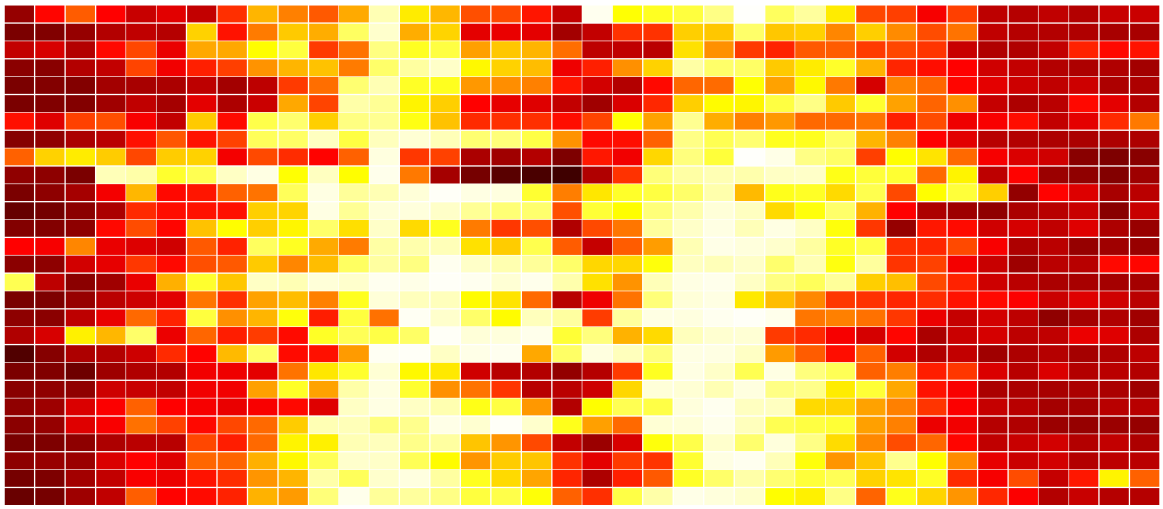

Single Nanoparticle Catalysis

Implementing Machine Learning Algorithms for the Identification of Kinetic Phase Transitions on the Surface of Individual Catalyst Nanoparticles during Surface Chemical Reactions through Nanoplasmonic Spectroscopy

Master's thesis in Complex Adaptive Systems

Henrik Klein Moberg

Department of Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2020

Master's thesis 2020

Single Nanoparticle Catalysis

Implementing Machine Learning Algorithms for the Identification of Kinetic Phase Transitions on the Surface of Individual Catalyst Nanoparticles during Surface Chemical Reactions through Nanoplasmonic Spectroscopy

Henrik Klein Moberg

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Physics
Division of Chemical Physics
Langhammer Lab
Chalmers University of Technology
Gothenburg, Sweden 2020

Single Nanoparticle Catalysis
Implementing Machine Learning Algorithms for the Identification of Kinetic Phase Transitions on the Surface of Individual Catalyst Nanoparticles during Surface Chemical Reactions through Nanoplasmonic Spectroscopy
Henrik Klein Moberg

Supervisors: Olga Serebrennikova, Ievgen Nedrygailov, Giovanni Volpe, Daniel Midtvedt
Examiner: Christoph Langhammer, Department of Physics

Master's Thesis 2020
Department of Physics
Division of Chemical Physics
Langhammer Lab
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Heatmap showing the QMS activity of seven individual nanoparticles across four measurements taken at an experiment conducted at 340^{o}C, calculated using a deep densely connected neural network.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2020

Single Nanoparticle Catalysis
Implementing Machine Learning Algorithms for the Identification of Kinetic Phase Transitions on the Surface of Individual Catalyst Nanoparticles during Surface Chemical Reactions through Nanoplasmonic Spectroscopy
Henrik Klein Moberg
Department of Physics
Chalmers University of Technology

Abstract

The importance of catalysis within industrial and scientific endeavours can hardly be overstated, playing a crucial role in most processes that contain chemical reactions of any kind. This ubiquitous nature of catalysts makes them very relevant for study, yet there is a clear divide between scientific understanding and practical purpose of catalysts - too little is known of the nature of catalysts at the nano-scale within atmospheric conditions. Research within single nanoparticle catalysis, a hitherto elusive field of chemical physics, aims to bridge this gap by analyzing the surface of individual nanoparticles, during reaction conditions that are relevant for practical applications, using a relatively cheap spectroscopy setup. However, the resulting data is so complex and multivariate that standard analysis methods are incapable of elucidating the relevant physical correlations between nanoparticle surface kinetics and spectroscopic results. To remedy this, this thesis' focus is the implementation of a deep densely connected neural network and a long short-term memory recurrent neural network to identify a specific surface kinetic behaviour, known as the kinetic phase transition, within each individual nanoparticle, given only the results of spectroscopy. We find that both networks are reliable, and with a few reservations and admonissions, give generalizable results that can be applied across samples of nanoparticles.

Keywords: nanoparticle, catalysis, nanoplasmonic spectroscopy, surface catalysis, heterogeneous catalysis, machine learning, recurrent network, deep learning

Acknowledgements

I would like to acknowledge Ievgen Nedrygailov, whose insightful descriptions of the underlying chemistry and chemical physics was absolutely vital for the successful completion of this project, and for providing several useful images and graphs that were featured in this thesis.
I would like to acknowledge Gustaf Sjösten, who worked as my de facto partner throughout this entire process, together with whom I co-developed the machine learning algorithms in close collaboration.
I would like to acknowledge Olga Serebrennikova, Giovanni Volpe and Daniel Midtvedt for consistently being willing to meet with me every week and discuss prospective solutions to various problems, as well as for mentoring me in overall appropriate scientific conduct.
I would like to acknowledge David Albinsson for his technical expertise, without which the production of this thesis would not have proceeded nearly as smoothly as it did.
I would like to acknowledge Christoph Langhammer, without the initiative of whom this project would never have existed.
Finally, I would like to acknowledge Myfab laboratories for allowing me access to their cleanroom.

Henrik Klein Moberg, Gothenburg, September 7, 2020

Contents

List of Figures

List of Tables

1

Introduction

Catalysis is the process in which the rate of a chemical reaction is improved by the use of a catalyst, which is not consumed during the reaction. Generally speaking, the rate of reaction improves as a result of new intermediate reaction steps with lower activation energies in the presence of a catalyst within a given reaction. In turn, this enables reactions to take place at lower temperatures and pressures (among other things), which also cheapens production cost in addition to improving production rates. [8] These subtle but vital effects have made catalysts so pervasive and ubiquitous that an estimated $\approx 90\%$ of commercially produced chemicals involve catalysts at some stage of the manufacturing process [9] and an estimated one-third of the GNP of the US involves a catalytic process somewhere in the production chain. Some of the major applications for catalysts today include energy processing, food processing and the production of (bulk and fine) chemicals, but perhaps the single most common use is the catalytic converter within automobile exhausts [9]. Motivated by this vital application, this project is based on an experimental setup in which CO and O_2 (automobile exhaust gases) quickly and cheaply convert to CO_2 through the use of a catalyst.

However, most inquiries and discoveries related to the chemical mechanisms of catalysts are done on bulk catalysts, thereby missing all the intricate subtleties of what occurs on the surfaces of each individual nanoparticle [18]. Indeed, where this project aims to innovate is by developing methods to isolate and investigate single nanoparticles during the catalytic process, to investigate how the size, shape and structure of individual nanoparticles impact their catalytic efficiency. There are reasons to believe that in any given ensemble of catalytic nanoparticles (i.e. a general catalyst), most nanoparticles contribute very little to the overall efficiency of the bulk material.[14] Therefore, unlocking the ability to diligently scrutinise the catalytic properties of individual nanoparticles would potentially make it possible to manufacture a bulk catalyst consisting solely of effective and contributing nanoparticles, leading to significant increases in overall catalytic efficiency. Seen in this light, it becomes clear how advantageous and profitable the successful analysis of catalytic nanoparticles could be to the industrial sector, and how important it could be to the scientific community by garnering more profound insights on catalysis at a fundamental level. Importantly, we note here that single nanoparticle catalysis has hitherto been an exceedingly elusive sub-field of chemical physics, owing to the decades-long issues of the ”materials- and pressure gaps” which soon will be further

elaborated upon.

1.1 Background and purpose

It has become clear that there are fascinating differences between catalytic nanoparticles which profoundly impact the overall efficiency of the entire catalyst, and that these differences can only be resolved by a single nanoparticle approach. Conventional methods of bulk or ensemble catalytic analysis are not capable of resolving the hidden intricacies and influences of nanoparticle surface kinetics for the purpose of catalytic efficiency [3]. For the past few years at Chalmers, a research team in the physics department has been using nano-plasmonic spectroscopy, which combines dark-field microscopy, mass-spectrometry and nanofluidics to study properties of catalytic nanoparticles, such as palladium, in operando. In this method, one studies the scattering spectra of the plasmonic catalytic nanoparticles resulting from irradiating them with light in the visible range of the spectrum. At the nanoscale, when the particle size is comparable to the wavelength of the incident light, the intensity of the scattered spectra from plasmonic nanoparticles is about hundred times higher, when compared to non plasmonic nanoparticles, due to the localized surface plasmon resonance effect. Therefore, in order to observe the catalyst nanoparticles using a conventional optical microscope, these particles are combined into complexes with plasmonic nanoparticles acting as an antenna, which amplifies an optical signal. The catalytic properties of nanoparticles are investigated in a nanofluidic reactor by measuring the optical spectrum of scattered light from the reaction surface over time. Gas molecules are adsorbed onto the surface of the catalyst nanoparticles and react with the reactants, which affects the spectrum of the plasmonic resonance. In this experimental setup, the reaction of interest is $2CO + O_2 \rightarrow 2CO_2$ catalyzed on palladium nanoparticles. To this end, a mass-spectrometer is used to analyse the reaction products along with the optical spectroscopy data, in order to determine the catalytic effect that the catalytic nanoparticles have on the reaction of the injected gases. Experiments take approximately 24 hours to complete, during which period the relative concentration of the injected gases $\alpha = P[CO]/(P[O_2]+P[CO])$ is measured alongside with the spectra of light scattered on nanoparticles. The experiments are usually conducted so that α goes from $1 \rightarrow 0 \rightarrow 1$ linearly for the whole experiment. During this time, gas concentration is withheld in pulses with a duration of about 15 minutes, and the nanofluidic reactor is purged with argon for 15 minutes until the next $CO + O_2$ gas mixture with different concentration is injected.

Given this background knowledge of the experimental methods, it is important to note that previous data analysis methods have hitherto failed to resolve the relevant information from optical data alone. That is; standard algorithms have so far failed to elucidate the correlations between optical data and and surface kinetic changes during catalytic activity. The main observable relevant to previous methods of data analysis has been the shift of full-width-half-maximum (FWHM) over time [13]. If applied to this experiment, we see that the FWHM signal becomes bistable at a given value of α, which is most likely related to kinetic phase transition (KPT) on

the surface of catalyst nanoparticles from reduced to oxidized state. Though, since the signal is roughly on the same order of magnitude as the noise, which arises primarily from gas flow pulses, mechanical vibrations (which cause the particles to shift position by around 10 nm) and thermal effects due to high temperatures around the nanotubes ($\sim$400-500 C^o), this is difficult to analyze in the best cases and outright impossible in the worst.

The data that can be used for analysis is captures of the scattering spectra of plasmonic nanoparticles at a given time interval, readings from mass-flow controllers about the gas composition at the reactor inlet, and mass-spectroscopic data on the rates of a catalytic reactor in the nanofluidic reactor. Currently, the spectra are recorded at a time interval of 20 seconds. The collected scattering spectrum is the average of 20 spectra that is subsequently measured by a machine every second.

In the experiment, three ”batches” of catalytic nanoparticles are located within the nanofluidic reactor. First, there is a ”row” of single particles, which is followed by an array of around a thousand particles which the gases pass through in order to reach a final row of particles on which one can measure spectra from single particles. This specific structure allows spectra to be measured from single particles in the two rows. Furthermore, one single particle is on the order of a few pixels in the measurements, whereas the measurement from one spectrum is gathered from selecting a ”window” of around 10-15 pixels wide around each particle. Finally, there is a team working on simulating the chemical process with CFD-methods, some of which might be of interest to the further validation of data analysis that has been conducted in this thesis.

1.1.1 Materials- and pressures gap

Finally, we return to the materials- and pressures gap that was eluded to in the introduction. The crux of the materials-and pressure gap is that analysis of the surface catalysis of single nanoparticles has been done succesfully in high vacuum conditions, but relevant for industrial applications are usually conditions around atmospheric pressures or higher. This technical issue gives rise to the so-called pressure gap, which is a consequence of the fact that high pressures render adsorbate-adsorbate reactions very important. The materials (and complexities) gap is less fundamental, and is instead the consequence of the gap between simplified linear models and the complex heterogeneous effects of real catalysts under experimental conditions. The behaviour of a single catalytic crystal is often very different to that of a whole (real) catalyst, making fruitful analysis of singular catalytic particles incredibly difficult. [1]

Several attempts at remedying this disparity have so far been partially or entirely unsuccessful, with the most promising approach being in the use of synchrotrons [14]. However, since synchrotrons are generally prohibitively expensive and difficult to access, this project therefore aims to investigate the reaction kinetics of single-nanoparticle catalysis through a non-invasive and cheap source: the optical spectrum of scattered light (achieved via plasmonic resonance). Most other attempts at

bridging the materials- and pressure gap has revolved at making evermore complex theoretical models and computational simulations of catalysts at the particle level, which has of yet not been able to resolve the aforementioned issues [1].

The brief reasonings above form the basis of this thesis' involvement in the story, which presents a new and novel method of bridging the gaps and analyzing the activity of catalysts at nanoparticle resolution. We know from theory that the correlations between wavelength shifts in optical data correspond to kinetic changes on the surface of catalytic nanoparticles, but the correlations are much too complex to resolve with any hitherto attempted analytical or computational method. The aim of this thesis has therefore been to apply modern machine-learning algorithms to correlate series of optical spectra, resultant of scattered light on nanoparticles under experimental conditions, into kinetic changes on the surface of these individual nanoparticles.

1.2 Aim

The aim of my thesis is to find the concentration of input gases which result in a kinetic phase transition on the surface of individual nanoparticles using only optical data. The reason why this is desirable is because it enables deeper analysis of the change in surface kinetics of individual nanoparticles in atmospheric (or higher) conditions, which is of particular relevance to most real-world applications. To this end, several machine learning methods have been tested and two have been fully implemented. Firstly, a Long Short Term Memory (LSTM) network is used to analyse the change in FWHM over time based on the team's previous hypothesis that the reaction kinetics fundamentally change after the individual nanoparticle's KPT, which can be measured as a shift in the behaviour of the FWHM signal. The network will therefore be trained upon simulated time-series data to recognize this shift in behaviour, and output a prediction of the mixing range of input gases that it is most likely to occur in. Secondly, a deep neural network is trained on the entire spectral data with the QMS signal as target, to attempt to qualitatively correlate optical spectra to QMS signals and thereby calculate QMS activity based solely on spectral data. Depending on the network's accuracy, the resultant calculated QMS signals can then be analyzed in further studies using standard methods within chemical physics, as will be discussed in chapter 2, to understand the behaviours of each individual nanoparticle.

2

Theory

This section is dedicated to the most relevant theory underlying the experiment and analytical method. In practice, this means an initial run-through of catalytic theory and how its efficiency may be measured, followed by a discussion on optics & plasmonic resonance particles and finally a brief review of the mathematical framework behind the machine learning methods that have been implemented for the purposes of data analysis. Note that all discussions on chemistry, chemical physics and optics will be relatively brief, in comparison to similar articles on the subject, as this thesis' main focus is not on the physics underlying the experiment, but rather on the resulting data analysis.

2.1 Catalytic Surface Reactions and Conventional Methods of their Bulk Measurement and Characterization

The most widely accepted definition of catalysis today was first formulated by Nobel laureate Wilhelm Ostwald around 1900 as "*A catalyst is a substance which affects the rate of a chemical reaction without being part of its end products*" [8]. We initiate this section on heterogeneous catalysis theory by briefly describing catalysis in general, followed by more specific discussions on the type of catalytic reactions of interest to this particular study, followed in turn by discussions of theory behind the optical method and setup for the experiment.

In general, chemical reactions involve the breaking of chemical bonds between or within molecules and the subsequent forming of new ones. This process is associated with a transfer of energy, and related to an activation energy E_{act} such that the reaction is (effectively) impossible if the excess energy of the reactants is less than E_{act} [16]. If the excess energy of the reactants is provided by thermal energy, then the probability of the reaction occurring is given by the Boltzmann distribution $P = e^{-E_{act}/k_bT}$, meaning that the overall reaction rate increases with higher temperature and decreases with higher activation energy [16]. By introducing a catalyst to a reaction, the activation energy is effectively decreased by providing the reactants with intermediate reaction steps that are energetically favourable to the uncatalytic reaction [8]. Important to note here is that most relevant technical applications of

catalysts take place on the catalytic surface, leading to catalytic reaction rates being very sensitive to the size, shape and structure of the catalyst's surface [8]. One fairly well-studied aspect of single particle catalysis is that of reactions with reactants in the liquid phase, which was first successfully attempted in 2008 by Mulvaney et.al [15]. In this work, surface plasmon spectroscopy was implemented to study the catalytic behaviour on the surface of single gold nanoparticles during oxidation of ascorbic acid. However, a much less studied and understood aspect of single particle catalysis is that of reactions at gas phase, which generally puts experimental instrumentations under much more stress as a result of higher temperatures and/or pressures [3]. This leads well into the next section, where we discuss the surface reaction of CO and O_2, the relevant reactants of this study, on catalytic palladium nanoparticles.

2.1.1 Oxidation and Reduction of Palladium Nanoparticles

As mentioned in the introduction, one of the most relevant reactions to study within the field of catalysis is the catalyzed production of carbon dioxide from carbon monoxide and oxygen gas. Relevant for this study is the aforementioned reaction over palladium nanoparticles, since these provide a platform suitable for investigating Localized Surface Plasmon Resonance (LSPR) behaviour through optical means, as is elaborated upon in section 2.2. In this section, we first introduce the elementary oxidation reaction before delving deeper into the deceptively complex intricacies taking place on the surface of the palladium nanoparticles, which profoundly affect the behaviour of scattering of optical light across the surface.

To be perfectly clear, the reaction we are ultimately interested in can be written as:

$$CO \ + \frac{1}{2} \, O_2 \ \rightarrow CO_2. \tag{2.1}$$

In order for this reaction to occur at relatively low temperatures and pressures (relevant to most industrial and automotive applications), it is necessary to include a catalyst in the reaction to reduce its required activation energy [2]. In the presence of a catalyst, reaction 2.1 is split into four elementary steps; adsorption of CO and O on the surface of the catalyst, a surface reaction and finally desorption of CO_2. Let $\odot$ represent an active site on the surface of a catalyst, and $X^{\odot}$ represent the adsorption of molecule X on the surface of the catalyst. The catalyzed reaction can then be expressed as [2]

$$\begin{aligned} CO + \odot &\rightarrow CO^{\odot}, \\ O_2 + 2\odot &\rightarrow 2O^{\odot}, \\ CO^{\odot} + O^{\odot} &\rightarrow CO_2^{\odot} + \odot, \\ CO_2^{\odot} &\rightarrow CO_2 + \odot. \end{aligned} \tag{2.2}$$

This is an example of the Langmuir-Hinshelwood mechanism, whereupon the maximal reaction rate is expected to occur when the surface coverage of carbon monoxide

θ_{CO} and oxygen θ_O (oxygen gas separates upon adsorption with the surface) are equal, according to equation 2.3.

$$r = k\theta_{CO}\theta_O, \tag{2.3}$$

where k is constant of the system. We note here that the same reaction can occur by the Eley-Rideal mechanism if CO in its gas form reacts with surface-adsorbed O, but this is a comparatively small portion of the total reaction rate and will therefore be neglected [12]. Now, as has already been hinted at previously, a clear issue arises from the fact that CO adsorbs onto a single site on the surface, whereas O_2 needs to dissociatively adsorb onto two sites, whereupon they are then separated by two lattice constants on the catalyst's surface [18]. Since the gases compete for the same active sites on the surface, the adsorption of CO can therefore become much greater than O_2 in relation to their input concentrations, which leads to lowered CO oxidation over time [18]. This phenomenon is known as CO poisoning, which is discussed further in section 2.3.

An allegorical issue arises in the opposite region, when the Pd surface becomes contaminated with oxygen and starts to form surface oxides (PdO) [18]. This form of O-poisoning can be produced, for example, as a result of the active catalyst depleting the gas phase of CO but not O_2 (i.e. by different input concentrations). These surface oxides that form as a result of O-poisoning can be highly reactive in certain conditions because defects in the oxides bind CO much more strongly than the rest of the surface [17]. In general, the reaction rate of CO oxidation on the surface of Pd is exceedingly high, to the extent that both reactants may become almost entirely depleted on the surface [18]. In this case, which is known as the mass-transfer limited (MTL) regime, the reaction rate instead becomes dependant of the gas-phase diffusion towards the catalyst [18]. As a direct result of this obscuring MTL effect, it remains unclear exactly why surface oxides have such high reactivity, and how different oxides may compare against eachother [18]. There are good reasons to believe that the MTL regime does not affect the results of this experiment, as is argued for in section 2.3 below.

The issue of CO poisoning can effectively be solved by increasing the temperature, resulting in the desorption rate of CO being higher than the desorption rate of O_2 [2]. Indeed, it is this phenomenon that leads to the so-called "cold-start" problem, where a vehicle's catalytic CO converter may malfunction at low temperatures [2].

As will be made clearer in section 3 below, the phenomenon of CO poisoning leads to two distinct regimes of the aforementioned reaction. Firstly, when the surface is poisoned with CO to the extent that the reaction rate is negatively proportional to input CO gas concentration. Secondly, when the surface coverage of CO is low enough that the reaction rate is positively and linearly proportional to input CO gas concentration. Between these two regimes, a kinetic phase transition occurs, which is categorized as a considerable change in a reaction kinetics in a relatively brief input concentration mixing range.[2] Gerhard Ertl et. al. famously showed that these kinetic phase transitions correlate with a structural reshaping of the top surface layers of platinum, which in turn changes the dynamics of surface reactions

[8]. Similar effects are likely to occur on the surface of palladium [18]. As will also be made clear in section 3, this results in spatio-temporal oscillations of the reaction rate even at constant rate of input gas concentration. A summary of these kinetic phase shifts on the surface of palladium is shown in figure 2.1 as functions of reaction temperature, pressure of input gas, and chemical potential of components $\Delta\mu$. It is outside of the scope of this thesis to describe in detail how these experimentally discovered phase transitions form, but the relevant takeaway is that the surface of palladium catalysts are very complex and multivariate, and may well pass through several phases throughout a single experiment.

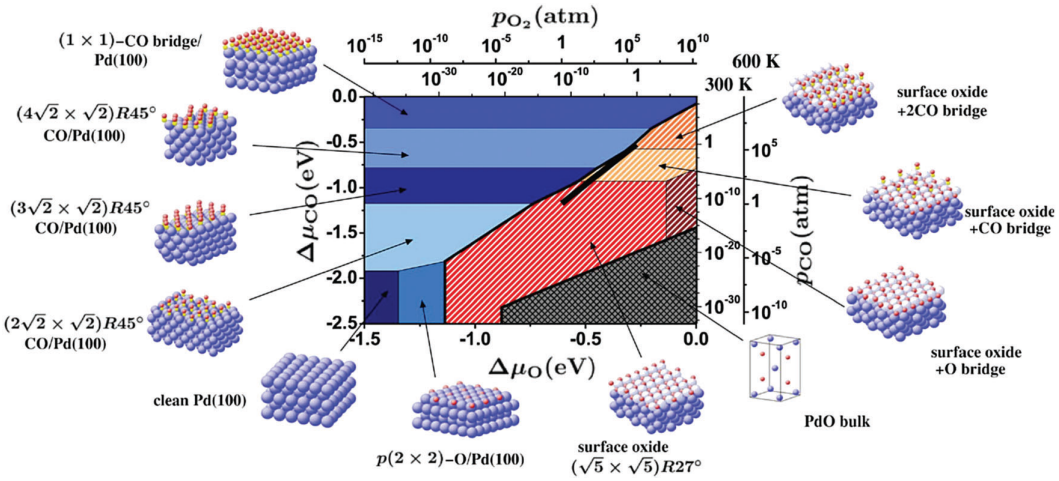

Figure 2.1: Surface phase diagram of the Pd(100) surface in equilibrium with the gas phase (with the restriction that CO cannot react with O2). The black bar indicates relevant reaction conditions (101 kPa, 300–600 K). Ball models visualize the different structures (metallic Pd, blue; oxidized Pd, gray/light blue; O, red; C, yellow). From *Surface science under reaction conditions* [18] - Published by The Royal Society of Chemistry.

2.1.2 Mass Spectroscopy - Quadrupole Mass Analyzer

In this section, a brief introduction to mass spectroscopy will be followed by a more specific discussion regarding the exact type of mass analysis used in this study: quadrupole mass analysis.

Generally speaking, mass spectrometry is an analytical technique used to measure the mass spectrum of a given sample of particles or molecules. Principally, this is done by measuring the mass-to-charge ratio of ions, whereupon each separate type of molecule or particle can be distinguished and its total incident mass represented as intensity as a function of mass-to-charge ratio [7]. The specific manner in which incoming particles are ionized may vary between solutions, but it may for example be achieved by bombarding them with electrons [7]. Upon bombardment, some particles may just become charged whilst other may become fragmented before becoming charged, which is important for distinguishing molecules that have the same mass-to-charge ratio [7]. When this has been accomplished, the now-ionized sam-

ples can be separated by their mass-to-charge ratio through various methods, which will soon be elaborated upon further [7]. Finally, the now-separated samples hit a detector connected to some form of counter, resulting in the aforementioned mass spectrum representing the number of incident particles of a given mass-to-charge ratio as a MS intensity signal [7].

The specific MS method used in this study is the Quadrupole MS method, whereupon its namesake comes from the detector's use of four parallel rods, through which incoming ions' trajectories are modified [7]. Two of the parallel rods are put under a combined DC and AC potential, with the other two remaining parallel rods also put under the same combined potential, such that each opposing pair of parallel rods have the same combined potential [7]. This technique utilizes the fact that only particles of a certain mass-to-charge ratio will successfully travel the entire length of the quadrupole for a certain ratio of DC and AC voltages, whereas all other particles simply collide with the surrounding rods due to their unstable trajectories [7]. This is shown schematically in figure 2.2. Indeed, through solving the relevant Mathieu's differential equations, the exact trajectories of particles of well-defined mass-to-charge ratio within a given voltage ratio can be calculated [7]. Based on these equations, the voltage controller can be set to "scan" specific ratios of voltages corresponding to the successful transfer of particles with given mass-to-charge ratios, by oscillating continuously between different ratios of voltages [7].

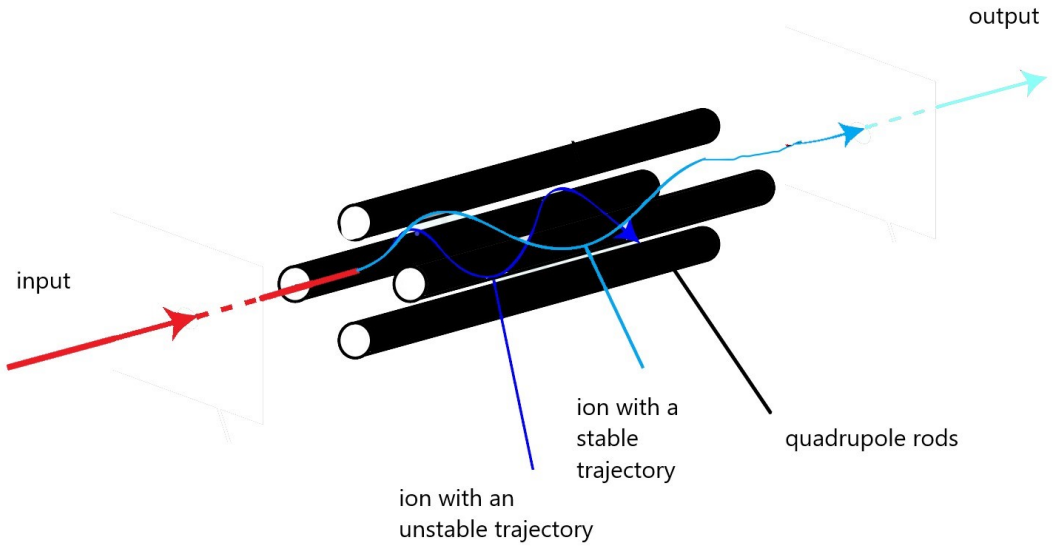

Figure 2.2: Schematic of the function of quadrupole mass spectrometers. Ions pass through four rods which are put under a combined DC and AC potential, with each pair of rods having the same combined potential. Only particles of a certain mass-to-charge ratio will succesfully travel through to the detector and be measured.

Furthermore, it is important to be aware that the QMS filters can only distinguish ions by their mass-to-charge ratio, which is a problem if sample particles or molecules have different mass but identical mass-to-charge ratio. [7] This problem can be

resolved, for example, by measuring fragments of molecules that form as a result of electronic bombardment, as was hinted upon previously. [7] For instance, if the sample molecules are N_2 and CO, which have different mass but identical mass-to-charge ratio, then they can be distinguished by bombarding CO and N_2 such that they separate into N, C and O particles, whereupon one can measure the QMS intensity signal of the individual particles instead of their aggregates [7]. In this specific case, then, none of the resulting individual atoms have identical mass or mass-to-charge ratio. However, other issues may occur as well, such as higher order ionizing events that result in atoms being detected as having higher charge, and therefore lower mass, than they actually do. It is therefore important to be aware of these potential drawbacks and complications when operating a mass spectrometer of any kind. [2]

2.2 Optics & Plasmonic Resonance

This section of theory introduces and discusses the optical theory upon which the experimental optical setup is based. Firstly, the discussion covers the fundamental principles of (plasmonic) particles which allow us to measure extremely subtle kinetic changes on the surfaces of particles through optical light. Secondly, we discuss and justify the practical methods used to implement and utilize these unique features of (plasmonic) particles to probe the surface of catalytic nanoparticles in this experiment.

2.2.1 Plasmonic Nanoparticles

Generally speaking, chemical processes are fairly well understood in the limits of sizes of atoms ($\sim$1 Å), and in the limits of bulk materials ($\gtrapprox$ 1μm). However, complex and considerably less well understood physics emerges in collections of small amount of atoms, which are colloquially known as nanoparticles because of their size usually being between 1 and 100 nm. Indeed, the inherent properties of these nanoparticles have fairly recently been shown to be very dependant on their particular size, shape and structure, greatly complicating their rigorous analysis. [10]

A famous and historically relevant example of plasmonic nanoparticles are Faraday's studies of how the size and shape of gold nanoparticles affect their observed colour, as illustrated in figure 2.3. The important takeaway from these results is that simply enlarging a gold nanoparticle does not necessarily change its optical properties in any meaningful way. However, adding an anisotropy to the nanoparticle, such as growing nanorods in this case, the observed properties may change considerably. This tends to be true in general for metal nanoparticles, but not necessarily for nanoparticles of other materials.

The mean free path of free electrons in gold particles is $\approx 50nm$, so interactions on nanoparticles smaller than this are only expected to occur on the surface [10]. Results of scattering and absorption spectroscopy has shown that standing resonance

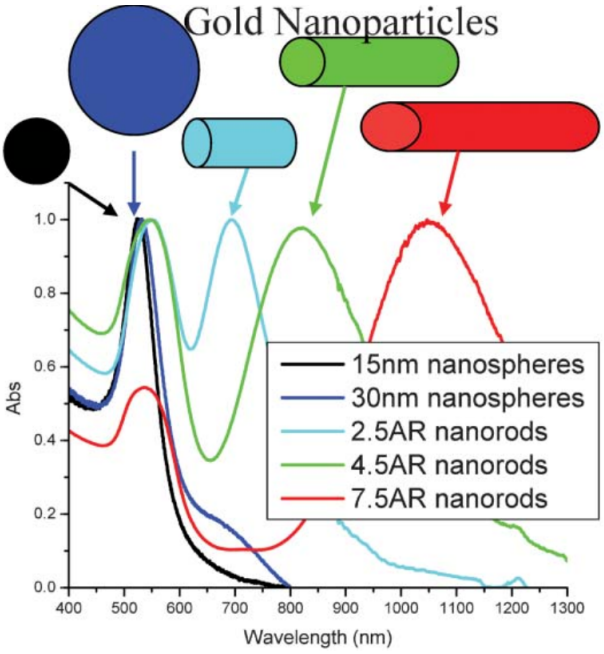

Figure 2.3: Absorption of electromagnetic radiation within gold nanoparticles as a function of wavelength, shown here for five different samples of gold nanoparticles. Note that differently sized gold nanospheres have similar optical properties, whereas differently sized gold nanorods have very different properties. Reproduced from [10] with permission from The Royal Society of Chemistry.

can occur for the electrons in the conduction band of a given nanoparticle, if the wavelength of the incident light is comparable to the nanoparticle's size [10]. This occurs because the wave front of incident light polarizes the electron density to the surface, thereby causing the electron density to oscillate with the frequency of incident light in a standing oscillation [10]. The observed colour is a consequence of these collective oscillations of electrons within the conduction band of the nanoparticle, also known as surface plasmon oscillations, as will be discussed further in section 2.2.2.

2.2.2 Localized Surface Plasmon Resonance (LSPR)

The recent developments of nanoplasmonic sensing as a versatile tool for in situ spectroscopy within catalysis was pioneered largely by the Langhammer team, after the method was first used in biosensing applications as early as 1998 by Englebienne [13]. However, before we delve deeper within the intricacies of in situ spectroscopy for catalysis applications, it's important to understand the basis upon which the theory is built; localized surface plasmon resonance (LSPR).

As discussed in section 2.2.1, it is known that a metallic surface might excite resonant collective oscillations in its electronic system upon interaction with light with a wavelength comparable to the nanoparticle's size. Simply put, these oscillations are known as LSPRs if they are localized to an entity of size on the order of nanometers, such as nanoparticles within the metal [13]. As a direct result of the resonant na-

ture of these oscillations, the metal nanoparticle interacts very strongly with light, which then enables possibilites to measure scattered or absorbed light that might otherwise have too low intensity to detect [13]. The polarization of the electronic system (as a result of high-wavelength incident light as discussed in section 2.2.1) amplifies the strength of local electric fields with respect to any given incoming field, which means, in principle, that any local perturbations which affect the polarization of the electronic system can be detected within the sensing volume [13]. In practice, according to Langhammer et al., these perturbations may be the addition or removal of moieties, rearrangement of material at the nanoscale, phase transitions, solid-state reactions or the presence of adsorbate molecules. In short, therefore, LSPR provides a method of translating relevant physical changes in a particle's size, shape or composition at the nanoscale to an (albeit small) shift in the resonance wavelength, which can be measured by relatively simple optical spectroscopy [13]. Important to note is that the sensing volume (correlated with the strength of the amplification of the local electric field) decays exponentially at low distance r and linearly at greater r. [13] This effect is illustrated in 2.4 below, where an electric field interacts with plasmonic sensing particles such that they effectively achieve a sensitive "sensing volume" extending a few tens of nanometres away.

Indeed, many applications have become possible due to the fact that the absorption coefficient of plasmon resonance absorption has been shown to be several orders of magnitude larger than strongly absorbing dyes [10]. Most relevant for this thesis is the fact that chemically bonded molecules on the surface alter the electron density of the surface, which in turn shifts the plasmon resonance wavelength. Important to note here is that whilst surface plasmon resonance absorption can be modelled and calculated accurately for arbitrary geometries (with methods such as the Discrete Dipole Approximation), it is far from trivial to investigate a nanoparticle's exact geometry during reaction conditions. As will be further developed upon in section 2.2.3, the precise shape, size and structure of nanoparticles have profound impacts on the nature of LSPR behaviour.

2.2.3 Nanoparticle Size, Shape and Structure Impact on LSPR

Historically, the most commonly used plasmonic materials have been either Au or Ag, likely as a result of their chemical inactivity. However, there have been succesful implementations of LSPR within nanoparticles of Cu, Pt, Pd, Ni, Al, alloys of noble metals and even some non-metallatic metals. Of course, not only do nanoparticles of each material differ amongst eachother, but there are even more important differences between nanoparticles of each material. Ergo, here we will take a moment to discuss how the size, shape and structure of a nanoparticle may affect the results of LSPR and what we might expect from our choice in this study to focus on catalytic Pd nanoparticles.

In general, the greater the size of a given nanoparticle; the worse is the assumption that the incoming electric field is constant across its entire surface. Therefore, nanoparticles of larger size exhibit a larger phase shift of the induced dipolar electric

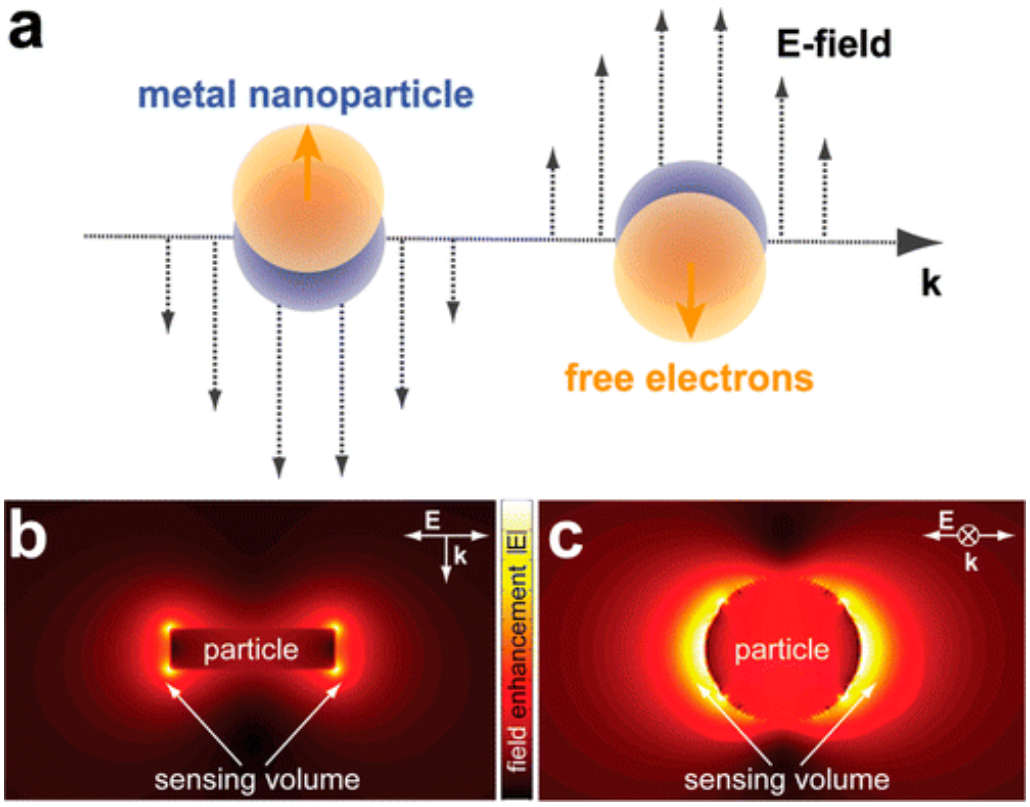

Figure 2.4: Top: The polarization of the surface of metal nanoparticles can start to oscillate upon interaction with near-visible light, leading to the phenomenon of LSPR. **Bottom:** Side and top view of plasmonic gold nanoparticle illuminated at the LSPR wavelength. The locally enhanced dipolar field is illustrated schematically by the intensity of light, as calculated by the FDTD method. This enhanced field region effectively acts as a nanosized sensing volume, where local changes in poralisation can be detected. Reprinted with permission from *Nanoplasmonic In Situ Spectroscopy for Catalysis Applications* [13]. Copyright 2012 American Chemical Society.

field relative to the incoming electric field, which in turn leads to a larger phase shift of its plasmonic resonance. In the same vein, larger particles are expected to produce broader resonance peaks as a result of more radiative damping and therefore shorter lifetimes. This is especially true for materials with allowed interband transitions in the visible spectra, which is the case for Pd nanoparticles. [2]

Additionally, important for this experiment is that palladium tends to oxidize more quickly and its oxides are more stables than on many other metal nanoparticles. The reactivity of CO oxidation on Pd nanoparticles therefore tends to increase dramatically in regions of high oxygen. Even on the crystalline surfaces of Pd that do not form oxides, thin films of adsorbed oxygen tend to cover these surfaces entirely, thereby still increasing reactivity of CO oxidation. [18]

Finally, we note that the crystalline morphology of the catalytic surface affect the chemisorption of molecules, which in turn affects the coverage of molecules on the surface over time and in turn the total reaction rate. The arguments given above are general and based on approximately spherical nanoparticles - for more complex analysis, numerical solutions to Maxwell's equations are needed. An example from Langhammer's team's work is the use of Finite Difference Time Domain (FDTD) simulations to simulate the optical response of plasmonic nanoparticles. [2]

2.2.4 Indirect Nanoplasmonic In Situ Spectroscopy for Catalysis Applications (INPS)

Applying the LSPR method described above, it is possible to measure the excitations within a single or an array of Au nanoparticles with near-visible light to study the physical changes within catalysts in the vicinity. The Indirect Nanoplasmonic in situ Spectroscopy (INPS) method used in this experiment utilized a dielectric spacer of silicon of a few nanometres in diameter, placed between the metal catalyst nanoparticle and the plasmonic gold nanoparticle. The reasons for this are twofold; firstly, the layer creates a barrier between the two materials such that they cannot directly interact (by i.e. alloy formation). Secondly, protecting the gold sensors from the surrounding environment and from reshaping as a result of high temperatures. Additionally, it is possible to tailor the surface chemistry or local electric field amplification by choosing a suitable material for the dielectric spacer with a correspondingly suitable refractive index. This discussion is, however, not vital for the experiment conducted for this paper. More relevant at this stage of this theoretical discussion is that of the physical signal readout of INPS, which forms the basis of all analytical methods (i.e. neural network architectures) utilized in this thesis.

As mentioned in Section 2.2.2, physical changes on the surfaces of a catalyst nanoparticle, in the vicinity of a plasmonic material, cause a shift in the spectral peak position of the resonance wavelength, illustrated in schematic 2.5 below.

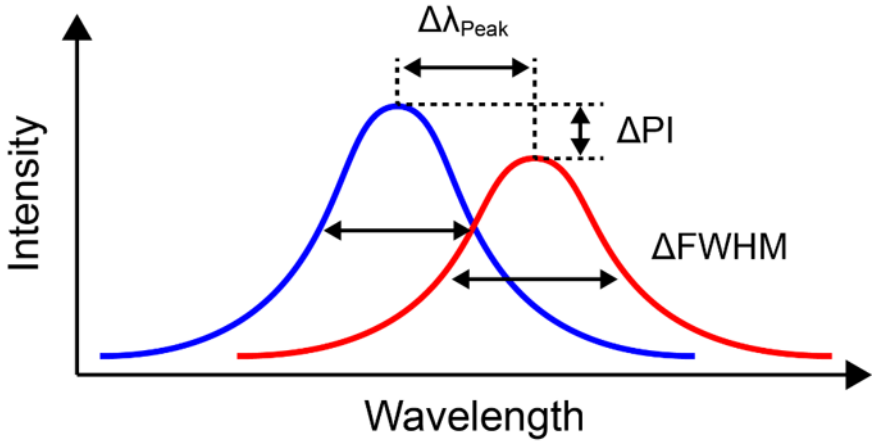

Figure 2.5: Idealised signal readout of a generic INPS experiment, in which physical surface changes on the sample material induce a change in the peak position of the resonance wavelength. Additional observables, such as the difference in particle intensity (ΔPI) and full-width at half-maximum (Δ FWHM) may provide additional information or provide a better signal-to-noise ratio.

As such, this shift in plasmonic resonance wavelength as a result of surface kinetic changes within a sensing volume of plasmonic sensing particles underlies the basis of the data analysis of this experiment. This peak shift can be extremely subtle, potentially only fractions of nanometres, but this sensitivity has been shown to be possible to achieve by Dahlin et. al through function fits of the data [6]. However, to fully realise this potential for each individual nanoparticle, we require a method of spectroscopy that is sensitive to such extraordinarily subtle changes in the optical spectra and robust to optical noise that might "spill" into the viewing field of the microscope.

2.2.5 Dark Field Scattering Spectroscopy & Single Particle INPS Spectroscopy

In this subsection, we discuss the specific method of optical spectroscopy used in the study; dark field scattering spectroscopy. Dark field spectroscopy, in contrast to more commonly applied bright field spectroscopy, is dependant on an experimental setup such that all incident light at the objective lens originates (as scattered light) from the relevant sample. As a direct consequence, the surrounding view around the sample is completely black, leading to the method's namesake. [2]

There are drawbacks to this method, in particular if the size of the samples are comparable to the wavelength of its emitted or scattered light. However, since we are interested in nanoparticles of size less than 100 nm, we are able to resolve each nanoparticle as single point-sources of light in an otherwise completely dark background. By then focusing this light into a spectrograph outfitted with a grating or prism of some kind, the incident light on the objective lens can be translated to the intensity of scattered light from the nanoparticles as a function of its wavelength.

Finally, we end this section of experimental theory by combining all theoretical frameworks above to converge to the specific method used in this study to measure scattered light from the plasmonic particles. The method used in this study comprises of far-field scattering dark field spectroscopy, to probe the nanoparticles' changes in surface kinetics as a function of scattered light. By utilizing the phenomenon of LSPR, the otherwise faint scattered light is amplified enough through plasmonic resonance that each nanoparticle can individually be resolved as bright spots in the spectroscope's dark field. This optical information can then be cross-checked with QMS data, which contains collective information about all particles inside the reactor chip, to translate and correlate resonance peak shifts to kinetic phase transitions on the surface of nanoparticles. However, since it is not obvious how these may be correlated, we dedicate next section to the description of measurement signal readouts and relevant methods of data analysis.

2.3 Measurement Signal Readouts & Data Analytics

In this section, we describe the nature of the relevant signals resultant from the experiment, as this is of vital importance to understanding which methods of data analysis are reasonable or even possible, and understanding the results in section 4. Following the same structure as the rest of this chapter, we firstly describe the nature of the QMS signal and how it might be interpreted, then describe the optical spectra of scattered light and how the corresponding FWHM signal is calculated. Following this, we describe the logic behind correlating individual optical spectra to collective QMS data and finally describe the underlying mathematical frameworks upon which neural networks are based.

2.3.1 Quadrupole Mass Spectrometer Readout

In this subsection, we focus on the most important aspect of the study for the purpose of this Master's thesis; data analytics. Firstly, we introduce and explain the nature of measurement readouts from the QMS signal and optical spectra, without going too in-depth about the specific experiment setup. Secondly, we describe previously attempted analytical solutions to resolve kinetic phase transitions in the data, particularly with where they succeed and where they are insufficient. Finally, we introduce our novel machine learning approach to the problem, and the theoretical basis upon which its likelihood to succeed is founded.

