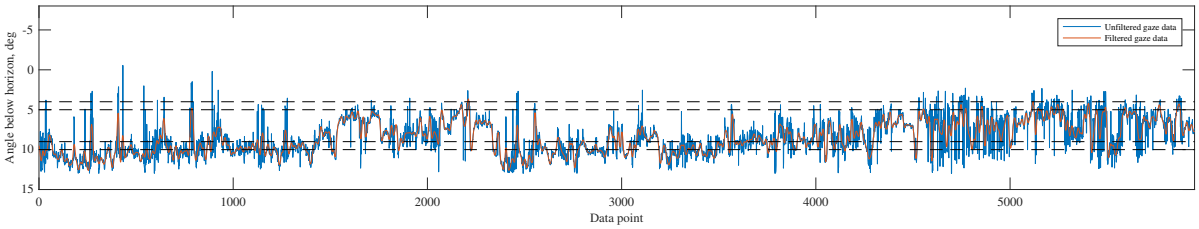

Figure D.191: Eye gaze timetrace, subject 11, Condition B45

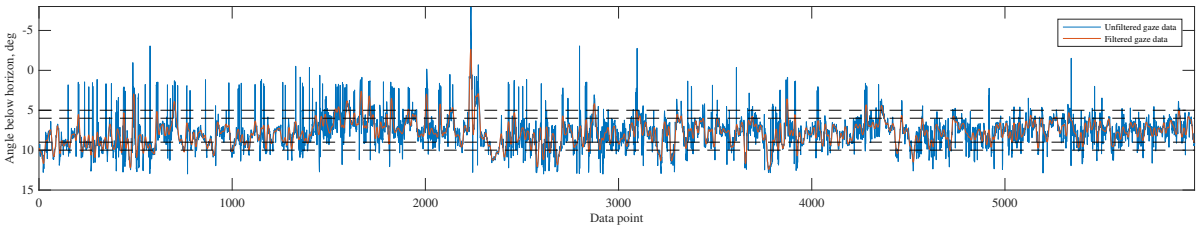

Figure D.192: Eye gaze timetrace, subject 11, Condition B56

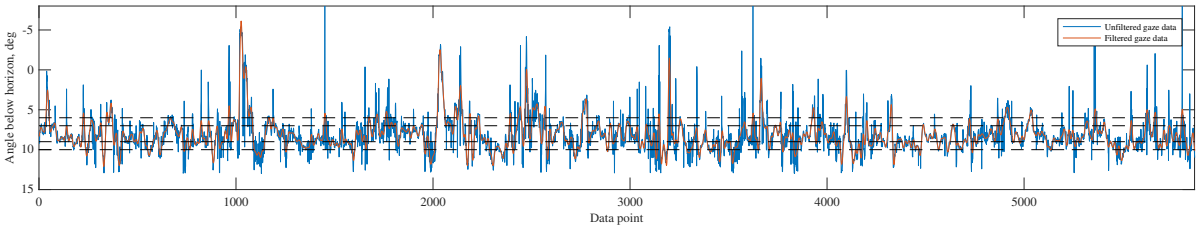

Figure D.193: Eye gaze timetrace, subject 11, Condition B67

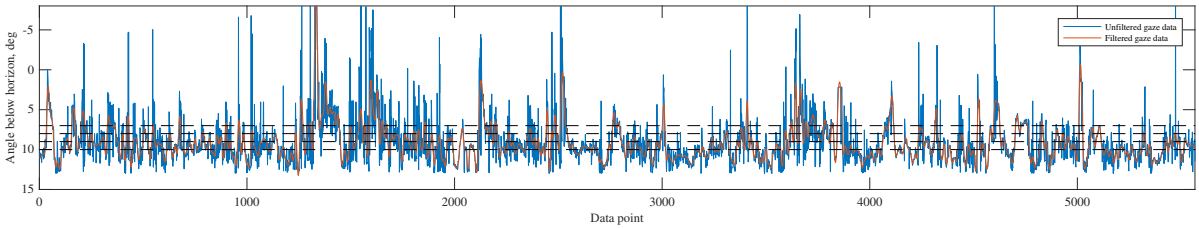

Figure D.194: Eye gaze timetrace, subject 11, Condition B78

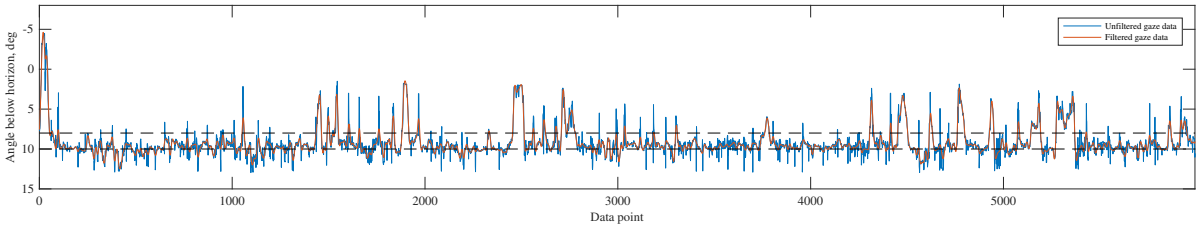

Figure D.195: Eye gaze timetrace, subject 11, Condition S810

Subject 12

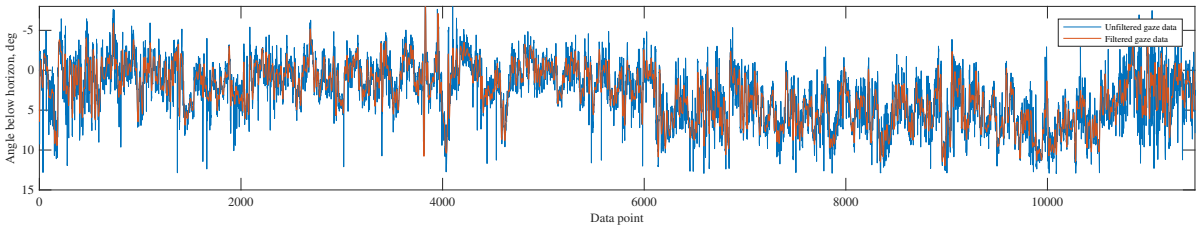

Figure D.196: Eye gaze timetrace, subject 13, Condition NO

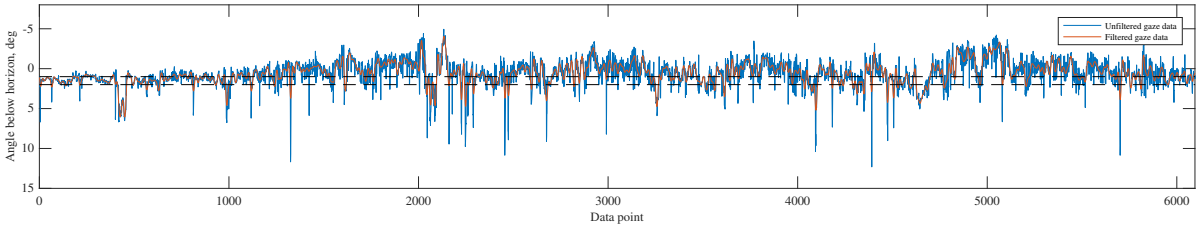

Figure D.197: Eye gaze timetrace, subject 13, Condition S12

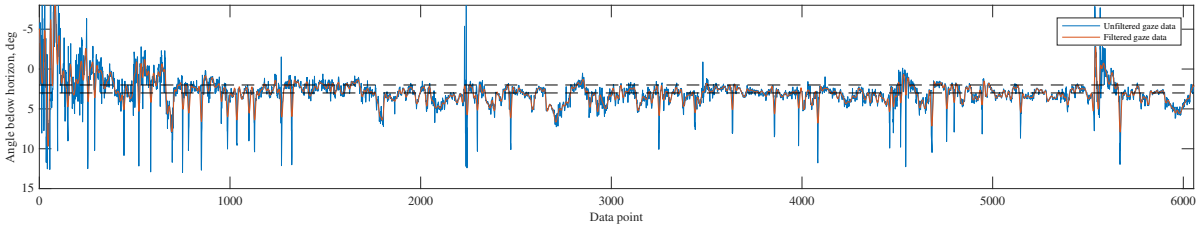

Figure D.198: Eye gaze timetrace, subject 13, Condition S23

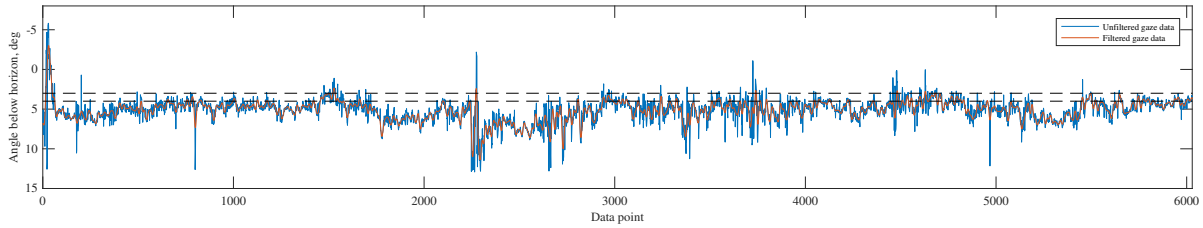

Figure D.199: Eye gaze timetrace, subject 13, Condition S34

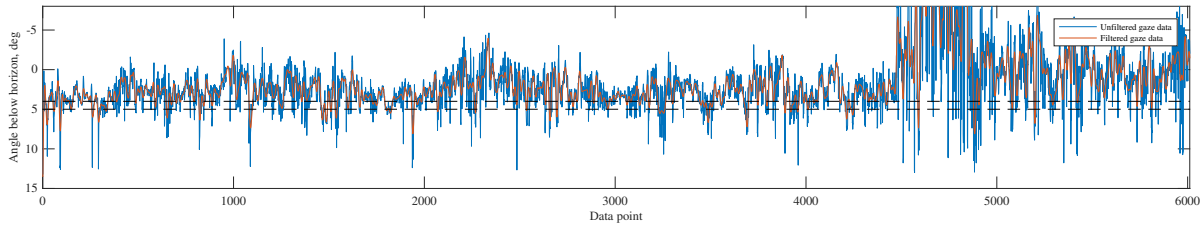

Figure D.200: Eye gaze timetrace, subject 13, Condition S45

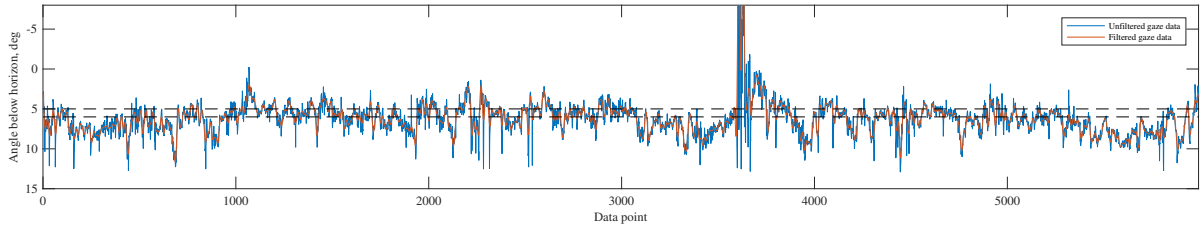

Figure D.201: Eye gaze timetrace, subject 13, Condition S56

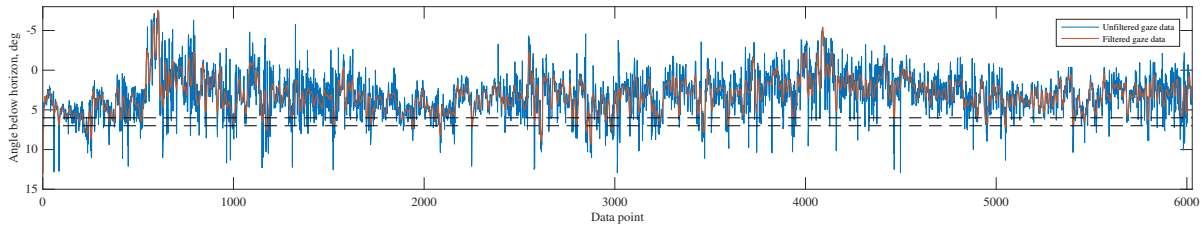

Figure D.202: Eye gaze timetrace, subject 13, Condition S67

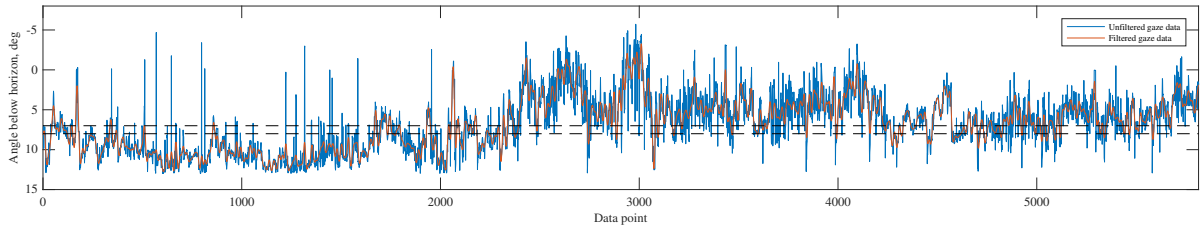

Figure D.203: Eye gaze timetrace, subject 13, Condition S78

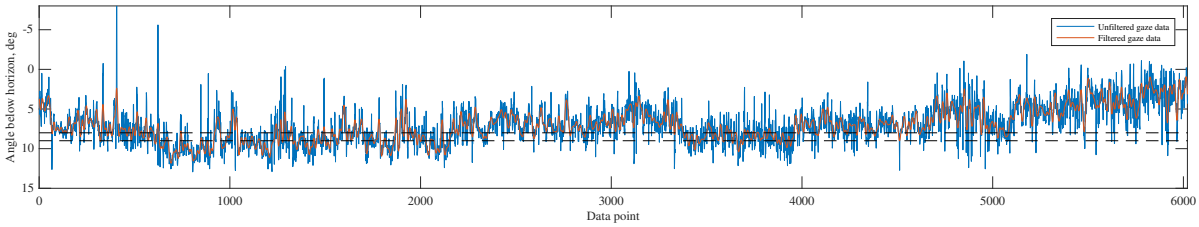

Figure D.204: Eye gaze timetrace, subject 13, Condition S89

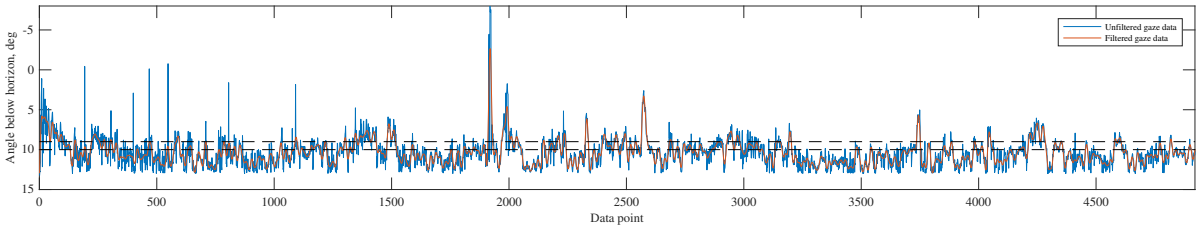

Figure D.205: Eye gaze timetrace, subject 13, Condition S910

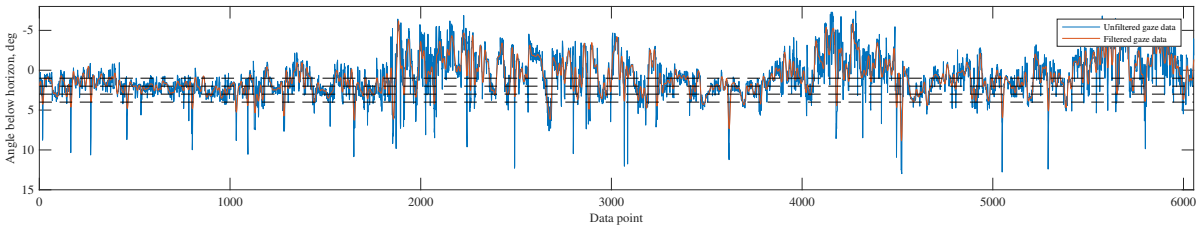

Figure D.206: Eye gaze timetrace, subject 13, Condition T34

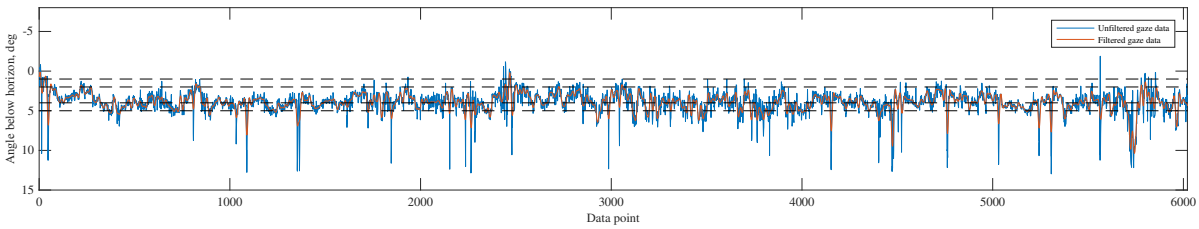

Figure D.207: Eye gaze timetrace, subject 13, Condition T45

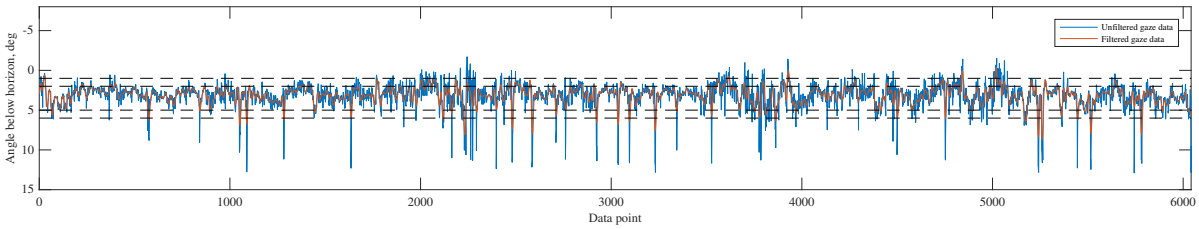

Figure D.208: Eye gaze timetrace, subject 13, Condition T56

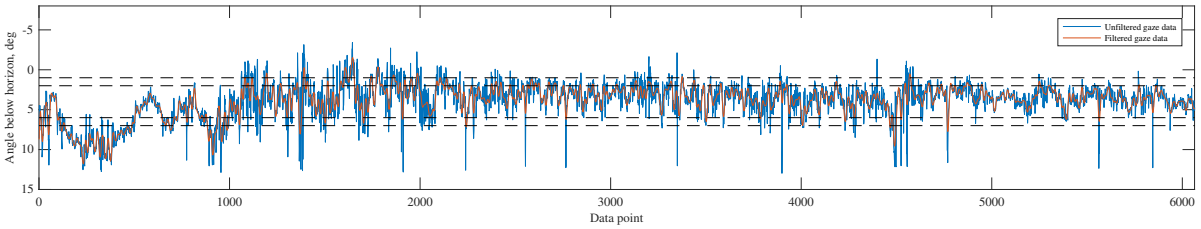

Figure D.209: Eye gaze timetrace, subject 13, Condition T67

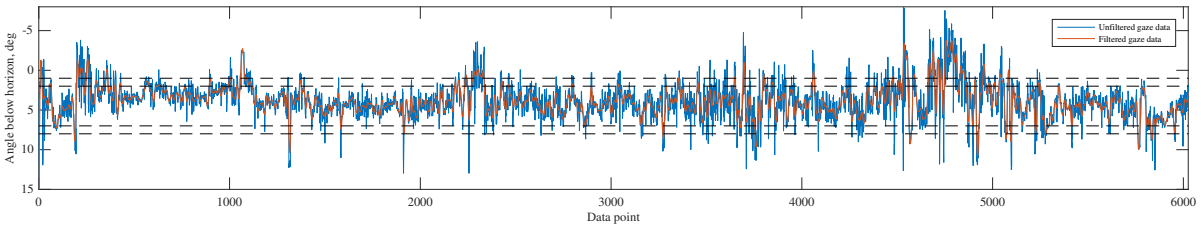

Figure D.210: Eye gaze timetrace, subject 13, Condition T78

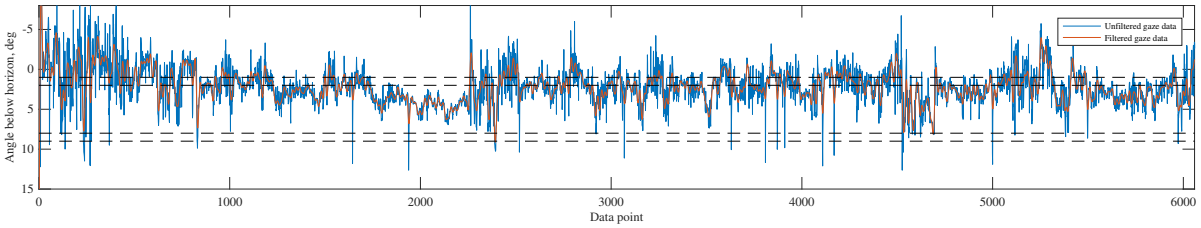

Figure D.211: Eye gaze timetrace, subject 13, Condition T89

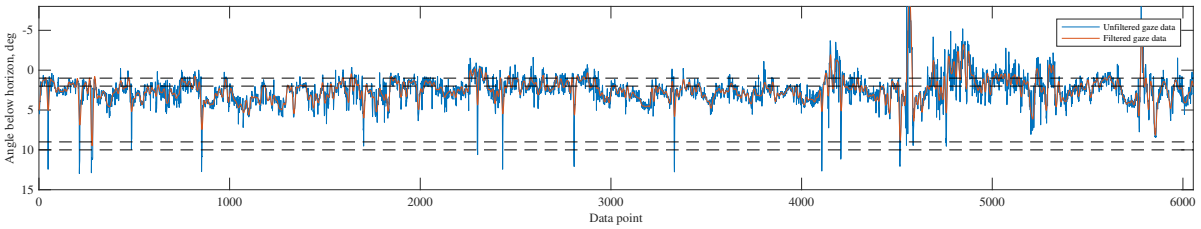

Figure D.212: Eye gaze timetrace, subject 13, Condition T910

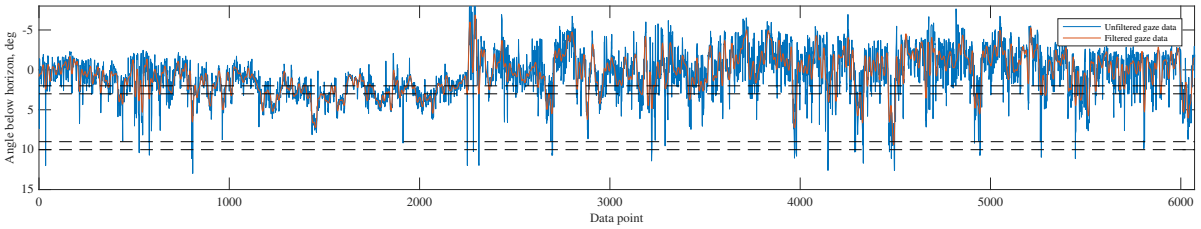

Figure D.213: Eye gaze timetrace, subject 13, Condition B23

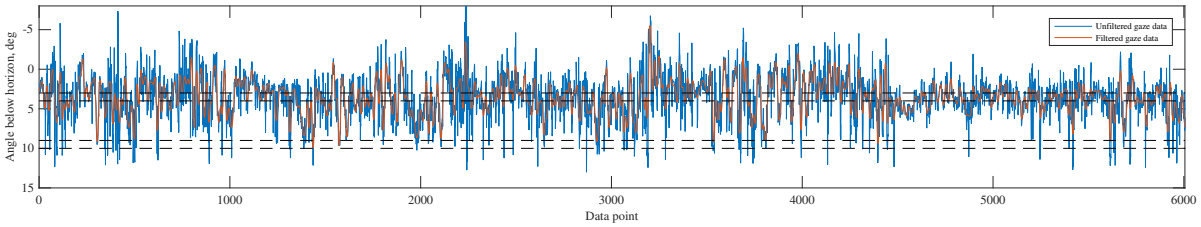

Figure D.214: Eye gaze timetrace, subject 13, Condition B34

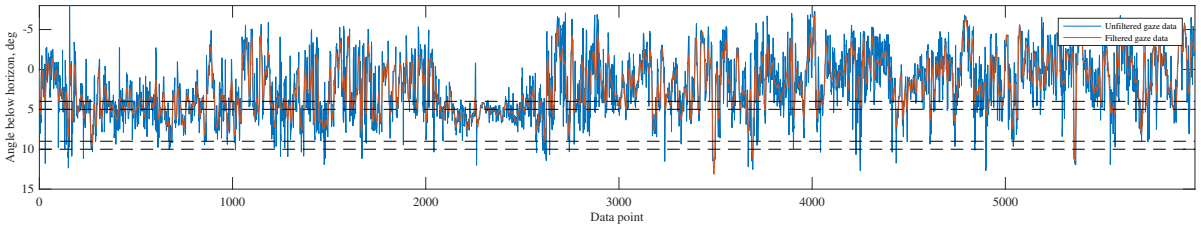

Figure D.215: Eye gaze timetrace, subject 13, Condition B45

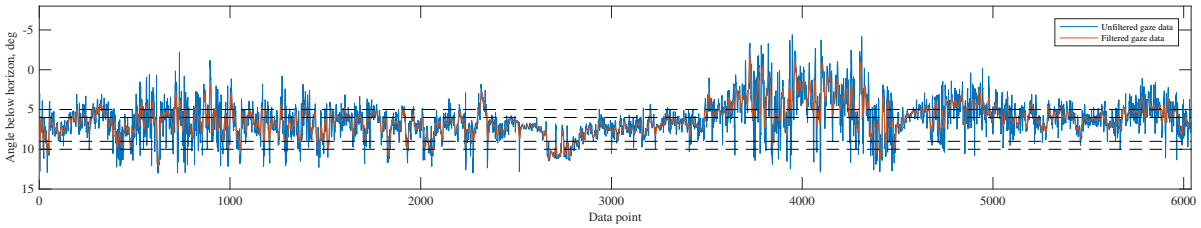

Figure D.216: Eye gaze timetrace, subject 13, Condition B56

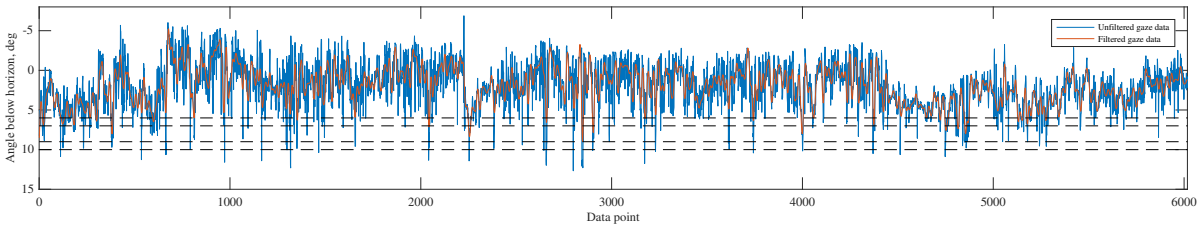

Figure D.217: Eye gaze timetrace, subject 13, Condition B67

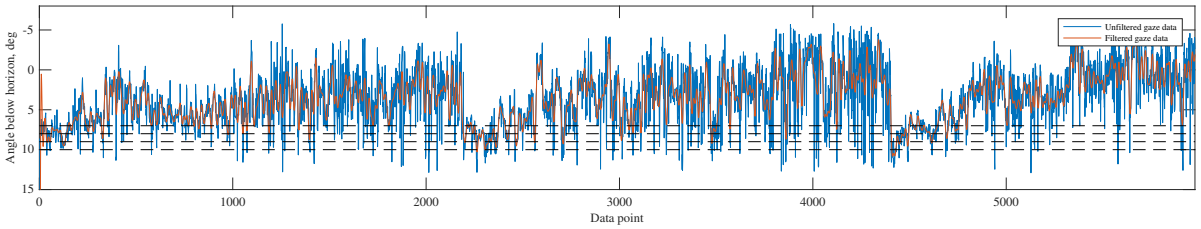

Figure D.218: Eye gaze timetrace, subject 13, Condition B78

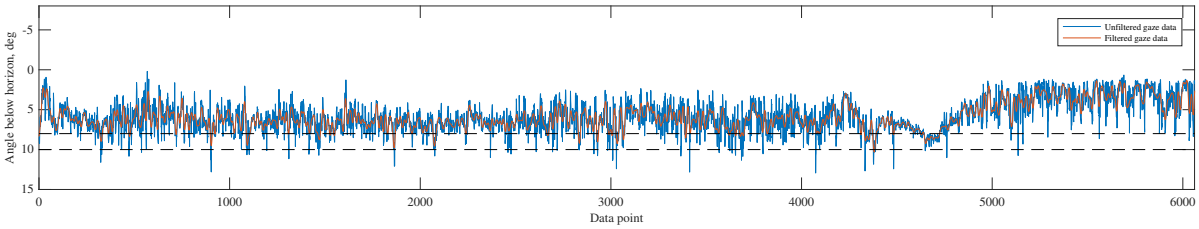

Figure D.219: Eye gaze timetrace, subject 13, Condition S810

E

Representative Model Fits - Subject 1

This chapter contains the Frequency Response Functions and fitted parametric model to the data of Subject 1.

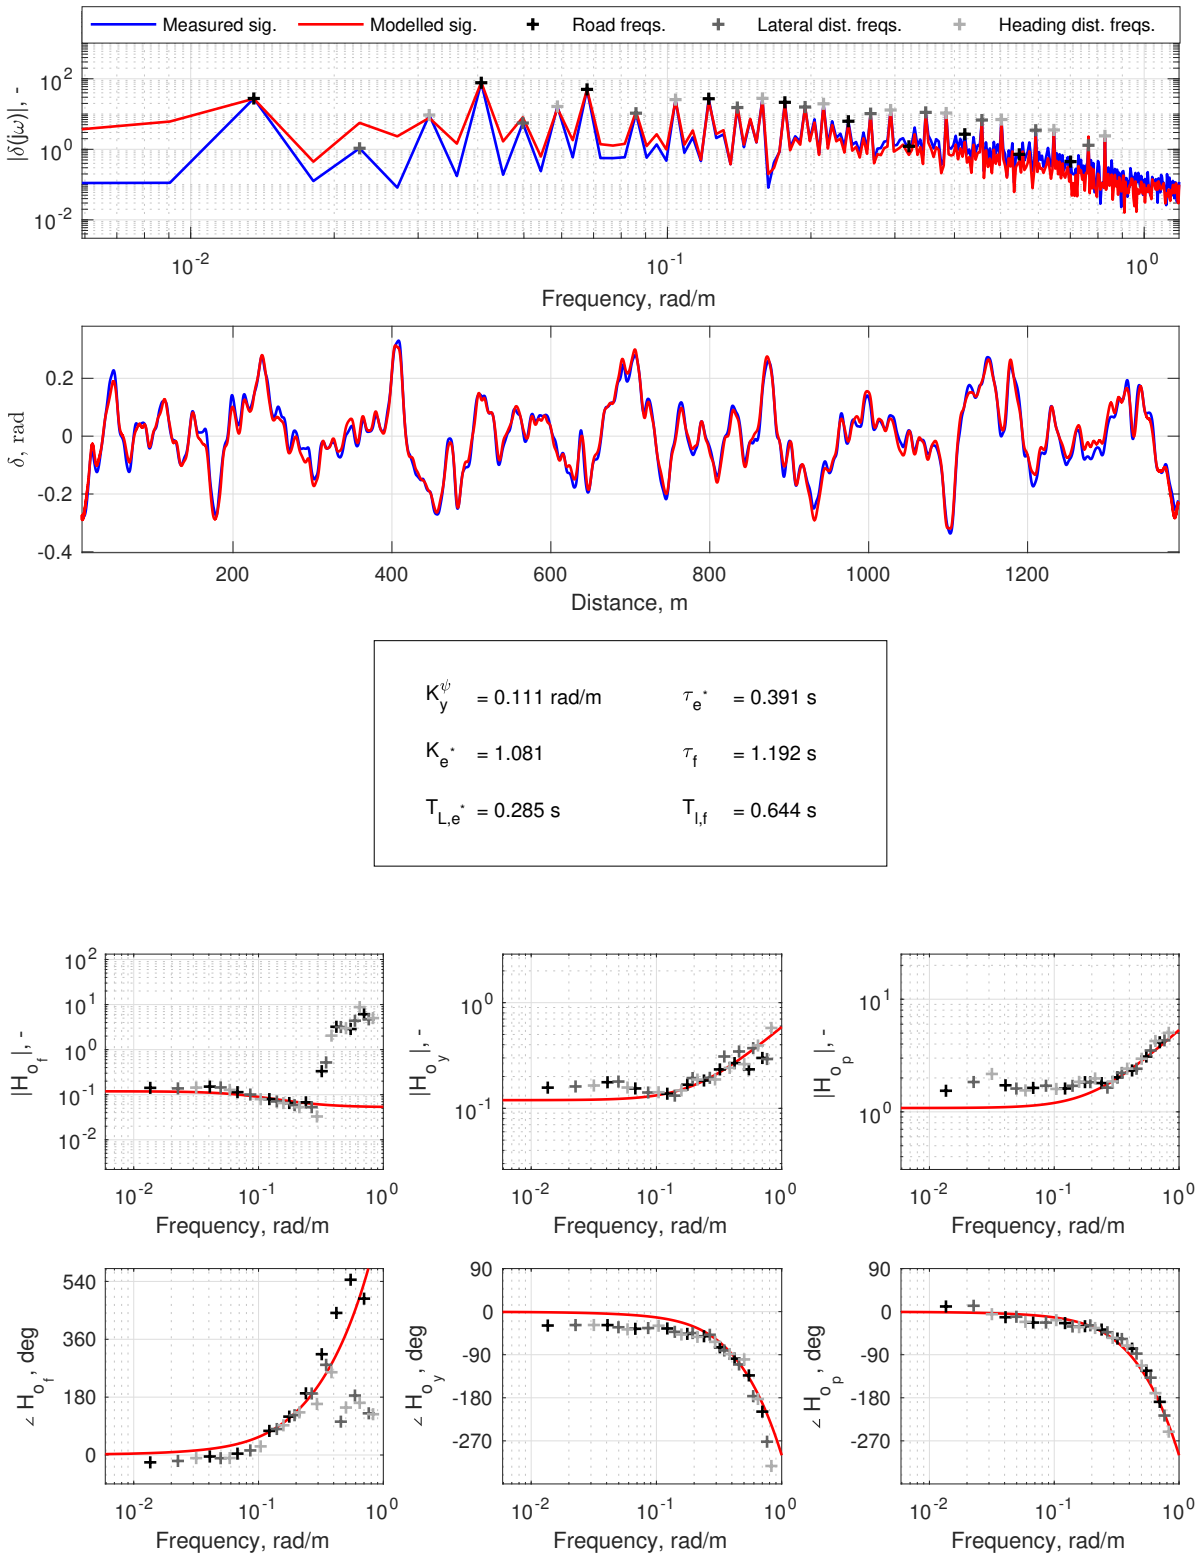

Figure E.1: FRF estimates and parameter estimation, subject 1, Condition NO

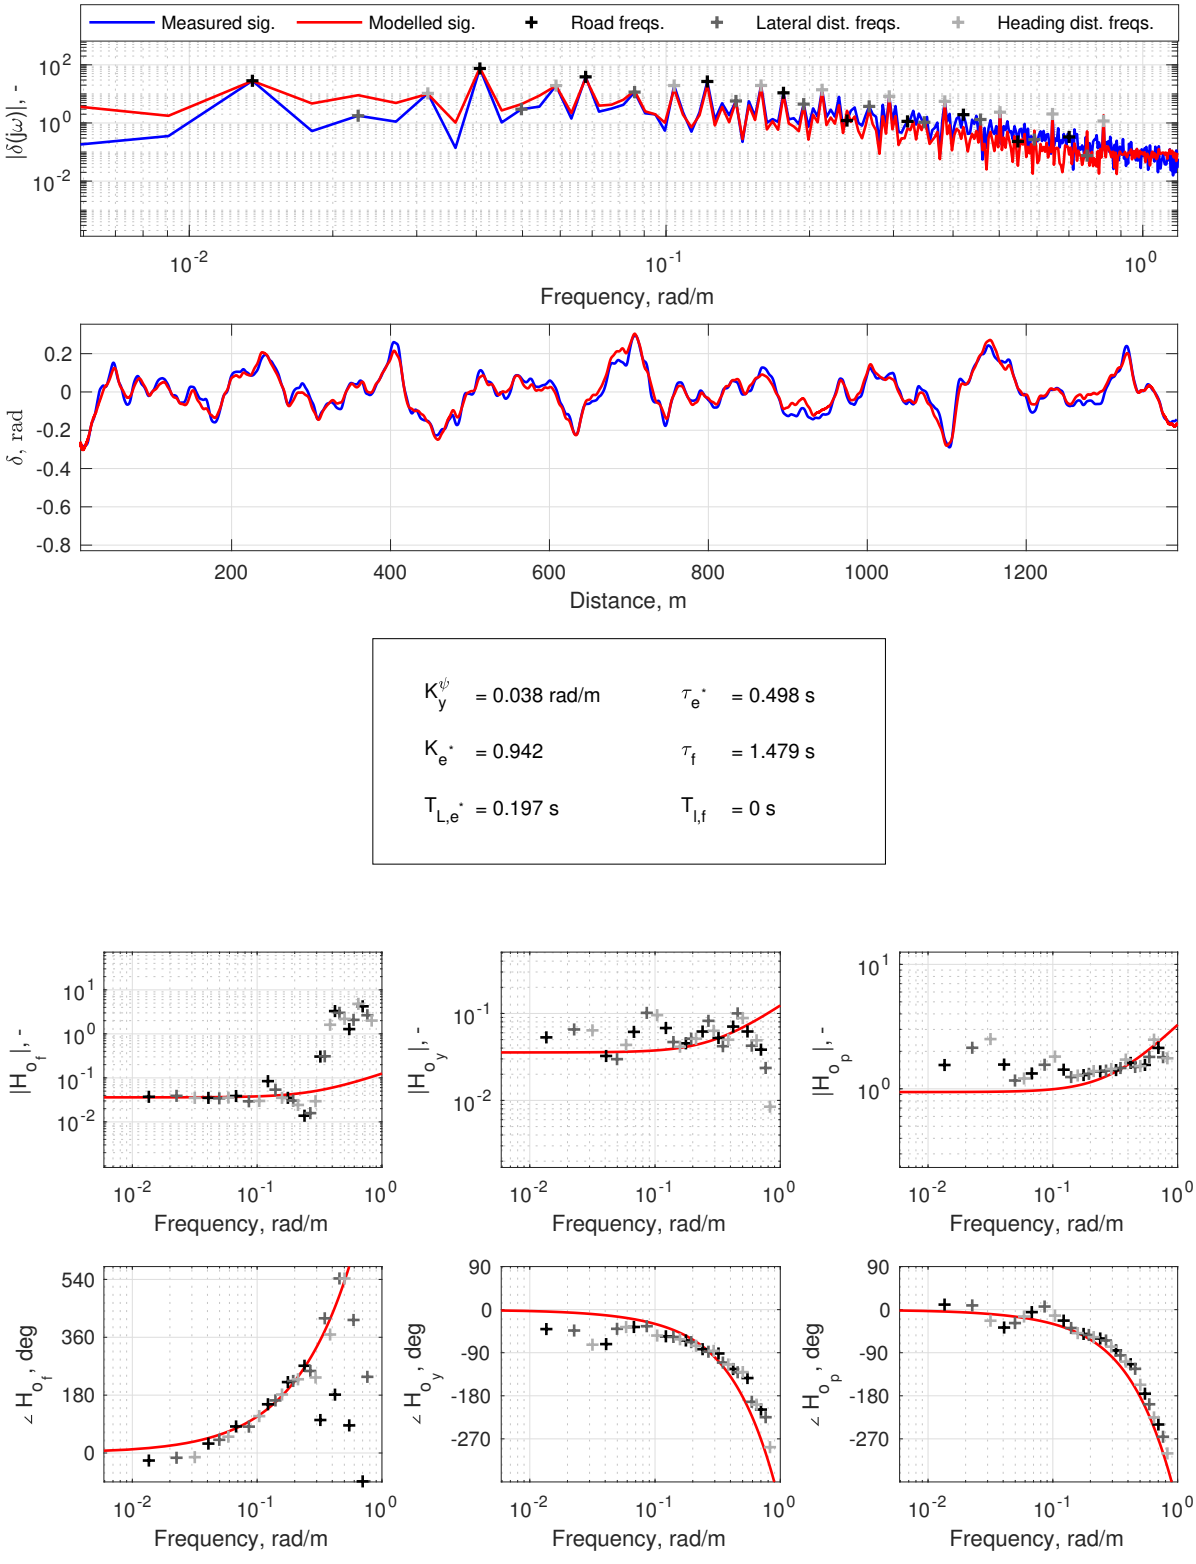

Figure E.2: FRF estimates and parameter estimation, subject 1, Condition S12

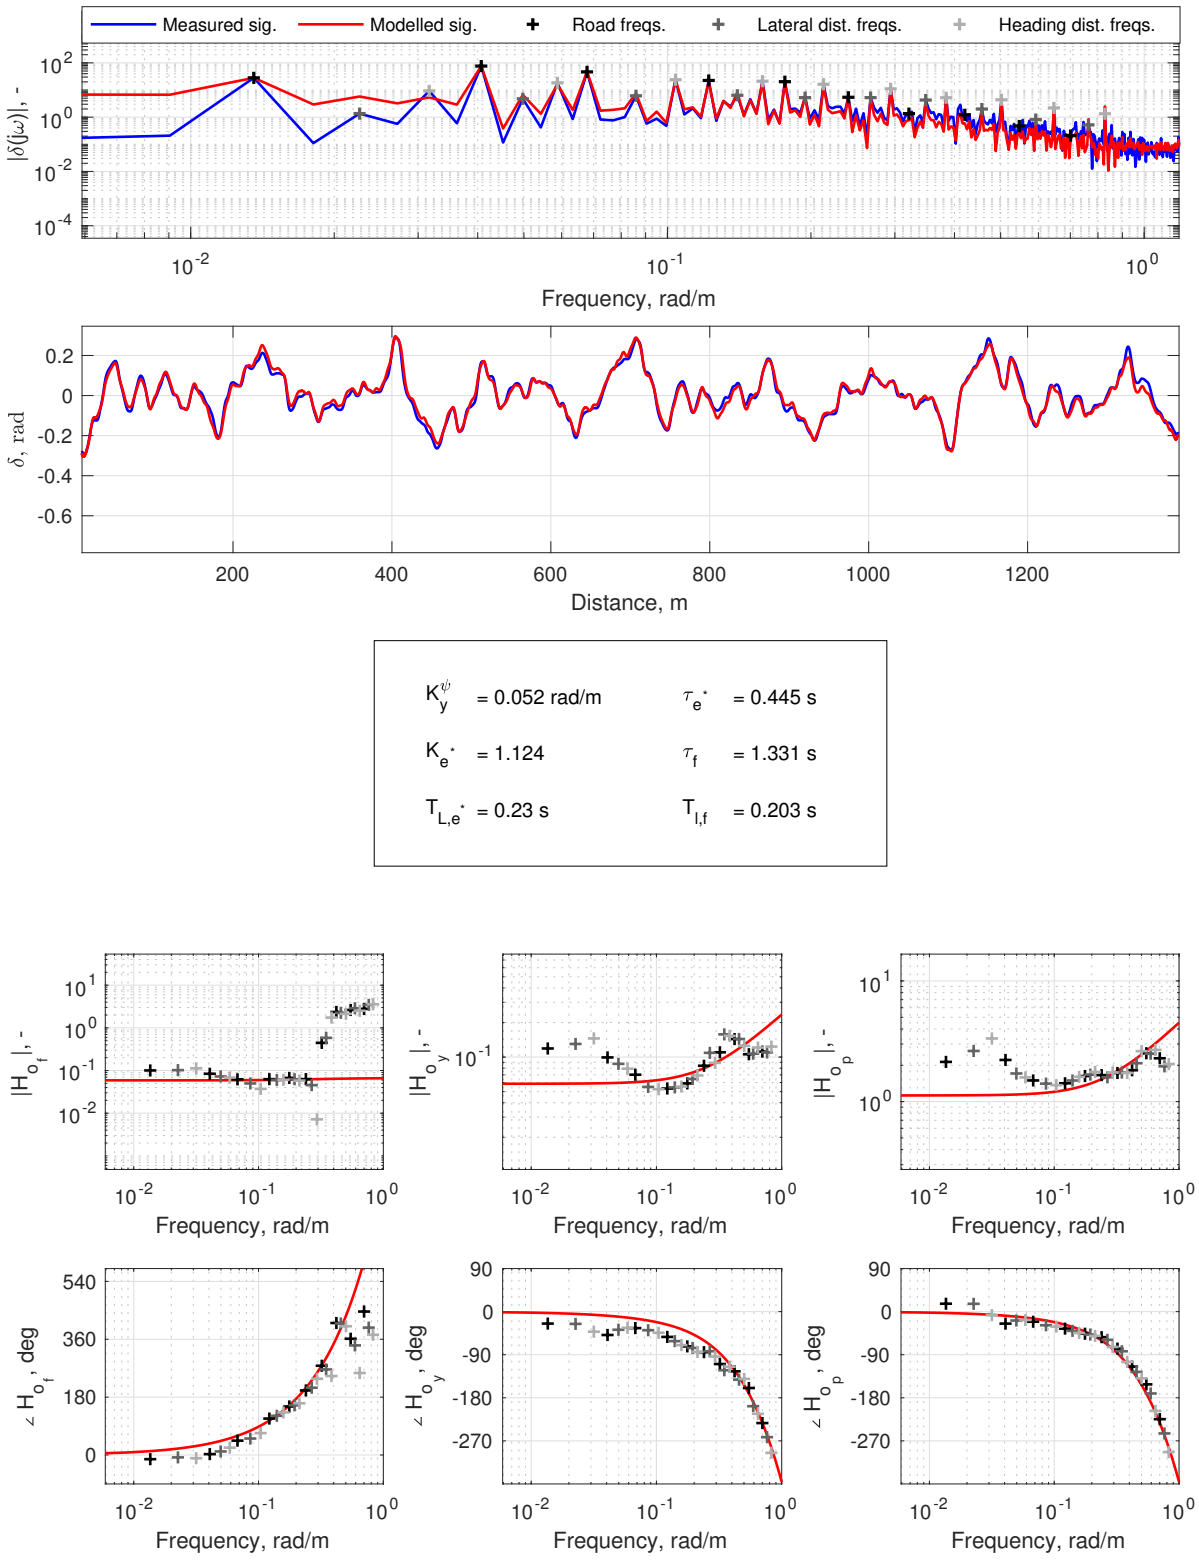

Figure E.3: FRF estimates and parameter estimation, subject 1, Condition S23

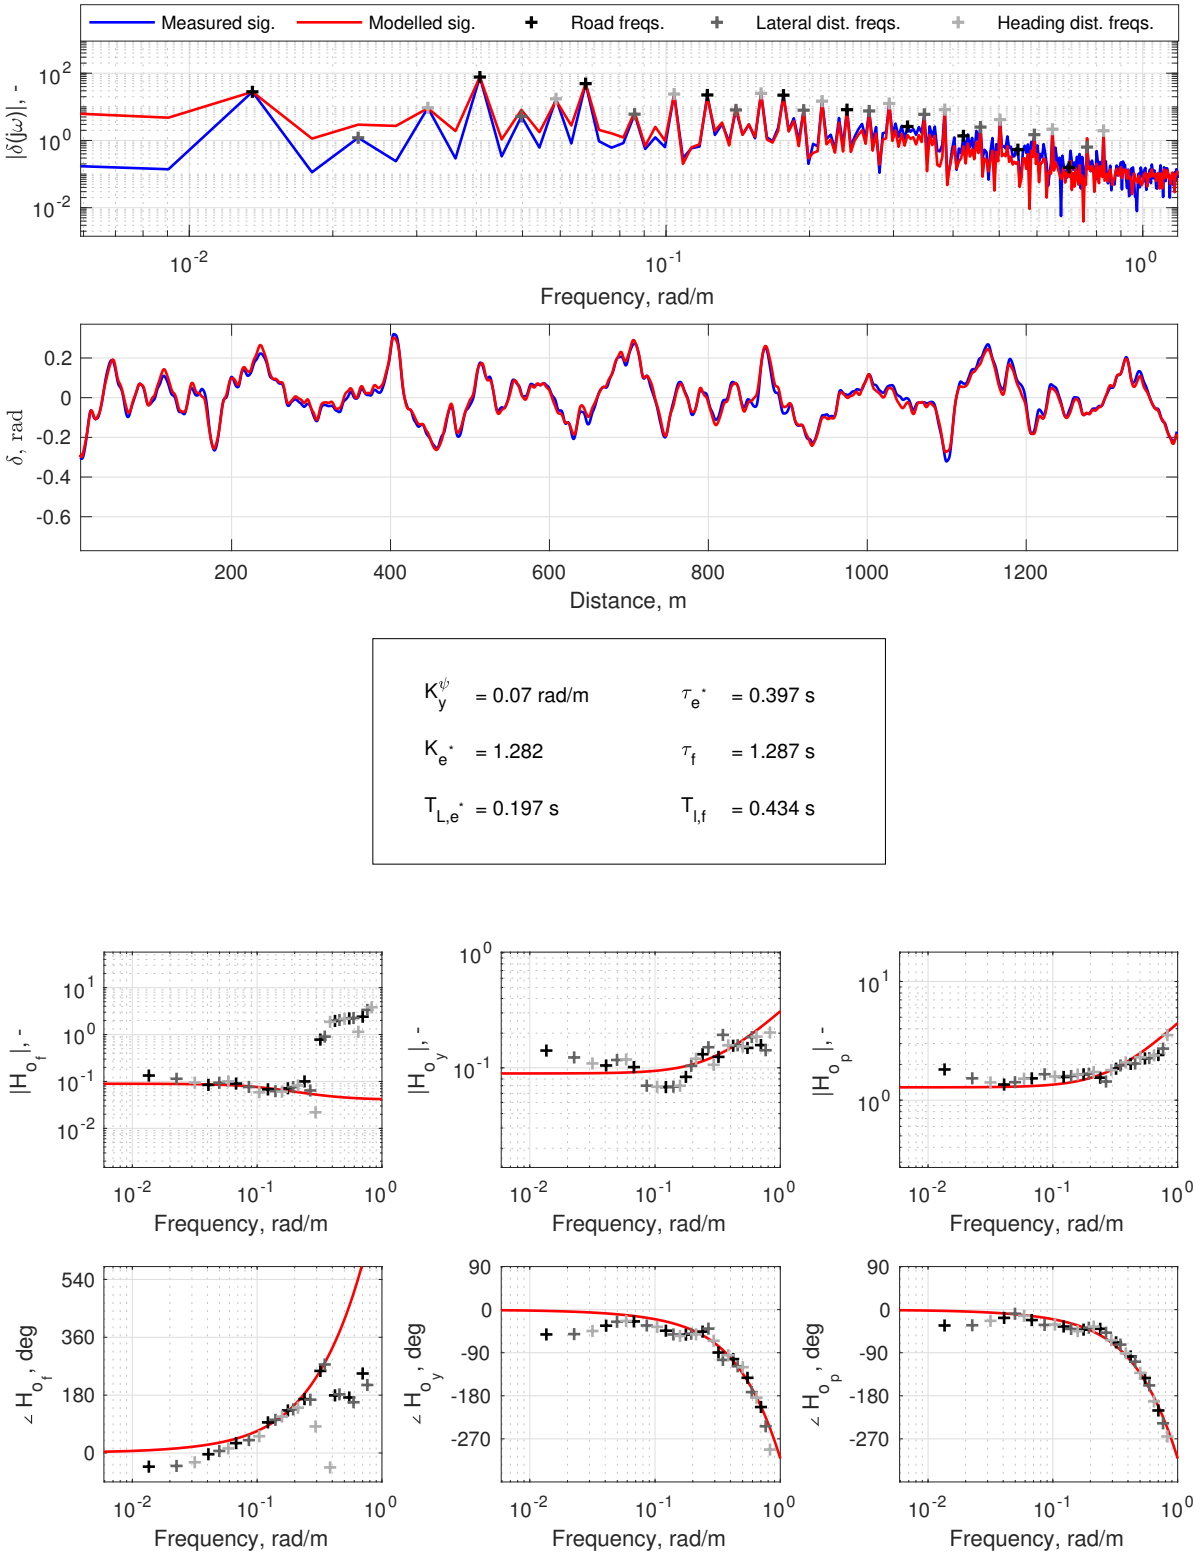

Figure E.4: FRF estimates and parameter estimation, subject 1, Condition S34

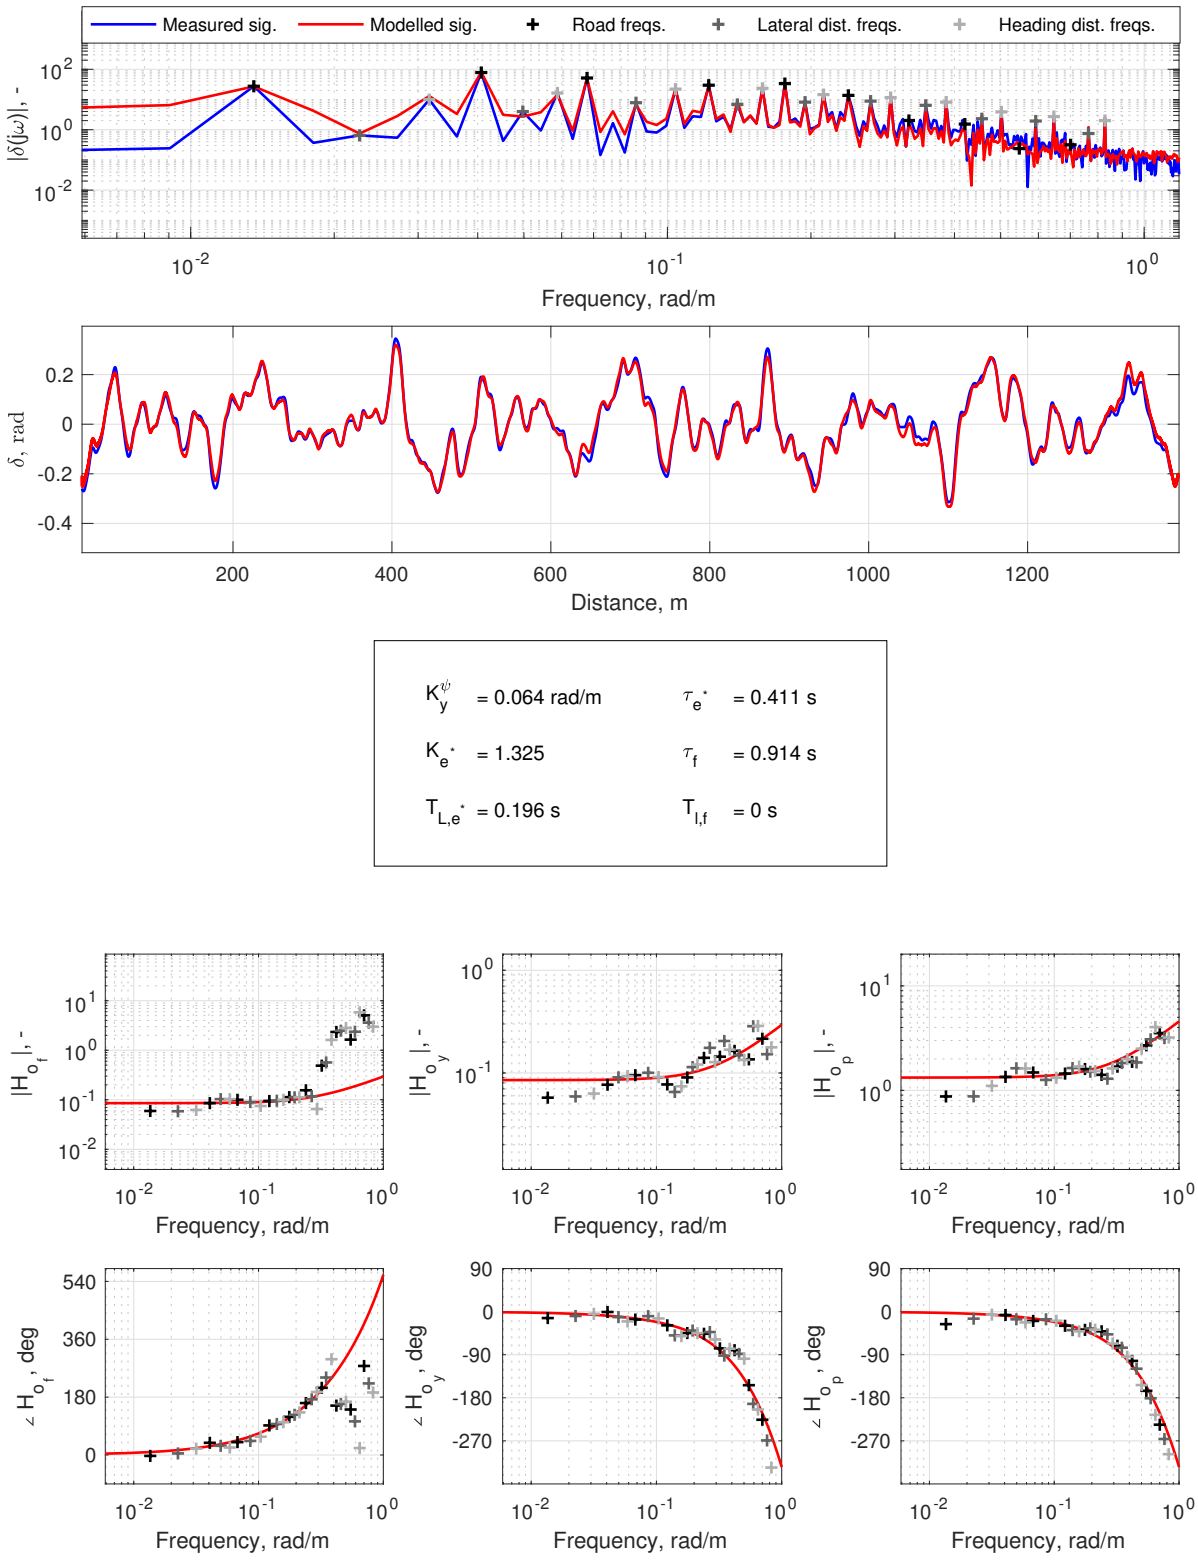

Figure E.5: FRF estimates and parameter estimation, subject 1, Condition S45

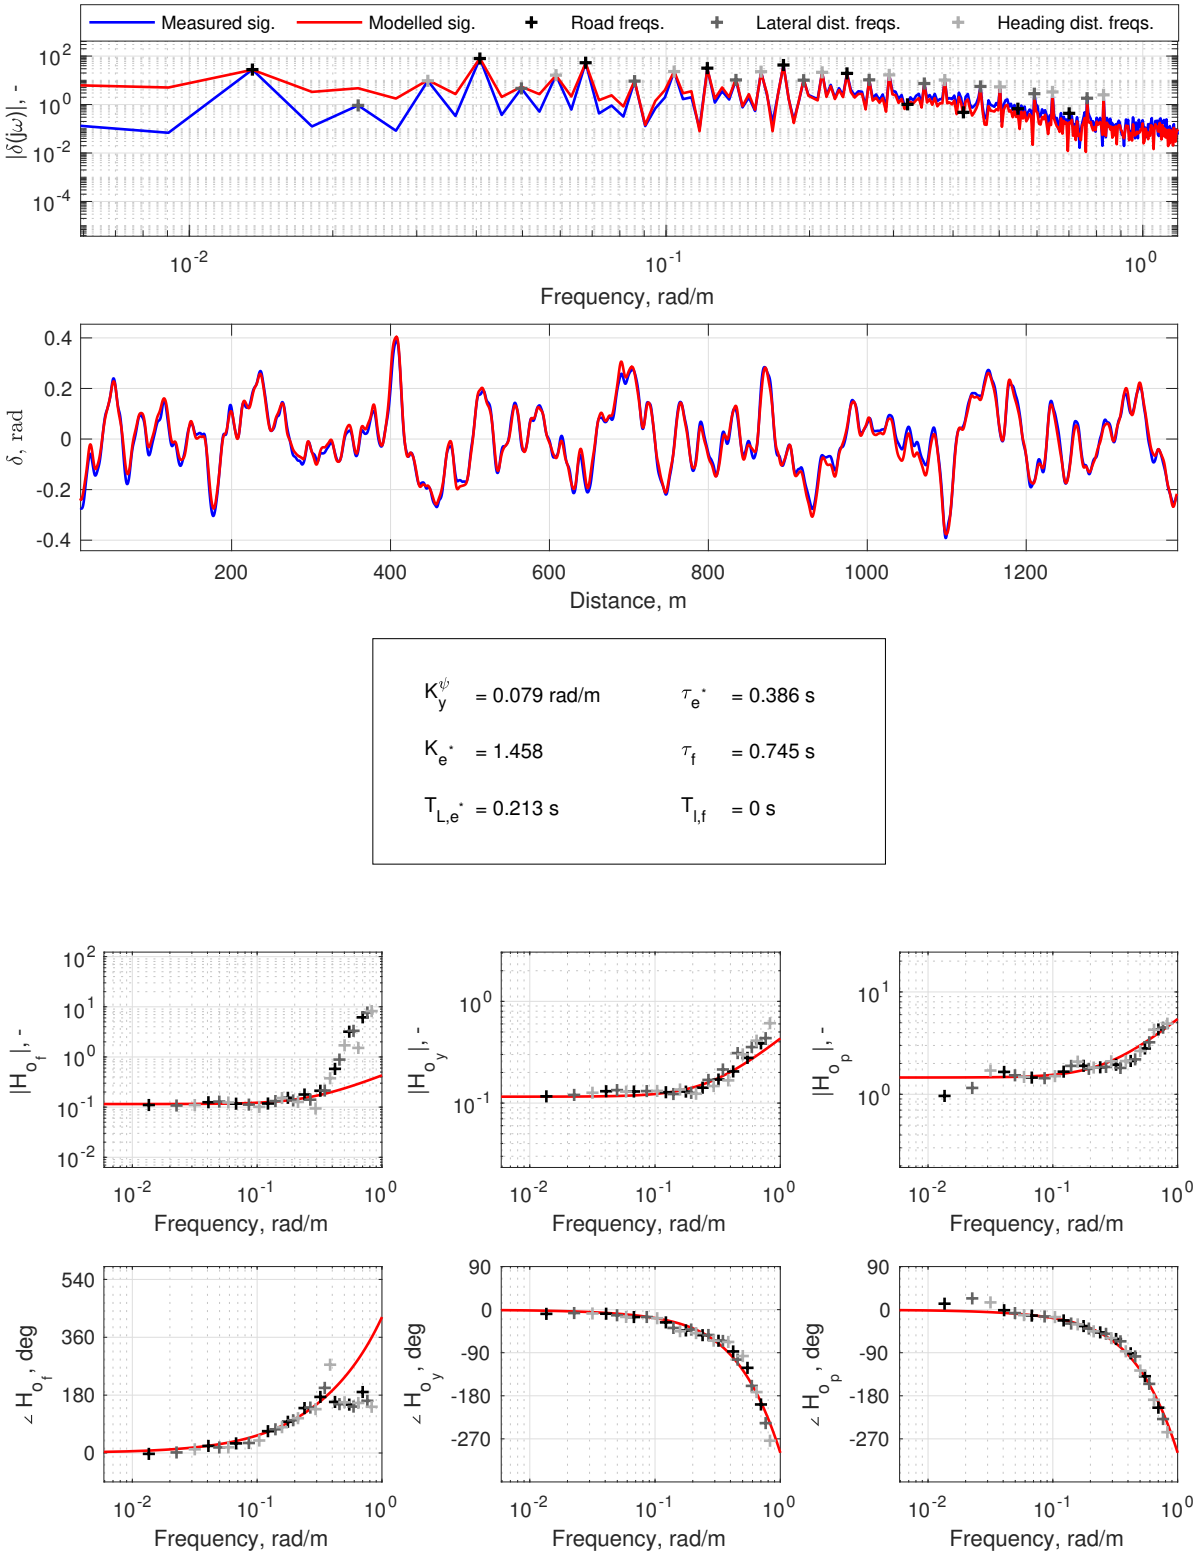

Figure E.6: FRF estimates and parameter estimation, subject 1, Condition S56

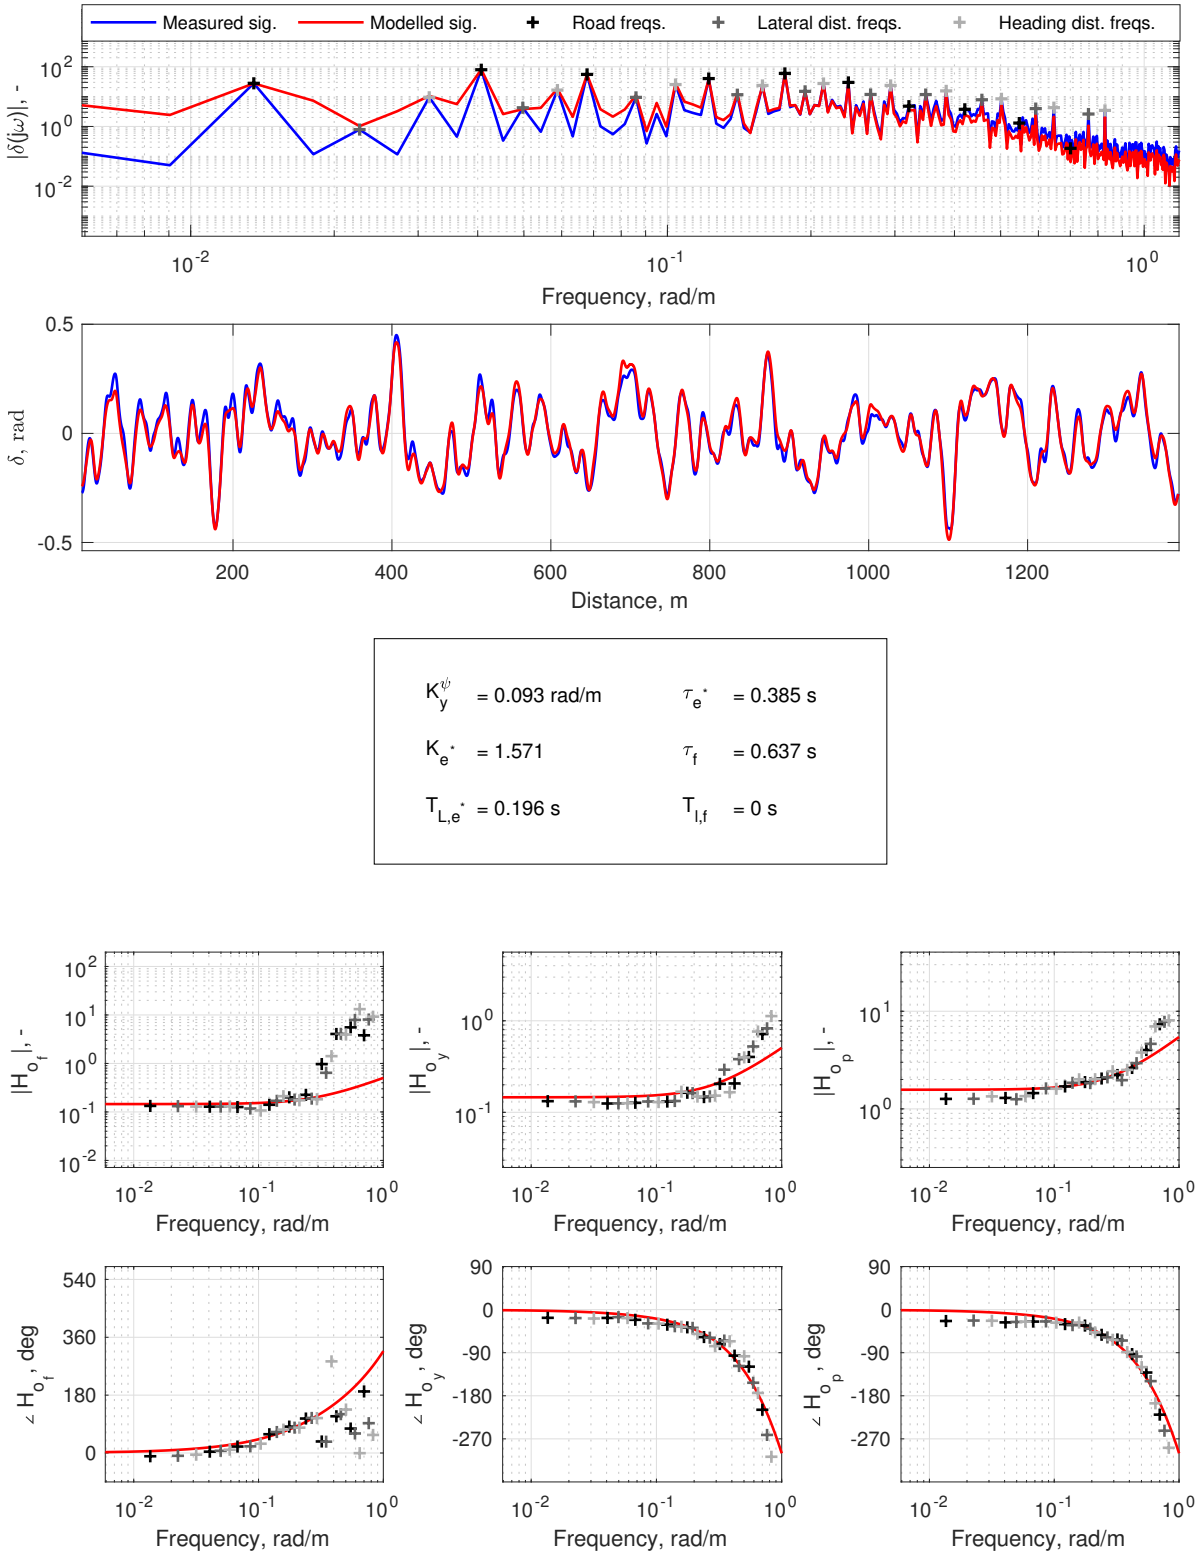

Figure E.7: FRF estimates and parameter estimation, subject 1, Condition S67