We begin by describing the analytical nature of the input gas concentrations. Denote concentration coefficient $\alpha = \frac{P_{CO}}{P_{CO}+P_{O_2}}$, such that $\alpha = 1$ when the only gas incident to the surface of Pd nanoparticles is CO, and equivalently $\alpha = 0$ when the same is true for O_2. In the experiment upon which this study is based, α goes from 1 to 0 and back to 1 in accordance with the top image of figure 2.6, corresponding to surface reactions with reactions within a high CO-low O_2 environment, high O_2-low CO environment and finally high CO-low O_2 environment again. Note that

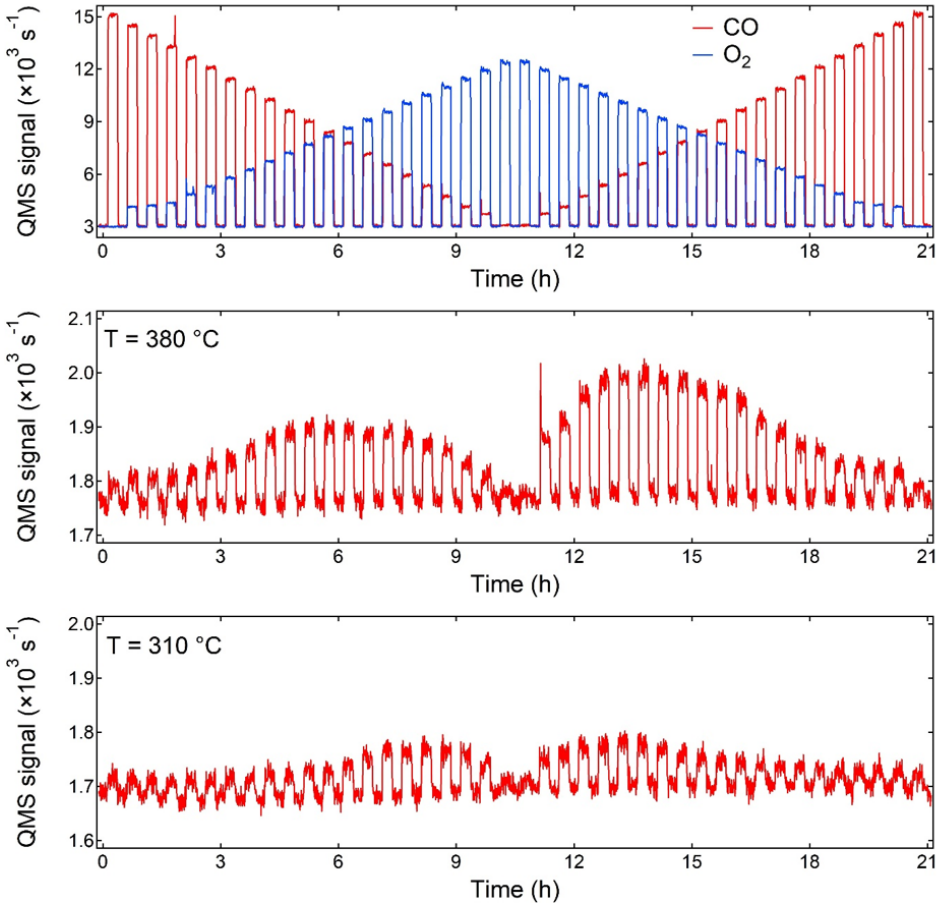

Figure 2.6: Top: QMS signal of CO and O_2, showing effectively the proportion of input reactant gases for any given time. Note that the proportion of gases follow the trend of $1 \geq \alpha \geq 0$ and then $0 \leq \alpha \leq 1$. In between each cycle, the nanochannels where the reactions take place are flushed by argon, corresponding to a QMS signal for CO and O_2 of 0. **Middle:** QMS signal of output CO_2 at a reaction temperature of $380^o\ C$. **Bottom:** QMS signal of output CO_2 at a reaction temperature of $310^o\ C$.

the microchannels where the catalytic reactions take place are purged with argon between each cycle of input gases, partially to reset the state of the palladium surfaces and partially to keep the pressure throughout the channels constant over time.

Turning our attention now to the middle image of figure 2.6, which shows an example of a QMS signal readout from an experiment made at $T = 380^o\ C$. Note that a higher QMS signal in a given time-period corresponds to a higher catalytic reaction rate in that time period. As expected, there is some mixing rate of input gases which correspond to a maximal reaction rate, but it is not necessarily when the gases are in identical proportion. Additionally, the reaction rate in the first region, as α goes from 1 to 0, is not necessarily equal to the rate in the second region as α goes from 0 to 1, owing to the aforementioned CO-poisoned environment and formation of surface oxides. Indeed, the spike in reactivity in the beginning of the second region is likely due to the formation of extremely reactive surface oxides. By comparing

these results with the bottom image in figure 2.6, we see that the formation of surface oxides are likely temperature-dependant, seemingly only able to form in significant volumes at high temperatures. Indeed, these results are consistent with the findings of R. Van Rijn et.al [17], where the formation of different types of surface oxides are shown to be highly temperature-dependant.

Harkening back to the mass-transfer-limited (MTL) regime discussed in section 2.1.1, where the reaction rate is limited by the rate of gas particle input, we clearly see that this regime has not affected the signals in 2.6. Firstly, in this MTL regime, the QMS of product CO_2 would be the same across temperatures, since increases in reaction rate at the catalytic surface would be irrelevant. Secondly, adding catalytic particles to the reaction should, in the MTL regime, not increase the amplitude of the QMS signal, but indeed it does, as can be tested easily by conducting the experiment with nanochips with differing amounts of nanoparticles.

Now, the QMS signals in figure 2.6 can be plotted in a more physically meaningful way; against the concentration gradient α, since α and time are directly correlated according to the top image in figure 2.6. The results are shown in figure 2.7 below.

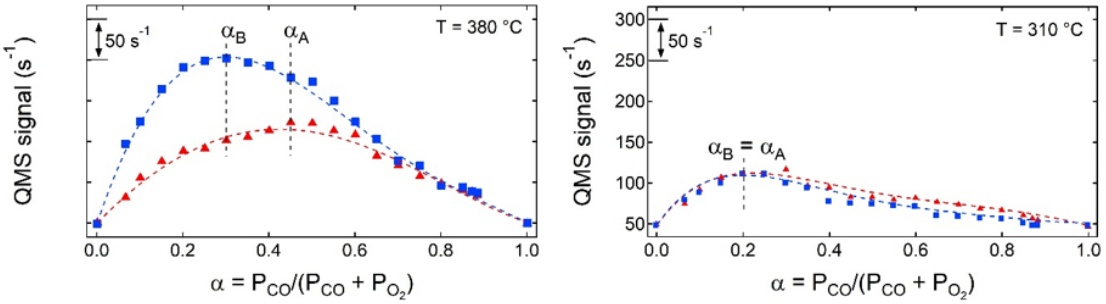

Figure 2.7: QMS signal for a measurement taken at T= 380^{o}C on the left and T= 310^{o}C on the right, plotted against concentration coefficient $\alpha = \frac{P_{CO}}{P_{CO}+P_{O_2}}$. The red and blue lines correspond to the first and second haves of the QMS signal, respectively. Note that the positions of α_A and α_B are more similary at low temperatures.

As described in section 2.1.1, there is a narrow mixing range of input gases which leads to an optimal reaction rate of the (combined) catalytic surfaces. This phenomenon corresponds to a physical kinetic phase transition on the surface of the bulk catalytic surface in the reactor, as this is where the (aggregate) surface transitions from a reduced to an oxidized surface or vice versa [2]. The mixing range where this occurs is represented in figure 2.7 by α_A and α_B, where α_A corresponds to the optimal reaction rate as α goes from 1 to 0, and similarly α_B corresponds to the optimal reaction rate as α goes from 0 to 1. Within certain temperatures and pressures, α_A and α_B may either overlap or occur in entirely difference mixing ranges. This can be understood quantitatively (albeit somewhat approximately) through figure 2.1, where the surface of palladium nanoparticles are clearly dependant on pressure and temperature.

2.3.2 Optical Readout

The optical spectra consist of 1024 wavelength measurements between $410-850$ nm, as shown in the first two images of figure 3.7. The light scattered from the palladium surfaces and amplified by plasmonic resonance results in a gaussian spectrum with a peak of approximately $410-850$ nm, depending on the exact kinetics of the surface. Since neural networks generally perform better with data with a mean of 0, we show in image 3 of figure 3.7 below the deviation from the mean spectra for each wavelength measurement by subtracting the mean intensity over time for each wavelength measurement. This is likely a much more relevant observable for the neural network to train on, since all relevant physical correlations resulting from utilizing nanoplasmonic spectroscopy is expected to be realised as shifts in peak position of the Gaussian spectra of scattered light, and we can clearly tell that the highest deviations from the mean behaviour of the gaussian spectrum occurs around the peak position.

Broadly speaking, past analytical efforts to translate optical data, both from this experiment and from other methods, have not entirely managed to resolve the complex problem of isolating kinetic phase transitions in individual palladium nanoparticles in a robust and consistent manner. Most commonly, FWHM is the most relevant observable for this type of experiment and analysis, owing to its stability and consistency even in the presence of considerable noise. Indeed, for a material like Platinum, the FWHM signal over time is so clearly correlated with the KPT of individual nanoparticles that hitherto unreleased results from the Langhammer team show that this type of analysis is certainly possible and desirable in many cases.

For a specific description of how the FWHM is extracted from the data, see section A.3.1.

2.3.3 Correlating Optical Spectra to QMS Data

At constant pressure, the peaks α_A and α_B at each temperature can be summarized through a phase diagram, as depicted in figure 2.8 below.

To reiterate, the goal of the data analysis is to find the value of α, that is the pressure ratio of input gases, which corresponds to a kinetic phase transition on the surface of each individual nanoparticle. We can find what is effectively the kinetic phase transition of the entire system, which is essentially an aggregate total reaction behaviour of the entire patch of nanoparticles within the microchannel, by finding the peak of the QMS data. However, of course, this tells us nothing about the behaviour of each individual nanoparticle - it is practically possible that none of the nanoparticles have a kinetic phase transition at the peak of the QMS signal, since it is effectively just an aggregate. Therefore, finding individual differences in the optical spectra for each individual nanoparticle as compared to the QMS signal is absolutely crucial in order to derive structure-function correlations between the optical data and QMS data.

Indeed, the ultimate goal is to develop graphs similar to 2.8, using only optical

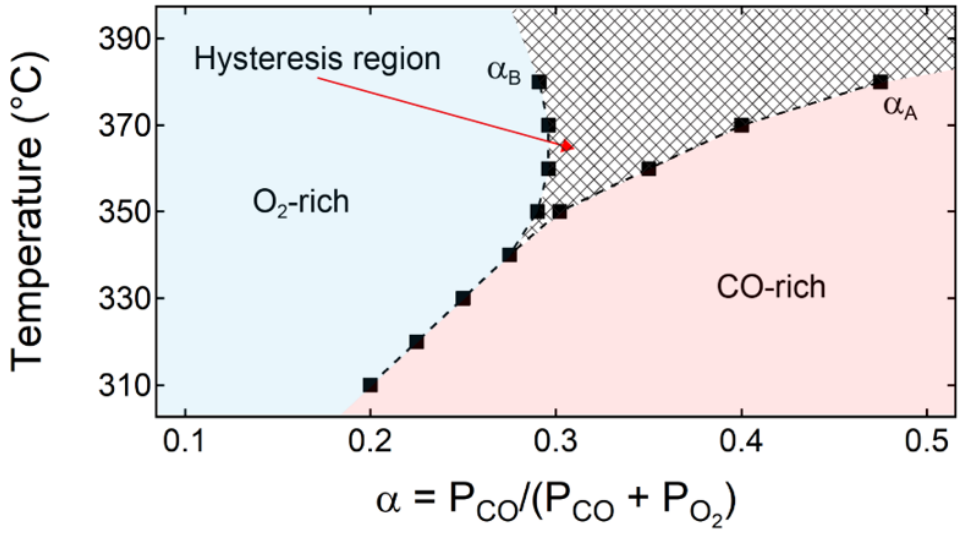

Figure 2.8: The concentration coefficient of QMS peaks α_A and α_B, plotted against the temperature of each measurement and concentration coefficient α forms a phase diagram. The hysteresis region in gray shows how CO- and O-poisoning might affect the results of our experiment as functions of temperature.

data, for each individual nanoparticle. This phase diagram can then be compared with figures similar to figure 2.1, though this is outside the scope of this thesis, to analyse how the surface of palladium nanoparticles shift over the duration of the experiment. This would, through investigation and analysis, ultimately let us see (through optical light) the surfaces of individual nanoparticles shift for different ratios of input gas α. To be abundantly clear, therefore, the development of an algorithm to translate optical to QMS data is effectively a middle-step to the actual goal, which is to develop structure-function correlations between optical data and surface kinetic changes on the individual nanoparticles.

In other works, as mentioned previously, correlating optical spectra of scattered light to the surface kinetics on individual nanoparticles is usually done by investigating the change over of FWHM of the optical signal [13], as illustrated in figure 2.5. Applying this method to this experiment reveals an interesting behaviour, as illustrated in figure 2.9 below.

Past analysis of the change in FWHM of the optical signal shows that there seems to be a correlation between the position of KPT and the formation or dissolution of bistable oscillations in the FWHM signal. By analysing the FWHM signal of a single nanoparticle, which is plotted in the bottom image of figure 2.9, it is clear that the QMS peaks seem to coincide with a point where the FWHM signal goes from very noisy to clearly oscillatory or vice versa. Note also that this is point is not entirely consistent with the position of α_B, where the peak is slightly shifted to the left of where the noisy period begins in the FWHM signal. The important question here is therefore whether this deviation is because the KPT for this particular nanoparticle happened to be slightly to the right of the aggregate signal (corresponding to a

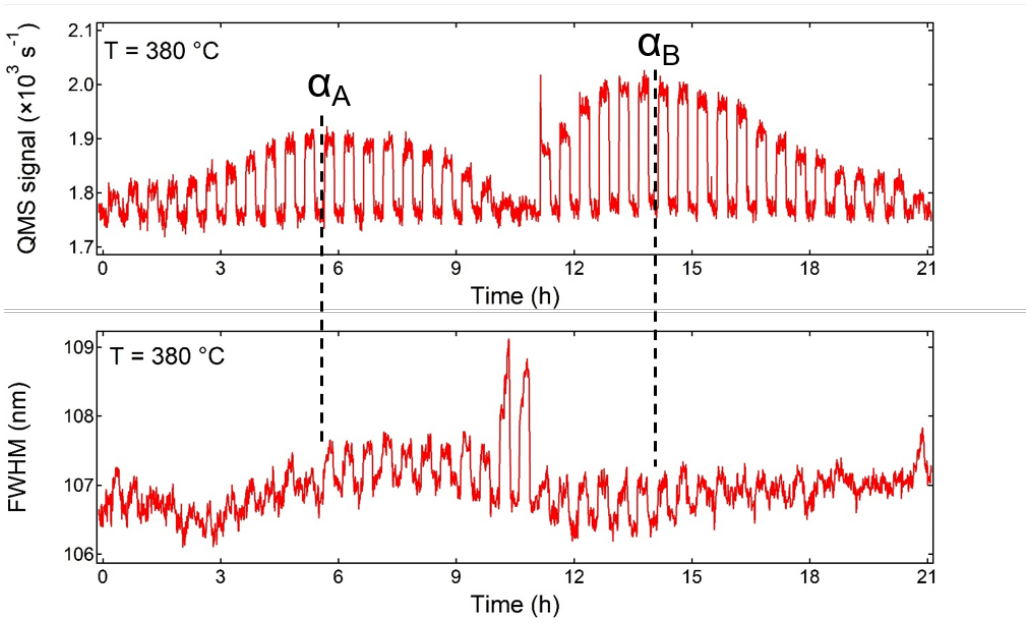

Figure 2.9: Top: QMS signal for a measurement taken at T= 380^{o}C. **Bottom:** FWHM signal for the same measurement, for a single individual nanoparticle. Note that the QMS peaks seem to roughly coincide with the formation or dissolution of oscillations in the FWHM signal.

KPT occuring at $\alpha > \alpha_B$), or whether we are imagining correlation which does not actually exist. From this data alone, it is really impossible to tell which of these is true. There are even more problems, however, as some FWHM signals do not seem to show any evident correlation at all, as is clear from figure 2.10.

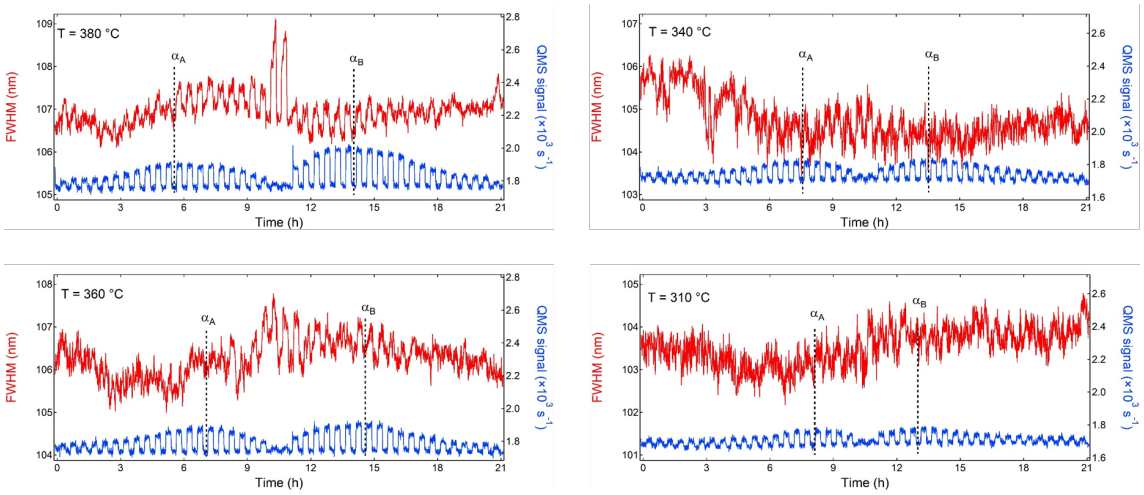

Figure 2.10: Four examples of the change in FWHM over time of individual nanoparticles, all from measurements taken at different temperatures. Note that the FWHM signal grows increasingly noisy with decreasing temperature.

As is clear from figure 2.10, many of the full-width-at-half-maximum spectra do not clearly follow the aforementioned trend of switching to a region of bistability

at their kinetic phase transition at all. This might be because the effect is less noticable at low temperatures, because higher pressures mean more noise, because some particles are fabricated differently and effectively nonresponsive to input gases, or a myriad of other reasons. This makes rigorous analysis based on the FWHM signal very complex, and concludes the full background on this thesis' involvement in the project. Of course, at this point, we could simply find this switching point only for the FWHM signals that are appropriately formed by eye, and this may indeed be significant on its own, but it is fruitful and desirable to be more ambitious than that.

Indeed, to this end, we develop a recurrent neural network to identify the positions of these switching points in a consistent manner, which even works for the entire temperature range relevant to this experiment. The results of this network are shown in section 4, and the in-depth description of its implementation can be found in section A.3. However, since this method of analysis is based on the assumption that this switching point coincides with the kinetic phase transition, which we only believe as a result of past results of data analysis, we propose here an alternate solution.

The main consideration of developing a new method of isolating the positions of individual nanoparticles' kinetic phase transition is to simply use more of the optical data by utilizing an algorithm that can analyse the problem using the entire spectra, rather than just the FWHM over time. Against the backdrop of the intricacies and difficulties of past methods of data analysis, we now briefly present again this thesis' novel approach to solving the problem; using a deep neural network to calculate each individual nanoparticle's theoretical QMS signal based only on its optical spectra over time, as found by the experimental method described in section 3. Before we delve too deep into the intricacies of our developed algorithm, as we do in section 3.3, we briefly cover the theoretical foundations of neural networks and machine learning in general in section 2.4.

2.4 Neural Networks and Machine Learning

In this section, we introduce briefly the concept of machine learning and the vital role of neural networks in many applications within data analytics of complex physical problems. Following this, we briefly cover the mathematical frameworks of (deep) neural networks, in order to give greater context and understanding to the specific method described in 3.3.

The history of machine learning, and the field of artificial intelligence as a whole, has been characterized by a cyclical repetition of short-term hype, disappointment and stagnation. The first time, recognized now as the first *AI winter*, happened in the 1960s with symbolic AI. The goal of AI in these early days was to pre-program a set of facts and rules into a computer, to be used as a tool for scientists in their analytical research. However, the hype for these machine became so great that one prominent scientist, Marvin Minsky, remarked in 1970 that "*In from three to eight years we will have a machine with the general intelligence of an average human be-*

ing". Clearly, this unsubstantiated claim never came to fruition, and funding dried up as hype died down when the field's excessive predictions proved widely inaccurate. This whole phenomenon repeated in 1985 with a new take on symbolic AI, *expert systems*, but hype died down again a few years later as the AI systems proved difficult to scale, expensive to maintain and very limited in scope. Many argue that we are currently experiencing the third reincarnation of this cycle, as wide and seemingly improbable predictions dominate most news coverage of AI. Indeed, many of these outlandish claims will never come to fruition. However, there are many reasons to believe that this time is different; particularly in how unprecedented the success of modern machine learning methods has been in a few vitally important areas. These include, but are not limited to: self-driving vehicles, text-to-speech conversion, natural language processing and translation, digital assistants, sophisticated ad targeting, speech recognition, image classification, superhuman gameplay (e.g. in chess and Go), and many others. In many of the above fields, modern AI has sparked nothing short of revolutions or paradigm shifts in the past decade alone. [5]

Despite unprecedented success in such a wide variety of topics, machine learning works on a remarkably simple basic principle: search for useful realisations of input data within some predefined mathematical space of possibilities, given a continuous feedback signal. Indeed, while older algorithms are often based on providing input data and rules to apply on said data to produce meaningful output data, machine learning works by providing input *and* output data, whereupon the machine learning algorithm produces meaningful rules of transformation between these data. Continuing this trend of relative simplicity, deep learning is also deceptively elementary in its principle. The difference between *deep* machine learning and *shallow* machine learning is the number of layers of data representation within the network; a shallow machine learning network typically only transforms input data into one or two representation spaces (layers), whereas a deep network jointly transforms input data into several layers. This may seem like a modest difference, but in practice its consequences are nothing short of profound, as a single feedback loop jointly adjusts all aspects of feature learning in the network simultaneously in purpose of a single end goal. According to François Chollet, author of the popular machine learning framework Keras upon which the implementation of our data analysis is built, these are the two essential characteristics behind the success of deep learning: *... incremental layer-by-layer way in which increasingly complex representations are developed* and *... [that] intermediate incremental representations are learned jointly.* [5]

2.4.1 Mathematical Framework of Neural Networks

In their most fundamental form, neural networks perform a sequence of tensor operations on input data, and returns the result as corresponding output data. The tensor operations are usually in the forms of weights and biases, such that the output of a single layer of a neural network can be calculated as:

$$\text{Output} = F(\text{Weights} \cdot \text{Input} + \text{Bias}), \tag{2.4}$$

where function F is known as an activation function, used to transform the output into a desired form, often in the ranges of [-1,1] or [0,1]. These weights and biases in eq. 2.4 are the trainable parameters of the model, continuously changing whilst the network is being trained. By using the output of one layer as input to another layer, we effectively perform deep machine learning within what is known as a deep neural network. [5]

Crucially, the fashion in which the values for weights and biases change during training is absolutely integral to the performance of a network. Without covering this topic in too much detail, we know that it is generally true that the minimum of any continuously differentiable function can be found by the limit $\lim f'(x) \to 0$. Indeed, denote the number of trainable weights in the network by W, and it is thereby possible to find the optimal weights for a network without doing any training at all, simply by solving the equation $\nabla f(W) = 0$ where $\nabla f(W)$ is a polynomial of degree $N = \#W$. However, this calculation becomes practically infeasible for any real neural network, where the amount of weights usually number in the thousands or millions. Instead, real neural networks use more clever ways of updating the values of weights, such as the method of stochastic gradient descent (SGD) which is most relevant for this project. In performing SGD, the network trains on a small amount of randomly chosen input data at a time, calculates the difference between predicted and target values (the 'loss"), and computes the gradient of the loss with respect to the network's parameters, in between each set of new input data. Since the activation functions of the network are by convention easily differentiable, this method has relatively low computational cost. All network parameters are then updated in the opposite direction of the gradient, thus reducing the loss on each small set of data a little bit. Indeed, because of the invaluable feature that all activation functions are simple and easily differentiable, all parameter values can be individually updated with a single function call, since it only becomes a matter of applying the derivative chain rule. [5]

The crux of this mathematical framework is that it provides a general method of doing complex high-dimensional function fits between some defined input and output ("target") data. Consider, as an unrealistic yet descriptive illustration, a set of (input,output) data (x, y) perfectly correlated through a mathematical model $p(x, y)$. Since this perfect model is often impossible to find for real applications, our goal is to develop a network through the method presented above with parameters θ, such that the network's predictive model $f(x, \theta)$ converges to the behaviour of $p(x, y)$. Finding the parameters θ is therefore a question of minimizing the expected value of the difference between predicted results and target data, mathematically described by function 2.5

$$E[L(f(x, \theta), y)]. \tag{2.5}$$

Here, L is the relevant loss-function, the choice of which will depend on the specific problem at hand. Using neural networks, therefore, our problem of identifying the KPT of individual nanoparticles is effectively reduced to two steps.

Firstly, we want to train a neural network to find parameters θ such that the expected value of its loss function is minimized, for the implicit purpose of making the network's predictions converge towards hypothetically perfect correlations $p(opticaldata, QMSbehaviour)$. In other words, we train the network to calculate the qualitative QMS behaviour of a given nanoparticle, given only its optical spectra of scattered light in some form. Secondly, we want to use the network's learnt parameters on nanoparticles that it has never been trained on, calculate their qualitative QMS behaviour, then finally analyse to what extent the results are correct and usable. For specifics of this implementation, see section 3.3.

Note that we have focused mostly on machine learning in general, as well as the basic implementation of a (deep) neural network. All mentions of the "network" will be in reference to this particular network; the deep densely connected neural network which constitutes one of our two machine learning approaches. The other approach, the recurrent neural network, will conversely always be referred to as such. For details of its specific implementations and underlying theory, see section A.3.

3

Experimental Setup & Analytical Method

The experimental method upon which this thesis is based revolves around four major components. Firstly, the design of nanochannels within a nanofluidic chip which allows the controlled mass transfer of gas to and from singular nanoparticles, including the accurate control of temperature and input/output pressure. Secondly, a system which allows us to monitor the optical spectra of scattered light from individual nanoparticles. Thirdly, the ability to monitor the mass spectrum of product gas created as from catalytic reactions. The theoretical background for these components has hitherto already been discussed in section 2 above, and lay the foundation for the practical implementations described in this section. Below, in figure 3.1, a schematic of the experimental setup is shown. In this section, we briefly present the three major components of the experiment before shifting attention to the larger focus of this thesis: describing the analytical methods of handling the data.

3.1 Nanoreactor Setup

The overall setup can schematically be seen in 3.1 below, and consists of a nanofluidic chip, inlet and outlet for gas flow, a temperature controller, a mass spectrometer and a dark field spectroscopy setup connected to an EMCCD camera. CO and O_2 is pumped into the nanofluidic chip at some controlled pressure (usually 1-3 bar), such that the concentration of gases, $\alpha = \frac{P_{CO}}{P_{CO}+P_{O_2}}$, goes from $1 \to 0 \to 1$ over the course of a whole $\approx$ 24hr experiment. Each value of α is held constant for approximately 15 minutes, after which the system is purged with argon to reset the surface state of the palladium nanoparticles and to keep pressure constant throughout the experiment. The flow of input gas to the system therefore happen in pulses, where each pulse entails a new value of α for the system, as illustrated in figure 2.6. The reactions then take place on the surface of the catalytic nanoparticles, which have been combined with plasmonic gold nanoparticles to enable nanoplasmonic spectroscopy, as will be further elaborated upon in section 3.1.1. The scattered light over time from the surfaces of these nanoparticles is recorded as functions of wavelength in the dark field scattering spectroscopy setup, as described in section 3.2. Finally, the gas flows through a nanochannel into a UHV chamber connected to a QMS, where the concentration of reactant product CO_2 is measured.

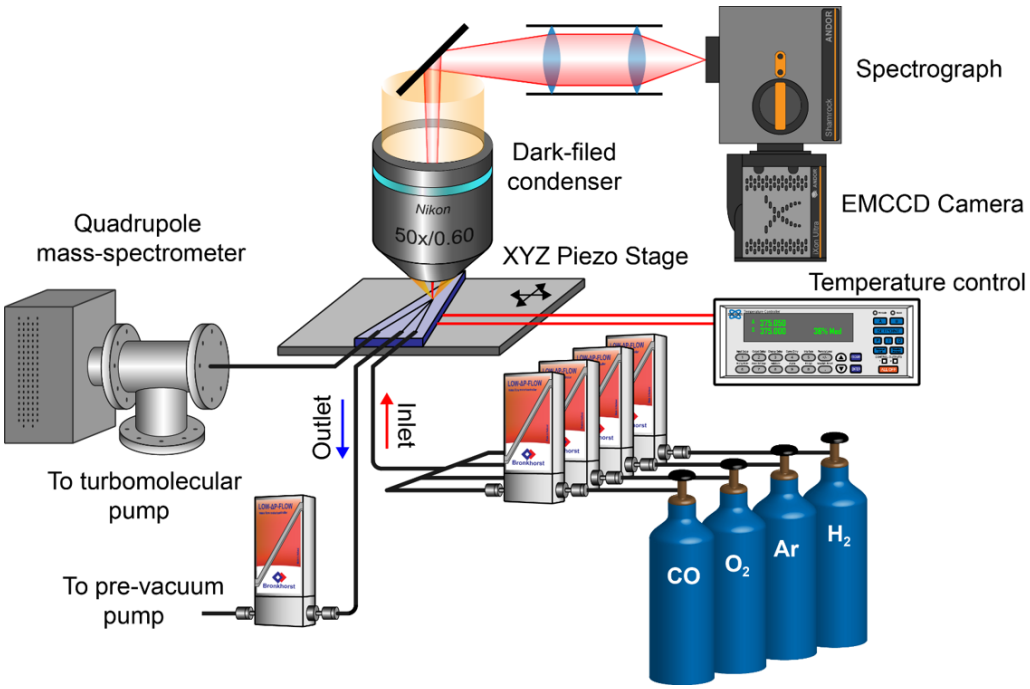

Figure 3.1: Schematic of the experimental setup. Gas reacts on a silicon chip at pressure between 1 and 5 bars and then passes into a quadrupole mass spectrometer. A spectrograph is used to measure scattered light during the reaction process, which is then averaged every 20 seconds to form a single (averaged) Gaussian spectrum. Unreactive argon is used to flush the active area between gas input phases, both to keep pressure constant at all times and to reset the particles' behaviour in between discrete values of α.

3.1.1 Nanofluidic Chip

The nanofluidic chip, shown in the top image of figure 3.2 and represented schematically in the center of the XYZ Piezo Stage in figure 3.1, is designed to allow controlled gas flow to each individual nanoparticle within the chip, whilst still allowing them to be optically accessible to the spectroscopic setup and allowing their reaction products to be fed into a mass spectrometer. The stage is motorized and controlled through a feedback loop to stabilize the sample. A schematic of the inside of the nanofluidic chip is shown in the bottom image of figure 3.2, where each orange dot represents a single nanoparticle. As hopefully evident in the bottom image of figure 3.2, each chip is lined with a patch of a large number of particles between two rows of ten isolated particles. The purpose of the patch of particles, which varies in size between chips, is to have enough reaction product to achieve a high signal-to-noise ratio in QMS data. The dark field scattering spectroscopy is then focused on this row of particles, with the camera set up such that the scattered light from each individual nanoparticle is recorded separately (albeit potentially with some small leakage of light, as will be discussed shortly). Finally, when the gas has passed through this microchannel, the gas flows into an Ultra-High Vacuum (UHV) chamber outfitted with a high-end QMS. The ultimate goal that this setup achieves is to measure the state of individual nanoparticles during catalytic reaction conditions, which are individual properties of each nanoparticle, and correlate this with mass spectrometry data which shows reaction product formation, which is a global property of the system, to derive structure-function correlations.

Additionally, note that the chips may be manufactured differently, resulting in slightly different physical properties. For instance, a higher particle count in the patch may increase the intensity of the QMS signal, but also increase light leakage into the scattered light of otherwise isolated nanoparticles. The dimensions of the microchannels may also vary, with wider and shorter microchannels producing higher flow rate and thereby number of reactant molecules on the surface of the catalytic nanoparticles per time. Another very importact consequence of this is that the specific properties of each single nanoparticle are very hard to control, and often essentially unknown. The desire is to manufacture each nanoparticle as a symmetric and homogenous disk, but they will all inevitably differ in unknown ways after the manufacturing process. We therefore expect all nanoparticles to be on a spectrum of efficiency as catalysts - nanoparticles within the chip are likely to range from highly efficient to effectively useless.

Finally, note that the nanofluidic chip is equipped with a PID-controlled heater, allowing a consistent working temperature range of 20-375^{o}C with a quick response time of <10s.

3.2 Optical Setup

As discussed previously, the vital key to a functioning optical setup is the successful translation of surface kinetic phase transitions as a function of surface catalytic reactions into optical signals, and subsequent amplification of these signals such

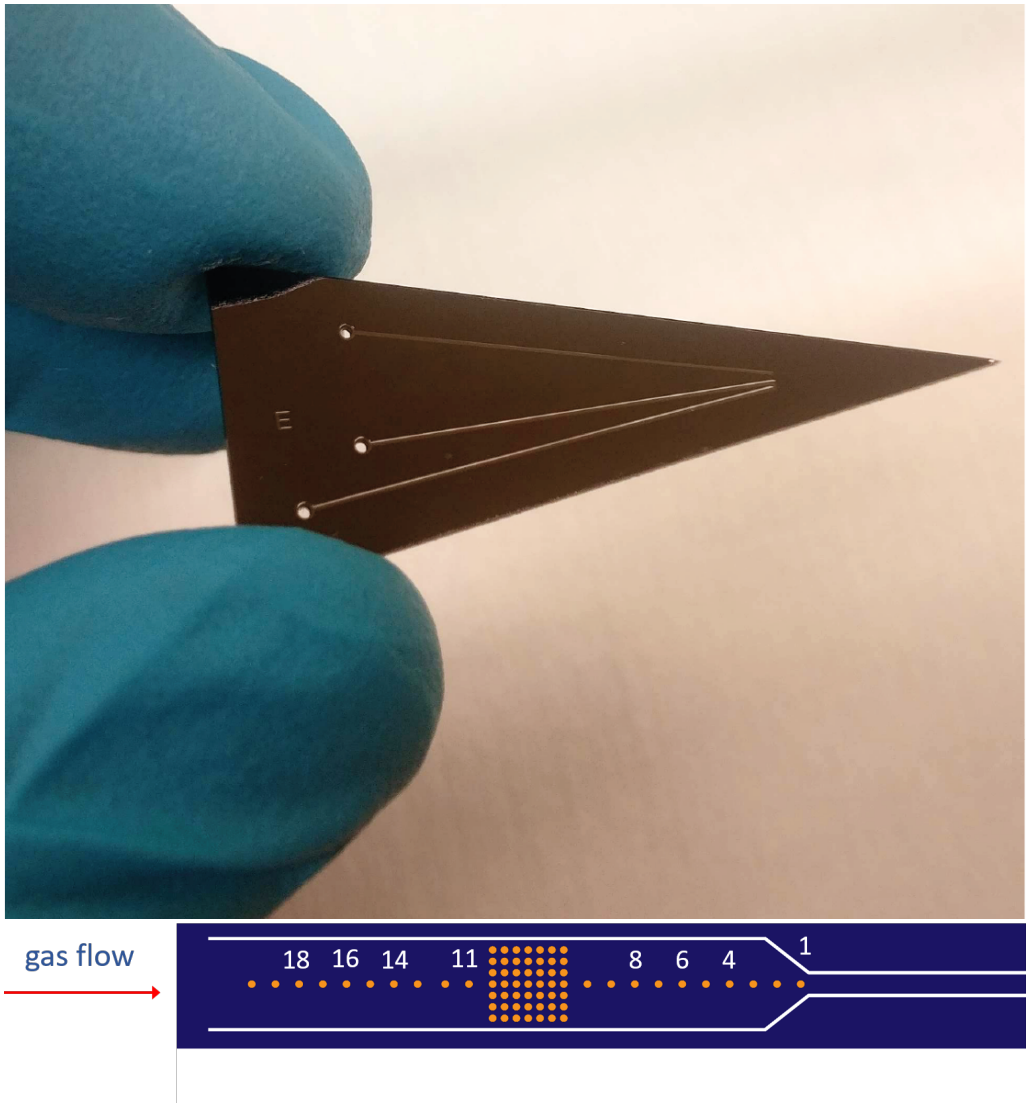

Figure 3.2: Top: Photograph of the nanofluidic chip. The microchannels leading towards the chip's edge are clearly visible, and where they converge are where the actual reactions occur on the nanoparticles. **Bottom:** The portion of the silicon chip where all microchannels converge, as seen in the top image. Gas flows into a microchannel, containing ten isolated nanoparticles to the left and 9 isolated nanoparticles to the right of a grid of nanoparticles (with varying amounts of nanoparticles). The optical spectrum is measured separately for each of the 18 particles, and the grid of particles is used to get a good QMS signal.

that we may measure and interpret them. In practice, this is achieved by combining the catalytic particle with a plasmonic one, since plasmonic resonance is sensitive to surface kinetics, and utilizing dark field scattering spectroscopy to isolate individual nanoparticles.

In figure 3.3, an example of the optical view through dark field scattering microscopy is shown. The micro-channel through which the gas flows and whereupon the nanoparticles lie leads into a much narrower nanochannel, in order to reduce the pressure of the gas before it enters the mass spectrometer (which is vital for its intended function). In the top right of the image, the combination of a catalytic palladium nanoparticle on top of a plasmonic gold nanoparticle is shown.

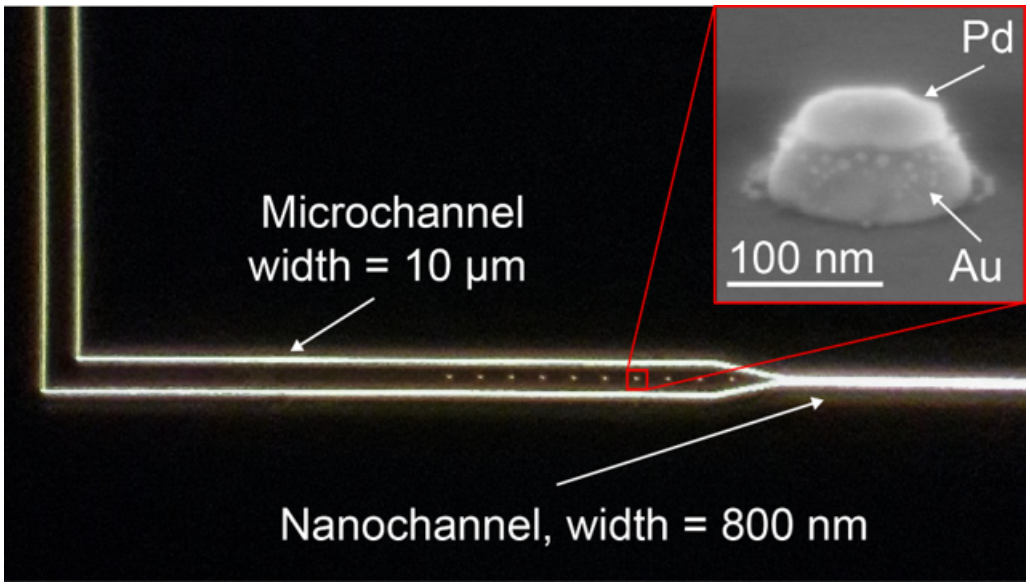

Figure 3.3: Zoom-in of the channels within the nanofluidic chip, showing how the nanoparticles lie within the relatively large microchannel, which leads into a much narrower nanochannel before the resulting gas is analyzed through a mass spectrometer. Note that the nanoparticle patch is not featured in this particular image.

During experiment, the scattered light from the walls of the micro/nanochannels are effectively removed by using an aperture that blocks incident light perpendicular to the microscope.

In post-processing the data, we normalize the intensity of the spectrum with regards to the intensity of emitted light from the initial light source (i.e. a lamp) and subtract the intensity of the background, which can be measured by measuring the intensity of a dark spot near the nanoparticle in question. Note, also, that the spectrograph used in this study measures intensity once per second and calculates the average spectrum over 20 seconds. The output of the optical setup is therefore an amount of averaged gaussian spectrums over several hours of measurement time, depending on the total measurement time and averaging frequency of the spectrograph.

3.3 Neural Network & Data Analysis

The approach of using a deep densely connected neural network can be split into two main steps. Firstly, we want to develop a deep neural network architecture which can translate optical data into QMS data. Secondly, we want to use this trained network on data it has never seen before, whether it is from different chips, different experiments or different particles, and analyze to what extent its predictions of the position of individual nanoparticles' kinetic phase transitions are accurate. The second step also has a dual purpose, as the quality of predictions in different circumstances may well reveal much about the underlying physical processes within the experiment, as will be discussed in section 4.

To achieve the first step, as alluded to in section 2.3.3, a deep neural network is trained on recreating the QMS signal of the entire experiment, given all optical data for several individual nanoparticles. Though this is effectively recreating a global property of the system given local properties of each nanoparticle, the aspiration is that the network, with enough correlatory data and proper preprocessing, will learn qualitatively which optical behaviours tend to correspond to which QMS activity. To reiterate, the goal for the network is not to calculate the activity of each individual particle, but rather to qualitatively translate optical spectra to QMS data, such that the position of its peak can be measured. That it may be impossible for the network to elucidate every single trend in data is therefore irrelevant, if only the qualitative behaviour of QMS data can be predicted with enough certainty to isolate the positions of kinetic phase transitions from optical data alone.

I will repeat here the argument given in section 2.3.3 for the sake of clarity, as to why we expect it to be possible to recreate theoretical QMS data from optical data alone. The basis for this aspiration is that we know the change in FWHM and QMS are correlated temporally through surface plasmon resonance wavelength shifts. Therefore, since the FWHM signal is of course derivative of the entire optical spectra, training a network on the optical spectra should allow it to find both the correlations between FWHM and QMS that we know exists, as well as hidden correlations that we do not yet understand. A deeper discussion of these hidden correlations can be found in section 4.5, and further discussion on this will follow towards the end of this section.

As mentioned in section 3.1.1, we expect some nanoparticles to be effectively nonresponsive to input gases and therefore irrelevant for the purposes of our data analysis. Therefore, a robust method is needed to discover which particles are appropriate to analyze before any proper analysis of the optical spectra can take place. In the future, we hope to use a method similar to what was used by Langhammer in [13], where the responsiveness of individual nanoparticles to hydrogen gas was qualitatively measured, and served as a basis for deciding which nanoparticles were manufactured to an acceptable standard. An example of this responsiveness can be found in figure 3.4 below, yet this responsiveness to hydrogen gas does not seem to translate to responsiveness to CO oxidation, and we can therefore not use these results for the purposes of this thesis. Since no equivalently convincing theoretical

basis exists for judging a nanoparticle's catalytic effectiveness to CO and O_2 gas, we propose here a different solution.

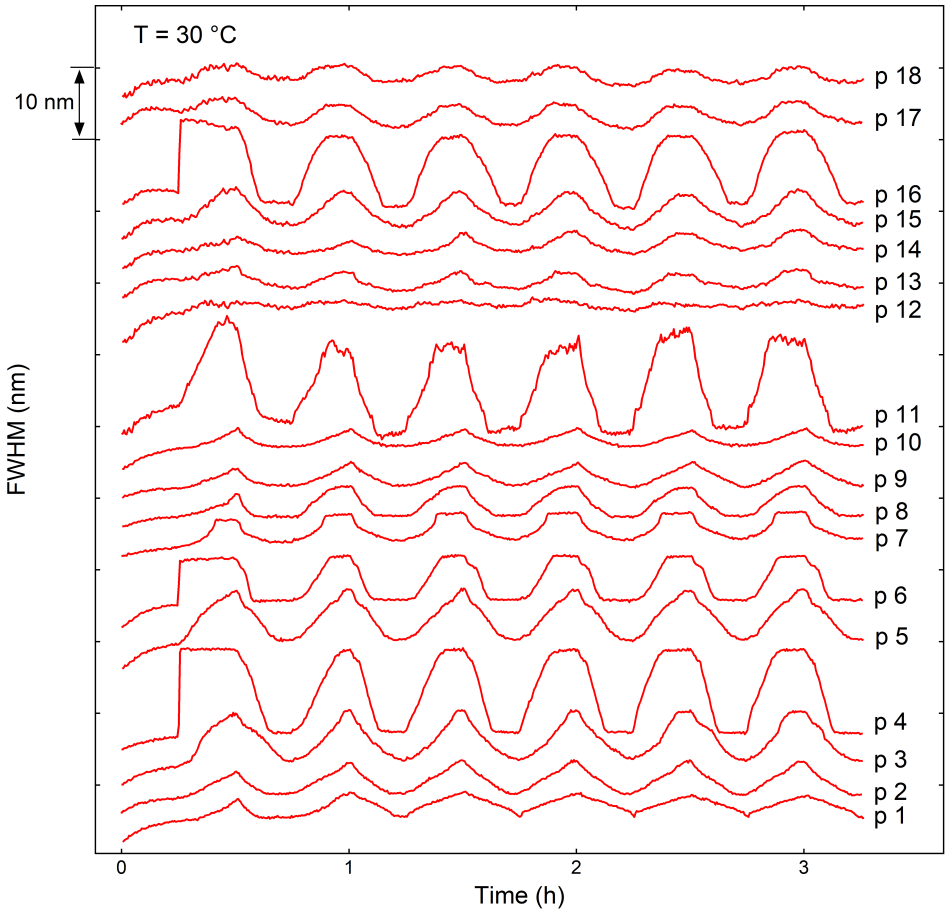

Figure 3.4: Change in FWHM for 18 particles within the nanocfluidic chip during the experiment conducted with hydrogen gas rather than CO oxidation. As is clearly evident, the responsiveness of particles varies widely, and some particles seem to be of considerably higher quality for the purposes of reacting with hydrogen gas. A similar method for analyzing the quality of nanoparticles' efficacy during CO oxidation would be desirable, but does not yet exist.

It is clear from images such as the middle one in figure 2.6, that the quantitative and qualitative difference in QMS signal between regions where the input is argon and where it is catalytically reactant gas is very high. However, low-quality nanoparticles, which in this context concerns nanoparticles which react poorly with input gas, are unlikely to reflect this difference. Instead, it is likely that nonresponsive particles have similar optical behaviour in both these regions. The corresponding hypothesis is therefore that the trained network will fail to distinguish these two regions whilst predicting on bad particles, resulting in an adequately robust and unbiased method of nanoparticle elimination for the purposes of this thesis. The solution to find bad particles is therefore clear; use the same neural network as already described in this section to recreate the entire QMS signal, including regions where the input gas

is argon, and remove the particles which show a small standard deviation in QMS signal between regions which correspond to argon gas and regions which correspond to CO and O_2 gas. An example of a bad and a good particle, in this respect, is shown in figure 3.5. For an explanation of the appearance of the QMS signal, see section 3.3.1 below.

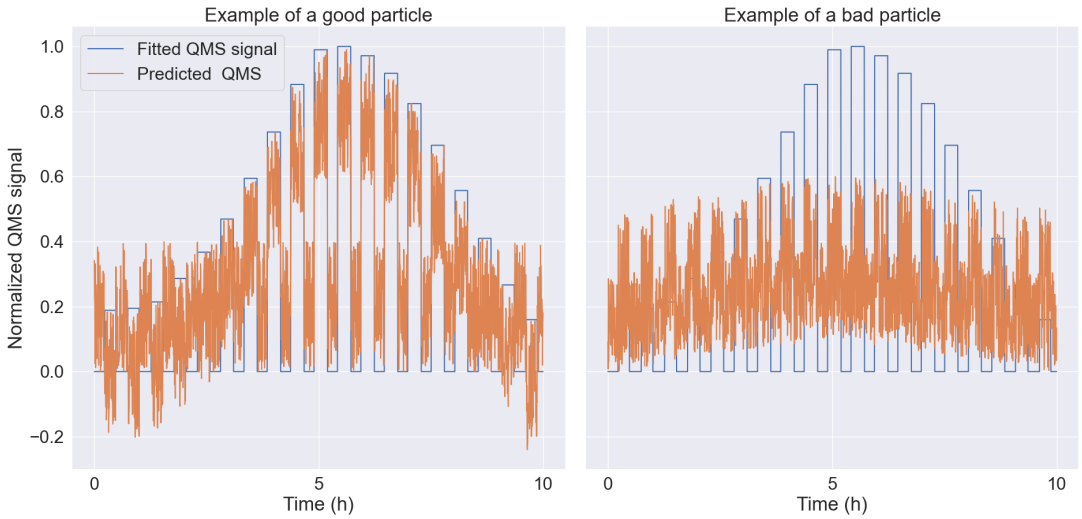

Figure 3.5: Left: Example prediction of a nanoparticle we consider good, in the sense that its predicted QMS signal (based on its optical spectra) shows a clear difference between regions of CO input and Ar input. **Right:** Example prediction of a nanoparticle we consider bad, in the sense that its predicted QMS signal does not show a clear difference between regions of CO input and Ar input. Note that these images are schematics, and not resultant or indicative of actual results.

Note that the most responsive particles found by this method were 1,2,3,4,5,6 and 8. Common for all these particles is that they are all found on the side of the nanochip that is behind the large patch of nanoparticles. The most likely reason for this is that the optical setup is calibrated such that they track particles in a straight line starting from the position of the first particle. Therefore, if the nanochip is slightly titled, which is certainly likely as it is placed by hand, later particles might be increasingly out of focus of the camera, therefore decreasing the quality of the optical spectra.

The network itself uses a method of cross-validation known as Leave-P-Out cross validation, specifically Leave-1-Out, where each network is trained on all particles except one. This constitutes the first step of the two-step deep neural network approach, where the network is trained on recreating the corresponding QMS signal through each nanoparticle's individual optical spectrum, thereby hopefully finding correlations between certain behaviours in the optical spectra and corresponding behaviours in QMS activity. When the training is complete, the model is tested on the optical spectra from the singular particle which was not trained on, and the results of these tests are shown in chapter 4. This process is reset and repeated until

all particles' KPT have been predicted.

Finally, after the process of training, we come to the second step of this approach; analyzing the accuracy of the network's predictions. Generally speaking, this type of problem is usually solved by either comparing predicted results with theory (i.e. a hypothesis), results from another experimental method, or by testing the algorithm on new samples of particles which act on the same basic fundamental principles but are de facto very different, such that we test whether the correlations learnt by the network are generalizable. In the case of this experiment, we rule out the first possibility since there is currently no such available theory. The second possibility is severely limited, since only moderately similar experimental results exist to compare with, but further discussion of this can be found in section 4.5. Indeed, this is why we turn mainly to the third option, where training the network on the function of particles within a single sample (nano-fluidic chip) and testing it on another will elucidate whether the learnt correlations of the neural network are generalizable, and therefore practically useable.

There are, however, several ways of checking that the network's results are likely to be incorrect. For instance, by noting that the QMS data is effectively an average behaviour of a large number of particles, we should expect the average prediction of several particles to consistently converge to the correct peak and shape of the QMS signal. Indeed, we should expect the average of several predictions to converge to the same values and temperature-dependent behaviour as the phase diagram in figure 2.8, letting us recreate the diagram through averaging the predicted QMS signals of several particles. There are other ways of analyzing the results of the network within a single sample, as will further be discussed in section 4.

Additionally, the performance of neural networks and indeed machine learning algorithms in general depend on the amount of data available to train on [5], wherein using several experiments simultaneously helps this goal. Therefore, we repeat each experiment four times and train the network on the results of all of these measurements, to give the network the possibility of learning what correlations are only particle-based rather than experiment-based.

3.3.1 Data Preprocessing

As mentioned in section 2.4, proper preprocessing of the data is key to a successful algorithm. In general, all preprocessing done on the data revolves around reducing the prominence of noise and measurement outliers, as well as amplifying the prominence of properties of data which are of interest whilst repressing the prominence of data which is not. This section revolves around the preprocessing which was done for the deep densely connected neural network, and the process of preprocessing for the recurrent network can be found in section A.3.

In practice, QMS preprocessing concerns four steps. Firstly, the whole signal is synced against the optical spectrum, since the acquisiton time of QMS and optical data is different, such that optical spectra and QMS signal are correctly correlated in time. Secondly, the signal is shifted vertically such that the entire QMS signal

has the same baseline. Thirdly, the whole signal is normalized between 0 and 1 and its temporal mean is subtracted from the whole signal. Finally, a step-function fit is made to the data, letting the value of each step be equal to the mean value of each corresponding region. This step function fit is crucial, as we are not interested in stochastic differences in neither QMS or optical data within a single value of α, but rather only interested in how the mean behaviour of each signal matters as functions of α. All of these preprocessing steps are done separately for the CO-rich and O_2 rich areas, and are visualized in figure 3.6 below.

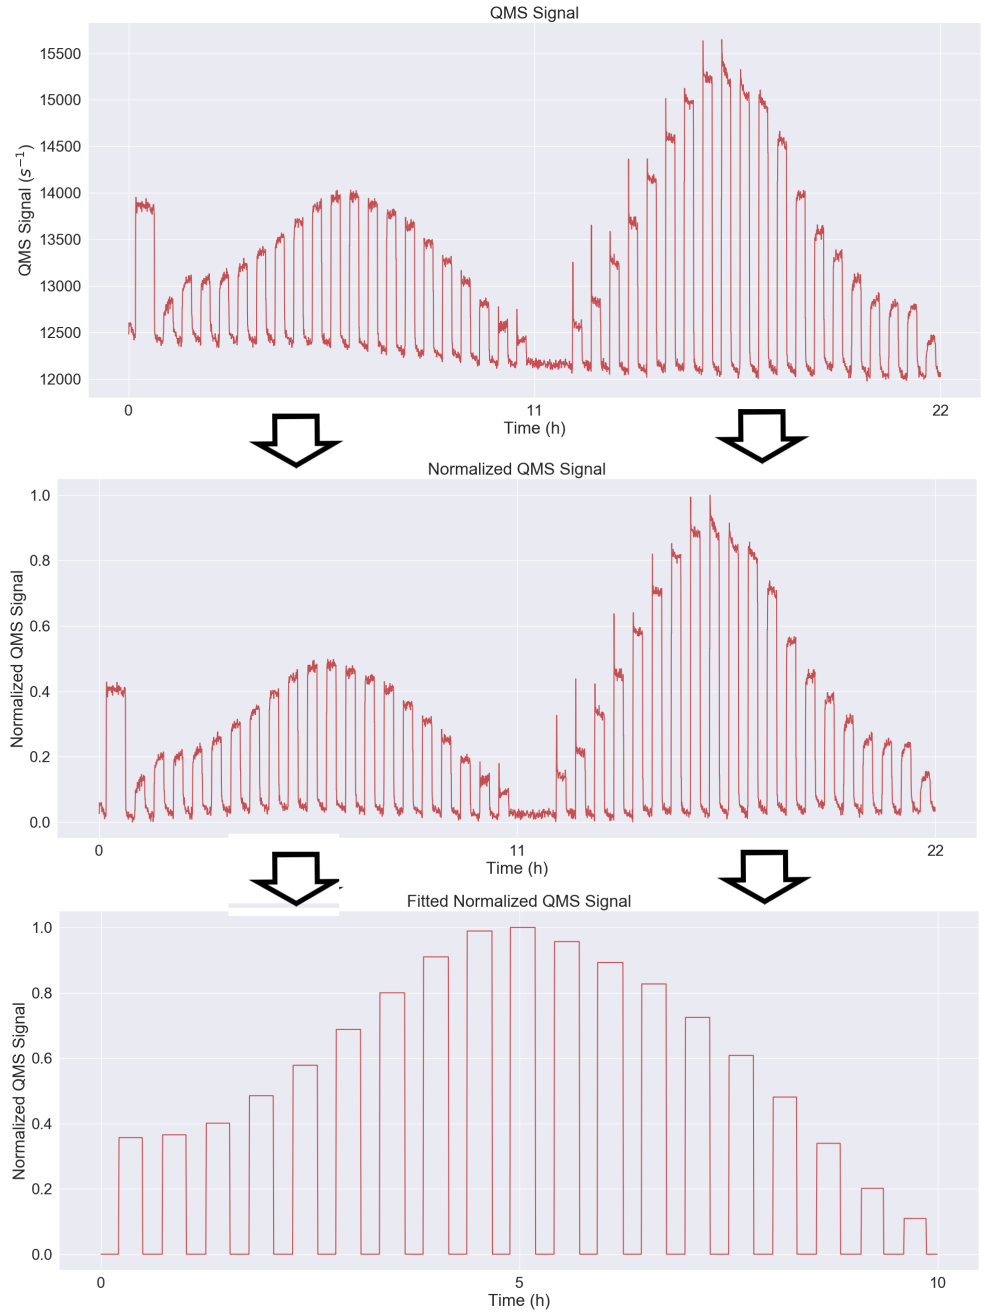

Figure 3.6: Preprocessing of the QMS signal. The signal is normalized and stabilized to a unified base, then sliced in twain in areas of CO-rich and O_2-rich behaviour. Finally, a step function fit is made to this data to eliminate stochastic differences within each singular value of α.

The preprocessing of optical data differs in a few key ways, and happens in four steps. Firstly, the intensity of the spectrum is normalized with regards to the intensity of emitted light from the initial light source and the intensity of the background, which can be measured by measuring the intensity of a dark spot near the nanoparticle in question. Secondly, measurement outliers, which are likely the result of camera

software errors, are removed from the spectra by setting them equal to the temporal mean of that particular wavelength measurement (temporal shifts over most wavelengths are very small). Thirdly, the temporal mean of all optical spectra is subtracted from the spectra, since neural networks learn much better when input and output data are represented as deviations from the mean. Finally, a step function fit similar to what was done when preprocessing QMS data is applied, where the value of intensity for each value of α is set to the temporal mean of the corresponding region in time of the corresponding value of α. In effect, this means each value of α corresponds to exactly one temporal shift per wavelength measurement, allowing the network to learn which sets of intensity measurement across 1024 wavelength measurements correspond to which behaviours of the QMS signal. The network should then, with enough training, be given new optical spectra and calculate (or predict, in ML terminology) its expected deviation in QMS behaviour from the mean (which in this case is our ground truth; the measured QMS signal). Indeed, viewed in this way, we expect that certain behaviours of the spectra in the bottom image of figure 3.7 correlate to certain levels of QMS activity.

Since describing the specifics of the deep neural network's implementation would require many more sections of theory in a thesis that is already quite long in length, we have omitted this discussion from the main paper. However, for the reader somewhat knowledgeable of the specifics of neural network implementation, we have included a brief description of the hyperparameters and training method of this network in the appendix, section A.1.

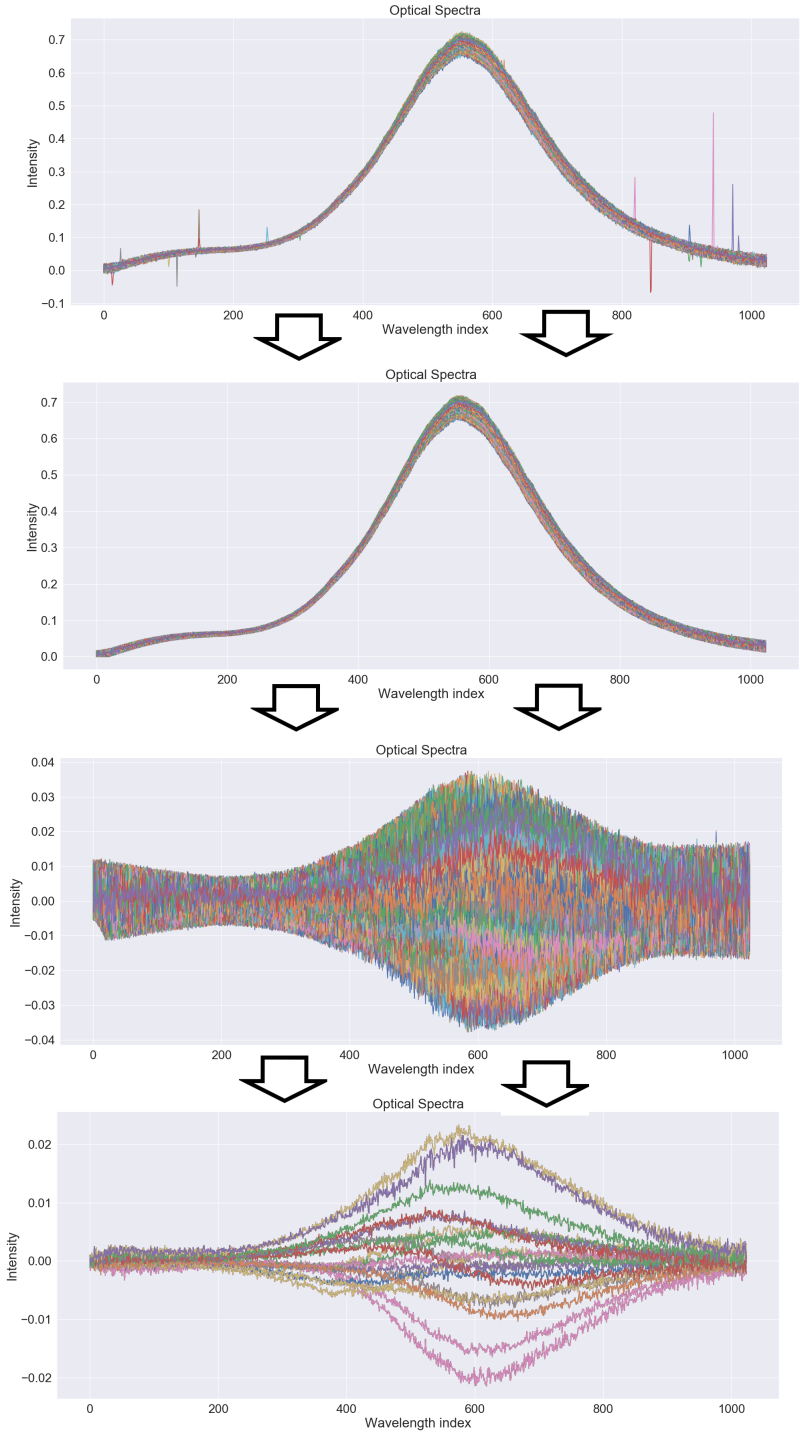

Figure 3.7: Preprocessing of the optical signal. The optical spectra is cleaned and normalized with regards to the intensity of emitted light, then its temporal mean is subtracted. A step function fit is made to the data, similar to what was described in the QMS preprocessing part, to eliminate stochastic differences between each singular value of α.

4

Prospects for Kinetic Phase Transition Detection on Individual Nanoparticles

In this chapter, we present the results of training the neural network to qualitatively translate optical spectra to QMS data, and discuss the accuracy and reliability of the results. To make this reasonably digestible, the results are effectively split in five parts. Firstly, example predictions are presented and used to provide an intuitive basis of understanding that will help us comprehend and analyze the final results. Secondly, we present the resulting phase diagrams of KPT as functions of temperature and α for both QMS and optical data, presented as one potential measure of the accuracy of our results. Thirdly, we present the results for each individual nanoparticle for a chosen set of experiments, around which the majority of analytical discussion will be based. Fourthly, we present a method of cross-validating our results by testing the neural network across different nanoparticle samples. Finally, we end this section with a discussion of the limitations, uncertainties and sources of error with the experiment and analysis.

4.1 Nature of Deep Dense Network Prediction Results

To reiterate the argument given in section 3.3, we expect the average prediction of all particles to converge to the shape and (most importantly) peak of the QMS signal of the corresponding experiment, since the QMS signal is effectively the result of the average behaviour of a large amount of particles. In figure 4.1, the top image shows an example prediction of a single particle across four experiments, and the bottom image the average prediction of seven particles across four experiments.

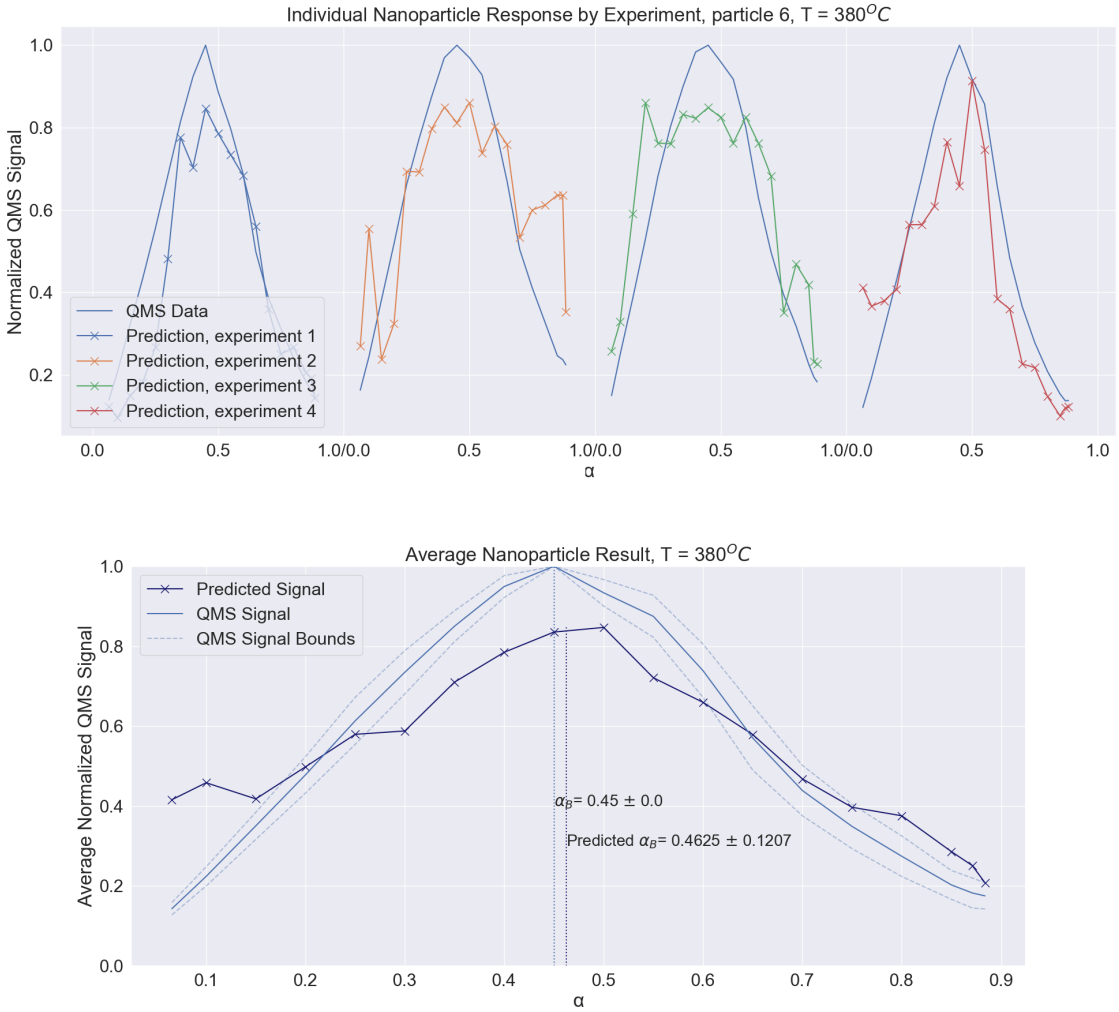

Figure 4.1: Top: Predicted QMS signal in crossed multicolour plotted against the real QMS signal in light blue of particle 6 within the four measurements conducted at T= 380^oC. **Bottom:** The light blue line is the QMS signal over four measurements, hence the slight uncertainty in QMS data, and the crossed dark blue line is the average prediction of seven particles over four measurements - so predicted data from 28 particles in total. Note that the values for α_A, both predicted and real, are shown with a standard deviation of plus-minus 0.0 and 0.1207, respectively. This is not an uncertainty in the data, but rather the standard deviation of the peak position. So in this particular case, the average position of all particles' predicted KPT is $\alpha_A = 0.4625$, and the QMS signals of all four measurements have their peaks at $\alpha_A = 0.45$.

As is clear from figure 4.1, each individual nanoparticle tends to have a distribution and peak slightly shifted from the distribution of the QMS signal, yet they are very similar in overall behaviour. Also, the average predicted KPT in the bottom image coincides quite closely to the QMS peak, given that the QMS is an average of approximately a thousand particles and the prediction is only an average of seven particles. The point of using several experiments that are conducted identically is

to provide the network more data to train on, and to make sure that results are reasonably consistent across experiments, such that we do not accidentally analyze correlations that are only unique to individual experiments rather than individual nanoparticles, for instance as a result of temporarily different noise profiles.

Yet, there is a clear problem with this type of result that we attempt to remedy by the end of this chapter. The crucial question that needs answering is whether the differences in distribution and peak between particles and their corresponding (global) QMS signal are because of problems in the data, problems in the algorithm, or because of actual physical differences on the surfaces of nanoparticles. Indeed, this problem is deceptively difficult, since neither theoretical nor experimental results exist that might verify or refute our findings. Therefore, we are left with the only option of testing our method on new and unrelated samples, to verify that the network's learnt correlations are generalisable and only represent physical correlations between scattered light resultant of nanoplasmonic spectroscopy and QMS data. As mentioned, this discussion will follow shortly after all the results have been presented.

4.2 Optical Phase Diagram of Nanoparticle Kinetic Phase Transitions

This type of result, where the position of KPT is predicted for each individual nanoparticle and then averaged, lets us recreate the phase diagram of figure 2.8, given only optical data, using the trained neural network. For clarity, we briefly summarize the method described in section 3.3: for each temperature, four experiments are chosen for the network to train and test on. Within these four experiments, nanoparticles that seem reasonably responsive to input gases (i.e. work as catalytic nanoparticles) are chosen, which for the nanochip used for these results were particles 1,2,3,4,5,6 and 8 out of 20, selected through the process described in section 3.3. The network is then trained on six out of the aforementioned seven particles and subsequently tested on the seventh, using the cross-validation method known as "Leave-1-Out". The position of the KPT (i.e. peak of the predicted QMS signal) is then calculated for the seventh particle, and the process is repeated until all particles have been tested on. The positions of their KPT are then averaged over experiments and over particles, which results in a single average value for the KPT, which should, according to our hypothesis, converge to the position of the peak of the actual QMS signal by the arguments given in section 2.3.3. By that logic, the phase diagram constructed from calculating the average KPT in this way should resemble that of the "real" phase diagram of the QMS, which gives us a potential measure of the accuracy of the results and a way of analyzing the particles' behaviour through optical data alone by comparing them with, for example, a surface phase diagram like 2.1. This optical phase diagram is shown in figure 4.2.

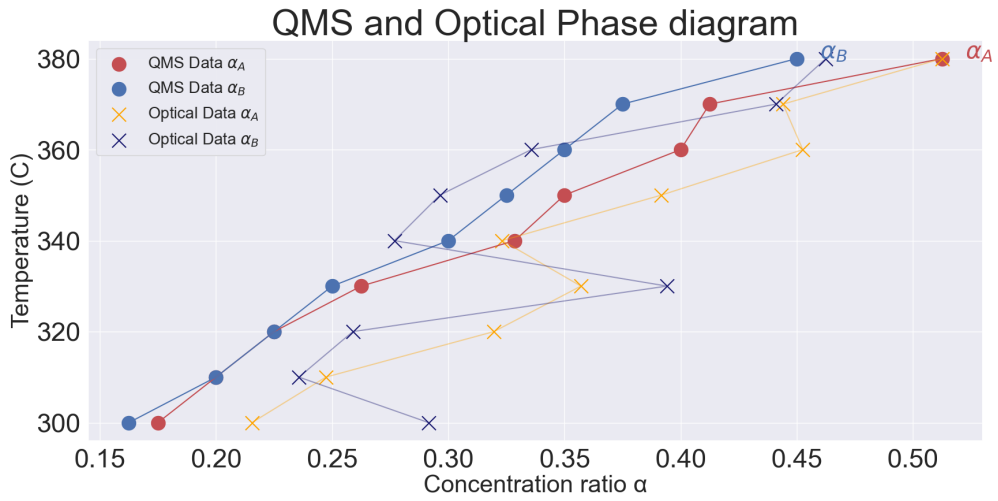

Figure 4.2: The alpha values of QMS peaks α_A and α_B are shown in red and blue dots, respectively. Equivalently, the average alpha value of KPT for all nanoparticles are shown in orange and dark blue crosses, respectively.

In figure 4.2, the positions of QMS peaks α_A and α_B are shown in red and blue dots, respectively. Equivalently, the average position of KPT for all nanoparticles are shown in orange and dark blue crosses, respectively. There are a few important things to note with these results. Firstly, note that the QMS data in figure 4.2 is considerably different from QMS data collected approximately 6-12 months earlier, as shown in figure 2.8. Most notable is that the hysteresis region is not nearly as prevalent, suggesting that the particles may morph significantly over time. Indeed, this behavour is also visible in the FWHM spectrum, as seen in figures 4.6 and 4.7, where the quality of the particles actually seems to improve over time. This is likely due to sample cleaning, as repeated experiments may remove anomalous surface features from contamination or as resultant of the manufacturing process.

There are a few issues we should note about these results. For instance, the experiment conducted at T= 360^{o}C degrees only consists of two measurements, since the other two were too faulty to use. Note also that two of the measurements taken at T= 350^{o}C are also partially faulty, their optical spectra being considerably more noisy than the other two, which likely is part of the reason for these two temperatures' worse prediction quality. Most importantly, however, note that the worst results are from the experiments taken at T= 330^oC and T= 300^{o}C, according to this particular way of presenting and scrutinizing the results. Based on the results in figures A.6 and A.9, we see that this discrepancy is mostly based on a few completely nonphysical outliers, which is likely as a result of the decreased signal-to-noise ratio of lower temperatures (likely partially due to lower reaction rates), making it more difficult for the network to learn the correct correlations. The results at these temperatures are therefore not nearly as disastrous as figure 4.2 makes it seem, and we will return to this. Note, finally, that all of this data was processed in the same way, but it is of course possible that other processing methods might be more apt for different temperatures (i.e. as signal-to-noise ratio decreases and the dominating

physical effects and noise source shifts).

Overall, the results are quite promising in terms of our hypothesis and arguments given in section 3.3. Apart from a few outliers, most notably T= 330^oC and T= 300^oC as discussed, the overall data trend is consistent for both prediction and ground truth. Note here the scale of the image in question; most average values are within approximately $\alpha \pm 0.05$ of their corresponding QMS peaks. Considering that the QMS peak is effectively resultant of the normalized aggregated behaviour of approximately a thousand particles, and the prediction is only the average of seven particles, this seems like a promising result of the average behaviour of each nanoparticle. To get an idea of how the individual nanoparticles' KPT might shift, we can plot the standard deviation of predicted KPT positions in the same image in figure 4.3.

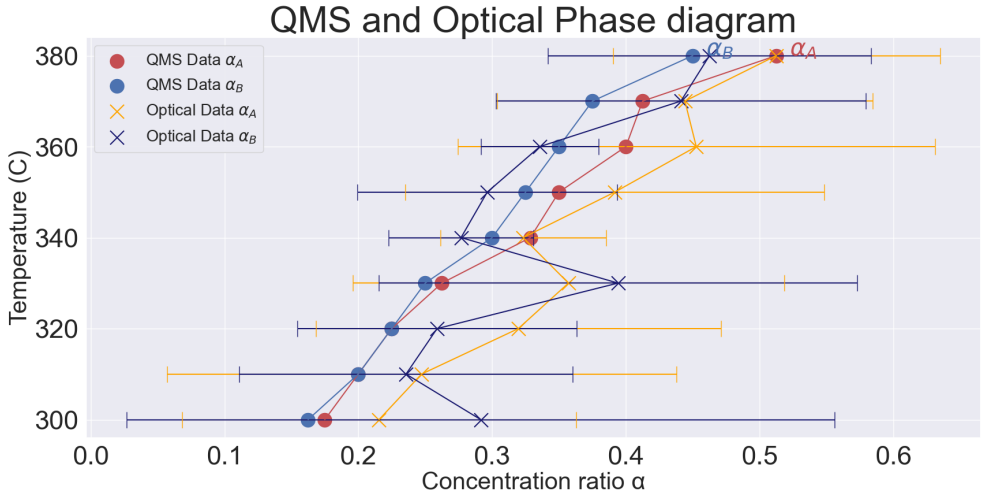

Figure 4.3: The alpha values of QMS peaks α_A and α_B are shown in red and blue dots, respectively. Equivalently, the average alpha value of KPT for all nanoparticles are shown in orange and dark blue crosses, respectively. Plotted in solid horizontal lines are the standard deviation of KPT positions within each set of measurements.

Generally speaking, the standard deviation is low ($\alpha \pm 0.1$) for temperatures whose average prediction coincides well with their corresponding QMS peaks, and high for those that coincide poorly. This is likely because the poor average predictions are not resultant of actual physical effects, but rather because of poor predictions for that particular experiment, ultimately resulting in widely varying predictions. This is further discussed in section 4.5 below. In the next section, we introduce representations of data which let us scrutinize the results of each individual nanoparticle in more detail.

4.3 Individual Nanoparticle Kinetic Phase Transitions

In figure 4.4, we show the results of training a neural network to qualitatively predict the QMS behaviour of individual nanoparticles and their corresponding kinetic phase transitions for the experiment described in section 3 at T= 380^{O}C. Instead of building a phase diagram, like in figure 4.2, for each individual nanoparticle, it is more indicative and informative to present the results of each measurement as a heatmap. By using a heatmap, we can simultaneously scrutinize whether all the particle's KPT are reasonable, whether their overall distributions are reasonable, and how they compare to the QMS peak (and distribution) and the mean position of KPT. The figure is set up such that, since each experiment was repeated four times, each particle has four black crosses, corresponding to that particle's KPT in those four experiments. The y-axis thereby lists each particle that was tested on four times, and the x-axis shows the discrete α-values corresponding to each pulse of gas input to the system. The colour of any individual square thereby shows the predicted QMS behaviour of the corresponding particle at the corresponding pulse of input gas, where light and dark colour correlate to high and low predicted QMS response, respectively.

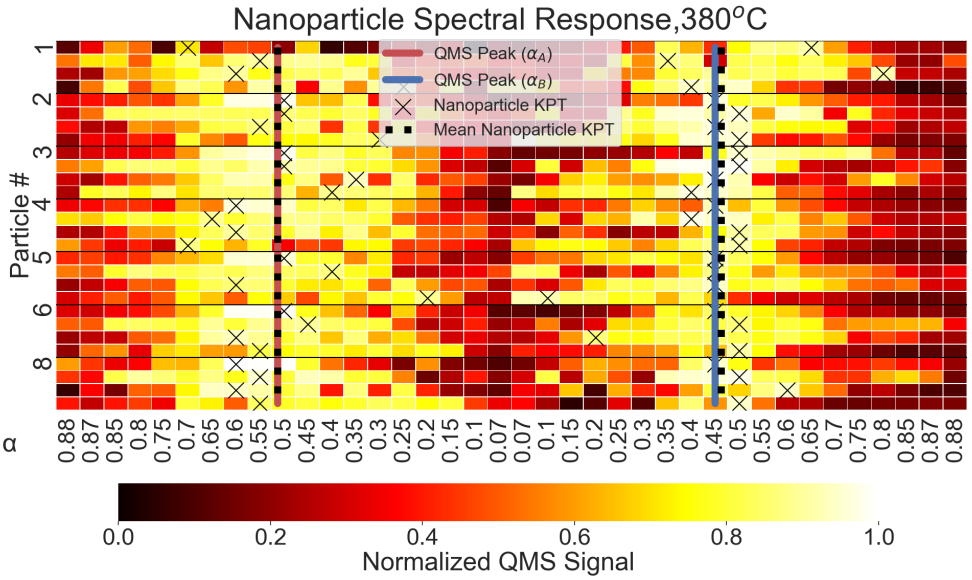

Figure 4.4: Heatmap showing the distribution of predicted qualitative QMS behaviour for each individual nanoparticle as a function of concentration coefficient $\alpha = \frac{P_{CO}}{P_{CO}+P_O}$. Bright whites and yellows corresponds to a high QMS response, indicating high surface catalytic reactivity, and conversely darker shades of red correspond to low QMS response. Each individual nanoparticle's kinetic phase transition (KPT) is calculated to occur at the α-value of its highest QMS response, which is represented by a black cross. The mean α-value for which individual KPTs occur is represented by a black dashed line, and the mean peaks of the actual QMS signal are shown in red and blue for the first and second regions of the QMS, respectively.

There are four main things that should be true with each set of predictions, if the predictions were hypothetically physically accurate. Firstly, the mean position of KPT for all the individual nanoparticles should converge to the position of KPT for the QMS signal (i.e. the whole system, including the patch of nanoparticles). The reasoning for this, as first described in section 3.3, is that the QMS signal is effectively an aggregate average of all nanoparticles' catalytic behaviour in the system (consisting of approximately 1000 particles), therefore the aggregated average of calculated positions of KPT for large numbers of particles should converge to approximately the same value. This point is limited by available data, as currently only seven particles across four are used to form this average. Secondly, the positions of KPT for the same particle should be reasonably consistent across experiments, assuming that the behaviour of particles is fairly consistent over time. This point is also dependant on the amount of data available for the network to train on, since the network will invariably learn some correlations that are unique to each experiment. The more experiments that are fed to the network, the more likely it is that the network finds what correlations are emergent of particles and which are not. Thirdly, the individual positions of KPT should not differ too much between nanoparticles, since we have already systematically removed particles that are very different (i.e. nonresponsive) within the experiment. In essence, this means that any results wherein the mean KPT position is accurate but each individual prediction varies wildly should raise doubts. Fourthly, the overall behaviour of the nanoparticle's predicted QMS signal should roughly follow that of the actual QMS signal, since we expect widely different behaviours to be mostly nonphysical. Keep in mind even the realization of all the aforementioned hypotheses does not constitute a rigorous proof that the positions of KPT are correct, but rather should be considered a guideline for judging the quality of predictions. Note also that the experiments are not necessarily temporally adjacent, some may be conducted even months apart. There is therefore no guarantee that particles are physically the same across experiments, instead we expect them to morph to some degree over time.

Nonetheless, with this shorthand, we return to focusing on figure 4.4. Focusing firstly on α_B in the right side of the image, we find that apart from particle 1 and a few singular outliers, the position of KPT seems to be reasonably consistent across experiments for each particle. The mean position of KPT also corresponds well with the position of QMS peak α_B and the distribution of predicted QMS behaviour roughly matches anticipated distribution of a QMS signal, as evident in this heatmap by following the colour trend of black/red-white/yellow-black/red. Indeed, this type of result is therefore indicative of what we expect to see. Turning our focus instead to α_A, it seems at first glance that the results are much less reliable, owing to the uneven placement of individual KPTs. However, upon closer inspection, we note two things: firstly, barring particle 5, almost all particles' KPT are predicted to be very closely together for each individual nanoparticle. Secondly, the distribution of QMS signal has a much broader "top" in the sense that the predicted QMS signal is consistently very high for a long time (corresponding to a long white/yellow band in the heatmap). This is likely a consequence of the optical spectra looking very similar for a range of different $\alpha-$values and results in the

position of the KPT being less unambiguous and harder to properly define within this band of high activity, ultimately leading to more spread predictions. Why this is physically true is unknown, but analysis of the change in FWHM, as shown in figure 4.6 shows the same general trend of α_A having more spread positions of KPT than α_B in general. These outliers might also to some extent be the byproduct of a non-optimized network, for reasons that will be covered in further detail in section 4.5. Note also that in this particular case, despite the outliers, the mean nanoparticle KPT position coincides exactly with the QMS peak, which strengthens the prospects that they are not merely coincidence.

To exemplify that this method also works at temperatures with considerably different surface physics (mostly due to low surface oxide formation), in figure 4.5 below we show the results for the same experiment conducted at T= 340^{o}C.

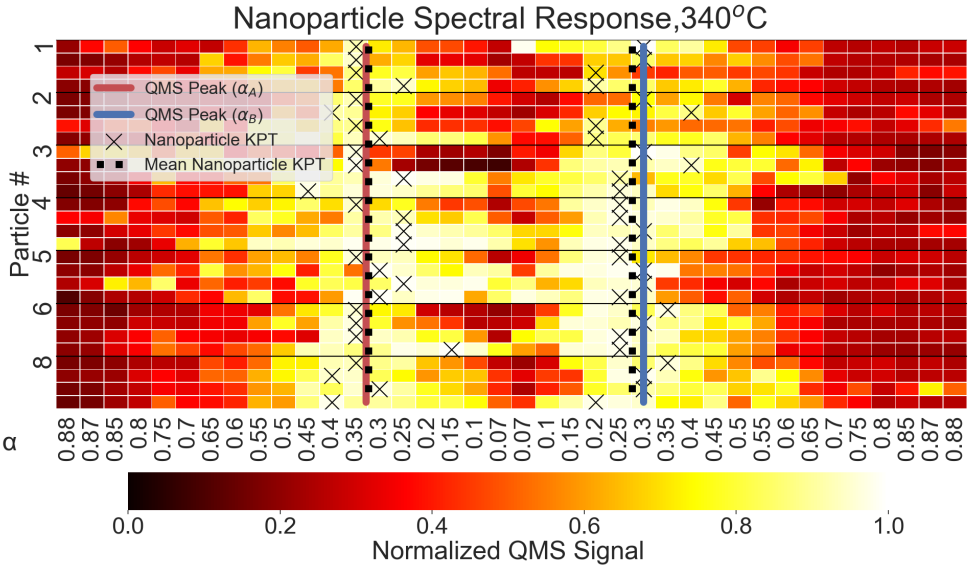

Figure 4.5: Heatmap showing the distribution of predicted qualitative QMS behaviour for each individual nanoparticle as a function of concentration coefficient α. Bright whites and yellows corresponds to a high QMS response, indicating high surface catalytic reactivity, and conversely darker shades of red correspond to low QMS response. Each individual nanoparticle's kinetic phase transition (KPT) is calculated to occur at the α-value of its highest QMS response, which is represented by a black cross. The mean α-value for which individual KPTs occur is represented by a black dashed line, and the mean peaks of the actual QMS signal are shown in red and blue for the first and second regions of the QMS, respectively.

Interestingly, at a glance, the results at this temperature seem more consistent with our hypotheses in comparison with higher temperatures. Since some of the measurements taken at T= 350^{o}C and T= 330^{o}C do not follow this trend, it is likely that these particular measurements were only temporarily of higher quality, rather than lower temperatures generally showing better results. A more detailed discussion on these discrepancies follow in 4.5.

To compare the results of our two networks, the LSTM network as described in section A.3 and the DNN described in section 3.3, see figure 4.6 below.

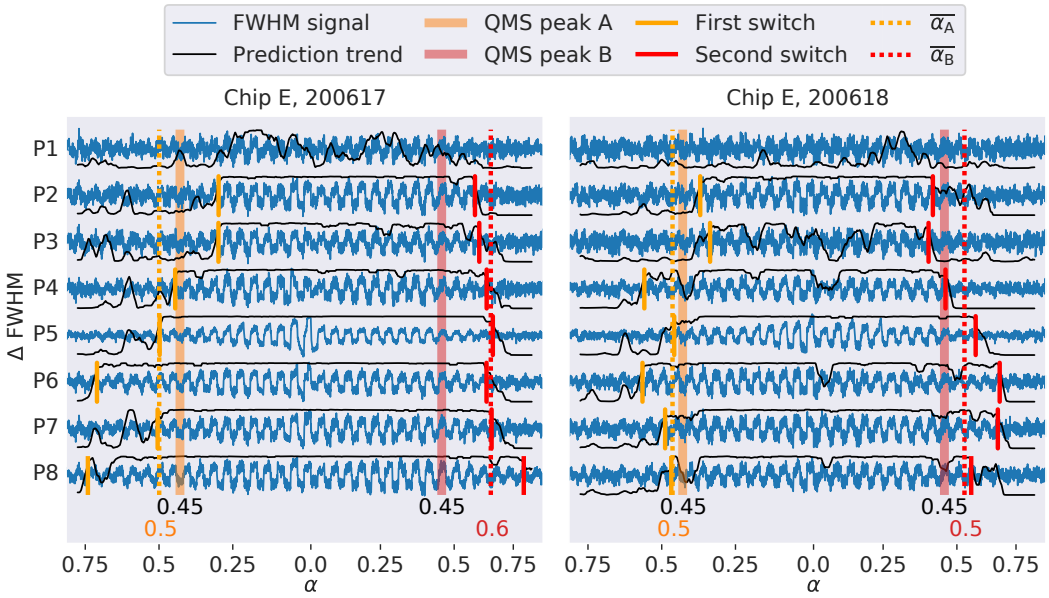

Figure 4.6: Results of the RNN which predicts the position of switching point between regions in the FWHM signal of clear oscillations and noisy oscillations. The change in FWHM over time is in blue, predictions of KPT positions α_A and α_B are in short yellow and red lines, respectively, whilst the real QMS peaks and average predicted KPT position are in long solid and dotted yellow and red lines, respectively.

Since these results were not produced impartially (the two methods were co-developed), it is difficult to make deep unbiased analysis. However, there are a few interesting trends that are worth noting. Notice, for instance, that the position of KPTs are considerably more spread in the first half (corresponding to α_A), especially for the measurement resulting in the image on the right. This is consistent with what was found in 4.4. Note, also, the difference between the two images for the first switch of particle 8 - the same particle is predicted to have a KPT (and, indeed, this prediction does not look unreasonable when analyzing the FWHM signal by eye) in $\alpha_A = 0.5$ and $\alpha_A = 0.7$. This is consistent with what we see in our heatmaps as well, where the predicted position of a single nanoparticle's KPT is not necessarily as consistent across measurements as we expected, but rather might shift considerably between (identical) measurements.

Note also that the average KPT coincides closely will the QMS peak in these results as well, with the exception that the average KPT is (almost) always at a higher α than the corresponding QMS peak. We believe the reason for this is because the first particles are consistently more likely to have their KPT at lower α (for unknown reason, though this is consistent with results in figure 4.11), but the network often has difficult finding this switching point for these particles. Therefore, the mean

KPT becomes skewed towards higher α in these particular results. Overall, however, the results seem reasonably consistent across the networks, in the sense that each nanoparticle's KPT is close in α to its corresponding QMS peak with few outliers.

In figure 4.7, we show the results of the RNN training on older data, taken four months before the results shown in figure 4.6.

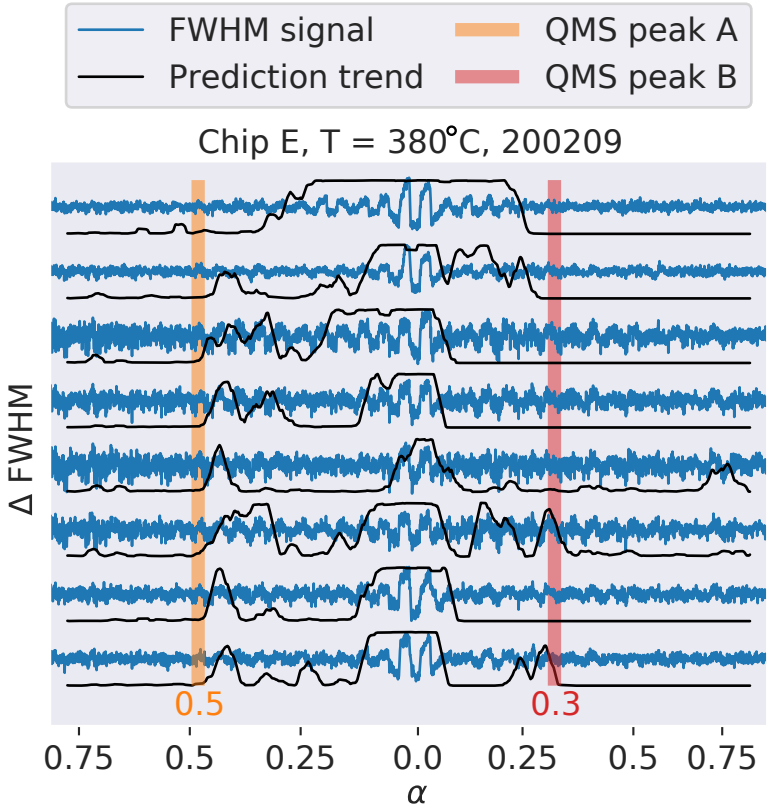

Figure 4.7: Results of the RNN which predicts the position of switching point between regions in the FWHM signal of clear oscillations and noisy oscillations. The change in FWHM over time is in blue, predictions of KPT positions α_A and α_B are in short yellow and red lines, respectively, whilst the real QMS peaks and average predicted KPT position are in long solid and dotted yellow and red lines, respectively. Note that the FWHM signal is far too noisy and unintelligible to make proper predictions on in this old data.

Note that the results are considerably different, and considerably worse for the purposes of our LSTM network, for the same experiment conducted 4 months earlier. This is yet another example of particles morphing or otherwise changing over long periods of time and, more interestingly, they seem to become cleaner and more responsive over time. This shows a fundamental problem with our RNN approach which does not seem to exist in the DNN (at-least, not to the same egregious extent) - that it is so dependant on the particular appearance and behaviour of the FWHM signal that it is probably not very generalizable. A deeper discussion of this will follow in section 4.4.

The rest of the results can be found in the appendix, section A.2.

4.4 Cross Validation Between Chips and Reliability of Results

In this chapter, we let the network predict the behaviour of particles from a completely different sample (nano-chip) than the one it was trained on. As mentioned before, we would ideally want to compare these results to a theoretical model or experimental results, to find some rigorous margin of error in our results. Since no such possibility exists, we instead attempt to predict the behaviour of particles across chips, to verify that the correlations learnt by the network are indeed indicative of the real physical optical responses. Indeed, through the results presented in the last chapter, we cannot be sure whether the network has actually learnt real general physical correlations between optical spectra and QMS signal, or whether it has only learnt how to recreate the particular QMS signals from this particular sample, which means our results would only really show that there is enough data in each optical spectra to recreate the QMS signal and nothing else. Testing the network on a different sample, which is fundamentally physically the same but different in specific behaviour, will elucidate more on which of these alternatives are true. This forms the basis of analyzing the validity of our results.

In figure 4.8 below, we show the results of training the network on one sample (nano-chip E) and testing it on another sample (nano-chip C). Note that we do not know which particles are of high quality in this new chip, so we choose to focus on the 8 particles that are furthest away from the gas input, since this is consistent with what particles were trained on in nano-chip E. The chips are manufactured identically in every way, except for the singular difference that the nano-channel is slightly narrower in nano-chip C.

Both images within figure 4.8 show the results of predicting on all eight particles of a single experiment, where the top and bottom images corresponds to the first and second halves of the QMS signal, respectively. As discussed in section 2.3, the first half corresponds to a CO-poisoned nanoparticle surface and the second half an O-poisoned (i.e. with oxide formation) nanoparticle surface. Note that the predicted distributions (in multicolour) overall seem very reasonable, and indeed are even better than expected, since both the distributions of the real QMS signals and predicted QMS signals have morphed considerably in comparison to what was shown in figure 4.1 for nano-chip E. To further exemplify this, in figure 4.9, we show the average prediction for the two images in figure 4.8.

Comparing with figure 4.1, which shows a representative example of the QMS signal for an experiment on chip E, we see that the distribution of the QMS signal is considerably different in figure 4.9. Indeed, we see that the results of figure 4.9 are more similar to other results in the field, which means that the new nano-chip is likely of better quality than the old one.

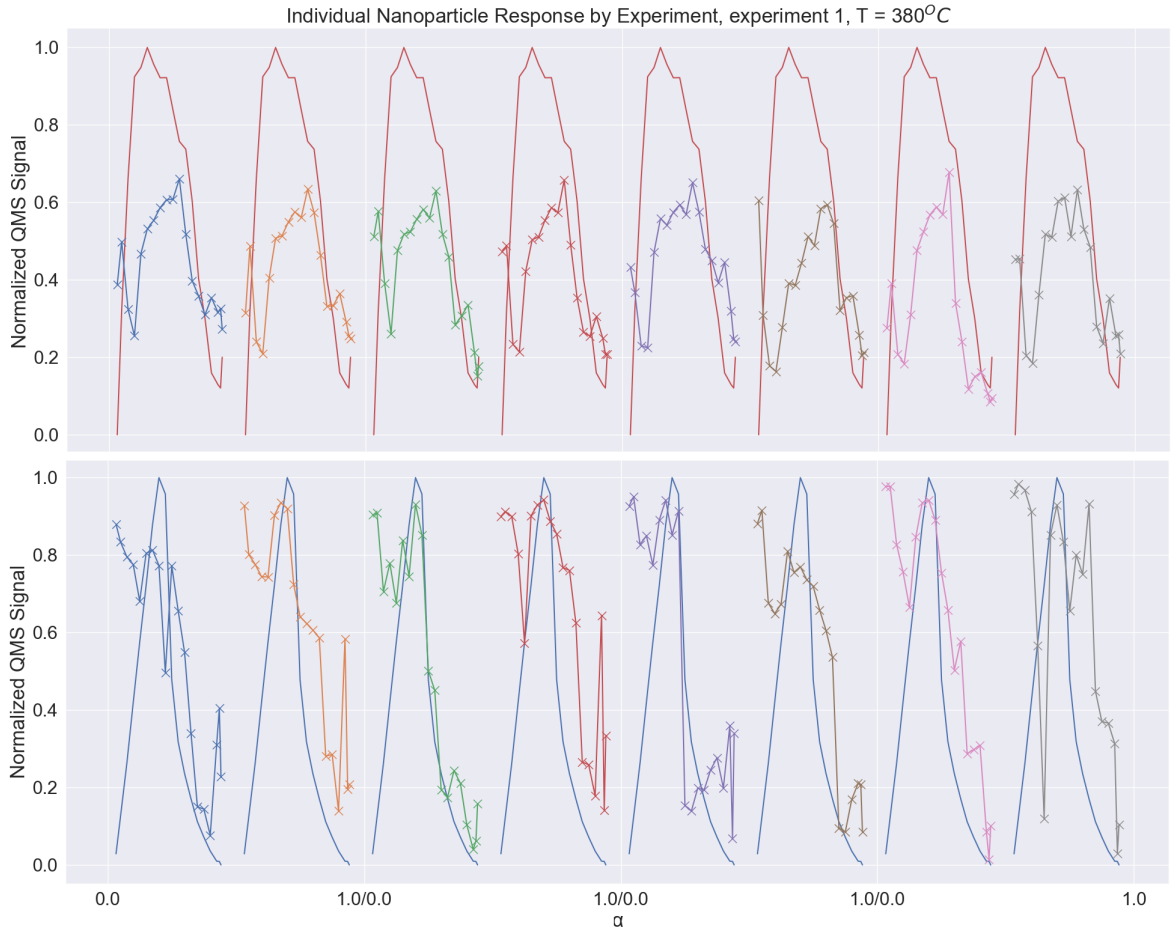

Figure 4.8: Top: Predicted QMS signal in crossed multicolour plotted against the real QMS signal in red for the measurement taken at T= 380^oC, corresponding to the period of the experiment in which the surfaces of nanoparticles' are CO-poisoned (α_A). **Bottom:** Predicted QMS signal in crossed multicolour plotted against the real QMS signal in light blue for the measurement taken at T= 380^oC, corresponding to the period of the experiment in which the surfaces of nanoparticles' are O-poisoned (α_B).