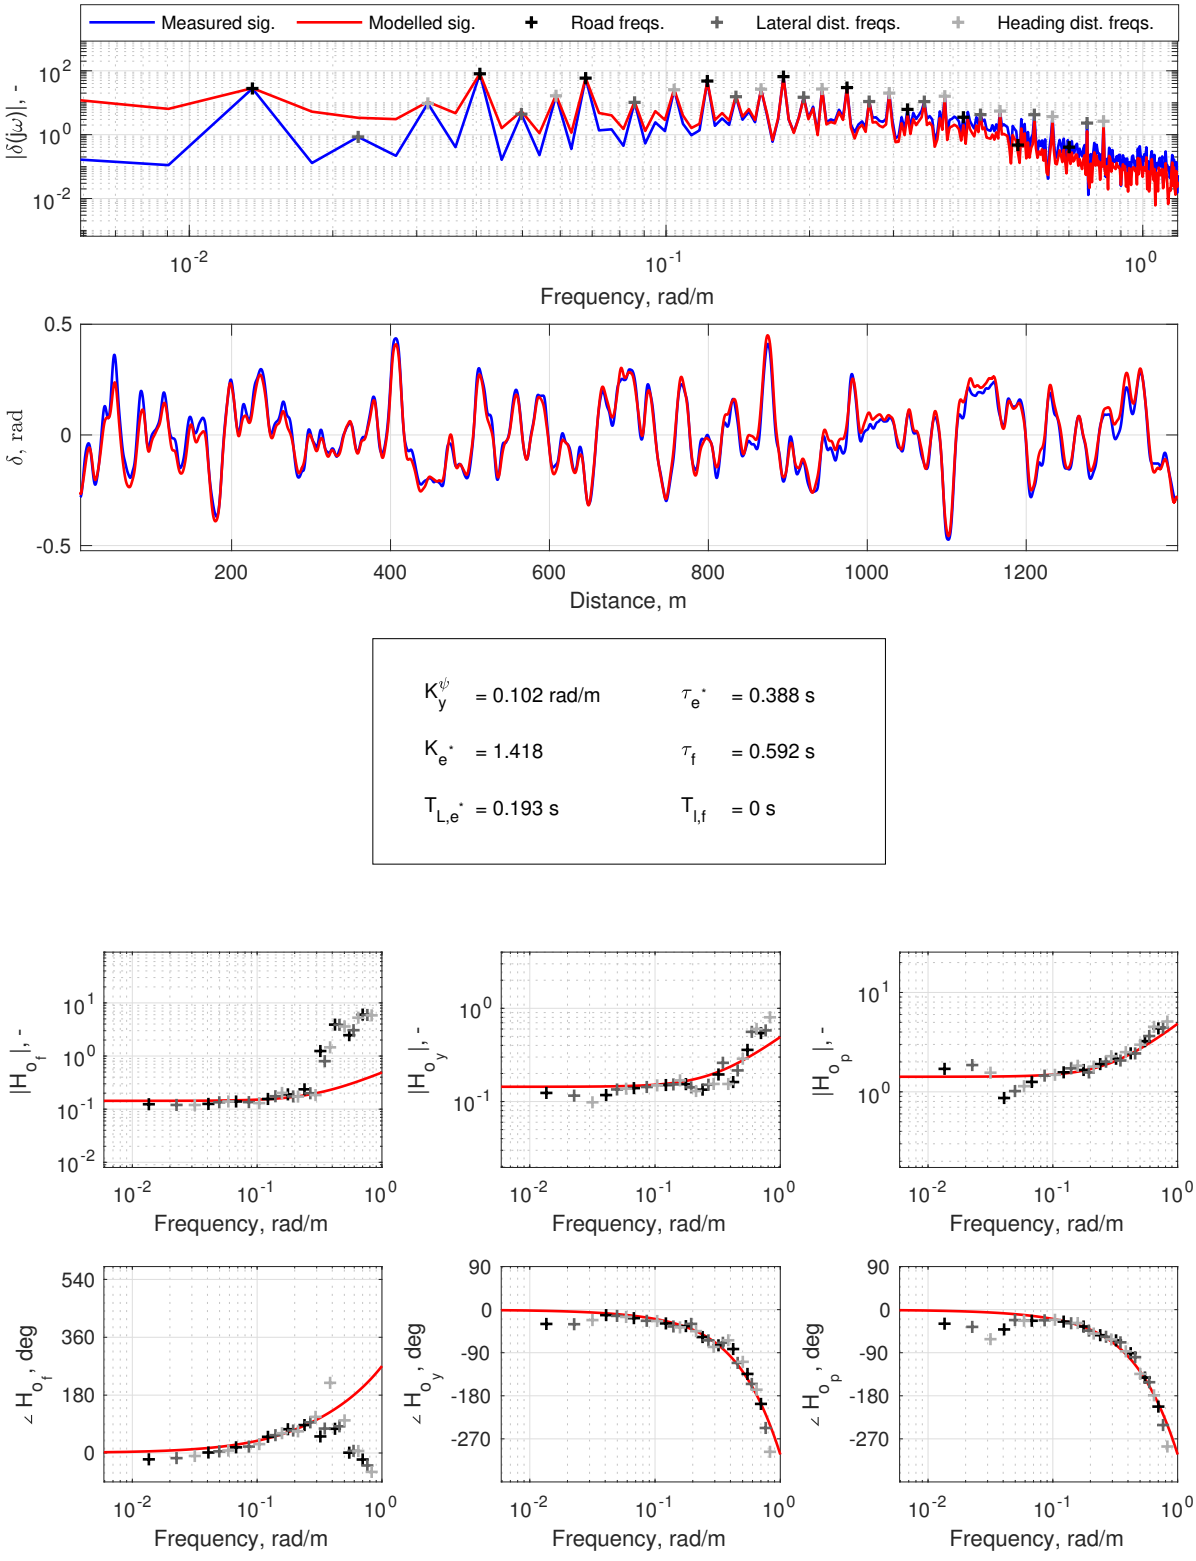

Figure E.8: FRF estimates and parameter estimation, subject 1, Condition S78

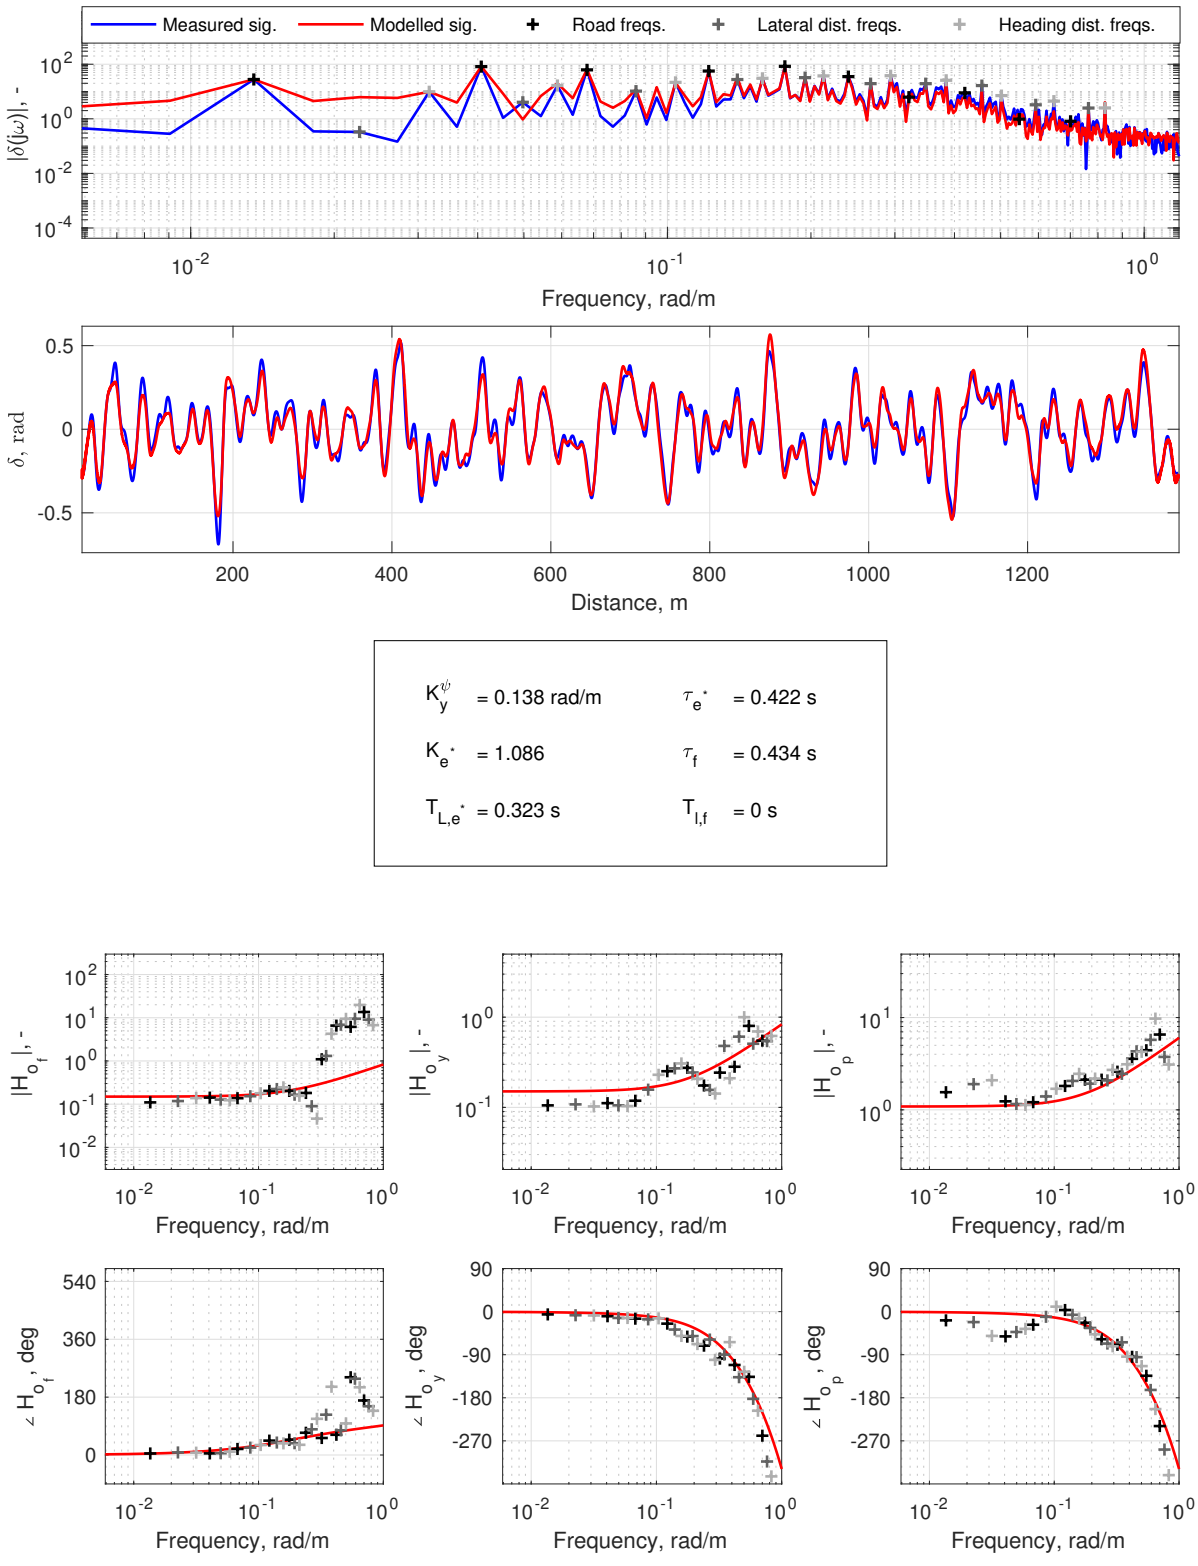

Figure E.9: FRF estimates and parameter estimation, subject 1, Condition S89

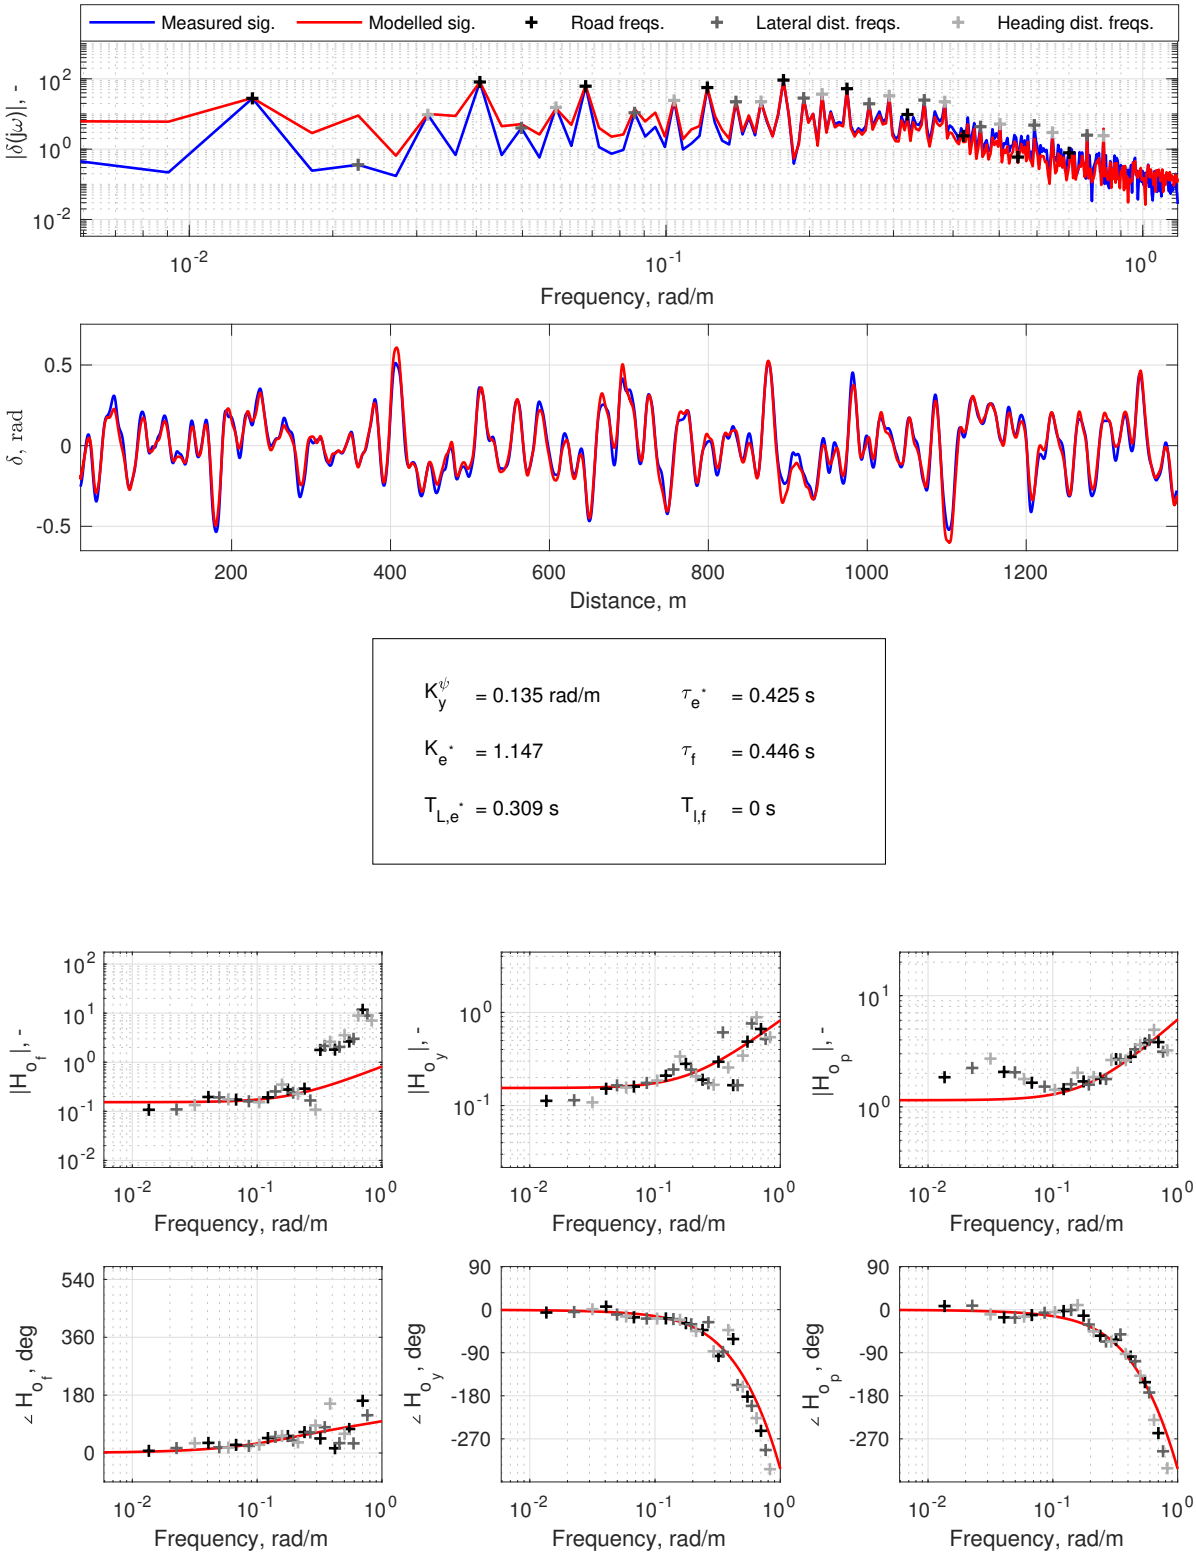

Figure E.10: FRF estimates and parameter estimation, subject 1, Condition S910

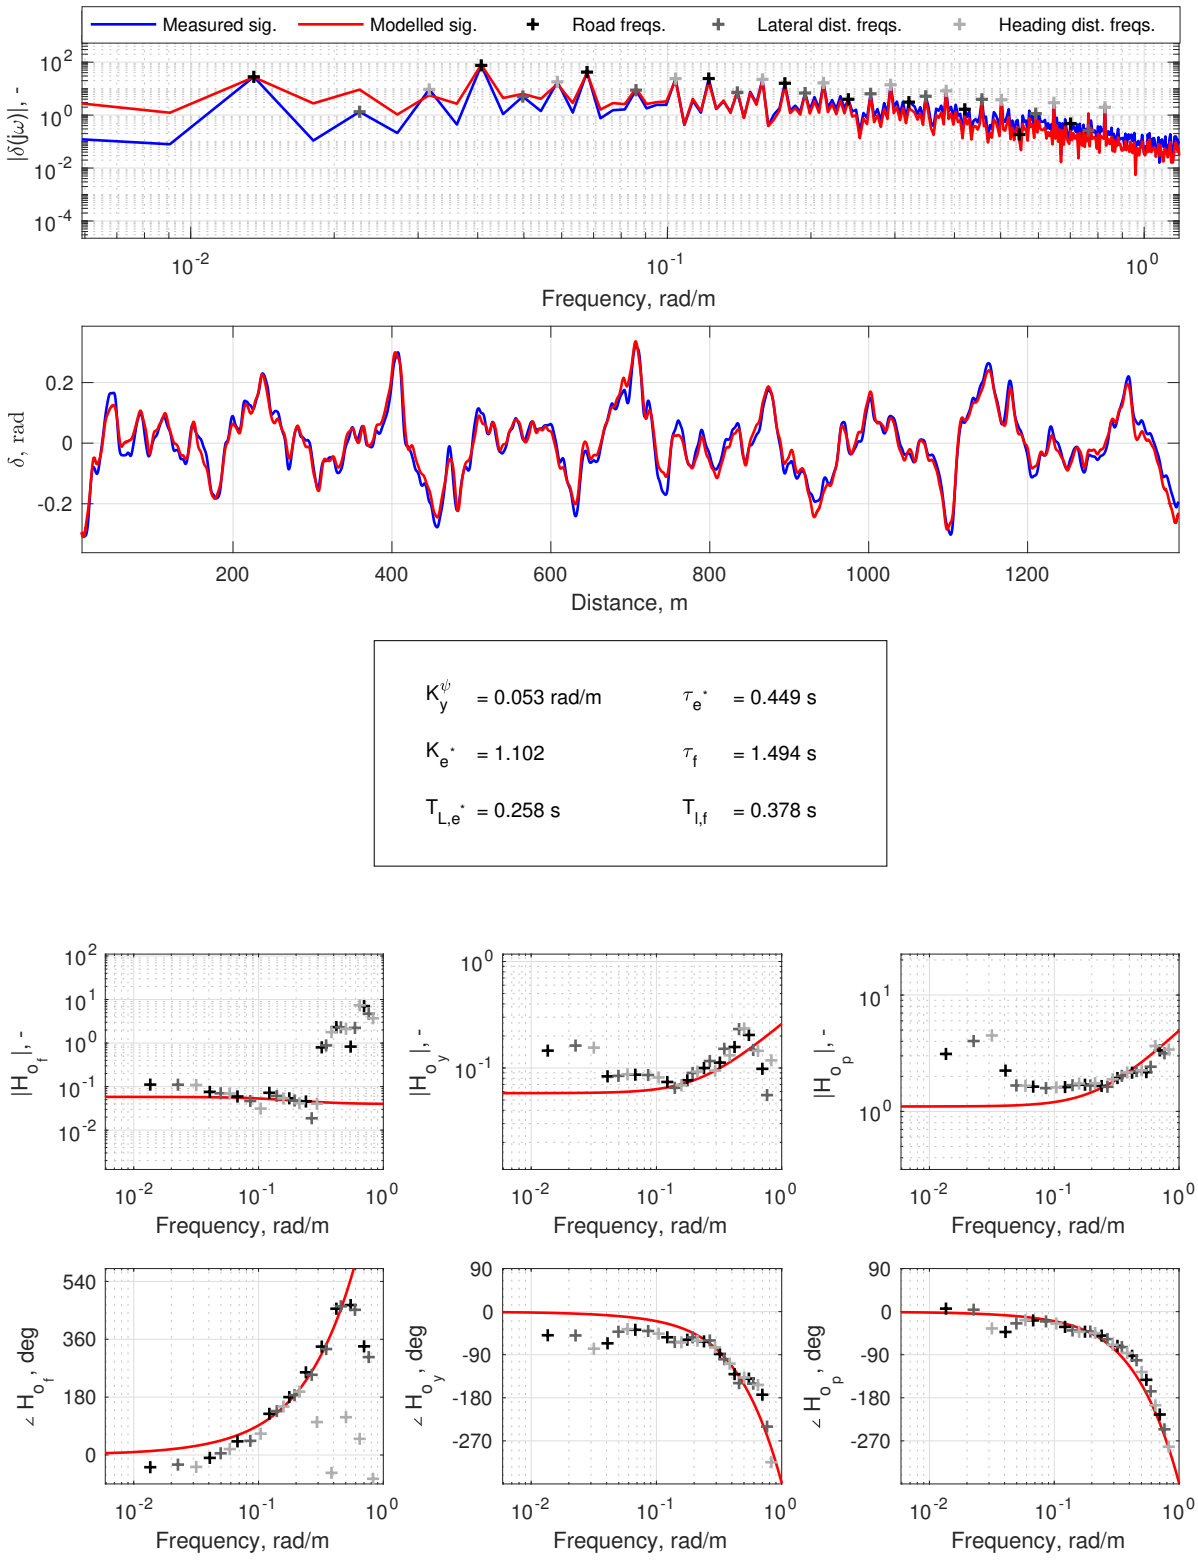

Figure E.11: FRF estimates and parameter estimation, subject 1, Condition T23

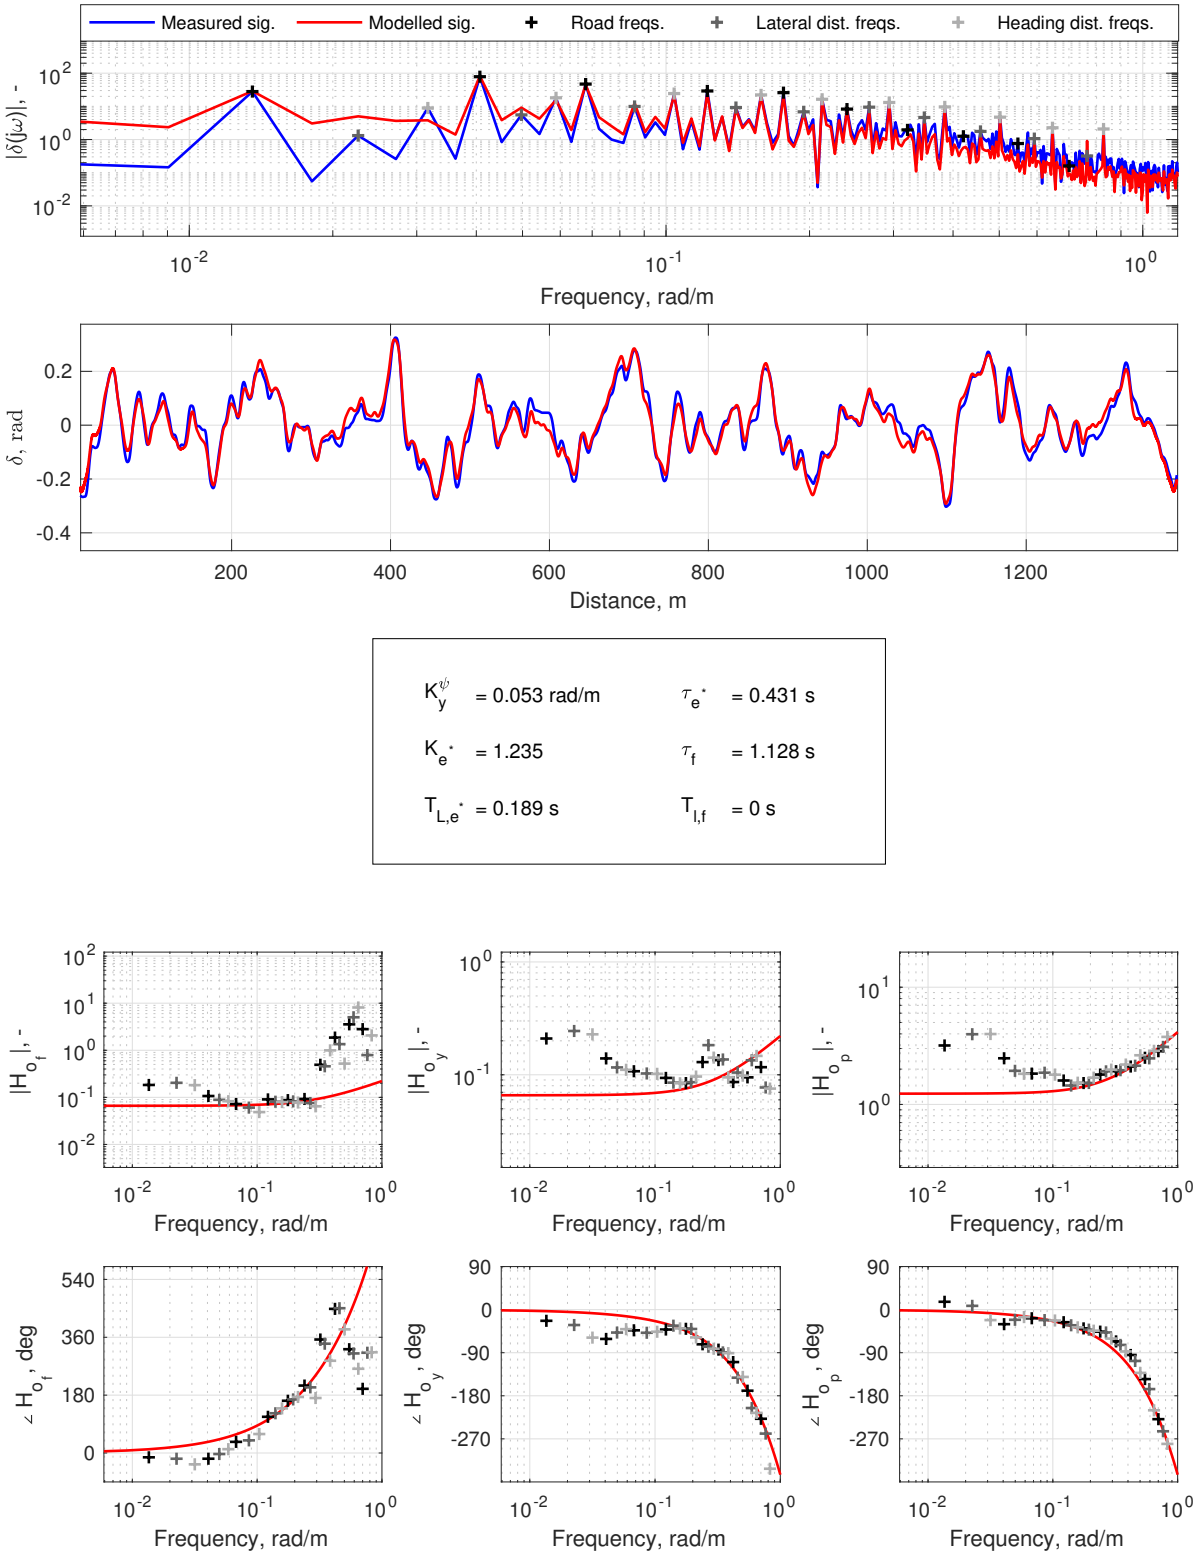

Figure E.12: FRF estimates and parameter estimation, subject 1, Condition T34

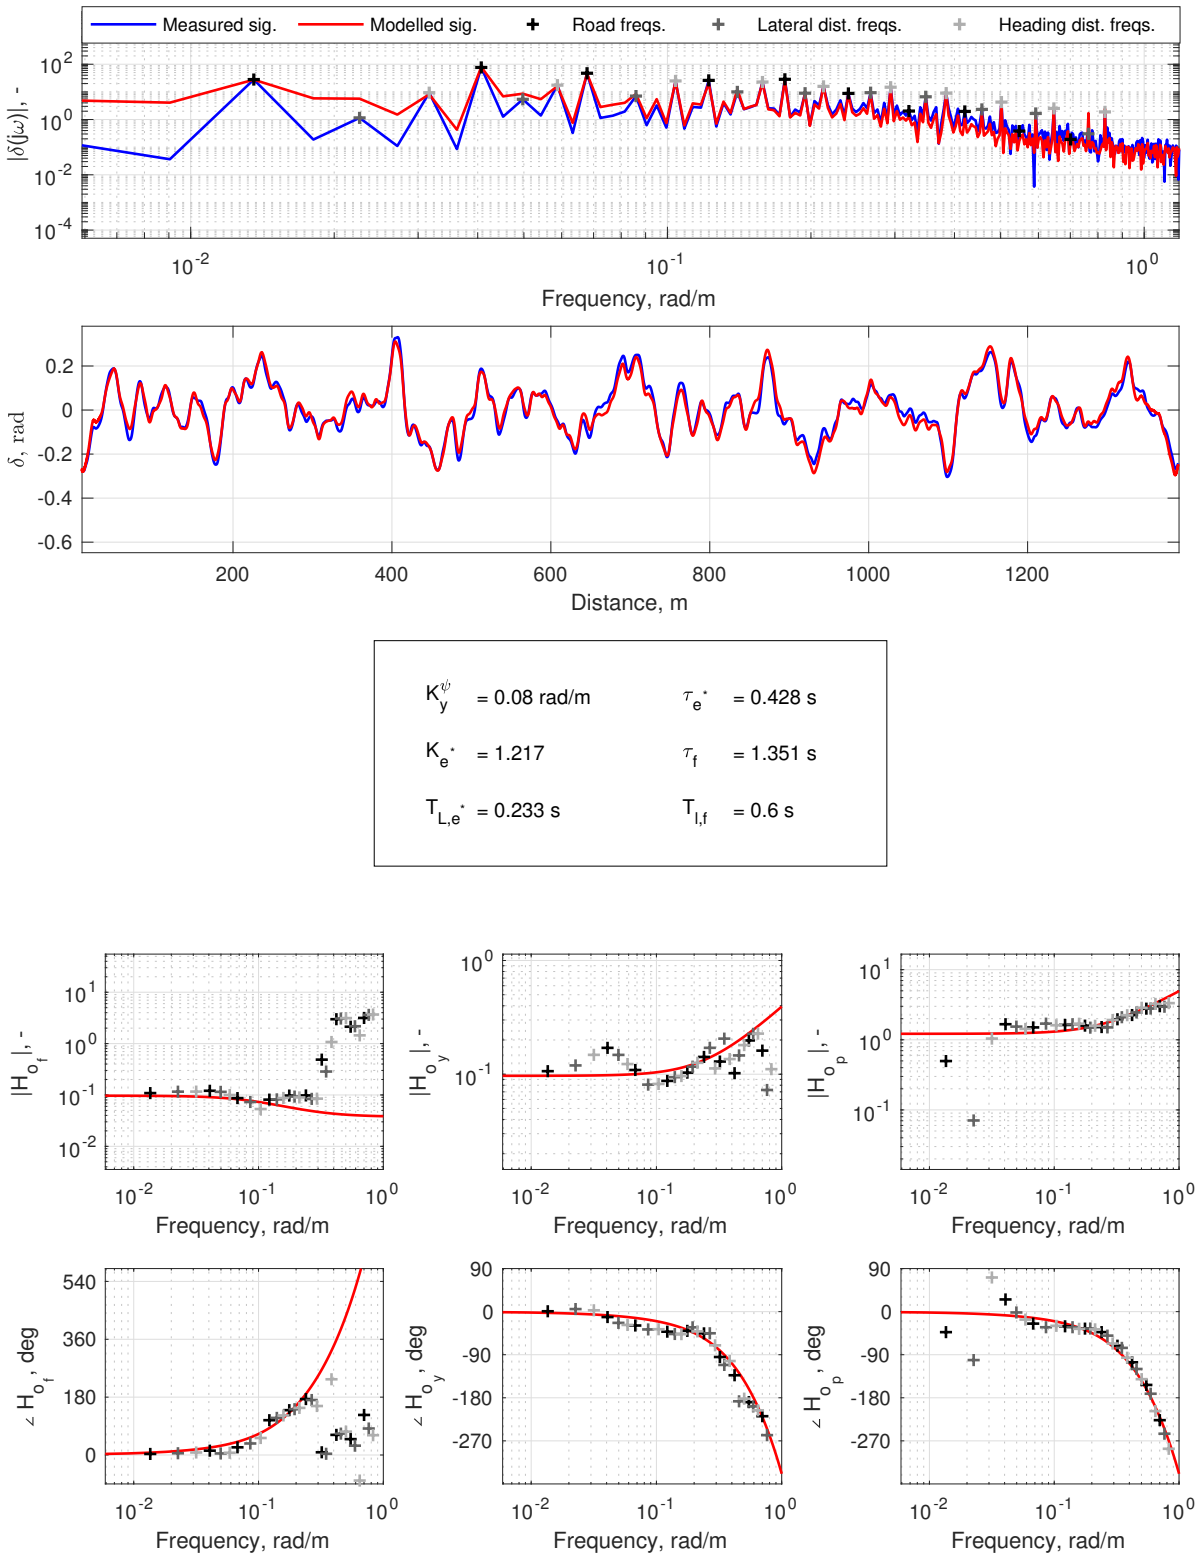

Figure E.13: FRF estimates and parameter estimation, subject 1, Condition T45

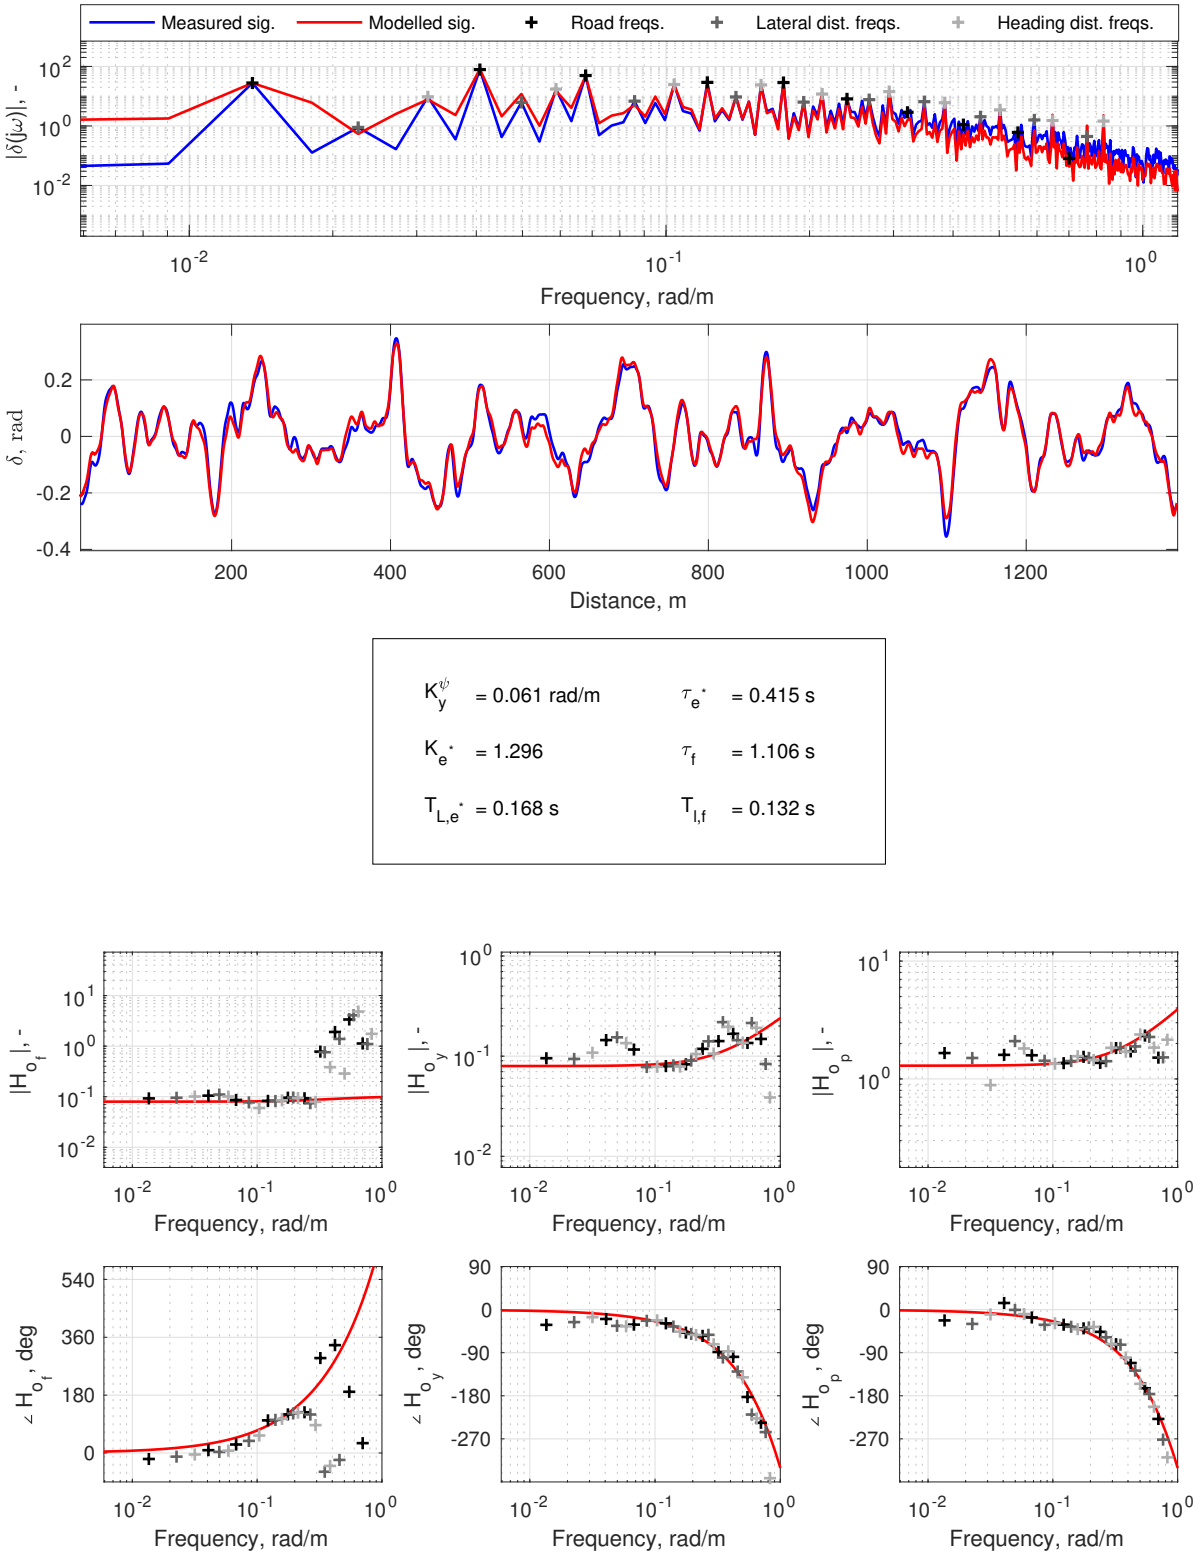

Figure E.14: FRF estimates and parameter estimation, subject 1, Condition T56

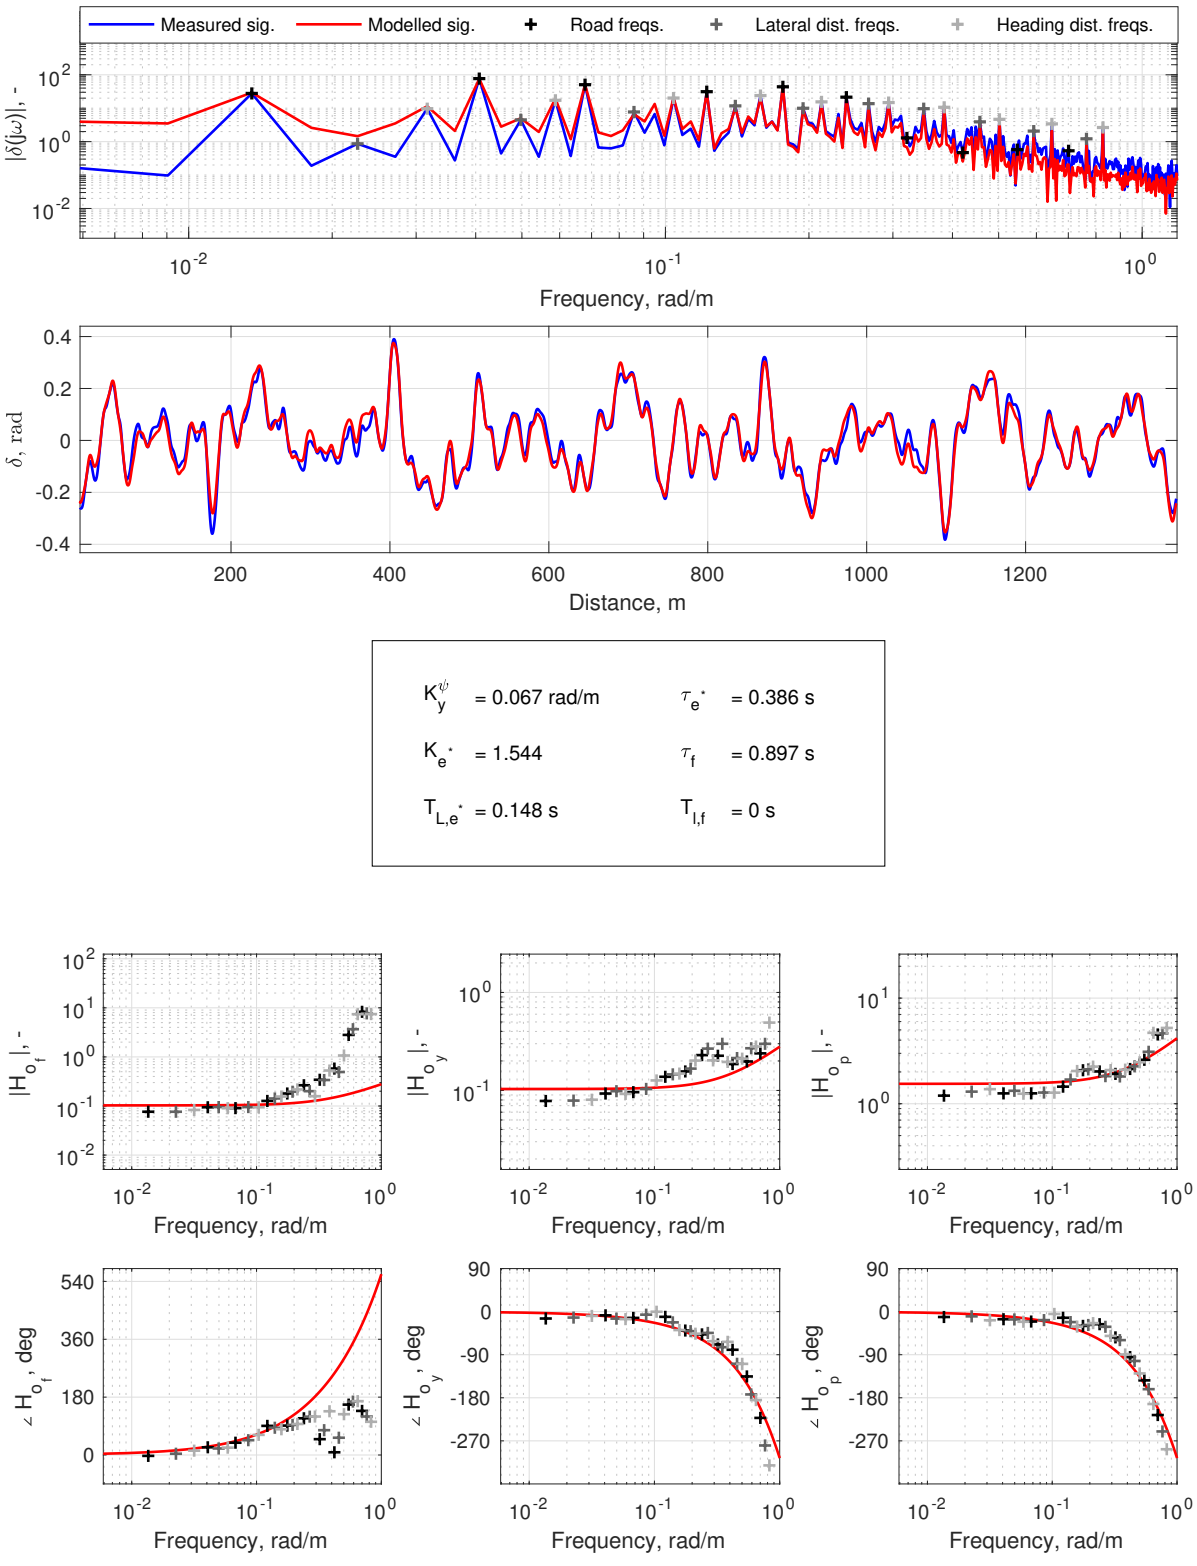

Figure E.15: FRF estimates and parameter estimation, subject 1, Condition T67

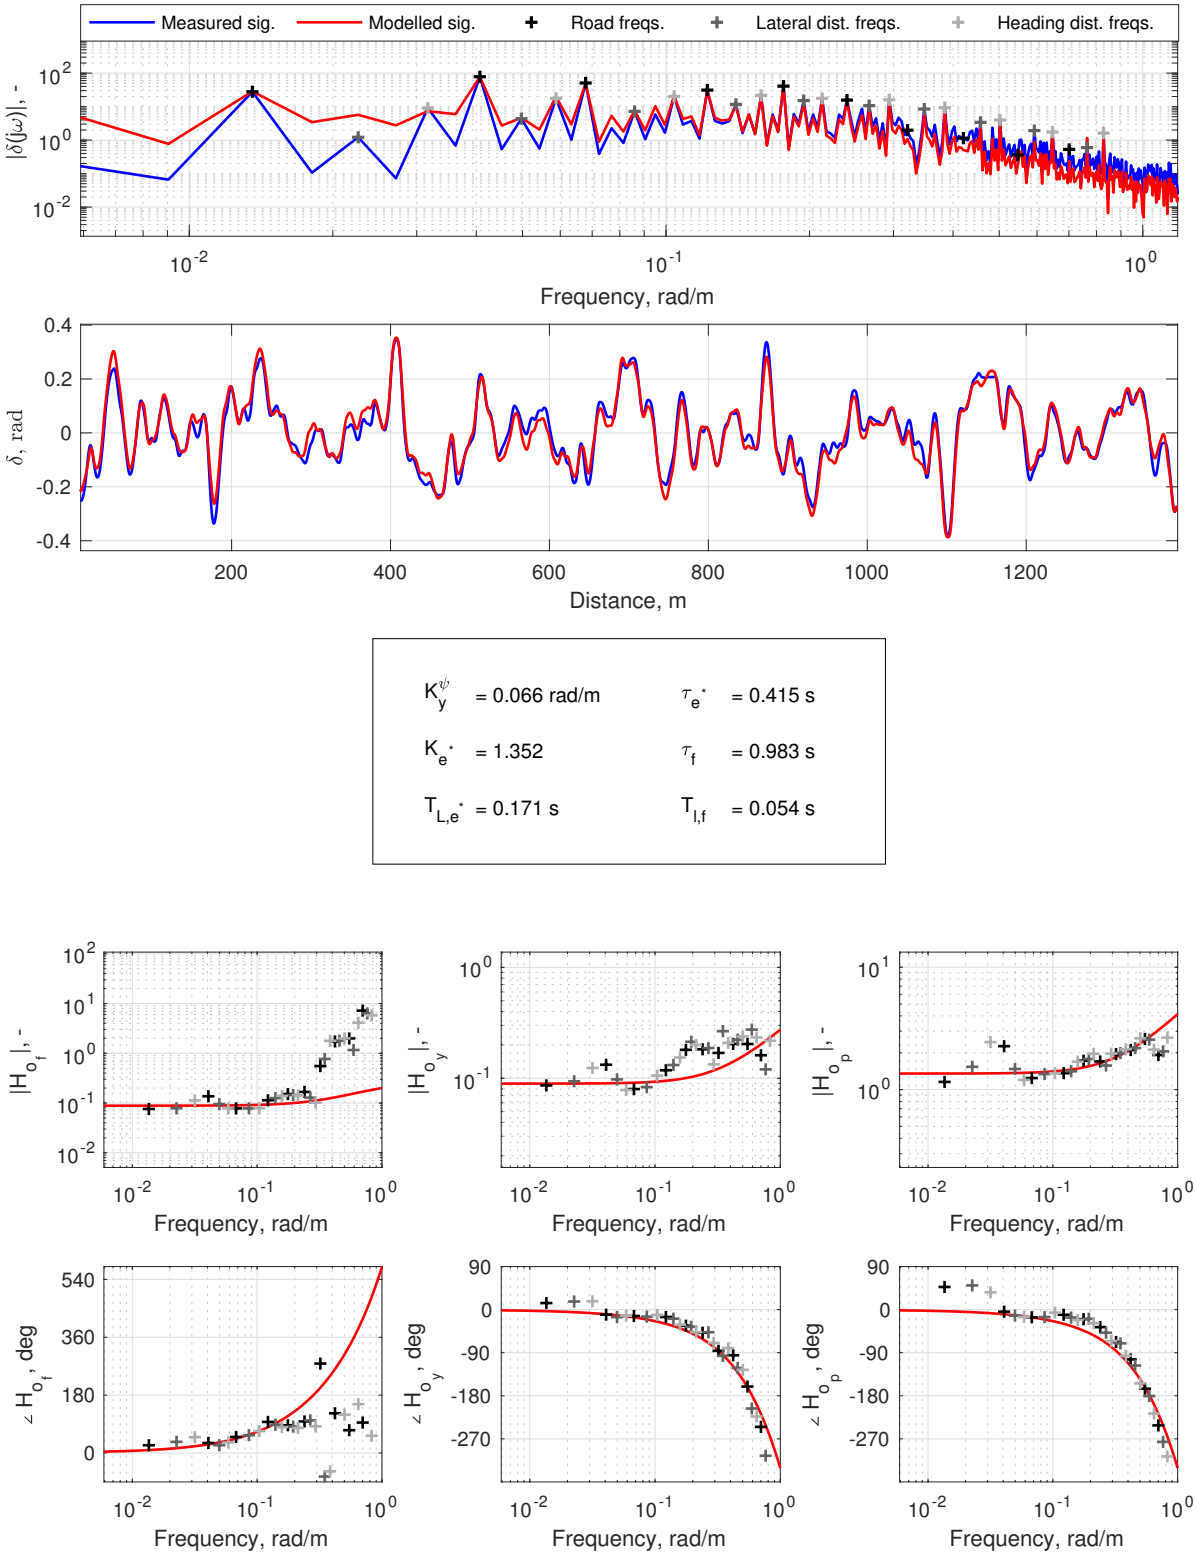

Figure E.16: FRF estimates and parameter estimation, subject 1, Condition T78

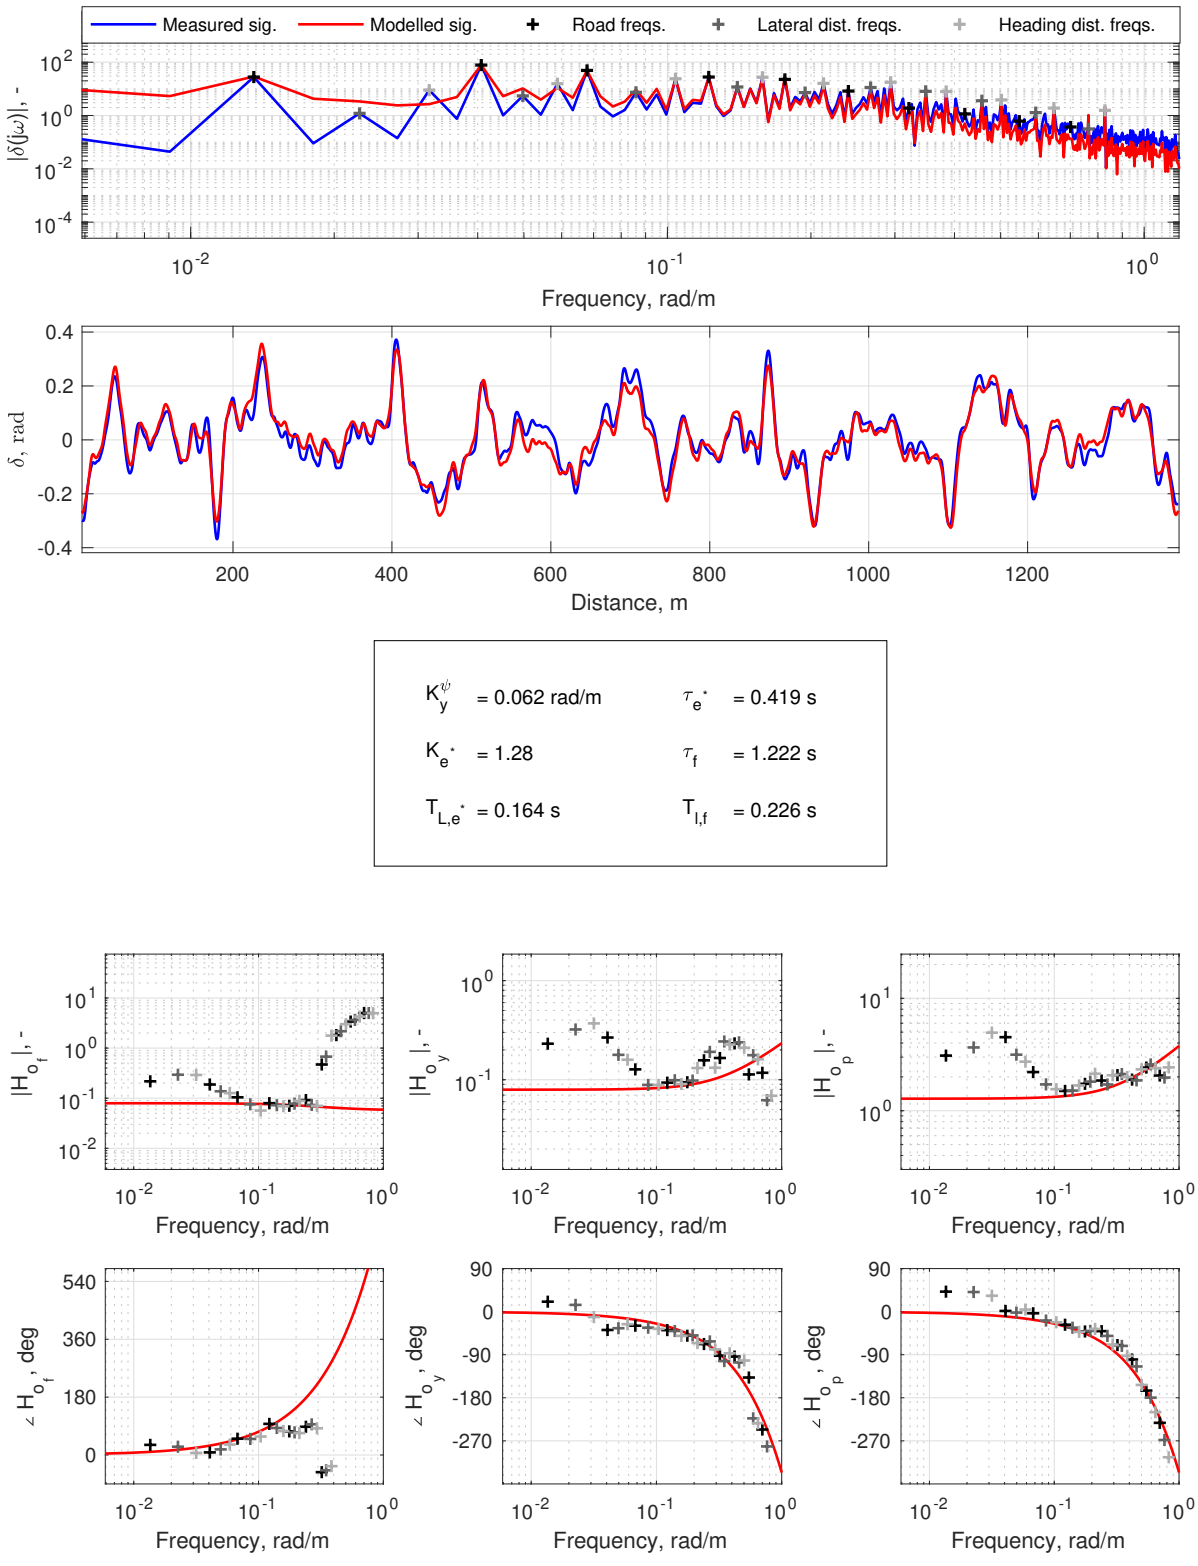

Figure E.17: FRF estimates and parameter estimation, subject 1, Condition T89

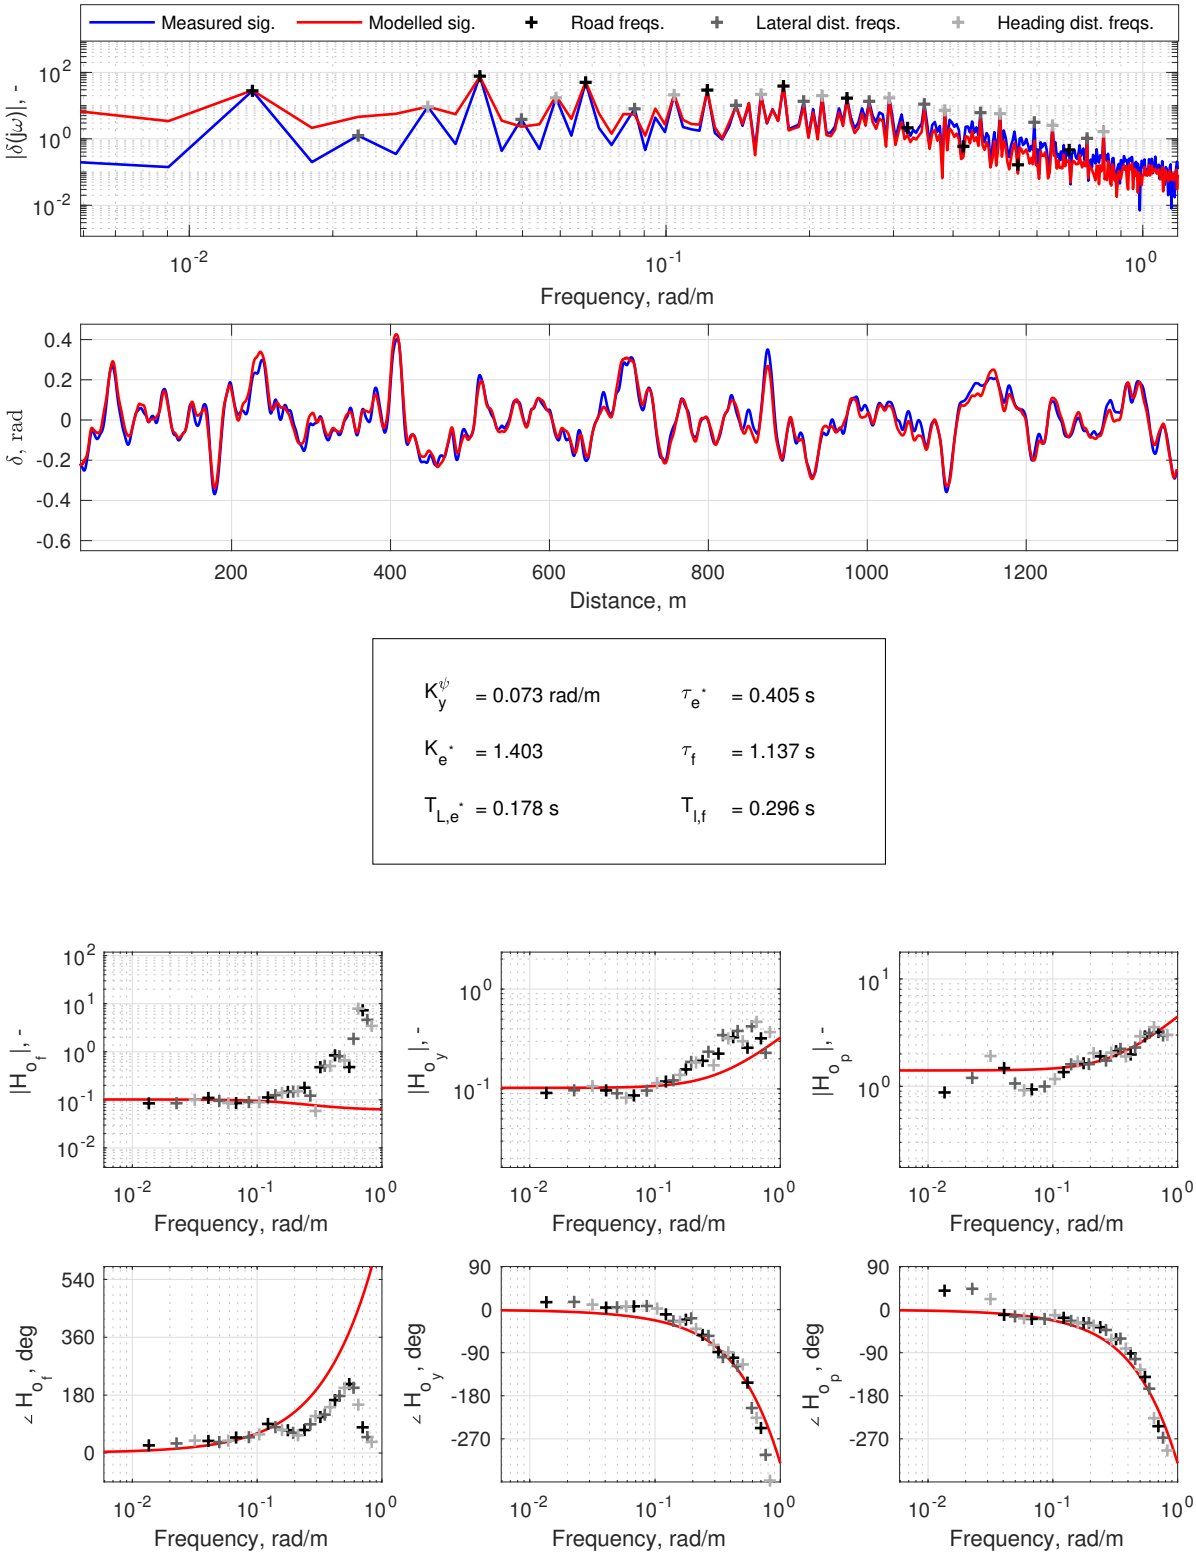

Figure E.18: FRF estimates and parameter estimation, subject 1, Condition T910

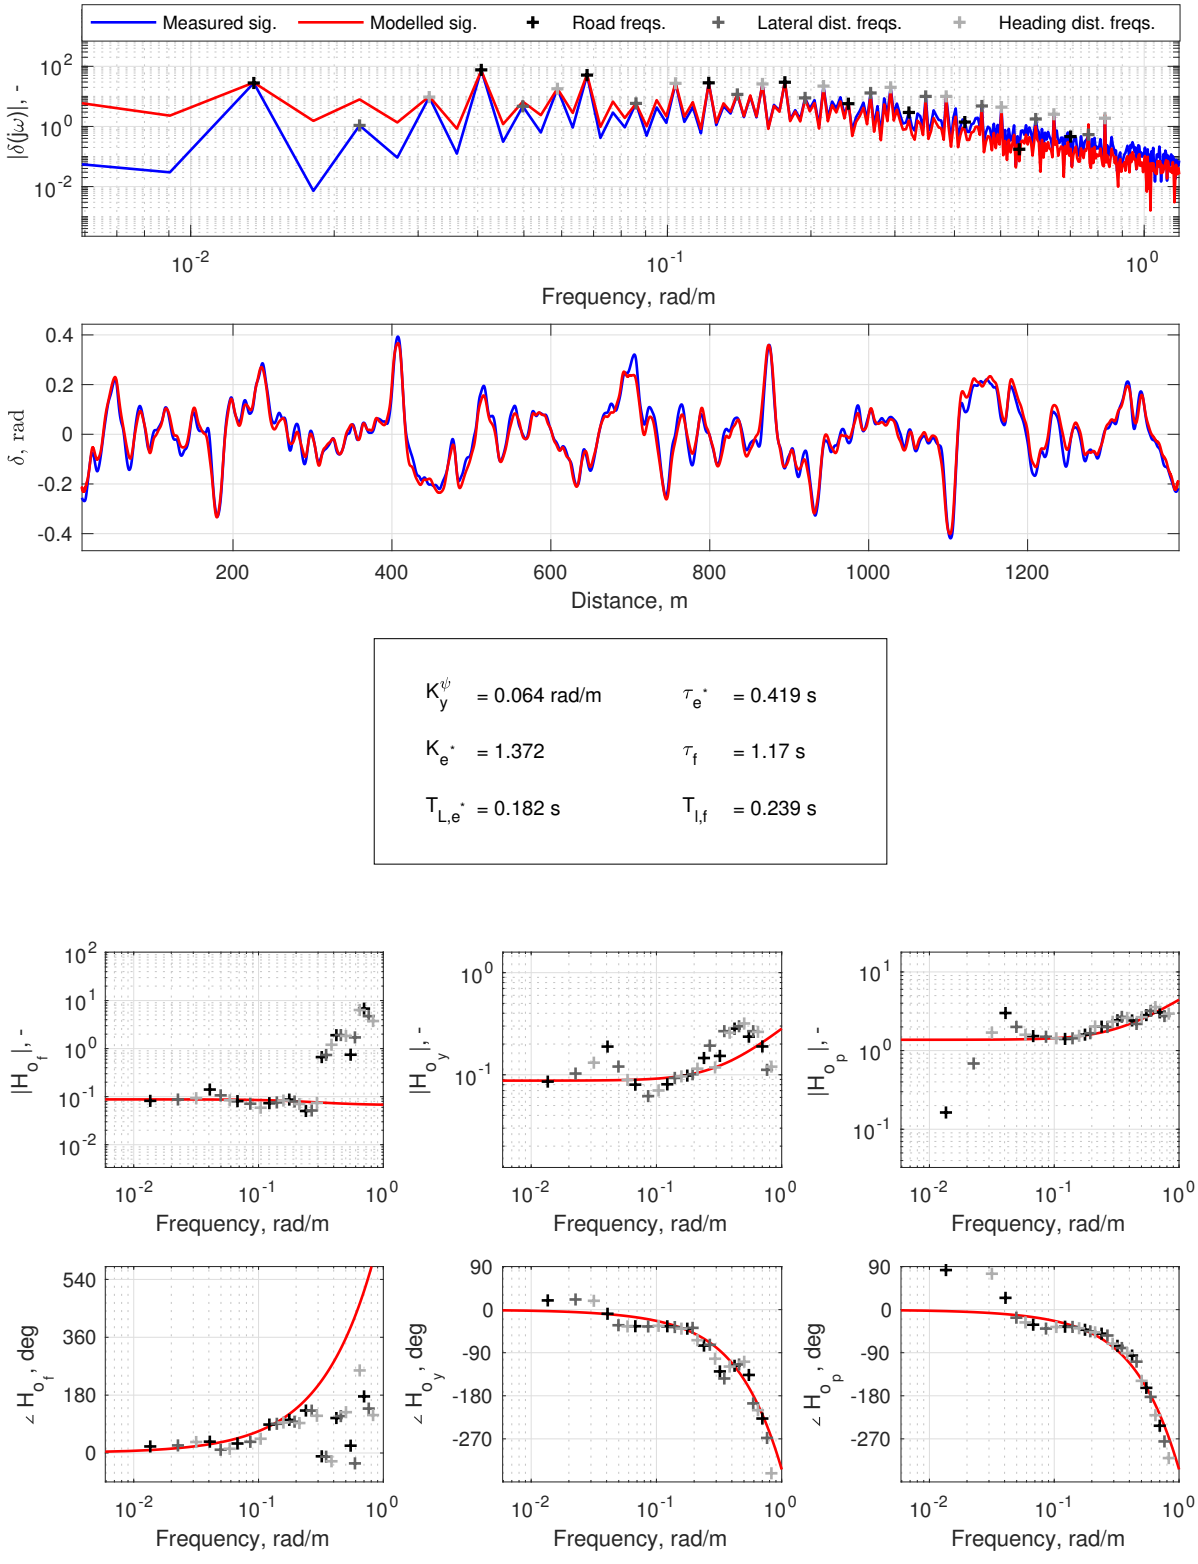

Figure E.19: FRF estimates and parameter estimation, subject 1, Condition B23

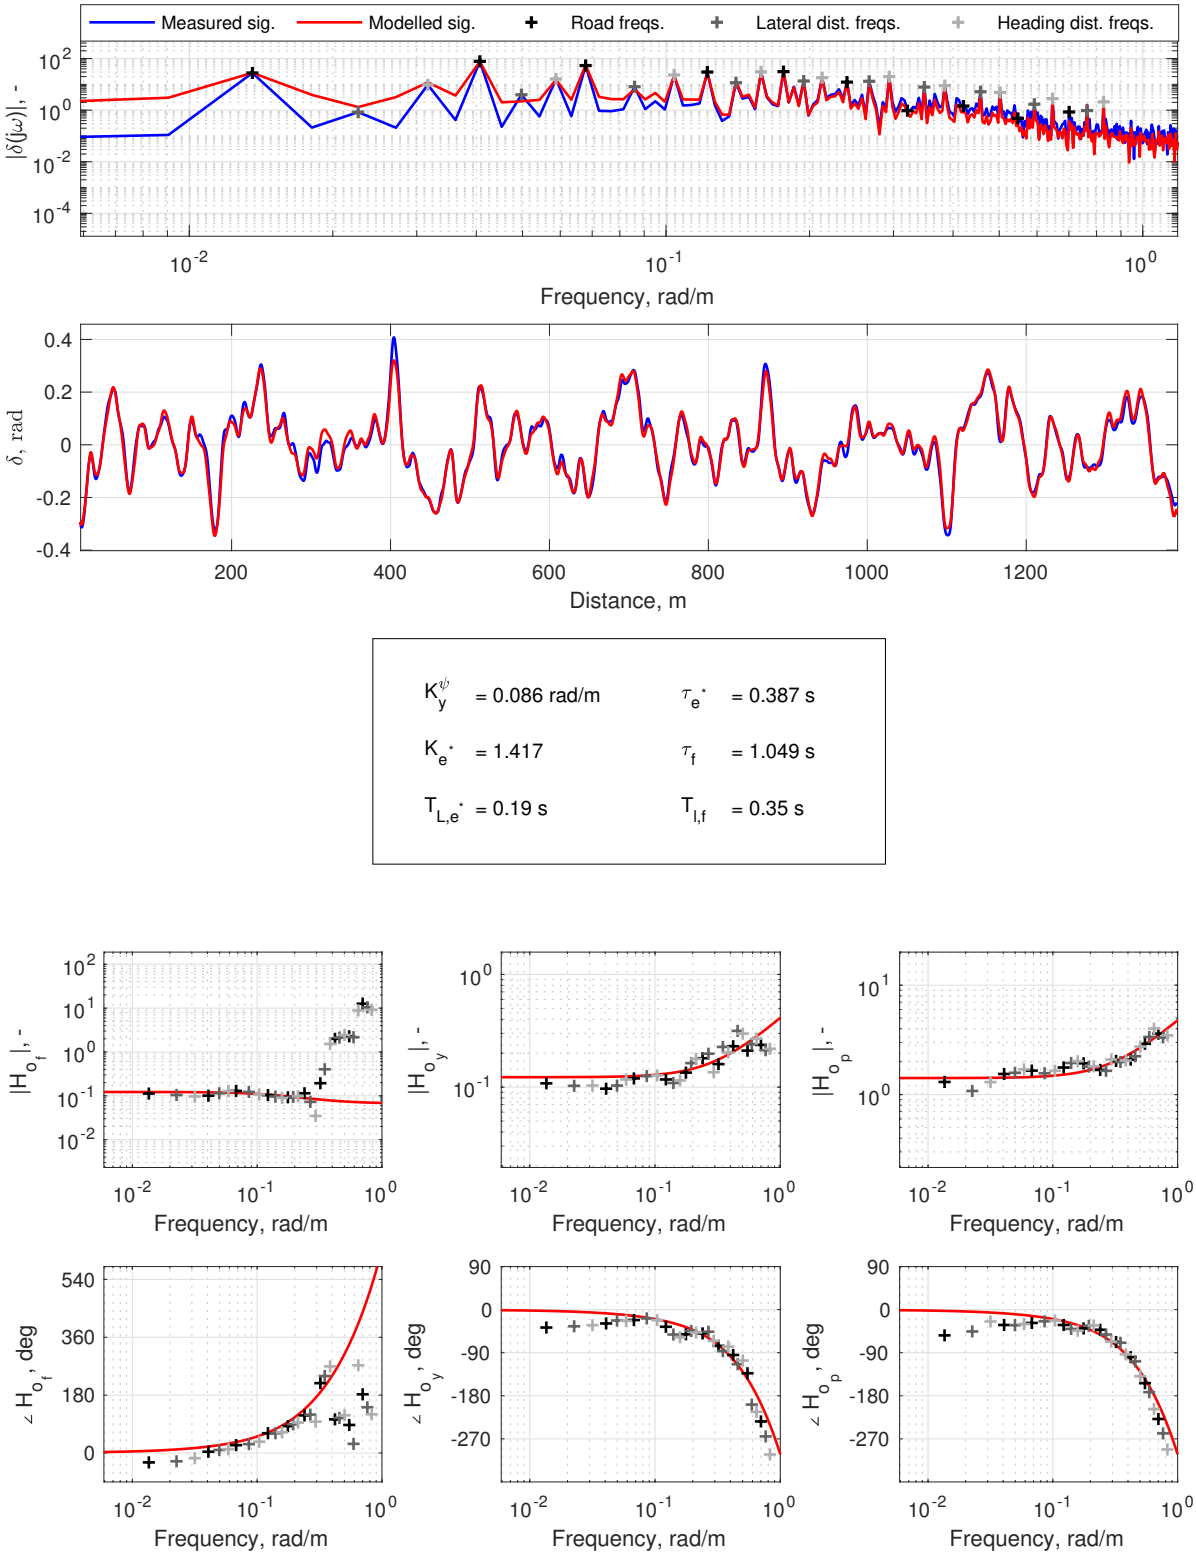

Figure E.20: FRF estimates and parameter estimation, subject 1, Condition B34

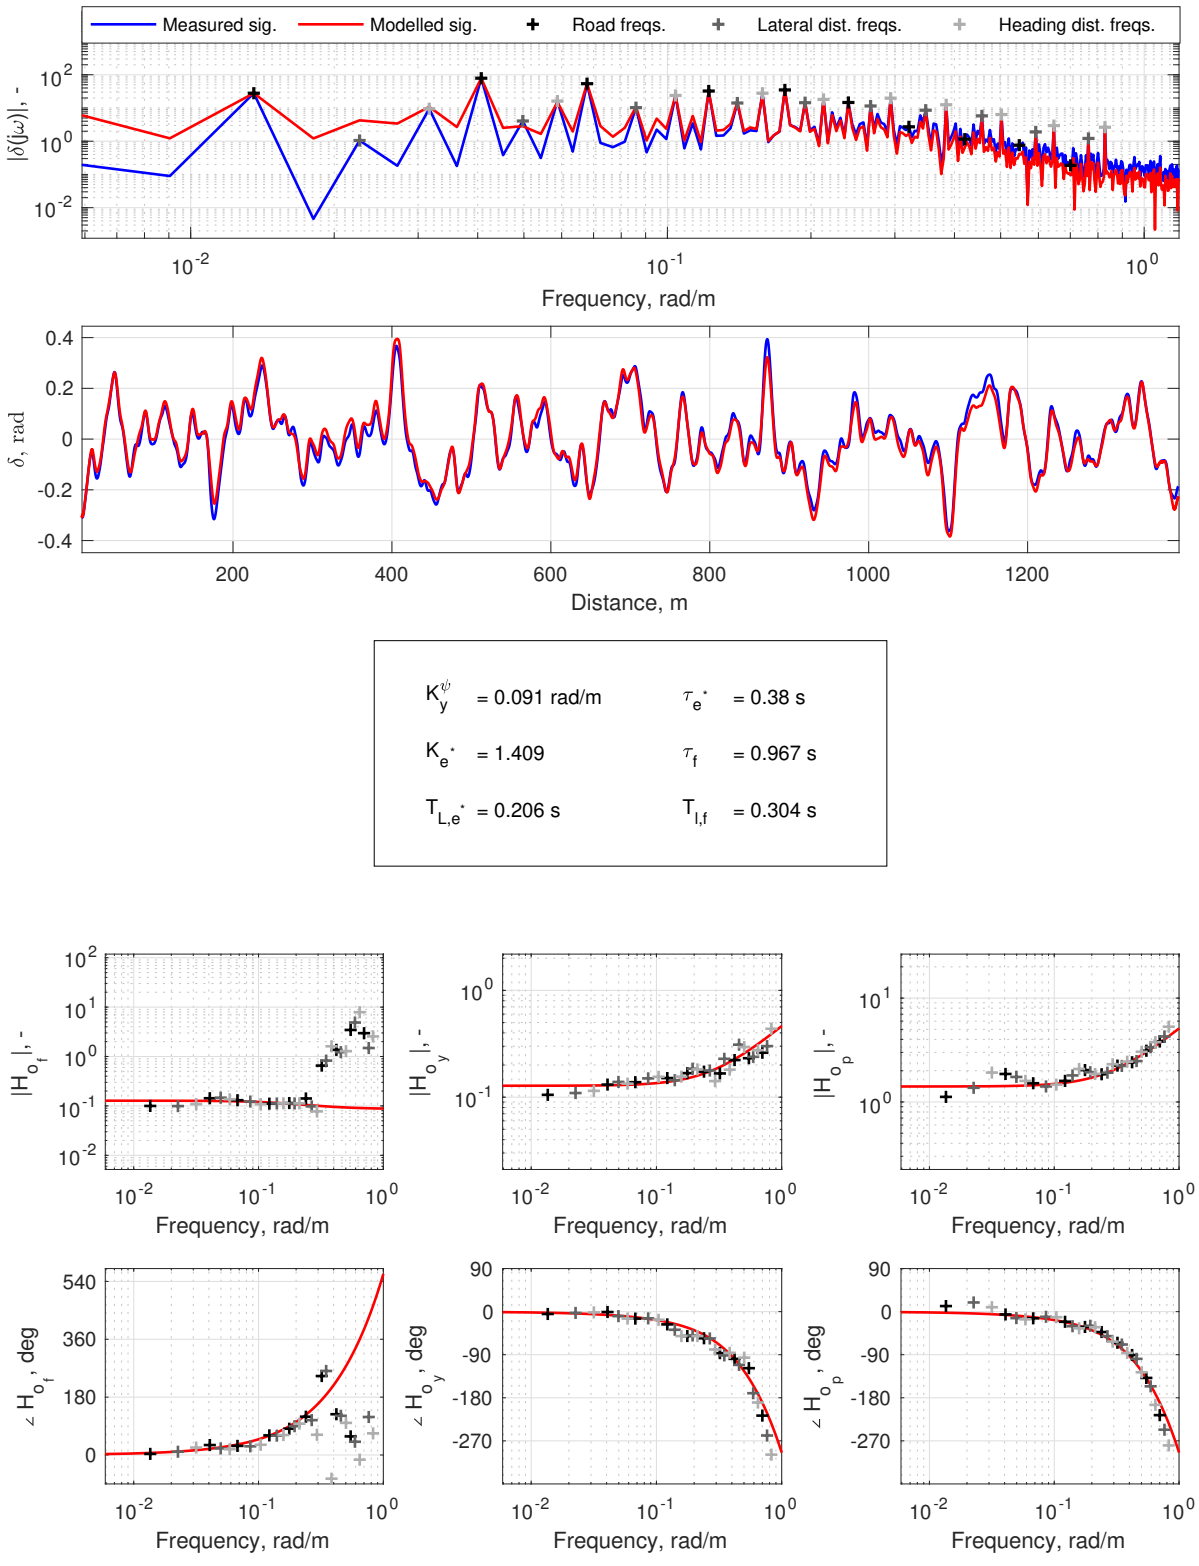