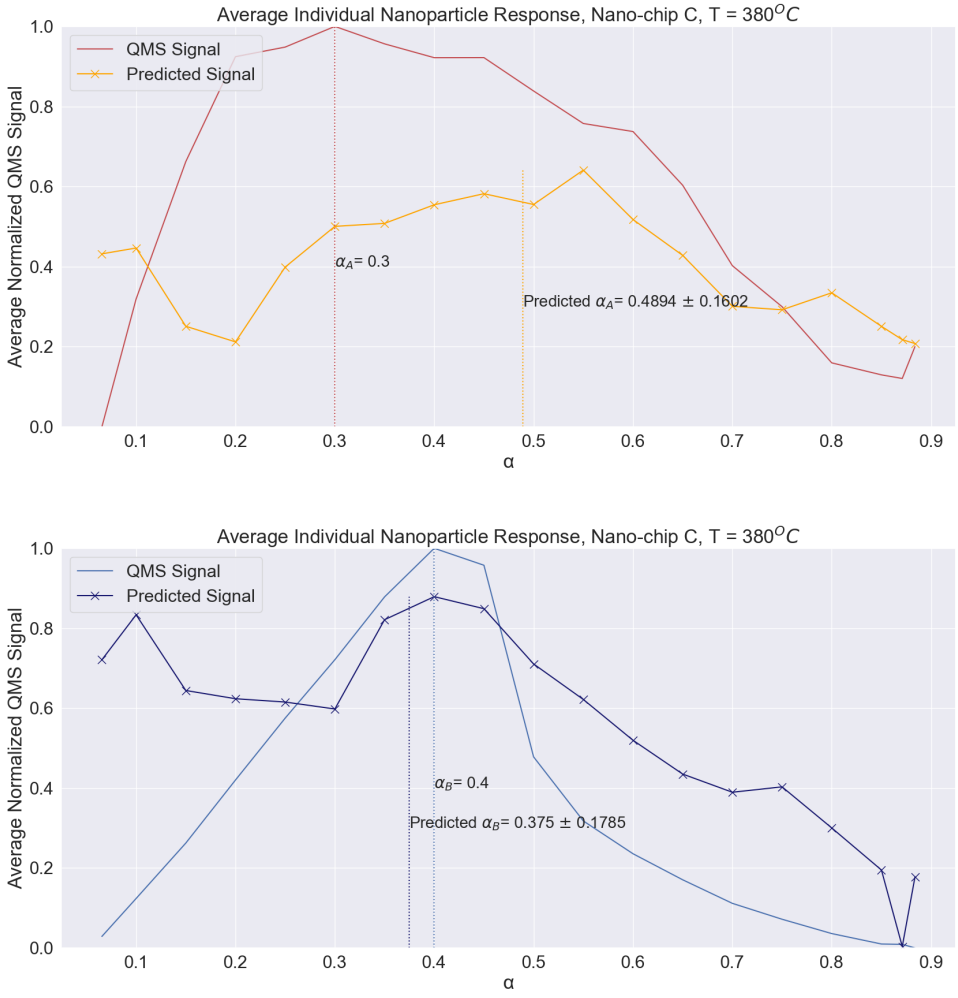

Figure 4.9: Top: Average predicted QMS signal in crossed orange plotted against the real QMS signal in red for the measurement taken at T= $380^{o}C$ on nano-chip C. The peak QMS signal for this experiment is $\alpha_A = 0.3$, and the average predicted KPT is $\alpha_A = 0.4894 \pm 0.1602$. **Bottom:** Average predicted QMS signal in crossed dark blue plotted against the real QMS signal in light blue for the measurement taken at T= $380^{o}C$ on nano-chip C. The peak QMS signal for this experiment is $\alpha_B = 0.4$, and the average predicted KPT is $\alpha_B = 0.375 \pm 0.1785$

Note, also, that the distributions in the top image of figure 4.8 are consistently quite low (only between about 0.2 and 0.6 in "normalized" y-value) in comparison to their corresponding QMS distributions. Though this is not physically meaningful, since the absolute values are essentially arbitrary and only their relative difference is of interest, it is worth mentioning and being transparent about since it means that the neural network's training has stopped prematurely. This is likely a result of 'real' correlations being harder to find in this region, either through higher noise or perhaps because there is less signal in CO-poisoned regions since there is comparatively less activity (oxides are highly reactive, see section 2.1.1).

Though the overall distributions are reasonable, the average predicted KPT position does not correspond terribly well with the corresponding QMS peak α_A, instead being within approximately $\alpha \pm 0.2$ of eachother.Further investigation into the consistency of these results is needed, but it may well be resultant of the network's relatively unsuccessful training in this region. Note that these results do not inherently mean that the network predicts exactly the correct positions of KPT, but the overall quality of predictions does make it likely that the network has learnt actual physical correlations, which is a crucial step in the overarching goal of finding structure-function correlations between surface kinetics and optical spectra.

Note that there are a few complications that make predicting across samples (chips) difficult. Firstly, we have shown several times in this thesis that the behaviour of each sample might change considerably over time. The exact reasons behind this are unknown, it might for instance be due to shifting nanoparticle morphology or gradual CO oxidation on their surfaces due to carbon formation. If this results in considerably different surface kinetics, then it is possible that the sample that the network was trained on essentially has different physical correlations between optical signals resultant of nanoplasmonic spectroscopy and its corresponding QMS behaviour. Indeed, as discussed in the next section, we believe this is the reason for the consistently high predictions in the first half of all predictions in the bottom image. It is also unlikely that training the network on a single sample has given it enough data to learn all possible correlations between optical data and QMS behaviour, therefore testing on severely different samples is unlikely to provide ideal results. Further discussions on these and other issues follow below.

4.5 Sources of Error & Uncertainty

Within this section, we discuss the sources of error and uncertainty both in the network's perfomance and in the experiment itself. We also make several notes about how these problems might be remedied, as well as how our results compare to analogous results found on Platinum nanoparticles.

4.5.1 Investigating Deviations of Nanoparticle Behaviour between Nano-fluidic Chips

Firstly, note that the resulting optical spectra of scattered light from nanoparticles on the new sample are considerably different in a few notably interesting ways, see figure 4.10. In the top-left and top-right images, we show the spectra of a single individual nanoparticle for nano-chip C and nano-chip E, respectively. Each coloured curve corresponds to a single step in time, and the particles plotted here are overall representative of all particles' behaviour within the corresponding nano-chip (to the eye).

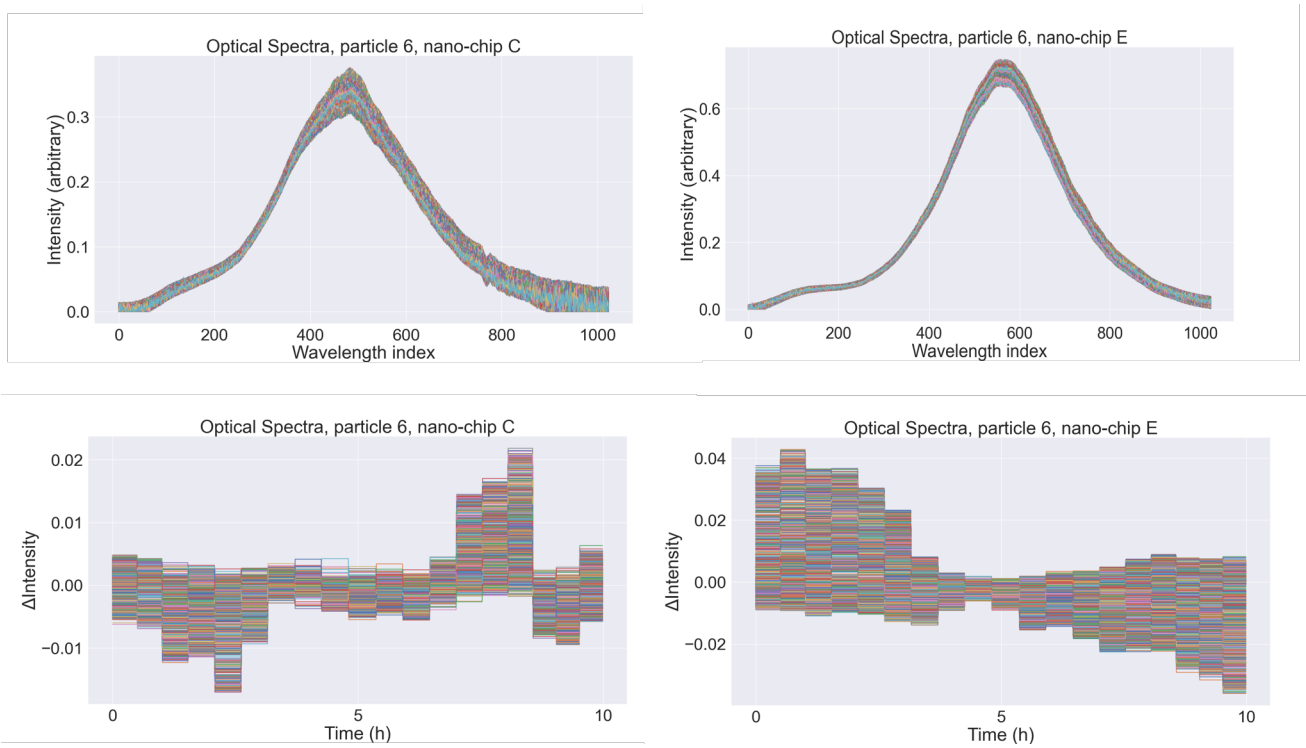

Figure 4.10: Top: Optical spectra of particle 6 within nano-chip C and particle 6 within nano-chip E. Note the overall higher noise within nano-chip C, especially in the fat tail towards higher wavelengths. **Bottom:** Deviation in optical spectra from its temporal mean, fitted to a step function where each step is the temporal mean value of each wavelength for each discrete value of α. For details, see section 3.3. Note that only the second half of the experiment is plotted here, corresponding to $0 < \alpha < 1$.

The most important thing to note within the gaussian spectra is the considerably noisier nature of the new sample (nano-chip C) compared to the old one (nano-chip E), especially within the fat tail towards higher wavelengths. Since this is consistent across all particles within the nano-chip C, it is likely resultant of some intrinsic noise source within the nano-chip. Since the chips effectively only differ in that the particles are different (though they are manufactured to be the same), and that the nano-channel within nano-chip C is narrower, it is likely that the increased noise is

due to light interference through scattering with the walls of the nano-channel.

The two bottom images are considerably more processed and difficult to understand at a glance, but the underlying information they hold is very important, since they are exactly what the optical data looks like when we feed it to the neural network. We show the deviation of the optical spectra from its temporal mean separately for each wavelength, where each wavelength is a separately coloured line, and fit it to a step function where each step is the temporal mean value for each wavelength across each discrete value of α. Therefore, each one of these steps corresponds to a single value of input gas ratio α, and is equal to the mean deviation from the original spectra over time. To be abundantly clear, in this figure, a value close to 0 means that all the gaussian spectra in that time period are similar to the mean gaussian spectra for the whole experiment, and a positive value means that the gaussian spectra are overall shifted upwards compared to the mean gaussian spectra. For the results of nano-chip E, the deviation in measured intensity for each individual wavelength tends to be high for the entire spectrum in the start, and then almost linearly decrease over the course of the experiment. In other words, this trend clearly shows that the deviations in measured intensity are high and positive when the surfaces of the nanoparticles are O-poisoned, and that they are low when the ratio of CO and O are roughly equal, and again that they are high and negative when the surface becomes CO-poisoned. See section 2.3.3 for further details.

The most important results here, however, are simply how different the same results are from particles in nano-chip C. Indeed, this is the main point we should take away from these results; despite how different the QMS signals (as shown in figure 4.9) and optical spectra (as shown in figure 4.10) look in the new sample, we still get fully reasonable predictions from the neural networks. Keep in mind that all spectra that the network has trained on show the same trend as in the bottom right image of figure 4.10, and despite the network never having seen the trend in the bottom left image of figure 4.10, it still corresponds to the correct QMS signal, apart from the beginning of the high predicted qualitative QMS behaviour in the first few values of α. Looking at the bottom left image of figure 4.10, we see that a likely explanation for this is because the behaviour of the optical data in the beginning of the experiment is similar to the behaviour around the QMS peak. More in-depth analysis would be needed to verify that this is not just the result of increased noise, as discussed previously. However, this type of result is indicative of what we initially hoped to see.

4.5.2 Comparison with Analogous Results on Platinum Nanoparticles

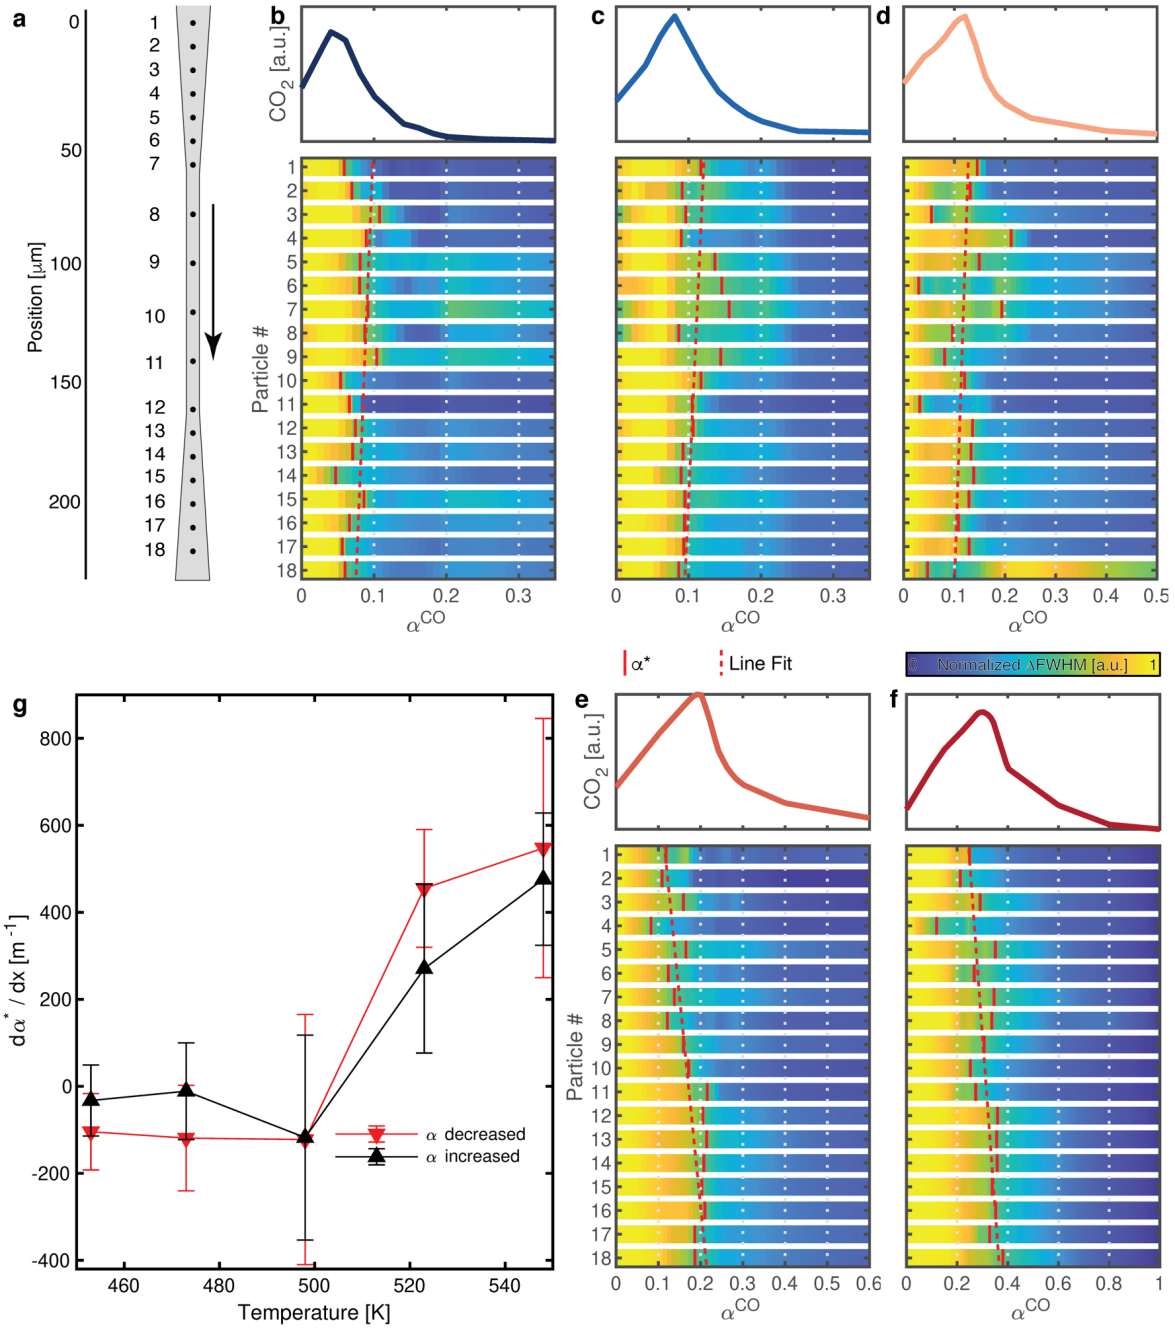

Figure 4.11: a) Schematic of a nanochannel with 18 particles labeled with numbers. The arrow indicates flow direction. b-f) CO_2 production (top panels) and optical response for the 18 particles (bottom panels) at five temperatures for a single sweep from $\alpha = 1$ to $\alpha = 0$. Each individual particle's data has been normalized to span between its minimum and maximum value. Red lines indicate the position of the kinetic phase transition extracted for each particle. g) The kinetic phase transition for each particle against position in the channel (triangles) and a linear fit to the data (dashed lines). h) The slope of the dashed lines in (f) as a function of temperature. Error bars correspond to 95% confidence intervals of the fitted lines. The up and down arrows in (g-h) correspond to data from the increase and decrease in α, respectively.

There will clearly be slightly differing results depending on how exactly the experiment and data analysis is conducted, in particular which particles are trained and tested on, which noise cancellation method is used (some aggressive methods may remove signal in addition to noise), and even which date the analysis is done, since the nature of the data changes over long periods of time as discussed in section 2.3. Therefore, a rigorous method of identifying and analysing the accuracy of calculated KPT positions is needed to really demonstrably isolate which methods are most physical and relevant. As mentioned before, there is effectively no prior theoretical or experimental work that can be used for this specific purpose, and this thesis' alternative method was to use another sample (nano-chip C) to test findings on. However, similar experiments done on nanoparticles of different material have found similar results, which gives us some reference on what results might be reasonable to expect from this experiment.

Above, in figure 4.11, is an example of results of a similar experiment, courtesy of the Langhammer team, where platinum rather than palladium nanoparticles were investigated. Without going into too much detail, note that the process of identifying the position of KPT for platinum nanoparticles is possible through analyzing the change in FWHM alone, such that more complex methods of data analysis are unnecessary in this case. In figure 4.11, the position of KPT are for each individual nanoparticle is shown as a red solid line, and the average trend of KPT position is shown in dotted red. I will not comment on the reliability or in-depth analysis of the results presented in this image, as this is done in great detail in the original report which is soon to be published, but note that these results mirror what has been found in this thesis. Generally speaking, the individual nanoparticles' KPT are found to be around the position of the QMS peak, with the average KPT position being approximately the position of the QMS peak. Indeed, even here, there are a few notable outliers whose KPT are far away from the QMS peak, on the order of approximately $\alpha \pm 10^{-}1$ or occasionally even greater. Note that though the experiment whose results are shown in figure 4.11 show positional dependence for the position of KPT, we have done no such investigation in this thesis for the aforementioned problems of later particles likely being too out of focus to properly analyse, but this is a possible focus for further analysis of these results.

4.5.3 Problems and Observations regarding Neural Network Performance

One problem with the results of our deep densely connected neural network is that the network might learn correlations that are only unique to the particular measurements that we train it on. Indeed, this is likely the case for the predictions of the experiment conducted at T= 360^{o}C, since only two subsequently conducted measurements were available for analysis. It is unlikely that this is a good enough basis for the network to distinguish particle-specific features from measurement-specific ones. This can be seen from figure A.3, where it seems like every particle from each experiment acts in roughly the same way, indicating that the network has learnt measurement-specific trends rather than particle-specific ones. Indeed, this would

be likely be the case even if any other method was used for analysis, but it is still worth noting, since this problem is effectively solvable by, for instance, using more data when training the neural network.

It is also possible that the results might be improved by choosing another set of particles. Since it is likely that the ones we have discarded are only of poor quality as a result of them being out of focus, as described in section 3.3, they might otherwise be very high-quality particles. The quality of particles might indeed also be temperature-dependant, as some are possibly more resistant to considerable surface changes during heating (i.e. a mostly homogenous disc versus a more fragile structure). Therefore, the results may well be improved with a more robust method of identifying the quality of particles, and with an experimental setup which can more consistently focus on each individual nanoparticle. Indeed, it is possible that including only a few faulty spectra or faulty particles in the training of the network may cause its performance to drop significantly, since we are effectively asking the network to learn correlations in some data which do not actually exist.

A brief note should also be made about the results of training and predicting with the network across experiments conducted at different temperatures instead of across particles, as some inquisitive readers will invariably wonder if this strategy might possibly be preferable. Another way of using the network to find the position of KPTs is to train it on all (good) particles from a set of measurements, then test it on results from the same particles but from other measurements. This is likely somewhat inferior to the method chosen in this thesis, since the network might "overtrain" on the training measurements if the training, validation and test measurements are very similar (which we assume them to be, otherwise the experiment setup would be too inconsistent for proper analysis). Therefore, the results often look very good with this alternative method, but this is less likely to correlate with real physical correlations.

The particles are effectively fabricated as homogenous discs (though, of course, there is always heterogeneity at higher resolution), whereupon they are likely to morph and change shape when heated and subsequently cooled [4]. Their shapes may therefore shift slightly over time, affecting both the behaviour of QMS signal (as discussed in section 3.3), and the behaviour of each individual optical spectrum. Therefore, training the network on several experiments that are conducted across long stretches of time may not be ideal - and indeed, this is what we have done for many of the results shown in the heatmaps within this thesis.

The optical spectra likely also contain some leakage of light from neighbouring particles and possibly from the walls of the nano-channel. This leakage is seemingly quite significant for particles nearest the patch, where the level of light leakage is likely highest. Mechanical vibrations are also a source of uncertainty in the measured optical spectra, as are thermal drifts, which may lead to the particles going out of focus of the camera or perhaps entirely out of the field of view. The tendency of particles to go out of focus can in general be remedied by simply normalizing the spectra (as the measured intensity of scattered light is likely to change but not its

peak position), but a particle that leaves the field of view entirely is unsalvagable. It might therefore be a better idea to focus the camera simultaneously on the entire channel whereupon the nanoparticles lie and thereafter isolate the particles using simple maximal intensity calibration or potentially through more advanced processing scripts such as something similar to ML-based particle-tracking algorithm DeepTrack [11].

It should also be mentioned that the measurements used to train the network have been hand-picked, since some experiments effectively become almost unusable, owing to particles drifting out of the view of the camera, construction work, contamination in the gas pipelines or other reasons. These problems are often solvable using different pre-processing techniques, such as by normalizing the spectra over time, such that the peak intensity of each gaussian spectrum is 1. This is especially important since the peak intensity (but not peak position) of each particle may shift considerably depending on the particle's particular drift within the camera's view. Normalizing all the spectra effectively reduces the impact of this drift on the optical spectra. However, this "aggressive" method of processing the data to salvage poor experimental results often removes relevant but unknown correlations in the data, effectively worsening the overall quality of predictions. However, this shows that more processing methods may well be considered to improve the quality of predictions, and we have by no means done an exhaustive test in this thesis.

It is likely that the network architecture is not optimized for each set of experiments (and, conversely, very optimized for a few particular sets), which leads to suboptimal results in some cases. Indeed, this could be solved relatively simply by increasing the computational power available to train the network, since this means that we could input many more relevant observables to the network at the same time (i.e. statistical features of the individual spectra, such as change in FWHM, peak position or spectral centroid over time). Since very limited computational power was available for this project, the network's design and subsequent performance was to some extent limited by this restriction.

As mentioned previously, it is clear that some predictions were clearly considerably worse than the others. In particular, according to our analysis of the results as described in the last two sections, it seems that the experiment at T= 330^oC led to the worst results in prediction. It is worthwhile to dwell on this observation for a while, to investigate what might have led to such results and what might be done to remedy it. Indeed, training the network on this temperature leads to consistently poor results, which means the problem is more likely in the data rather than the network (or, less likely, because the network is poorly optimized for this particular temperature for unknown reasons). Though it is difficult to tell by eye that something is wrong with the spectra, note that the four measurements used in our predictions are taken approximately a month apart. Training on each set of measurements separately is revealing, as training only on the earliest set of measurements (consisting of two measurements) garners considerably worse predictions than training only on the later set of measurements. The result of training only on the later two measurements is shown in figure 4.12, and shows that the results can be

considerably improved by picking out the best measurements. Why these particular measurements are so much better is unclear, but likely because they were in better focus, had an overall less impactful noise profile, or because of other temporary experimental discrepancies.

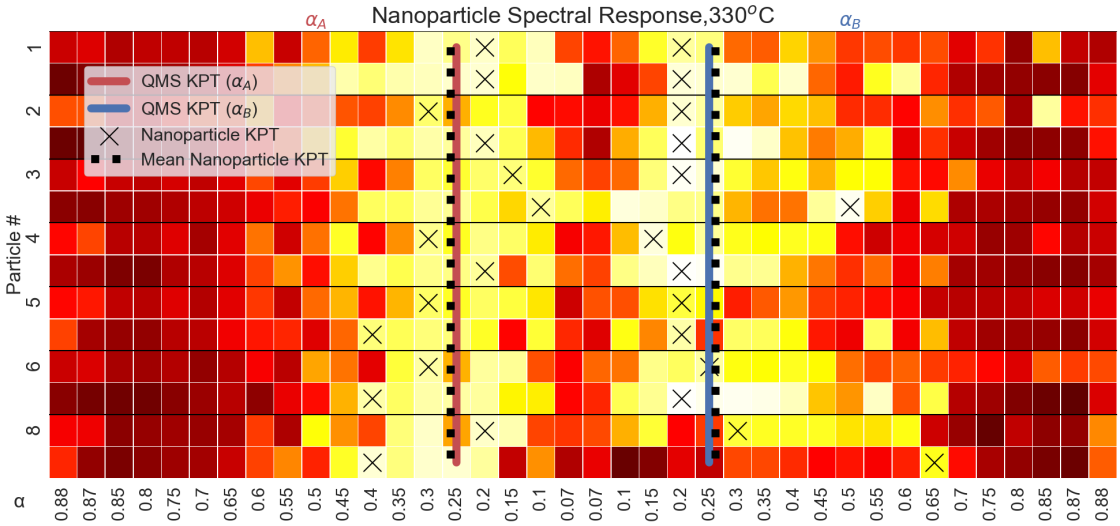

Figure 4.12: Heatmap showing the distribution of predicted qualitative QMS behaviour for each individual nanoparticle as a function of concentration coefficient α. Bright whites and yellows corresponds to a high QMS response, indicating high surface catalytic reactivity, and conversely darker shades of red correspond to low QMS response. Each individual nanoparticle's kinetic phase transition (KPT) is calculated to occur at the α-value of its highest QMS response, which is represented by a black cross. The mean α-value for which individual KPTs occur is represented by a black dashed line, and the mean peaks of the actual QMS signal are shown in red and blue for the first and second regions of the QMS, respectively.

For the sake of comparison, we have included the heatmap of the other experiment in figure A.6 in the appendix, which is notably worse. This certainly shows that more effort is needed to investigate why some measurements within the same experiment seem to be so notably worse with our approach.

Finally, I believe it is very important to note an obvious issue that arises with our chosen method of retrieving the position of KPT for each individual nanoparticle, and how its simplicity might hurt the accuracy of our predictions in an unreasonable fashion. Note that some predicted QMS behaviours have an overall reasonable distribution, but quite flat "tops", in the sense that their predicted QMS activity is almost constant at a high level for a large range of α-values, see the prediction in green in the top image of figure 4.1 for an example. This is not an unreasonable behaviour by any means, but the nature of our method of defining the KPT as the position of maximal QMS activity means that the prediction in this case becomes rather arbitrary. Though we have chosen this method in this thesis for the sake of absolute transparency, it may be a better idea to pre-define an area where the KPT should occur, such that we do not claim that the network predicts a KPT at an α

where it is clearly unreasonable. Indeed, something similar to this trade-off should be implemented before the positions of KPT are used for any practical purposes, and should be investigated further.

4.5.4 Investigating a Novel Approach to Noise-cancellation

Here, we present and discuss one particular attempt at remedying some of the issues mentioned above. One novel idea within this process was to conduct a single measurement at each temperature, just before all other measurements at that temperature are taken, without inputting any gas to the system. We would thereby have a unique noise profile for the experiment conducted at every temperature, excluding inevitable stochastic differences between otherwise identical measurements. This noise profile laid the basis for the simulated FWHM signal used to train the LSTM network, as described in section A.3. However, we also wanted to investigate the impact of subtracting these noise profiles from the "real measurements", to see if this might affect the predictions. It is possible that the results may improved further if the noise was processed in some way, for instance by cleaning it up, fitting gaussian curves to it or conducting a deeper analysis of the noise itself, yet even what expeditious method of noise cancellation we did try yielded very interesting results. Two examples of this can be found in figure 4.13 below.

It is likely that one of two reasons lie behind this apparent improvement of results. Firstly, it is possible that the subtraction of a noise-cancellation spectrum from all other spectra results in all spectra looking almost indiscernibly similar. If this is true, then the network effectively overtrains by testing and validation on data that is far too similar to the training data, thereby artificially imnproving the appearance of the results without it having any corresponding physical meaning. However, this can likely be ruled out by testing on a physically different chip, since the network would have had to learn real physical correlations to properly resolve the significantly different spectra (which are nonetheless based on the same physics). Secondly, it is possible that the noise cancellation does what we intend; subtract off a large portion of what is effectively just noise from the original optical spectra, thereby leaving mostly relevant physical signals to be analyzed. This would likely be futile if standard methods of analysis were applied, but the resulting higher general signal-to-noise ratio might make it easier for the network to learn the relevant correlations. Indeed, the network does train considerably faster (up to an order of magnitude faster in certain cases) using this method, though this may also be as a result of the aforementioned overtraining. These results should therefore not be viewed as pivotal to this thesis, but rather as an interesting pre-study in how we might use noise-canceling measurements in the future to possible improve our results.

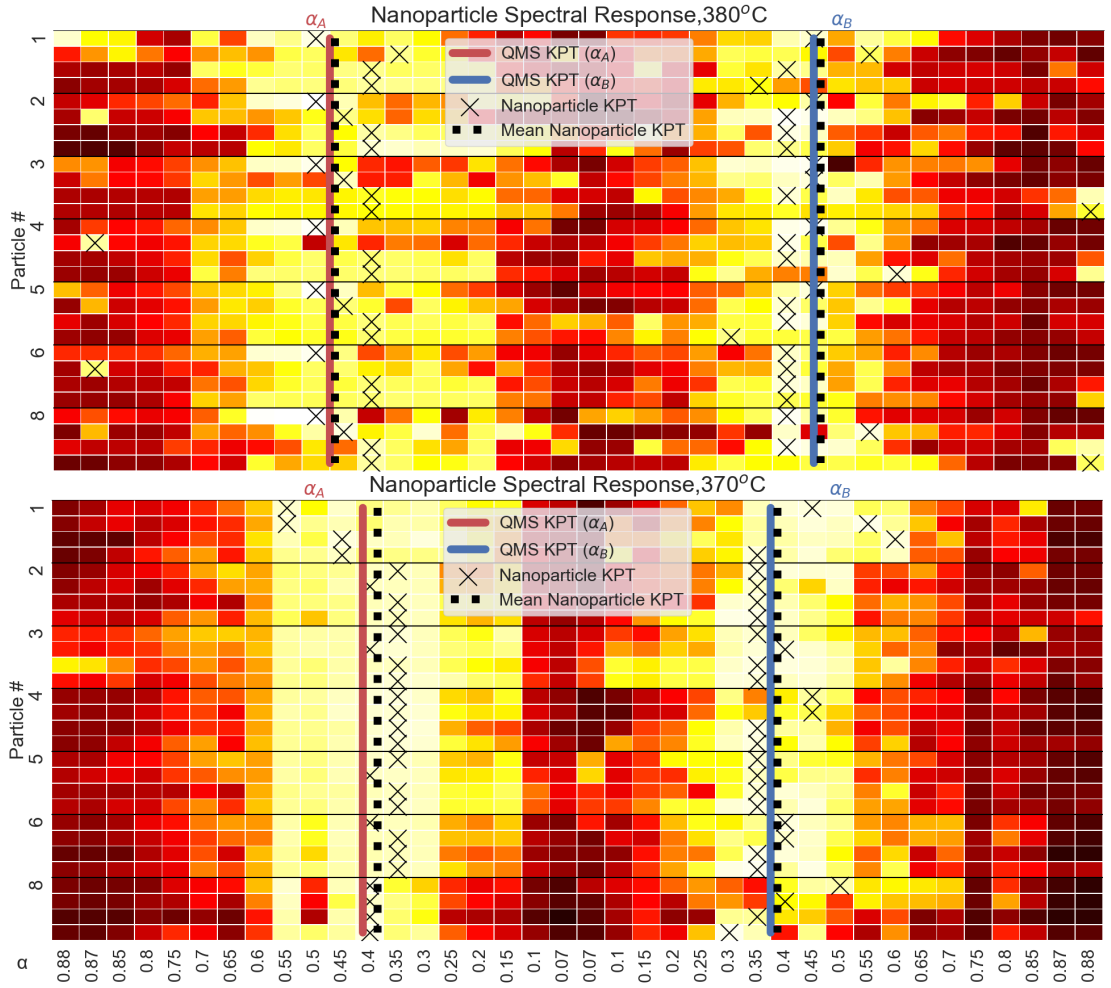

Figure 4.13: Heatmaps drawn from results of conducting a noise-canceling measurement at each temperature, where the experiment is run with only argon gas. The results after subtracting this noise-canceling measurement from all other measurements are certainly much cleaner, but it is unclear whether this truly is indicative of better results.

5

Conclusions

In this thesis, we have attempted to develop a method which correlates optical data to QMS data, for the purpose of identifying kinetic phase transitions on the surfaces of individual nanoparticles, and as a step in the greater goal of deriving general structure-function correlations between optical readout and surface kinetics of said nanoparticles. To achieve this, we used two approaches; a recurrent neural network trained on the change in FWHM of the optical spectra over time, and a deep densely connected neural network which correlates a nanoparticle's optical spectra to its corresponding QMS behaviour.

Without a proper ground truth to compare with, it is impossible to definitively determine whether the predictions of individual nanoparticles' KPT are accurate. However, by cross-validating predictions across samples (nano-fluidic chips), we show that the network is capable of learning generalized information about the nanoplasmonic behaviour of each nanoparticle and correlate this to fully reasonable QMS activity. Indeed, we also show that this is true despite the behaviour of the second sample of nanoparticles being considerably different both in terms of optical and QMS response compared to the sample that was trained on.

Several things can still be implemented to improve the predictions of the network, i.e. training on more data (including all temperatures at once), including tensors with more observables or relevant statistical variables. Another possibility is attempting more "aggressive" noise-canceling methods, such as fitting Gaussian distributions to each timestep of the optical spectra or using noise-canceling measurements, as briefly described in the end of section 4.5, in more rigorous ways. Indeed, a thorough investigation in the noise profile of the new sample should be conducted, such that we can fully isolate and remedy the issues that lead to clear mispredictions of the QMS activity in low α for the O-poisoned region of the experiment, as exemplified in the bottom image of figure 4.8.

As for the future of the project and the practical implementation of these results, there are ultimately two distinct paths depending on what the main focus is determined to be, and how confident we are in the accuracy of the predicted positions of kinetic phase transitions.

One option, if we do not feel confident, is that we spend more effort investigating the fundamental underlying physics behind the optical spectra within the kinetic

phase transition. In this case, we need to further analyse the results of the network and inversely investigate which parts of the optical spectra correlate to which QMS activity according to the network and why. This would effectively translate the goal of the analysis into interpreting the particular function of the trained network, in order to understand which aspects within the optical spectra correlate to a kinetic phase transition.

Another possibility is that we are satisfied with the network's performance and want to use its KPT predictions within chemical physics. The specifics of this approach should, of course, be decided in co-ordination with experts in that field, but I propose that we analyze the nanoparticles' surface behaviour based on their predicted QMS activity. The benefit of this is of course because plenty of theory already exists on this front - we can develop phase diagrams like figures 2.8 and 4.2 for each individual nanoparticle given only its optical spectra, then cross-examine this with surface phase diagrams like figure 2.1 to see which state the surface of each individual nanoparticle is in in any given moment.

Regardless of what intriguing future awaits for this project, we look at the results of this thesis with cautious optimism, and hope we might develop it further to unlock the true potential of this novel and ambitious approach to heterogeneous catalysis at the nano-scale.

Bibliography

[1] Mind the gap. *Nature Catalysis*, 1(11):807–808, Nov 2018.

[2] David Albinsson. Combining nanoplasmonics and nanofluidics for single particle catalysis. November 2018.

[3] Svetlana Alekseeva, Ievgen I. Nedrygailov, and Christoph Langhammer. Single particle plasmonics for materials science and single particle catalysis. *ACS Photonics*, 6(6):1319–1330, Jun 2019.

[4] S. W. Chee, J. M. Arce-Ramos, W. Li, A. Genest, and U. Mirsaidov. Structural changes in noble metal nanoparticles during CO oxidation and their impact on catalyst activity. *Nat Commun*, 11(1):2133, May 2020.

[5] François Chollet. *Deep Learning with Python*. Manning, November 2017.

[6] Andreas B. Dahlin, Jonas O. Tegenfeldt, and Fredrik Höök. Improving the instrumental resolution of sensors based on localized surface plasmon resonance. *Analytical Chemistry*, 78(13):4416–4423, 2006. PMID: 16808449.

[7] Vincent Stroobant Edmond de Hoffmann. *Mass Spectrometry: Principles and Applications*. WILEY, 2004.

[8] Gerhard Ertl. Reactions at surfaces: From atoms to complexity, 2007.

[9] Carl A. Busacca et. al. Transition metal catalysis in the pharmaceutical industry. 2005.

[10] Susie Eustis and Mostafa A. El-Sayed. Why gold nanoparticles are more precious than pretty gold: Noble metal surface plasmon resonance and its enhancement of the radiative and nonradiative properties of nanocrystals of different shapes. *Chem. Soc. Rev.*, 35:209–217, 2006.

[11] Saga Helgadottir, Aykut Argun, and Giovanni Volpe. Digital video microscopy enhanced by deep learning. *Optica*, 6(4):506–513, Apr.

[12] Richard K. Herz and Samuel P. Marin. Surface chemistry models of carbon monoxide oxidation on supported platinum catalysts. *Journal of Catalysis*, 65(2):281 – 296, 1980.

[13] Christoph Langhammer and Elin M. Larsson. Nanoplasmonic in situ spectroscopy for catalysis applications. *ACS Catalysis*, 2(9):2036–2045, 2012.

[14] Florian Meirer and Bert M. Weckhuysen. Spatial and temporal exploration of heterogeneous catalysts with synchrotron radiation. *Nature Reviews Materials*, 3(9):324–340, Sep 2018.

[15] Carolina Novo, Alison M. Funston, and Paul Mulvaney. Direct observation of chemical reactions on single gold nanocrystals using surface plasmon spectroscopy. *Nature Nanotechnology*, 3(10):598–602, Oct 2008.

[16] Daniel V. Schroeder. *An introduction to thermal physics*. San Francisco, CA : Addison Wesley, [2000] ©2000, [2000]. Includes bibliographical references (pages 397-405) and index.

[17] R. van Rijn, O. Balmes, A. Resta, D. Wermeille, R. Westerström, J. Gustafson, R. Felici, E. Lundgren, and J. W. M. Frenken. Surface structure and reactivity of pd(100) during co oxidation near ambient pressures. *Phys. Chem. Chem. Phys.*, 13:13167–13171, 2011.

[18] Matthijs A. van Spronsen, Joost W. M. Frenken, and Irene M. N. Groot. Surface science under reaction conditions: Co oxidation on pt and pd model catalysts. *Chem. Soc. Rev.*, 46:4347–4374, 2017.

A

Appendix

A.1 Implementation of the Deep Densely Connected Neural Network

This brief section is dedicated to describing the implementation of the neural network. This is intentionally kept brief and separate because I will stress that the specifics of the deep neural network's implementation are deceptively uninmportant - many different realizations of the network will yield similar results. Indeed, the specifics of data preprocessing, data acquisition and training method have been shown to influence the results far more, and are therefore given more attention in the main report. However, since some readers will invariably and rightfully wonder about the specifics of our method, I recognize that it is still worth including. I will also stress that this section should only be read if the reader has some prerequisite knowledge of the specifics of neural network implementations, since the underlying theory would be much too involved to include here. For a personal recommendation, I will refer to *Deep Learning With Python* [5] by François Chollet.

The deep densely connected neural network is implemented in Keras and composed of five layers, the first four of which are dense layers with 16 neurons each with ReLU activations functions. The fifth and last layer is a dense layer of a single neuron, using a sigmoid activation function to force values into the [0,1] range. The specifics of these numbers and activation functions are, as stated, not terribly important - they are found mostly through trial and error (along with a few basic assumptions, such as keeping hidden unit counts low and the overall efficiency and favourability of ReLU functions within hidden layers), and tweaking the parameters in subtle ways has negligible effects on the overall results.

The following hyperparameters were used for training the network:

- Loss function: Mean squared error
- Optimizer: SGD, learning rate: 0.01
- Batch size: 10

As mentioned before, holdout validation and "Leave-1-Out" cross validation is used. With that said, training the network occurs in two main steps - firstly, the optical

data is trained on in a network with the parameters given above, with a patience of 25. If the network's validation loss is not reduced to below 0.3 within this time, the network cancels its training and resets. The reason for this is because the network has a tendency to get stuck in local optima in this early stage of training, and neither regularization nor dropout remedies the issue. Fairly often, however, the network "breaks through" the local optimum and its validation loss decreases significantly, often to below 0.2. Once the network reaches this optimum, which is likely near the global optimum, the network's parameters begin to change slightly in each new epoch where nothing is learnt - the learning rate and batch size dynamically begin to decrease with each passing epoch as training continues. This is to allow the network to really edge out the last few correlations which might be difficult to reach with a high batch size and learning rate (due to aggressive "movements" within stochastic gradient descent). This process effectively makes it unlikely that the network gets stuck in local optima, but the same result can likely also be reached with some hitherto unattempted combination of regularizers or other methods of steering parameter change within the network's training process.

A.2 Individual Nanoparticle Kinetic Phase Transitions

Heatmaps showing the distribution of predicted qualitative QMS behaviour for each individual nanoparticle as a function of concentration coefficient α. Bright whites and yellows corresponds to a high QMS response, indicating high surface catalytic reactivity, and conversely darker shades of red correspond to low QMS response. Each individual nanoparticle's kinetic phase transition (KPT) is calculated to occur at the α-value of its highest QMS response, which is represented by a black cross. The mean α-value for which individual KPTs occur is represented by a black dashed line, and the mean peaks of the actual QMS signal are shown in red and blue for the first and second regions of the QMS, respectively.

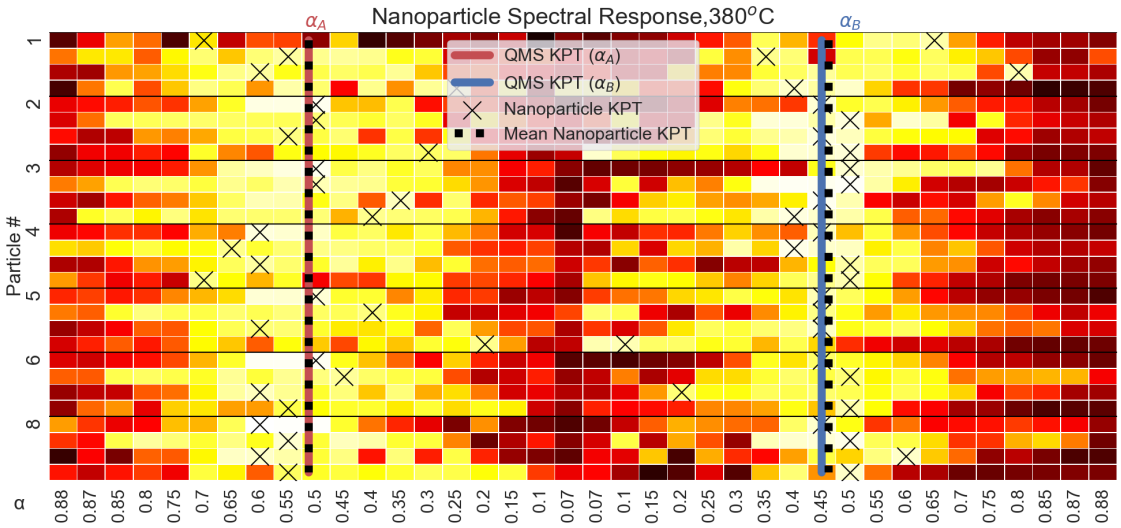

Figure A.1: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 380^{o}C.

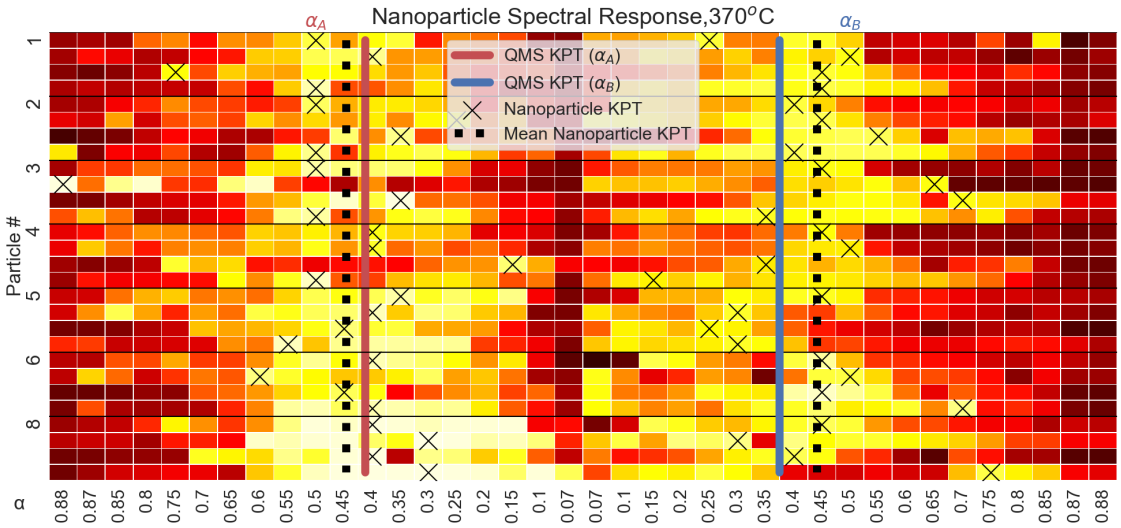

Figure A.2: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 370^{o}C.

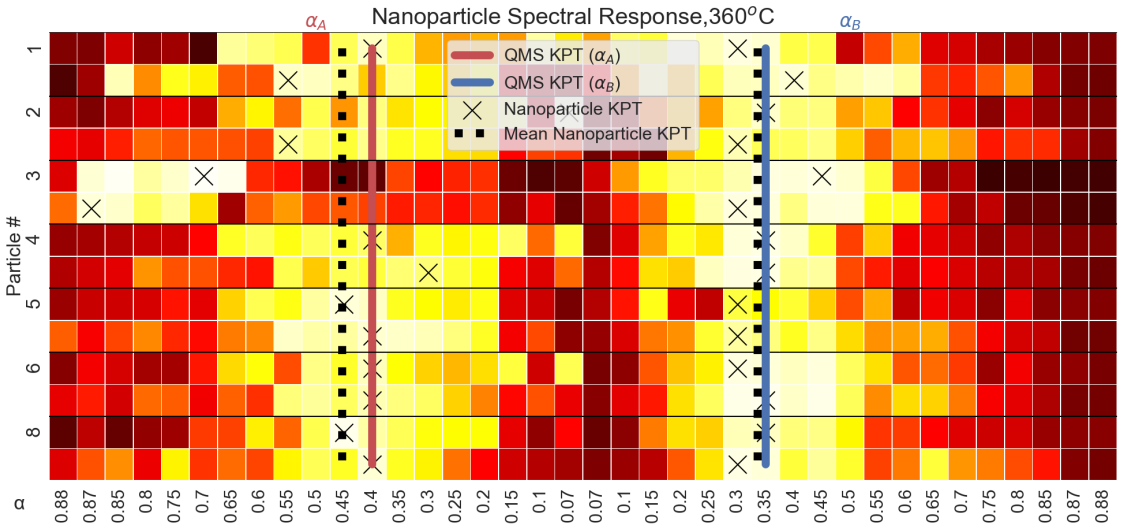

Figure A.3: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 360^{o}C.

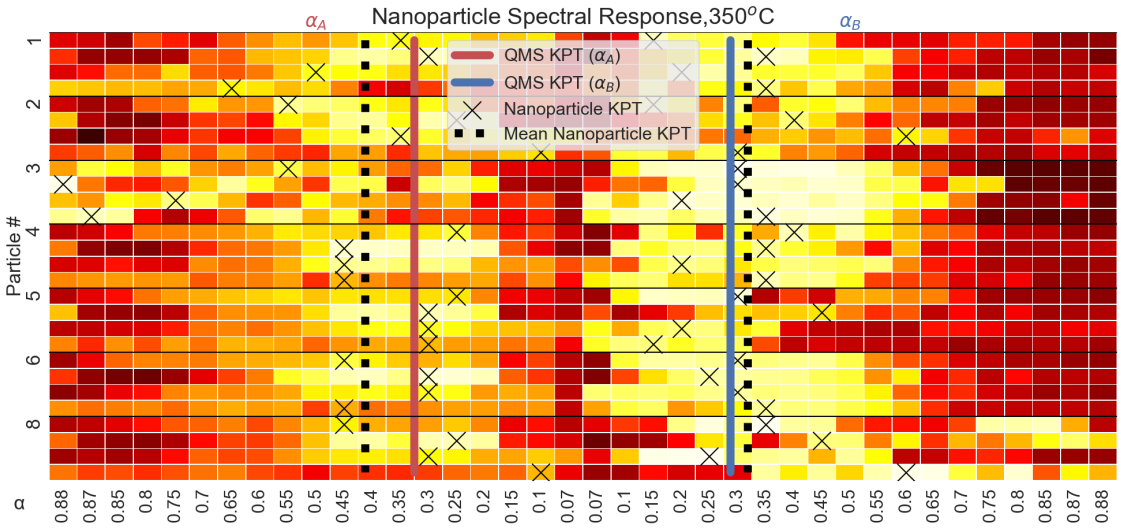

Figure A.4: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 350^{o}C.

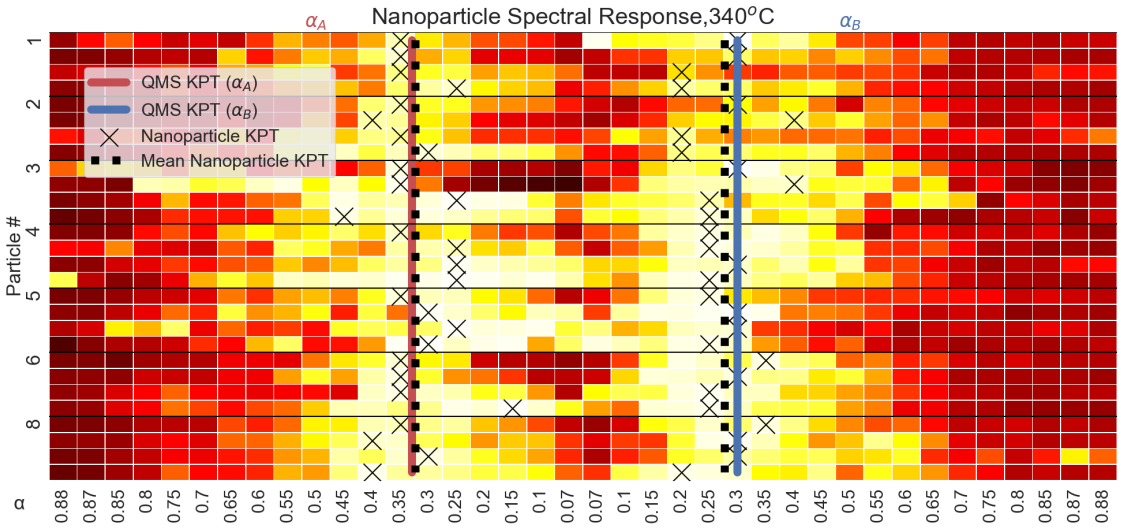

Figure A.5: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 340^{o}C.

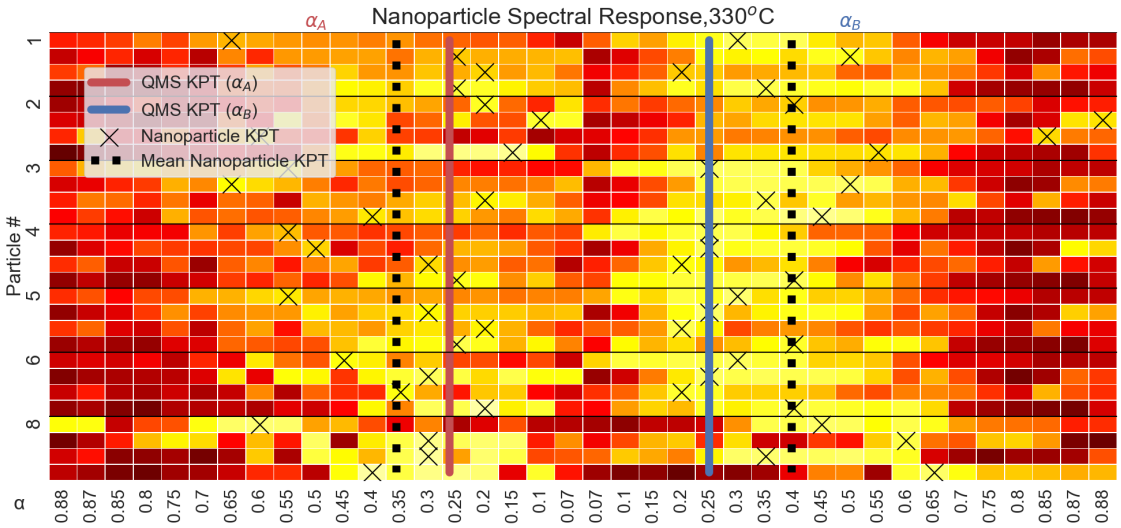

Figure A.6: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 330^{o}C.

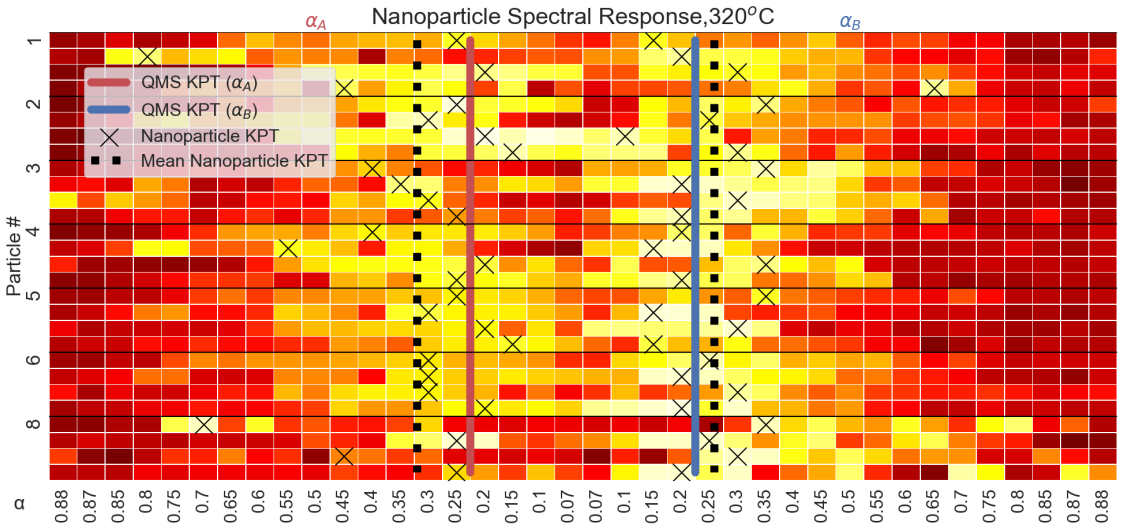

Figure A.7: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 320^{o}C.

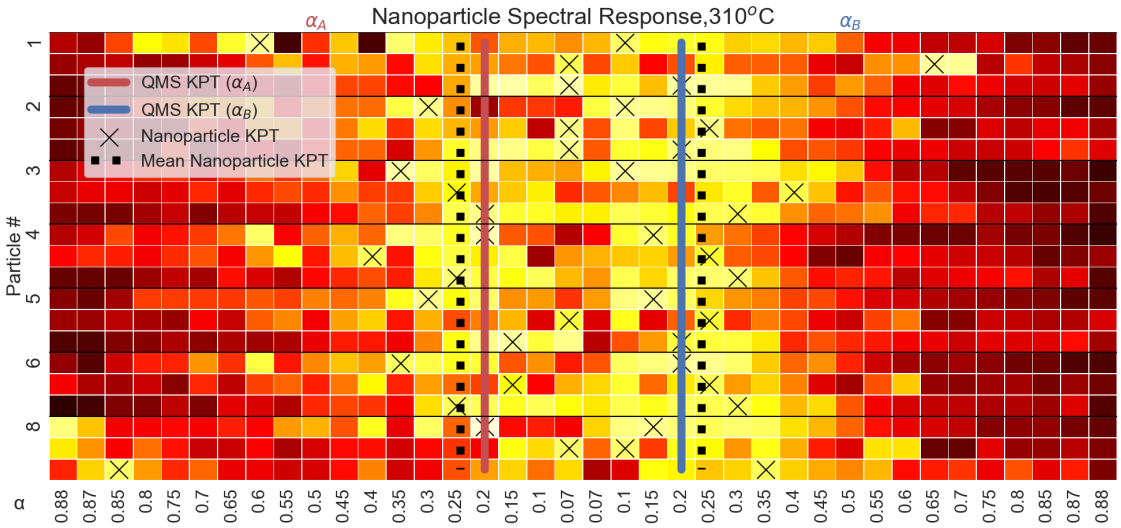

Figure A.8: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 310^{o}C.

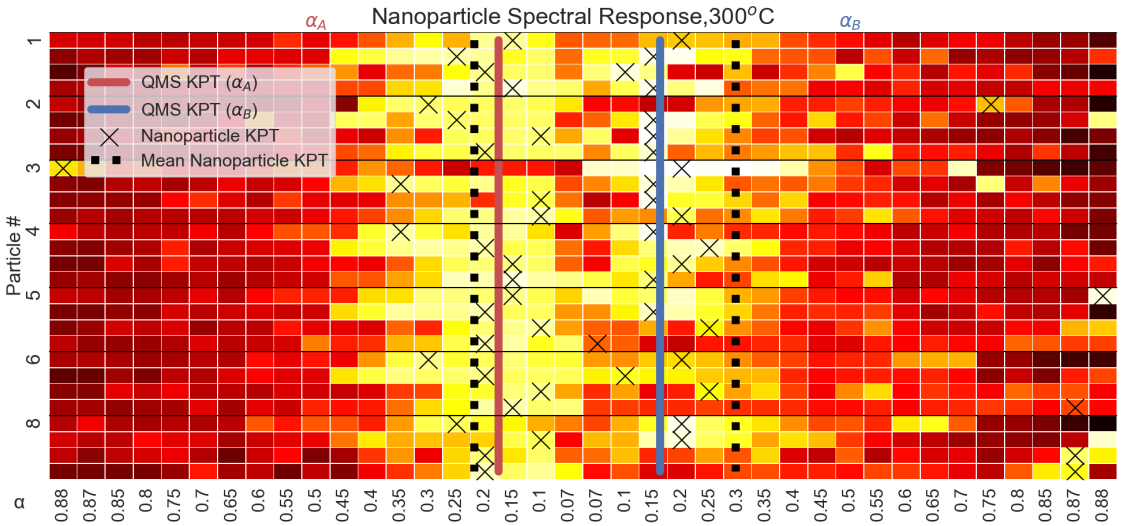

Figure A.9: Heatmap showing the distribution of kinetic phase transitions on individual nanoparticles for T= 300^{o}C.

A.3 FWHM - LSTM Network Approach

In figure A.10 we study the FWHM signal for one of the particles for one particular measurement. Here, we also plot the aggregated QMS response signal for all of the particles present in the chip for this measurement. We have also drawn vertical lines to indicate where the QMS response corresponds to the largest activity of the particles. Between these lines, the FWHM curve appears to exhibit an oscillatory behaviour compared to the signal outside of this interval. This observation suggests that there might be a correlation between the region between the QMS peaks and the oscillatory region of the FWHM signal.

With the above observation in mind, we propose that an alternative to finding the phase transitions of individual particles with our dense neural network is to perform a *switching point analysis* on the FWHM of the spectra generated for each particle. The idea is that at two individual times during the measurement the FWHM of the Gaussian spectra for each particles changes makes a transition from being mostly noise, to clearly oscillatory. We hypothesize that this transition occurs roughly simultaneously as the particles undergo phase transition between different states. Therefore, if one is able to find this switching point in the FWHM signal, that might give an indidication as to where each particle undergoes phase transition and thus making it possible to characterize each individual nanoparticle from their nanoplasmonic spectra alone. This method might then serve as one form of a cross validation for the main method of the dense deep neural network that uses the entire spectra (and not only the FWHM readout) in order to make similar predictions.

Furthermore, we wish to find an automatic way to determine where the transition takes place (if it exists at all). The method that we used is based on an *LSTM* (Long Short Term Memory) neural network, which will be described in this appendix. Before we go into the details of the actual network, we will describe how the FWHM signal was created and how we used it in order to make phase transition predictions.

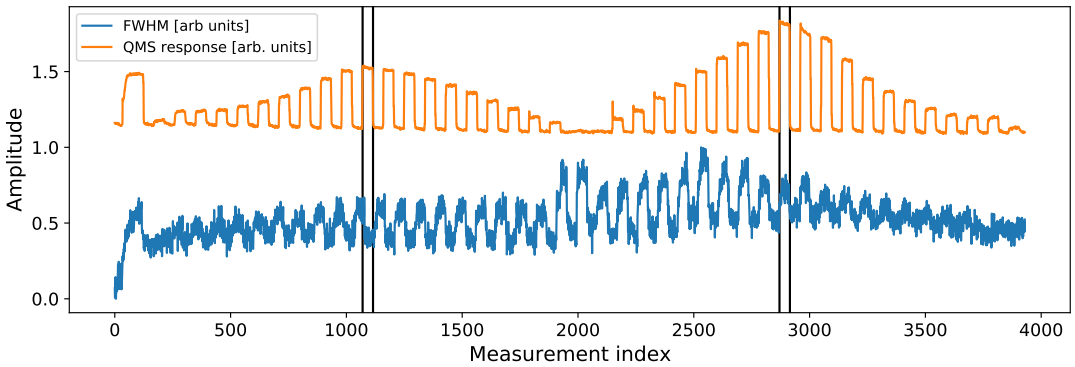

Figure A.10: Aggregated QMS response signal from all particles during one measurement and FWHM signal for one of the particles. The vertical lines indicate the pulses where the QMS response is the greatest.

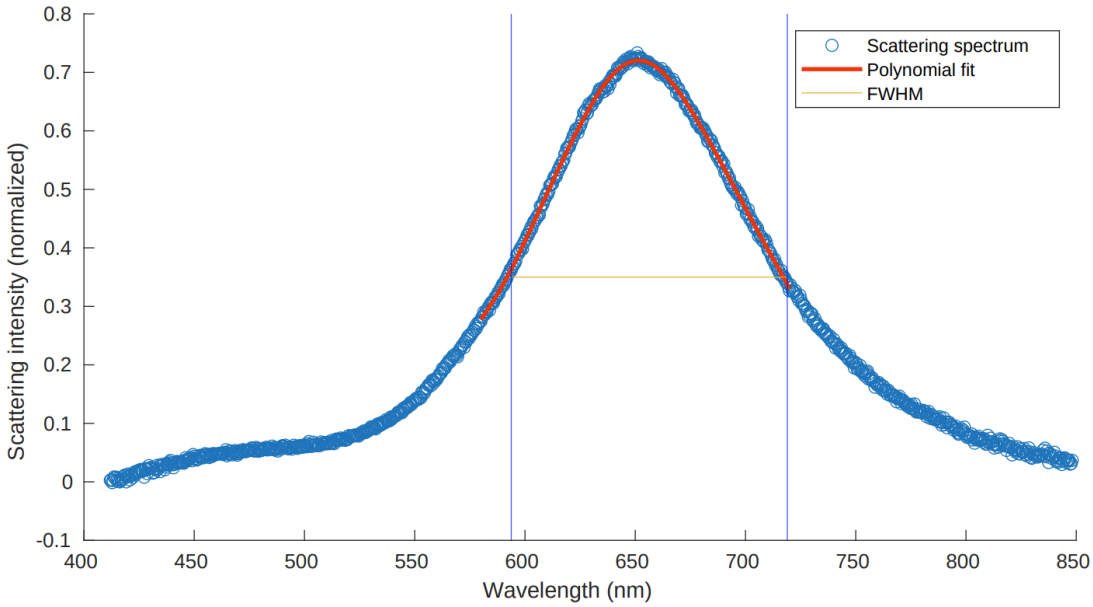

Figure A.11: Polynomial fit to a scattering spectrum and corresponding FWHM.

A.3.1 Extracting FWHM from the spectra

For each measurement made on one temperature on one chip, the average scattering spectrum for each 20 second interval is recorded for each particle in the chip that does not belong to the middle patch. The middle section of these spectra are roughly Gaussian. Therefore, if one can fit a Gaussian curve to the spectra, see figure A.11, one can record the *Full Width at Half Maximum* (FWHM) of each individual spectrum at every time step that is recorded. The FWHM will then change during the entire measurement to produce a time series that can look for example like the blue curve in figure A.10.

Currently, the FWHM signal is generated with a Matlab-script with the following algorithm.

1. A guess of the wavelength corresponding to the peak of the scattering spectrum is provided to the program. By stuyding the spectra it became obvious that, for each spectrum, the peak is always close to around 650 nm, so that is the initial guess.

2. A polynomial of degree 10 is fitted to the scattering spectrum in a manually selected wavelength range around the peak guess. We used a range of ± 50 nm for the fit.

3. The FWHM is then calculated for the polynomial fit as the horisontal distance between the two points where the polynomial amplitude is half the peak amplitude.

A.3.2 Network architecture and hyperparameters

In order to find the hypothesized switching points of the FWHM signal, the idea was to treat the FWHM signal as a time series and to then use an LSTM (*Long Short Term Memory*) network to classify each time step as either oscillatory or non-oscillatory, thus converting the task into a binary classification problem for each time step of the FWHM signal.

An LSTM network, in contrast to a dense NN, is well suited to make time series predictions because it uses previous data points of a time series in order to make predictions on subsequent time steps. First, the entire time series is fed into the LSTM layer, after which each time step is processed chronologically from left to right with LSTM cells, see figure A.12 for the cell architecture. In this figure, x_t is time step t of the entire signal, where $t \in [1, L]$, where L is the length of the entire time series. h_{t-1} is a hidden state containing information about the previous input x_{t-1} and c_{t-1} is a cell input activation vector. Once all the hidden states have been calculated for the entire time series (for all layers), the weights will be updated for each layer with backpropagation.

Since the oscillations we are trying to identify (with a period of roughly 100 time steps) are a local behaviour with respect to the entire time series length (roughly 4000 time steps), the LSTM is particularly useful for this kind of oscillation detection since it contains *forget*-gates, meaning that it is concerned with what happened recently in time (hence denoted as long *short term* memory) without being concerned with time steps occuring far earlier in time. An LSTM also has the benefit of counteracting the vanishing/exploding gradient problem that "vanilla" RNN often suffer from.

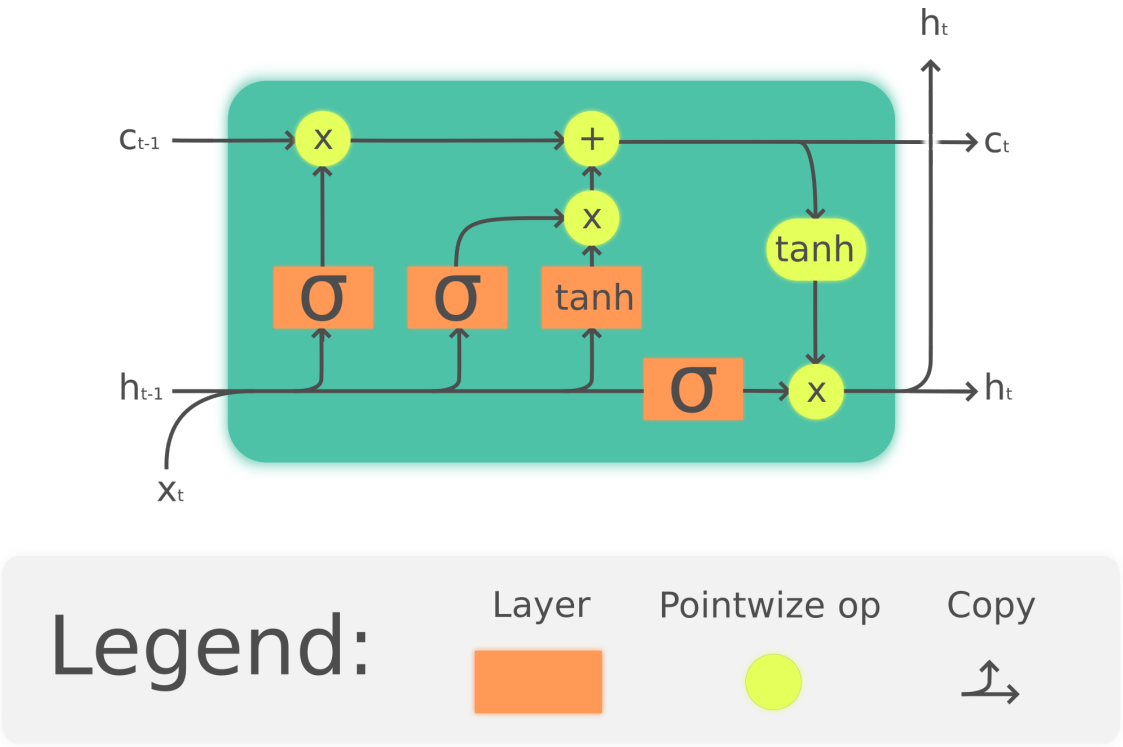

Figure A.12: An LSTM cell. Source: Long Short Term Memory, Wikipedia, available at `https://en.wikipedia.org/wiki/Long_short-term_memory` (Licensed under the Creative Commons Attribution 4.0 International license)

The neural network we used was composed of a two layer network, implemented in Keras, with the first layer being an LSTM with the output dimensionality of each hidden state h_t set to 16. Into this network, an entire FWHM-signal is input at a time. The output layer is a dense layer with a sigmoid output to convert the sequences from the LSTM into a classification accuracy for each time step.

The following hyperparameters were used for training the network:

- Loss function: Binary crossentropy
- Optimizer: Adam, learning rate: 0.01-0.1
- Batch size: 50

We then generated labeled training data for the network, but before we explain the generating procedure, we first present the preprocessing of the FWHM signals.

A.3.3 Preprocessing of the FWHM

The shape of one of the FWHM signals can be seen in figure A.10. From studying several of these signals, it became clear that the overall shape of these curves did not seem to have any significant correlation between different particles and/or

measurements. Therefore, our preprocessing approach includes removing a low frequency trend from the FWHM signal in order to ”stabilize” the curve. In this way, it's possible to let the network to primarily focus on the oscillatory behaviour of the signal, and not the overarching shape of the curve. This behaviour is also more difficult to replicate since it varies between different particles and experiments. Furthermore, since we hypothesize that the switching point can be localized by means of studying only the oscillatory behaviour of the signal, the loss of the overarching shape should not greatly affect the result of this particular network (even though we do not disaffirm the usefulness of using a less preprocessed FWHM signal for making predictions with another method). See figure A.13 for a comparison of the FWHM signal before and after subtraction of the low frequency trend. The exact preprocessing procedure is described below.

1. Remove outliers from the signal by setting a tolerance value that the signal must not overshoot and replace the outlier by the tolerance ceiling (or floor).
2. Subtract a low frequency trend (with a period of $T = 100$ time steps) from the entire signal.
3. Normalize the signal and shift the curve to mean 0.5.

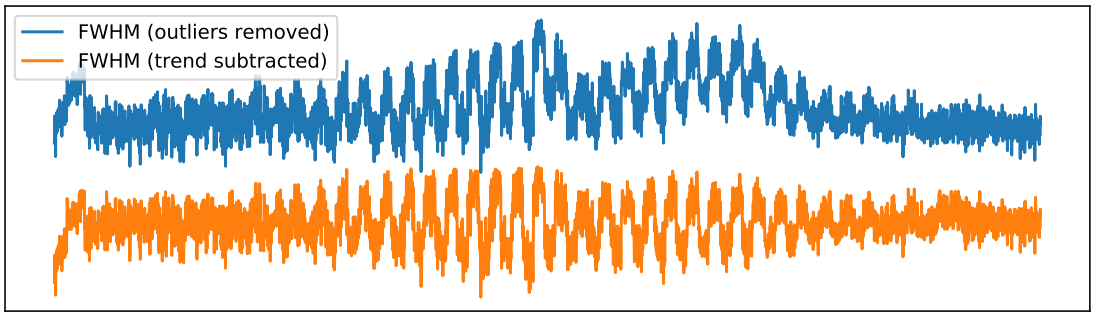

Figure A.13: FWHM signals before and after subtraction of a low frequency trend with a period of $T = 100$ time steps.

A.3.4 Data generation

We generated unique training samples in order to replicate both the oscillations and the noise of the FWHM signals resemble the preprocessed signal in the final form of the FWHM, visualized by the orange curve in figure A.13. Then, by assigning binary labels to this data, indicating whether it's oscillatory or noise, the network should be able to generalize to finding these oscillations in the real FWHM signals. It should be noted that since we hypothesize that the switching point might occur as the FWHM signal trancends from noisy to ”clearly oscillatory” the results of the network will (naturally) largely depend on what type of oscillations we present to the network. Since it is unclear at this point what exact underlying mechanism causes this behaviour, we try to *emulate* the data instead of trying to *simulate* it. Hence, we will try to replicate the FWHM signal purely by mathematical means

with no physics involved.

In figure A.14 we show ten generated random samples and their corresponding labels. The orange boxes in the figures indicate where the signal is labelled as oscillatory.

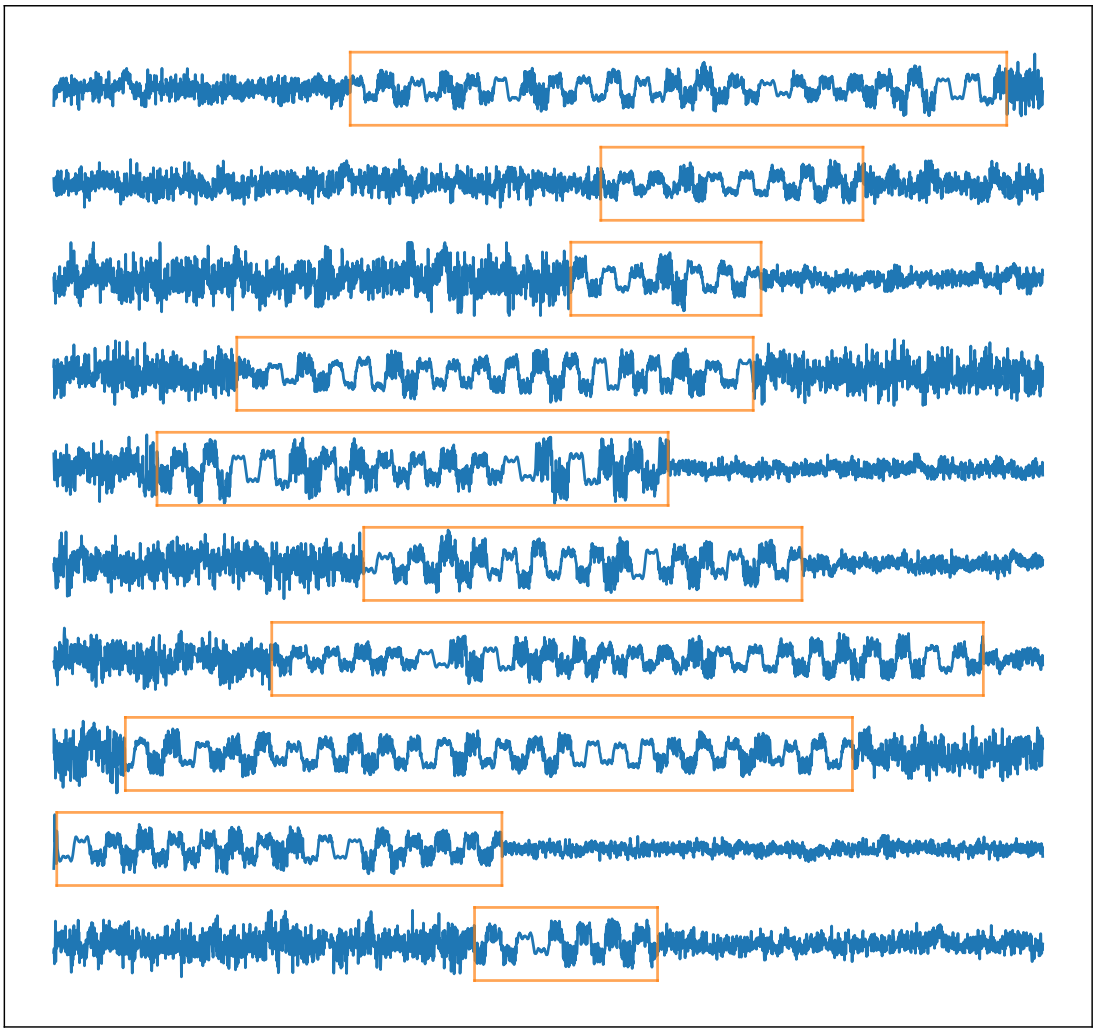

Figure A.14: Generated training samples and orange boxes indicating where each sample is labelled as oscillatory.

Since the FWHM signal is bistable where it oscillates, it seemed most suitable to generate an oscillation patterns formed by an approximation of a square wave (with stochastic deviations from this approximation in terms of weights r_1, r_2, r_3 in as well as by a uniform noise overlay). Let r_i denote uniform random numbers between 0 and 1 and R be an array of uniform random numbers in the same range. The formula used for generating these oscillations $f(t)$ uses the first three terms in the

fourier expansion of a square wave according to the function

$$f(t) = [(1 - 0.5r_1)\sin(t) + (1 - 0.5r_2)\frac{1}{3}\sin(3t) + (1 - 0.5r_3)\frac{1}{5}\sin(5t)] + 2R. \quad \text{(A.1)}$$

The exact formula was conceived by trial and error in an effort to accurately attempt to replicate different kinds of oscillation patterns occuring in the FWHM signals.

Furthermore, from studying the FWHM signal, it became clear that the oscillations occur in periods of approximately 90 time steps. Therefore, the period of generated oscillations were set to $80 + 20 \cdot rand$, where *rand* is a uniform random variable in the range $[0, 1]$. The noisy segments of the training samples were based on gas-free measurements from $T = 380°\text{C}$, which should consist exclusively of noise. In order for the network not to overtrain on these samples (consisting of approximately 45 000 datapoints in total), these were randomly shuffled for each training set in intervals of 50 time steps.

The training samples were further constructed such that, for each sample, the oscillations would be sustained in a uniform random range between $[500, 2500]$ data points, whereas the entire input signal to the network was fixed at 3000 data points, so that on average, the oscillations would be sustained for half of each training sample, but with a stochastic variable determining the length of each individual oscillation interval. The starting point of the oscillations was then based on $i_{\text{start}} = rand(1, 3000 - L)$, where L is the length of the oscillation interval. Two purely noisy segments would then fill in the gaps around the oscillation interval (see figure A.14). Conclusively, randomizing both the length and the start point of the oscillations should make sure that the network does not overtrain on a particular starting and ending point, thus being able to predict general oscillatory behaviour in the FWHM signals.

Finally, the amplitude of the oscillation segments was normalized between $[0, 1]$ and then stochastically scaled between $[0.75, 1]$ before having its mean set to 0.5 (the same as for the preprocessing of the FWHM signals). The noisy signals were scaled between $[0.5, 1]$ and their means were also set to 0.5. The different scalings of the oscillatory region and the noise was made to ensure successful training of the network. Setting them to the same amplitude made it very difficult for the network to learn (if at all). Indeed, scaling the segments differently might seem to make the network biased towards recognizing amplitude as the key factor determening the labels, but as made clear from the predictions, this was not the case.

A.3.5 Training procedure and history

In figures A.15 and A.16, we show the training history for the final network, which reached a BCE-loss of 0.052. Note the different scales on the y-axis in the two figures. The network was trained on 1000 new samples each epoch and was validated on 100 samples (that do not change during training). The low number of validation samples was chosen in order to speed up the training, which is why the validation loss is much more volatile than the training loss, as can be seen by the validation

loss spikes in both figures. Since new training data is generated each epoch, the training loss is a form of validation in and of itself, and the validation loss serves more as a sanity check, which is why we deem it acceptable with a low amount of validation samples. Furthermore, the network was trained for over 3600 epochs, but due to the volatile nature of randomized data, the training was reset to the last best checkpoint in the case that the training loss was not improved for at least 2 epochs. Therefore, the final network was trained on approximately 2-3 million samples (we do not have the exact figure). After the loss reached 0.10, the learning rate was reduced from 0.1 to 0.01, which is why the training history was split into two figures. Also, after the learning rate was reduced, the training was terminated when there was no signficant improvement in validation accuracy for a couple of hundred epochs. The large number of epochs needed to train the network is in large part due to noisy oscillations being quite similar to the noise in the training data, which makes it difficult for the network to quickly distinguish between oscillations and noise. If the possible noise amplitude present in the oscillations of the training data is lowered slightly, the training time is signficantly reduced, but in turn, the network will not recognize noisy oscillations to the same extent.

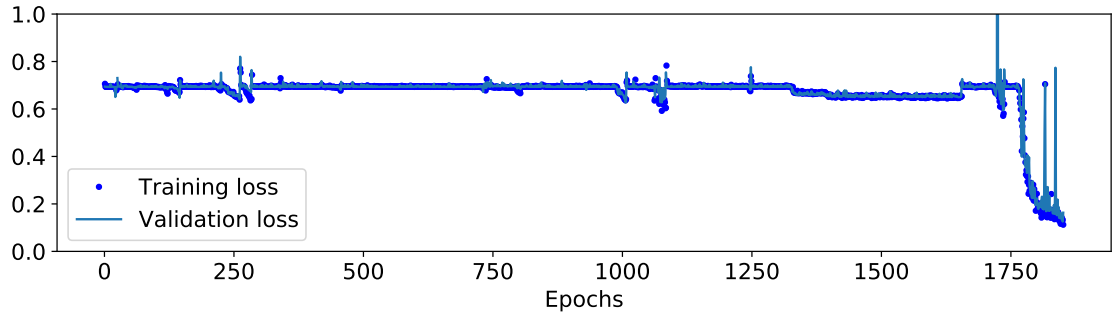

Figure A.15: First part of training history with a learning rate of 0.1.

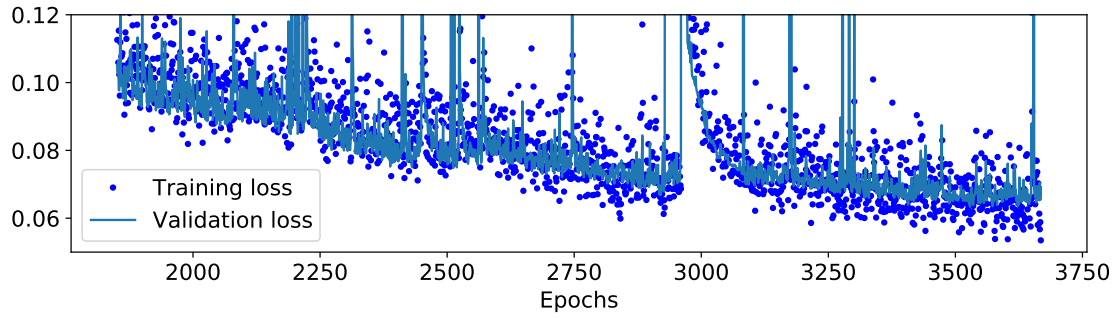

Figure A.16: Second part of training history with a learning rate of 0.01.

A.3.6 Predictions of Chip E

In what follows, we present the predictions results from the trained network. We start by showing the quantitative predictions on training samples that the network

has never seen before, to visualize how the switching point might be determined from the binary predictions. Then, we show predictions on FWHM signals that were extracted from several different measurements that were collected over the span of more than three months.

A.3.6.1 Training predictions

After training, we made predictions on newly generated data to visualize the networks performance on the training samples. These results can be seen in figure A.17. From this figure, we see that the predictions consequently align with the labels of the data. Though in some cases, the network might be delayed by around 50-100 measurements in finding the oscillatory region, both at the start and at the end of this region. This is important to keep in mind regarding the actual accuracy of the real predictions. We also note that once the oscillations have been identified, the prediction might flutter somewhat during the oscillatory region. Since it is clear from the predictions where the oscillations begin and end however, this should not be a problem. It does, however, signify that the training data (or the learning) could be improved, since the network is not perfectly certain as to what kind of signal should be labelled as an oscillation or as noise.

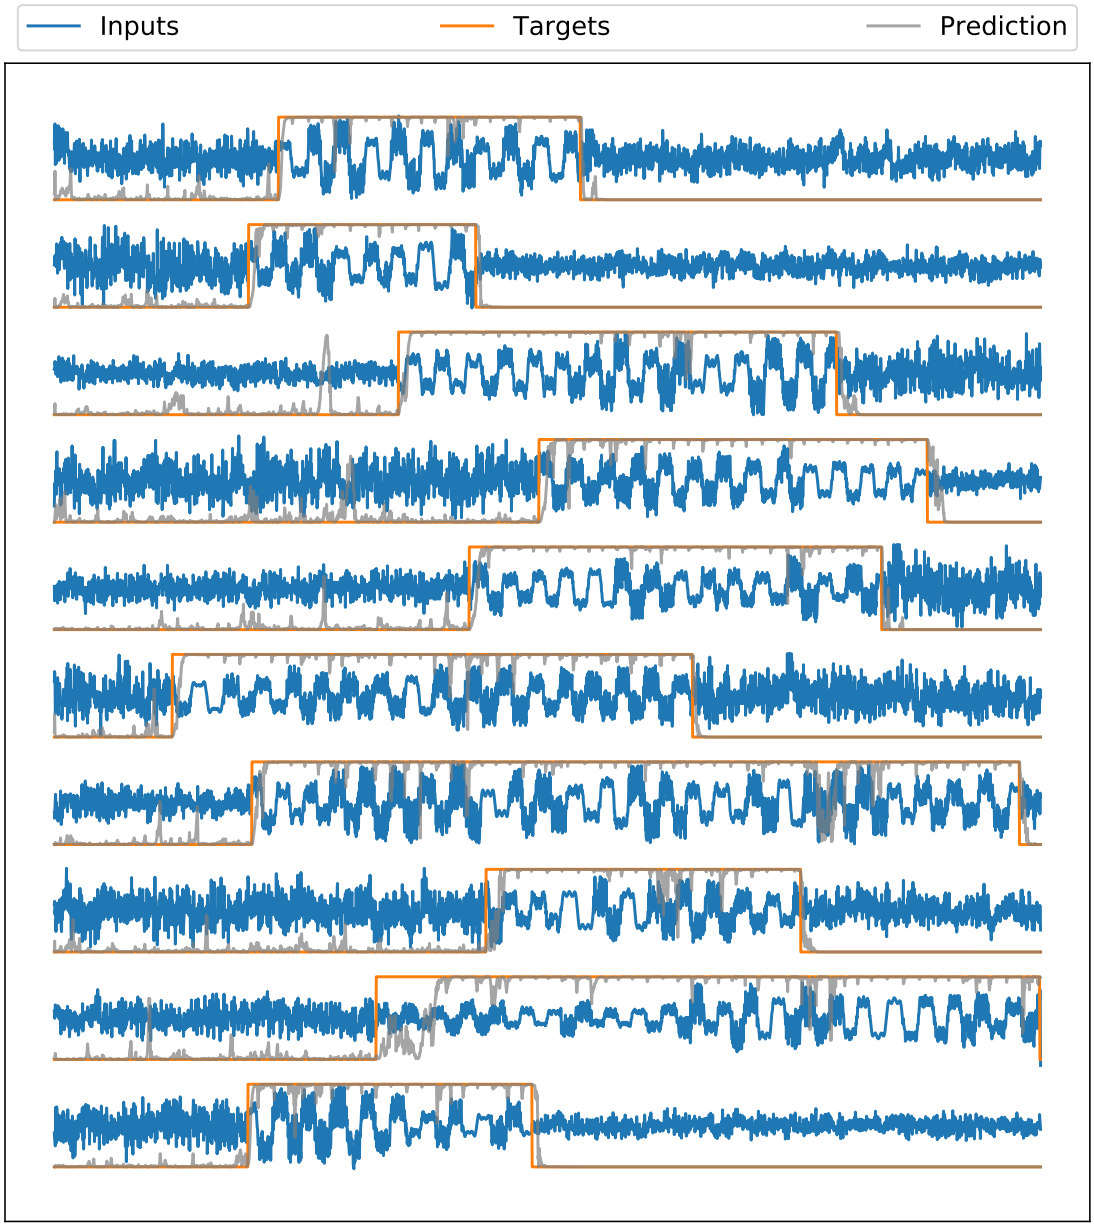

Figure A.17: Training predictions

A.3.6.2 General results

Before presenting the individual results for each particle, we show a plot in figure A.18 where the mean switching point for each measurement was found and compared to the mean QMS peaks for each temperature. In this figure, we identified the switching points for two measurements for each particle and present the average of the two switching points for temperatures between 300 and 380°C. The individual predictions of these measurements are then presented in the sections succeeding this one. We observe that the switching points are clearly shifted towards higher α-values compared to the QMS peaks, but they follow roughly the same trend of

reaching lower alpha values for lower temperatures. Studying for instance the predictions in figure A.20 we see that, for some measurements, there is a trend of the switching points occuring for larger α-values for particles further farther away from the gas inlet. Additionally, the switching points for particles 1-4 were excluded more frequently than for particles 5-8 due to that the FWHM signal of particles 1-4 is in general more noisy than for particles 5-8 (thus making the switching points more difficult to identify). This is a possible explanation for the shift of the average switching points compared to the QMS peaks.

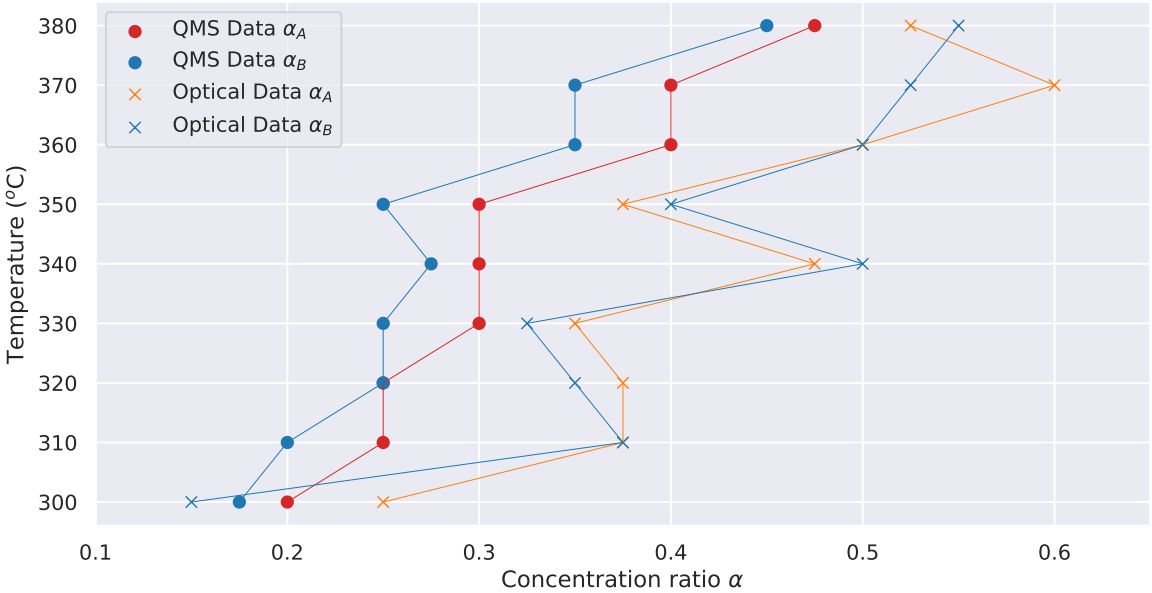

Figure A.18: Average switching points (optical data) of particles 1-8 compared to the corresponding average QMS peaks for measurements made with temperatures between 300 and 380°C. For each temperature, two measurements were used to calculated the QMS peaks and the average switching points.

With the above reasoning in mind, we also present the average switching points for the same measurements for only particle 5 in figure A.19. Here, we see that the predicted oscillatory region has a higher correlation with the QMS peaks for some individual particles. By studying the individual results, it is clear that the FWHM signal for each particle can differ quite substantially even within the same measurement. This result points toward that the LSTM network might be useful for some individual particles, but the method needs to be improved to accurately account for each particle within a measurement.

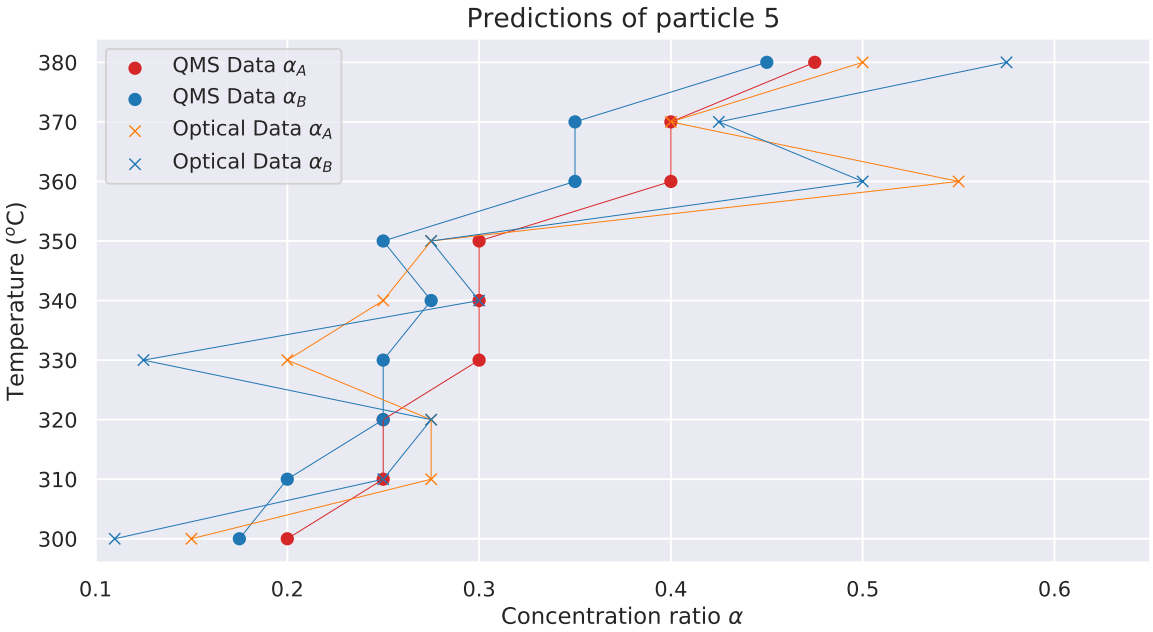

Figure A.19: Average switching points (optical data) for only particle 5 compared to the corresponding average QMS peaks for measurements made with temperatures between 300 and 380 °C. For each temperature, two measurements were used to calculated the QMS peaks and the average switching points.