Figure E.21: FRF estimates and parameter estimation, subject 1, Condition B45

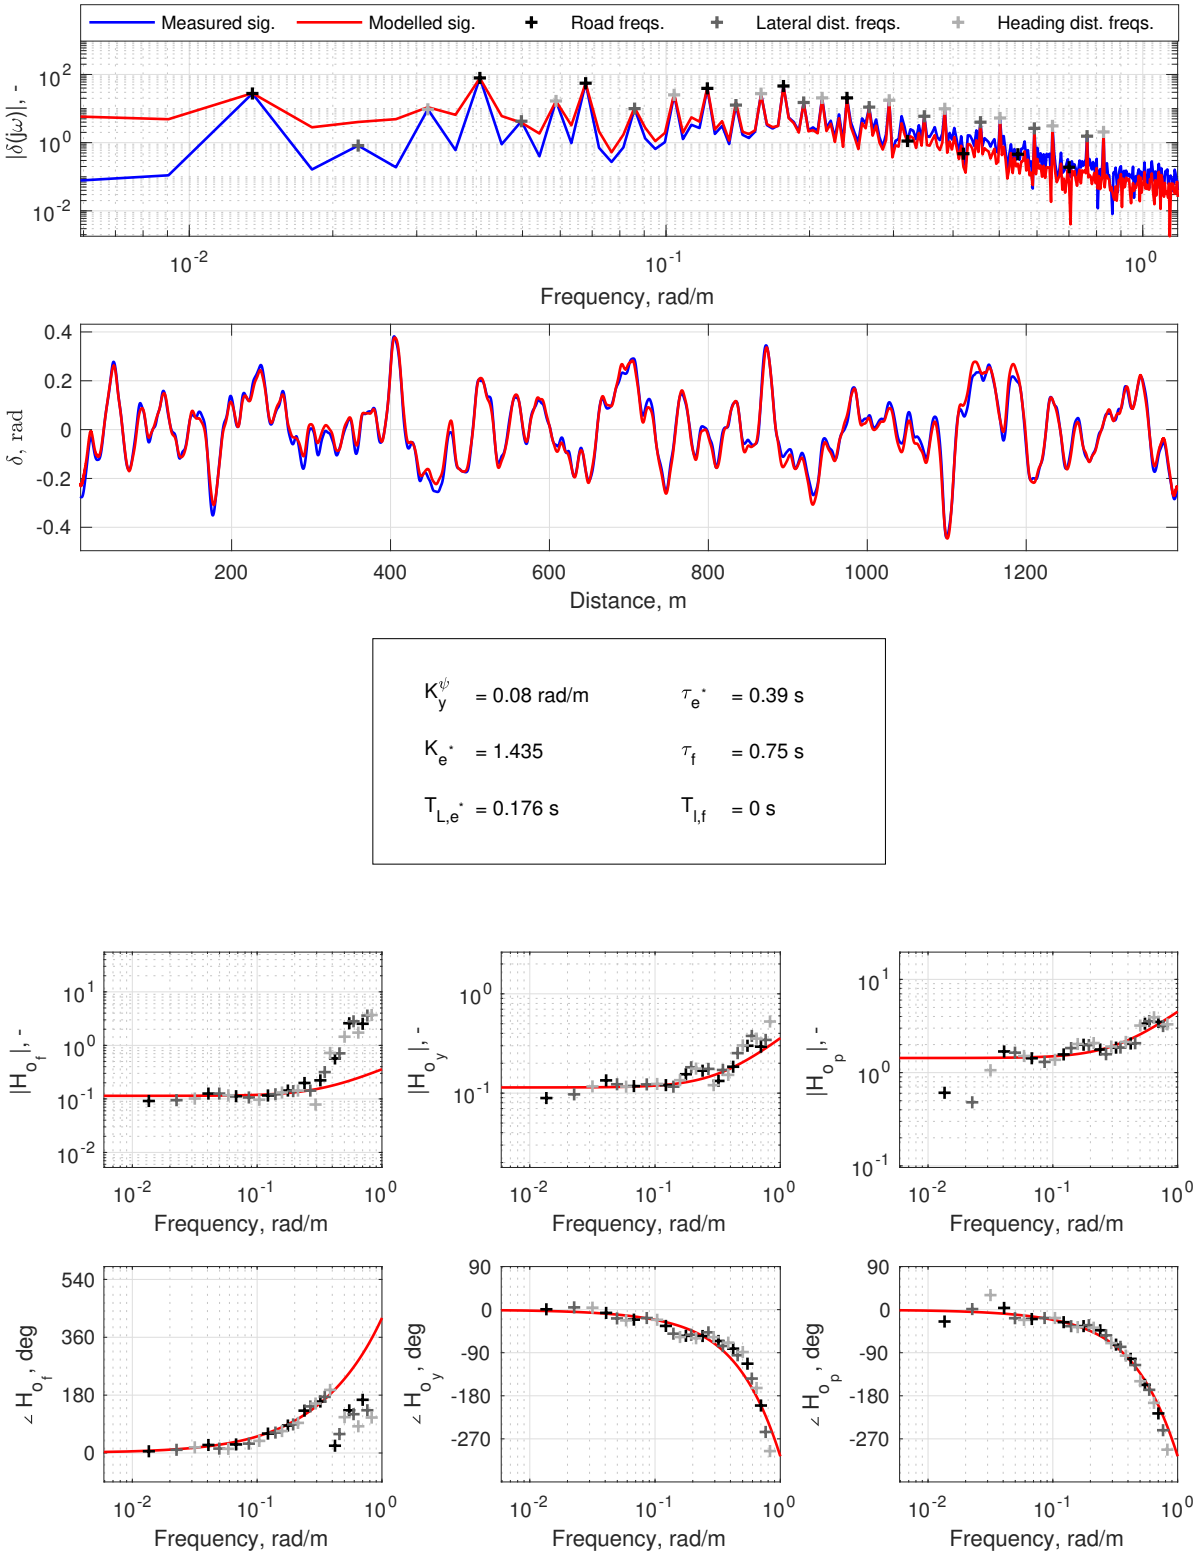

Figure E.22: FRF estimates and parameter estimation, subject 1, Condition B56

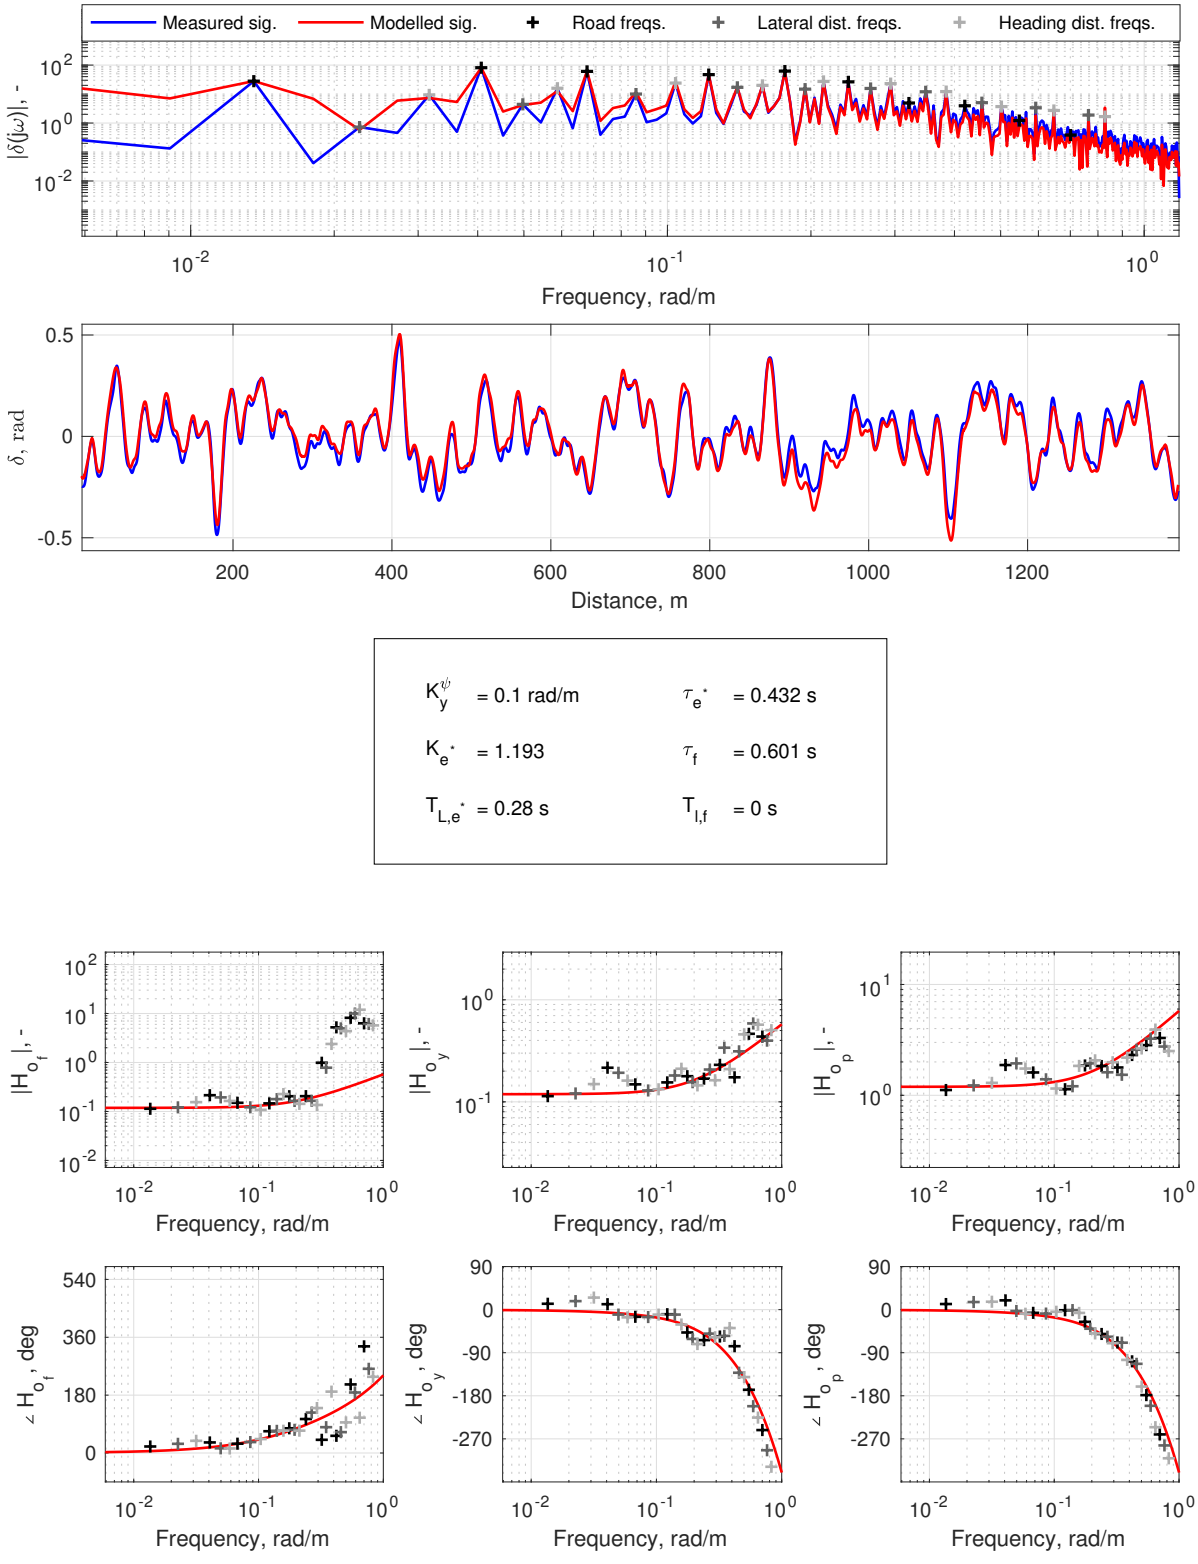

Figure E.23: FRF estimates and parameter estimation, subject 1, Condition B67

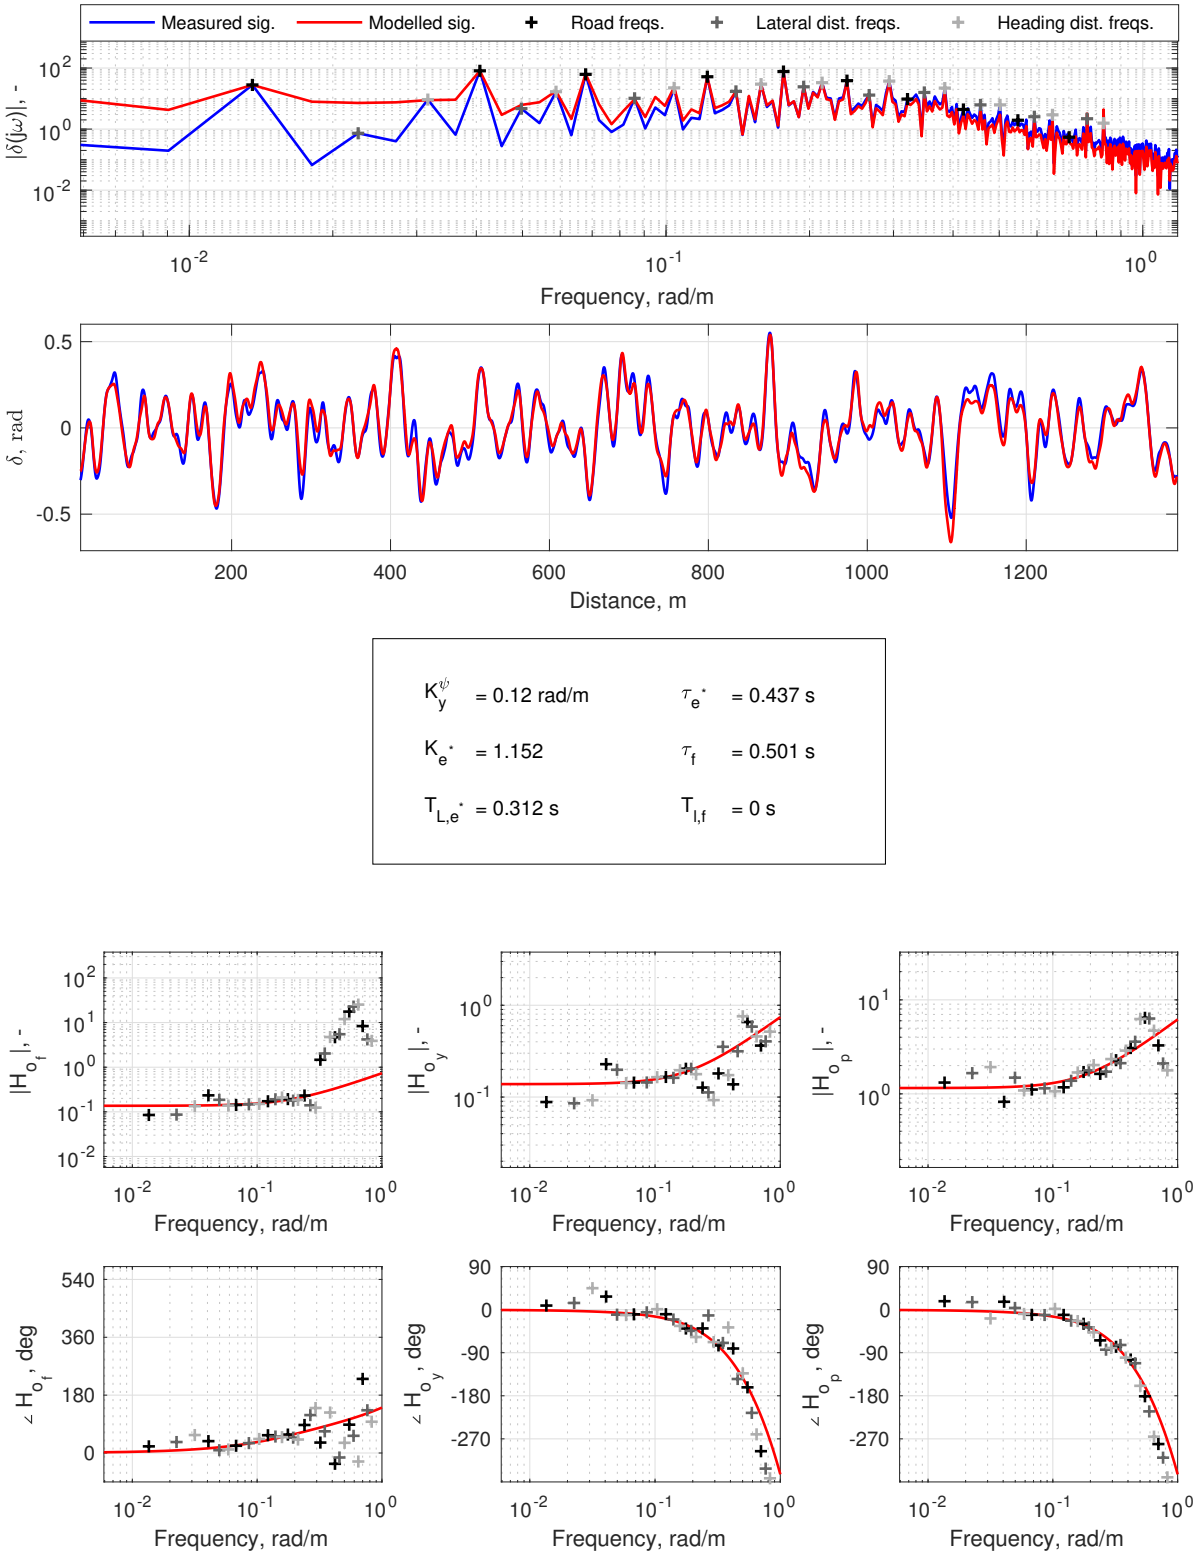

Figure E.24: FRF estimates and parameter estimation, subject 1, Condition B78

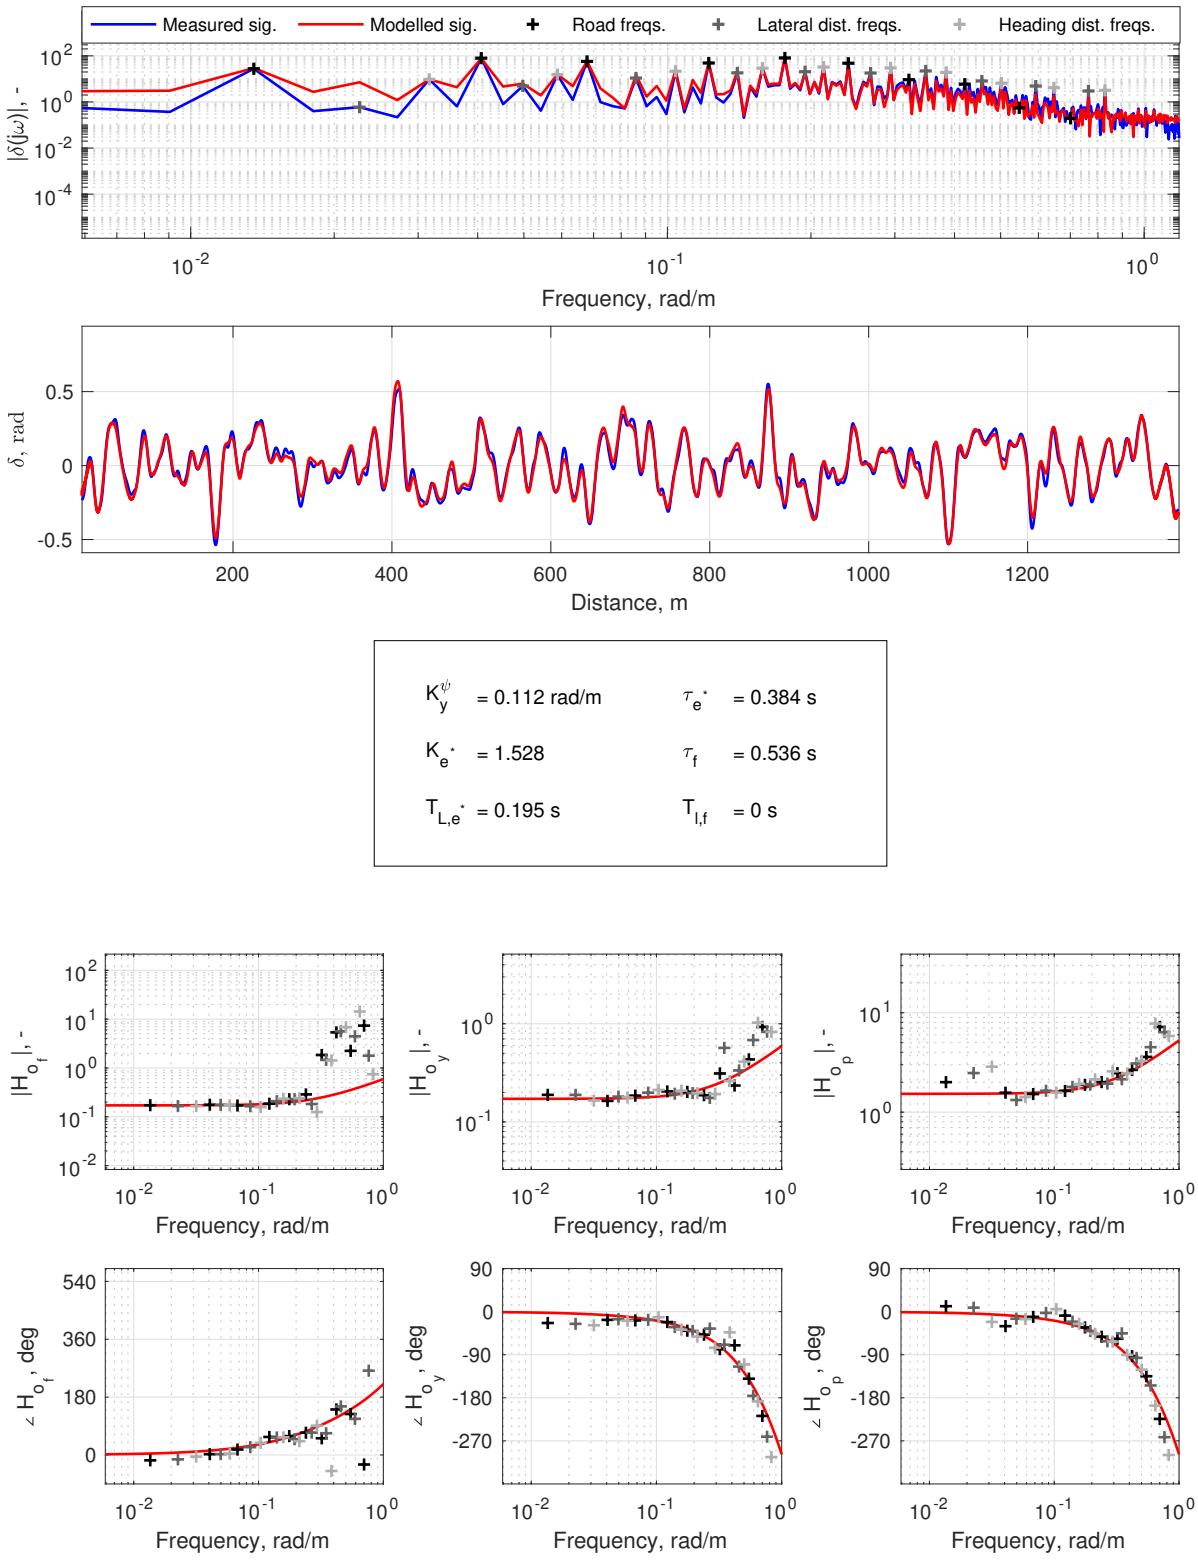

Figure E.25: FRF estimates and parameter estimation, subject 1, Condition B89

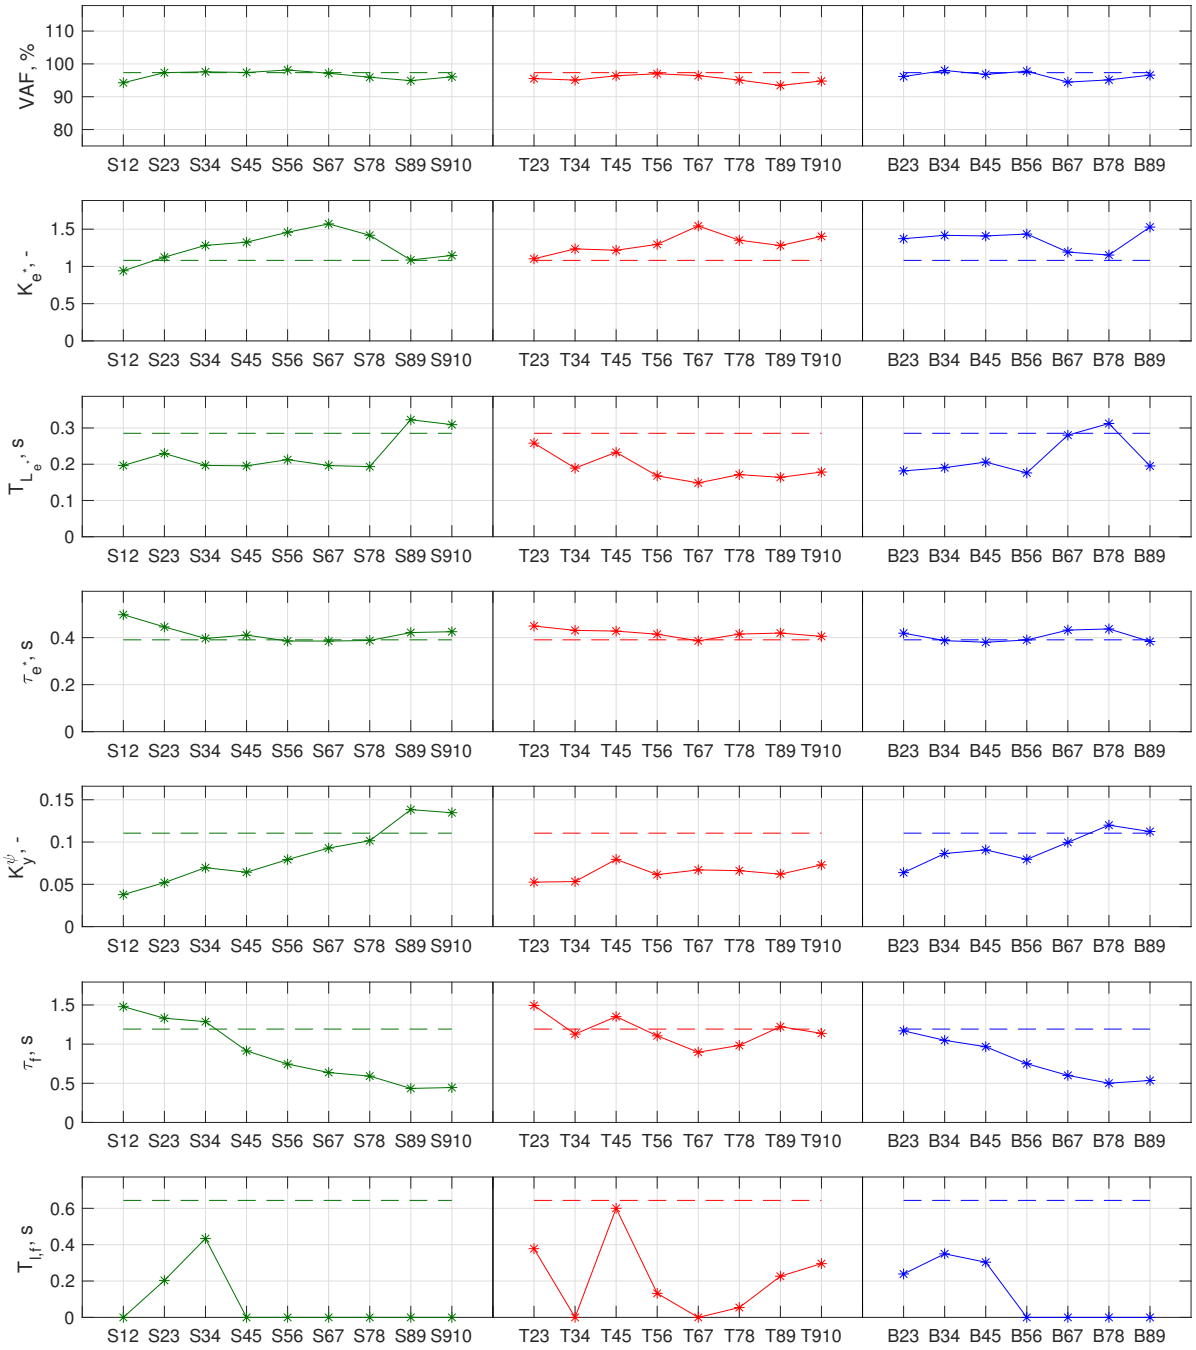

Figure E.26: Parameter estimations over all conditions, subject 1

F

Representative Model Fits - Subject 3

This chapter contains the Frequency Response Functions and fitted parametric model to the data of Subject 3.