In the following sections we show the predictions of individual FWHM signals that were collected for measurements collected at temperatures between 300 and 380°C. As previously mentioned, we attempt to identify one oscillatory region for each particle by providing switching points based on the predictions of the network. The switching points are identified as the time step where the moving average (with a window of 50 time steps) of the predictions reach a cut off value of 0.8. In the cases where the predicted oscillatory period of the signal is sustained for less than one half of the signal between the QMS peaks, no switching points will be identified, even though the average of the signal might reach 0.8. The reason for this cutoff is that for particles that do show clear oscillatory behaviour (by manual inspection) for high temperatues, the predicted oscillatory period is generally sustained for at least 1000 time steps. Therefore, claiming to have found a switching point when the predicted oscillatory period is too short (for instance several individual unconnected peaks), seems nonsensical. For particles where no switches are identified, there is a clear ”(no switch)” label attached to the prediction of the particle in question.

Additionally, a form of hysteresis was implemented for the predictions. If the averaged predictions reach a value below a hysteresis threshold of 0.1 outside of the QMS peaks, the switching point search will be nullified and chosen as the next time the network reaches the cutoff value of 0.8. This hysteresis is imporant, because the predictions often peak before the oscillatory prediction is sustained, due to isolated parts of the noisy period where some kind of oscillatory behaviour is found by the network.

The predictions for each measurement are shown in the figures below, with the first being figure A.20. They all contain the predictions of particles 1-8. In the title of each figure, there is information on which temperature and date for which the data was collected. In all figures, we also provide the value of the average switching point of all the particles in the current figure. This average is calculated with exclusion of the particles for which no switching points were identified. The horisontal axis for each figure shows which α that the corresponding time step corresponds to in the QMS-signal. The QMS peaks (solid vertical lines over the entire figure) signify where the QMS signal reached its two global optima.

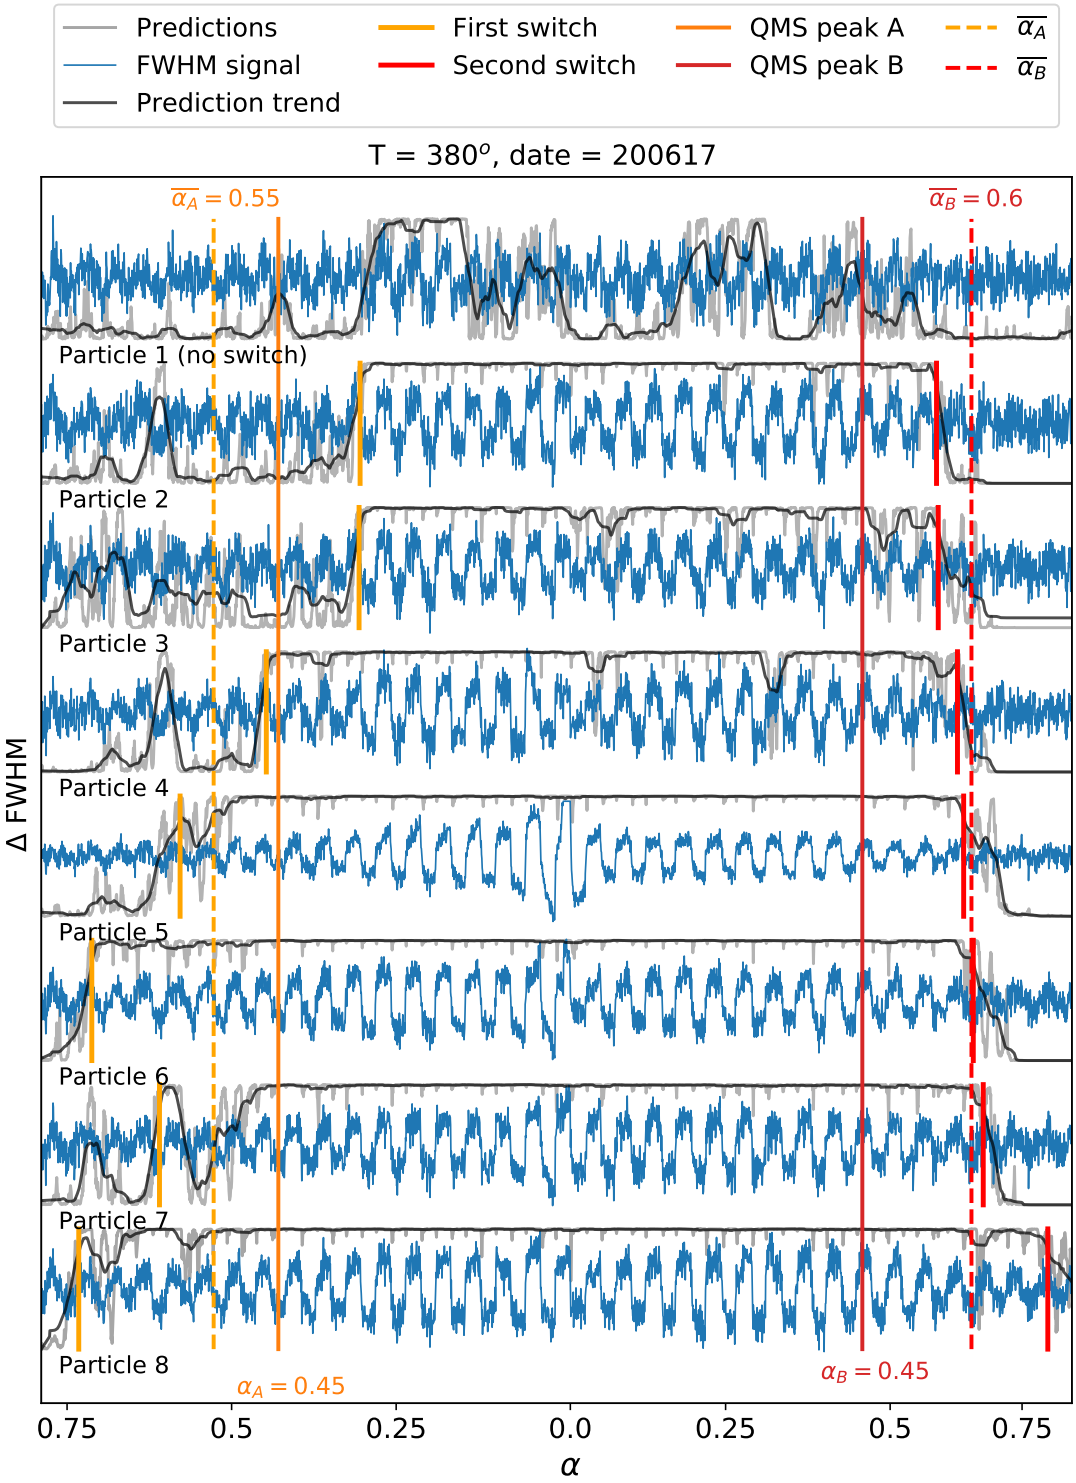

Figure A.20: Switching point predictions of FWHM signal with $T = 380°C$ collected on 200617.

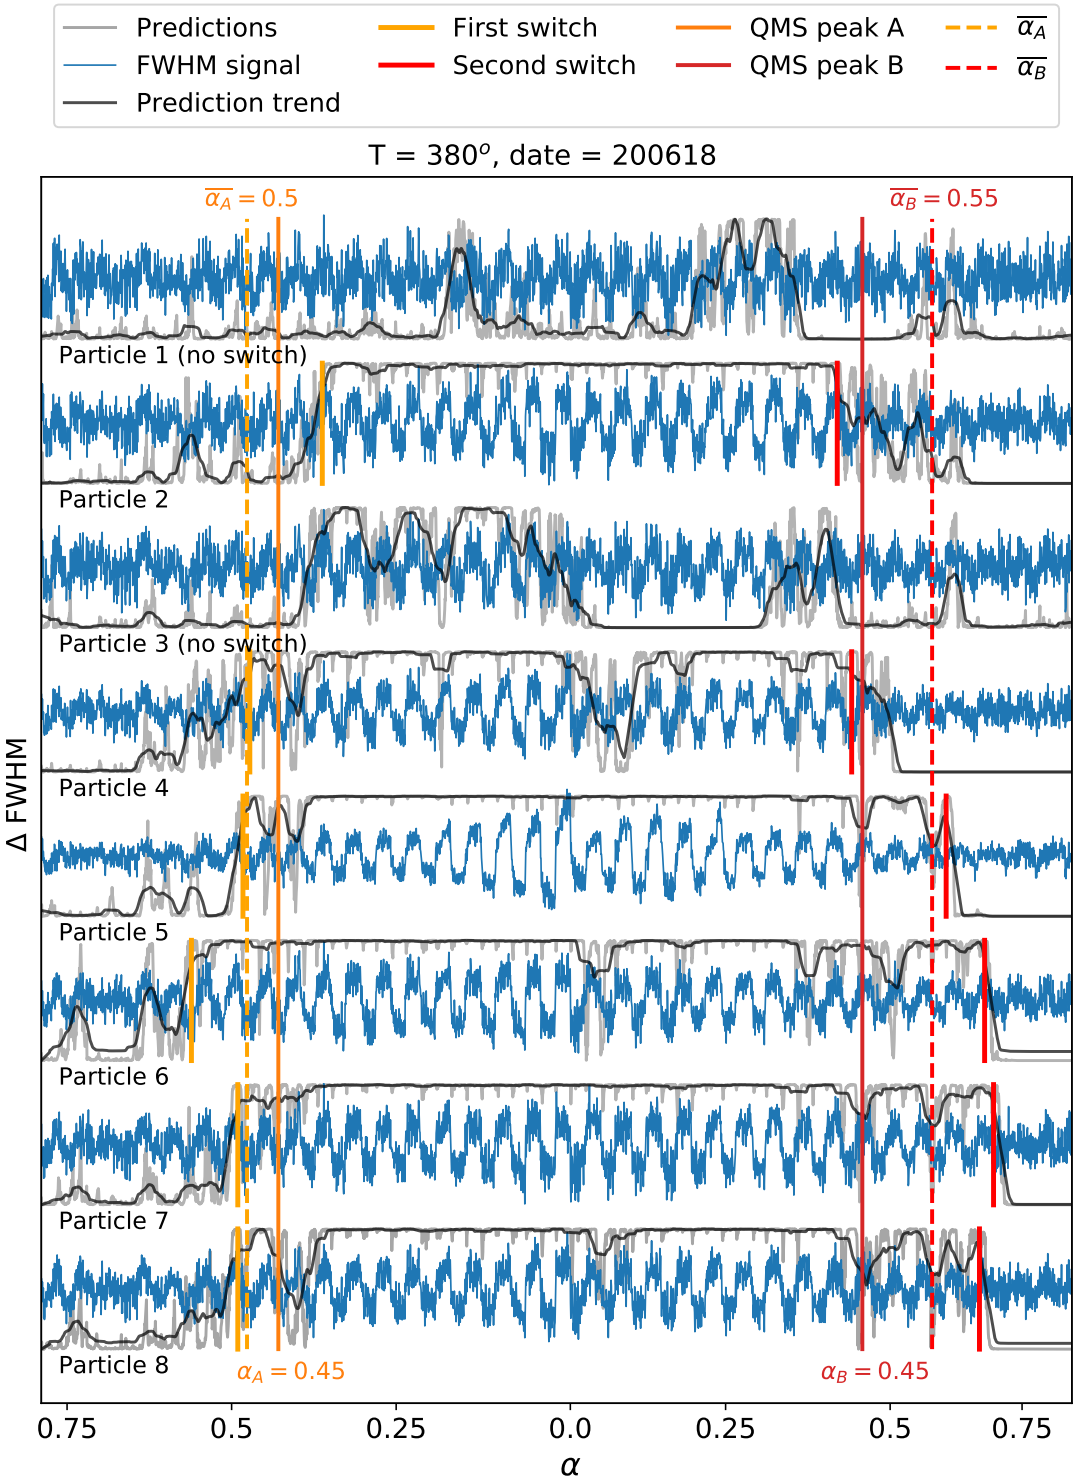

Figure A.21: Switching point predictions of FWHM signals with $T = 380°C$ collected on 200618.

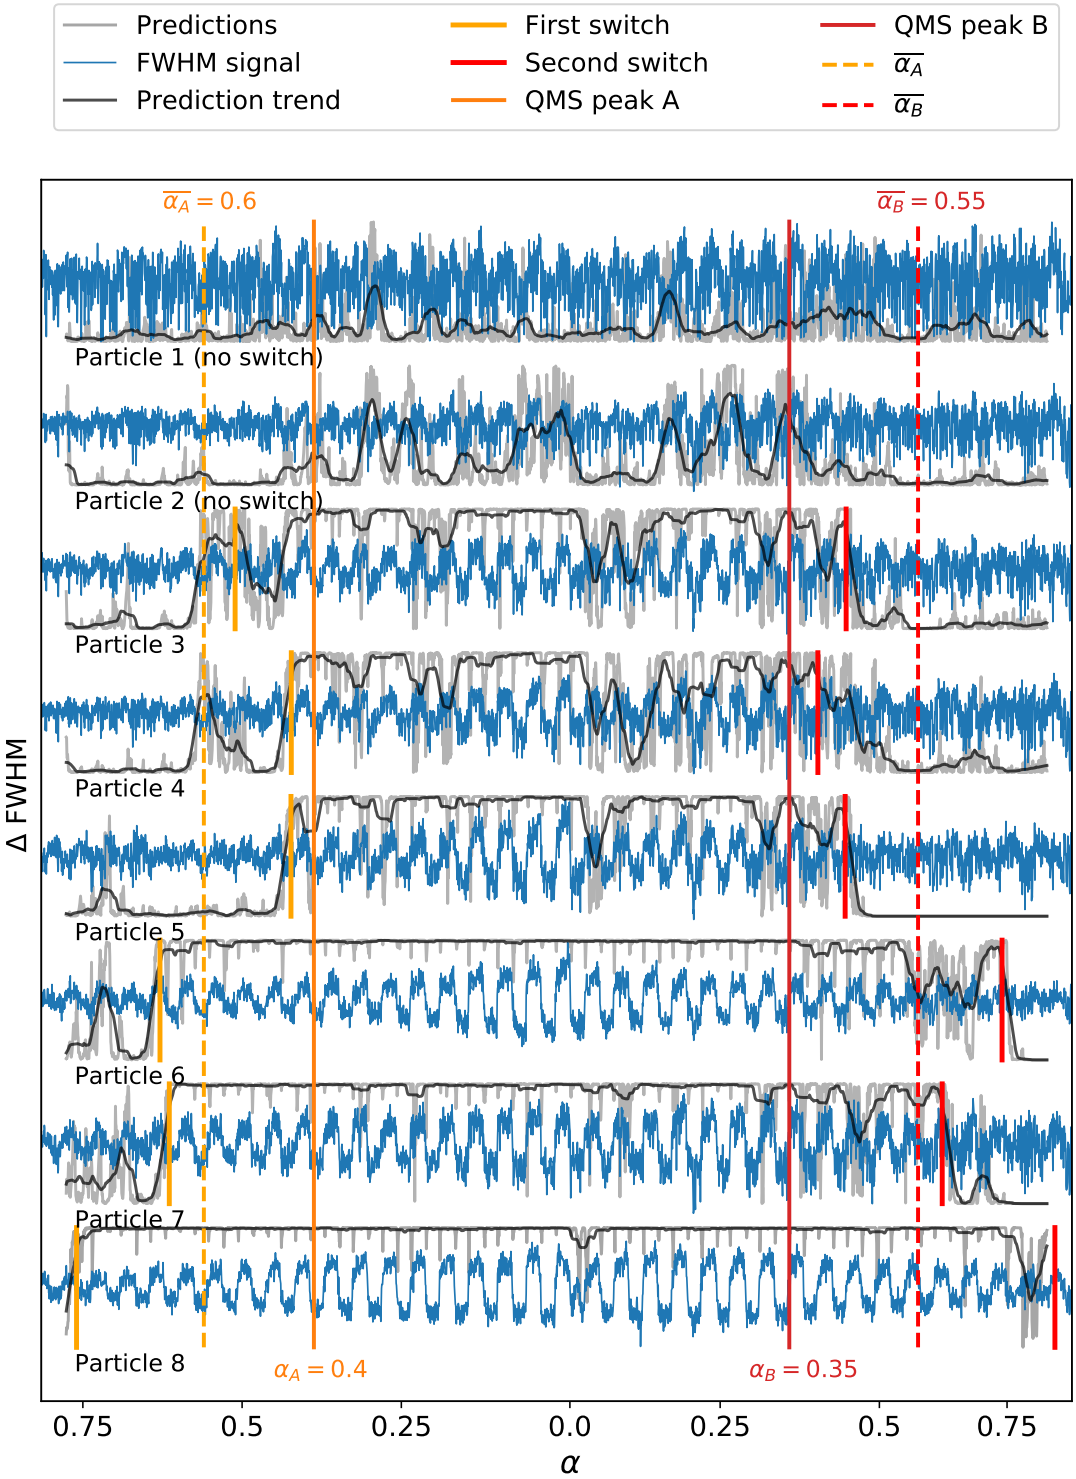

Figure A.22: Switching point predictions of FWHM signal with $T = 370°C$ collected on 200622.

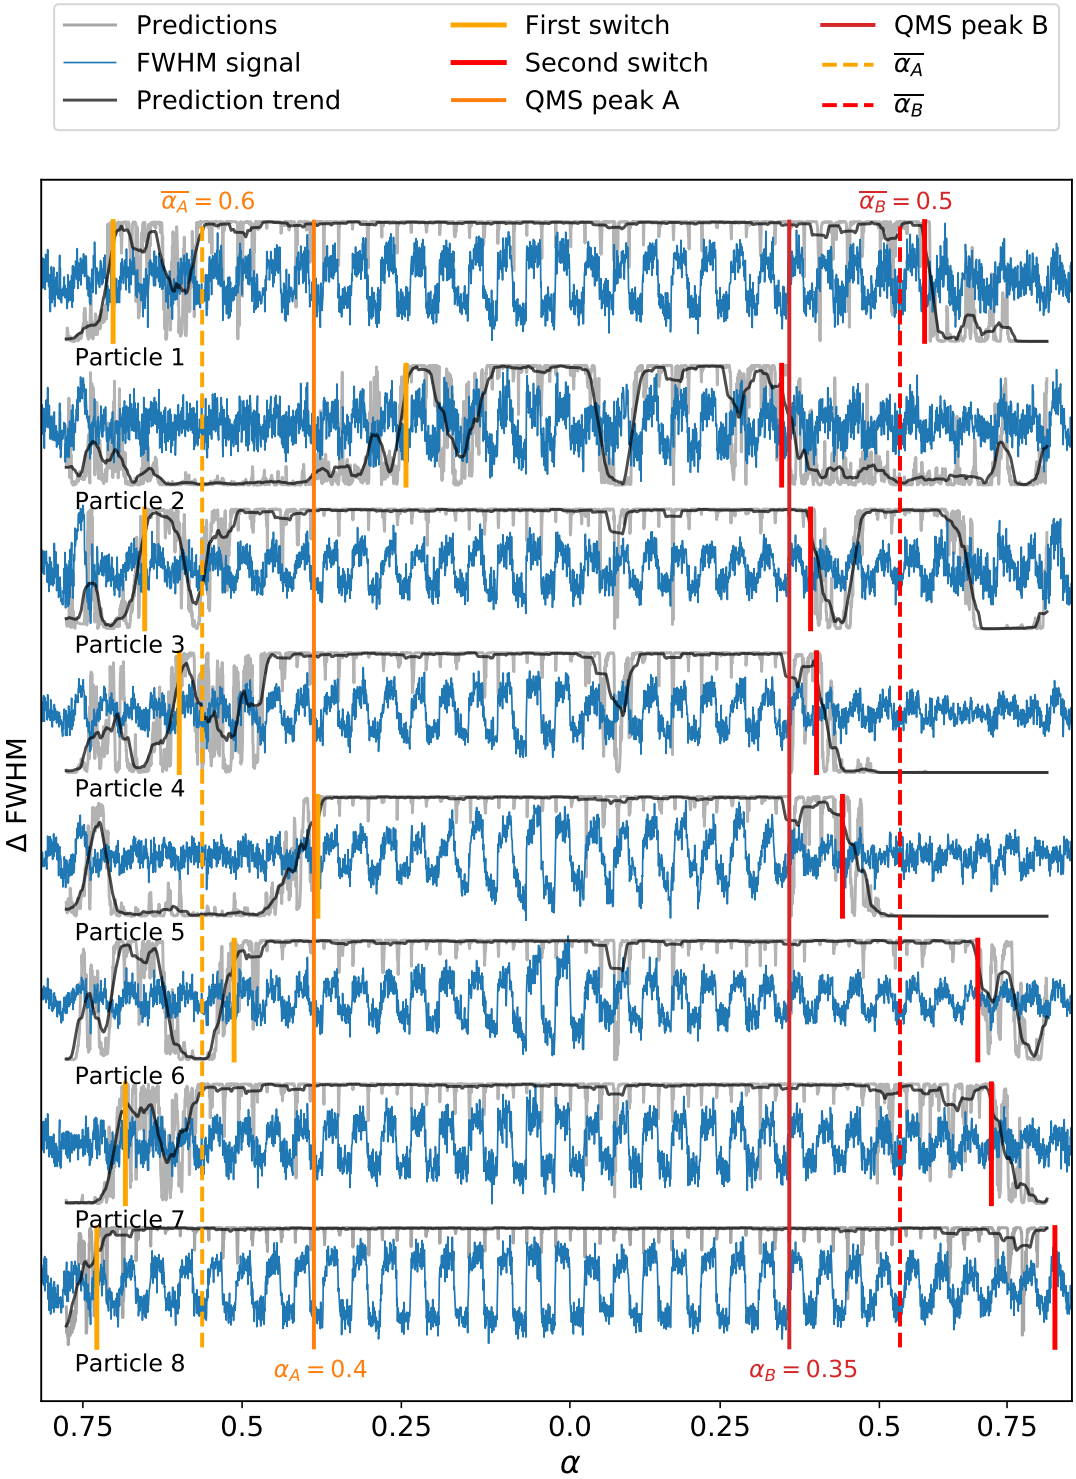

Figure A.23: Switching point predictions of FWHM signal with $T = 370°C$ collected on 200623.

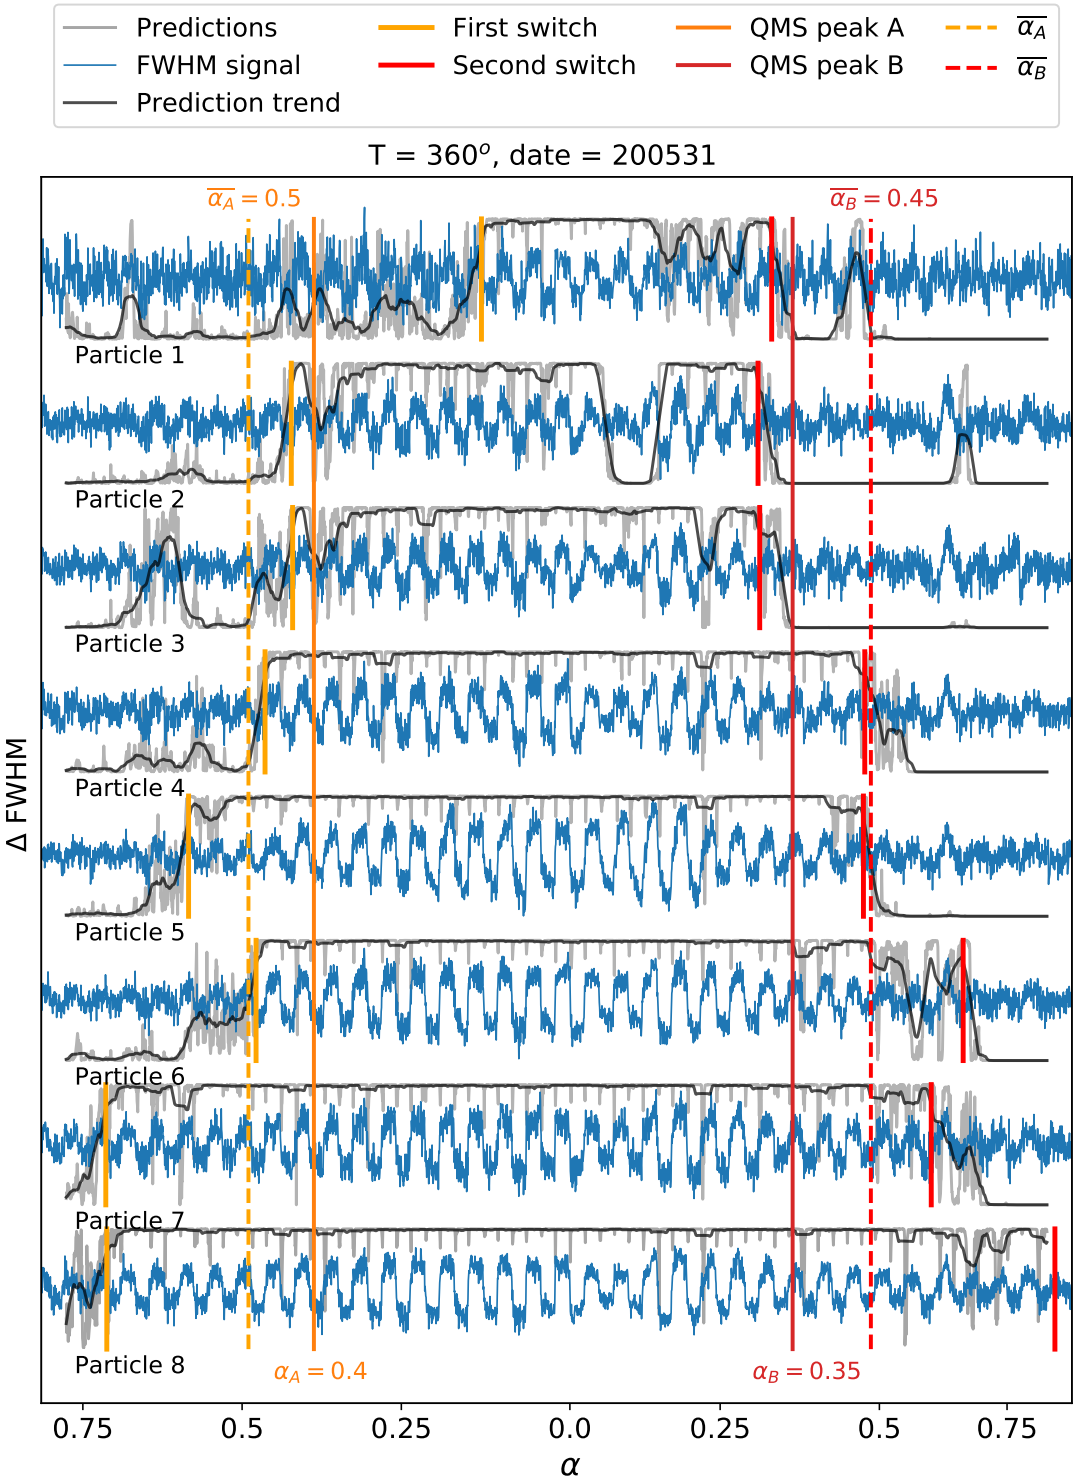

Figure A.24: Switching point predictions of FWHM signal with $T = 360°C$ collected on 200531.

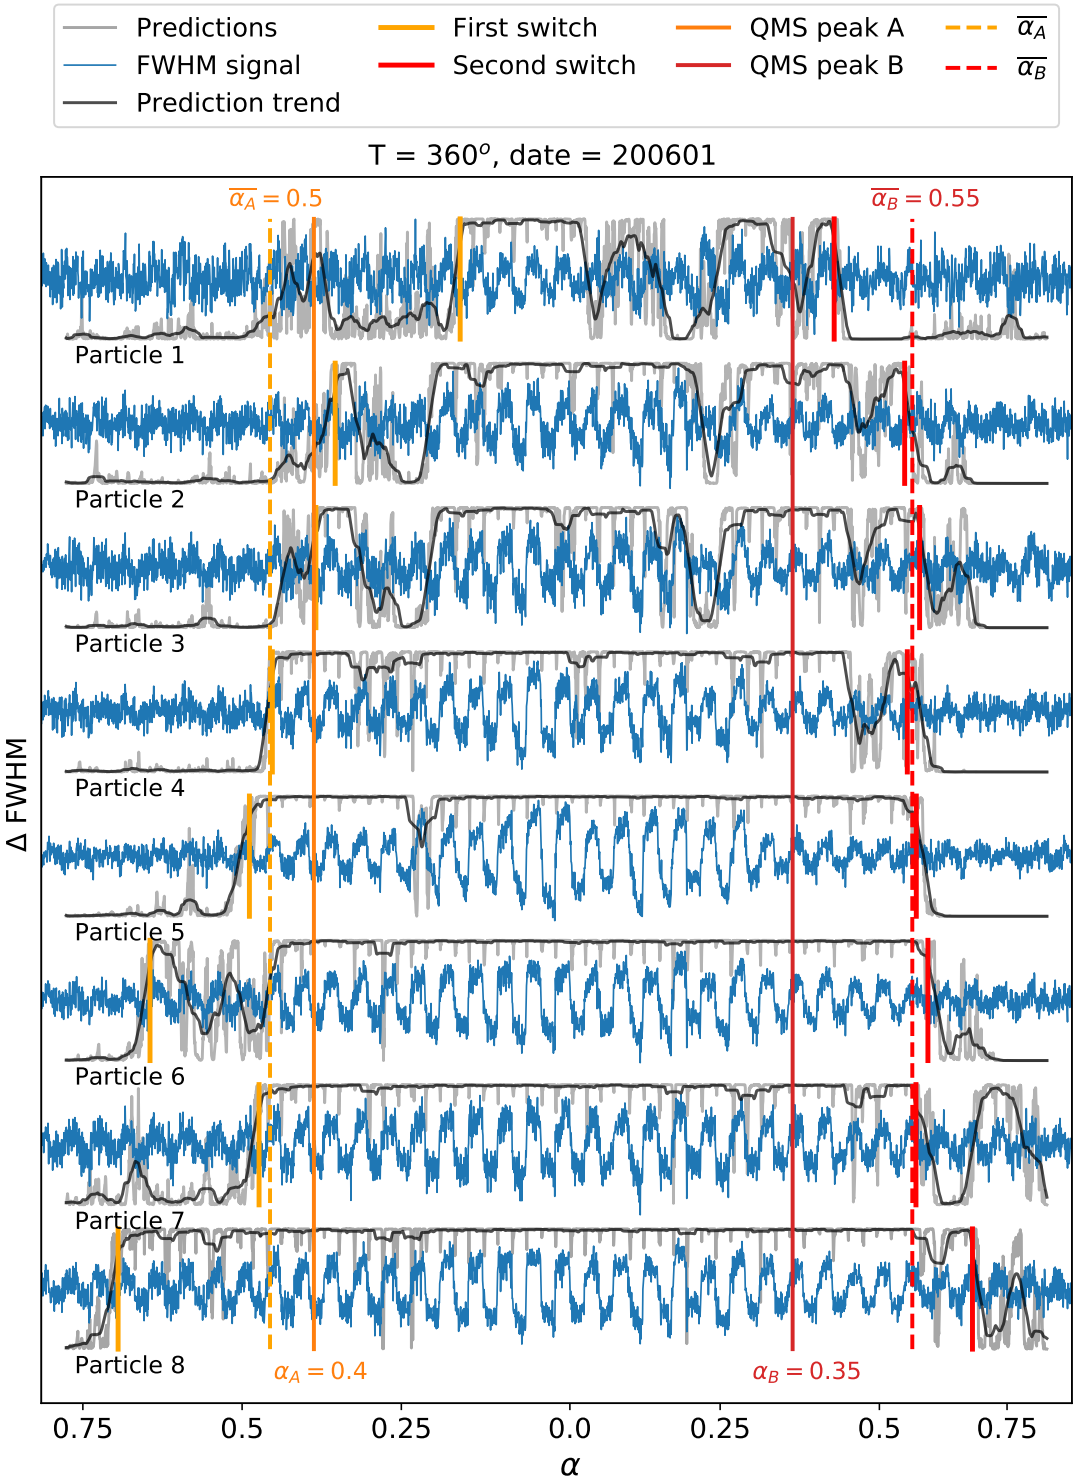

Figure A.25: Switching point predictions of FWHM signal with $T = 360°C$ collected on 200601.

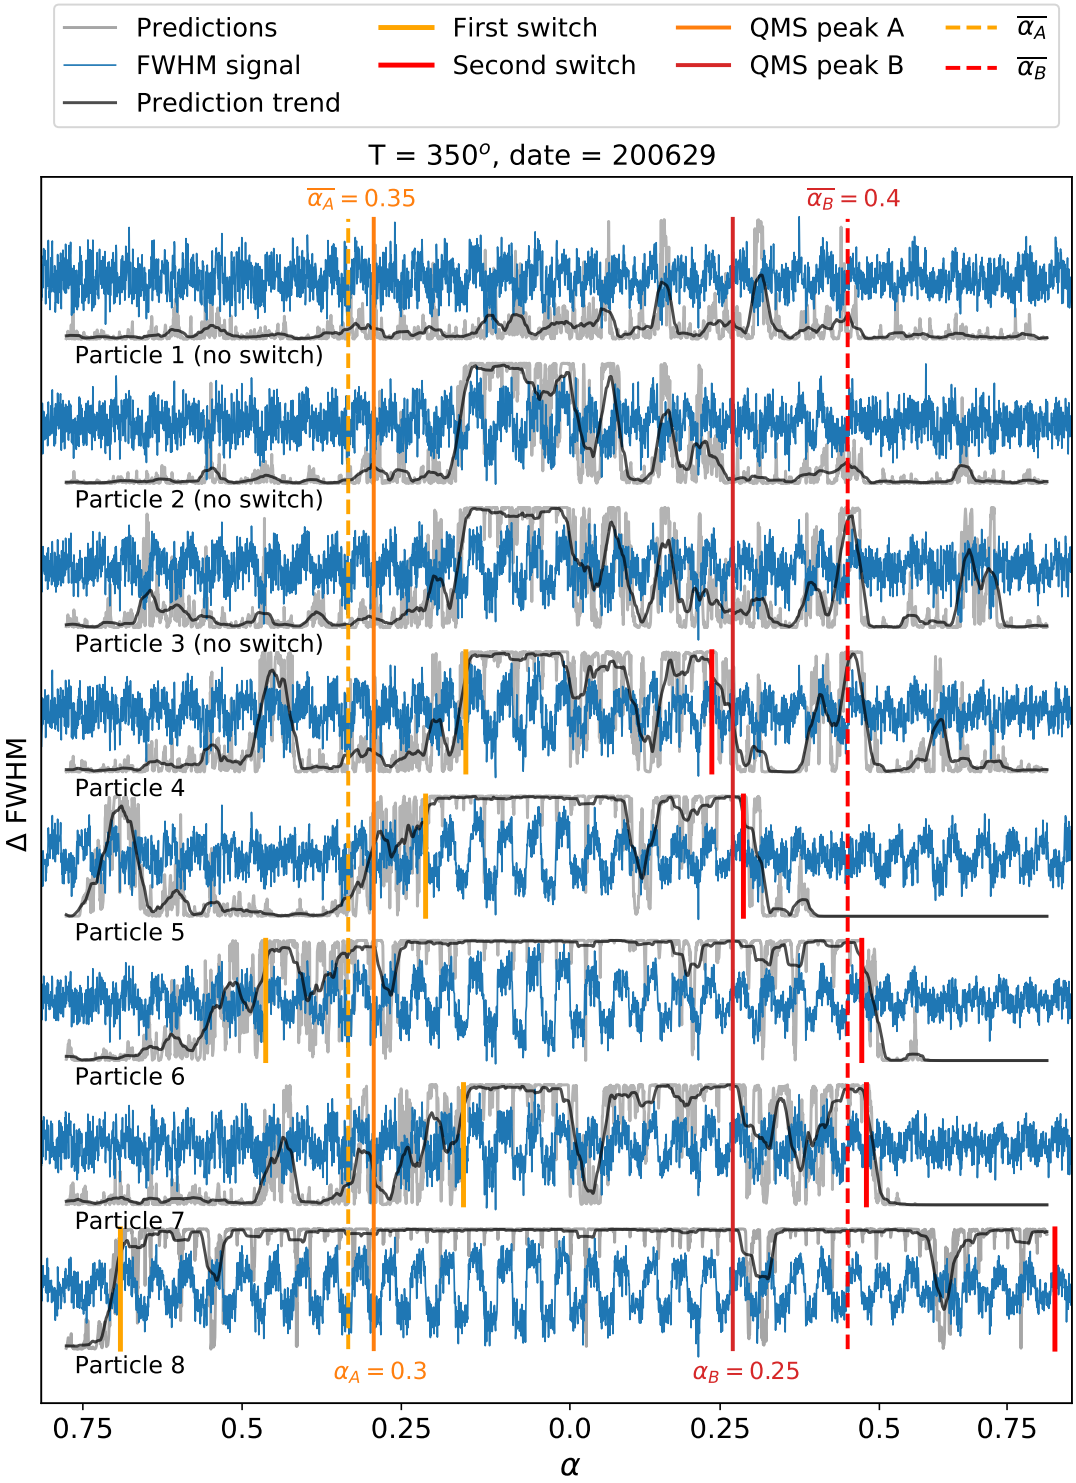

Figure A.26: Switching point predictions of FWHM signal with $T = 350°C$ collected on 200629.

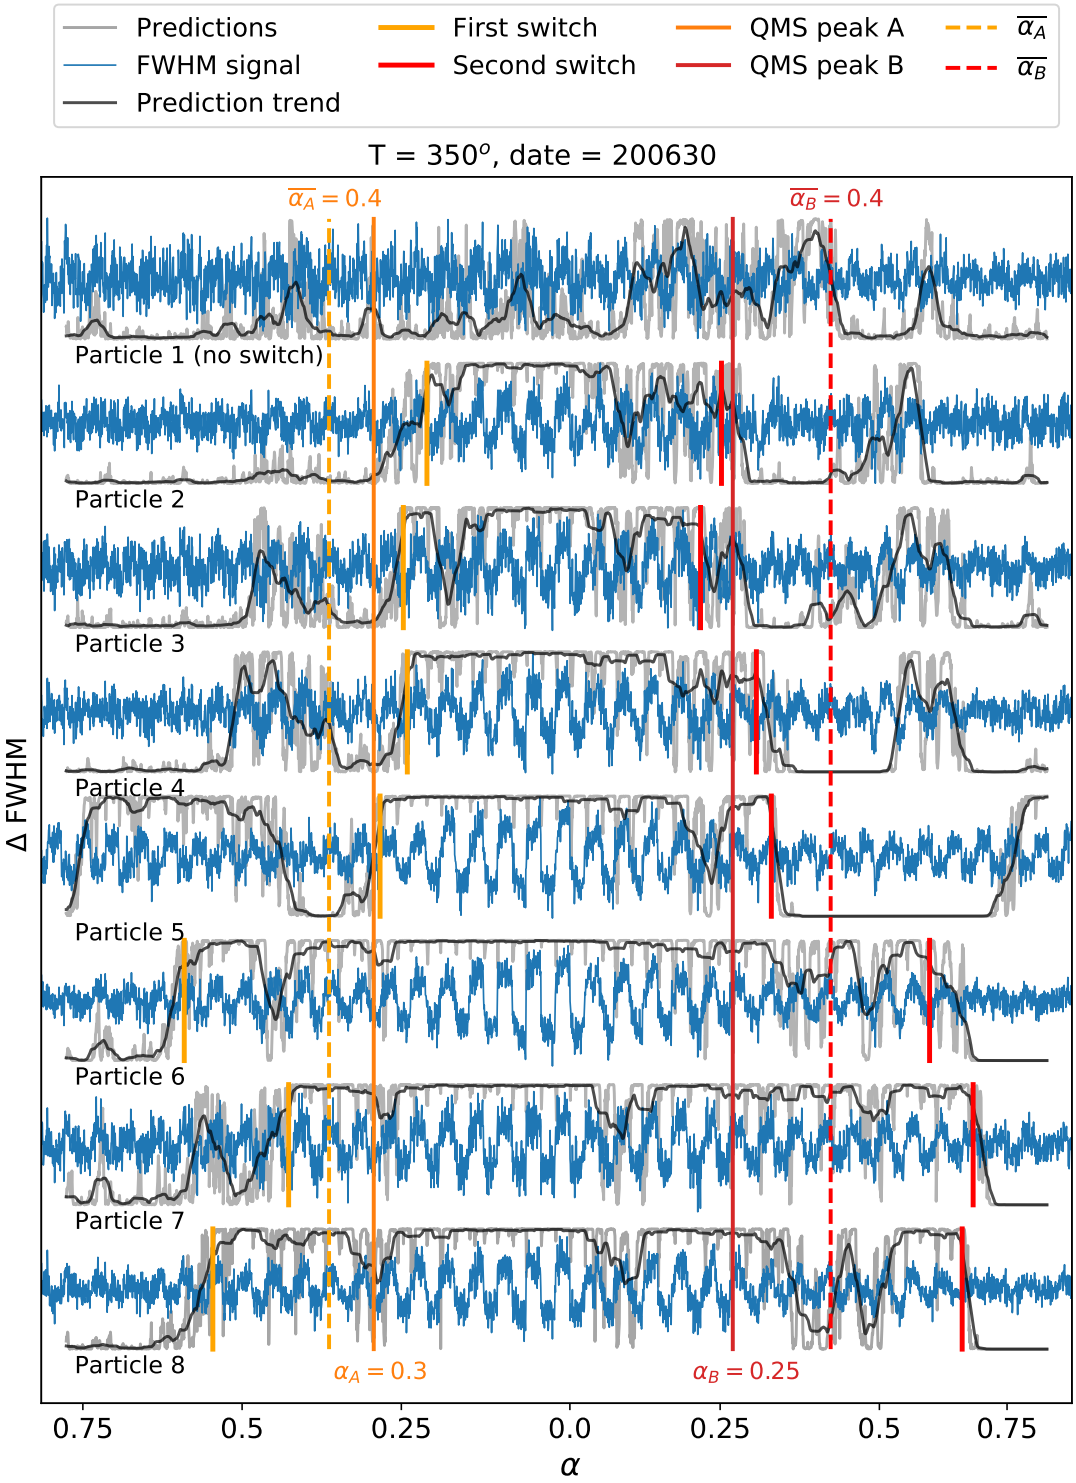

Figure A.27: Switching point predictions of FWHM signal with $T = 350°C$ collected on 200630.

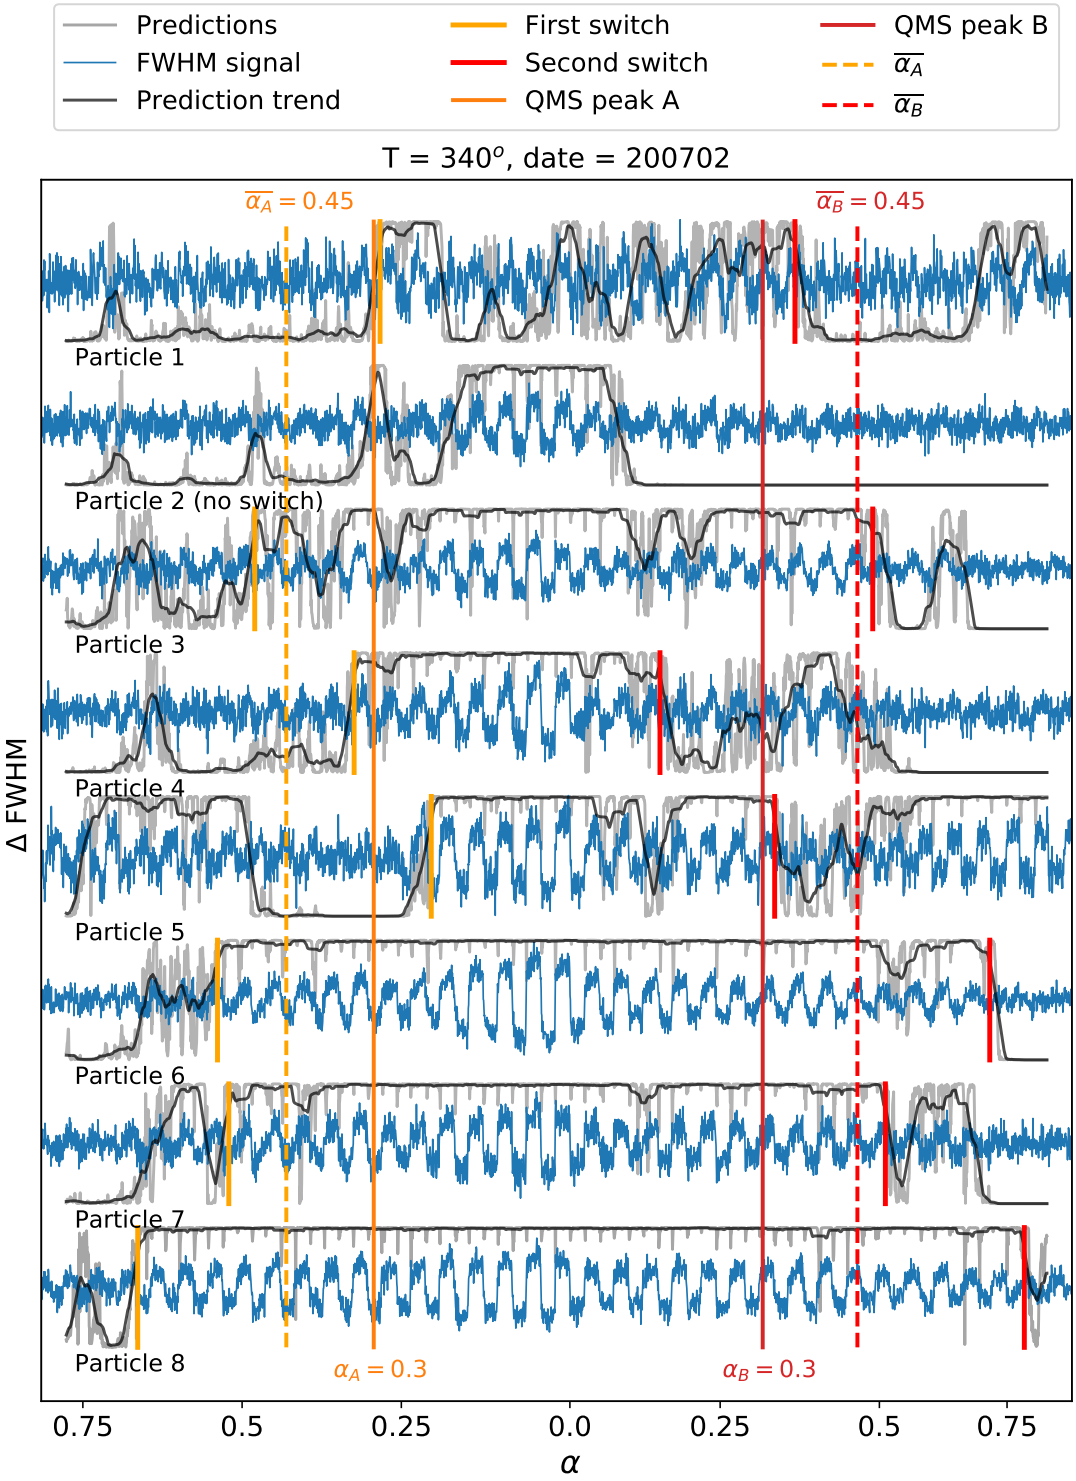

Figure A.28: Switching point predictions of FWHM signal with $T = 340°C$ collected on 200702.

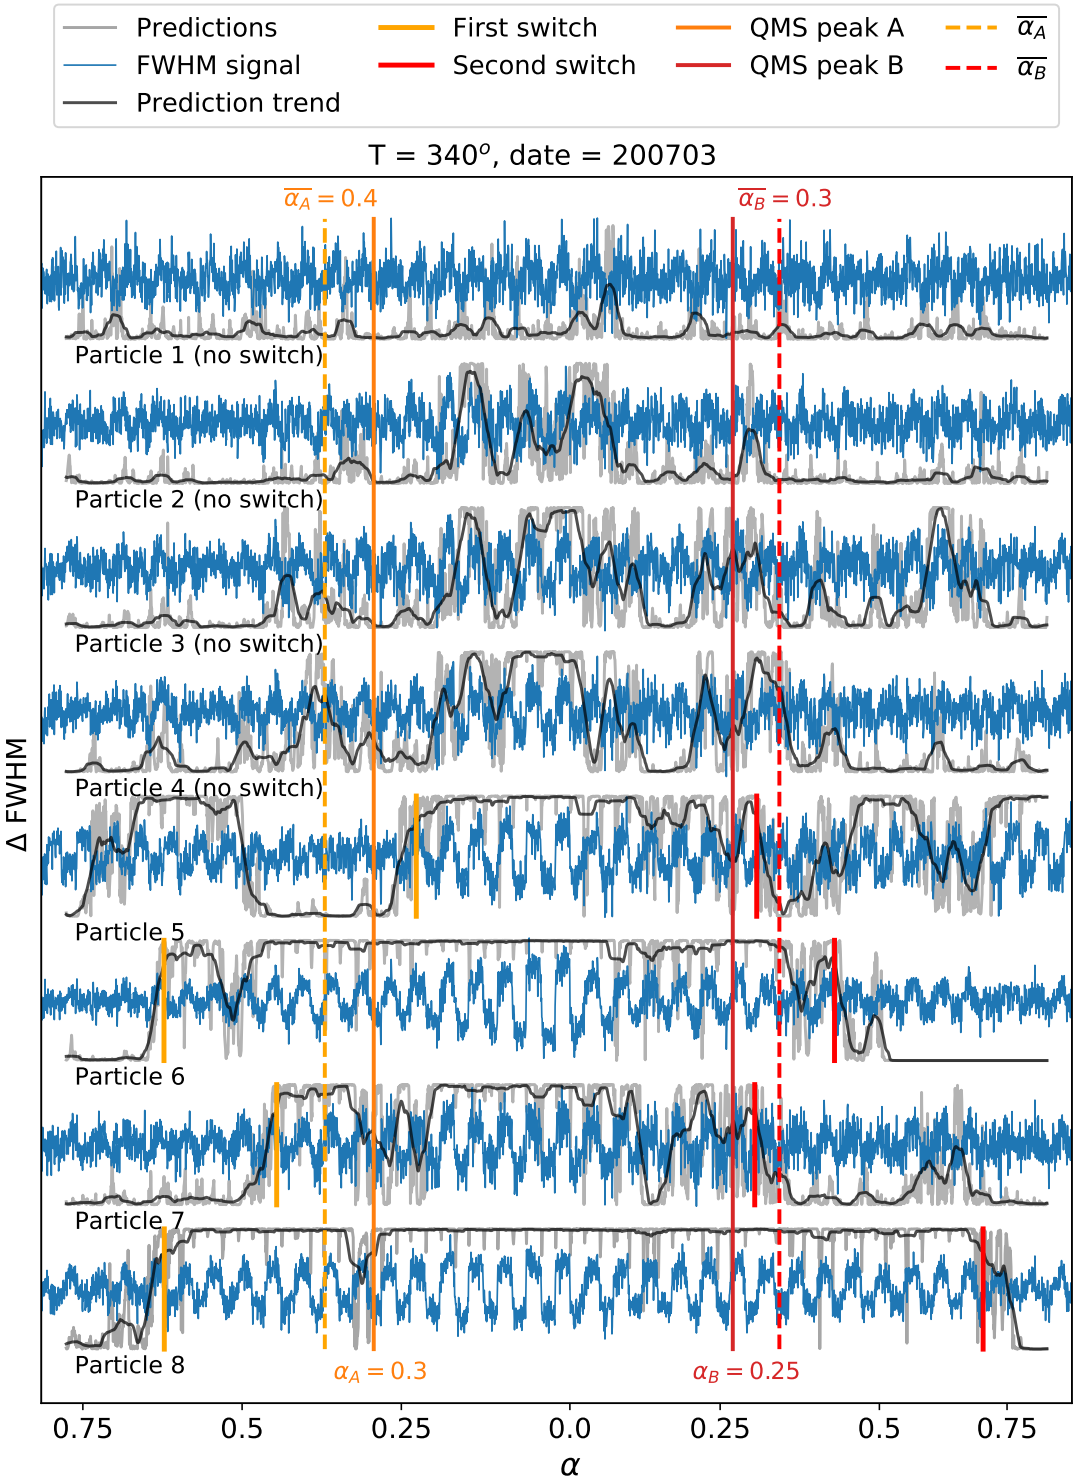

Figure A.29: Switching point predictions of FWHM signal with $T = 340°C$ collected on 200703.

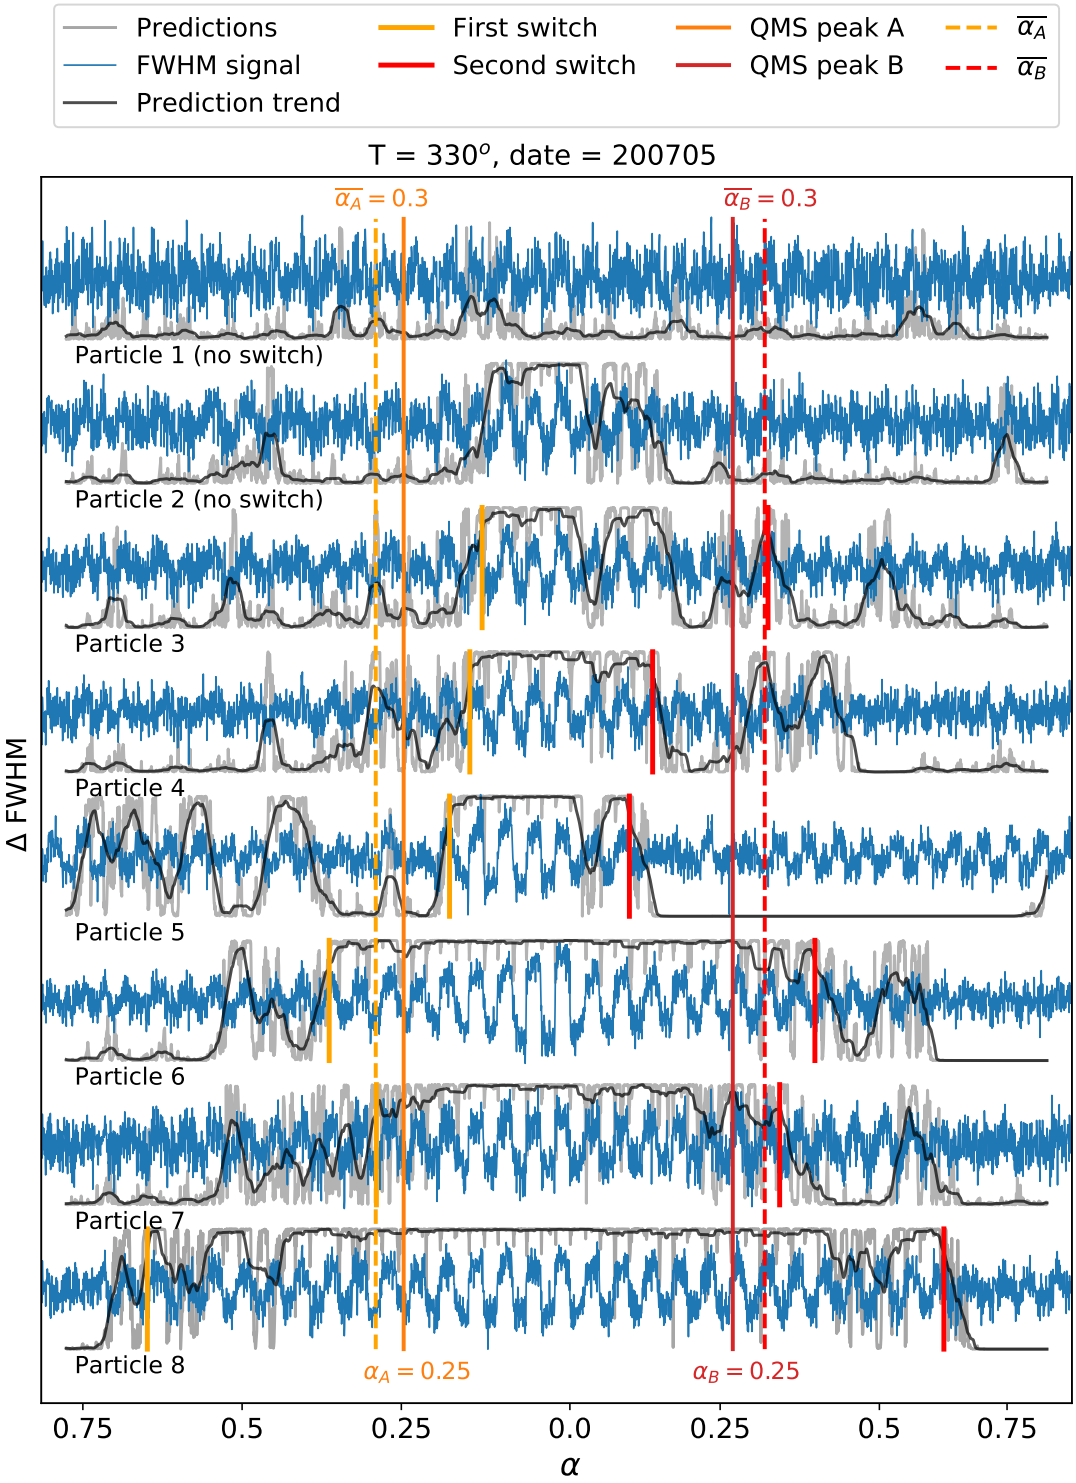

Figure A.30: Switching point predictions of FWHM signal with $T = 330°C$ collected on 200705.

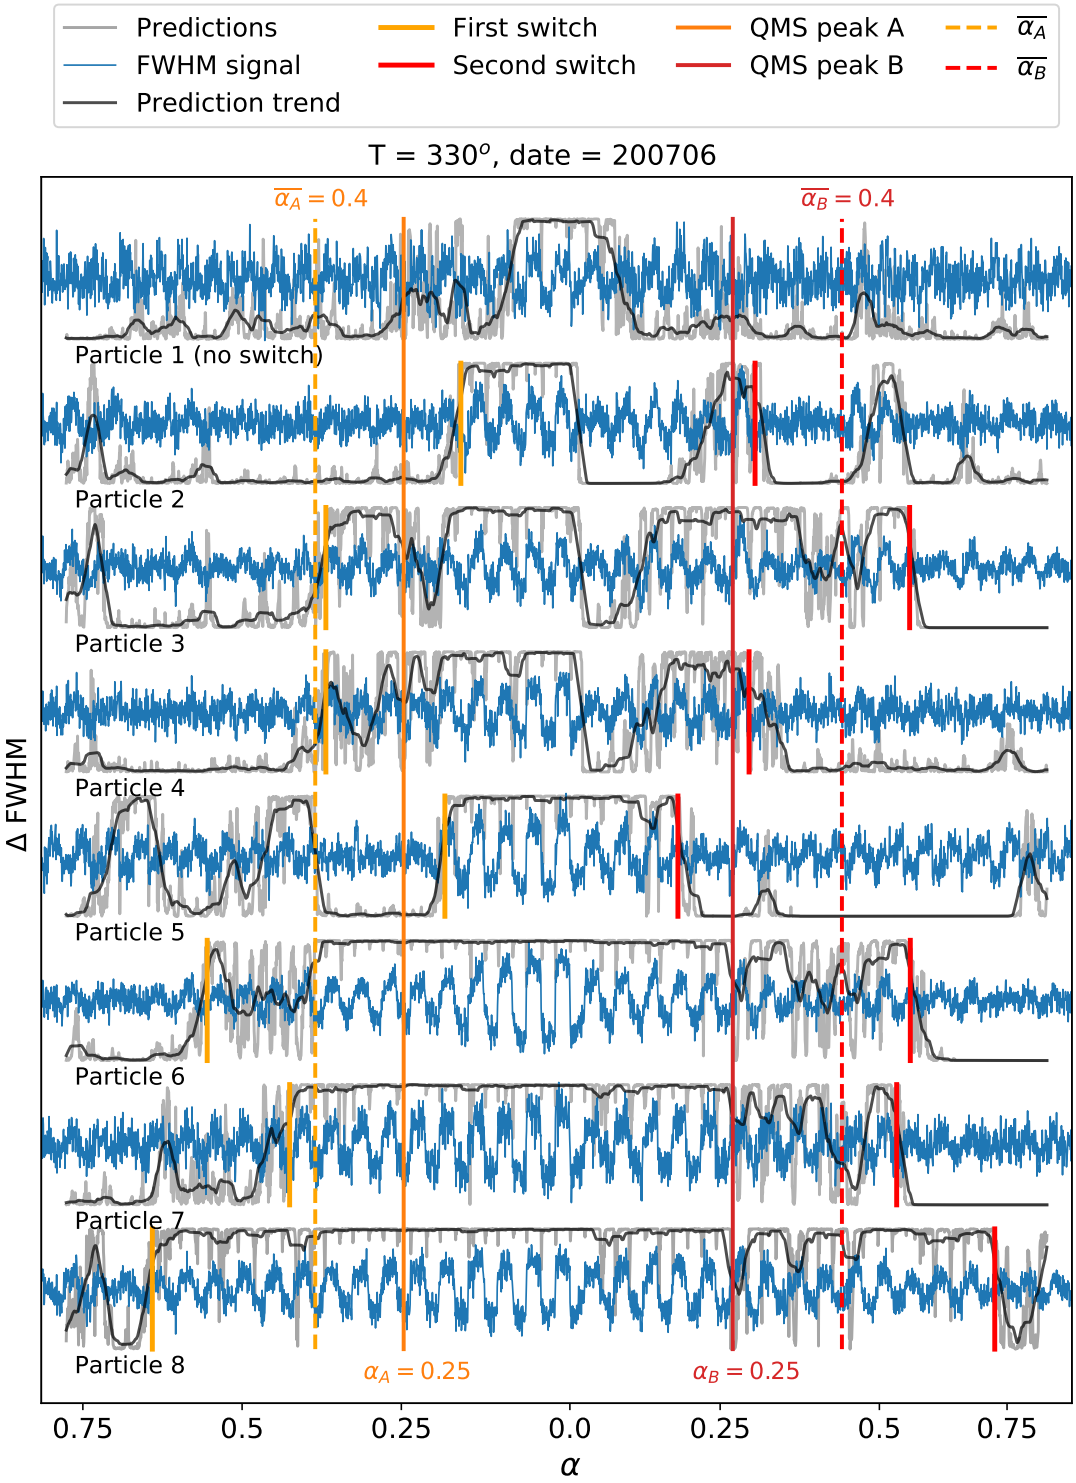

Figure A.31: Switching point predictions of FWHM signal with $T = 330°C$ collected on 200706.

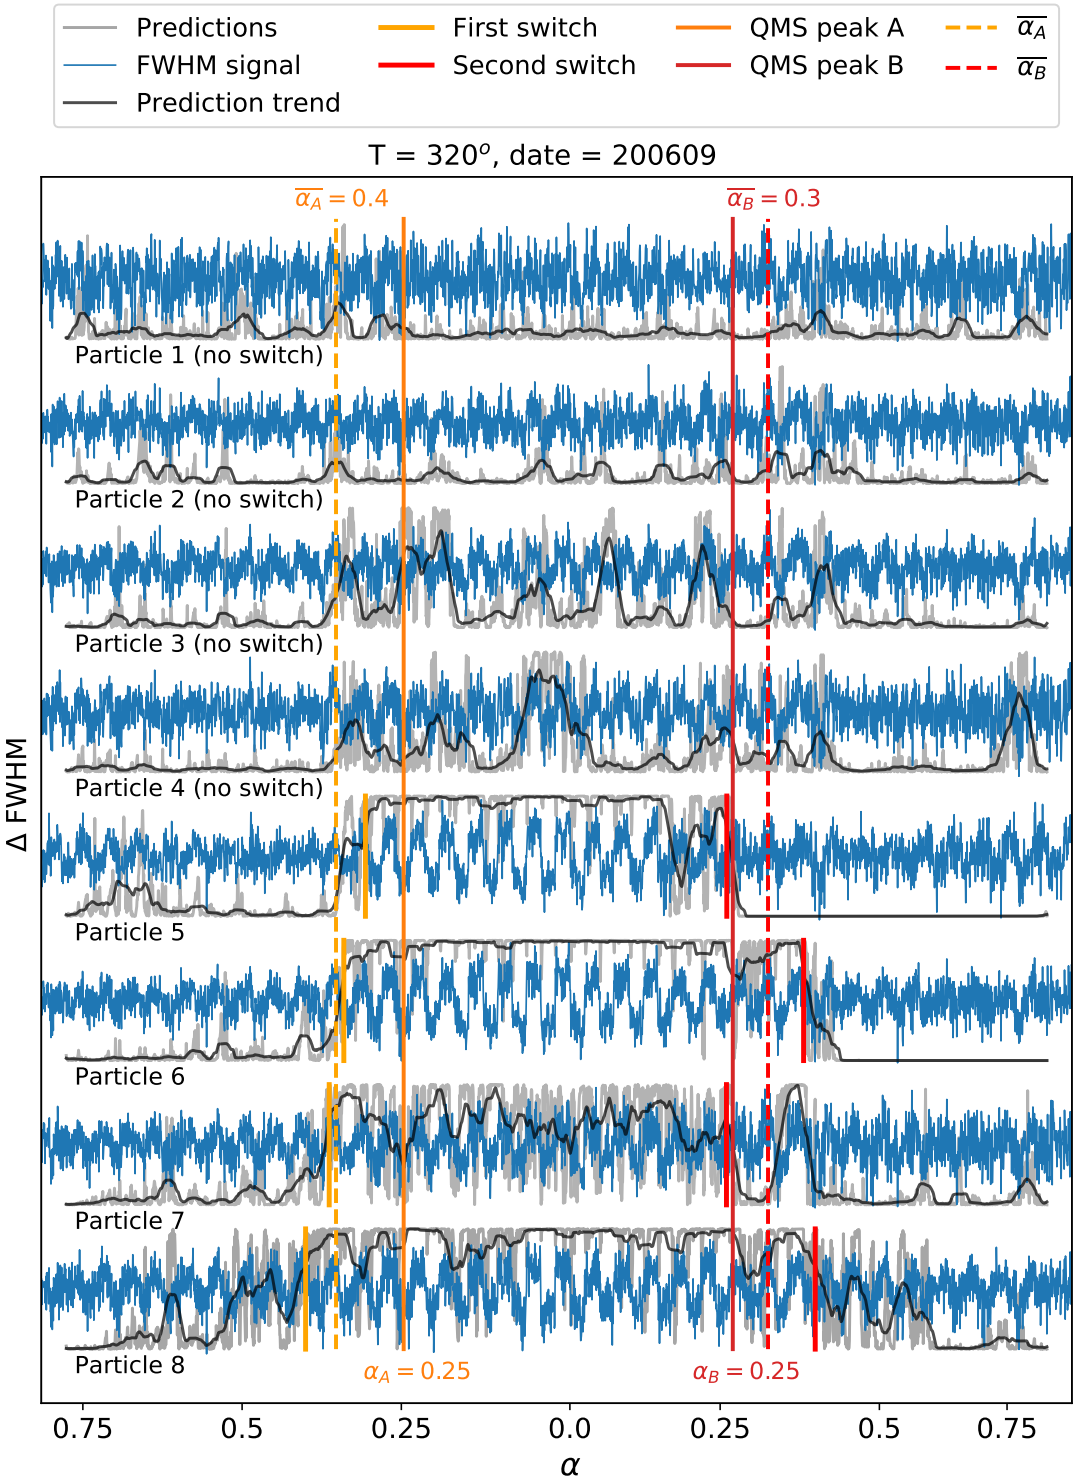

Figure A.32: Switching point predictions of FWHM signal with $T = 320°C$ collected on 200609.

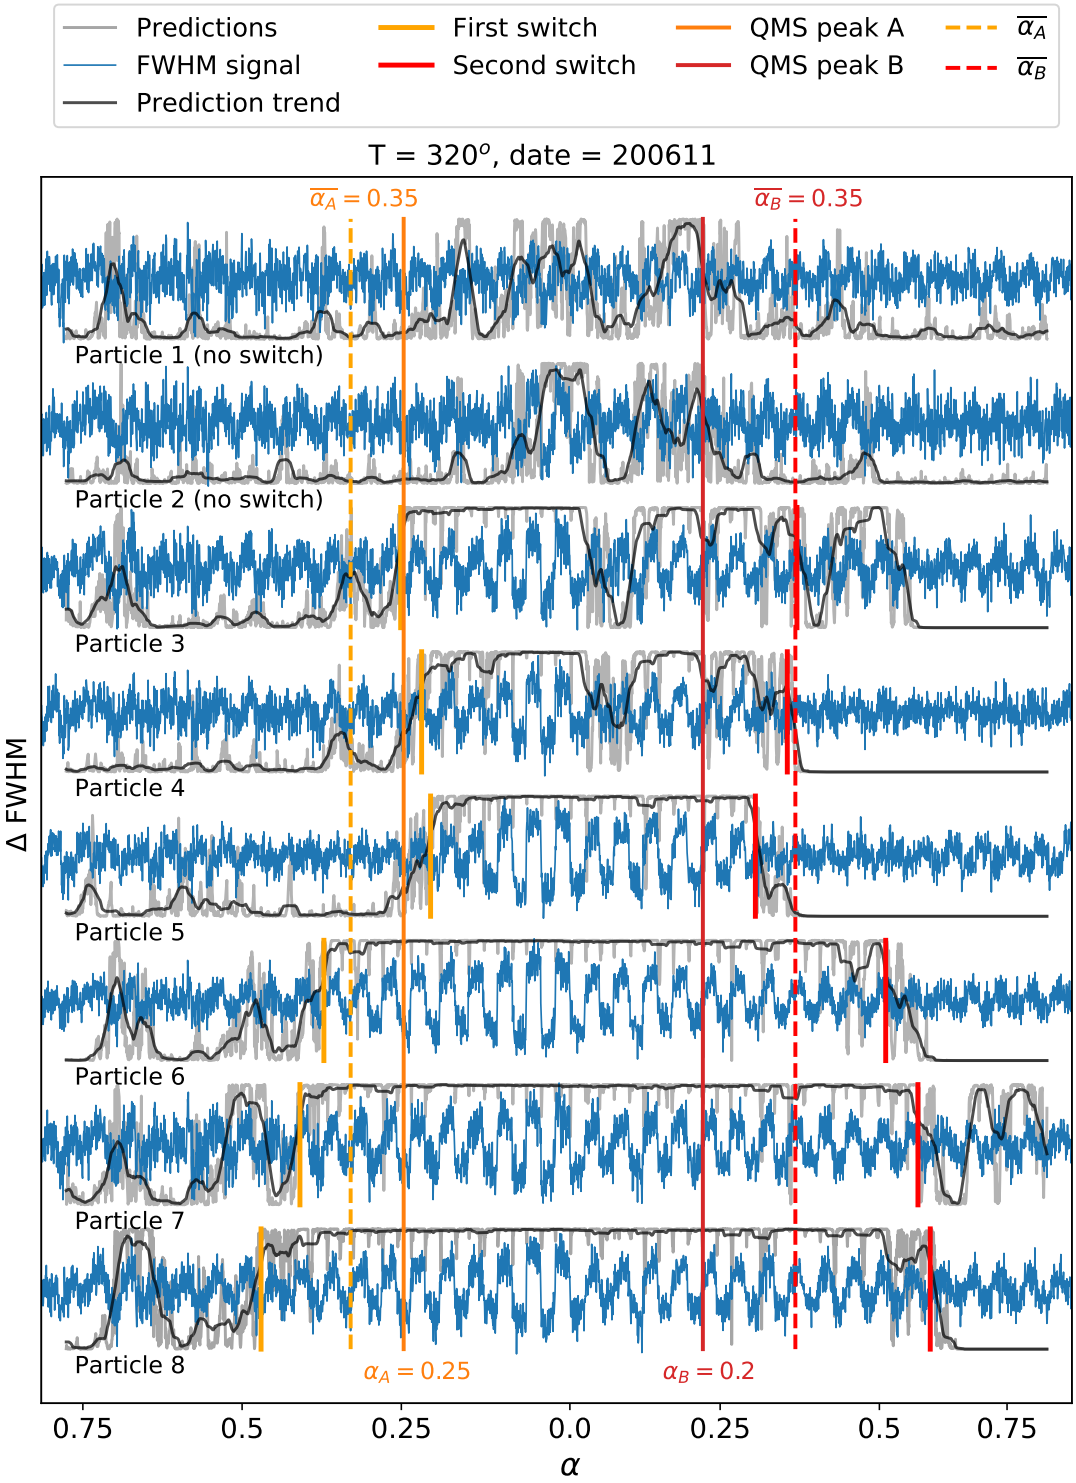

Figure A.33: Switching point predictions of FWHM signal with $T = 320°C$ collected on 200611.

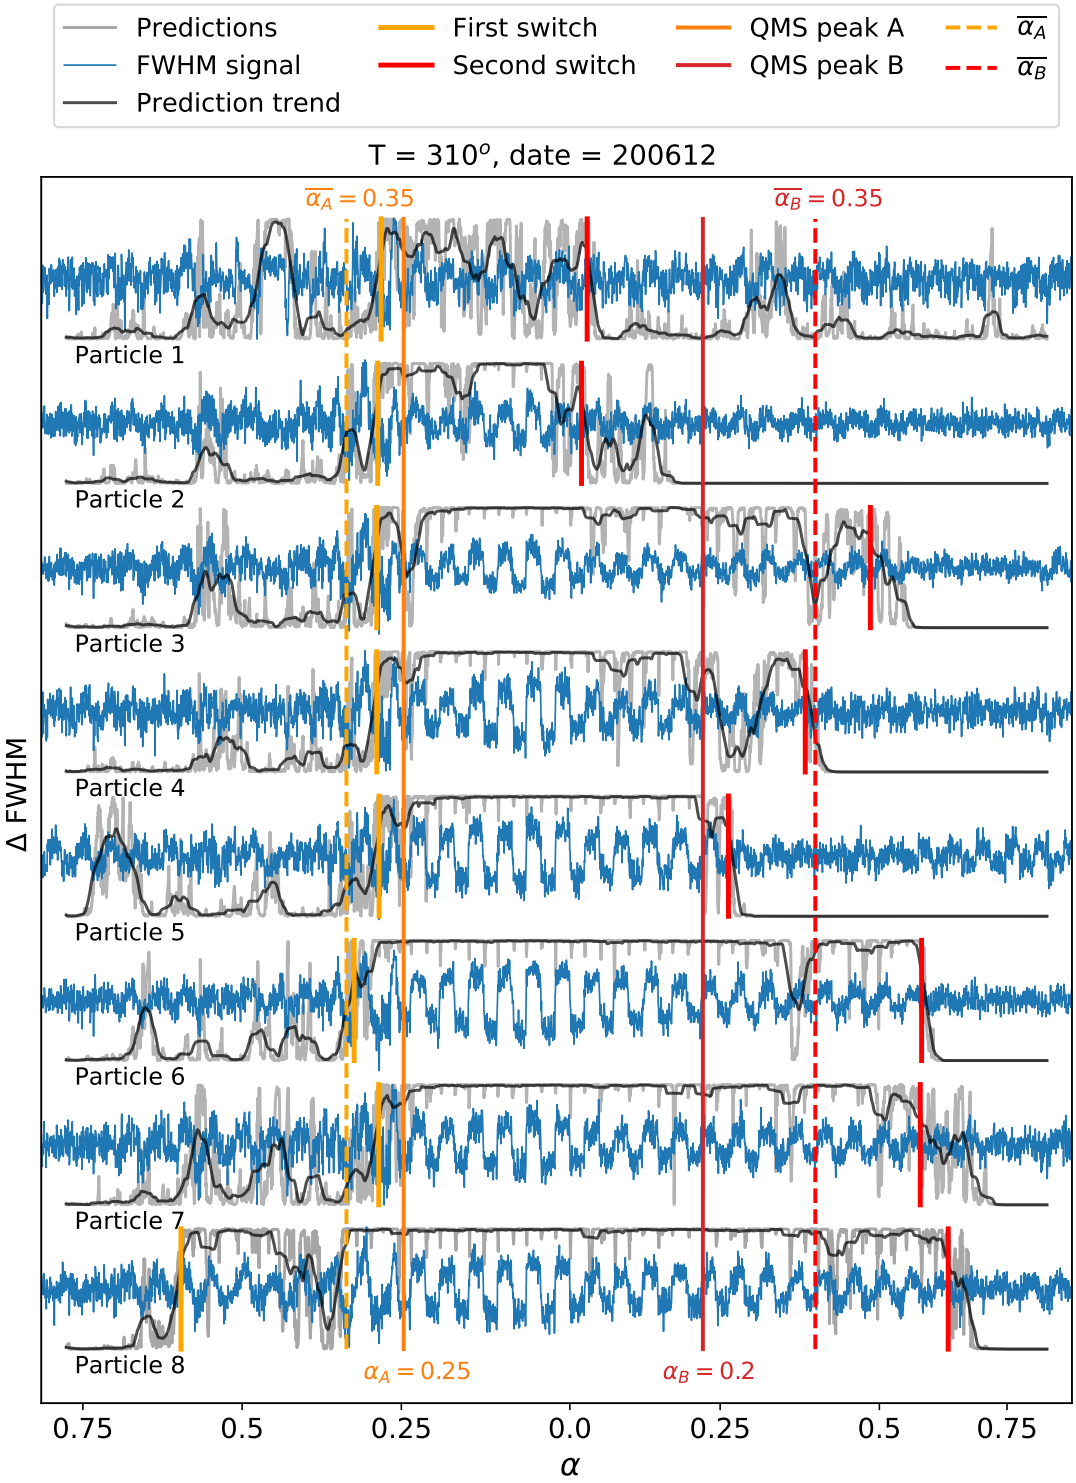

Figure A.34: Switching point predictions of FWHM signal with $T = 310°C$ collected on 200612.

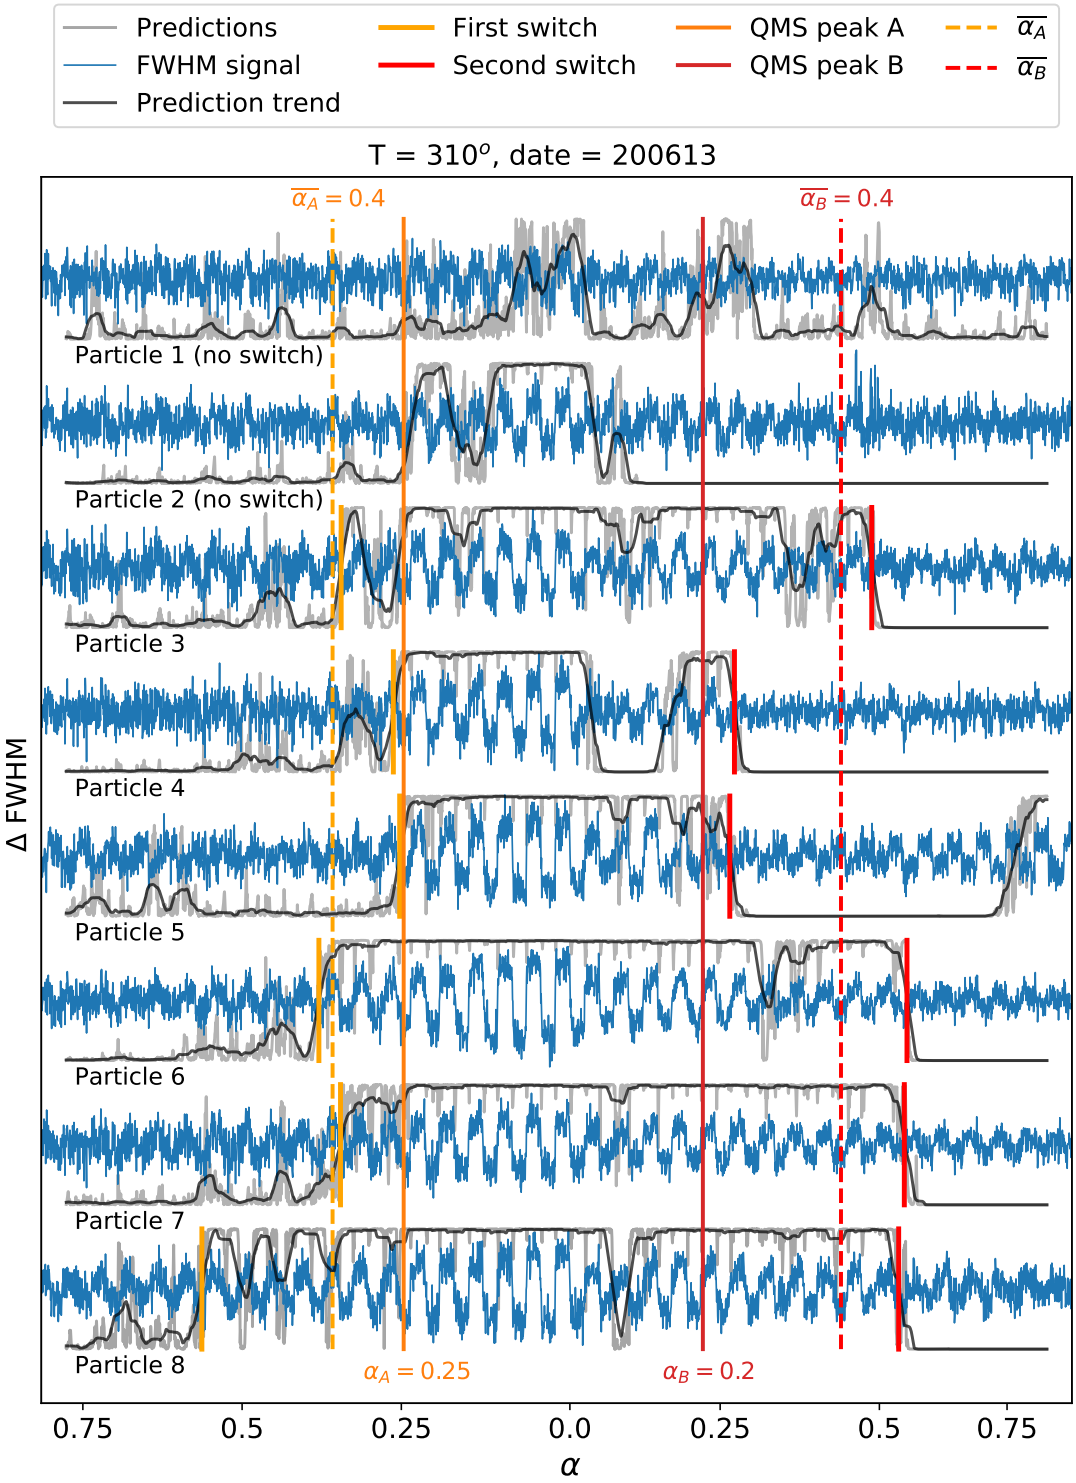

Figure A.35: Switching point predictions of FWHM signal with $T = 310°C$ collected on 200613.

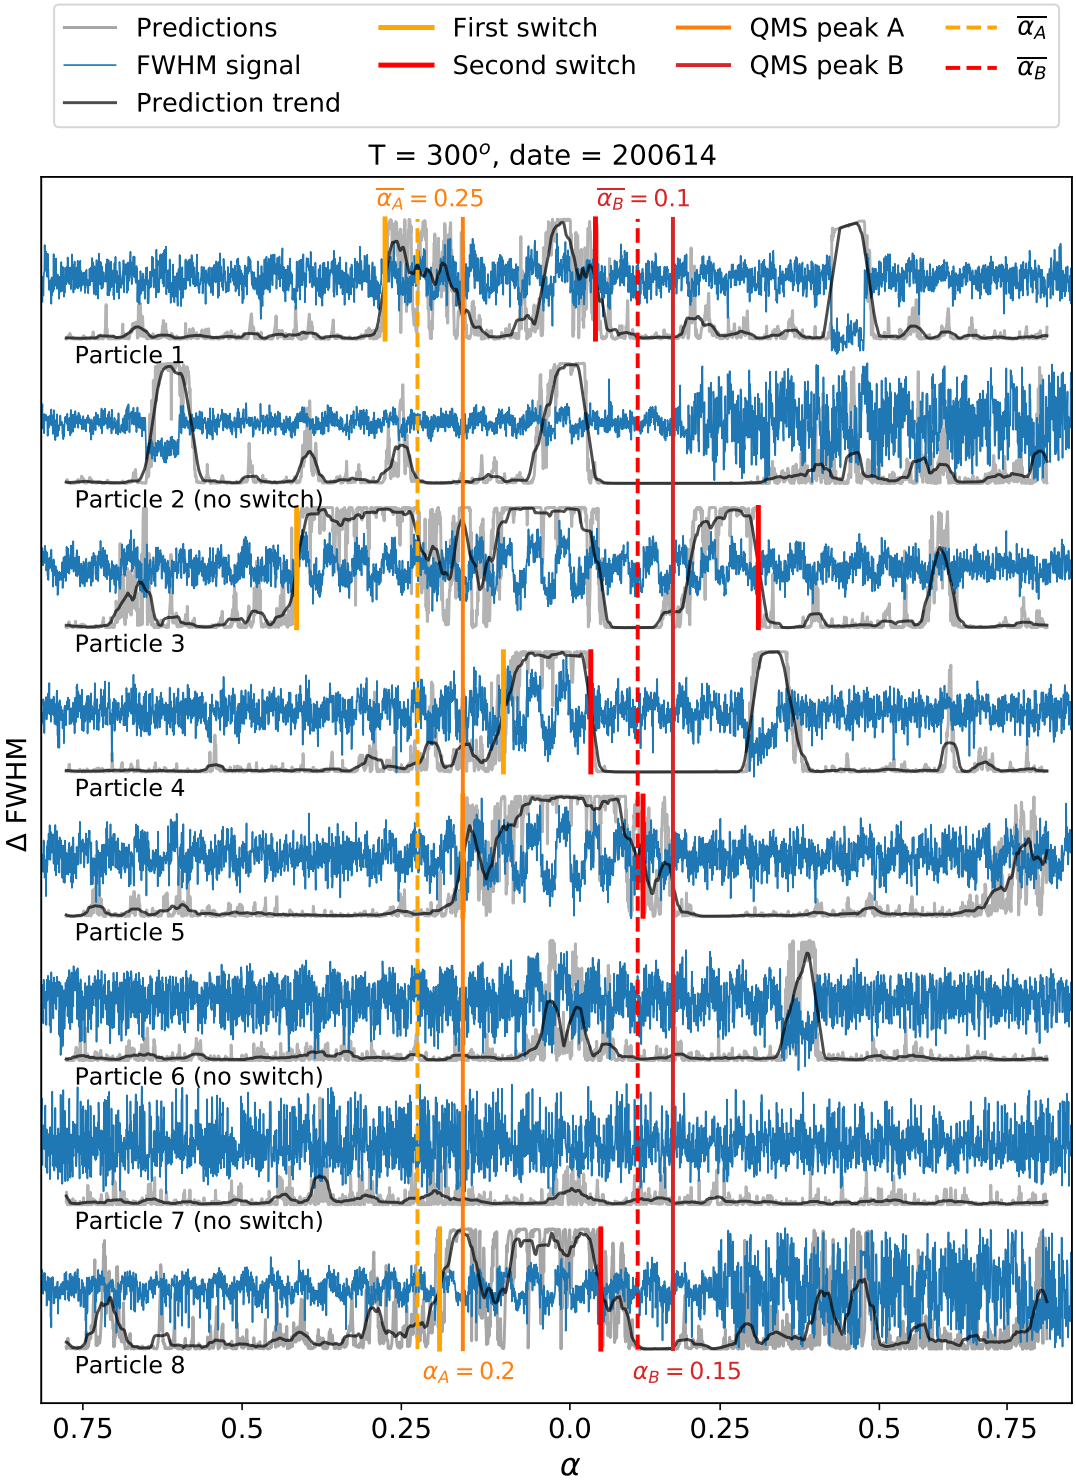

Figure A.36: Switching point predictions of FWHM signal with $T = 300°C$ collected on 200614.

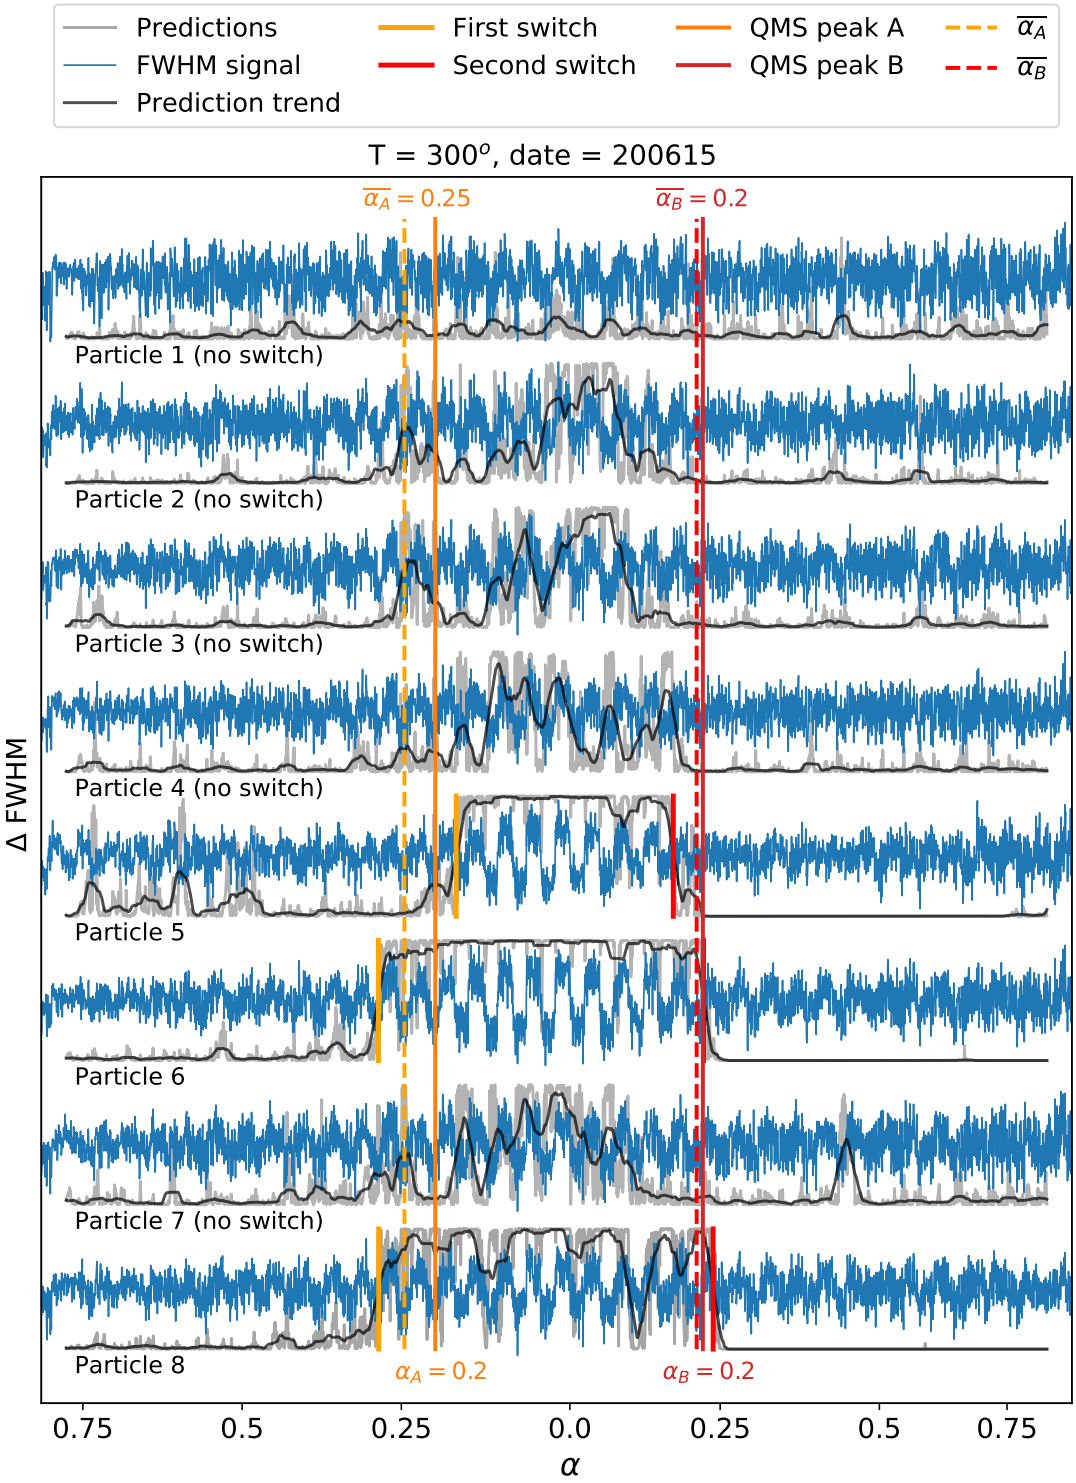

Figure A.37: Switching point predictions of FWHM signal with $T = 300°C$ collected on 200615.

A.3.6.3 Predictions of data from February

In this section we provide predictions of oscillations for measurements made on the same chip a few months prior to the predictions shown in the previous section. We have observed that the FWHM signals for these measurements differs a lot from the measurements made a few months later. By manual observation, we see that these signals do not show oscillatory behaviour to nearly the same extent as the newer measurements, indicating that the particles undergo some form of transformation with time. As can be seen in figure A.38 (and in the rest of the figures in this section), the network consequently identifies much shorter oscillatory regions compared to the newer measurements, see for instance figure A.20 for comparison. Due to the uncertain output of the network for these old measurements, we will not provide switching points based on these predictions (in the few cases where it would even be possible). By inspecting the predictions of for instance figure A.38, we can however tell that the predictions seem to be correlated to the region between the peaks of the QMS signal, but the predictions are far from convincing enough for us to make precise switching point identifications.

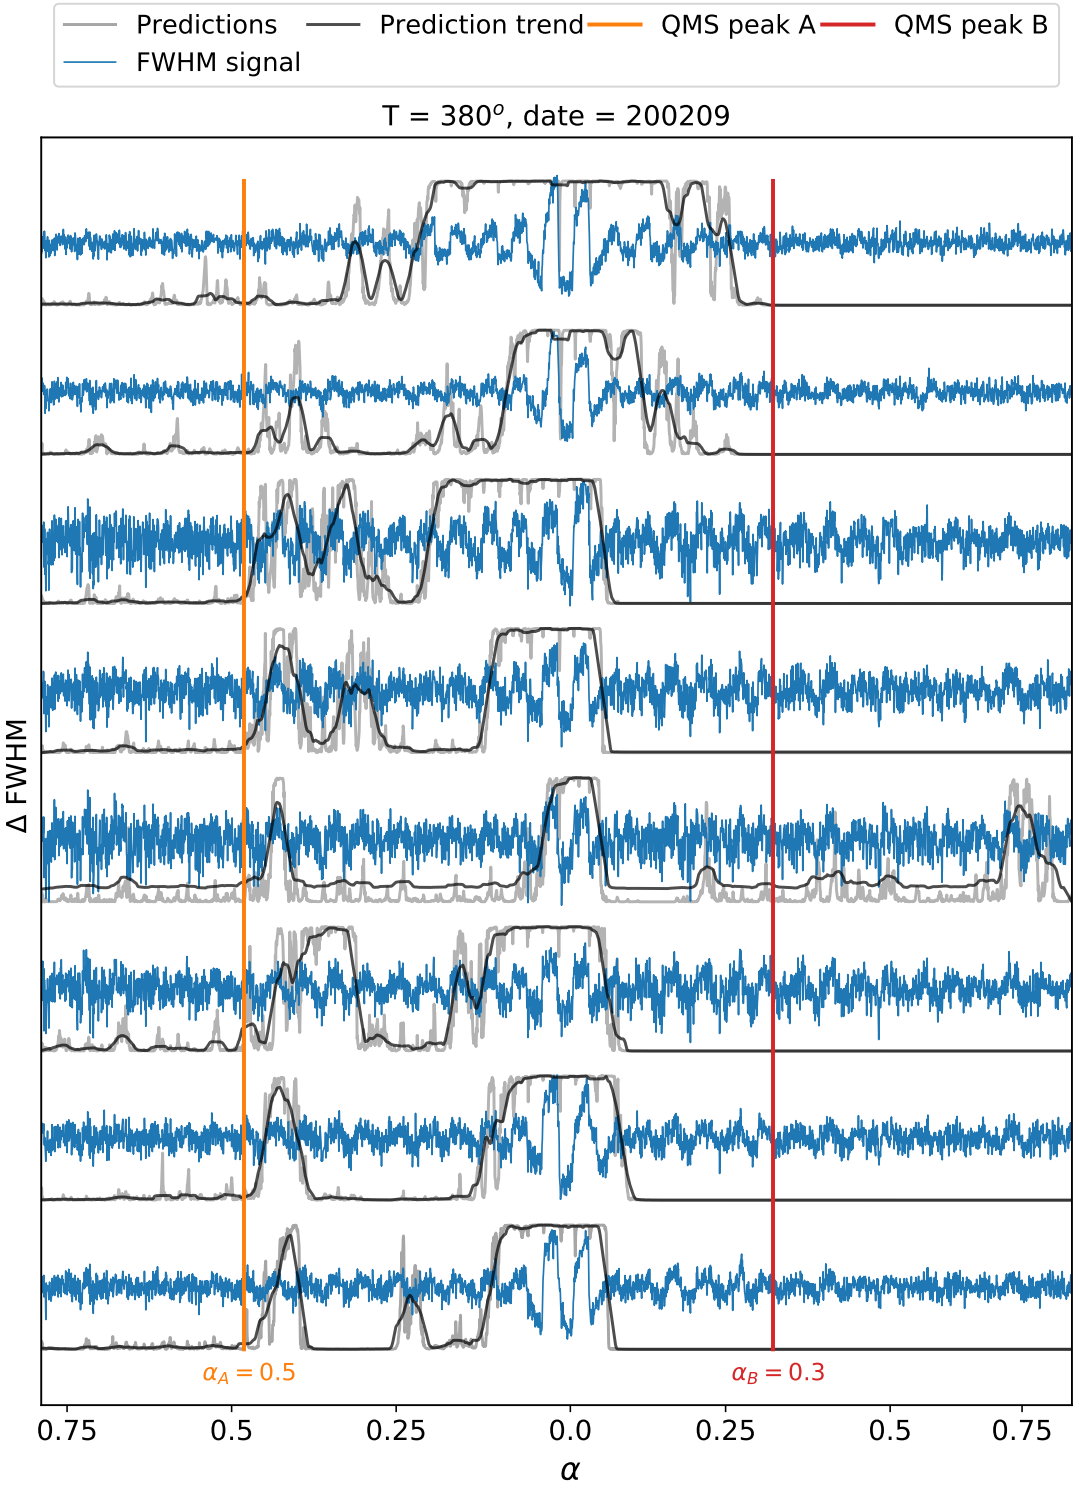

Figure A.38: Oscillatory predictions of FWHM signal with $T = 380°C$ collected on 200209.