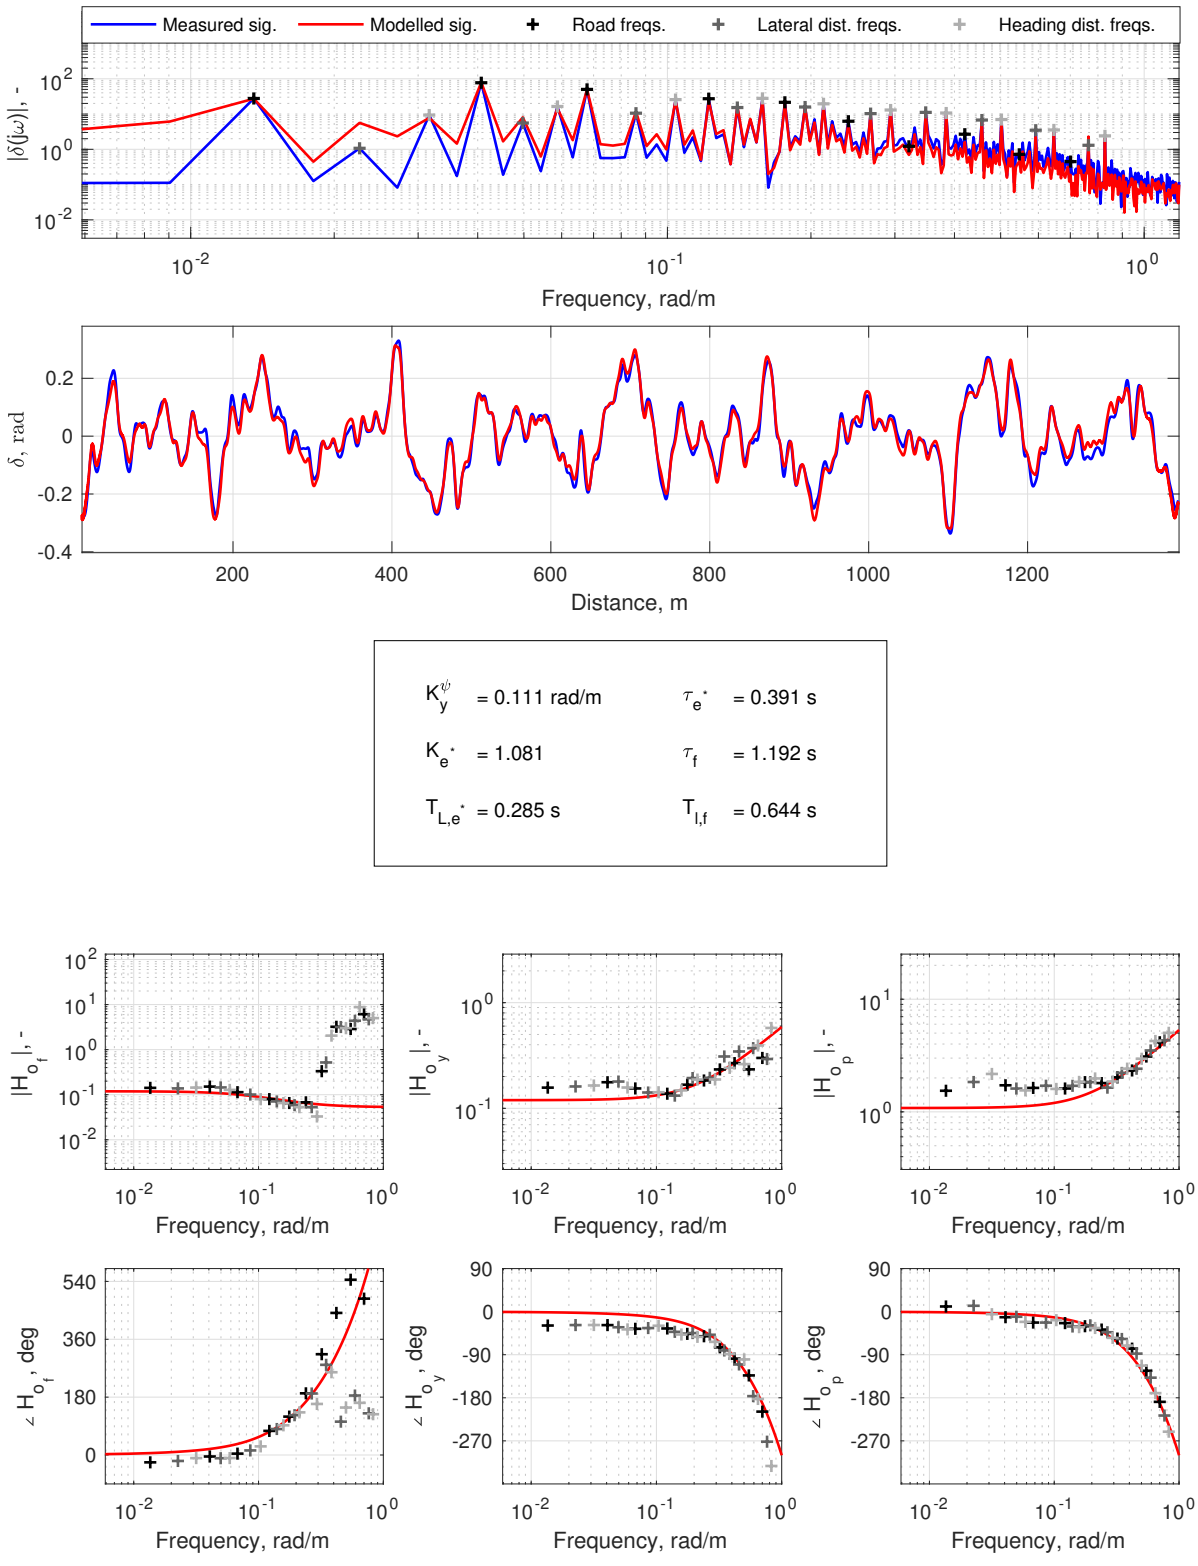

Figure F.1: FRF estimates and parameter estimation, subject 3, Condition NO

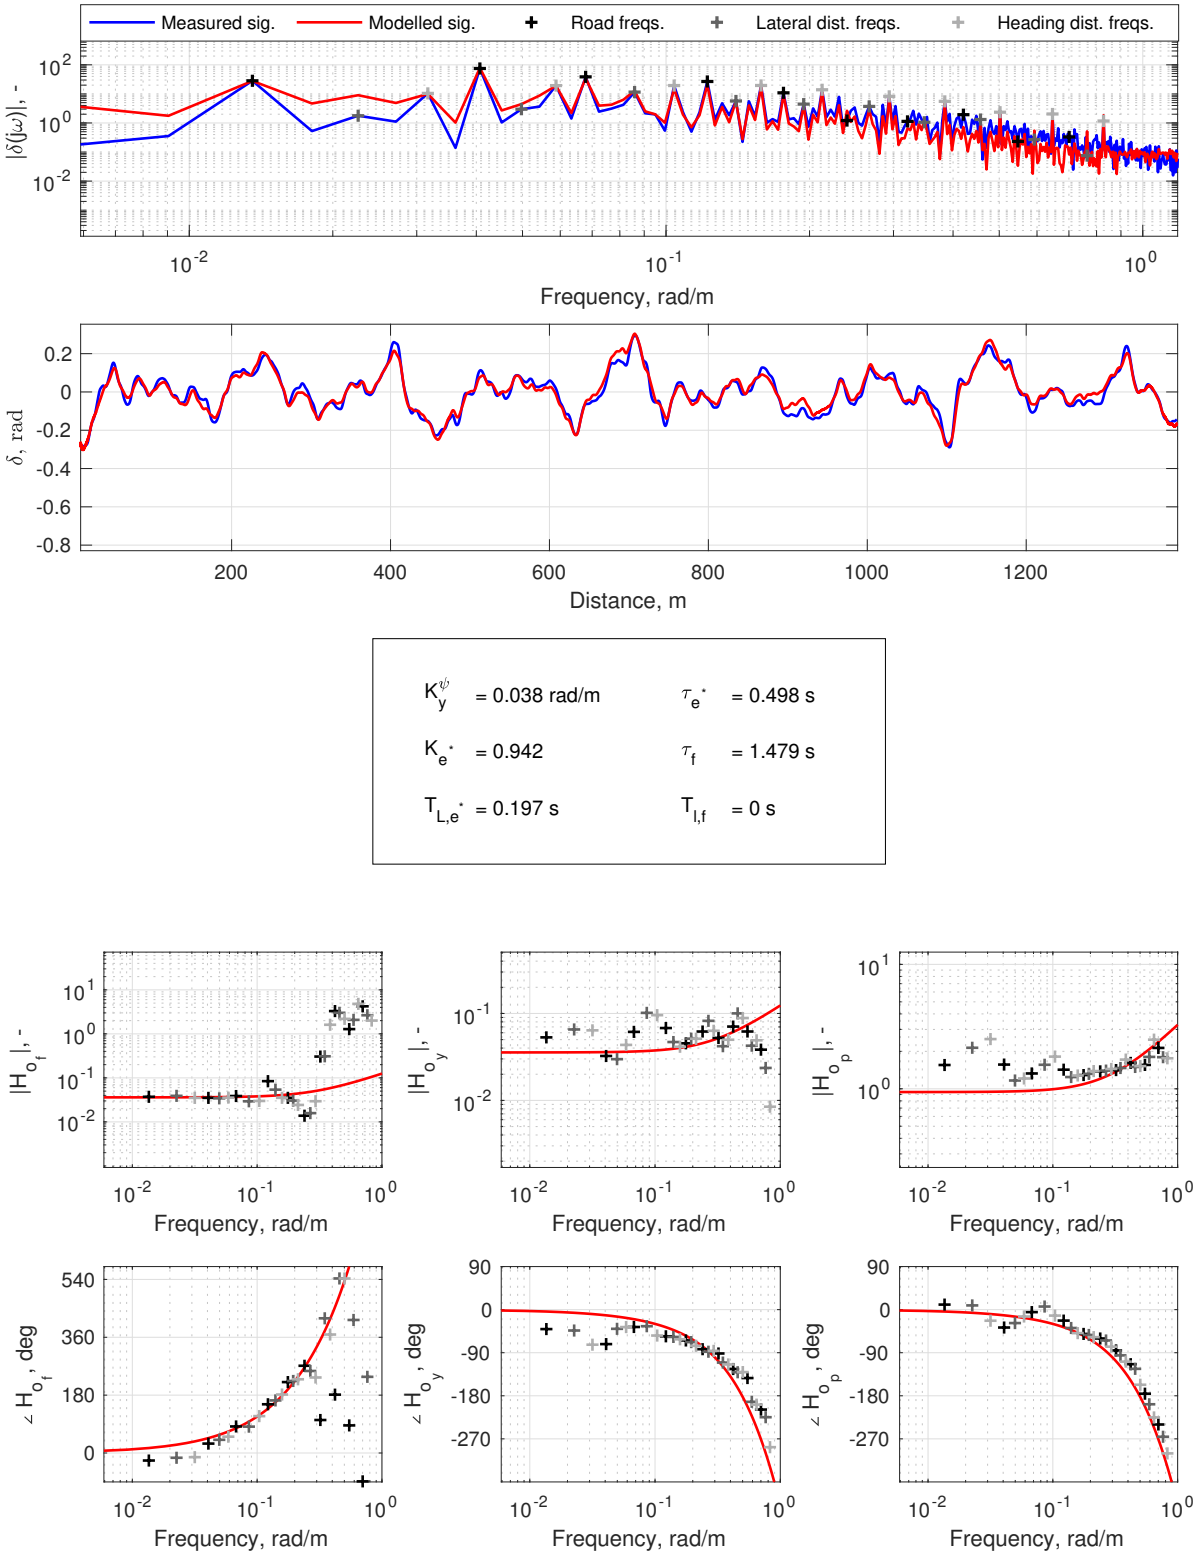

Figure F.2: FRF estimates and parameter estimation, subject 3, Condition S12

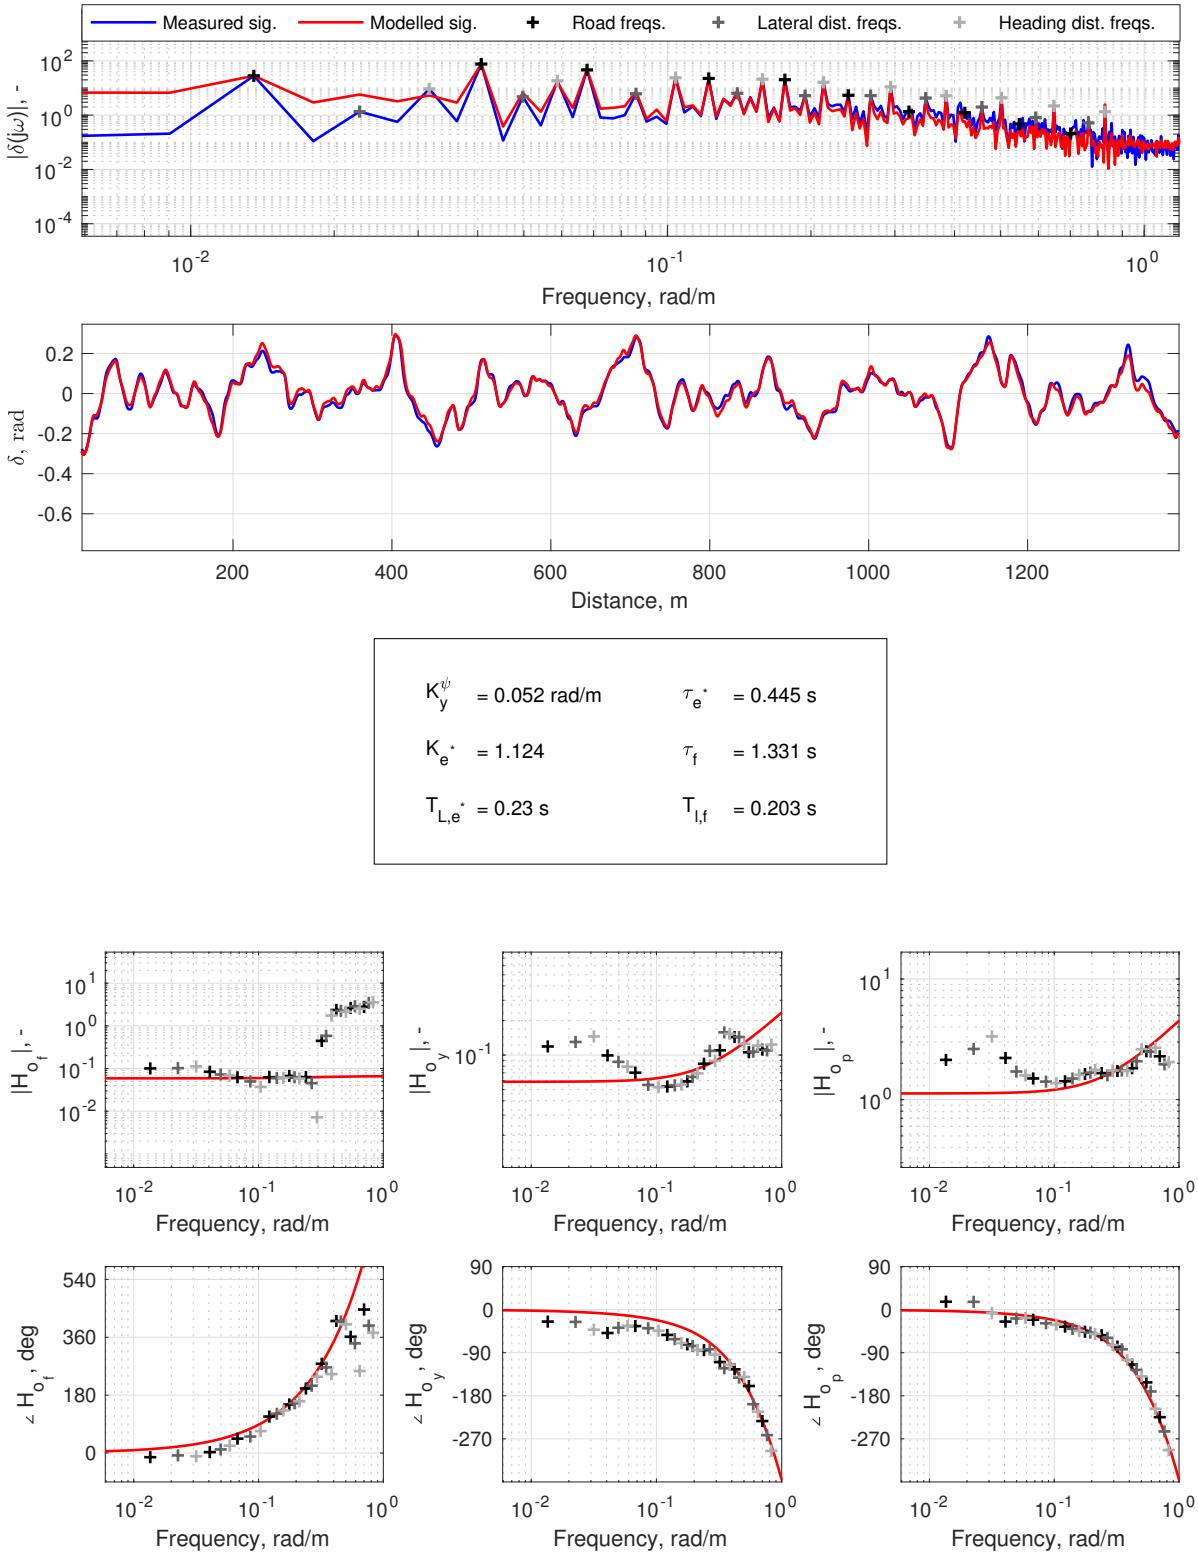

Figure F.3: FRF estimates and parameter estimation, subject 3, Condition S23

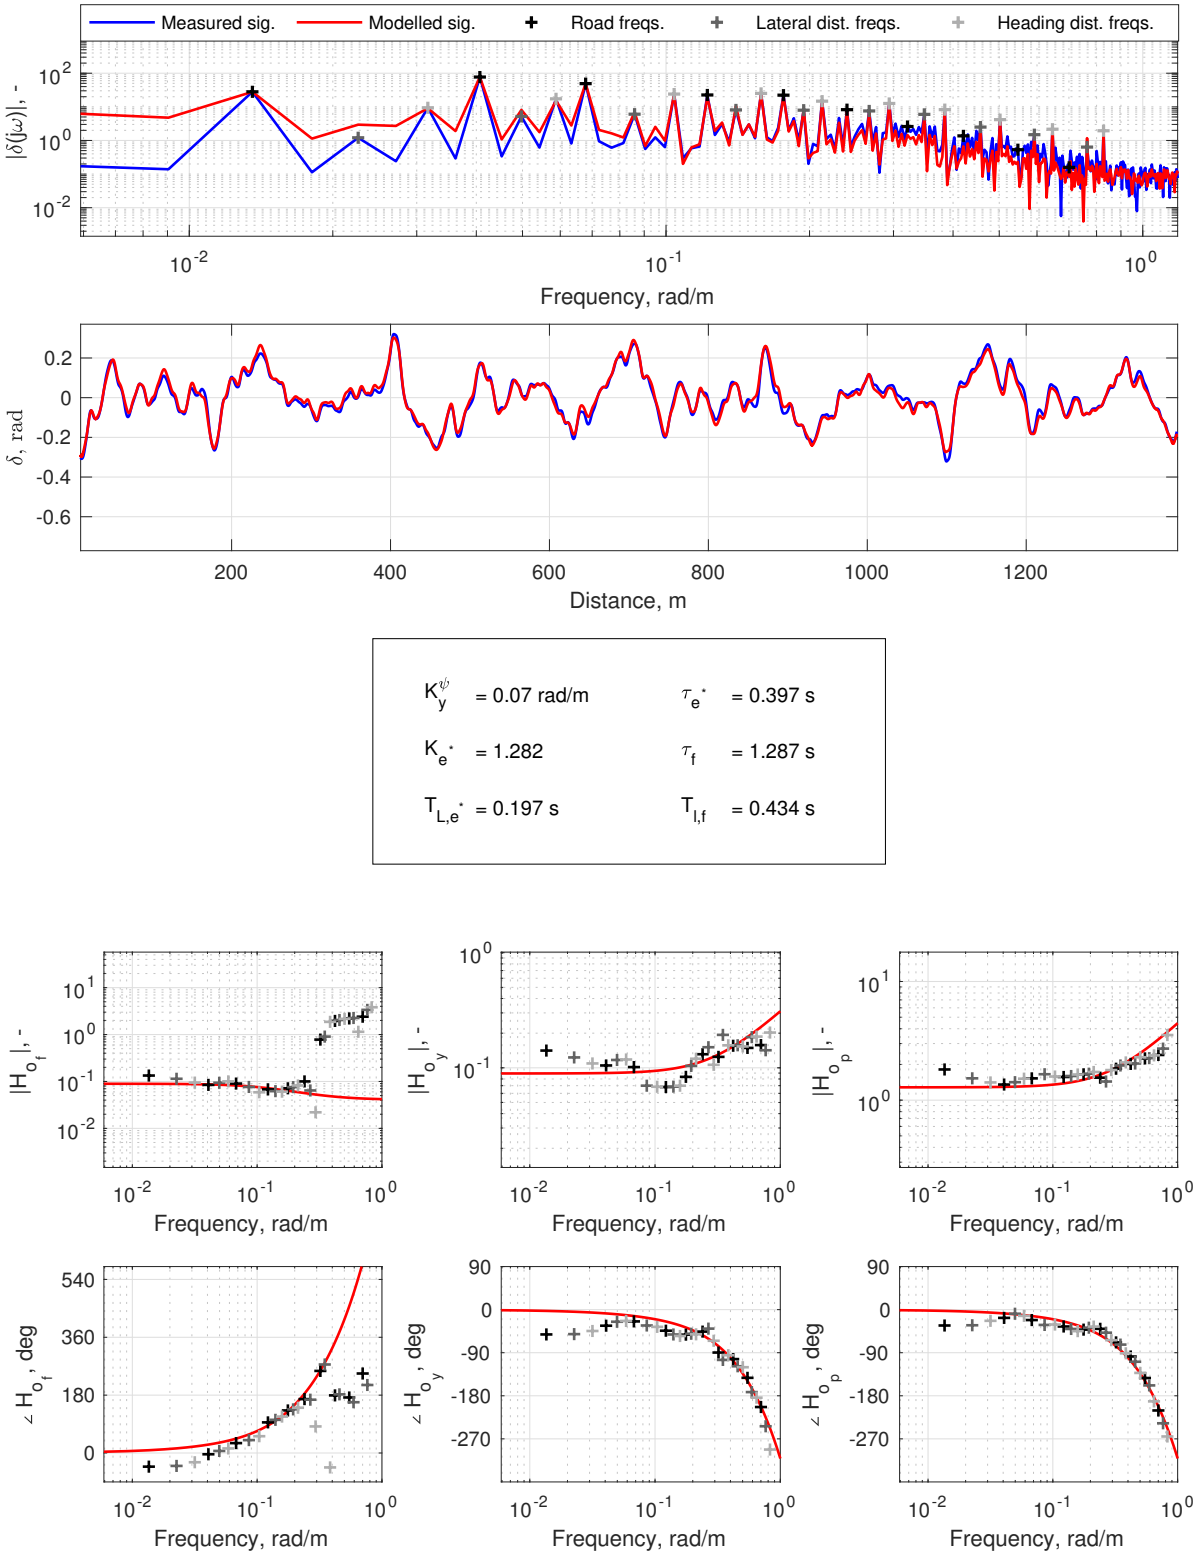

Figure F.4: FRF estimates and parameter estimation, subject 3, Condition S34

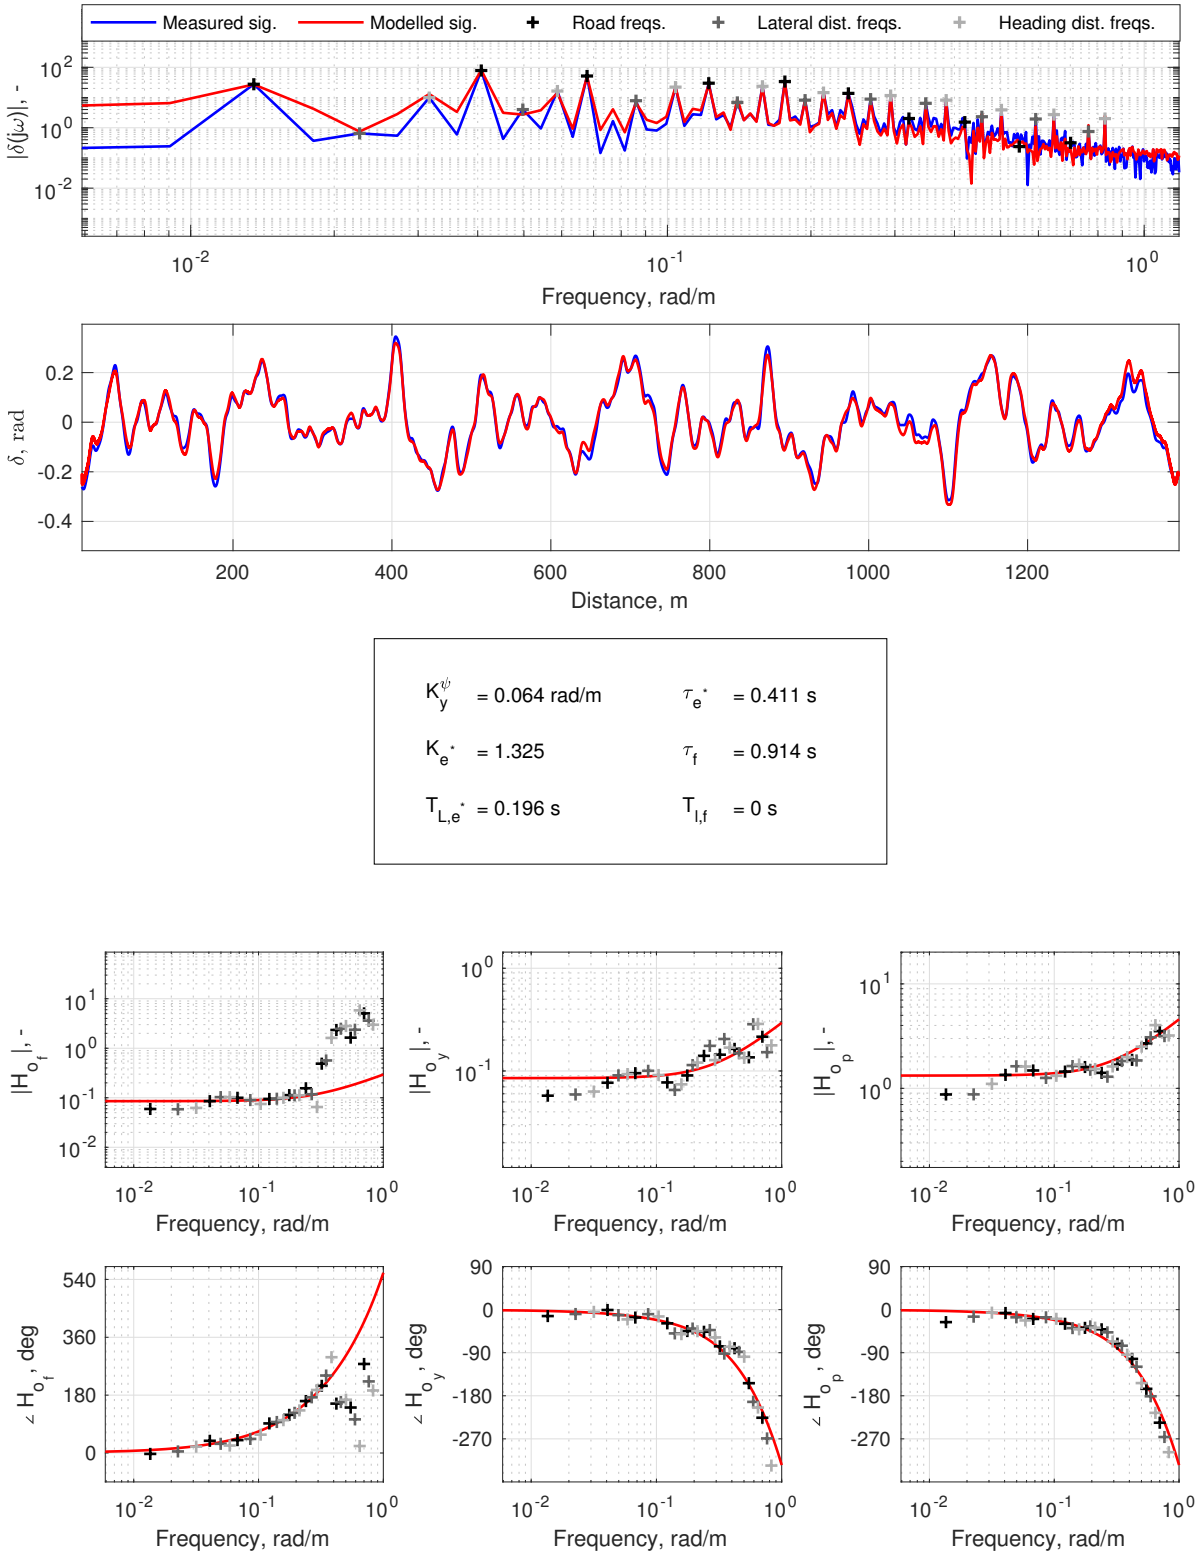

Figure F.5: FRF estimates and parameter estimation, subject 3, Condition S45

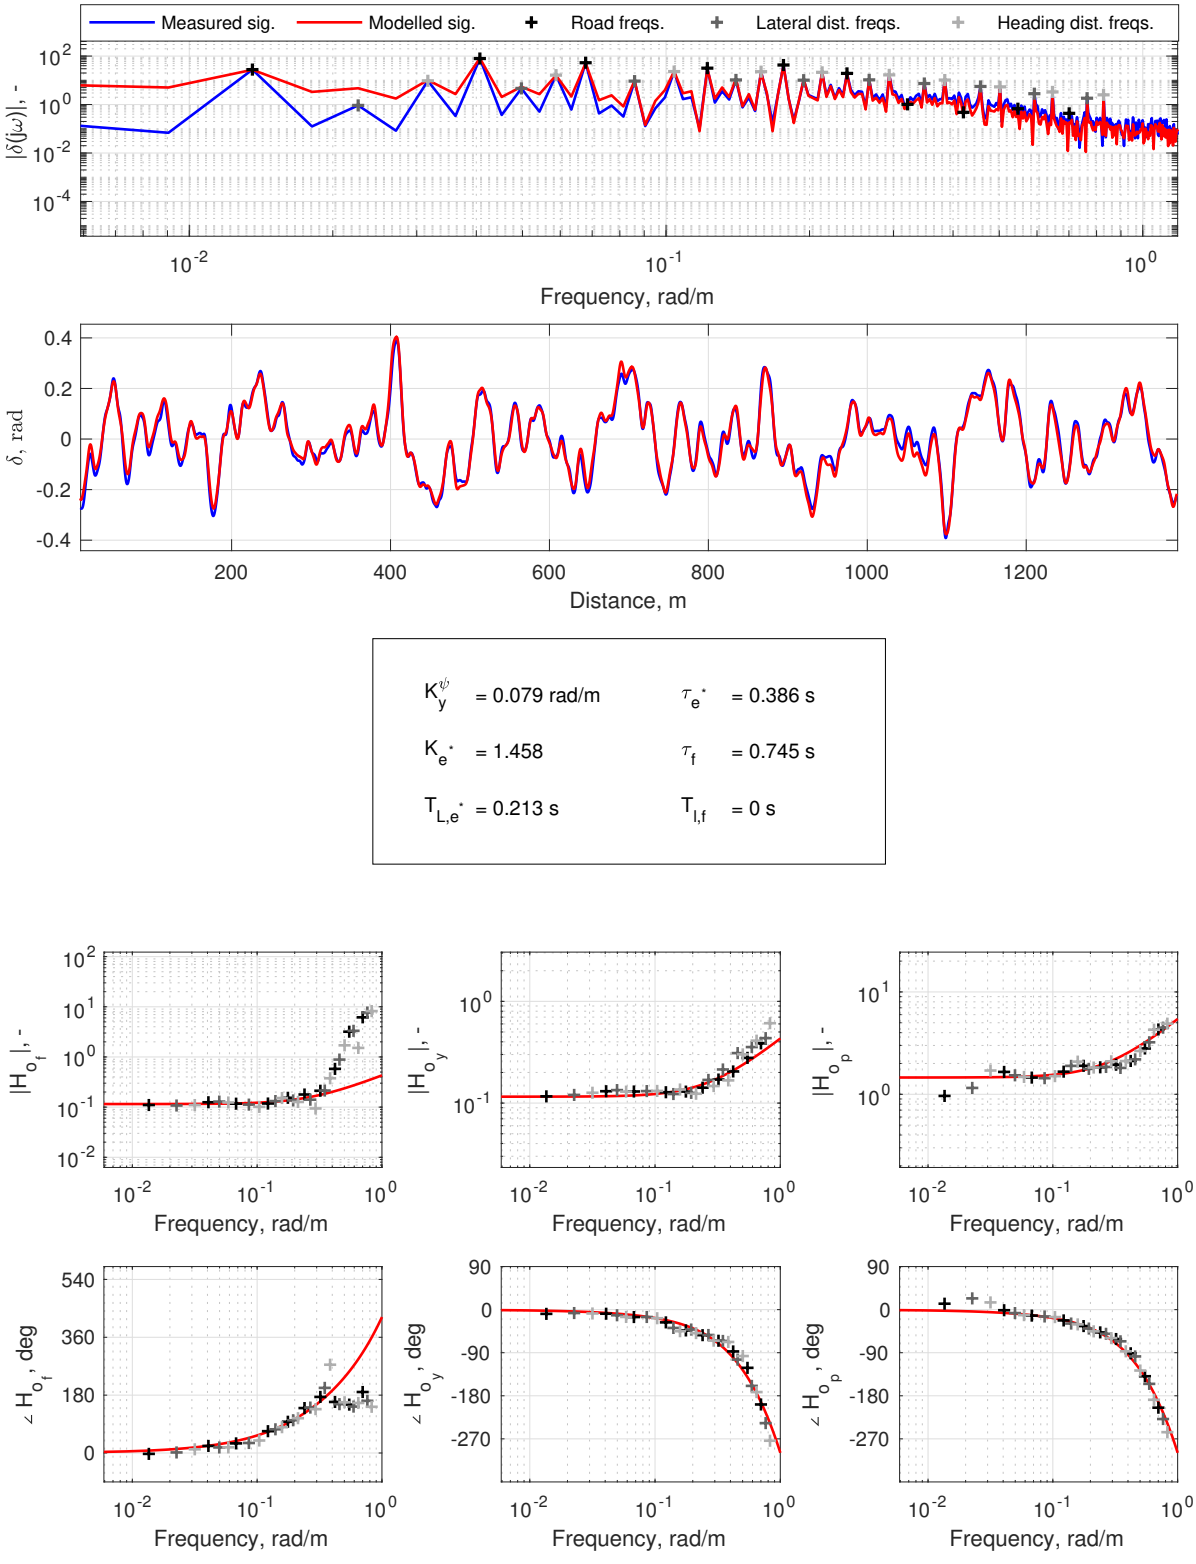

Figure F.6: FRF estimates and parameter estimation, subject 3, Condition S56

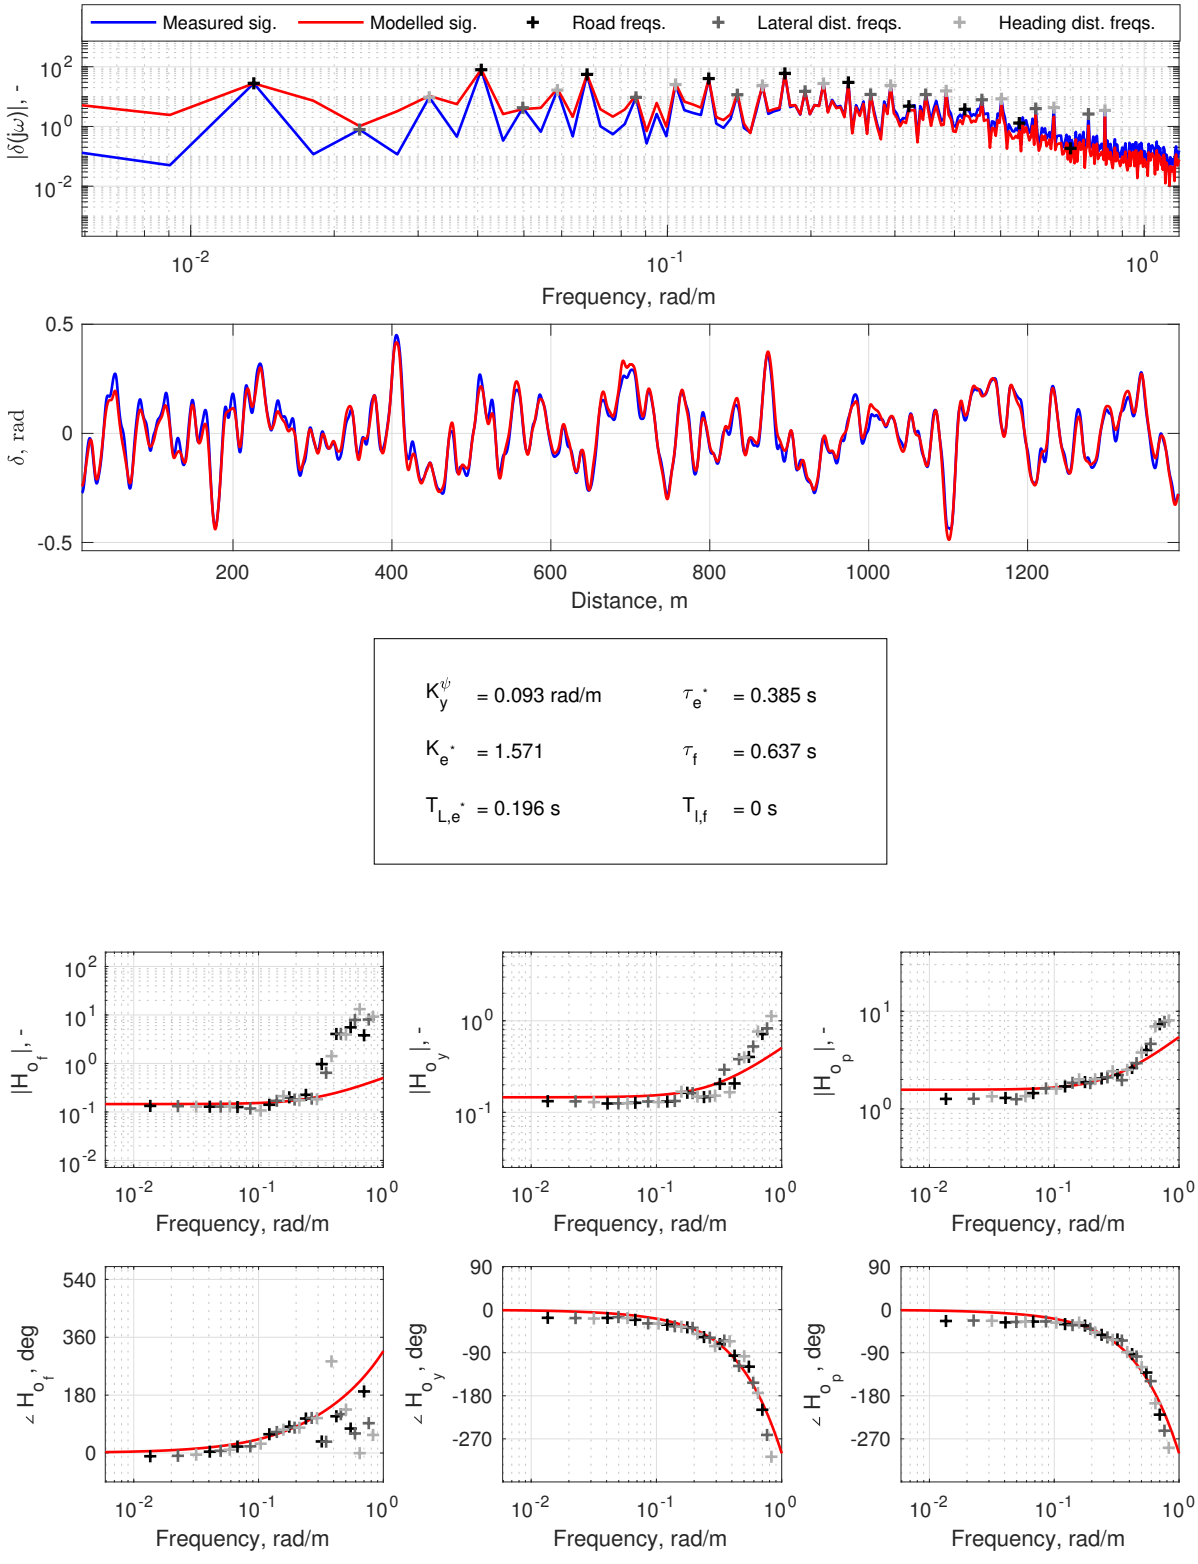

Figure F.7: FRF estimates and parameter estimation, subject 3, Condition S67

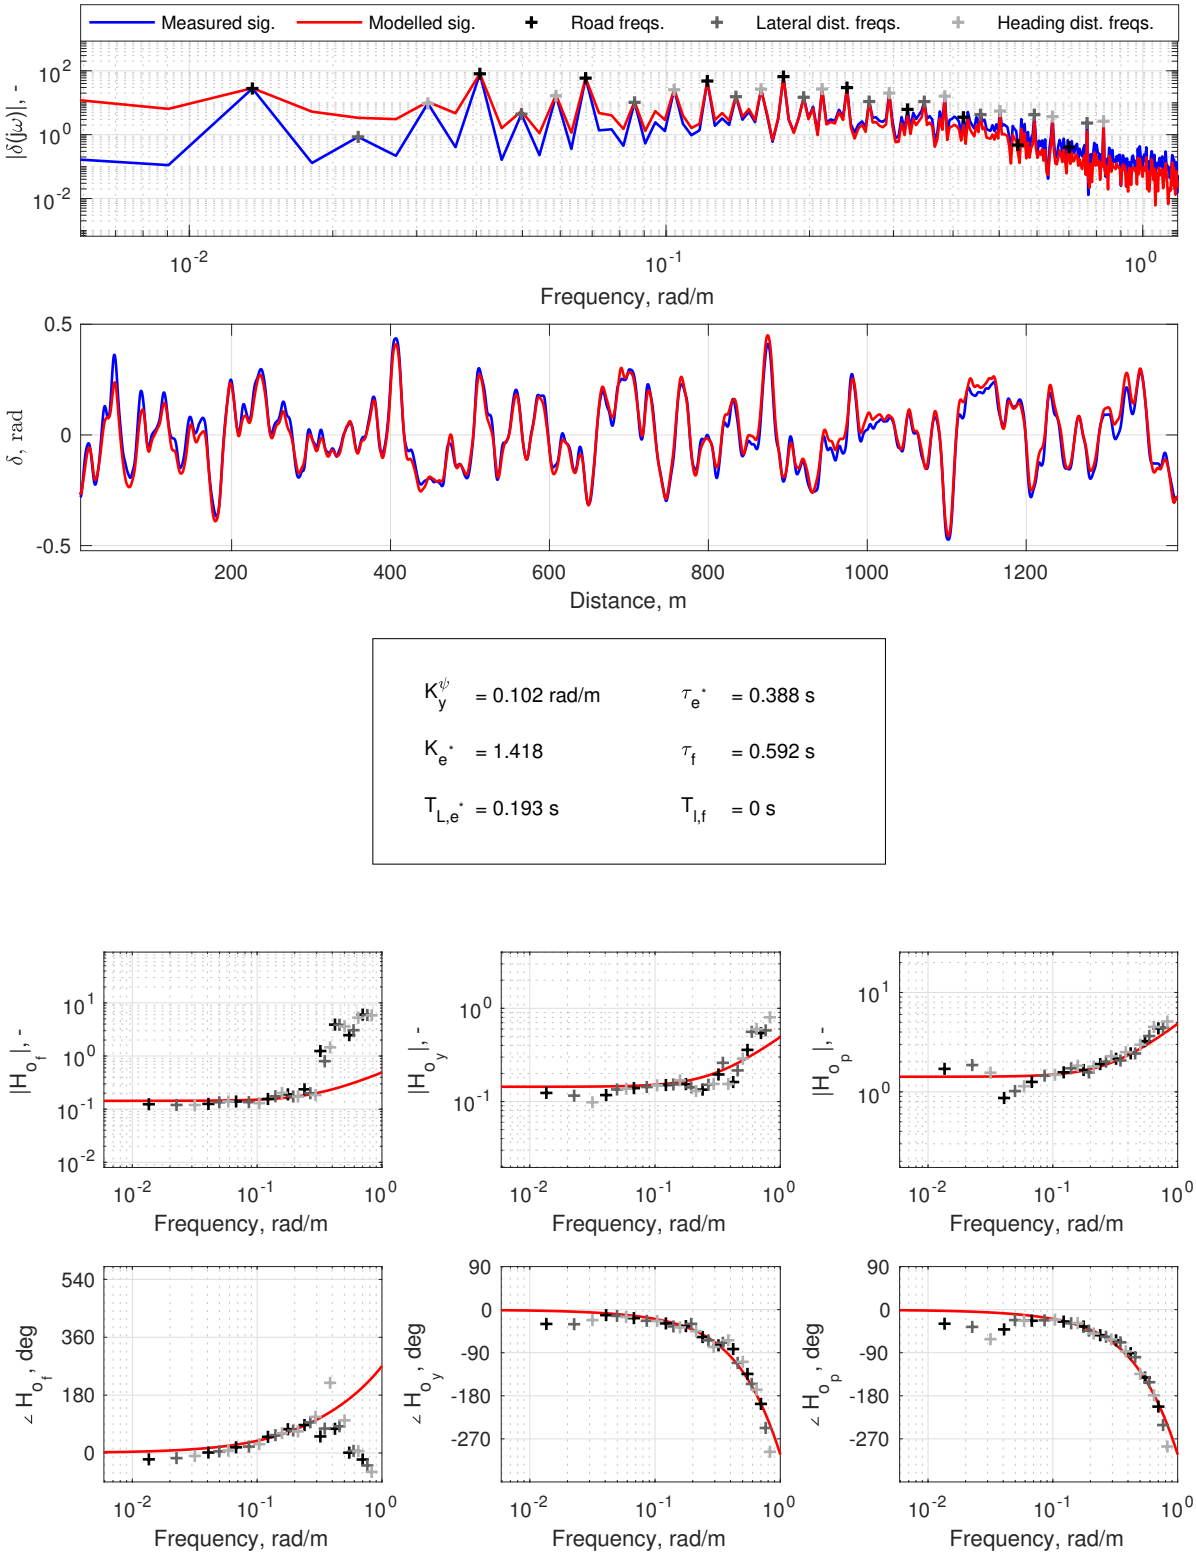

Figure F.8: FRF estimates and parameter estimation, subject 3, Condition S78

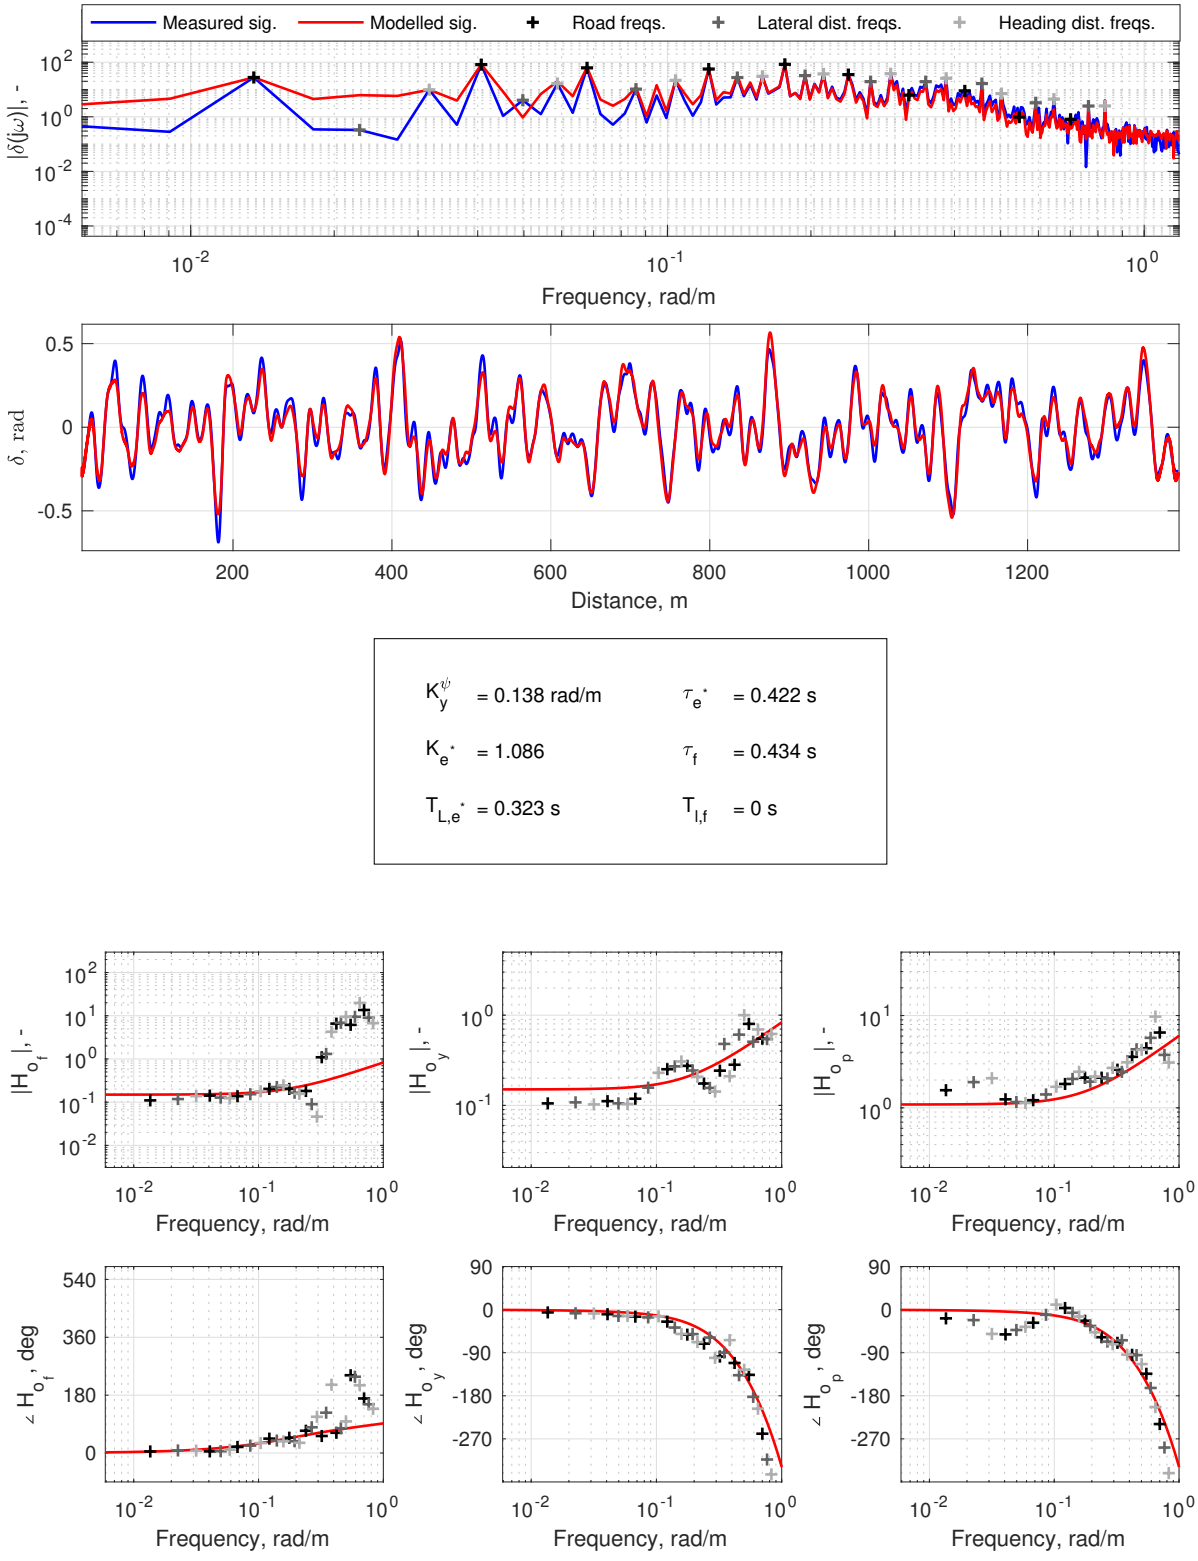

Figure F.9: FRF estimates and parameter estimation, subject 3, Condition S89

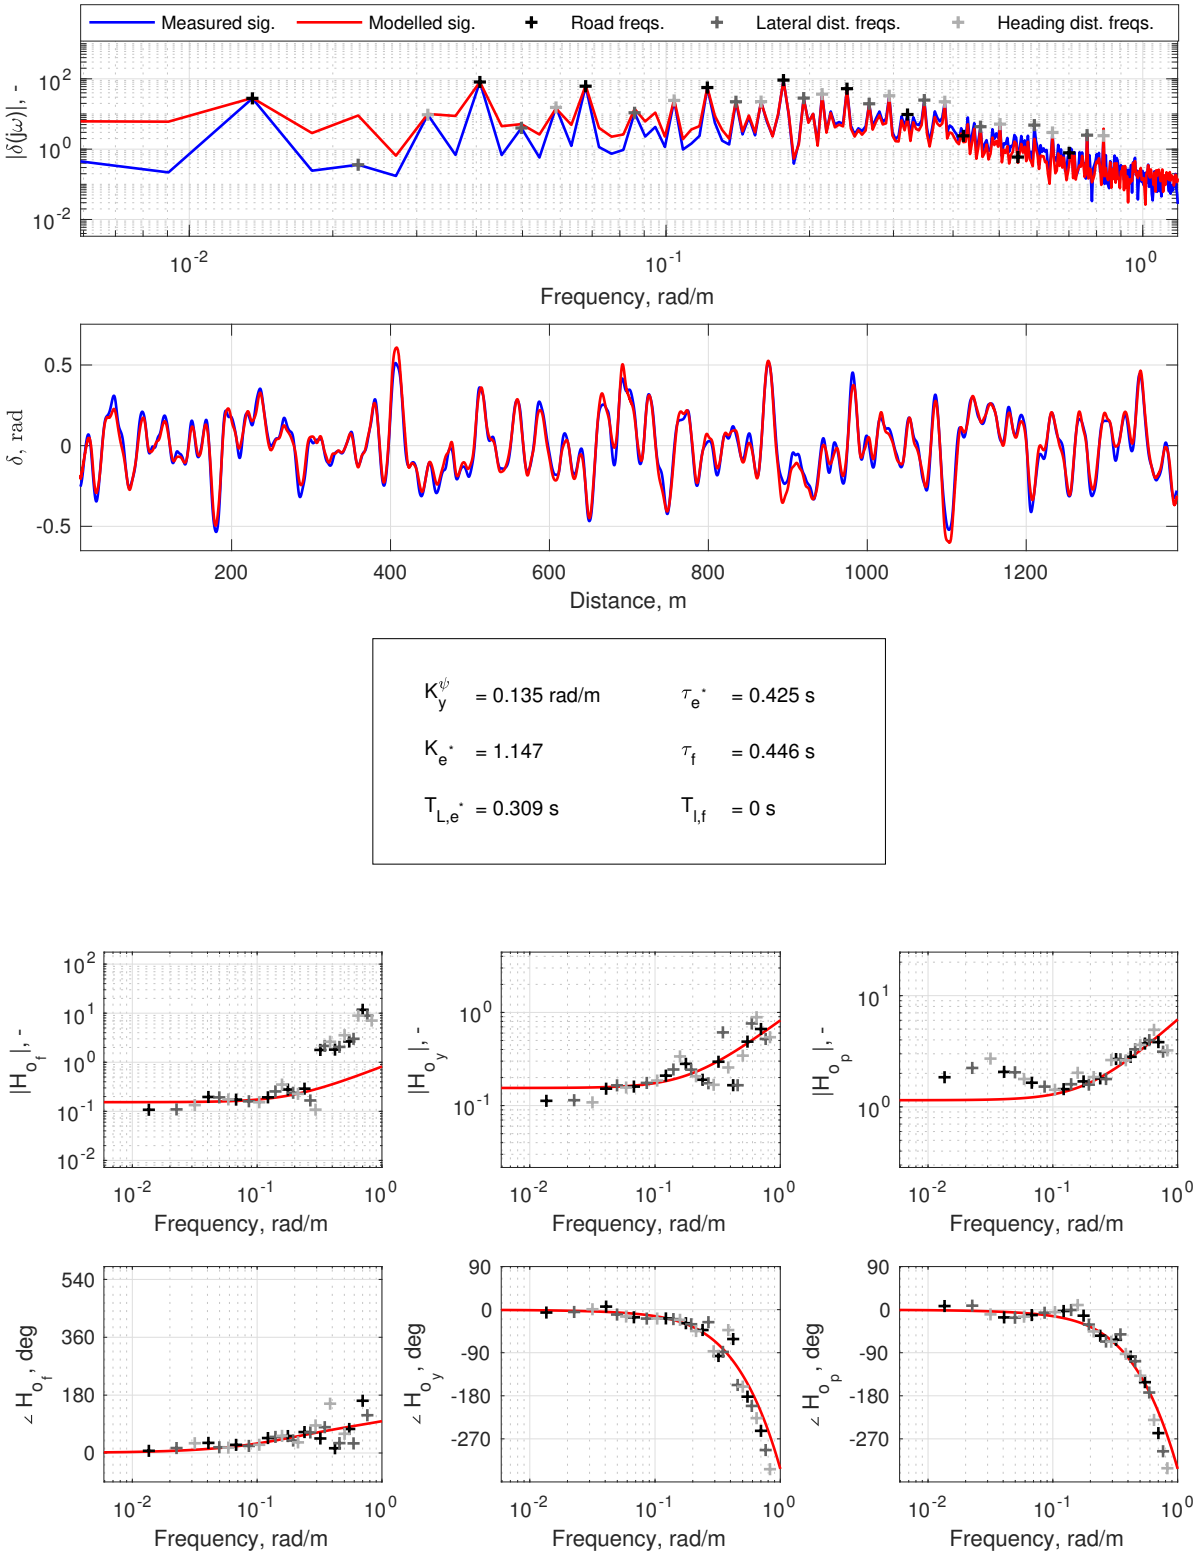

Figure F.10: FRF estimates and parameter estimation, subject 3, Condition S910

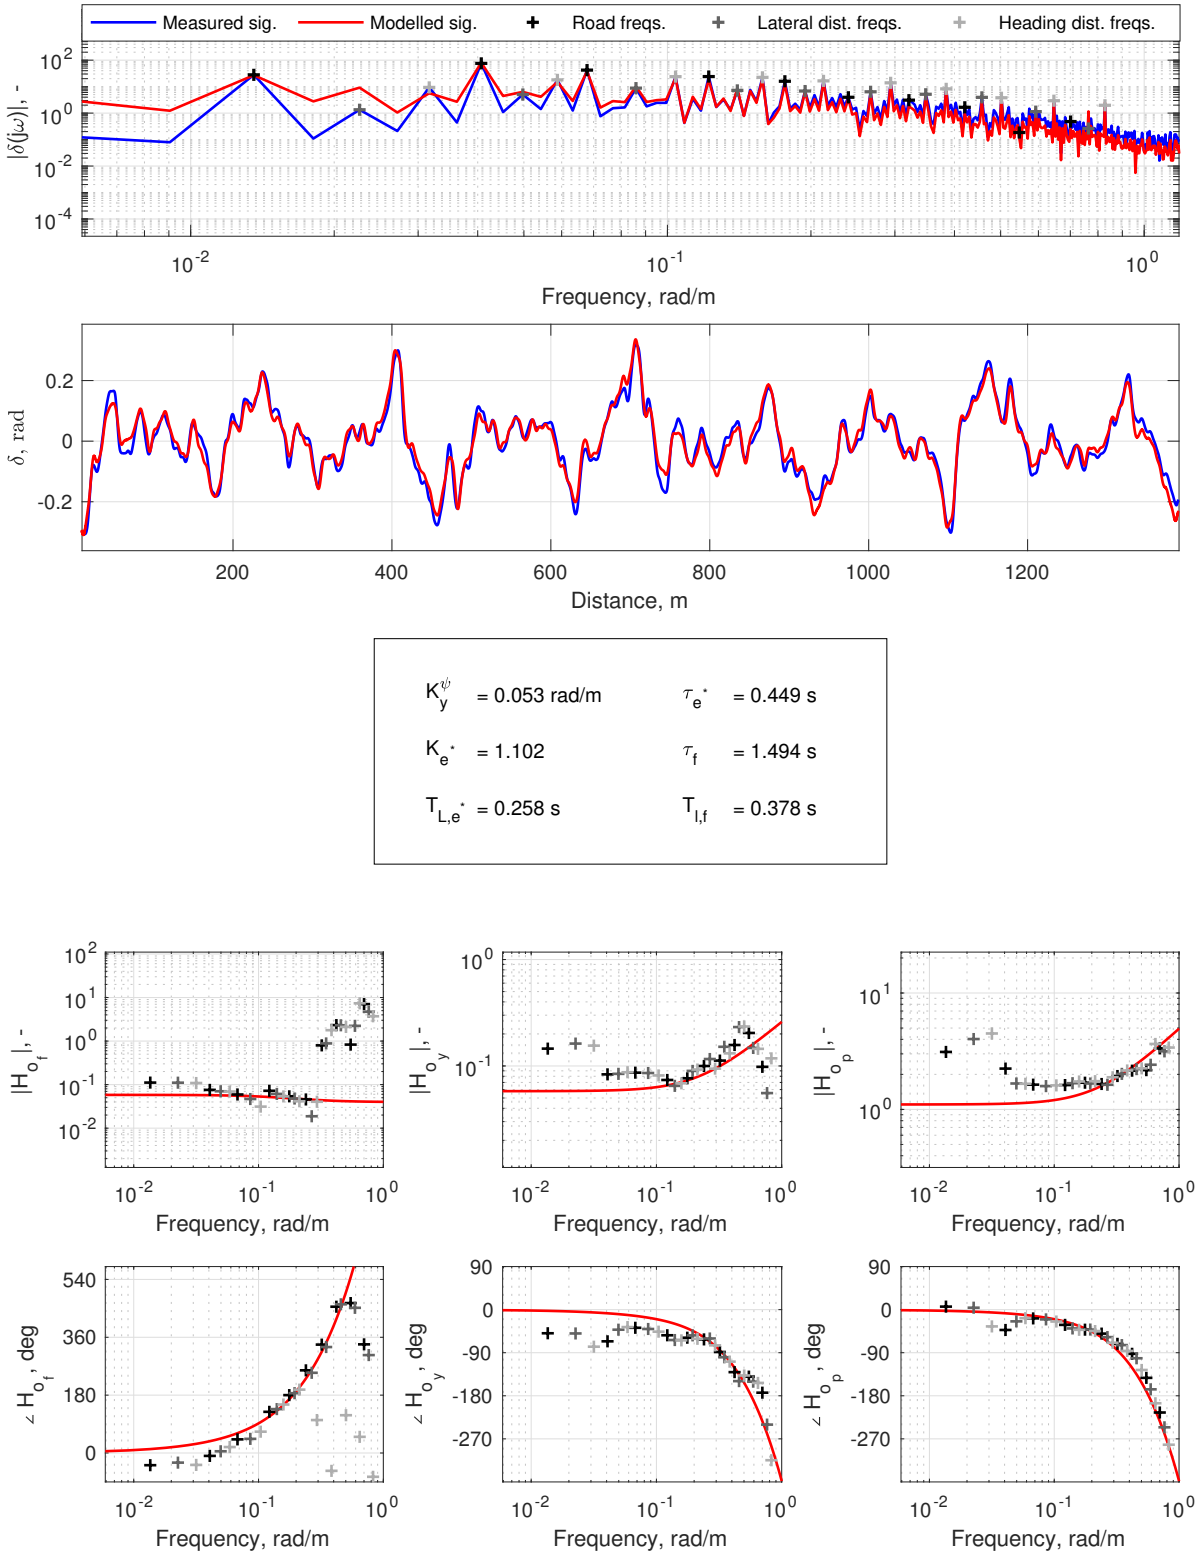

Figure F.11: FRF estimates and parameter estimation, subject 3, Condition T23

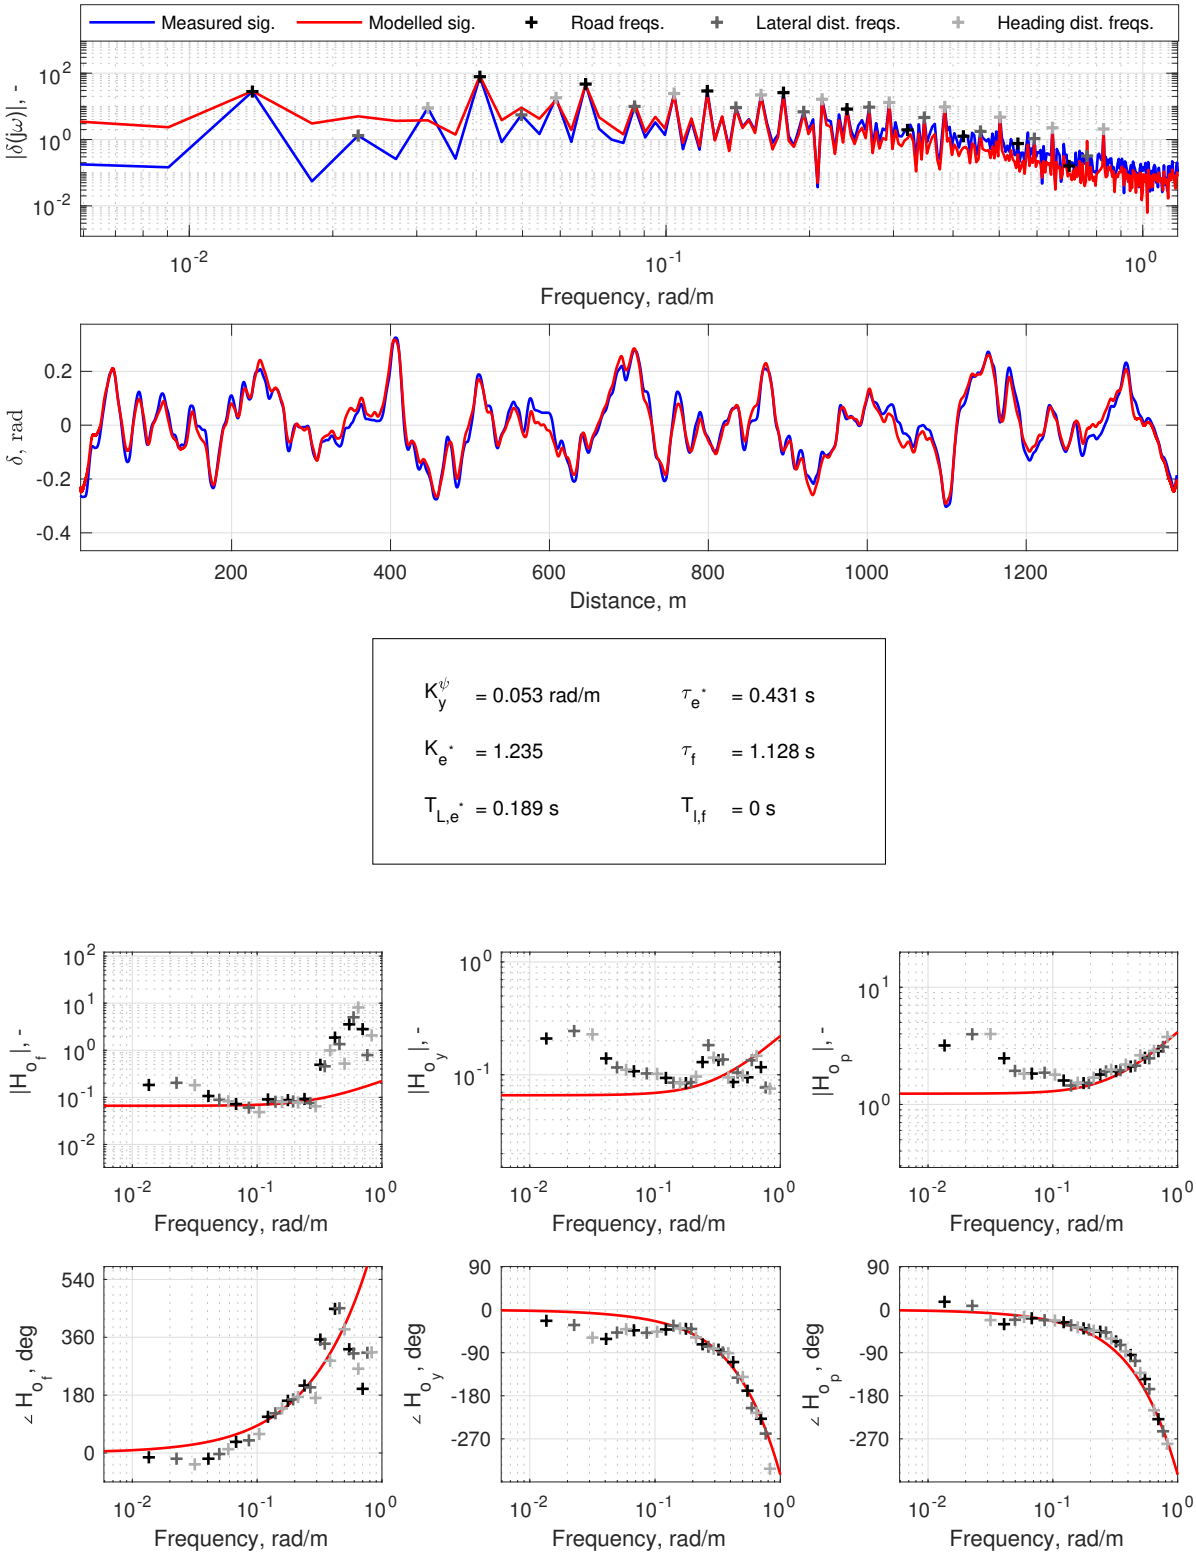

Figure F.12: FRF estimates and parameter estimation, subject 3, Condition T34

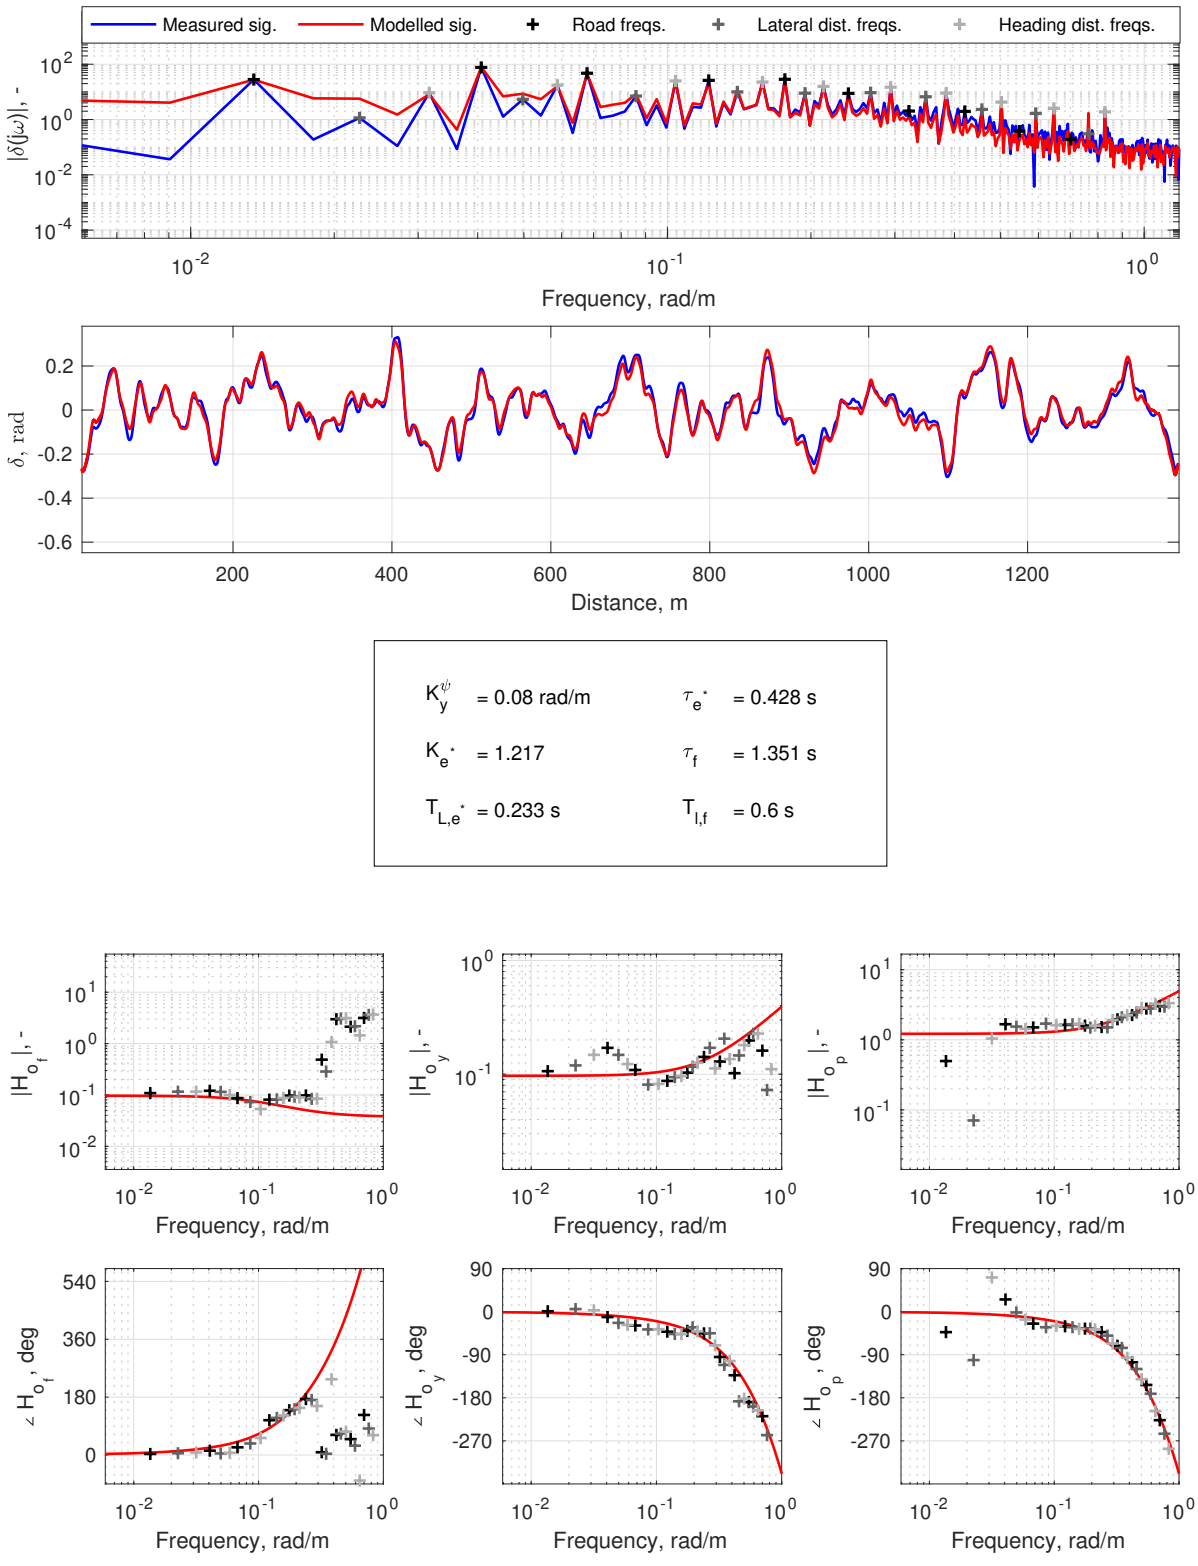

Figure F.13: FRF estimates and parameter estimation, subject 3, Condition T45

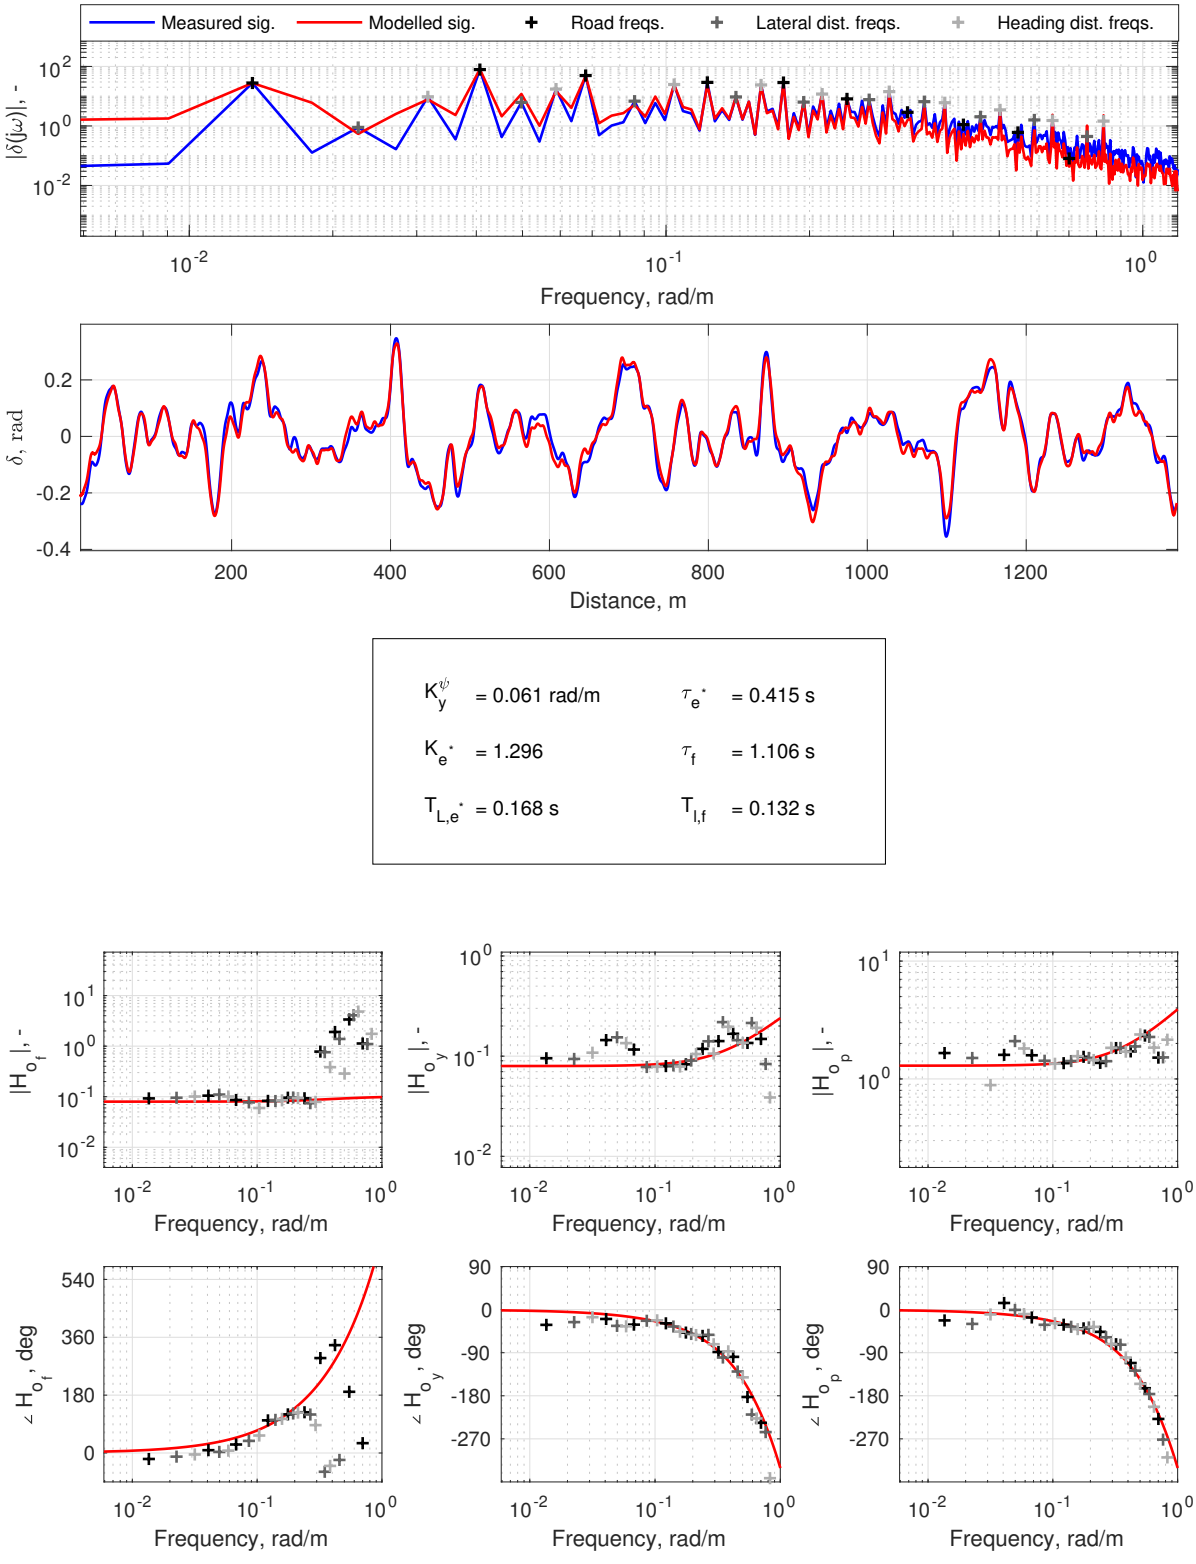

Figure F.14: FRF estimates and parameter estimation, subject 3, Condition T56

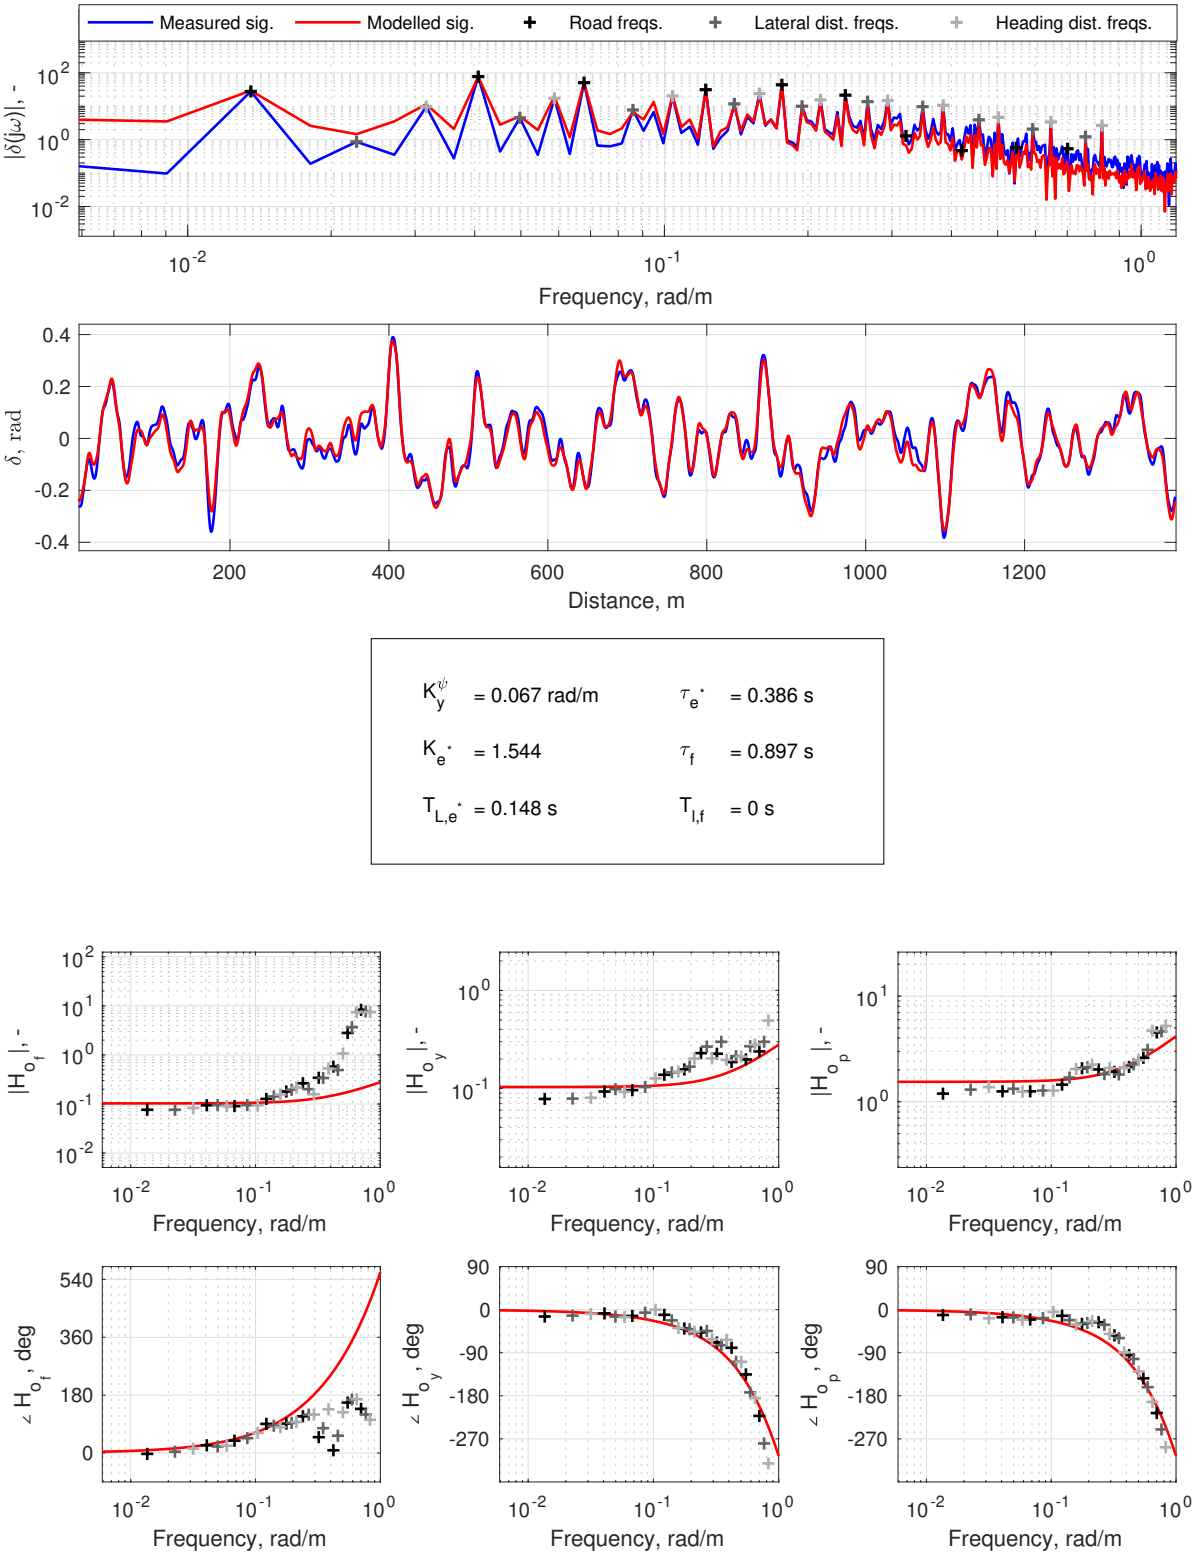

Figure F.15: FRF estimates and parameter estimation, subject 3, Condition T67

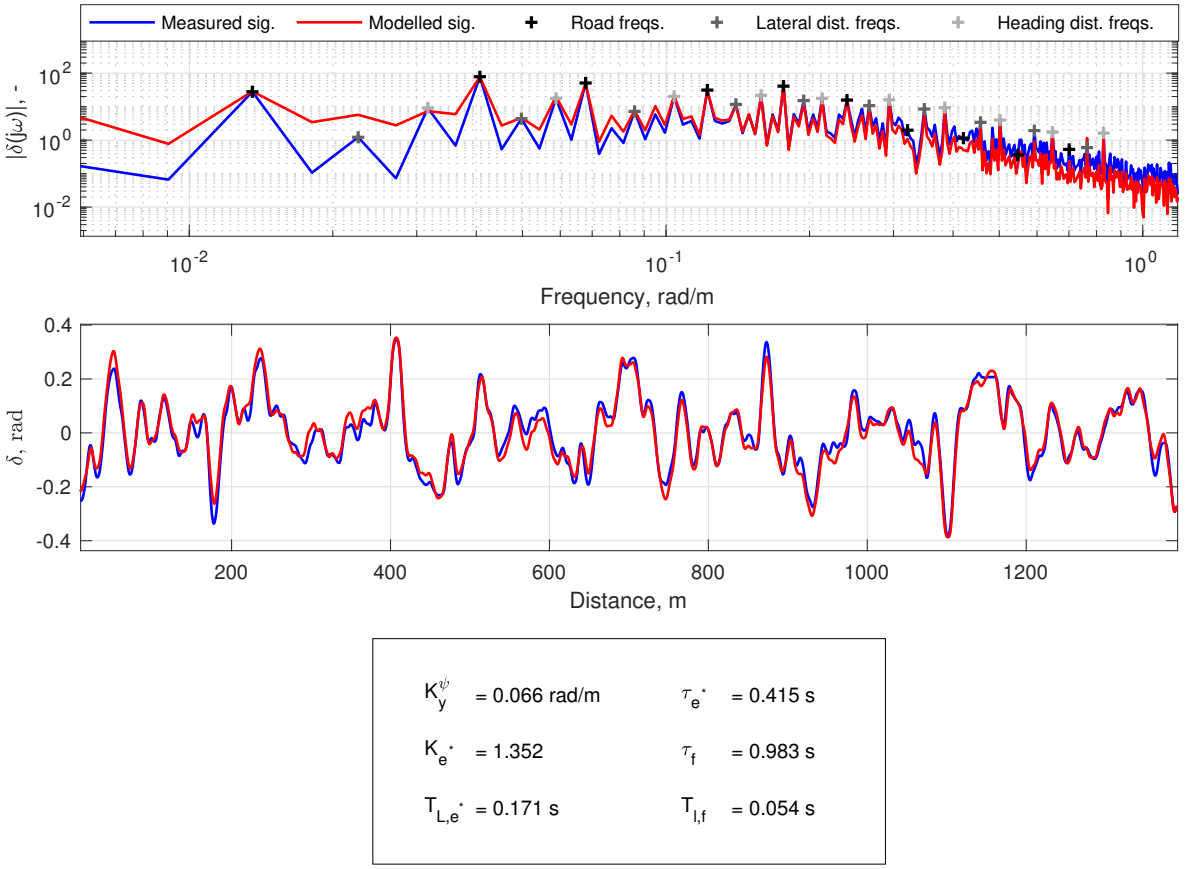

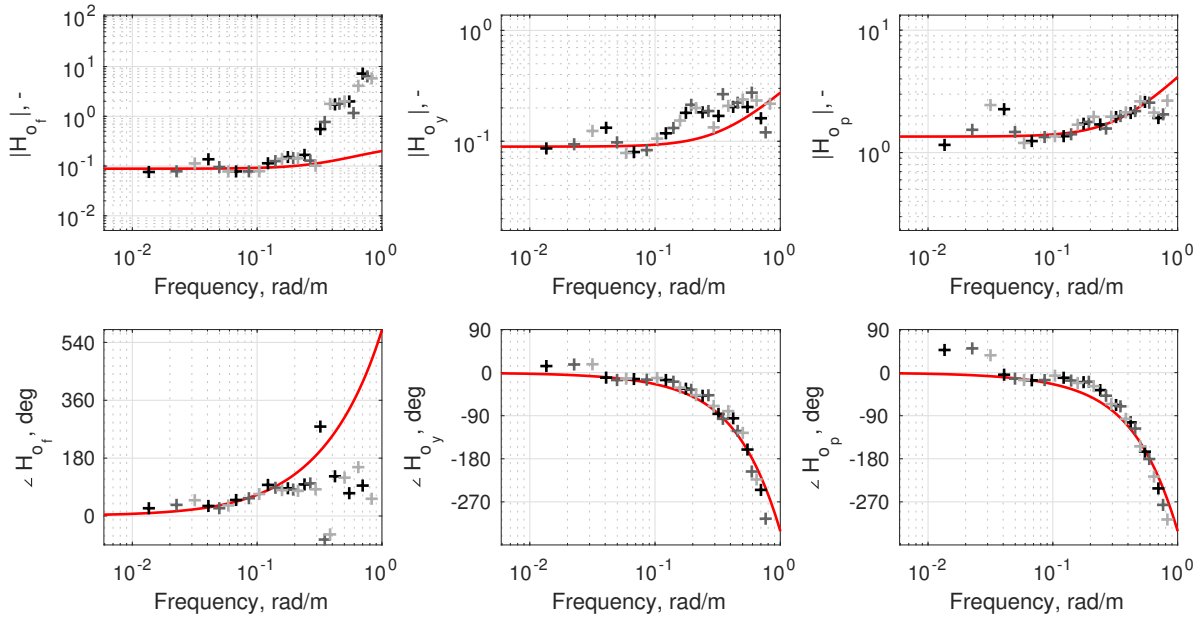

Figure F.16: FRF estimates and parameter estimation, subject 3, Condition T78

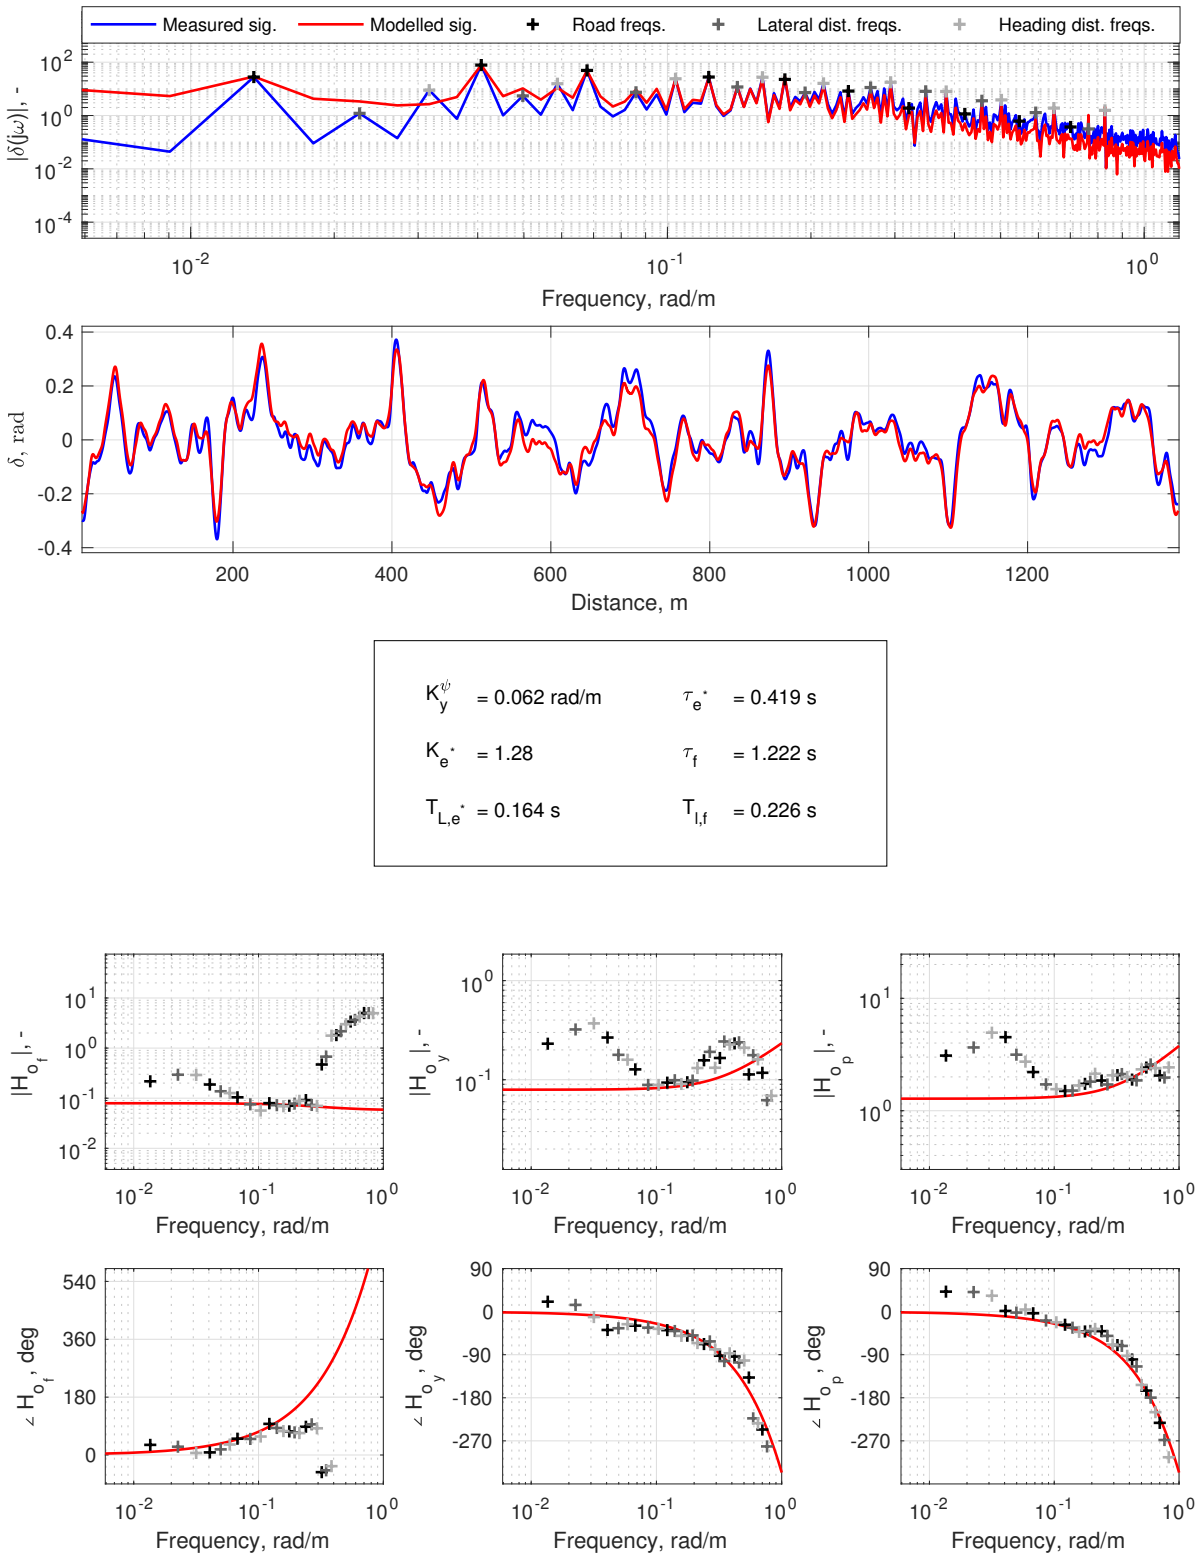

Figure F.17: FRF estimates and parameter estimation, subject 3, Condition T89