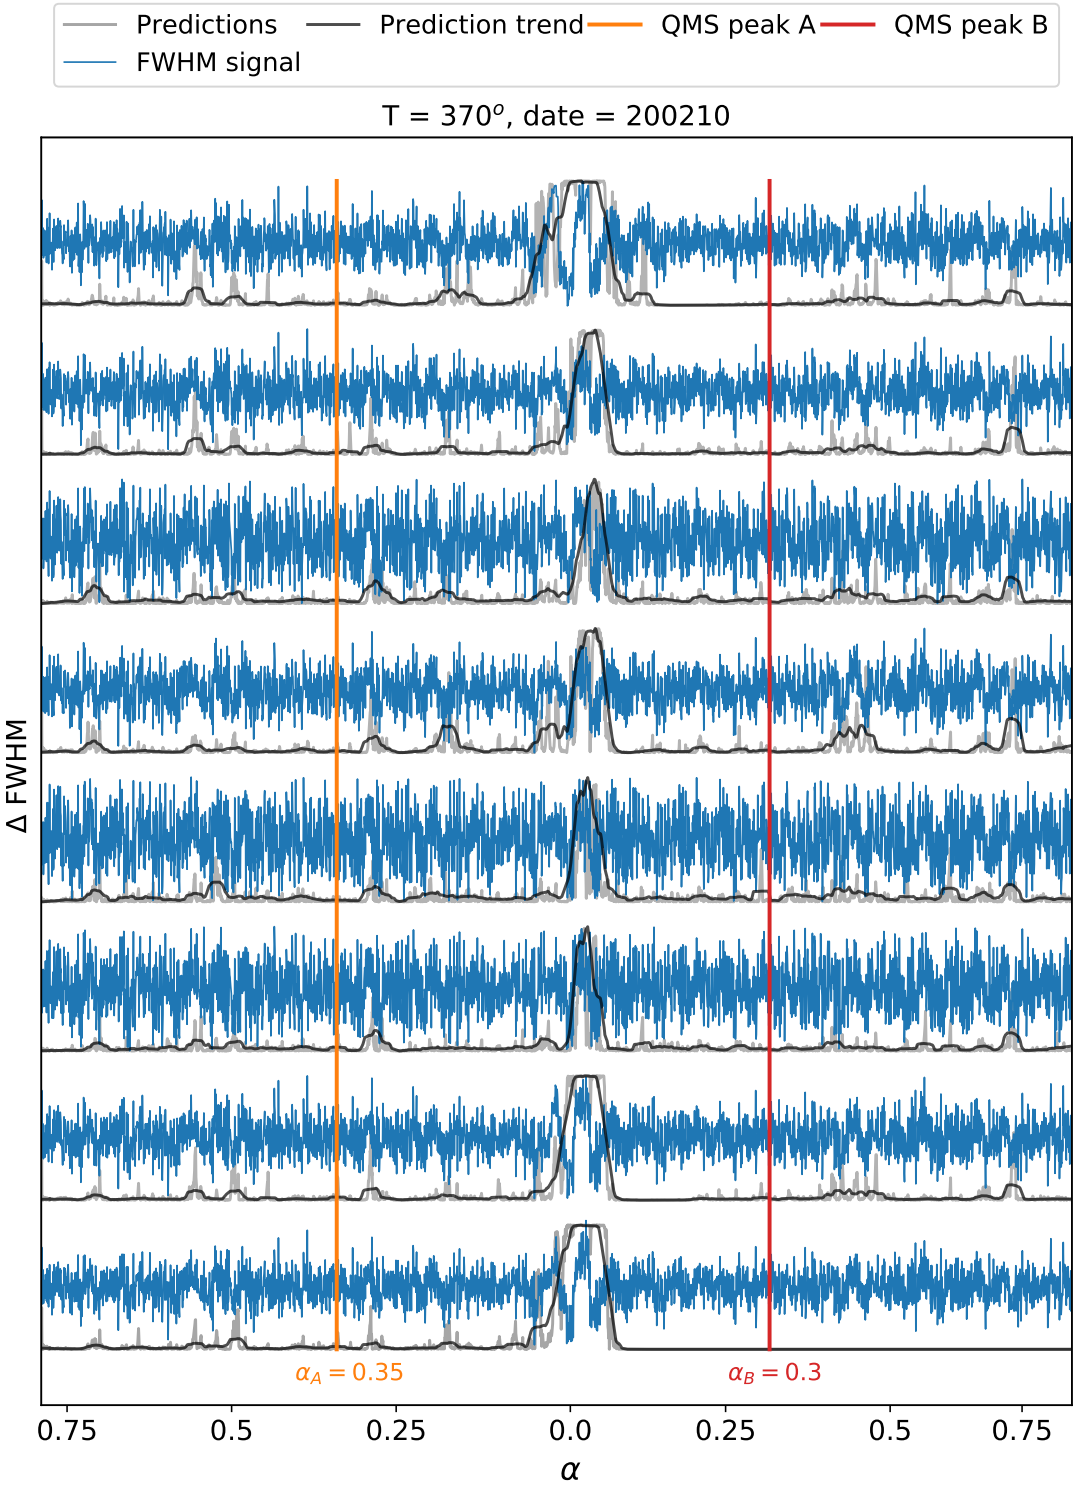

Figure A.39: Oscillatory predictions of FWHM signal with $T = 370°C$ collected on 200210.

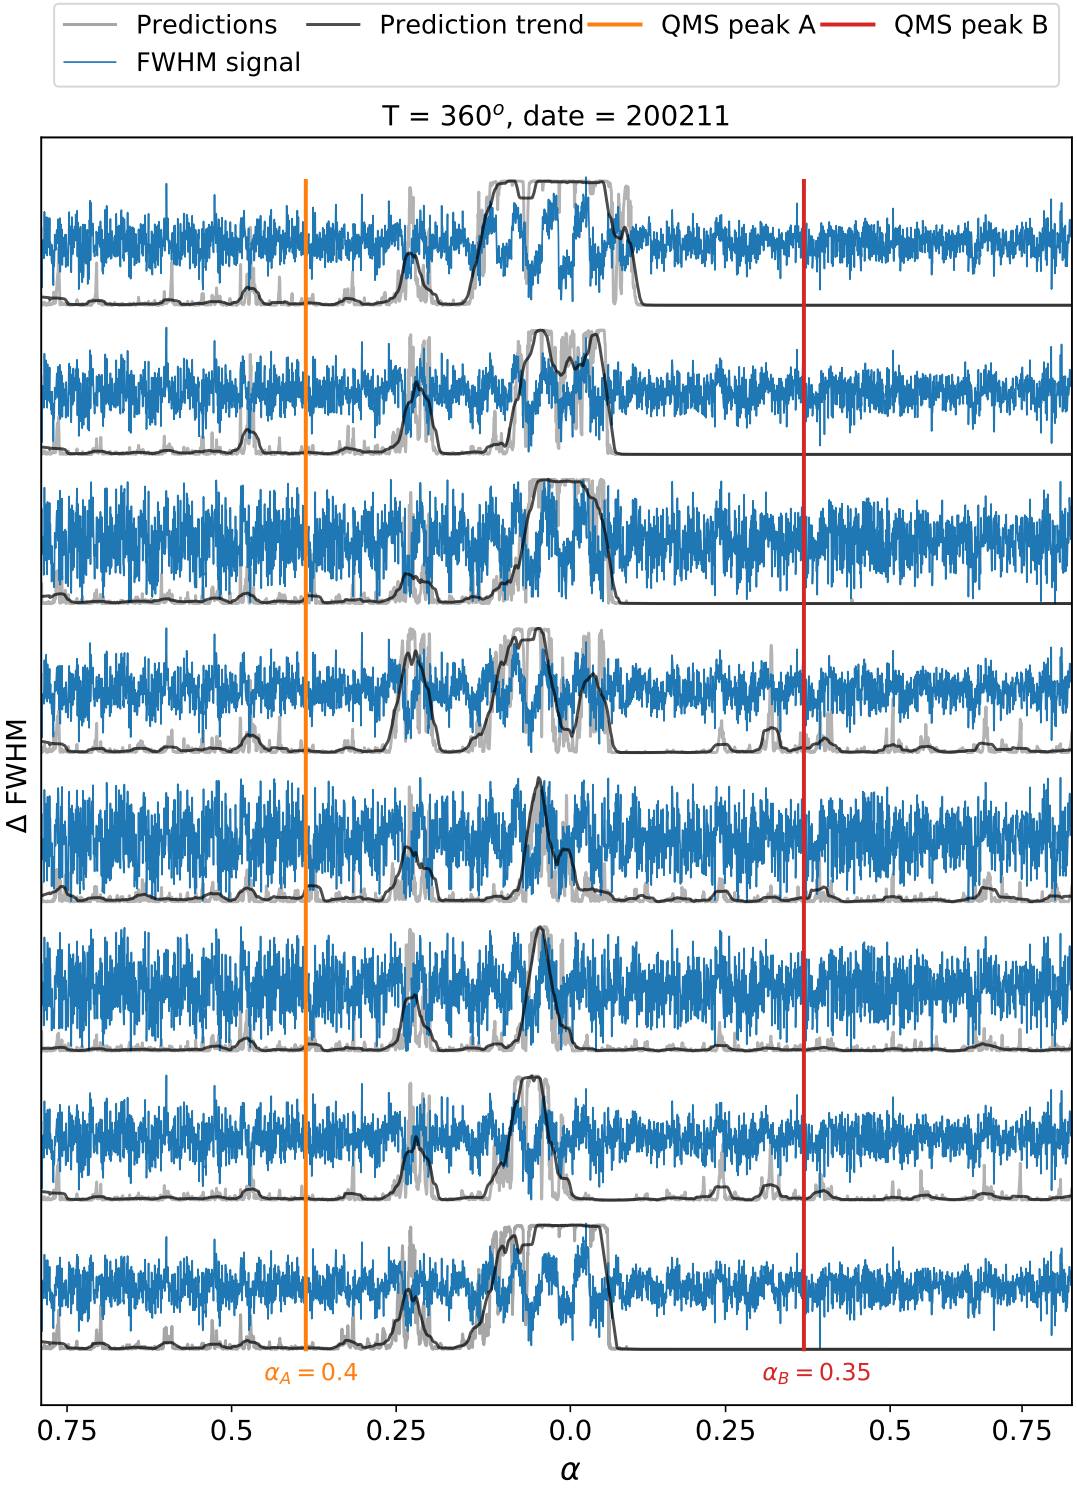

Figure A.40: Oscillatory predictions of FWHM signal with $T = 360°C$ collected on 200211.

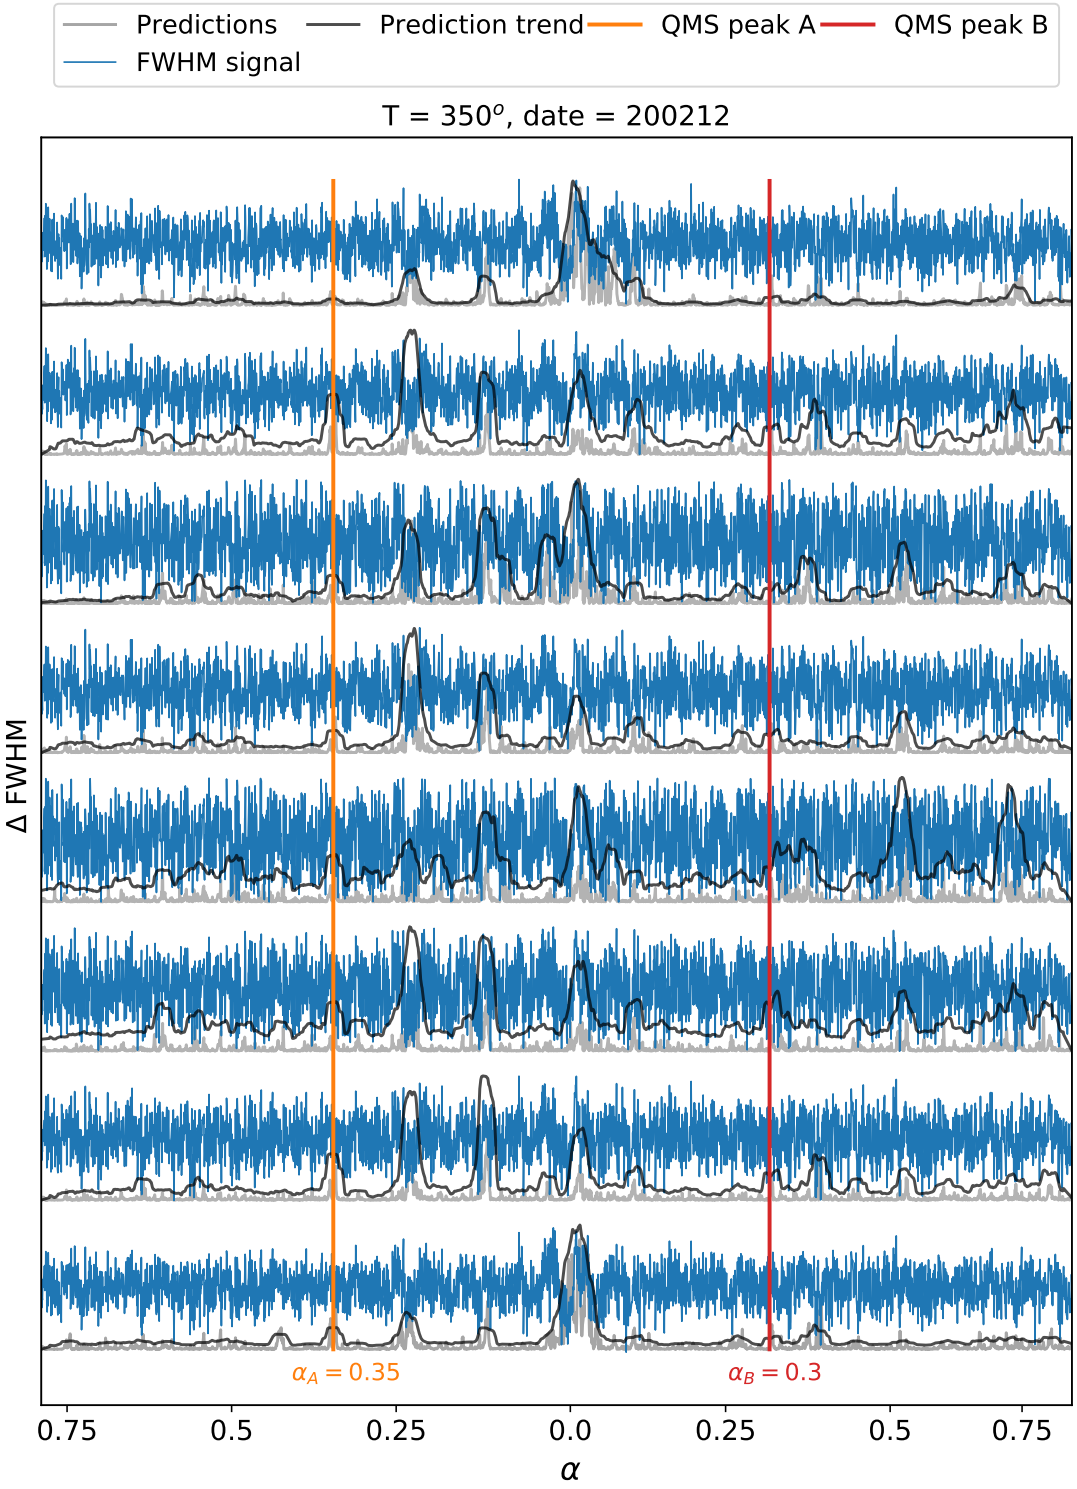

Figure A.41: Oscillatory predictions of FWHM signal with $T = 350°C$ collected on 200212.

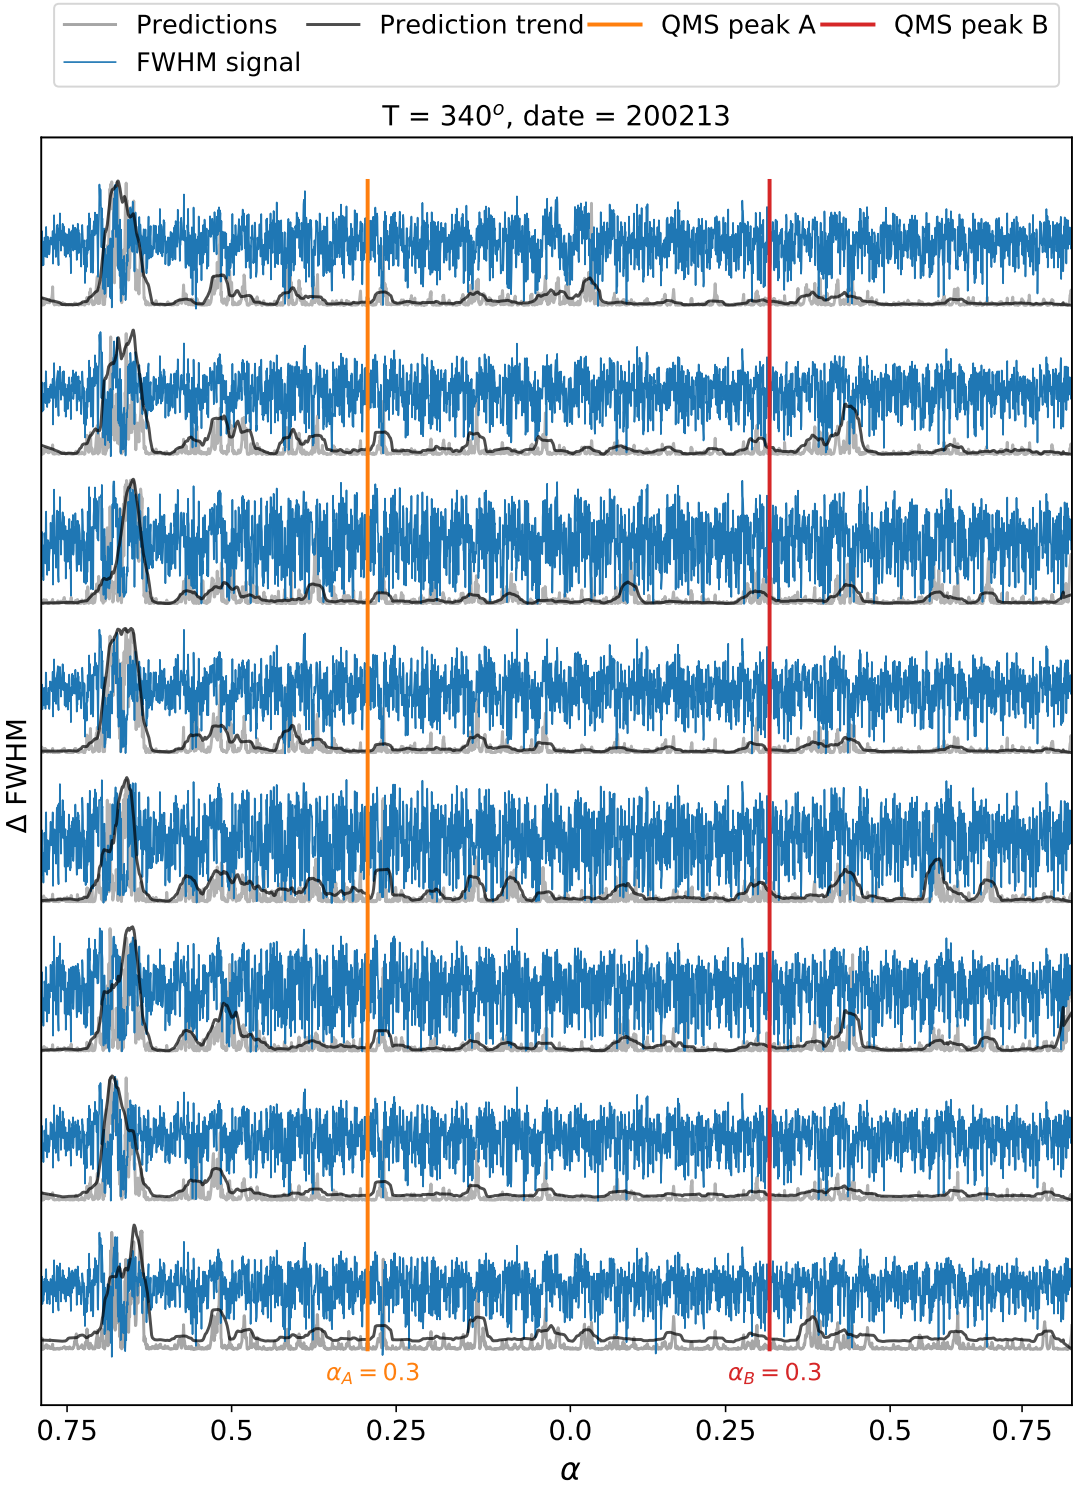

Figure A.42: Oscillatory predictions of FWHM signal with $T = 340°C$ collected on 200213.

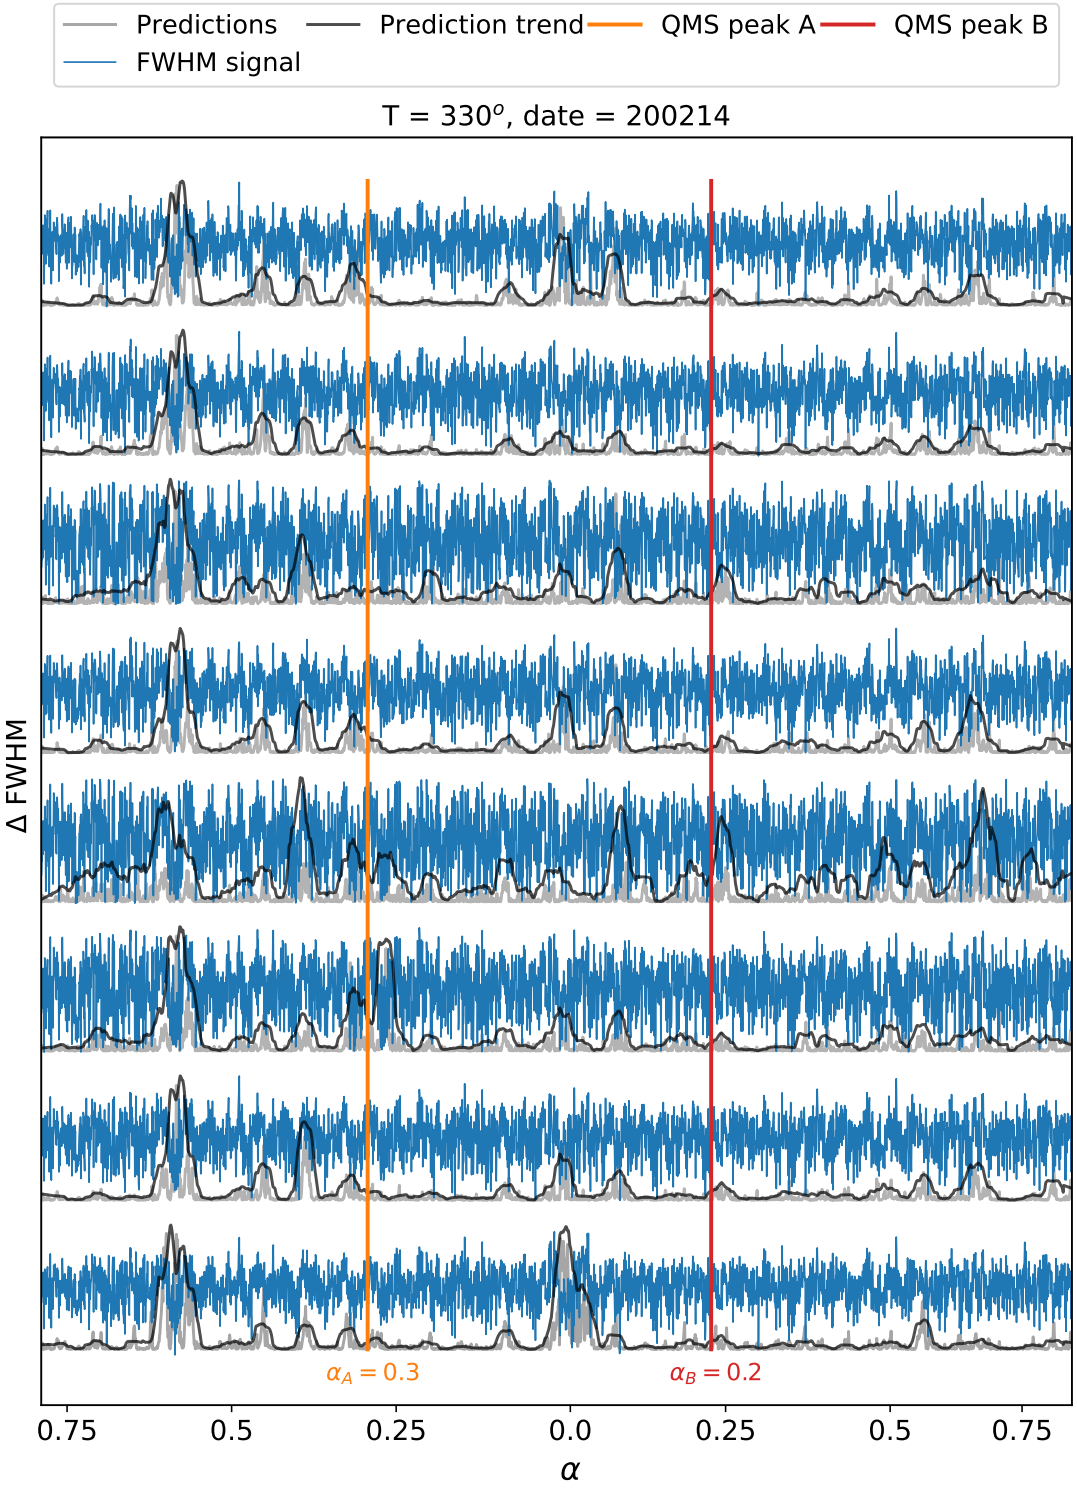

Figure A.43: Oscillatory predictions of FWHM signal with $T = 330°C$ collected on 200214.

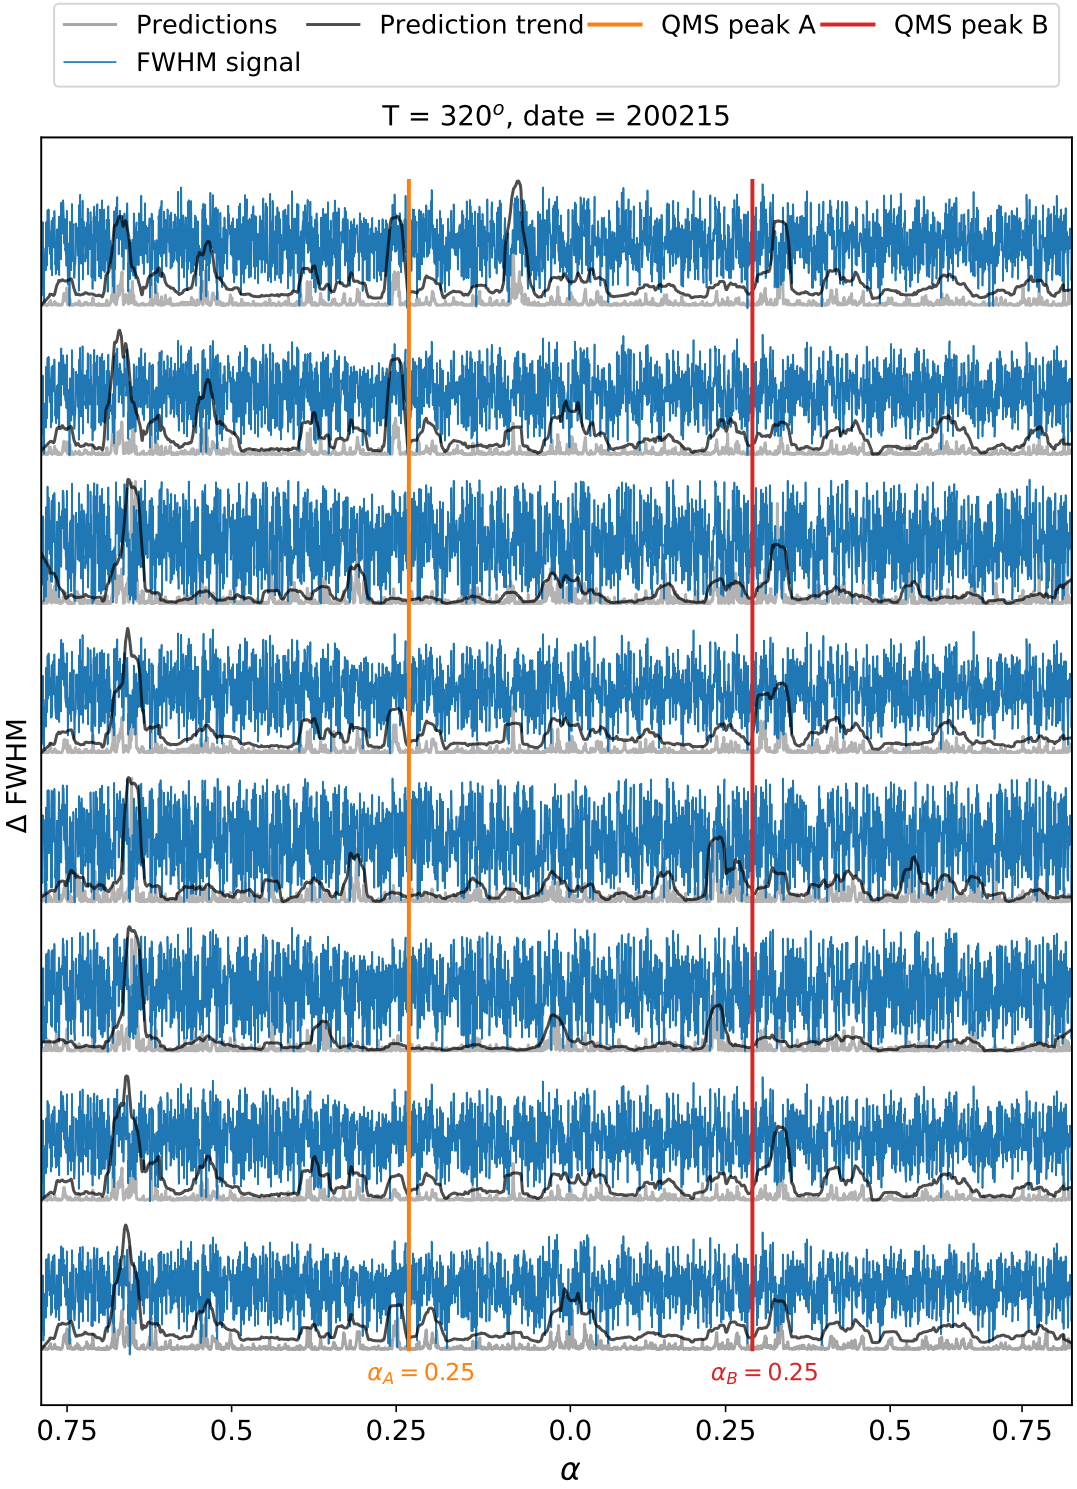

Figure A.44: Oscillatory predictions of FWHM signal with $T = 320°C$ collected on 200215.

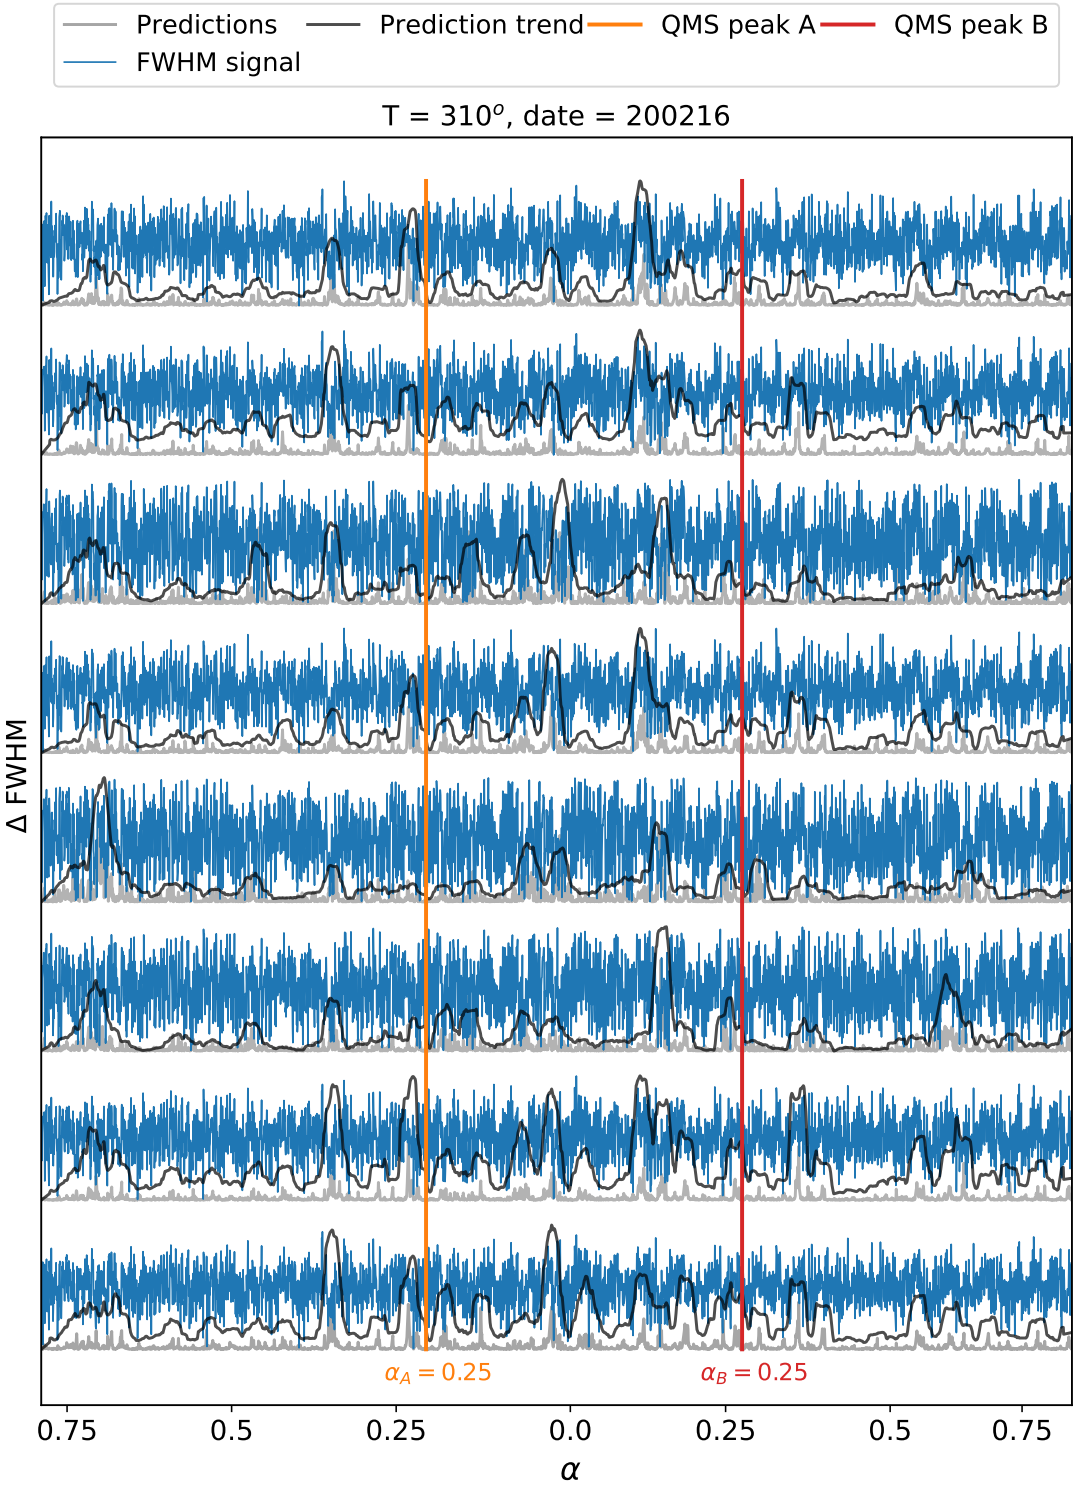

Figure A.45: Oscillatory predictions of FWHM signal with $T = 310°C$ collected on 200216.

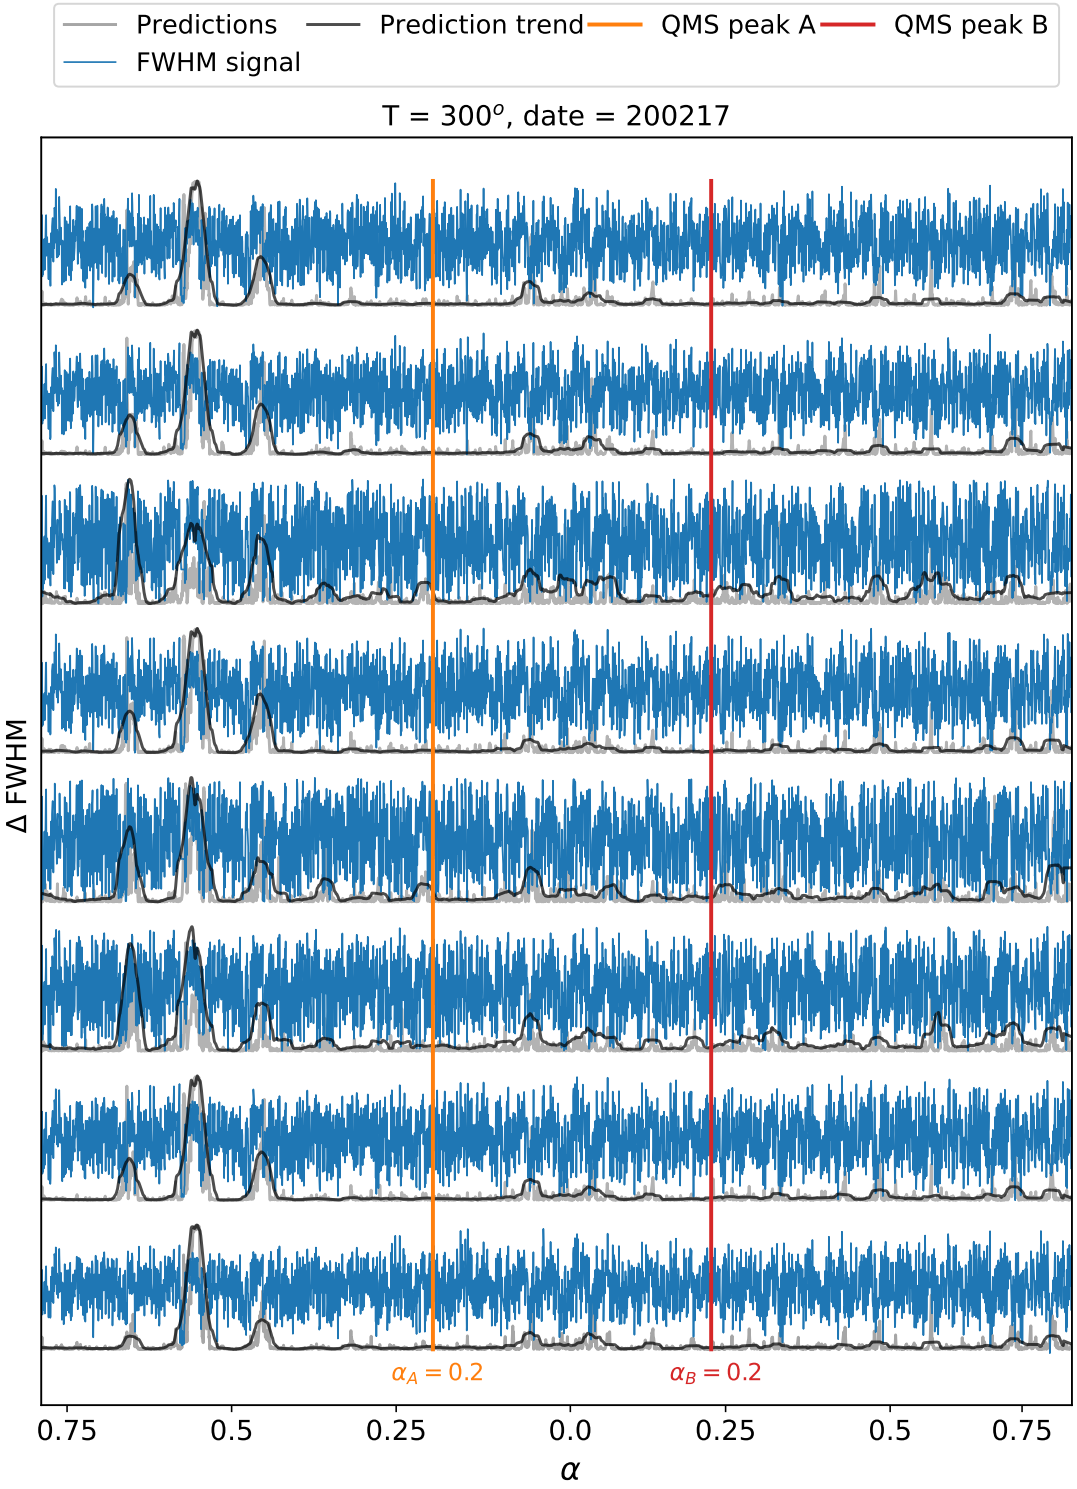

Figure A.46: Oscillatory predictions of FWHM signal with $T = 300°C$ collected on 200217.

A.3.7 Predictions of Chip C

In this section we present the predictions of our ”test” chip to further investigate if our network is able to generalize across different chips. The result is presented in the same way as for chip E. For this chip, we only have four measurements; three from 380°C and one from 340°C. From the measurements of 380°C, (see figure A.47) the average prediction of the switching points corresponds very well for the last two measurements and a bit worse for the first one. The results from the last two measurements of 380°C quite clearly indicates that the network is able to find an individual oscillatory regions that (on average) corresponds well to the QMS peaks, even for a chip for which we had no previous information on how the FWHM signals would look. Even the prediction from $T = 340$°C is decently accurate with respect to the QMS peaks, but for this measurement, the switching points were only found for three of the particles.

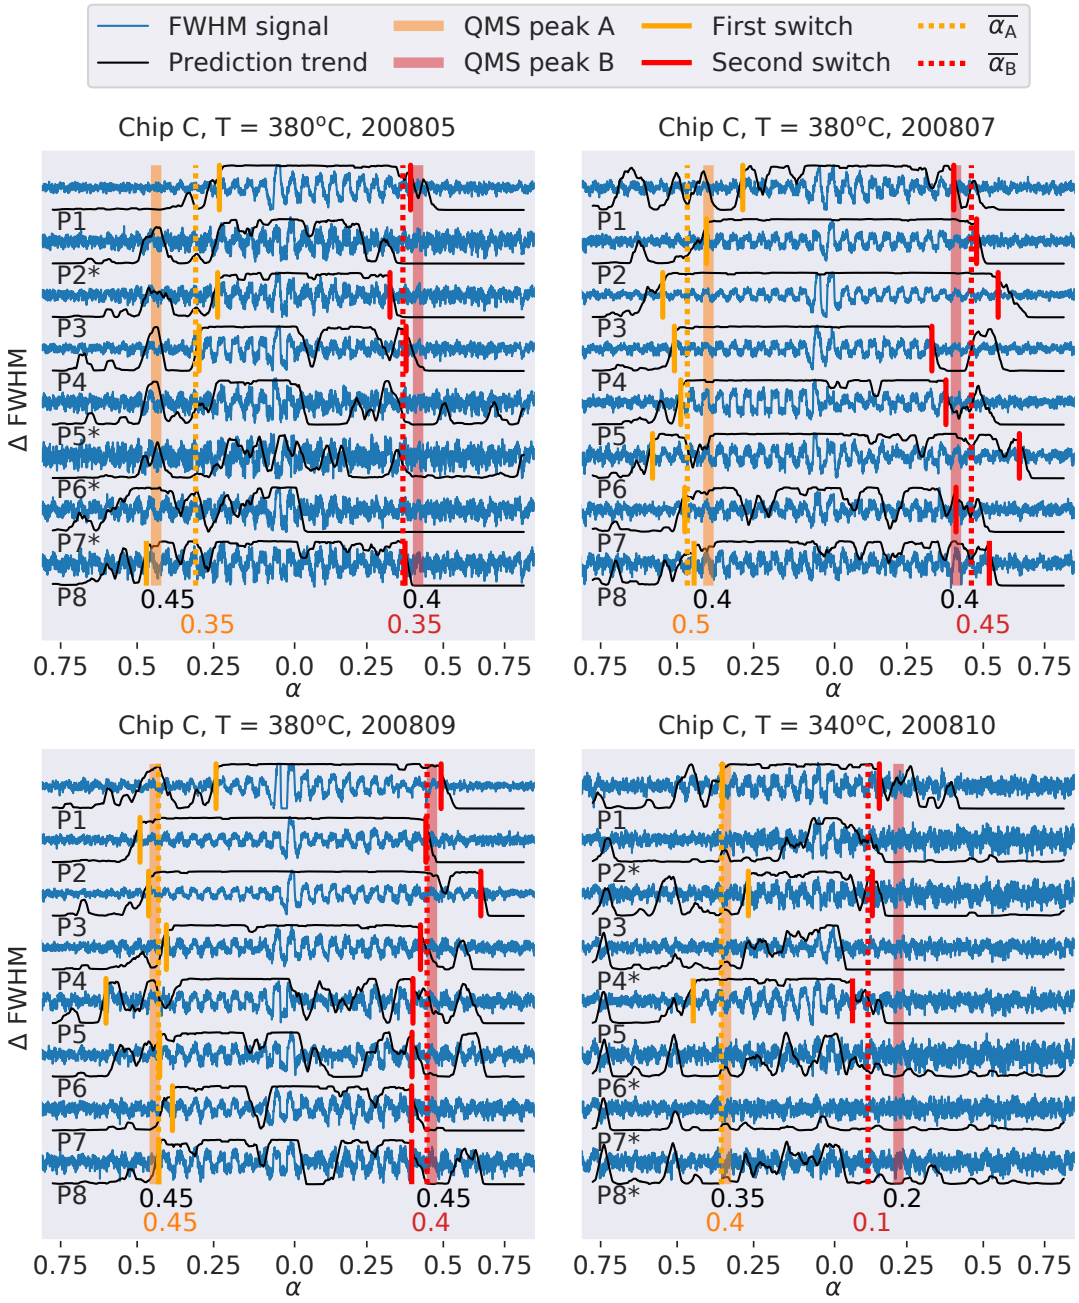

Figure A.47: Oscillatory predictions of FWHM signal from chip C with $T = 380°C$ and 340°C collected in August 2020. For particles marked with a an asterisk, no switching point was found.