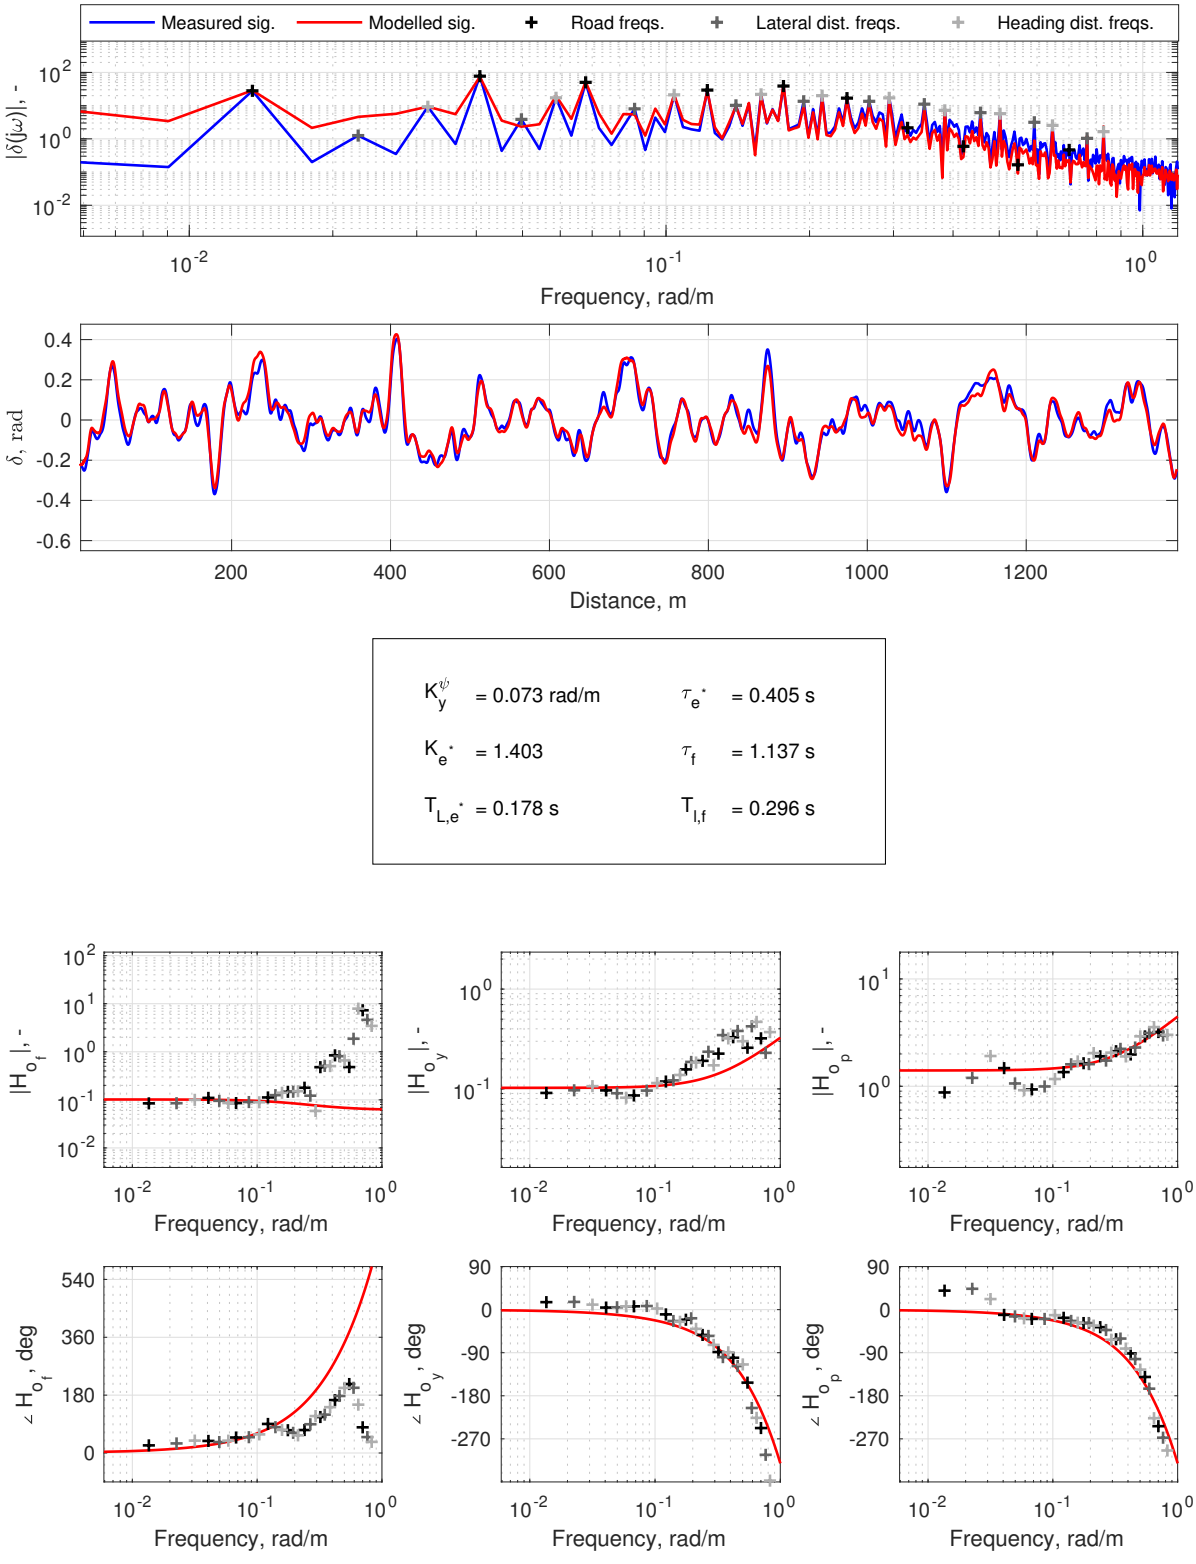

Figure F.18: FRF estimates and parameter estimation, subject 3, Condition T910

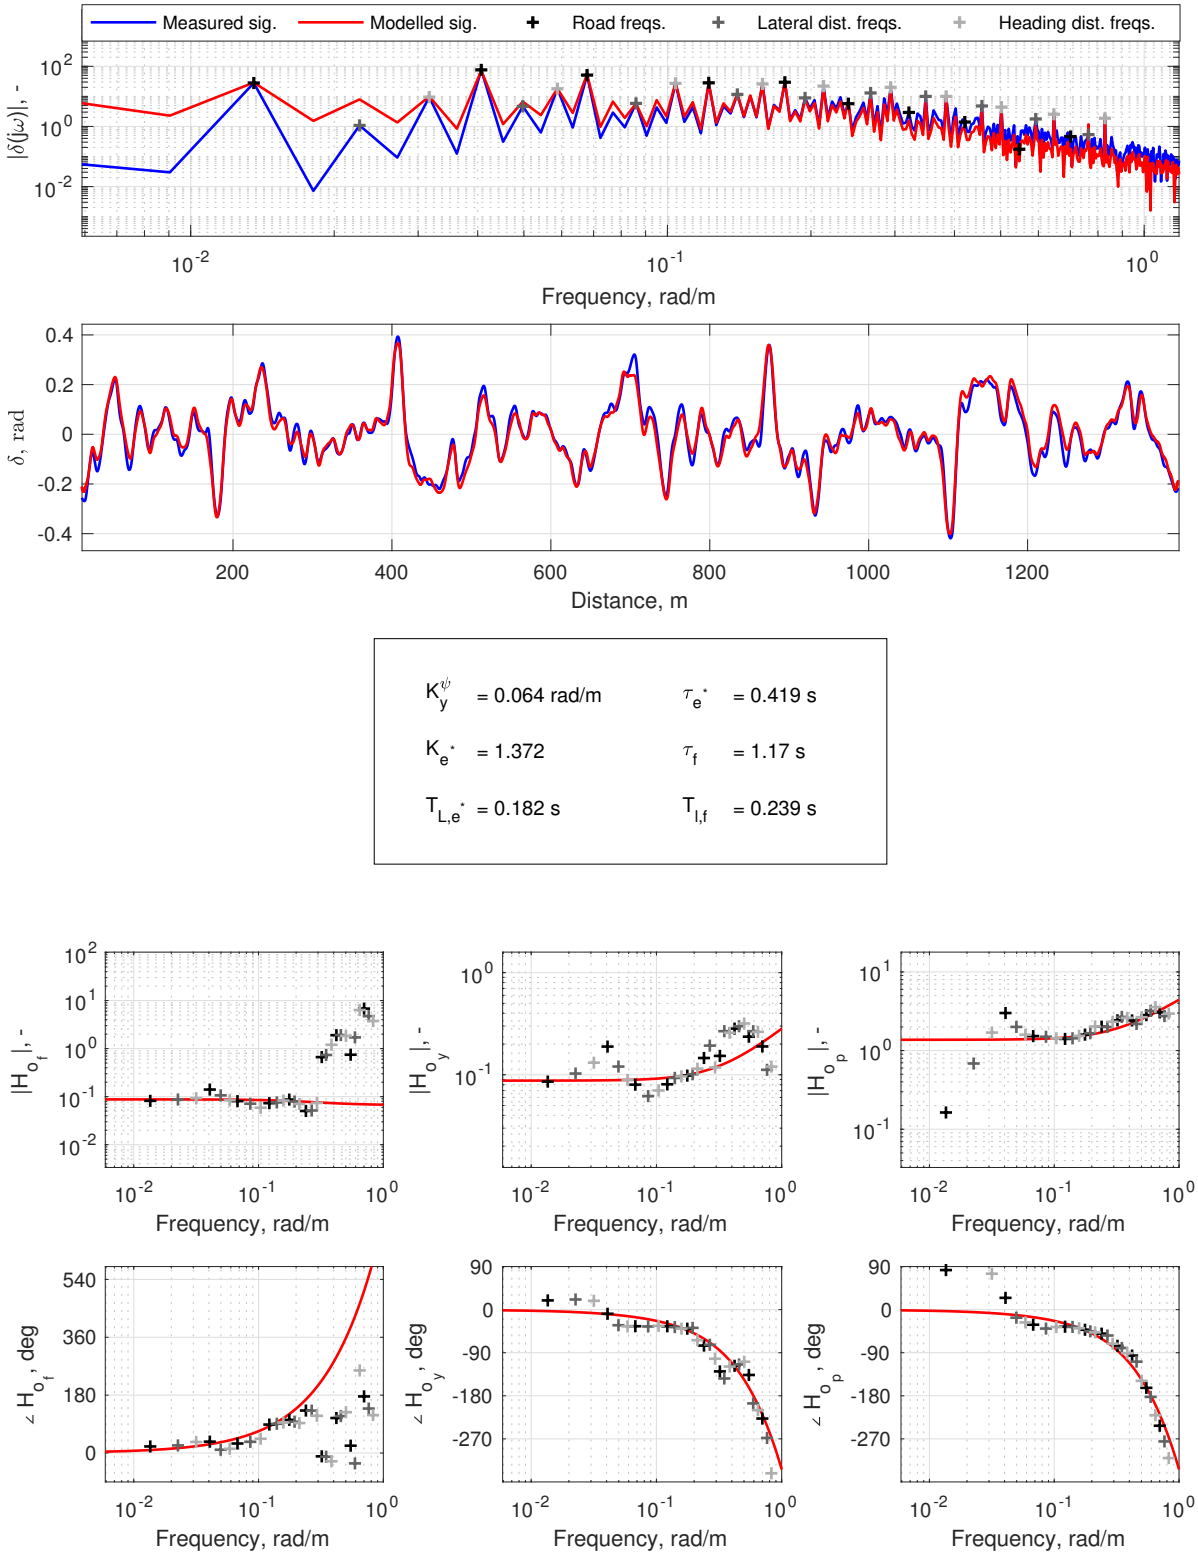

Figure F.19: FRF estimates and parameter estimation, subject 3, Condition B23

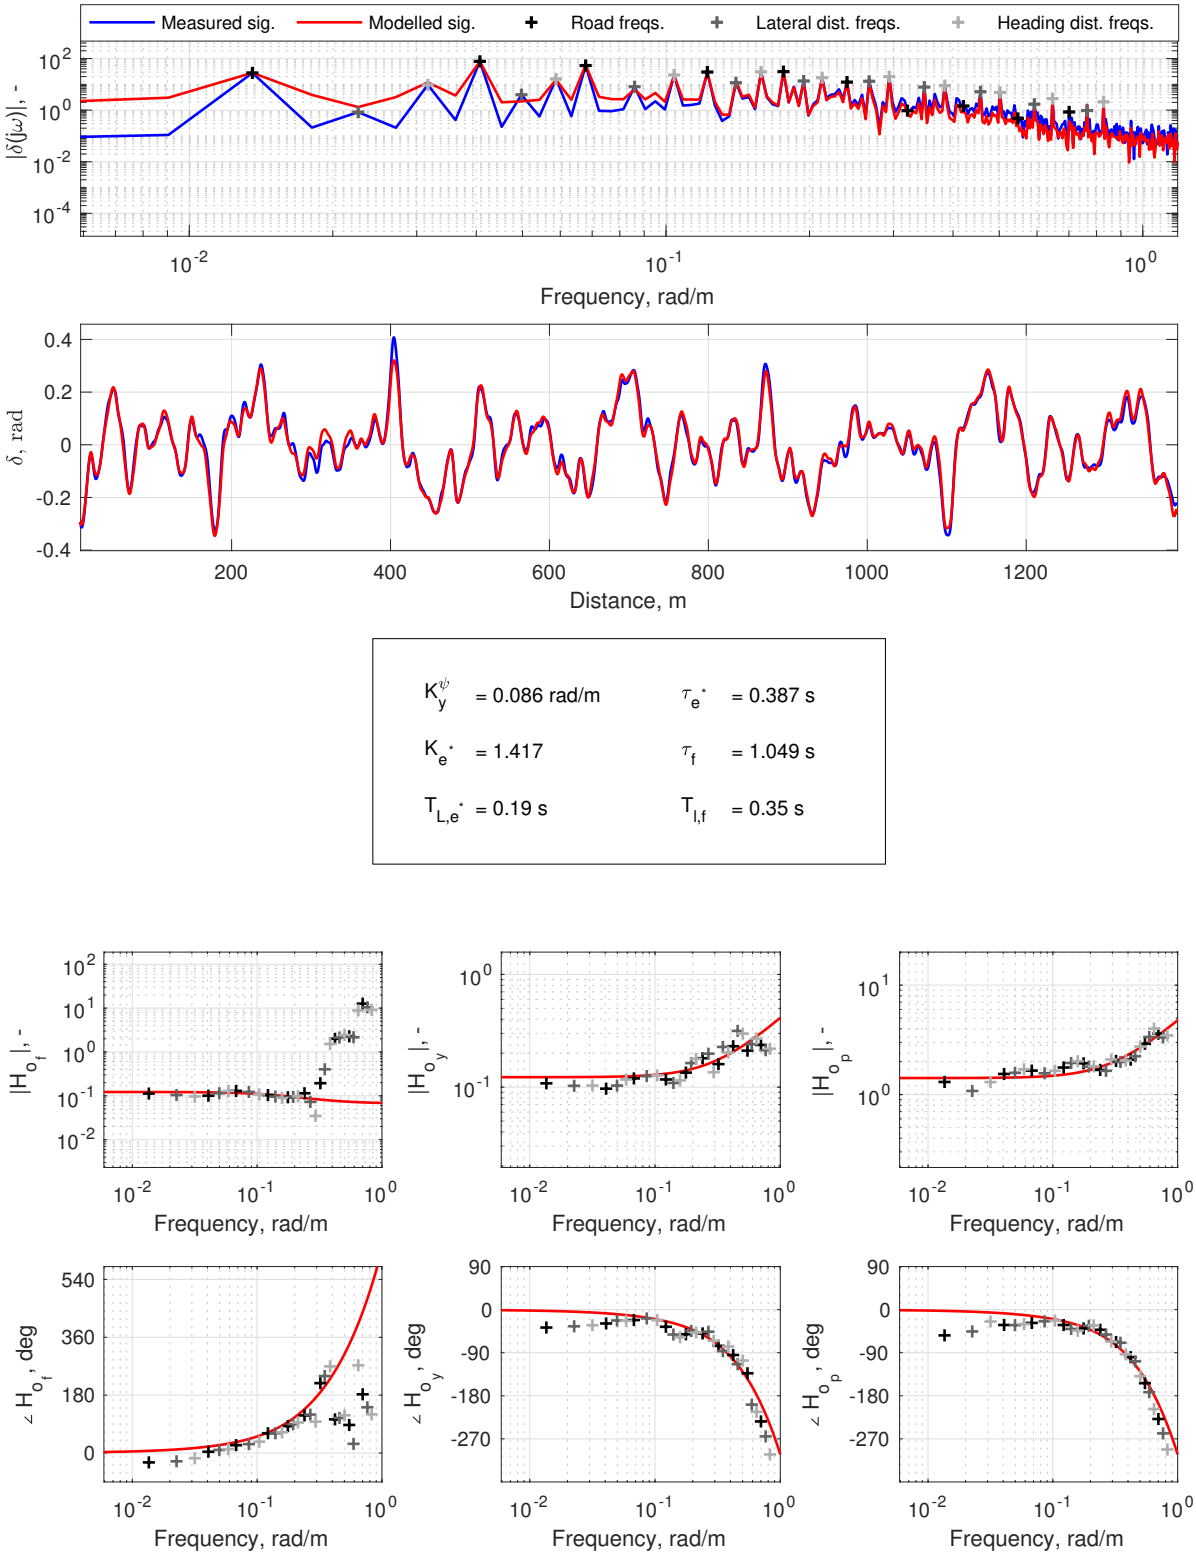

Figure F.20: FRF estimates and parameter estimation, subject 3, Condition B34

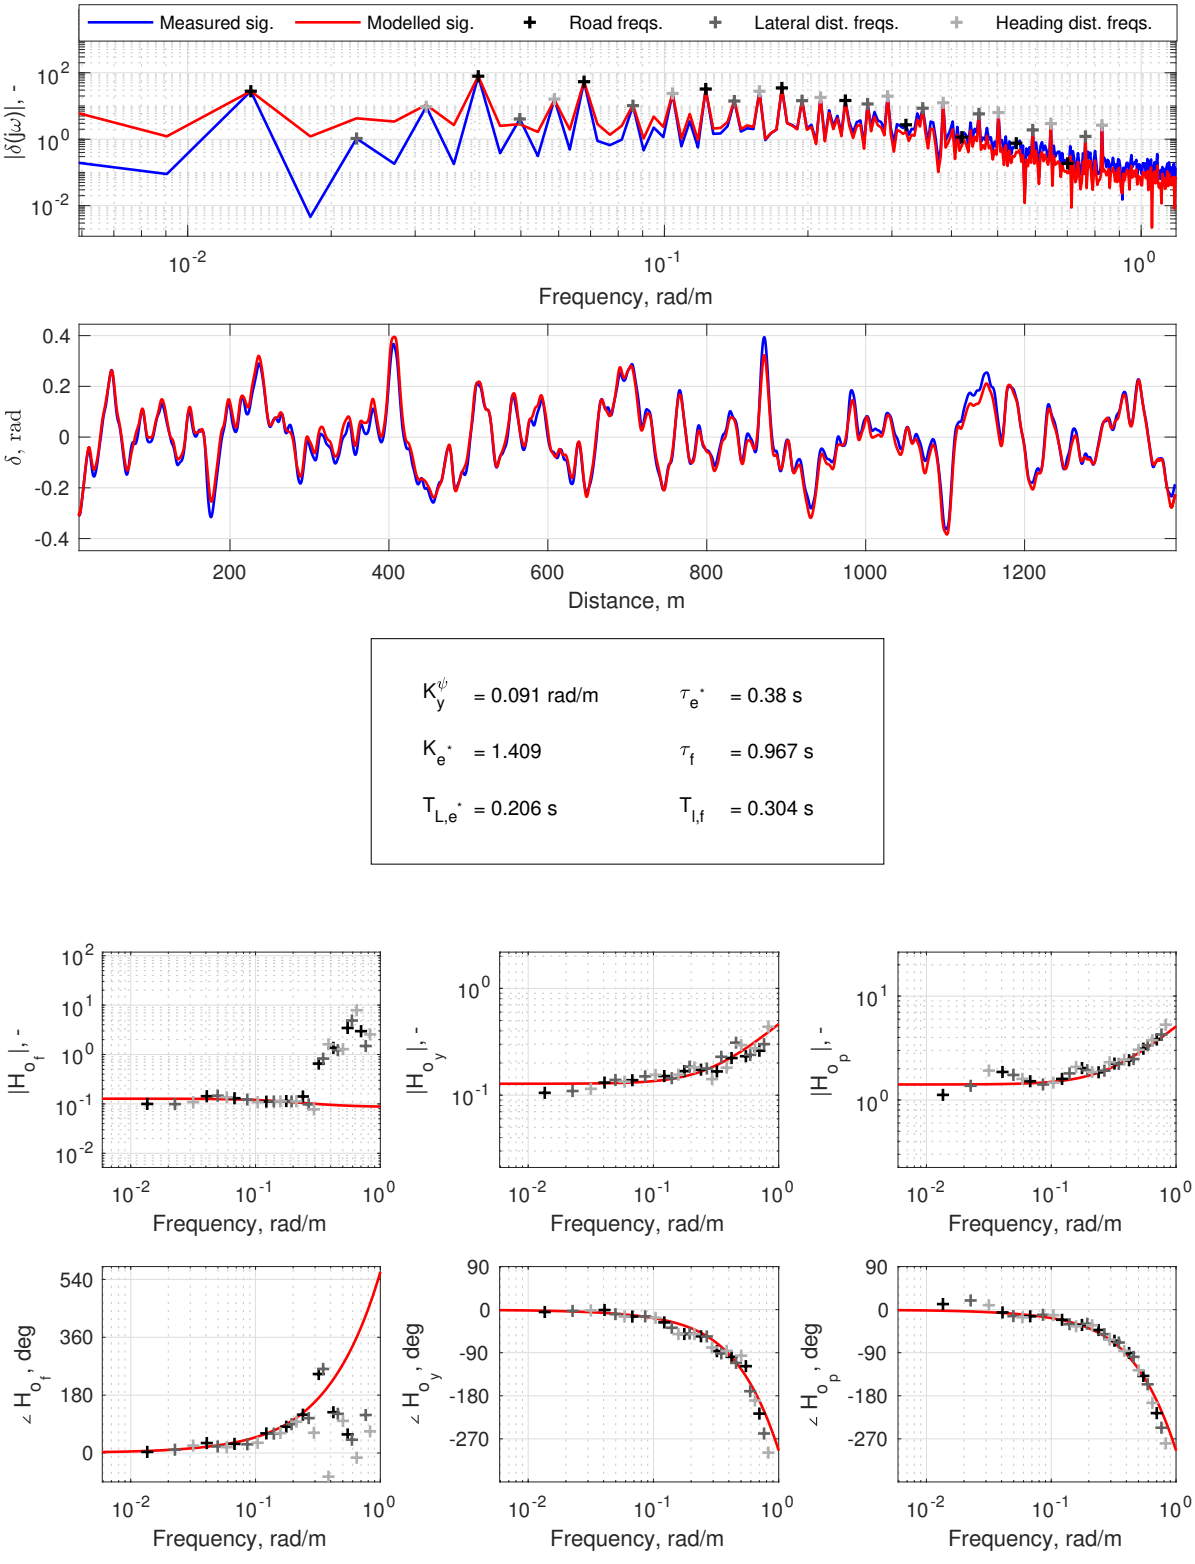

Figure F.21: FRF estimates and parameter estimation, subject 3, Condition B45

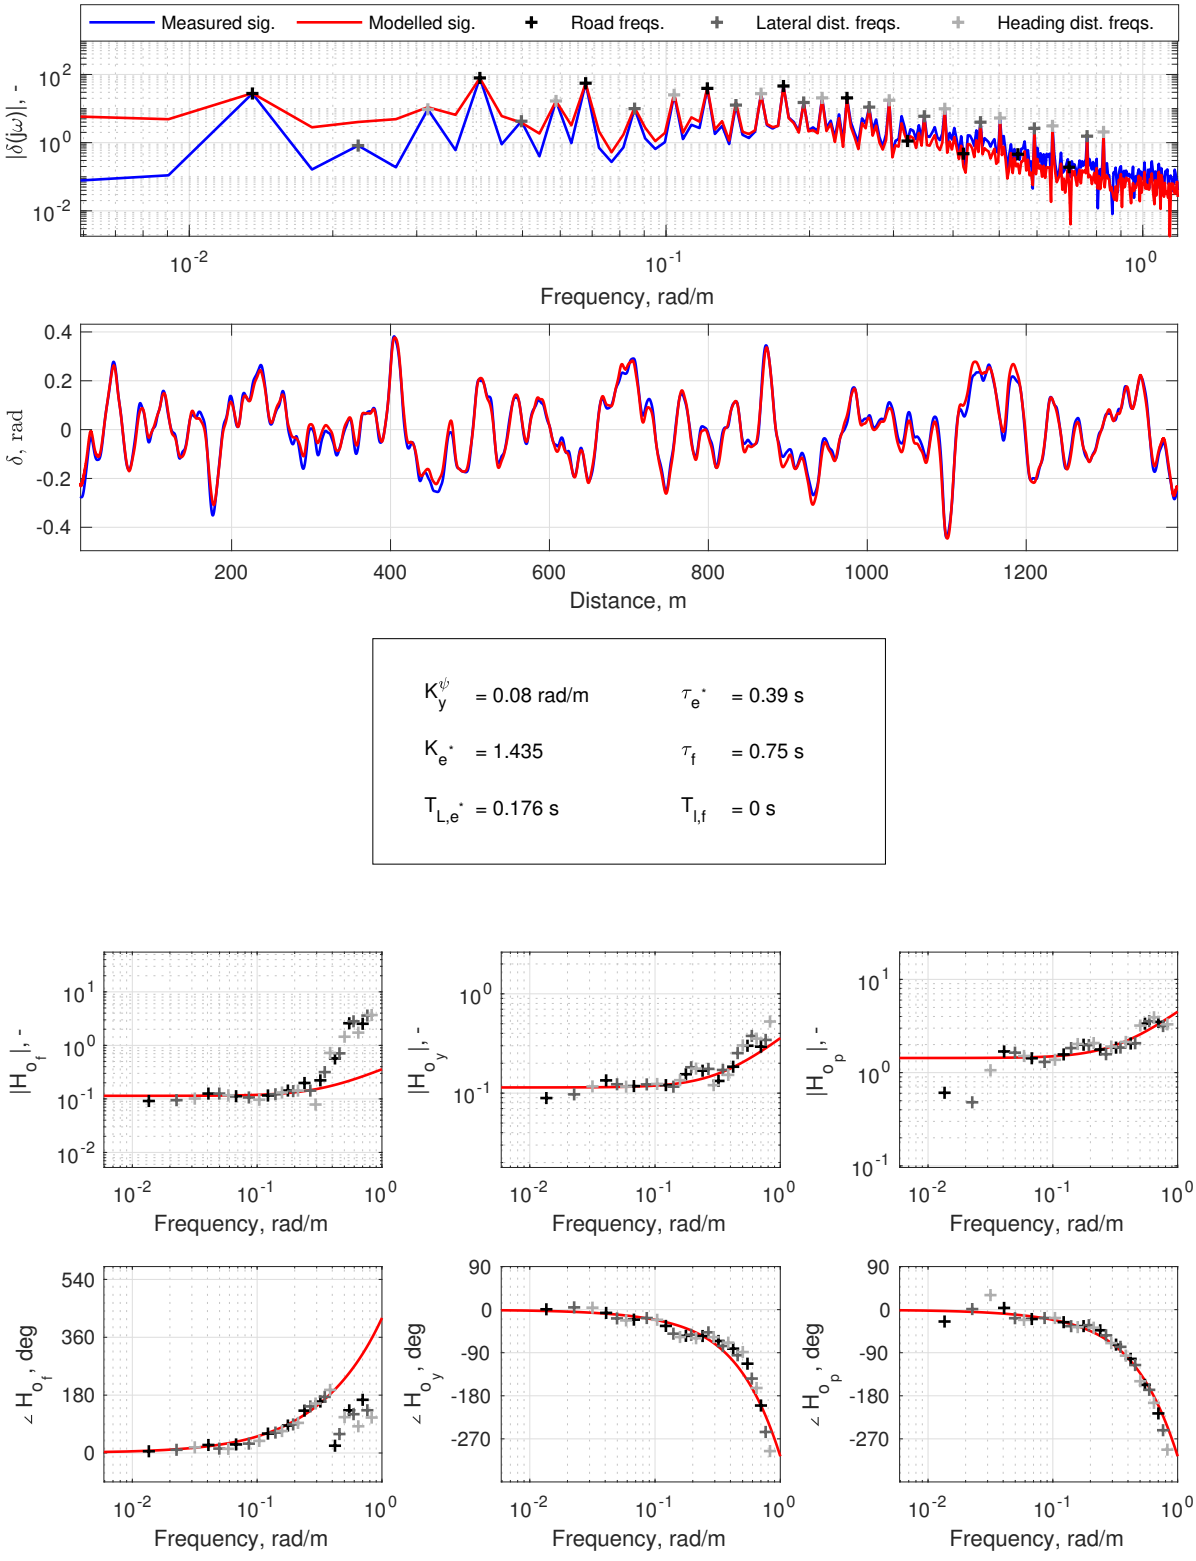

Figure F.22: FRF estimates and parameter estimation, subject 3, Condition B56

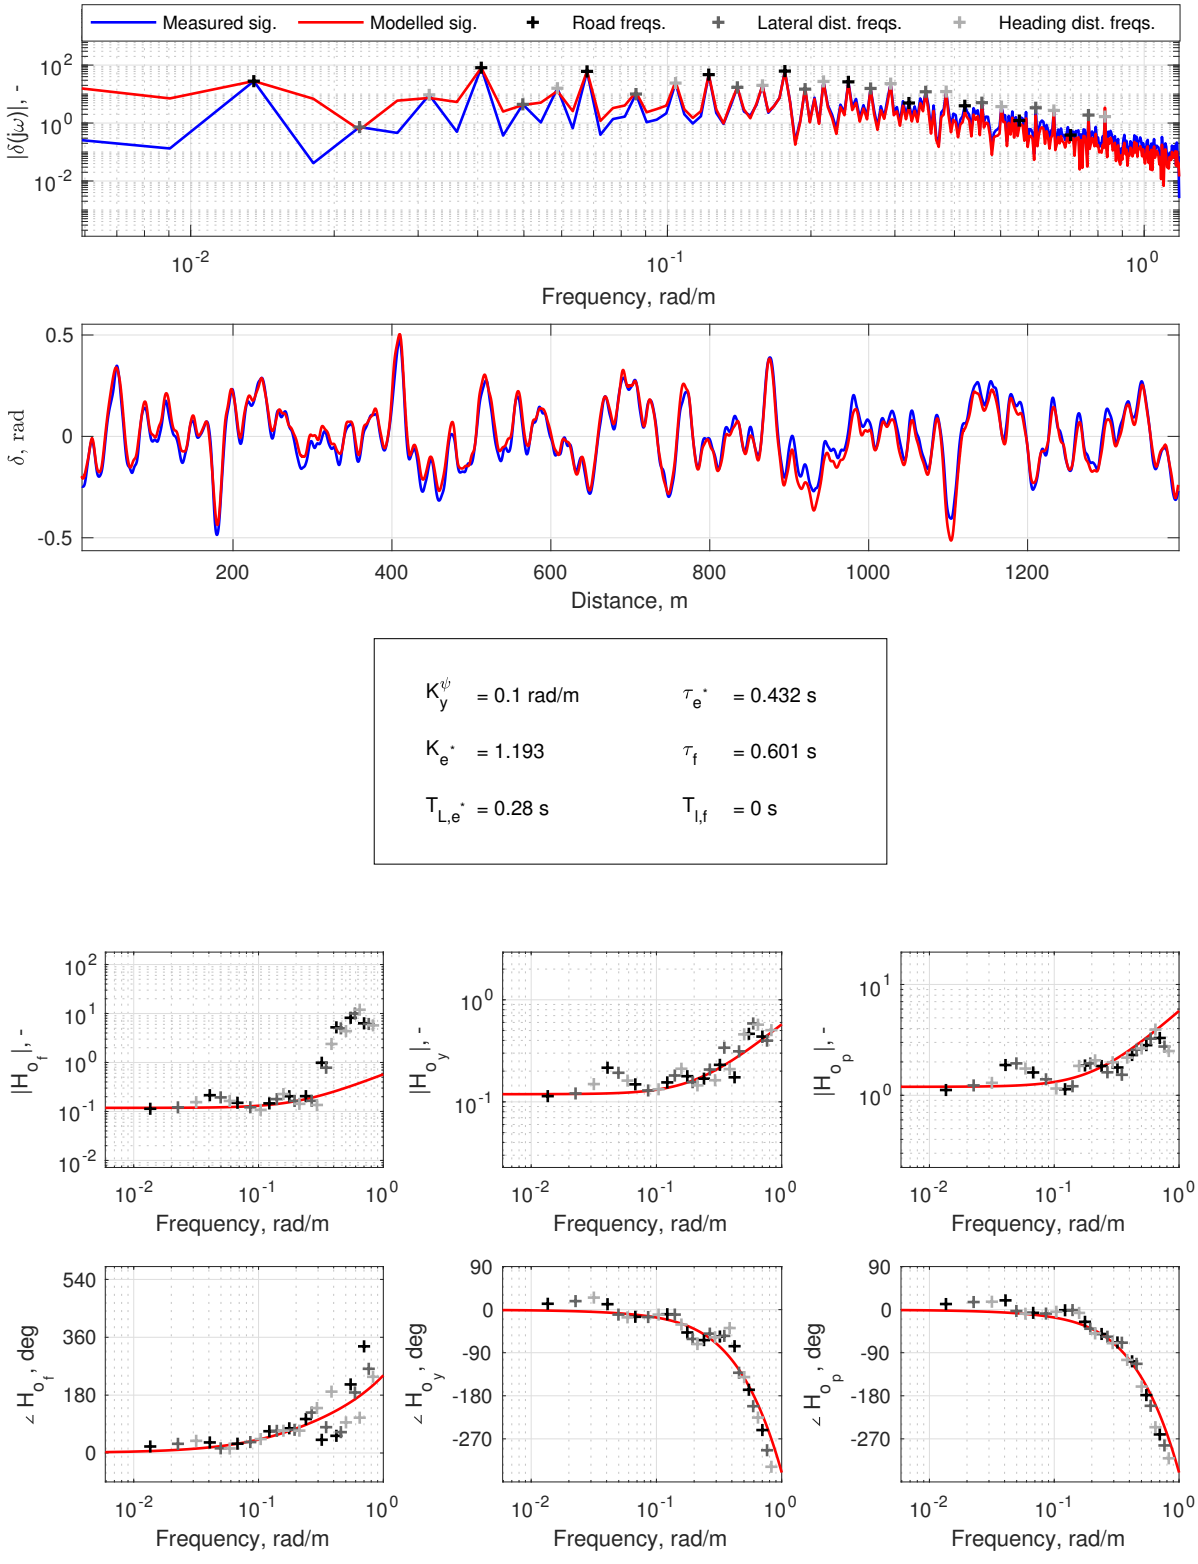

Figure F.23: FRF estimates and parameter estimation, subject 3, Condition B67

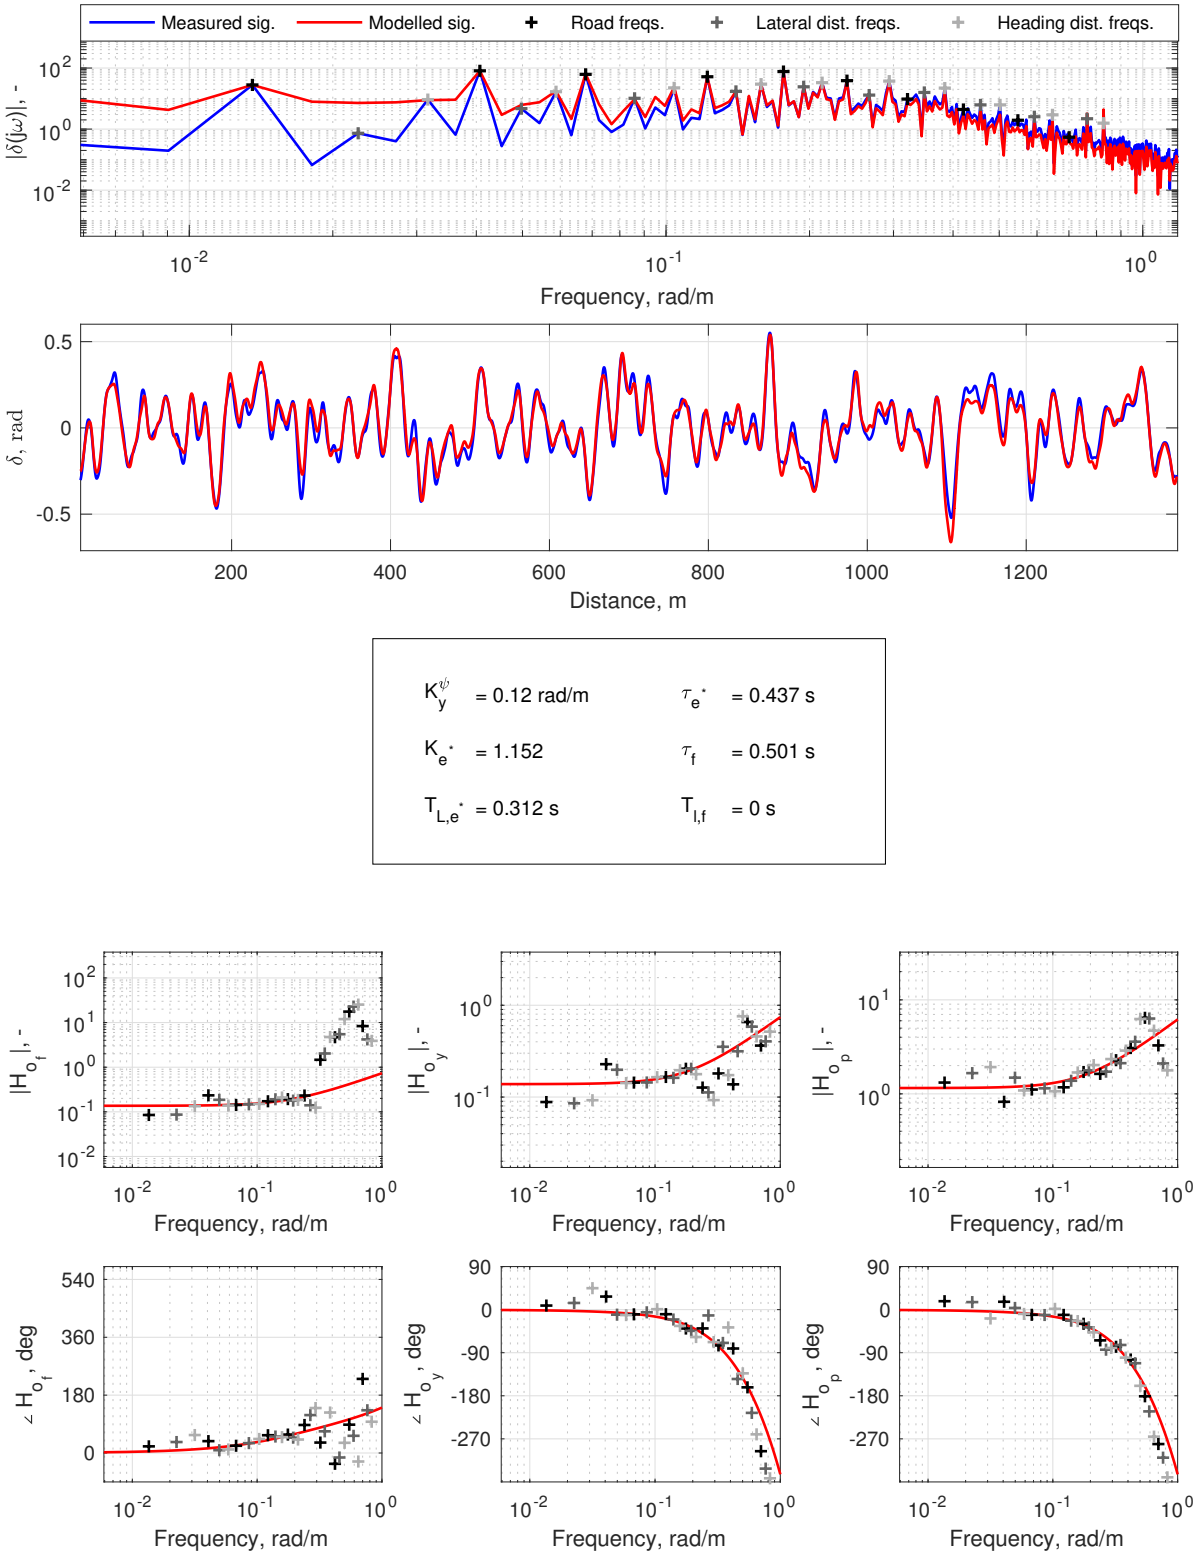

Figure F.24: FRF estimates and parameter estimation, subject 3, Condition B78

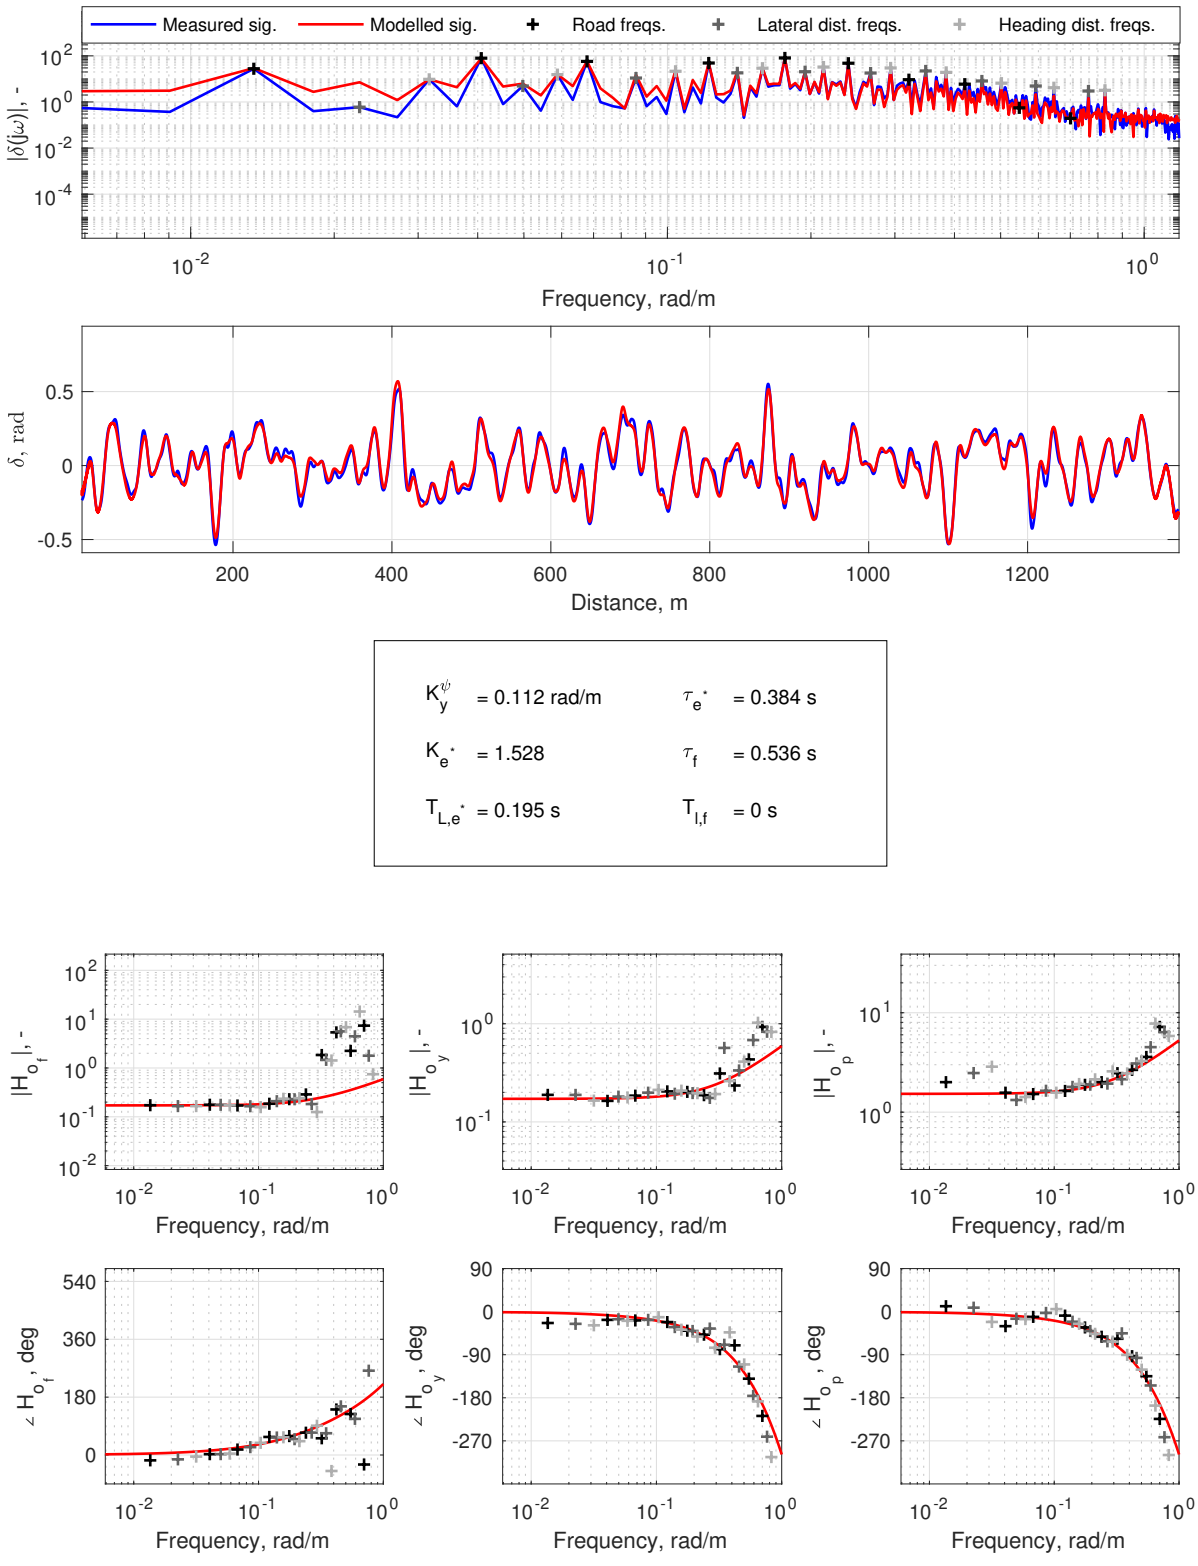

Figure F.25: FRF estimates and parameter estimation, subject 3, Condition B89

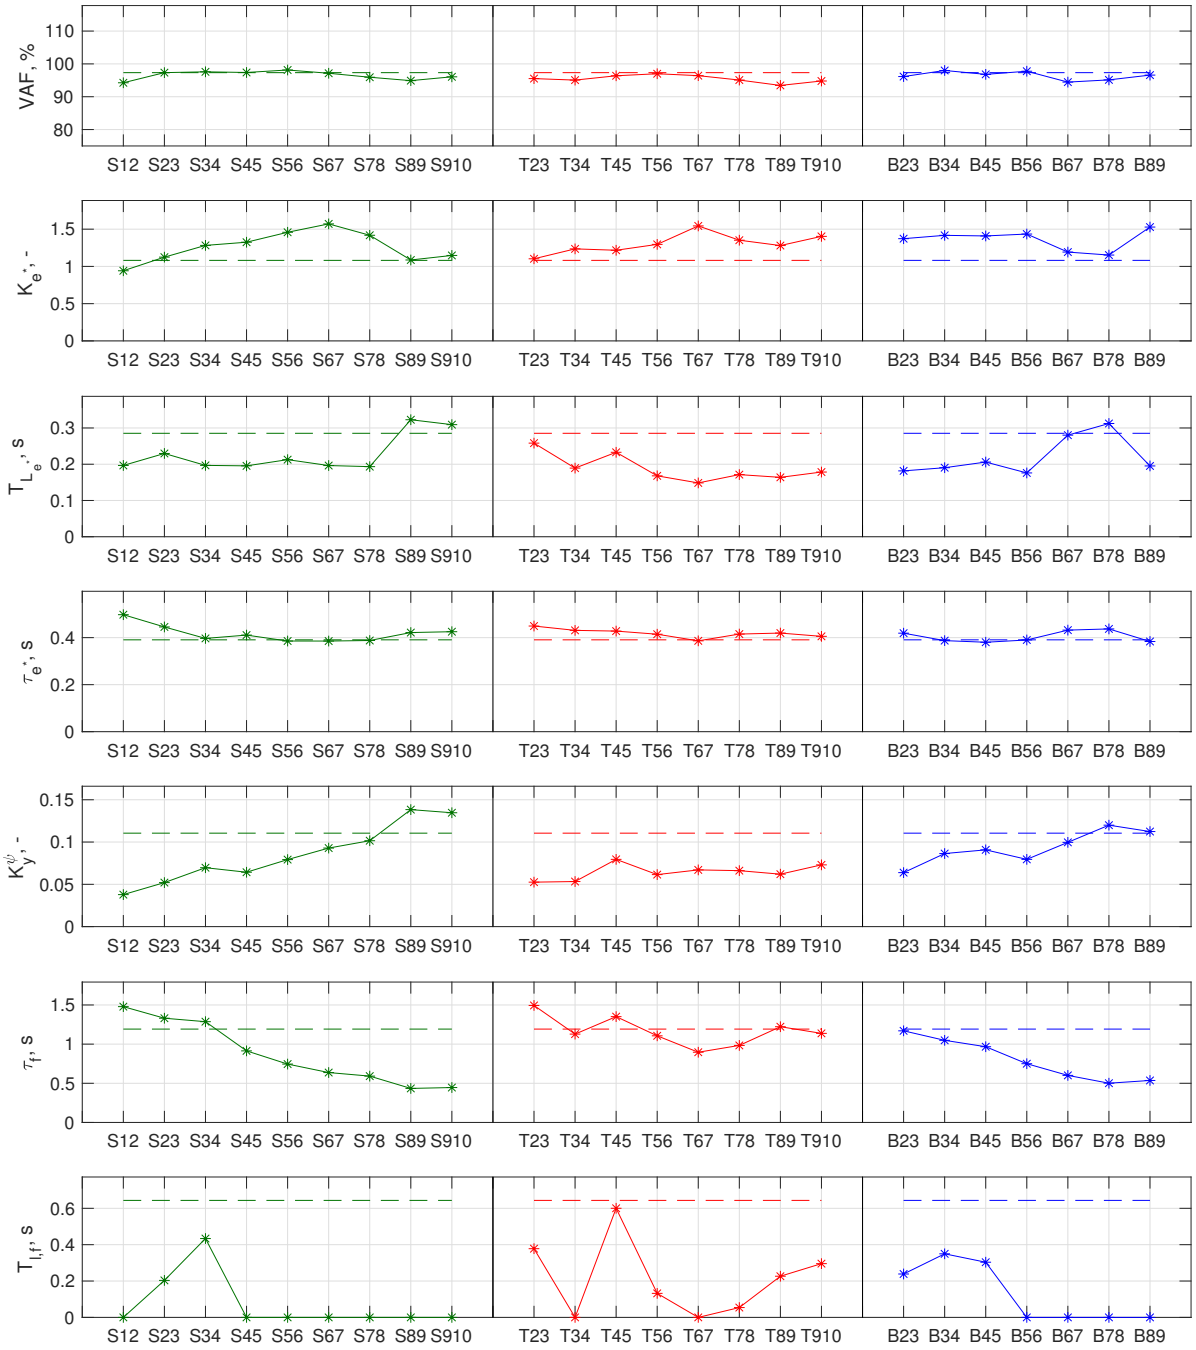

Figure F.26: Parameter estimations over all conditions, subject 3